

Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News*

Cosmin Ilut[†]
Duke & NBER

Matthias Kehrig[‡]
UT Austin

Martin Schneider[§]
Stanford & NBER

October 2013

Abstract

This paper studies models of firms that adjust asymmetrically to private signals and uses Census data to assess predictions at the establishment level. When firms' hiring decisions respond more to bad signals than to good signals, both aggregate volatility and the cross sectional dispersion of employment growth are countercyclical. Responses to measured technology shocks and the time series behavior of skewness are consistent with the mechanism. Even in the absence of physical adjustment costs and with homoskedastic shocks to fundamentals, information processing under Knightian uncertainty generates asymmetric adjustment.

*Preliminary and incomplete – comments welcome. We thank Nick Bloom and conference participants at the 2013 SED Annual Meeting and the 2013 Empirical Macro Workshop for helpful comments. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. The latest version of this paper can be downloaded at <http://papers.ssrn.com/abstract=2304442>.

[†]cosmin.ilut@duke.edu.

[‡]matthias.kehrig@austin.utexas.edu.

[§]schneidr@stanford.edu.

1 Introduction

Recent work has documented strong comovement between (*i*) the volatility of macroeconomic aggregates and (*ii*) the dispersion of micro-level variables from which the macroeconomic aggregates are constructed. For example, clusters of large changes in aggregate employment growth tend to go along with high dispersion of employment growth across firms. Models of the distribution of firms often account for this fact by assuming correlated shocks to first and second moments of fundamentals, such as productivity.

This paper shows that an *endogenous* link between aggregate volatility and cross sectional dispersion emerges naturally if firms respond asymmetrically to private signals about future profits. Suppose hiring after good news is slower than firing after bad news. A bad aggregate shock that lowers the mean of the distribution of private signals then has two effects. On the one hand, the mean signal is lower so hiring falls on average. On the other hand, the typical signal now brings bad news and thus generates a stronger employment response. As a result, dispersion also increases.

Asymmetric adjustment to private signals leads to a number of predictions beyond the comovement of aggregate volatility and cross sectional dispersion. In particular, it should induce negative skewness of employment growth in both the cross section and the time series. Moreover, firms' actions *in anticipation* of relevant fundamental shocks should depend on the sign of the shock realization: for example, bad (good) productivity realizations should be preceded, on average, by large drops (small increases) in hiring. We verify both sets of predictions in Census data.

The mechanism we study relies on two key elements. First, firms' actions respond to dispersed private signals about future profitability. For example, if total factor productivity at the firm level is persistent, the current TFP realization serves as a signal about future TFP. More generally, firms may receive intangible signals about future productivity or demand that need not be correlated with current profits. Dispersion in intangible signals may then arise simply because the signals are imperfect. For example, if every firm receives a noisy private signal about aggregate TFP, then the dispersion of signals is driven by the dispersion of the noise.¹

The second key element of our mechanism is that firms respond more to good news than to bad news. It does not matter exactly why this asymmetric adjustment obtains. One possibility is that the logistics of the hiring process directly make hiring more costly than firing. This could be, for example, because hiring new workers is subject to costly search,

¹The assumption that firms' actions respond to private signals also requires that there is no public signal that aggregates all dispersed information. For example, to the extent there are prices that all firms can easily observe, those prices must be sufficiently affected by noise.

whereas firing is free. A second candidate for asymmetric adjustment comes from information processing: if firm decision makers are averse to Knightian uncertainty (ambiguity) and are uncertain about the quality of signals, then it is also optimal to respond more to bad news.

Asymmetric adjustment to ambiguous signals arises even if there are no adjustment costs and all shocks are homoskedastic. Indeed, ambiguity averse firms evaluate hiring decisions *as if* taking a worst case assessment of future profits. With ambiguity about signal quality, the worst case then depends on what the signal says: for a good signal, the worst case interpretation is that it is noisy, whereas for a bad signal the worst case is that it is very precise. Updating from ambiguous signals thus endogenously generates asymmetric actions.

We use confidential Census data on U.S. manufacturing establishments 1972-2009 to test a number of predictions of our mechanism. First, the cross-sectional employment distribution is negatively skewed, that is, there are more firms that contract employment relative to the mean than there are that expand employment. This comes from the fact that firms with bad signals lay off lots of workers while firms that received good signals hire only very few workers. We find that in almost all years the distribution of employment growth rates across firms is negatively skewed.

Second, the cross-sectional employment dispersion endogenously varies over the business cycle. In bad times, there are more firms receiving bad signals to which they respond strongly by decreasing employment significantly. At the same time, there are less firms receiving good signals to which they respond only very weakly by hiring few workers, so the inter-quantile ranges increase. In the data, the cross-sectional inter-quartile range of employment growth rates is 30% more spread-out in recessions than in booms.

Third, the dynamic response of an individual firm to good versus bad shocks over time is markedly different although it faces a homoskedastic time series of innovations to total factor productivity. That is, the firm's TFP innovations are not skewed, but the firm's employment response is negatively skewed. When a firm receives a standard negative technology shock, it decreases hiring by 2.5%; after receiving a positive technology shock of the same size, the firm increases employment only by 0.5%.

In addition to the dynamics of the employment distribution of the cross-section and that of individual firms, this model explains a number of aggregate facts as well: Fourth, the growth rate of aggregate employment is negatively skewed: recessions bring precipitous aggregate employment contractions while booms are associated with slow aggregate hiring.

Fifth, downturns are not only associated with precipitous firing, but also more (time-series) volatility in the aggregate: This countercyclical volatility clustering means the aggregate employment growth that is smooth in expansionary years compared to recessionary years when it is volatile.

The above-described patterns of cross-sectional dispersion and aggregate volatility have been widely noted and studied, and we offer a new explanation of these facts. What sets our approach apart from other research is that it provides a unified explanation for micro- and macro-level facts and that this explanation is rooted in firms’ endogenous behavior. Firms will hire and fire such that their policy endogenously generates negative skewness in employment growth, countercyclical employment growth dispersion, negative skewness in both micro-level and aggregate employment growth rates and volatility clustering in aggregate employment growth.




Figure 1: Time-Varying Employment Volatility in the Aggregate and in the Cross-Section

Our paper is related to a recently growing literature studying the time-varying nature of economic fluctuations. Some research focuses on time variation in aggregate volatility, that is, the cyclical up- and downturns of the U.S. economy were large up to the early 1980’s – the “Great Inflation” –, declined in size since then – the “Great Moderation” –, but may have risen again recently in the “Great Recession.” Most research in this area – see for example [Stock and Watson \(2002\)](#); [Justiniano and Primiceri \(2008\)](#); [Fernández-Villaverde et al. \(2011\)](#); [Carvalho and Gabaix \(2013\)](#) – attribute the low-frequency movements in aggregate volatility to changes in the exogenous shock processes. [Sims and Zha \(2006\)](#); [Fernández-Villaverde et al. \(2012\)](#) consider long-run changes in monetary or fiscal policies as an alternative source of changing exogenous shocks.

At the same time, another strand of research has studied how differences between firms change over the business cycle: the dispersion across firms is larger in recessions than in booms and has been found in the dispersion of productivity levels ([Kehrig \(2013\)](#)), productivity innovations ([Bloom et al. \(2012\)](#)), employment dispersion and returns to capital

(Eisfeldt and Rampini (2006)).² Arellano et al. (2010); Christiano et al. (forthcoming); Bloom et al. (2012) consider these fluctuations in dispersion as exogenous uncertainty and causal for business cycles.

What is common across the two strands of the literature outlined above is their focus as time-varying volatility being exogenous and causal. We ask if time-varying volatility could be an endogenous outcome and if the aggregate and firm-level volatility changes could be linked. To that end, we propose a model where agents are ambiguity averse. The model generates changes in volatility and skewness in aggregate and idiosyncratic employment growth from symmetric and homoskedastic shocks.

2 The Model

Here we present a simple model with a continuum of firms that at the beginning of each period get an idiosyncratic noisy signal about end-of-period productivity z_t^i . After observing the signal, each firm chooses employment in a competitive labor market, where the wage is \bar{w} . At the end of the period, productivity is realized. Thus the firm problem is a repetition of static hiring decisions based on a signal-extraction problem within the period.

Firm i 's log productivity is described as

$$z_t^i = a_t + b_t^i - \frac{1}{2} (\sigma_a^2 + \sigma_b^2)$$

where a_t is an aggregate shock, normally distributed with mean \bar{a} and variance σ_a^2 , and b_t^i is an idiosyncratic, firm-specific shock, normally distributed with mean 0 and variance σ_b^2 . The variances σ_a^2 and σ_b^2 , are constant over time and known to the firm. The end-of-period profit is

$$\exp(z_t^i) (L_t^i)^\alpha - \bar{w} L_t^i$$

where L_t^i is employment chosen by firm i . The noisy signal about firm's i productivity is

$$s_t^i = z_t^i + \sigma_e \varepsilon_t^i$$

where ε_t^i is a standard normal innovation, *iid* across time and firms.

We assume that firms are ambiguous about the information precision of the signals that they received. In particular, similarly to the models used in Epstein and Schneider (2008) and Ilut (2012), we assume that each firm is ambiguous about the value of σ_e and has a set

²A notable exception is investment dispersion which is procyclical (Bachmann and Bayer (2011); especially across establishments *within* multi-establishment firms Kehrig and Vincent (2013)).

of beliefs given by

$$\sigma_e \in [\underline{\sigma}_\varepsilon, \bar{\sigma}_\varepsilon]$$

The objective of the firm is choose L_t^i to maximize the multiple priors utility:

$$\max_{L_t^i} \min_{[\underline{\sigma}_\varepsilon, \bar{\sigma}_\varepsilon]} (L_t^i)^\alpha E^{\sigma_e} [\exp(z_t^i | s_t^i)] - \bar{w} L_t^i$$

where the minimization operator reflects the firm's ambiguity-aversion. The key characteristic of firms' preferences is that, faced with uncertainty over the signal-to-noise ratio, they act *as if* the worst-case σ_e characterizes the true DGP. The worst-case σ_e minimizes the conditional expectation of end-of-period profits, which are only a function of the expected z_t^i conditional on the observed signal s_t^i . Details and axiomatic foundations for the multiple priors utility are provided in [Gilboa and Schmeidler \(1989\)](#) for the static case and in [Epstein and Schneider \(2003\)](#) for the dynamic version. [Ilut and Schneider \(2011\)](#) provide a general framework to solve business cycle models featuring recursive multiple priors utility.

Because of the normality of innovations, the problem above is equivalent to

$$\max_{L_t^i} \min_{[\underline{\sigma}_\varepsilon, \bar{\sigma}_\varepsilon]} \exp \left[E^{\sigma_e} (z_t^i | s_t^i) + \frac{1}{2} \text{var}^{\sigma_e} (z_t^i | s_t^i) \right] (L_t^i)^\alpha - \bar{w} L_t^i \quad (1)$$

2.1 Asymmetric responses

To characterize the optimal solution, it is analytically helpful to define the relative precision of signal γ_t as a function of a given $\sigma_{e,t}^2$:

$$\gamma_t = \frac{\text{var}(z_t^i)}{\text{var}(z_t^i) + \sigma_{e,t}^2}$$

After observing the signal, the posterior conditional mean and variance of z_t^i are given by

$$\begin{aligned} E(z_t^i | s_t^i) &= \gamma_t \left[s_t^i + \frac{1}{2} \text{var}(z_t^i) \right] - \frac{1}{2} \text{var}(z_t^i) \\ \text{var}(z_t^i | s_t^i) &= (1 - \gamma_t) \text{var}(z_t^i) \end{aligned}$$

The firm problem in (1) then simplifies to

$$\max_{L_t^i} \min_{[\underline{\sigma}_\varepsilon, \bar{\sigma}_\varepsilon]} \exp(\gamma_t s_t^i) (L_t^i)^\alpha - \bar{w} L_t^i$$

The solution to this problem results in a hiring policy that is based on the worst case

precision γ_t^* characterized by:

$$L_t^i = \left[\frac{\alpha}{\frac{w}{w}} \exp(\gamma_t^* s_t^i) \right]^{\frac{1}{1-\alpha}}; \quad \gamma_t^* = \begin{cases} \bar{\gamma} & \text{if } s_t^i < 0 \\ \underline{\gamma} & \text{if } s_t^i \geq 0 \end{cases} \quad (2)$$

The interpretation of the optimal solution is that the firm acts *as if* the signal precision is high for bad news and low for good news.

The labor decision is then to maximize expected profits under the worst-case precision γ_t^* . This results in an asymmetric hiring decision rule such that the firm receives a negative signal $s_t^i = -x$, then employment contracts by more than it would expand if the firm receives a positive signal of the same magnitude $s_t^i = x$. Figure 2 plots in the top panel such an asymmetric response to signals together with a normal distribution for signals in the bottom panel.




Figure 2: Firms respond asymmetrically to symmetric shocks

2.2 Implications for employment distribution across firms

2.2.1 Negative Skewness

To be added.

2.2.2 Countercyclical Dispersion

What are the implications of this firm behavior for the overall economy? We first focus on the implications of the cross-sectional dispersion of employment growth rates. The literature has noted that the cross-sectional dispersion is more spread out in recessions than booms (see among others Bloom et al. (2012)) and has interpreted this as a result of (firm-level) uncertainty shocks. Since firms respond to their signals in a nonlinear fashion and firms receive different signals, our model has the potential to address cross-sectional distribution dynamics. We illustrate the point by examining the inter-quartile range. As one can see from (2), cross-sectional dispersion will be driven by the variance of signals (which we assume to be constant) and parameters that involve the distribution of γ , i.e. the measure of firms with negative signals (which changes endogenously).

From the decision rule in (2), it follows that the cross-sectional quantiles of $\log L_t^i$ are monotonic in those of TFP signals. In particular, the top and bottom quartiles are given by

$$Q_3^L = c\gamma^*(Q_3^s)Q_3^s; \quad Q_1^L = c\gamma^*(Q_1^s)Q_1^s$$

where Q_3^L and Q_3^s denote the top quartile and Q_1^L and Q_1^s the bottom quartiles for the $\log L_t^i$ and s_t^i cross-sectional distributions. The $\log L_t^i$ quartiles are a function of the signal quartiles where the proportionality factor has two parts: the first one is a constant part, denoted above by c , capturing the terms α, \bar{w} in (2). The second is given by the worst-case signal precision used to evaluate the signal in forming the conditional probability. This factor endogenously changes as the sign of the observed signal changes. In particular this weight is larger if the corresponding signal quartile Q_m^s is negative:

$$\gamma_t^*(Q_m^s) = \begin{cases} \bar{\gamma} & \text{if } Q_m^s < 0 \\ \underline{\gamma} & \text{if } Q_m^s \geq 0. \end{cases}$$

Moreover, because of the normality of the signal distribution we know that

$$Q_3^s = E(s^i) + 0.67\sqrt{Var(s^i)}; \quad Q_1^s = E(s^i) - 0.67\sqrt{Var(s^i)}$$

so we can analytically compute the interquartile range $IQR_L \equiv Q_3^L - Q_1^L$. In Figure 3 we graphically present the implied IQR by varying the mean of the signal. There we represent $E(s^i)$ conveniently as a multiple of the form $x\sqrt{Var(s^i)}$ and the horizontal axis varies the factor x .

Very good average signals When the mean signal is larger than $0.67\sqrt{Var(s^i)}$, then both quartiles are positive and the hiring at both quartiles is based on the low precision interpretation of the signals

$$IQR_L = c\underline{\gamma}(Q_3^s - Q_1^s) = c\underline{\gamma} \left(1.34\sqrt{Var(s^i)} \right).$$

In this region, IQR_L is independent of $E(s^i)$ and it is proportional to $\sqrt{Var(s^i)}$. As shown further below, this proportionality factor is lower than in the case in which hiring at both quartiles responds based on the high precision of signals.

Neutral average signals As $E(s^i)$ decreases below $0.67\sqrt{Var(s^i)}$, but it is still larger than $-0.67\sqrt{Var(s^i)}$, then the response of Q_3^L to Q_3^s is lower than that of Q_1^L to Q_1^s :

$$\begin{aligned} IQR_L &= c\underline{\gamma}Q_3^s - c\bar{\gamma}Q_1^s \\ &= c(\underline{\gamma} - \bar{\gamma})E(s^i) + c(\underline{\gamma} + \bar{\gamma})0.67\sqrt{Var(s^i)} \end{aligned}$$

Since $\underline{\gamma} < \bar{\gamma}$, we see that the IQR_L increases as $E(s^i)$ decreases.

Very bad average signals Finally, when the mean signal is worse enough so that it is lower than $-0.67\sqrt{Var(s^i)}$, then the response of both Q_3^L and Q_1^L are strong:

$$IQR_L = c\bar{\gamma}(Q_3^s - Q_1^s) = c\bar{\gamma} \left(1.34\sqrt{Var(s^i)} \right)$$

In this region, IQR_L is higher by a factor of $\bar{\gamma}/\underline{\gamma}$ than in the region when $E(s^i) \geq 0.67\sqrt{Var(s^i)}$.

The asymmetric decision rule implied by the ambiguity-averse behavior of individual firms will endogenously lead to **countercyclical** employment dispersion.

Similar arguments can be made by any inter-quantile ranges, the inter-decile range for example is flat if $|E(s^i)| > 1.28\sqrt{Var(s^i)}$ while it decreases in $E(s^i)$ as long as $|E(s^i)| < 1.28\sqrt{Var(s^i)}$.

2.3 Implications for aggregate employment

What kind of behavior does an aggregate economy display when its firms are ambiguity averse? It is well known that aggregate employment – like most other time series – exhibits time-varying volatility, i.e. periods of small fluctuations alternate with periods of large fluctuations, and negative skewness, i.e. employment contractions are sharper and deeper than employment expansions. We will aggregate employment and study the effects of aggregate




Figure 3: Cross-sectional employment growth dispersion is countercyclical

technology shocks a_t . A fluctuating a_t will influence the signals and the share of firms that receive good versus bad signals. The fundamental asymmetry how individual firms treat good versus bad signals is at work here, too.

The average employment integrates individual firms' strong negative and weak positive responses to their observed signals:

$$L_t = \int L_t^i di = c \left(\int_{-\infty}^0 \exp(\bar{\gamma}s_t^i) f(s_t^i) ds_t^i + \int_0^{\infty} \exp(\underline{\gamma}s_t^i) f(s_t^i) ds_t^i \right) \quad (3)$$

Equation (3) shows the main effect of changes in the aggregate component of TFP. On the one hand, consider the case in which a_t falls by some amount x . Then more firms will receive negative signals, and since these firms respond strongly to such signals, the aggregate L_t falls a lot. On other hand, if a_t increases by x , then more firms will receive good signals and these firms will endogenously respond less to their observed signals. Importantly, this means that aggregate L_t responds proportionally less to increases than to similar decreases in a_t . This aggregate differential proportionality factor means that asymmetric individual decision rules will generate two important properties of the higher moments of L_t : First, as L_t responds stronger to negative a_t realizations, there will be **countercyclical aggregate volatility clustering**. Second, due to the same mechanism, given symmetric shocks to a_t there will be **negative skewness in time-series** of the aggregate L_t .

Figure 4 shows visually the three main implications of the model: countercyclical dispersion in the cross section, countercyclical aggregate volatility and negative skewness in the aggregate. We feed in a particular sequence of innovations to a_t , the aggregate TFP, marked by the red dotted line. Note how the dotted red line is a homoskedastic time series. We then

compute the model implied aggregate L_t (solid blue line) and cross-sectional $IQR_{L,t}$ (solid red line). The top panel shows the countercyclical nature of IQR. The flat parts of IQR on the downside (period 12/13) or upside (period 17/18) reflect the flatness for large shifts in the mean signal evident in Figure 3. The middle panel shows how two periods (marked by the two-blue circles) that are characterized by symmetric TFP shocks are marked by countercyclical aggregate volatility. The bottom panel highlights that for symmetric increases and decreases in aggregate TFP, the resulting aggregate hiring is negatively skewed in the time-series, with sharper contractions than booms.




Figure 4: Time-Varying Volatility and Negative Skewness in Simulated Data

Note: This figure displays aggregate employment growth and the cross-sectional inter-quartile of employment growth in simulated data. As highlighted in the middle panel, recessions are times when aggregate volatility is high and cross-sectional dispersion wide. As the bottom panels shows downturns in aggregate employment are deep while upturns are mild thus creating negative time-series skewness.

3 Empirics

3.1 Data sources

Building on the main idea how higher moments of firm-level and aggregate employment may be a consequence of asymmetric decision rules, we use micro-level and aggregate data to establish an array of empirical facts and test the main implications of our model.

We use confidential data on manufacturing establishments collected by the Census Bureau which comprise the Annual Survey of Manufactures (ASM), the Census of Manufactures (CMF), the Plant Capacity Utilization Survey (PCU) and the Longitudinal Business Database (LBD). Following the literature, we identify an establishment as an individual decision maker. We combine the Census data with industry-level data from several publicly available sources: price deflators from the NBER-CES Manufacturing Industry Database (NBER-CES)³, various asset data from the the Capital Tables published by the Bureau of Labor Statistics (BLS)⁴ and the Fixed Asset Tables published by the Bureau of Economic Analysis (BEA)⁵. Unless otherwise noted, all datasets are at annual frequency.

From the Census of Manufactures (CMF) and the Annual Survey of Manufactures (ASM) we construct a large dataset of plants in the U.S. manufacturing sector. This panel spans the years 1972-2009, which allows us to study business cycle properties. Every year, there are about 55k observations for a total of 2.1 million. For more details, see the data description in [Kehrig \(2013\)](#).

We focus on the changes in a firm’s total employment (the sum of production and non-production workers): $n_t^i \equiv d \log(L_t^i)$. This choice is helpful empirically as it free of any specific metric, eliminates any ‘fixed effects’ specific to the the firm or the industry and it is motivated by a dynamic version of the model presented above. The structure of that dynamic model is similar to the static model, except that we make the stock of employment a state variable. There, the firm enters the period with some stock of labor, it first observes an idiosyncratic signal about next period’s TFP innovation and then chooses net hiring.

3.2 Asymmetric responses

We test our main prediction by regressing employment growth on innovations in total factor productivity. We focus on TFP innovations rather than TFP levels because we think of the

³The NBER-CES Manufacturing Industry Database is a joint program of the National Bureau of Economic Research and the Census Bureau; <http://www.nber.org/nberces/>.

⁴1987-2011 Capital Data for Manufacturing Industries <http://www.bls.gov/mfp/mprdownload.htm>.

⁵Tables 3.1S, 3.1E, 3.3S, 3.3E, 3.7S, 3.7E, 3.8S and 3.8E at <http://www.bea.gov/national/FA2004/SelectTable.asp> .

former as being responsible for *changes* in employment which should manifest themselves in employment growth rates. The main mechanism of the model relies fundamentally on whether an establishment’s employment growth responds asymmetrically to signals about future productivity innovations. To test this we first require a measure of innovations. Following the literature (see for example Kehrig (2013)), we first estimate establishment-level TFP, and then estimate a random walk process for it to recover innovations, denoted here by ω_t^i .

Although we, as econometricians, do not observe the idiosyncratic signals received by firms, we can use the fact that these unobserved signals show up on average in future innovations. Thus we run the following regression:

$$n_t^i = \alpha + \beta_{pos}\omega_{t+1}^i + \beta_{neg}\omega_{t+1}^i\mathbf{1}\{\omega_{t+1}^i < 0\} + \theta X_t^i + c_i + y_t + \epsilon_t^i \quad (4)$$

where we allow for firm fixed effects, a time trend, other industry characteristics, the lagged TFP level and the log-employment level $\log(L_t^i)$. The regression allows for a piece-wise linear differential effect of negative versus positive innovations on hiring. In data generated from the model above where firms respond weakly to positive technology shocks but strongly to negative ones the asymmetry should manifest itself in positive estimates of both β_{pos} and β_{neg} . That is, firms do respond to technology shocks in general, but when they are negative they respond very strongly. Table 1 reports the main results for this regression. We find significant evidence of asymmetry. In our preferred specification, in which TFP follows a random walk, a typical *positive* TFP shock increases employment by 0.5%, while a typical *negative* TFP shock decreases employment by 2.5%.

Could this asymmetry in employment result from asymmetric shocks in the first place? That is, are negative TFP innovations just larger than positive ones? As detailed in the next subsection we find that the cross-sectional dispersion of TFP *innovations* is not negatively skewed.

Could this result follow from a labor adjustment cost? If firms have a hard time finding workers after a positive technology shock, but can easily fire workers, a similar asymmetry might arise. Although we do not have detailed information about the establishment’s costs that are associated with hiring efforts and firing problems, we do utilise information reported in the Plant Capacity Survey (PCU). This subset of the ASM sample reports data on capacity utilisation and reasons why the establishment is not producing at full capacity. Among several options (too little demand, lack of materials, weather disruptions, ...) establishments can also report “lack of sufficient/sufficiently skilled labor” as a reason for producing less at full capacity. We interpret a positive answer to that question as indication of hiring obstacles

when the firm wants to hire. We thus modify regression (4) as follows

$$n_t^i = \alpha + \beta_{pos}\omega_{t+1}^i + \beta_{cstr}\omega_{t+1}^i\mathbf{1}\{\omega_{t+1}^i > 0\} + \beta_{neg}\omega_{t+1}^i\mathbf{1}\{\omega_{t+1}^i < 0\} + \theta X_t^i + c_i + y_t + \epsilon_t^i \quad (5)$$

If constraints in labor markets were entirely responsible for the asymmetric employment response, then we would expect β_{pos} to increase such that $\beta_{pos} \approx \beta_{neg}$ and $\beta_{cstr} < 0$. Then, only the constrained firms would be the ones responsible for the asymmetric employment response. The results of this modified regression are displayed in columns (2) and (4). Although β_{pos} is slightly higher than before, β_{pos} and β_{neg} are significantly different from each other and the asymmetry persists. As expected, $\beta_{cstr} < 0$, so that constrained firms do hire less after positive technology shocks; in fact, their employment stagnates and is indistinguishable from zero.

Table 1: Asymmetric employment responses

Specification	FE	FE	FE	FE
Sample	ASM	ASM	PCU	PCU
<i>A. Unweighted regressions</i>				
Firms w/ pos. shock	+0.8%*** (0.1%)	+0.8%*** (0.1%)	+0.8%*** (0.2%)	+0.8%*** (0.2%)
Firms w/ pos. shock & hiring constraint		+0.3% (0.4%)		+0.1% (0.3%)
Firms w/ neg. shock	-2.0%*** (0.1%)	-2.0%*** (0.2%)	-2.3%*** (0.4%)	-2.8%*** (0.4%)
<i>B. Employment-weighted regressions</i>				
Firms w/ pos. shock	+0.4%*** (0.1%)	+0.5%*** (0.1%)	+0.7%*** (0.2%)	+0.8%*** (0.2%)
Firms w/ pos. shock & hiring constraint		-0.4% (0.3%)		-0.2% (0.4%)
Firms w/ neg. shock	-2.5%*** (0.3%)	-2.5%*** (0.3%)	-2.7%*** (0.7%)	-2.8%*** (0.8%)
<i>N</i>	1,416k	1,416k	116k	116k

Asymmetric shocks or asymmetric responses? In addition to the evidence above, we present statistical moments of a firm’s TFP innovations, TFP growth rates and employment growth rates. If the firm displays asymmetric employment growth and it results from asymmetric TFP innovations, then one should see negatively skewed TFP innovations and employment growth for each firm (on average). For each firm we construct its time-series skewness and then average over all firms.

$$Volatility = \frac{1}{N} \sum_{i=1}^N Vol^i = \frac{1}{N} \sum_{i=1}^N \frac{1}{\tau^i} \sum_{t=1}^{\tau^i} (x_t^i - \bar{x}^i)^2$$

$$Skewness = \frac{1}{N} \sum_{i=1}^N Skew^i = \frac{1}{N} \sum_{i=1}^N \frac{1}{\tau^i} \sum_{t=1}^{\tau^i} \frac{(x_t^i - \bar{x}^i)^3}{(Vol^i)^{3/2}}$$

where τ^i is the number of periods firm i is observed. As reported in Table 2, an important property is that for a typical firm its employment growth skewness is much larger than for the TFP innovations. In fact, the TFP innovations do not appear to be significantly skewed at all while the employment growth rates are negatively skewed. This is consistent with the direct evidence presented in Table 1 above.

Table 2: Time-series volatility and skewness of a typical firm

Skewness	Variable		
	$d \log(TFP_t^i)$	ω_t^i	n_t^i
Unweighted	-0.05	-0.02	-0.18
Employment-weighted	-0.12	-0.04	-0.50

3.3 Implications for the employment distributions across firms

3.3.1 Negative Skewness

The asymmetric decision characterizing our model predicts that in a given year, firms with negative signals respond a lot, but firms with positive signals respond little. Thus we should see that the cross-sectional distribution of employment growth is negatively skewed. To

compute the third moments we use:

$$3^{rd}Moment_t = \frac{1}{N} \sum_{i=1}^N (n_t^i - \bar{n}_t)^3$$

$$Skewness_t = \frac{\frac{1}{N} \sum_{i=1}^N (n_t^i - \bar{n}_t)^3}{\left[\frac{1}{N} \sum_{i=1}^N (n_t^i - \bar{n}_t)^2 \right]^{3/2}}.$$

Figure 5 reports the historical evolution of these measures. We see that skewness is almost always negative with a sample average of -0.4 .⁶

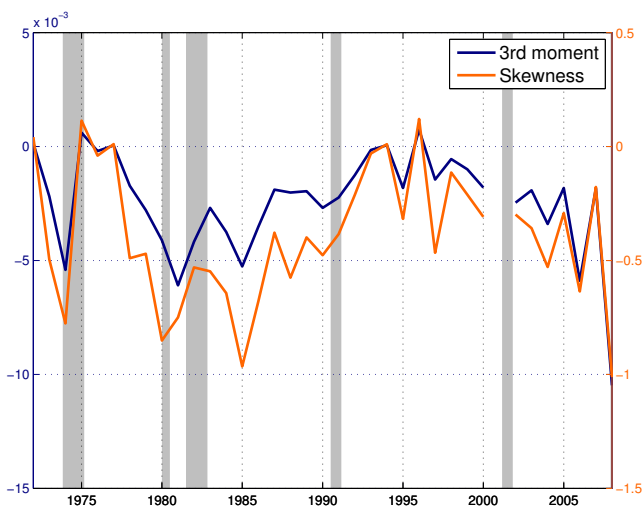


Figure 5: Skewness of cross-sectional distribution of employment growth

3.3.2 Countercyclical Dispersion

The other important moment of the cross-sectional distribution of employment growth is its dispersion. For this we compute the interquartile range, with a sample average of 13%, and find strong evidence of countercyclicity. In particular, being one quarter of the year in a NBER recession increases IQR to 17%. In fully recessionary years, IQR doubles. Figure 6 reports the historical evolution of the IQR and the standard deviation of the cross-sectional distribution of employment growth.

⁶We drop the data in the year 2002 because Census's sampling design changed that year making statistics based on growth rates prone to error.




Figure 6: Dispersion of cross-sectional distribution of employment growth

3.4 Implications for Aggregate Employment

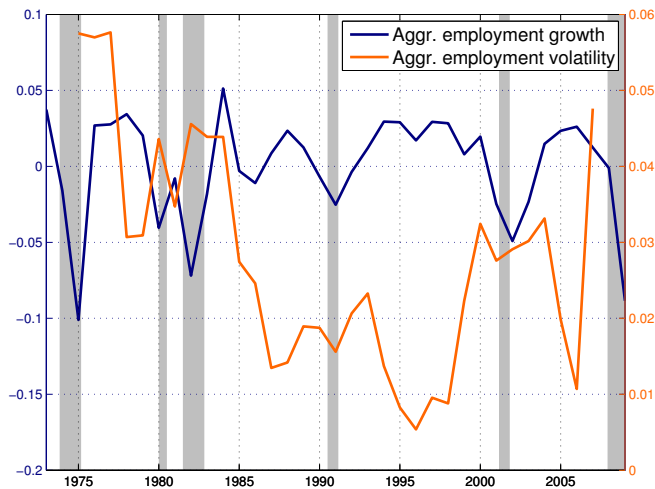


Figure 7: Aggregate Employment growth and volatility

Note: Aggregate employment growth (detrended) and volatility of aggregate employment (computed as 4-year rolling windows of the detrended growth rate)

The model predicted negative skewness in the growth rate of aggregate employment growth and countercyclical volatility in aggregate employment growth. Both are verified in

the data:

$$Skewness = -1.21$$

$$Corr(\text{Volatility, Aggr employment growth}) = -0.23$$

The average growth rate in contractionary years is -3.1% whereas the average growth rate in expansionary years is $+2.3\%$.

3.5 Robustness

3.5.1 Sectoral Specificities and Sampling

Is this all an artefact of manufacturing? After all, this is a declining sector, so negative skewness of establishments, negative skewness in the cross-section and volatility clustering in recessions could all be features of a sector in decline rather than the outcome of asymmetric decision rules. We can check most of the model's prediction in the Longitudinal Business Database (LBD) which covers essentially all active establishments in the U.S. economy. We will now check this much larger dataset (≈ 160 million observations 1976-2011) for the following features:

1. Negative skew of employment growth at the establishment level.
2. Negative skew on average in the year-by-year cross-section.
3. Year-by-year cross-sectional employment dispersion should be countercyclical.
4. Aggregate employment should exhibit negative skew and countercyclical volatility.
5. How do the above data facts look like for entrants/exiters?

We will check for these moments once for the entire LBD (≈ 4.5 m obs/year) to check whether the above data features are specific to manufacturing. We will then check for these moments again for all manufacturing establishments in the LBD (≈ 300 k obs/year) to check whether these moments are specific to the continually sampled establishments in the ASM (only ≈ 50 k obs/year). It turns out, points 1., 3., 4. and 5. do not pose any concerns and point 2. poses only little concern.

1. Negative time-series skewness at the firm level

Table 3: Firm-level skewness in the U.S. economy and manufacturing

	LBD all	LBD mfg only	ASM
Unweighted Skew	-0.10	-0.19	-0.29
Empl.-weighted Skew	-0.04	-0.31	-0.53

Note: The average establishment has negative skew over time that is significantly different from zero. We ran a Kolmogorov-Smirnov test to test for a mean-zero distribution which gets safely rejected.

2. Negative skewness across firms

Things look consistent with our story, but due to the government shutdown the results could not be disclosed in time.

3. Countercyclical employment dispersion across firms

Things look consistent with our story, but due to the government shutdown the results could not be disclosed in time.

4. Aggregate employment skewness and volatility clustering

Things look consistent with our story, but due to the government shutdown the results could not be disclosed in time.

5. Attrition Bias We focus on establishments that are older than 2 years in order to avoid picking up strong growth at the beginning of their life. Furthermore, there is a lot of selection going on in an establishments first life year. How do results look like if we were to extend our analysis to include entrants and exiters? We follow Bloom (2009); Davis et al. (1998) and define a growth rate as

$$\tilde{n}_t^i \equiv \frac{L_t^i - L_{it-1}}{(L_t^i + L_{it-1}^i)/2}$$

This growth rate is bounded between 2 (an entrant) and -2 (an exiter). When we carry out our analysis, we basically find the same patterns; so we omit the results here.

3.5.2 Endogenous Quits

When workers quit firms that did not get a profitability shock and firms do not replace these workers (say, because of labor market frictions), then we would measure a declining employment growth while there were no shocks. So we might think that dispersion is larger

than it actually is and skewness to be more negative than it actually is. Since quit rates are procyclical, this bias would be procyclical too. Our measured dispersion in booms would be hence too wide and skewness too negative which means that our measurements are a lower bound on actual dispersion and skewness caused by productivity shocks.

As for the aggregate time series though, quits are high in a boom when we measure slow growth and low in a recession when we measure strongly negative growth. This could mean that the time-varying volatility and negative skewness we measure in the aggregate time series could well be unmeasured quits firms do not make up for. If we make the extreme assumption that with unchanged technology the economy does not replace workers that leave employment and that no quitting workers take another job in the manufacturing sector, then we could recompute the time-varying volatility and skewness of aggregate employment growth that is corrected for quits. The skewness increases from -1.97 to -1.88 , so even the worst possible case does not overturn our results.

-- put picture of aggr empl growth w/ and w/o quits here --

4 Conclusion

To be done.

References

- Cristina Arellano, Yan Bai, and Patrick J. Kehoe. Financial markets and fluctuations in uncertainty. *Federal Reserve Bank of Minneapolis Staff Report*, April 2010.
- Rüdiger Bachmann and Christian Bayer. Investment dispersion and the business cycle. *Working Paper*, 2011.
- Rüdiger Bachmann, Steffen Elstner, and Eric Sims. Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics*, 5(2): 217–249, April 2013.
- Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, May 2009.
- Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen Terry. Really uncertain business cycles. *NBER Working Paper No. 18245*, July 2012.
- Vasco M. Carvalho and Xavier Gabaix. The great diversification and its undoing. *American Economic Review*, 103(5):1697–1727, August 2013.
- Lawrence J. Christiano, Roberto Motto, and Massimo Rostagno. Risk shocks. *American Economic Review*, forthcoming.
- Steven J. Davis, John C. Haltiwanger, and Scott Schuh. *Job Creation and Destruction*. MIT Press, 1998.
- Andrea L. Eisfeldt and Adriano A. Rampini. Capital reallocation and liquidity. *Journal of Monetary Economics*, 53(3):369–399, April 2006.
- Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, November 2003.
- Larry G. Epstein and Martin Schneider. Ambiguity, information quality, and asset pricing. *Journal of Finance*, 63(1):197–228, February 2008.
- Jésus Fernández-Villaverde, Pablo Guerrón-Quintana, Juan F. Rubio-Ramírez, and Martín Uribe. Risk matters: The real effects of volatility shocks. *American Economic Review*, 101(6):2530–2561, October 2011.
- Jésus Fernández-Villaverde, Pablo Guerrón-Quintana, Keith Kuester, and Juan F. Rubio-Ramírez. Fiscal volatility shocks and economic activity. *Working Paper*, 2012.
- Domenico Ferraro. Asymmetric unemployment dynamics. *Working Paper*, 2013.
- Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.
- Cosmin Ilut. Ambiguity aversion: Implications for the uncovered interest rate parity puzzle. *American Economic Journal: Macroeconomics*, 4(3):33–65, July 2012.

- Cosmin Ilut and Martin Schneider. Ambiguous business cycles. *Working Paper*, 2011.
- Alejandro Justiniano and Giorgio E. Primiceri. The time varying volatility of macroeconomic fluctuations. *American Economic Review*, 98(3):604–641, June 2008.
- Matthias Kehrig. The cyclicity of productivity dispersion. *Working Paper*, 2013.
- Matthias Kehrig and Nicolas Vincent. Financial frictions and investment dynamics in multi-plant firms. *Working Paper*, 2013.
- Christopher A. Sims and Tao Zha. Were there regime switches in U.S. monetary policy? *American Economic Review*, 96(1):54–81, March 2006.
- James H. Stock and Mark W. Watson. Has the business cycle changed and why? *NBER Macroeconomics Annual*, 17:159–218, 2002.
- Stijn van Nieuwerburgh and Laura Veldkamp. Learning asymmetries in real business cycles. *Journal of Monetary Economics*, 53(4):753–772, 2006.
- Laura L. Veldkamp. Slow boom, sudden crash. *Journal of Economic Theory*, 124(2):230–257, 2005. Learning and Bounded Rationality.