NBER WORKING PAPER SERIES

SLOW TO HIRE, QUICK TO FIRE:
EMPLOYMENT DYNAMICS WITH ASYMMETRIC RESPONSES TO NEWS

Cosmin Ilut
Matthias Kehrig
Martin Schneider

Working Paper 20473
http://www.nber.org/papers/w20473

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2014

We thank George-Marios Angeletos, Nick Bloom, Pete Klenow and conference participants at the NBER Summer Institute, SED Annual Meeting, the conference on Causes and Macroeconomic Consequences of Uncertainty, the 2013 Empirical Macro Workshop, the Ifo Conference of Macroeconomics and Survey Data and the Annual Meeting of German Economists Abroad for helpful comments. The views, opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the National Bureau of Economic Research or the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News
Cosmin Ilut, Matthias Kehrig, and Martin Schneider
NBER Working Paper No. 20473
September 2014
JEL No. D2,D8,E20,J2

ABSTRACT

This paper studies the distribution of employment growth when firms adjust asymmetrically to dispersed but correlated signals. If hiring decisions respond more to bad news than to good news, both aggregate conditional volatility ("macro-volatility") and the cross sectional dispersion of employment growth ("micro-volatility") are countercyclical, as in the data. Fluctuations in both macro and micro volatility emerge endogenously in response to news shocks and do not require exogenous changes in volatility. Establishment level Census data confirm other implications of the mechanism: in particular, employment growth (contrary to TFP shocks) is negatively skewed in both the cross section and the time series, and employment responds more to bad than to good TFP shocks.

Cosmin Ilut
Department of Economics
Duke University
223 Social Sciences Building
Box 90097
Durham, NC 27708
and NBER
cosmin.ilut@duke.edu

Matthias Kehrig
University of Texas at Austin
Department of Economics
2225 Speedway, Stop C3100
Austin, TX 78712
matthias.kehrig@gmail.com

Martin Schneider
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305-6072
and NBER
schneidr@stanford.edu

1 Introduction

There is ample evidence of countercyclical movement in both (i) conditional volatility of macroeconomic aggregates ("macro-volatility") and (ii) cross sectional volatility of micro-level variables from which macroeconomic aggregates are constructed ("micro-volatility"). For example, Figure 1 shows that sharp movements in aggregate employment growth around recessions coincide with high dispersion of employment growth across firms measured by the interquartile range.

While the correlation between micro and macro volatility has been documented in recent literature (see e.g. the survey by Bloom (2014)), its causes have not been discussed much. In fact, existing studies often focus on either the macro or micro level in isolation and explore the effects of exogenous micro or macro volatility shocks. To the extent comovement is considered, it is derived from correlated shocks to macro and micro volatilities in fundamentals such as total factor productivity (TFP).

This paper shows that an *endogenous* link between aggregate and cross sectional volatility emerges naturally if firms respond asymmetrically to dispersed but correlated signals about profitability. We then use establishment level Census data to not only verify the behavior of micro and macro volatility but also confirm other predictions of the mechanism, in particular negative skewness of employment growth in the cross section and time series as well as stronger employment responses to bad TFP innovations.

To see how the mechanism works, suppose that the response of employment growth to news about profitability is concave: there is less hiring after good news than there is firing after bad news. Suppose further that aggregate shocks shift the mean of all firms' signals: for example, a spell of bad aggregate shocks generates signals that are on average worse. With concave decision rules, the typical firm's response to its signal during this spell of bad aggregate shocks is then stronger than during a spell of good aggregate shocks.

It follows that both macro and micro volatility of employment growth are countercyclical. Indeed, firms' stronger responses to bad signals generate not only stronger average responses – that is, sharper movements in aggregate employment growth – but also stronger responses to idiosyncratic components in signals and hence higher cross sectional volatility. Importantly, changes in volatility here derive only from firms' endogenous nonlinear responses to mean shifts, not from exogenous changes in volatility.

The paper uses minimal structure to capture the two key elements of the mechanism: a distribution of dispersed but correlated shocks ("signals") and a concave hiring response to those shocks. Our findings are thus compatible with several scenarios for what generates the two elements. One possible signal structure is that firms respond to firm-specific productivity innovations that have both aggregate and idiosyncratic components.[1] Alternatively, firms may receive noisy private signals about profitability. They may care about those signals because they need to fore-

[1]Firms may care about firms-specific productivity for technological reasons, even if they do not need to infer aggregate productivity for decision making. Nevertheless, firm productivity effectively serves as a noisy signal of aggregate productivity under this scenario.

cast profitability; the dispersion of signals then depends in part on the dispersion of noise. The difference between the two scenarios is not important for the basic mechanism.

Similarly, we do not take a stand on why adjustment is asymmetric. One possibility is that the logistics of the hiring process could directly make hiring more costly than firing: hiring new workers might entail costly search, whereas firing is free. At the same time, physical adjustment costs to labor are not necessary for asymmetric adjustment. An alternative story is based on 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.[2]

Asymmetric responses to dispersed signals generate a number of predictions beyond the comovement of macro and micro volatility. To check those predictions, we rely on confidential Census data on U.S. manufacturing establishments. The size of the Census databases allows detailed micro level analysis of moments, controlling for time period, industry and other firm level properties. Figure 1 shows that micro and macro volatility behave as expected; in particular, the inter-quartile range increases by about 20% in the typical recession.

Figure 1: Time-varying employment volatility in the aggregate and in the cross-section

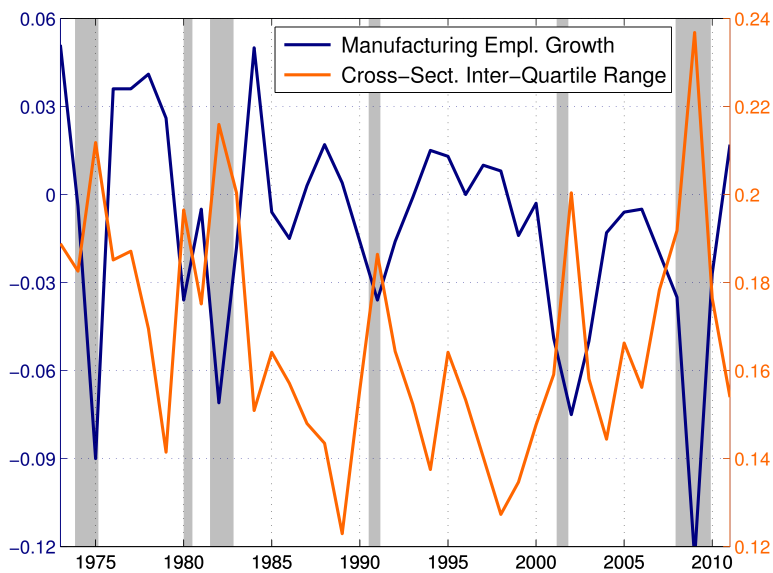

Note: This figure plots aggregate employment growth in the manufacturing sector (left axis in blue) together with the year-by-year inter-quartile range (orange, right axis) of employment growth rates across firms in the sample of manufacturing establishments described in Section 3.

[2]Intuitively, 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.

Our mechanism predicts that employment growth should be negatively skewed in both the cross section and the time series. In the data, the cross section of employment growth is indeed negatively skewed in all years, with an average skewness of -0.48. The asymmetry is significant: the average firm that contracts (expands) employment exhibits a net firing (hiring) rate of 2.2% (1.5%). At the same time, employment growth for the typical firm is negatively skewed in the time series. For example, the employment-weighted average skewness is -0.38; results are robust to different weighting schemes and subsamples.

In principle, negative skewness of employment growth could obtain either because firms respond linearly (that is, symmetrically) to skewed (asymmetric) shocks, or because firms respond asymmetrically to (possibly even symmetric) shocks. One piece of evidence in favor of asymmetric responses is the countercyclical movement of cross sectional dispersion – it would not follow from skewness of shocks alone.[3] More direct evidence can be obtained, however, by measuring specific shocks and estimating firms' employment responses to those shocks.

We estimate firm responses to innovations in TFP, as measured by the Solow residual. If firms respond more (less) to bad (goods) signals about TFP, then bad (good) innovations to TFP should go along with, on average, larger drops (smaller increases) in hiring. To check this, we first construct establishment level Solow residuals and then run both nonparametric and nonlinear parametric regressions of employment growth on TFP innovations, controlling in various ways for other state variables that could be relevant to the firm.

We find that the relationship between employment growth and innovations to the Solow residual is usually concave. On average, a firm faced with a typical negative (positive) TFP innovation – which corresponds to 18% lower (higher) output holding inputs fixed – decreases (increases) employment by 1.8% (1.2%). We find similar asymmetric hiring responses in the vast majority of parametric and non-parametric specifications that we consider, conditioning in particular on size, industry and time period. We conclude that at least one important source of shocks to firms, namely TFP innovations, is propagated via asymmetric adjustment in firm behavior.

The estimated coefficients suggest that concavity in firm responses is large enough to account for sizable movements in volatility. For example, our parametric regressions imply that a one standard deviation shock to average TFP shifts the cross sectional IQ range of employment growth by at least 10%. We cannot expect that shifts in average TFP account for all the movements in micro volatility, since firms in general experience other shocks as well. The results do suggest however that shifts in average TFP innovations can play a quantitatively important role.

The estimated concave responses also contribute significantly to negative skewness in employment growth. In fact, the negative skew in employment growth stands in sharp contrast to the fact that TFP innovations have a nearly symmetric but slightly positively skewed density. It can therefore not be explained by a linear response to skewed TFP shocks. At the same time, the fitted value of employment growth explained by the concave response to TFP exhibits cross sectional

[3] Indeed, if an aggregate shock shifts the mean of a skewed distribution of shocks and the employment response is linear, the volatility of responses need not change at all.

negative skewness of -1.17. The average firm that contracts (expands) employment in response to a TFP innovation exhibits a net firing (hiring) rate of 1.4% (0.8%).

At the industry level, we further show that variation in the degree of concavity for the estimated response functions helps account for differences in the cross sectional negative skewness in employment growth. In particular, one might expect that industries with more concave hiring rules exhibit more negative skewness in employment growth. Among 86 NAICS-4 manufacturing industries, we find that that industries with a near-linear decision rule are almost not skewed at all while skewness drops to -0.6 for those industries with the most concave hiring rules.

The paper is structured as follows. Section 2 illustrates the mechanism using a simple organizing framework. Section 3 introduces the data and describes the distribution of employment growth, with an emphasis on negative skewness. Section 4 turns to the cross sectional relationship between employment growth and TFP innovations.

Related literature There is now a large literature on business cycle models with exogenous shocks to idiosyncratic or aggregate volatility. On the one hand, Bloom (2009); Bloom et al. (2012); Christiano et al. (2014); Arellano et al. (2010); Berger and Vavra (2014); Vavra (2014); Schaal (2012) derive implications of cross sectional risk shocks for business cycles, borrowing and lending, consumption, price dispersion and the effectiveness of policy. On the other hand, Fernández-Villaverde et al. (2011), Gourio (2012) and Basu and Bundick (2011) study the effects of changes in conditional higher moments of aggregate TFP. Our paper does not formulate a full-fledged business cycle model; our goal is to zero in on firms' asymmetric response as a mechanism to connect micro and macro volatility. We also emphasize that our mechanism is not incompatible with the presence of exogenous volatility shocks – instead, we think of it as an amplification mechanism that would be interesting to explore in the context of business cycle models.

While some studies draw connections between micro and macro moments, they look at different stylized facts. For example, Carvalho and Gabaix (2013) show that the sectoral composition of the U.S. economy could have contributed to the Great Moderation, holding fixed the distribution of shocks within sectors. Our paper instead focuses on the business cycle frequency and heterogeneity within industries. Nimark (2014) derives an endogenous link between large (positive or negative) movements in aggregates and the dispersion of survey forecasts from a learning mechanism in which outliers are more salient. While learning could help explain the adjustment behavior we emphasize, it is critical for our mechanism that adjustment be larger for bad shocks, not simply for large shocks.

A recent literature has asked why dispersion in measured productivity is countercyclical. Kehrig (2013) documents countercyclical TFP dispersion and explains it with cyclical entry/exit and overheads in production. Bachmann and Moscarini (2011); Kuhn (2014) show how firms' pricing decisions can lead to countercyclical measured productivity dispersion. Decker et al. (2014) relate countercyclical productivity dispersion to cyclical investment in intangible production knowledge. While our paper shares with these studies an emphasis on endogenous movements in micro volatility, our focus is on employment as well as on the connection between micro and macro moments as well as the role of skewness.

Several recent papers have studied the macroeconomic effects of asymmetric firm decision rules, in particular the resulting asymmetry in the business cycle itself. For example, George and Kuhn (2014) consider costly capacity choice and investment irreversibilities and Ferraro (2013) studies a search and matching framework with heterogeneous workers. They show that asymmetric responses can help explain “deepness asymmetry” (deep recessions vs. meek booms) and “steepness asymmetry” (rapid recessions vs. smooth booms) – concepts introduced Sichel (1993) and studied for aggregate employment by McKay and Reis (2008).[4] Our mechanism is consistent with asymmetry of the business cycle and creates a connection between aggregate asymmetries and micro and macro behavior of volatility and skewness.

2 Macro & micro moments with asymmetric adjustment

To illustrate the basic mechanism, the following minimal framework is sufficient. Consider hiring decisions by a continuum of firms. Every firm receives a signal s about future profitability. Firms' signals can be decomposed into a common component a and an idiosyncratic component ε:

$$s = a + \varepsilon. \tag{1}$$

The idiosyncratic component is independent and identically distributed across firms with mean zero and distribution function G_ε. We can thus think of (1) as representing both the distribution of an individual firm's signal and the cross sectional distribution of signals.

Firms respond to signals about future profitability by changing employment. Assume that all firms follow the same decision rule

$$n = f(s), \tag{2}$$

where n denotes employment growth rate and the function f is smooth, strictly increasing and strictly concave. The assumption of concavity reflects asymmetric adjustment: firms respond less to good signals than to bad signals. We take asymmetric adjustment as given; it can be due, for example, to an asymmetric hiring cost function through which it is more costly to hire than to fire. Appendix B shows that it can be due to information processing, even if there are no adjustment costs.

Our setup does not take a stand on the relationship between the signal and actual profitability and is therefore compatible with several scenarios. One possibility is that there is some true profitability innovation π, say, which itself has aggregate and idiosyncratic components, and firms respond to private noisy signals about their own profitability. In this case, a would contain the

[4] Negative skewness in growth rates of macroeconomic aggregates such as consumption and output is also important in the recent literature on the effect of disasters for asset pricing, following Rietz (1988) and Barro (2006).

aggregate component of π plus correlated noise and ε would contain the idiosyncratic component of π plus uncorrelated noise. The relative share of noise and "fundamental" is not important for our argument. The function f describes both how beliefs are updated given the realized signal s and the optimal employment choice given those updated beliefs.

In general, firm decisions may depend on state variables other than the signal on the next innovation to profitability. For example, firms might respond differently to shocks depending on fixed characteristics like industry, or variable characteristics like size, the level of profitability or the deviation of employment from a target level. For now, these additional features of the firm decision at a point in time are all subsumed in the function f. In other words, the class of firms we study in this section is identical along all these dimensions.

In order to describe common shocks to firms, we condition on the aggregate component a. Let $G_n(n|a)$ and $G_s(s|a)$ denote the conditional cdf's of employment growth and signals, respectively, given a. We will refer to high values of a as representing "good times," that is, times when firms on average receive good news about profitability. An implication of our assumptions is that the conditional variance $var(s|a)$ is independent of a, so good times are only reflected in a high mean signal. This is helpful to zero in on the endogenous link between macro and micro volatility that is driven by asymmetric adjustment. We discuss interaction of asymmetric adjustment and exogenous shifts in volatility further below.

2.1 Implications for micro and macro volatilities of employment

The main intuition for countercyclical volatilities can be seen in Figure 2. The top panel plots the concave response function f, with the signal realization on the horizontal axis and employment growth on the vertical axis. The bottom panel shows three densities of signal realizations $g_s(s|a)$, distinguished by a shift in the mean signal. Taking the middle (blue) density as a reference point, a shift to the left (red) density is an arrival of bad times, whereas a shift to the right (green) density is an arrival of good times.

The figure shows how asymmetric adjustment helps make both macro and micro volatility countercyclical. Consider first macro volatility: the solid horizontal lines in the top panel represent mean employment growth for the three densities. They show how bad news generate larger aggregate responses: the change in aggregate growth in response to bad times (the difference between the horizontal red and blue solid lines) is larger than the change in growth in response to bad times (the difference between the horizontal green and blue solid lines).

To illustrate changes in micro volatility, the dotted lines in both panels show interdecile ranges for signals (along the horizontal axis) and employment growth (along the vertical axis in the top panel). The point here is that concavity of the response function accentuates dispersion in signals in bad times while it attenuates it in good times. The countercyclical cross-sectional volatility is easiest illustrated for interquantile ranges, but this property is inherited by other measures of dispersion such as the cross-sectional variance.

Beyond the specific example of Figure 2, the properties of countercyclical macro and micro

Figure 2: Employment growth and signals

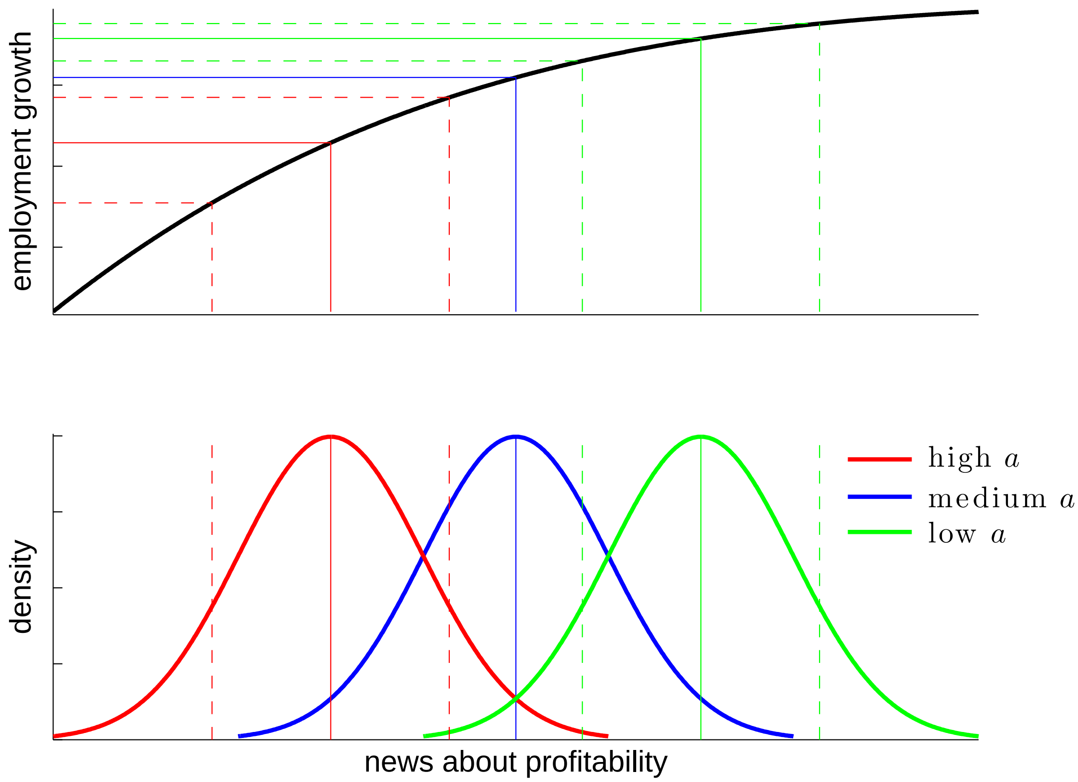

Note: Figure plots the (homoskedastic) distribution of signals about future profitability (bottom panel), once centered around a negative aggregate mean signal (red), once around a neutral aggregate signal (blue) and once for a positive aggregate signal (green). The top panel displays how the concave hiring rule (black) transforms symmetric and homoskedastic signals into asymmetric and heteroskdastic employment responses.

volatility require only the signal structure and concavity of the response function. We summarize this result in the following proposition, with proofs detailed in Appendix A:

Proposition 1 *(Macro & micro volatilities) For any two aggregate shock realizations $a < a'$,*

1. *the sensitivity of the aggregate action with respect to the aggregate shock is higher at a:*

$$\left.\frac{d}{d\tilde{a}} E\left[n|\tilde{a}\right]\right|_{\tilde{a}=a} > \left.\frac{d}{d\tilde{a}} E\left[n|\tilde{a}\right]\right|_{\tilde{a}=a'},$$

2. *the cross sectional variance is higher at a:*

$$var\left(n|a\right) > var\left(n|a'\right),$$

3. *the interquantile range for any two quantiles $\underline{x}$ and $\overline{x}$ is higher at a:*

$$G_n^{-1}\left(\overline{x}|a\right) - G_n^{-1}\left(\underline{x}|a\right) > G_n^{-1}\left(\overline{x}|a'\right) - G_{\Delta e}^{-1}\left(\underline{x}|a'\right).$$

Figure 3 shows the connection between micro and macro volatilities in simulated time series data. The red dotted line is a particular sequence of aggregate news a_t, drawn from a symmetric, homoskedastic distribution. For each realization of a_t, we compute the model implied aggregate employment growth $E\left[f\left(a_t+\varepsilon\right)|a_t\right]$ – shown as the solid blue line – and the interquartile range of the cross sectional distribution $G_n^{-1}\left(.75|a\right) - G_n^{-1}\left(.25|a\right)$, shown as the solid green line. The latter two (blue and green) lines move against each other, reflecting the countercyclicality of micro volatility: the interquartile is wide when aggregate employment growth is low.

Moreover, comparison of the red and blue lines shows that the latter has larger movements when employment growth is low, whereas the movements are quite similar when employment growth is high. Asymmetric adjustment thus translates homoskedastic shocks into heteroskedastic responses – if we measured the variance of aggregate employment growth over subsamples, we would obtain larger numbers in low growth periods. At the same time, the figure shows how asymmetric adjustment translates symmetric shocks into negatively skewed responses. This is another general property of our setup to which we turn next.

2.2 Implications for micro and macro skewness of employment

Figure 2 above suggests that, if asymmetric adjustment connects micro and macro volatilities, it also induces skewness in the cross section and the time series. In particular, the distribution of employment growth responses should be more negatively skewed than the underlying distribution of signals. We define skewness in the standard way as a ratio of third and second moments. For a random variable x, we write

$$\gamma\left(x\right) = \frac{E\left[\left(x - E\left[x\right]\right)^3\right]}{var\left(x\right)^{\frac{3}{2}}}. \tag{3}$$

Figure 3: Time-varying volatility and negative skewness in simulated data

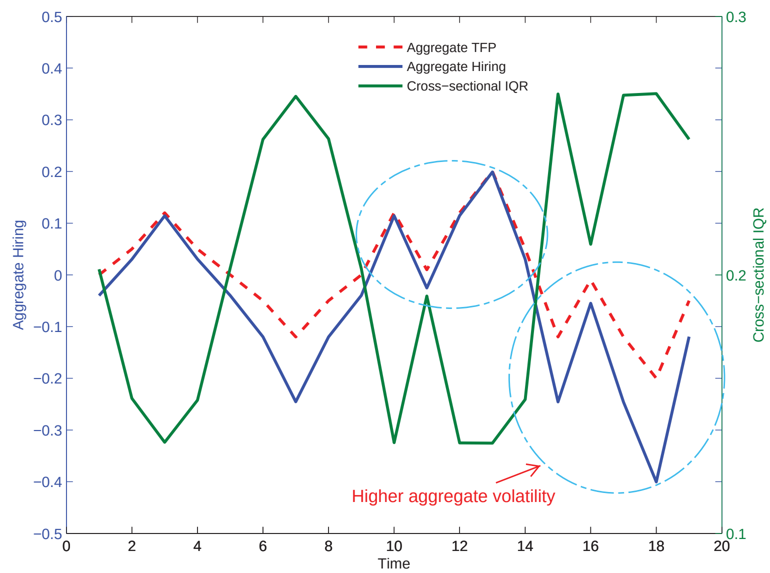

Note: This figure displays aggregate employment growth and the cross-sectional inter-quartile of employment growth in simulated data. Downturns in aggregate employment are deep while upturns are mild thus creating negative time-series skewness. In recessions the cross-sectional dispersion widens and, as highlighted by the blue circles, aggregate volatility is higher than in booms.

The following proposition states formally that a concave response function induces skewness.

Proposition 2 (Micro and macro skewness)

1. *For any aggregate shock a the skewness of the cross sectional distribution of employment growth, $\gamma(n|a)$, is lower than that the skewness of the cross sectional distribution of signals, $\gamma(s|a)$.*

2. *The skewness of aggregate employment growth, $\gamma(E[n|a])$, is lower than the skewness of the aggregate signal $\gamma(a)$.*

Proposition 2 makes a statement about skewness *in general*, but it is silent on the cyclical movements in skewness. It is then natural to ask whether asymmetric adjustment implies systematic movements in skewness together with movements in micro volatility. To see that movements in skewness are not implied by concavity alone, consider the example of a negative exponential response function $f(s) = -e^{-s}$ and normally distributed noise. The conditional mean given a is $e^{-a+\frac{1}{2}var(\varepsilon)}$ and the kth centered moment conditional on a is given by

$$E\left[\left(f(s) - E[f(s)|a]\right)^k |a\right] = e^{-ka} E\left[\left(-e^{-\varepsilon} + e^{\frac{1}{2}var(\varepsilon)}\right)^k\right].$$

It follows in particular that skewness as defined in (3) is independent of aggregate news a.

The special feature of the negative exponential response function is that curvature as measured by the coefficient of absolute risk aversion $-f''(s)/f'(s)$ is everywhere the same. The next proposition provides a more general connection between curvature and skewness.

Proposition 3 *(Cyclicality of skewness) For any two aggregate shocks $a < a'$, the skewness $\gamma(n|a')$ of the cross sectional distribution of employment growth at a' is higher than (lower than, equal to) the skewness $\gamma(n|a)$ of cross sectional distribution of employment growth at a if the coefficient of absolute risk aversion $-f''(s)/f'(s)$ is decreasing in s (increasing in s, constant).*

Intuitively, changes in skewness derive from changes in curvature in the relevant range of signals, in contrast to changes in volatility that derive simply from changes in the slope of the response function. It is therefore possible to have countercyclical volatility combined with countercyclical skewness, that is, in bad times the distribution of employment growth can be more dispersed but less negatively skewed. This is true for example if the response function is quadratic and thus exhibits increasing absolute risk aversion. We conclude from these results that our basic mechanism does not imply definite predictions about changes in cross sectional skewness over time.

2.3 Employment growth and innovations to the Solow residual.

The basic mechanism of asymmetric adjustment works whenever there is a distribution of firm-specific signals that shifts in mean over the business cycle. Unfortunately, firms' signals are not directly observable. However, with some extra structure on what shocks firms respond to, we can derive a relationship between those shocks and the distribution of employment growth.

We follow a large literature in focusing on shocks reflected in the firm's Solow residual. In particular, we assume that the detrended Solow residual, denoted by Z_t^i, evolves according to an AR(1) process

$$\begin{aligned} Z_t^i &= \rho Z_{t-1}^i + z_t^i \\ &= \rho Z_{t-1}^i + u_t^a + u_t^i \end{aligned}$$

where the innovation z_t^i has both an aggregate (u_t^a) and an idiosyncratic (u_t^i) component, both with mean zero.

When firms make hiring decisions relevant for date t production, they observe past profitability Z_{t-1}^i as well as a signal about the new innovation

$$s_t^i = u_t^a + u_t^i + v_t^a + v_t^i \tag{4}$$

where the idiosyncratic components are independent conditional on u_t^a. In terms of the notation used above, the aggregate news a_t here corresponds to the common shock u_t^a plus the correlated noise v_t^a, while the idiosyncratic component ε subsumes both the idiosyncratic profitability innovation u_t^i

and the idiosyncratic noise v_t^i. Moreover, the firm's decision will typically take into account Z_{t-1}^i as well as possibly the level of employment (that is, firm size), or other features of the production function captured for example by the firm's industry. Those characteristics are captured by the function f.

Suppose an econometrician runs a nonparametric regression of employment growth n_{t+1} on the innovation to the Solow residual. Suppose further that the econometrician controls for calendar time, the role of industry and firm specific variables, allowing for nonlinear impact of those variables. He will then recover the conditional expectation given the *true* innovation $u_t^a + u_t^i$:

$$g\left(u_t^a + u_t^i\right) = E\left[f(s_t)|u_t^a + u_t^i\right]. \tag{5}$$

The only random variable in the expectation is the idiosyncratic noise v_t^i. Both components of the true innovation are fixed since they are observed by the econometrician, whereas the correlated noise v_t^a is fixed since it is common to all firms.

Suppose that the variance of the noise v_t^i is independent of u_t^a. The proof of Proposition 1, part 1 then implies that the conditional expectation function g is concave (convex, linear) if the response function f is concave (convex, linear). If the econometrician recovers a concave regression line g, he can therefore rule out that the actual response function is linear or convex. A concave regression line is evidence in favor of the asymmetric adjustment that underlies our mechanism. We emphasize that this conclusion does not depend on homoskedasticity of the innovations u_t^i. In particular, it is true even if the variance of u_t^i depends on u_t^a, for example because innovations are more dispersed in bad times.

Suppose now that the conditional variance of signals depends on the aggregate innovation u_t^a. In particular, one might expect that firms receive more precise signals in good times. A second order Taylor expansion of the regression line delivers

$$g\left(u_t^a + u_t^i\right) \approx f\left(u_t^a + u_t^i\right) + \frac{1}{2} f''\left(u_t^a + u_t^i\right) var\left(v_t^i|u_t^a + u_t^i\right).$$

The curvature properties of the average response g now reflect not only those of the decision rule f but also the interaction of asymmetric adjustment (captured by the second derivative f'') and the variance. Nevertheless, several properties of the relationship make looking at g informative. First, concavity of g cannot be due to movements in signal precision alone. Indeed, if f is linear, then g must also be (approximately) linear. Observing concave g is thus evidence of asymmetric adjustment. Second, suppose that the decision rule is quadratic and the variance of noise is decreasing and convex in u_t^a, as would be the case if higher u_t^a increases the number of iid signals observed by the firm. It then follows that a convex (concave) decision rule implies a convex (concave) average response. The converse holds if the curvature properties of the decision rule do not change over the domain. We conclude that measuring a concave average response is thus indicative of a nonlinear and concave decision rule under plausible assumptions.

3 Employment growth in cross section and time series

The concave decision rule illustrated in the previous section implies countercyclical volatility in the cross-section and in the aggregate over time as well as negative skewness in the cross section and over time. In this section, we first introduce our micro data sources and then check volatility and skewness properties for the distribution of employment growth.

3.1 Data sources

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) and the Longitudinal Business Database (LBD). Following the literature, we identify an establishment as an individual decision maker, which from now on we call a firm. 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), various asset data from the the Capital Tables published by the Bureau of Labor Statistics (BLS) and the Fixed Asset Tables published by the Bureau of Economic Analysis (BEA). 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 1972-2011, which allows us to study business cycle properties over six recessions including the "Great Recession" 2008/09. Every year, we observe about 55k establishments which total up to 2.1 million observations. We focus on the establishment as the unit of analysis and use the term "establishment" and "firm" interchangeably as many other papers do. Since both technology and employment are real phenomena taking place at the establishment level and since we do not focus on variables related to organizational matters of the (hierarchically higher) firm, this choice seems natural to us. Most of the information contained in the non-Census datasets (BEA, BLS, NBER-CES) other than the manufacturing data are merely needed to estimate productivity.[5]

We combine the ASM and the ASM portion of the CMF data (identified by establishment type `ET`=0), so that we have a consistent longitudinal panel. By focusing on the ASM portion in all years, we automatically eliminate all administrative observations (identified by `AR`=1) which are imputed off industry means and thus corrupt moments of the distribution we are interested in.

Entry and exit in our data poses another problem. This can be due to economic birth and death of establishments or due to the rotation of the ASM sample in years ending with 4 and 9. Both features are cyclical and thus not only affect the distribution in general, but have the potential to affect the cyclicality of employment dispersion we are interested in. In order to avoid results being driven by cyclical entry and exit, we construct employment growth rates as defined by Davis and Haltiwanger (1990)[6] and confirm our results using their growth rate. Note that this method is

[5] For more details about the primary data and their transformation needed to estimate productivity, see the description in the appendix to Kehrig (2013).

[6] This growth rate as in Davis/Haltiwanger is defined as $n_t^{i\ \mathtt{DH}} = 2(L_t^i - L_{t-1}^i)/(L_t^i + L_{t-1}^i)$ though Bloom et al. (2012) propose a variant where L_t^i is replaced by L_{t+1}^i.

unable to distinguish between establishments that exit the economy or establishments that continue to exit but drop out of the ASM sample.

3.2 Employment growth dynamics

We focus on the changes in a firm's total employment (the sum of production and non-production workers): $n_t^i \equiv \Delta \log(L_t^i)$. This choice is helpful empirically as it is free of any specific metric. We focus on hiring rather than hours worked because the data report employment for all type of workers while hours worked are only reported for production workers. This is not only a fraction of employment but it also reflects how the firm chooses overtime hours relative to normal hours. Furthermore, a firm's employment variables in the Census data are considered of very high quality and there are virtually no missing values.

The asymmetric decision characterizing our model setup predicts that firms with negative signals reduce employment strongly, but firms with positive signals increase employment only slightly. As a consequence, this decision rule makes predictions about moments of the employment growth distribution across firms as summarized by Propositions 1-3: First, employment growth rates should be more spread-out when most firms receive negative signals, i.e. in recessions. Second, employment growth should be negatively skewed on average. Third, this skewness is not necessarily cyclical. To check for these data features, we compute the second and third moments of employment across all firms in a given year and study the properties of the resulting annual time series which are summarized in Table 1.

Table 1: Cross-sectional Moments

Moment	IQR_t	StD_t	$Skewness_t$
Average	0.167	0.252	−0.477
Average in NBER boom	0.158	0.246	−0.495
Average in NBER recession	0.200	0.272	−0.417
$Corr(dE_t, ...)$	−0.580***	−0.622***	−0.182
	(0.143)	(0.085)	(0.126)
$Corr$(#boom qtrs/year$_t$, ...)	−0.517***	−0.336***	−0.135
	(0.097)	(0.127)	(0.112)
$Corr$(#boom qtrs/year$_{t-1}$, ...)	−0.761***	−0.554***	−0.084
	(0.071)	(0.102)	(0.148)
Avge. no. of observations/year		40,400	

Note: Data are averages of those moments of the cross-sectional employment distribution plotted in Figures 4 and 5. dE_t denotes the growth rate of aggregate manufacturing employment, #boom qtrs/year denotes the number of expansionary quarters in the entire U.S. economy as defined by the NBER. Standard errors for the correlation coefficients are computed using a GMM procedure that corrects for heteroskedasticity and autocorrelation as in Newey and West (1987) and is adapted from Hansen et al. (1988). Census disclosure rules require rounding the number of observations to the nearest hundred.

Cross-sectional employment dispersion

Figure 4 displays the year-by-year evolution of the inter-quartile range (IQR) and the standard deviation of the cross-sectional distribution of employment growth. Both of them show that firms differ a lot in their employment changes: on average, the firm at the top quartile grows employment by 16.7% more than the firm at the bottom quartile. The standard deviation across all firms is 25.2%. Both dispersion measures increase significantly in recessions. The inter-quartile range, for example, reaches 20.0% on average around a NBER recession compared to an average 15.7% in boom times.

In the Great Recession the IQR experienced its sharpest increase ever, rising by more than half its normal value before again returning to its long-run average in 2011. The standard deviation exhibits similar patterns. We also confirm that these patterns are present within 4-digit NAICS industries in order to avoid our results being driven by cyclical composition changes between industries with different but constant long-run dispersion. We also confirm that employment-weighted dispersion follows a similar pattern, so we know that outlier firms with very few employees are not driving our results.

Figure 4: Dispersion of cross-sectional distribution of employment growth

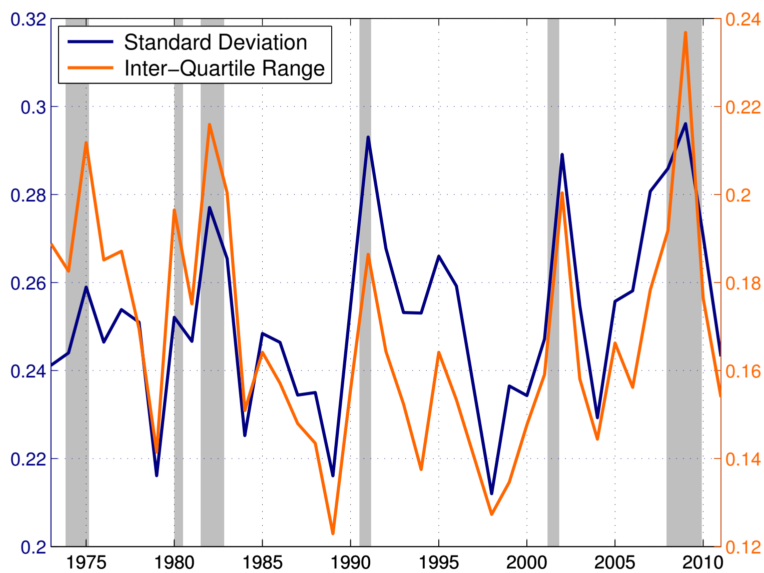

Note: This figure displays the unweighted year-by-year standard deviation and inter-quartile range of employment growth across firms for all t=1973-2011. That is, $StD_t = \sqrt{1/N_t \sum_{i \in t} (n_t^i - \overline{n}_t)^2}$ and $IQR_t = n_t^{q75} - n_t^{q25}$ where $\overline{n}_t$ denotes the mean employment growth, N_t the number of firms and n_t^{qx} the x-percentile. Following common procedure in the literature, we truncate the 1% tails of the overall panel to remove outliers and replace the 1997 dispersion values by the average of its adjacent years. In 1997, sampling switched from SIC to NAICS industry codes which resulted in a reclassification of some establishments into/out of manufacturing thus possibly corrupting cross-sectional moments.

Cross-sectional employment skewness

Our mechanism predicts that the employment distribution should be negatively skewed. Firms that shrink their employment should have an average rate of change that is larger in absolute value than that of firms which grow employment. To check for this data feature, we compute two skewness measures across all firms in a given year and study the year-by-year properties of these time series.

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

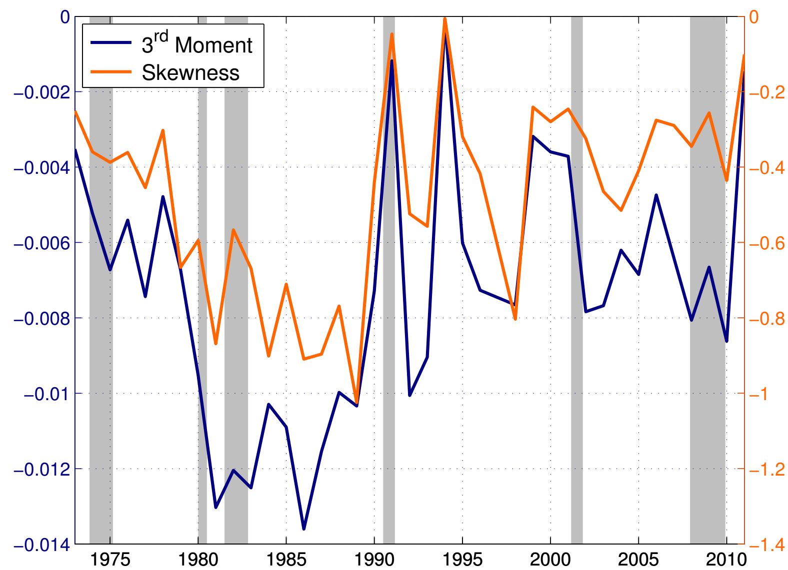

Note: This figure plots the unweighted year-by-year skewness measures across of employment growth across all firms for all t=1973-2011: $3^{rd}Moment_t = \frac{1}{N}\sum_{i=1}^{N}(n_t^i - \overline{n}_t)^3$ and $Skewness_t = 3^{rd}Moment_t/StD_t^3$. Further details as described in the notes to Figure 4.

Figure 5 displays the historical evolution of the third centered moment and the coefficient of skewness:

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

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

The measure $Skewness_t$ is preferred because it is dimensionless in comparison to the measure $3^{rd}Moment$ which depends on the average variance. We see that both measures are negative throughout the sample, and the average coefficient of skewness is -0.477. This skewness is significantly negative and can be interpreted using moments from the employment distribution displayed in Table 1. Using these moments we simulate a cross-sectional distribution of employment growth in an average year. An average skewness of -0.477 is then consistent with the average hiring firm expanding employment by 1.5% and the average firing firm reducing employment by 2.2%. The negative cross-sectional skewness is a new fact and a key supporting element of our asymmetric hiring rule. As we did above with the dispersion measures, we confirm that no outlier industry or firms at the tails drive the negative employment skewness.

Neither measure displays clear cyclical behavior. While Table 1 suggests that skewness is slightly countercyclical, the correlation coefficient with aggregate employment is not significantly different from zero. Through the lenses of Proposition 3, the slight countercyclicality of skewness hints at the shape of the hiring function: it suggests that the coefficient of risk aversion $-f''/f$ is slightly increasing in the signal. For example, a concave linear-quadratic function would deliver such a cyclicality. The plausible presence of other unobserved shocks that influence hiring, even if also associated with concave decision rules but of different shapes, may contribute to a complex cyclical pattern of the cross-sectional skewness.

A firm's employment dynamics over time

Our mechanism also implies that employment growth at the individual firm level should be negatively skewed over time. We thus construct each firm's time-series skewness and then average over all firms to obtain the time series skewness of a "typical firm:"

$$\begin{aligned} FirmSkewness &= \frac{1}{N}\sum_{i=1}^{N} Skew^i \\ &= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\tau_T^i}\sum_{t=\tau_1^i}^{\tau_T^i}\frac{(n_t^i-\overline{n}^i)^3}{(Vol^i)^{3/2}} \end{aligned} \tag{6}$$

where $\tau_1^i, \tau_2^i, ..., \tau_T^i$ indicate the periods where firm i is observed and Vol^i is the time series variance of the firm's employment growth. Panel A of Table 2 shows that employment growth of a typical firm is indeed negatively skewed, employment contractions (relative to the long-run average growth rate of the firm) are larger in absolute value than are employment expansions. This pattern is particularly true for larger firms as the employment-weighted measure shows.

The second column of Table 2 shows that our panel comprises about 150k firms for which we have enough observations to compute their time-series skewness. Naturally, data limitation due to sampling and entry/exit have more bite when one tries to construct longitudinal moments of individual firms. As a robustness check we limit our attention to a strongly balanced panel that contains those 1,900 firms which both exist and are sampled continually from 1972 until 2009.[7]

If the ASM sample overrepresents young firms relative to old ones and the young firms tend to grow smoothly while older ones shrink in a volatile fashion, then this could bias our firm skewness measure towards negative infinity and vice versa for the opposite patterns. Panel B shows that

[7] We limit the sample to 2009 here since we want to compare the moments to those of TFP innovations which we can only construct until 2009 due to data limitations.

continually sampled firms are even more negatively skewed than those in the full panel. This suggests that the ASM sampling happens to miss those (few) years in a firm's life when it contracts a lot. Again, the employment-weighted skewness of the typical firm is more negative than the unweighted skewness confirming our above conclusion that larger firms tend to exhibit a slightly more asymmetric pattern as small firms.

Table 2: Time-series skewness of a typical firm

Firm Skewness of...	n_t^i	No. Firms
A. Full Sample		
Unweighted	-0.187	149,800
Employment-weighted	-0.379	149,800
B. Balanced Panel		
Unweighted	-0.386	1,900
Employment-weighted	-0.554	1,900

Note: Panel A. of this table displays the average time-series employment skewness of a firm as described in (6). That is, we first compute the time-series skewness for each of those 149,800 firms for which we have at least five observations. Then, we compute the average across those firm measures to get the time-series skewness for the "typical firm." In order to obtain the employment-weighted measures, we also compute the average employment stock for each firm and weight each firm's skewness measure by its employment share when computing the employment-weighted skewness of the "typical firm." Panel B. displays the same measures except restricted to those 1,900 firms that we observe in all 38 years from 1972-2009. Census disclosure rules require rounding the number of observations to the nearest hundred.

4 The joint distribution of employment and TFP

In this section we show that the firm level employment growth response to TFP shocks is concave. We first construct innovations to the Solow residual, a measure of TFP, and then run regressions of employment growth on those innovations, using both nonparametric and parametric specifications that control for a battery of other variables. Finally, we evaluate the quantitative significance of the estimated concavity and TFP innovations for the properties of employment growth distributions.

4.1 Solow residuals

Everything described in this subsection is done at the industry level to take into account that long-run productivity growth and productivity persistence differs across industries. Unless otherwise noted, we consider industries at the 4-digit NAICS level. This allows for enough specificities to account for heterogeneity, but leaves us with enough observations in each industry to reliably estimate our conjectured asymmetry relationship and related distributional moments.

Deriving Solow residuals

We begin by deriving the Solow residual for every firm i and year t from the following standard Cobb-Douglas production function (in logs):

$$y_t^i = sr_t^i + \beta^k k_t^i + \beta^l l_t^i + \beta^m m_t^i + \beta^e e_t^i$$

where y_t^i is production (sales corrected for inventory changes and resales), k_t^i, l_t^i, m_t^i, e_t^i are real inputs of capital (structures and equipment), hours worked (production and non-production hours), materials and energy, respectively, and sr_t^i is the Solow residual. The production elasticity of production input $X = k, l, m, e$, labeled β^X, is equated to the revenue share of factor X. This is the only step where we define an industry at the 6-digit NAICS level, so the β's are specific to the 6-digit industry.

Several advantages make this approach very suitable in our context. First, it is fairly free of structural assumptions and thus very general: we only need to assume that firms maximize profits and take factor prices as given. Using more involved structural estimates such as Olley and Pakes (1996) and others would require us to make timing assumptions about the arrival of information and the choice of inputs. Some of these would conflict with our own setup where choices are based on current signals that herald future productivity.

Identifying long-run growth

We assume that Solow residuals contain an aggregate growth trend, a common and a firm-specific fixed effect and a stationary component:

$$sr_t^i = gt + \overline{A} + \alpha^i + Z_t^i \tag{7}$$

where g denotes the long-run (industry-specific) growth rate, $\overline{A}$ initial technology level (which will show up as long-run productivity differences across industries), α^i a firm-specific fixed effect such that $\sum_i \alpha^i = 0$. The distribution of Z_t^i over time is assumed to be stationary and have mean zero. Ultimately, we will be interested in TFP shocks, that is, innovations to Z_t^i. Given the assumptions on Z_t^i, we can identify g and $\overline{A}$ as follows

$$E\left[d\overline{sr}_t\right] = E\left[g + d\overline{Z}_t\right] = g$$
$$E\left[\overline{sr}_t\right] = gt + \overline{A} + E\left[\overline{Z}_t\right]$$

where $\overline{sr}_t \equiv N_t^{-1} \sum_i sr_t^i$ is the average Solow residual in t, $d\overline{sr}_t$ its growth rate, $\overline{Z}_t \equiv N_t^{-1} \sum_i Z_t^i$ is the common cyclical component of technology and N_t and is the number of firms in year t.

Identifying TFP innovations

To identify technology shocks, we assume that Z_t^i follows an AR(1) over time: $Z_t^i = \rho Z_{t-1}^i + z_t^i$.[8]

[8] We have also examined whether or not the stochastic productivity component follows a random walk with drift: $sr_t^i = \alpha_i + Z_t^i = \alpha_i + g + Z_{t-1}^i + z_t^i$ but the estimated residuals are clearly autocorrelated which suggests this model is misspecified.

After detrending the Solow residual in equation (7) and imposing the assumption of an AR(1) of Z_t^i, we get:

$$\begin{aligned} X_t^i \equiv sr_t^i - gt - \overline{A} = \alpha^i + Z_t^i &= \alpha^i + \rho Z_{t-1}^i + z_t^i \\ X_t^i &= \alpha^i(1-\rho) + \rho X_{t-1}^i + z_t^i \\ &= \alpha^i(1-\rho) + \rho X_{t-1}^i + u_t^a + u_t^i. \end{aligned} \tag{8}$$

Then we can estimate the objects α^i, ρ and z_t^i in equation (8) with a panel fixed effects regression.[9]

Moments of the such-obtained technology shock measure can be found in Table 3. The standard deviation of a TFP shock is 0.179 and so is, by coincidence, the inter-quartile range; this implies that an establishment with a 1 standard deviation technology shock produces almost 20% more output than the mean establishment. Productivity dispersion in our data are a bit lower than what Syverson (2011) reports as we focus on TFP *innovations* and also trim the 1% tails in the distribution. The distribution of TFP innovations is almost not skewed at all. Employment growth is negative on average (the manufacturing sector shrank over time): a manufacturing firm in the U.S. reduced its employment by 1.2% per year. Dispersion and negative skewness have been discussed above.

Table 3: Summary Statistics

Moment	TFP Innovation z_t^i	Employment Growth n_t^i
Mean	0	-0.012
Std. Dev.	0.179	0.258
IQR	0.179	0.172
IDR	0.414	0.491
Skewness	0.065	-0.453
Observations	$1,536k$	$1,536k$

Note: Data moments for employment n_t^i and TFP innovation z_t^i of the pooled panel 1973-2009. Slight differences between the employment dispersion and skewness measures in this table and those plotted in Figures 4 and 5 are due to pooling both firms and year observations in this table.

4.2 Asymmetric shocks or asymmetric responses?

Given our estimated TFP shocks, we ask to what extent negative skewness of employment growth is simply due to skewness in those shocks. In principle, skewed TFP could generate skewed employment growth even if firms' decision rules are linear. This subsection shows, however, that this possibility can be ruled out based on both cross section and time series evidence.

[9] The fixed effect causes a well-known bias in the estimates of (8). Monte Carlo studies suggest the bias with 37 periods should be rather small. To check for robustness of our results, we also estimate (8) as suggested by Arellano and Bond (1991) and Blundell and Bond (2000); the resulting measure of TFP innovations is not only extremely similar to the one we obtain, the resulting employment growth-TFP innovation relationship is equally asymmetric; these results are detailed in the Appendix C, see in particular Table 8.

First, consider the summary statistics of cross sectional TFP innovations and employment growth displayed in Table 3. The distribution of TFP innovations is not skewed negatively; the coefficient of skewness is positive and close to zero. In fact, the density of TFP innovations is close to symmetric, as shown by the histogram in Figure 6.

Table 4: Time-series skewness of a typical firm

Firm Skewness of...	Variable z_t^i	n_t^i	No. Firms
A. Full Sample			
Unweighted	+0.047	−0.187	149,800
Employment-weighted	+0.027	−0.379	149,800
B. Balanced Panel			
Unweighted	+0.037	−0.386	1,900
Employment-weighted	+0.041	−0.554	1,900

Note: This tables displays the same time series skewness averaged across all firms as described in the notes to Table 2. The TFP innovation z_t^i is defined in (8).

Second, Table 4 compares the skewness of firm level TFP and employment in the time series dimension. It shows that the TFP innovations, z, do not exhibit negative skewness. At the same time, employment is clearly negatively skewed. We see this as evidence that the observed negative skewness of a firm's employment growth does not result from TFP innovations that are negatively skewed over time.

4.3 Employment and TFP: non-parametric evidence

We now examine the average response of employment growth to TFP innovations. We start with a nonparametric approach. Its key advantage is flexibility in the shape of the relationship between employment growth and TFP – we only require that the response function is smooth. A disadvantage of a nonparametric approach is that it is hard to control for other state variables that may be relevant to the firm in a similarly flexible way. We address this issue here by running many nonparametric regressions that condition on firm properties as well as time periods. We return to the issue of controls below when we run parametric regressions.

The decision rule $f(\cdot)$ from Section 2 relates employment growth n to a contemporaneous signal s that carries information about future productivity, so employment precedes TFP innovations by one period. To map this setup into Census data, we use the fact that TFP innovations in year t are computed from the flow values (output, materials consumption etc.) during the year, while the employment stock reflects quarterly snapshots for production workers and annual snapshots (in March) for non-production workers. A large portion of employment stock in year t is therefore decided in the previous year $t-1$ based on the signal in $t-1$ which reflects information about the TFP innovation in t. As a result, employment growth in the data should respond to TFP innovations measured in the same period.[10]

Concave responses: an example plot and summary statistics

As a representative example of our nonparametric regressions, Figure 6 displays the non-parametrically estimated employment response given technology shocks of mid-sized firms (union of the fourth and fifth octile) in the machinery industry (NAICS 333) in the years 1985-1990. This is a subsample that accounts for some differences across industries, size and time, yet it leaves enough observations to reliably estimate the employment growth-TFP innovation relationship.

Figure 6: Employment growth and TFP innovations

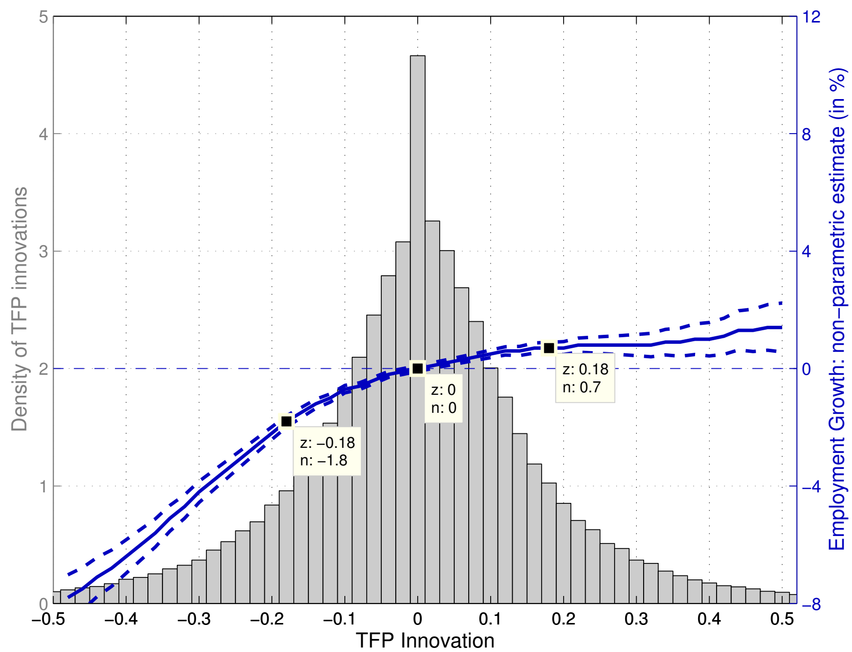

Note: Non-parametric regression of employment growth on TFP innovations for mid-sized firms in the Machinery industry (NAICS 333), years 1985-1990. We condition on size and, a five-year time period and a specific industry to avoid composition driving an observed concavity.
The non-parametric estimate is displayed on the right scale (in dark blue). The solid line displays the mean estimate, dashed lines standard error bands. Employment growth has been cleansed of a time trend. TFP innovations are obtained as described in Section 4.1, and their density is displayed on the left scale (in gray). Indicated data points display employment growth at a typical positive/negative (+1/−1 standard deviation) technology shock relative to no shock.

The estimated regression line corresponds to the function $g(\cdot)$ in Section 2.3, which reflects the recovered employment growth given technology innovation $z_t^i = u_t^a + u_t^i$. The solid dark blue line

[10]One margin which a firm might use to faster adjust its employment to signals are temporary production workers. According to Ono and Sullivan (2013), however, the share of temporary workers makes up less than 5% of the production workers and they find evidence that temporary workers are usually hired to substitute regular production workers when the firm expects permanently lower output in the future.

displays the mean employment change. We can clearly see that employment growth responds more strongly to negative TFP innovations than to positive ones. Furthermore, one can see that the shape of the asymmetry is strictly concave over the domain of TFP innovations.

To quantify the strength of the asymmetry and to report results across many industry size classes and time periods, we use three sets of summary statistics. First, we compare the slope of the nonparametric regression line, at $+1/-1$ standard deviation around the mean. A concave response function exhibits a smaller slope above the mean. For example, in Figure 6, the marginal employment response at a positive standard deviation TFP shock is 2.4% while at a negative standard deviation TFP shock it is 17.5%.

Second, we compare the differences between employment growth at one standard deviation above (below) the mean to mean employment growth. This metric represents the difference between responses to a good and bad shock of "typical size"; it does not rely on local behavior at two specific points. As with the first pair of statistics, a larger value for a positive shock indicates concave response. For example, in Figure 6, a typical positive TFP innovation ($z = +0.179$) leads to about 0.7% more hiring compared to a firm that received no TFP innovation ($z = 0$). This contrasts with a stronger -1.8% employment contraction of a firm after a typical negative TFP innovation ($z = -0.179$). As the standard error bands show, this difference in the average responses is statistically significant at the 95% level.

Our final statistic captures the increase in volatility contributed by nonlinearity in the estimated employment response function g:

$$\phi_g \equiv 1 - \frac{g'(0)^2 \, var(z)}{var\left(g(z)\right)} \tag{9}$$

If the estimated function g is close to linear and $g(z)$ is thus a scaled copy of TFP z, then ϕ_g is close to zero. In contrast, for concave g, the variance of $g(z)$ is larger than that of $g'(0)\,z$, and ϕ_g is positive. For example, Figure 6 implies that $\phi_g = 0.489$; almost half of the variation in $g(z)$ thus comes from concavity of the decision rule.[11] This number is similar to the average $\phi_g = 0.451$ that we obtain below – in this sense, the subsample of firms shown in the plot is a good representation of the asymmetry in the overall population.

Concavity is pervasive across time, industry and size

To show that concavity is a robust fact, we partition the data into subsamples that condition on size, industry and time period, only one control at a time. This procedure limits the subsamples to be analyzed to 37 (= eight time periods + eight size bins + 21 industries) compared to 1,344 ($= 8 \times 8 \times 21$). Table 5 summarizes our findings.

We report the slope of the hiring policy at -1 standard deviation versus the slope at $+1$ standard deviation and the absolute values relative to the mean establishment. If the concave hiring rule was

[11] By itself, the statistic ϕ_g measures the contribution of nonlinearity rather than concavity; for example $\phi_g > 0$ could be generated by convexity in g. At the same time, when the other statistics indicate concavity, ϕ_g is useful to gauge the effect of concavity on volatility.

Table 5: Asymmetry across Time and Size

Sample	(I) Std. Dev. of TFP Innovation	(IIa) Empl. growth elasticity wrt TFP of firm at... −1 StDev	(IIb) +1 StDev	(IIIa) Empl. growth diff. (in %) betw. mean and firm at... −1 StDev	(IIIb) +1 StDev
All	0.179	0.10	0.03	–1.1	0.6
1972-75	0.171	0.10	0.09	–1.4	1.2
1976-80	0.158	0.09	0.06	–1.1	0.8
1981-85	0.169	0.20	0.14	–1.9	0.7
1986-90	0.174	0.11	0.01	–1.0	0.4
1991-95	0.170	0.08	0.02	–0.9	0.4
1996-00	0.186	0.04	0.02	–0.5	0.3
2001-05	0.196	0.03	–0.06	–0.4	0.0
2006-09	0.216	0.03	–0.06	–0.0	0.4
Size Bin 1 (small)	0.206	0.40	–0.30	–4.4	–2.5
Size Bin 2	0.187	0.20	0.01	–1.3	0.3
Size Bin 3	0.183	0.06	0.04	–0.8	0.5
Size Bin 4	0.178	0.06	0.04	–0.7	0.6
Size Bin 5	0.173	0.07	0.05	–0.9	0.8
Size Bin 6	0.170	0.06	0.04	–0.7	0.6
Size Bin 7	0.173	0.03	0.03	–0.5	0.5
Size Bin 8 (large)	0.180	0.01	0.16	–0.4	0.6

Note: Analogously to Figure 6 we non-parametrically estimate detrended employment growth as a function of TFP innovations for each sample denoted in the left column. We then evaluate the slope of the non-parametric estimate at +1 and −1 standard deviation in Columns (IIa) and (IIb). Column (IIIa) displays how much a firm with a negative StD technology innovation shrinks employment relative to a firm without a TFP innovation; column (IIIb) does the same for a firm that experienced a positive standard deviation technology innovation.

Table 6: Asymmetry across Industries

Sample	(I) Std. Dev. of TFP Innovation	(IIa) Empl. growth elasticity wrt TFP of firm at... −1 StDev	(IIb) +1 StDev	(IIIa) Empl. growth diff. (in %) betw. mean and firm at... −1 StDev	(IIIb) +1 StDev	(IV) Asymmetry measure ϕ_g
311 Food	0.185	0.035	0.009	−0.305	0.229	0.847
312 Beverages and Tobacco	0.238	0.046	0.015	−0.335	0.139	0.963
313 Textile Mills	0.157	0.075	0.033	−0.996	0.615	0.405
314 Textile Product Mills	0.183	0.102	0.056	−0.742	0.448	0.642
315 Apparel	0.259	0.023	−0.010	−1.513	−0.076	0.887
316 Leather Products	0.177	0.094	0.016	−0.905	0.564	0.816
321 Wood Products	0.155	0.087	0.040	−1.488	1.082	0.103
322 Paper	0.128	0.046	0.027	−0.232	0.144	0.728
324 Petroleum and Coal	0.174	0.079	0.045	−0.550	0.474	0.520
325 Chemicals	0.216	0.045	0.015	−0.390	0.247	0.544
326 Plastics and Rubber	0.152	0.183	0.097	−0.908	0.753	0.399
327 Nonmetallic Minerals	0.169	0.070	0.015	−0.923	0.557	0.420
331 Primary Metals	0.199	0.178	0.037	−1.135	0.759	0.169
332 Fabricated Metals	0.168	0.165	0.050	−1.420	0.982	0.233
333 Machinery	0.181	0.140	0.078	−1.129	0.887	0.517
334 Computers	0.259	0.043	0.013	−0.654	0.342	0.599
335 Electrical Equipment	0.177	0.097	0.050	−0.802	0.683	0.576
336 Vehicles	0.179	0.118	0.076	−1.403	1.055	0.418
337 Furniture	0.137	0.123	0.096	−1.359	1.015	0.259
339 Miscellaneous	0.177	0.076	0.049	−1.056	0.522	0.325
Raw Average	0.182	0.088	0.037	−0.869	0.515	0.498
Empl.-weighted average	0.189	0.096	0.041	−1.004	0.610	0.503

Note: Analogously to Figure 6 we non-parametrically estimate detrended employment growth as a function of TFP innovations for industry denoted in the left column. We then evaluate the slope of the non-parametric estimate at +1 and −1 standard deviation in Columns (IIa) and (IIb). Column (IIIa) displays how much a firm with a negative StD technology innovation shrinks employment relative to a firm without a TFP innovation; column (IIIb) does the same for a firm that experienced a positive standard deviation technology innovation. Columns (IV) reports our preferred asymmetry measure ϕ_g defined in equation (9). We drop the industry "Printing and Related Support Activities" (NAICS 323) as it is no longer considered a part of manufacturing and sampling after 1997 gets spotty in the manufacturing data.

merely an artifact of composition, we would see symmetric responses in (some) subsamples. But the overall asymmetry of employment growth remains. Table 5 displays the slope of the employment growth function at $-1/+1$ standard deviation technology shock as well as the employment growth level at $-1/+1$ standard deviation technology shock vis-à-vis the mean.

The evidence suggests that with the exception of very large establishments the basic asymmetry is evident in all subsamples and thus largely independent of the chosen controls (time, size, industry). That almost all cuts of the data still exhibit the asymmetric employment growth-TFP innovation relationship is reassuring and informative at the same time. [12] While almost all subsamples exhibit the asymmetry, the quantitative magnitudes of employment responses after positive and negative TFP innovations are different. Mindful of these quantitative differences, the evidence presented in Table 5 points us towards those factors we have to control for when we want to examine this asymmetry more closely.

4.4 Employment and TFP: parametric evidence

The nonparametric approach of the previous section has the advantage of being mostly free from any assumptions about the sign and shape of the relationship. At the same time, the concave shape as displayed for example in Figure 6 can be approximated quite well by a parsimonious parametric form. Doing so also helps to flexibly control for other state variables that could be relevant to the firm without falling prey to the curse of dimensionality or loosing tractability.

In the following, we study the shape of the employment growth-TFP innovation asymmetry and run the following regression

$$n_t^i = \beta^\zeta \zeta_t^i + \beta^{time} t + \beta^Z Z_{t-1}^i + \beta^L l_t^i + c \tag{10}$$

where t is a time trend, Z_{t-1}^i the lagged TFP level (reflecting our assumption of TFP following an auto-regressive process), l_t^i the logarithm of employment and c a constant. ζ_t^i is a function of the technology innovation aimed at a capturing any possible asymmetries. We examine the following specifications

$$\beta^\zeta \zeta_t^i = \begin{cases} \beta^1 z_t^i + \beta^2 (z_t^i)^2 & \text{Specification (I): concave,} \\ \beta^1 z_t^i + \beta^2 (z_t^i)^2 + \beta^3 (z_t^i)^3 & \text{Specification (II): non-monotonic asymmetry,} \\ \beta^1 z_t^i + \beta^2 z_t^i \mathbb{I}\left\{z_t^i < 0\right\} & \text{Specification (III): piece-wise linear,} \\ \beta^1 z_t^i + \beta^2 (z_t^i)^2 \mathbb{I}\left\{z_t^i \geq 0\right\} & \text{Specification (IV): linear-concave.} \end{cases}$$

Specification (I) implies that employment growth is increasing and concave (increasing and convex) in TFP innovations if $\beta^1 > 0, \beta^2 < 0$ ($\beta^1 > 0, \beta^2 > 0$), specification (II) allows for more flexibility in fitting non-monotone relationships, specification (III) assumes a linear relationship,

[12] Appendix C further shows that an asymmetric hiring rule is not an artifact of composition effects caused by firm life cycle patterns or by labor-saving capital-embodied technological change.

Table 7: Employment Asymmetry – parametric specifications

Regressions (I/V): $n_t^i = \beta_0 + \beta^i + \beta_1 t + \beta_2 Z_{t-1}^i + \beta_3 l_t^i + \beta_4 z + \beta_5 z^2$

Regressions (II/VI): $n_t^i = \beta_0 + \beta^i + \beta_1 t + \beta_2 Z_{t-1}^i + \beta_3 l_t^i + \beta_4 z + \beta_5 z^2 + \beta_6 z^3$

Regressions (III/VII): $n_t^i = \beta_0 + \beta^i + \beta_1 t + \beta_2 Z_{t-1}^i + \beta_3 l_t^i + \beta_4 z + \beta_7 z \mathbb{I}\{z < 0\}$

Regressions (IV/VIII): $n_t^i = \beta_0 + \beta^i + \beta_1 t + \beta_2 Z_{t-1}^i + \beta_3 l_t^i + \beta_4 z + \beta_8 z^2 \mathbb{I}\{z > 0\}$

	A. Unweighted Regression				B. Employment-weighted Regression			
Explanatory variables	(I) quadr.	(II) cubic	(III) piece-wise linear	(IV) neg. linear pos. quadr.	(I) quadr.	(II) cubic	(III) piece-wise linear	(IV) neg. linear pos. quadr.
z	0.0827***	0.0959***	0.0347***	0.1148***	0.0775***	0.1005***	0.0373***	0.1108***
	(0.0019)	(0.0027)	(0.0034)	(0.0025)	(0.0053)	(0.0093)	(0.0147)	(0.0051)
z^2	–0.0998***	–0.0956***			-0.0934***	–0.0881***		
	(0.0062)	(0.0063)			(0.0190)	(0.0198)		
z^3		–0.0901***				–0.1685***		
		(0.0163)				(0.0440)		
$z \times \mathbb{I}\{z < 0\}$			0.0947***				0.0795***	
			(0.0057)				(0.0217)	
$z^2 \times \mathbb{I}\{z \geq 0\}$				–0.1849***				–0.1980***
				(0.0106)				(0.0254)
Trend	–0.0033***	–0.0033***	–0.0033***	–0.0033***	–0.002***	–0.002***	–0.002***	–0.0020***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Z_{t-1}^i	0.0282***	0.0278***	0.0281***	0.0276***	0.0162***	0.0154***	0.0161***	0.0156***
	(0.0014)	(0.0014)	(0.0014)	(0.0014)	(0.0034)	(0.0035)	(0.0034)	(0.0034)
l_t^i	0.2965***	0.2966***	0.2966***	0.2966***	0.1906***	0.1906***	0.1906***	0.1905***
	(0.0020)	(0.0020)	(0.0020)	(0.0020)	(0.0148)	(0.0148)	(0.0148)	(0.0147)
Constant	–1.2887***	–1.2884***	–1.2853***	–1.2888***	–1.2397***	–1.2405***	–1.2380***	–1.2396***
	(0.0386)	(0.0385)	(0.0386)	(0.0385)	(0.1099)	(0.1100)	(0.1108)	(0.1093)
Sample	ASM	ASM	ASM	ASM	ASM	ASM	ASM	ASM
N	1,536k	1,536k	1,536k	1,536k	1,501k	1,501k	1,501k	1,501k
R^2	0.168	0.168	0.167	0.168	0.1258	0.1259	0.1257	0.1259
Pos. Resp.	+1.16%	+1.36%	+0.62%	+1.46%	+1.06%	+1.38%	+0.64%	+1.32%
Neg. Resp.	−1.80%	−1.97%	−2.31%	−2.05%	−1.61%	−1.90%	−2.01%	−1.91%
Diff. sign.?	Yes***	Yes***	Yes***	Yes***	Yes***	Yes**	Yes***	Yes***

Note: *, **, *** significantly different from 0 at the 10%, 5%, 1% level, respectively.

but has potentially different slopes for positive and negative innovations, specification (IV) is linear for negative innovations and concave for positive ones. We omit a specification where negative innovations enter in concavely and positive ones linearly as that could be thought of as the difference between specifications (I) and (IV).

For each specification (I)-(IV) we run a fixed effects panel regressions to account for persistent firm-specific factors in hiring. Since the non-parametric analysis above suggested that larger firms have a less asymmetric hiring policy, we also run a set of the same panel regressions (I)-(IV) weighting each observation by its average employment. This should tell us if the quantitative relevance of asymmetric hiring goes away if we consider such a "relevance-based" hiring policy.

Table 7 displays the estimates of a FE panel regression. As expected, TFP innovations have a positive effect on hiring and negative innovations reduce hiring. Ignoring the non-linear terms, a typical positive (negative) TFP shock increases (reduces) employment growth by 1.5 percentage points. But we do find considerable evidence for hiring asymmetries: Across all specifications, employment contractions after negative TFP shocks are larger in absolute value than expansions after a similarly sized positive TFP shock. In the first three rows, all estimates of the non-linear terms are statistically significant. To assess if an actual employment expansion is significantly different from an employment contraction, we evaluate the typical employment response at $+1/-1$ standard deviation TFP innovation: $\beta^{\zeta} \times \pm StD(\zeta_t^i)$. The rows "Pos. Resp." and "Neg. Resp." at the bottom of Table 7 display this typical employment response. For example, the unweighted estimates for specification (I) imply an employment response after a positive TFP innovation of $0.0827 \times 0.18 - 0.0998 \times 0.18^2 = 1.16\%$.

We label responses as significantly asymmetric at the X% level if the point estimate for a positive response lies outside the X% confidence interval of the symmetric negative response and vice versa. Across all specifications we observe that the hiring asymmetry is significant at least at the 95% level and that the hiring asymmetry is quantitatively relevant: the typical negative response is at least 1.5 times as strong in absolute value as the typical positive response and varies from -1.8% to -2.3%. The fit of all specifications seems fairly similar as suggested by a similar $R2$.

The employment-weighted regressions overall imply the same results. The hiring asymmetry is statistically significant and quantitatively almost as sizable as the one implied by the unweighted regression (it's smaller by about a tenth).

We conclude by discussing the other controls. We include the lagged TFP level because the same TFP shock has probably a different employment effect at a higher TFP level than at a lower one, especially since our regression setup considers the more persistent effects captured by changes in the employment growth rate. The lagged TFP level does enter positively which confirms out intuition. Employment matters positively for hiring which reflects the fact that large firms tend to hire more in general. As the negative coefficient on the time trend suggests, the already negative employment growth rate in the manufacturing sector is accelerating over time – probably a consequence of increased outsourcing of manufacturing jobs abroad.

4.5 Comparing skewness and concavity across industries

Given estimated hiring functions for each 4-digit NAICS industry, we now ask whether those functions relate to cross-sectional skewness at the industry level. In general, it is tricky to interpret the connection between industry moments and concavity of the TFP response. Indeed, the TFP response reflects only the response to TFP whereas industries may also differ in the properties of TFP and other shocks. For example, cross sectional dispersion may be higher in some industry not because asymmetry is stronger, but simply because the shocks in that industry are larger.

Nevertheless, there is a sense in which a negative relationship between our asymmetry measure and skewness is interesting evidence. Indeed, suppose our prior belief is that other shocks tend to be symmetric (as is TFP) and that the response to other shocks is concave because downwards adjustments are generally easier than upwards adjustments. Seeing a more concave response to TFP go along with more negative skewness in employment growth then supports our mechanism.

For the industry comparison, we use the 86 NAICS-4 digit industries in our data and focus on ϕ_g as a measure of asymmetry. According to the proposed mechanism, we expect to see ϕ_g vary negatively with cross-sectional "relative skewness," that is, the difference between cross-sectional skewness of employment growth and that of TF innovations. Although TFP innovations are essentially not negatively skewed in any industry, we still want to avoid making inference from employment skewness that could stem from skewness of the underlying shocks.[13] Table 10 in the appendix reports the detailed industry-level asymmetry and cross-sectional relative skewness which we also plot in Figure 7.

The scatter plot confirms our prediction: Symmetric industries (ϕ_g close to 0) have a coefficient of skewness of -0.2 on average while industries with a very asymmetric hiring rule (ϕ_g close to unity) have a coefficient of skewness that is three times as large. In industries at the symmetric end of the distribution, the average hiring (firing) firm increases (reduces) employment by 1.5% (-1.75%) while these numbers change to +1.5% (-2.5%) in the most asymmetric industries. We weigh each industry by its employment share at the beginning of the sample to make sure we consider "relevant" industries equivalently. The correlation between industry asymmetry and industry relative skewness is -0.326 (0.141) which is significant at the 95% level. This negative relationship is robust to dropping outlier industries and choosing different weighting schemes such as the share of shipments rather than employment.

4.6 Quantitative significance

In this section we ask whether shifts in average TFP together with the estimated concave response function from section 4 can generate interesting movements in micro and macro volatility as well as significant negative skewness in cross section and time series. While our focus on TFP limits the analysis to only one of the shocks experienced by firms, and can therefore not account for all fluctuations in volatilities, it is nevertheless interesting to ask whether the shocks and responses we

[13] All our results presented in this section continue to hold when we focus on employment skewness alone.

Figure 7: Industry Evidence: Hiring Asymmetry and Employment Patterns

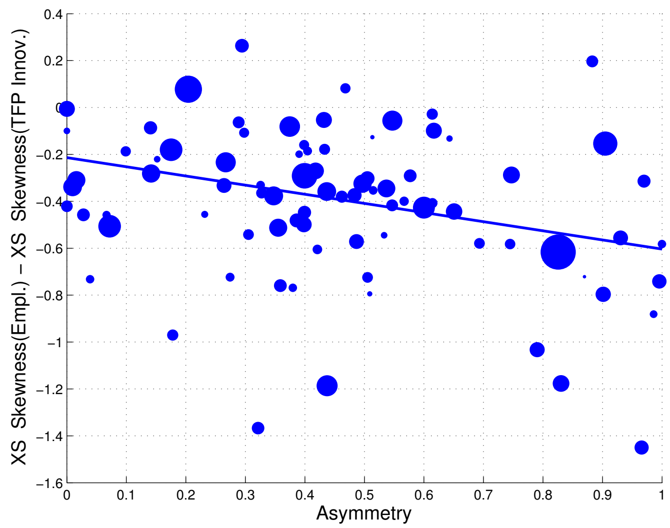

Note: The figure plots the cross-sectional skewness of employment growth relative to that of TFP shocks against the asymmetry of the hiring rule (ϕ_g) within that NAICS-4 digit industry. The correlation coefficient between asymmetry and cross sectional skewness is -0.326^{**} (0.141) in the employment-weighted specification and amounts to -0.284^{***} (0.105) in the unweighted scatter plot.

have identified are quantitatively important. We thus derive ballpark estimates of their contribution to the moments discussed in Section 3.

Notation follows the general setup of Section 2. Let s_t^i denote the signal received by firm i. For the benchmark specification discussed here, we assume that s_t^i simply equals the measured TFP innovation at the firm level, denoted by z_t^i as in Section 4. The effect of noise in signals is considered in Appendix C.4. We show there that additional dispersion across firms due to noisy signals generally enhances the effects of a concave response on volatility movements and skewness; the results here can thus be viewed as a lower bound.

Under the benchmark specification, correlation of signals across firms is due to the aggregate innovation u_t^a, whereas all dispersion is due to an idiosyncratic innovation u_t^i :

$$s_t^i = u_t^a + u_t^i. \tag{11}$$

We assume that the innovations are mutually uncorrelated and normally distributed: $u_t^a \sim \mathcal{N}\left(0, \sigma_a^2\right)$ and $u_t^i \sim \mathcal{N}\left(0, \sigma_b^2\right)$. This functional form assumption simplifies the calculations; it is motivated by the evidence on symmetric TFP distributions in Section 4. Based on the estimated volatilities above, we choose $\sigma_a = 0.045$ and $\sigma_b = 0.1733$.

For the benchmark experiments, we use the linear-quadratic hiring response:

$$f(s_t^i) = 0.0827 s_t^i - 0.0998 \left(s_t^i\right)^2 \tag{12}$$

where the estimated response coefficients are taken from Specification (I) in Table 7. Since s_t^i contains no additional noisy signals on z_t^i, the recovered hiring function indeed corresponds to the underlying $f(s_t^i)$ rather than its conditional expectation. We thus compute moments of the cross-sectional and time series distributions of fitted employment growth $f(s_t^i)$ by feeding the signals through this hiring function.

Countercyclical cross-sectional dispersion

Our estimated concave responses imply that average TFP shocks induce strong countercyclical movements in the cross sectional interquartile range (IQR) of fitted employment growth $f(s_t^i)$. Indeed, a negative one standard deviation shock to average TFP increases the IQR by 11% relative to the mean and by 22% relative to a positive one standard deviation shock. This is same order of magnitude required to explain movement in micro volatility in the data: for example, in periods associated with NBER recessions, the IQR is larger by 25% compared to its average during boom years.

Negative skewness of cross-sectional distribution

The estimated concave response also induces significant negative skewness in fitted employment growth $f(s_t^i)$. The skewness of TFP innovations is zero in this example, consistent with the near symmetry of TFP innovations in the data. In contrast, the coefficient of skewness for employment growth $\gamma(f(s_t^i)|u_t^a)$, as defined in (3), is equal to -1.17 when evaluated at the mean aggregate shock $u_t^a = 0$. As an alternative way of representing the degree of negative skewness, consider average firing compared to average hiring at the mean aggregate TFP innovation:

$$\frac{E(n_t^i|n_t^i < 0, u_t^a = 0)}{E(n_t^i|n_t^i \geq 0, u_t^a = 0)} = -1.73 \tag{13}$$

Average firing is thus more aggressive by a factor of 1.73 than the average hiring. In a model with a linear decision rule, or equivalently with symmetric responses to good and bad signals, the measure in (13) should be equal to -1. In the data we find that the average employment growth decrease is stronger than the average employment increase by a factor of 1.46. Conditional on the measured TFP shocks and the concave hiring rule, the negative skewness of the cross-sectional employment growth is significant and somewhat more pronounced than the empirical one based only on the raw data.

Negative skewness of time series distribution

To evaluate negative skewness in the time series, we simulate for each firm i a series of innovations for u_t^i of length $T = 40$, corresponding to the length of our data. We then compute the time-series skewness of the resulting n_t^i, which for each firm has a mean $\overline{n}^i$ and volatility Vol^i, and take the average of this skewness over $N = 150,000$ firms, the cross-sectional size of our unbalanced panel. We obtain that the average skewness, as defined in formula (6), equals

$$\frac{1}{N}\sum_{i=1}^{N}\frac{1}{T}\sum_{t=1}^{T}\frac{(n_t^i-\overline{n}^i)^3}{(Vol^i)^{3/2}} = -1.12, \tag{14}$$

a value that is more pronounced than the empirical one reported in Table 4, which for the balanced panel, employment-weighted, version of our database equals -0.554.[14]

Appendix C.4 redoes the above calculations under various alternative specifications. In particular, we show that the benchmark specification produces statistics that are more conservative than alternatives in which we add noise to the signal in (11), as in formula (4). The main reason is that the additional noise increases the volatility of the firm-specific signals, which produces stronger effects through the concavity of the hiring rule and thus results in statistics that are larger in absolute value.[15]

To summarize, the implied properties of the cross-sectional and time-series distributions of employment growth, conditional only on the measured dispersed TFP innovations and the estimated linear-quadratic concave hiring rule, are in line with what we document empirically using the raw data for employment growth. This suggests that one mechanism, namely an asymmetric hiring response to dispersed signals, has the quantitative potential to explain these properties jointly and endogenously.

[14] We prefer to compare the model-implied time-series skewness with the balanced panel since we do not account for entry/exit in this quantitative simulation, a factor that seems to be a cause for the weak negative skewness for the full panel reported in Table 4.

[15] We have also explored an alternative piece-wise linear function, as estimated for specification III in Table 7, for which we also find stronger quantitative significance. Appendix C.4 details these various alternative specifications.

References

Cristina Arellano, Yan Bai, and Patrick J. Kehoe. Financial markets and fluctuations in uncertainty. *Federal Reserve Bank of Minneapolis Staff Report*, April 2010.

Manuel Arellano and Stephen Bond. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies*, 58(2):277–297, April 1991.

Rüdiger Bachmann and Giuseppe Moscarini. Business cycles and endogenous uncertainty. *Working Paper*, 2011.

Robert J. Barro. Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866, August 2006.

Susanto Basu and Brent Bundick. Uncertainty shocks in a model of effective demand. *Working Paper*, 2011.

David W. Berger and Joseph Vavra. Consumption dynamics during recessions. *Working Paper*, 2014.

Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, May 2009.

Nicholas Bloom. Fluctuations in uncertainty. *The Journal of Economic Perspectives*, 28(2):153–175, 2014.

Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen Terry. Really uncertain business cycles. *NBER Working Paper No. 18245*, July 2012.

Richard Blundell and Stephen Bond. GMM estimation with persistent panel data: An application to production functions. *Econometric Reviews*, 19(3):321–340, 2000.

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*, 104(1):27–65, January 2014.

Steven J. Davis and John C. Haltiwanger. Gross job creation and destruction: Microeconomic evidence and macroeconomic implications. In Olivier Jean Blanchard and Stanley Fischer, editors, *NBER Macroeconomics Annual*, volume 5, chapter 3, pages 123–168. MIT Press, January 1990.

Ryan Decker, Pablo N D'Erasmo, and Hernan Moscoso Boedo. Market exposure and endogenous firm volatility over the business cycle. *Working Paper*, 2014.

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.

Domenico Ferraro. Asymmetric unemployment dynamics. *Working Paper*, 2013.

Chacko George and Florian Kuhn. Aggregate implications of capacity constraints under demand shocks. *Working Paper*, 2014.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

François Gourio. Disaster risk and business cycles. *American Economic Review*, 102(6):2734–2766, October 2012.

Lars Peter Hansen, John C. Heaton, and Masao Ogaki. Efficiency bounds implied by multiperiod conditional moment restrictions. *Journal of the American Statistical Association*, 83(403):863–871, September 1988.

Matthias Kehrig. The cyclicality of productivity dispersion. *Working Paper*, 2013.

Florian Kuhn. Endogenous profitability risk. *Working Paper*, 2014.

Alisdair McKay and Ricardo Reis. The brevity and violence of contractions and expansions. *Journal of Monetary Economics*, 55(4):738–751, May 2008.

Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, May 1987.

Stephen Nickell. Biases in dynamic models with fixed effects. *Econometrica*, 49(6):1417–1426, November 1981.

Kristoffer P. Nimark. Man-bites-dog business cycles. *American Economic Review*, 104(8):2320–2367, August 2014.

G. Steven Olley and Ariel Pakes. The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297, November 1996.

Yukako Ono and Daniel Sullivan. Manufacturing plants' use of temporary workers: An analysis using Census microdata. *Industrial Relations*, 52(2):419–443, April 2013.

Thomas A Rietz. The equity risk premium: a solution. *Journal of monetary Economics*, 22(1): 117–131, 1988.

Edouard Schaal. Uncertainty, productivity and unemployment during the great recession. *Working Paper*, 2012.

Daniel F. Sichel. Business cycle asymmetry: A deeper look. *Economic Inquiry*, 31(2):224–236, April 1993.

Chad Syverson. What determines productivity? *Journal of Economic Literature*, 49(2):326–365, June 2011.

Sara van de Geer and Marten Wegkamp, editors. *Selected Works of Willem Van Zwet*. Springer, 2011.

Joseph Vavra. Inflation dynamics and time-varying uncertainty: New evidence and an Ss interpretation. *Quarterly Journal of Economics*, 129(1), February 2014.

A Proofs for Section 2

Proof of Proposition 1. *Part 1.* Write the derivative of aggregate employment growth at the point a as

$$\frac{d}{d\tilde{a}}E[f(s)]|_{\tilde{a}=a} = \frac{d}{d\tilde{a}}\int f(\epsilon+\tilde{a})g(\epsilon)d\epsilon \; |_{\tilde{a}=a} = \int f'(\epsilon+a)g(\epsilon)d\epsilon.$$

Since $f'' < 0$, we have that for every realization of ε,

$$a' > a \Rightarrow f'(a'+\epsilon) < f'(a+\epsilon).$$

Part 2. Define the function

$$h(a) := var\left(f(s)|a\right) = E\left[f(s)^2|a\right] - E\left[f(s)|a\right]^2.$$

Then we have

$$\begin{aligned} h'(a) &= 2\left(E\left[f(s)f'(s)|a\right] - E\left[f(s)|a\right]E\left[f'(s)|a\right]\right) \\ &= 2cov\left(f(s), f'(s)|a\right) \end{aligned}$$

which is negative for all a since f is strictly increasing in s and f' is strictly decreasing in s.

Part 3. The conditional cdf of employment growth at some point $\bar{f}$ can be written as

$$\begin{aligned} G_f\left(\bar{f}|a\right) &= \Pr\left(f(a+\varepsilon) \le \bar{f}\right) \\ &= \Pr\left(a+\varepsilon \le f^{-1}\left(\bar{f}\right)\right) \\ &= G\left(f^{-1}\left(\bar{f}\right) - a\right). \end{aligned}$$

The inverse conditional cdf is therefore

$$G_f^{-1}(\bar{x}|a) = f\left(a + G^{-1}(\bar{x})\right).$$

An increase in a now means shifting the pair of points $a + G^{-1}(\underline{x})$ and $a + G^{-1}(\bar{x})$ at which we evaluate f by the same amount; concavity then means the interquantile range shrinks.■

Proof of Proposition 2. The proof uses the following result (Theorem 3.1 in van de Geer and Wegkamp (2011)).

Lemma. Let x denote a random variable, ϕ denote a nonconstant convex function, $\mu_j(y)$ denote the jth centered moment and $\sigma(y)$ the standard deviation of a random variable y. Then for all $k = 1, 2, ...$, provided the moments exist, we have

$$\frac{\mu_{2k+1}(x)}{\sigma^{2k+1}(x)} \le \frac{\mu_{2k+1}(\phi(x))}{\sigma^{2k+1}(\phi(x))}.$$

Part 1. Define $\phi(x) = f^{-1}(x) - a$. If x represents the distribution of s, then $\phi(x)$ represents the distribution of ε. The function ϕ is convex by concavity of f, and the result follows directly from the lemma.

Part 2: Write aggregate employment growth as a function of aggregate news

$$g(a) = E[f(a+\varepsilon)|a].$$

The function g is also concave. Indeed, for any $a \neq a'$,

$$\begin{aligned}
g(\lambda a + (1-\lambda) a') &= E\left[f\left(\lambda a + (1-\lambda) a' + \varepsilon\right)|\lambda a + (1-\lambda) a'\right] \\
&= E\left[f\left(\lambda (a+\varepsilon) + (1-\lambda)(a'+\varepsilon)\right)|\lambda a + (1-\lambda) a'\right] \\
&> \lambda E\left[f(a+\varepsilon)|\lambda a + (1-\lambda) a'\right] + (1-\lambda) E\left[f\left(a'+\varepsilon\right)|\lambda a + (1-\lambda) a'\right] \\
&= \lambda E[f(a+\varepsilon)|a] + (1-\lambda) E\left[f\left(a'+\varepsilon\right)|a'\right] \\
&= \lambda g(a) + (1-\lambda) g\left(a'\right)
\end{aligned}$$

where the last line uses the fact that ε is independent of a.

Define now $\phi(x) = g^{-1}(x)$. If x represents the distribution of aggregate employment growth $E[f(a+\varepsilon)|a]$, then $\phi(x)$ represents the distribution of a. Since ϕ is convex, the result follows directly from the lemma.

Proof of Proposition 3.

Define $y = f(a'+\varepsilon)$ and $x = f(a+\varepsilon)$. The random variable x represents the cross sectional distribution at a, whereas the random variable y represents the cross sectional distribution at a'. Since f is strictly increasing, we can write $y = \phi(x)$, where

$$\phi(x) = f\left(a' - a + f^{-1}(x)\right).$$

By the lemma in the proof of Proposition 2, skewness is higher at a' if ϕ is convex. The first and second derivatives are

$$\begin{aligned}
\phi'(x) &= \frac{f'\left(a' - a + f^{-1}(x)\right)}{f'\left(f^{-1}(x)\right)} > 0 \\
\phi''(x) &= \frac{f'\left(f^{-1}(x)\right) f''\left(a' - a - f^{-1}(x)\right) - f'\left(a' - a - f^{-1}(x)\right) f''\left(f^{-1}(x)\right)}{f'\left(f^{-1}(x)\right)^3}.
\end{aligned}$$

Since $f' > 0$ and $f'' < 0$, we have $\phi''(x) > 0$ if and only if

$$-\frac{f''\left(a' - a - f^{-1}(x)\right)}{f'\left(a' - a - f^{-1}(x)\right)} < -\frac{f''\left(f^{-1}(x)\right)}{f'\left(f^{-1}(x)\right)}.$$

This is true if absolute risk aversion is decreasing everywhere. The relationship between increasing absolute risk aversion and lower skewness follows by reversing the inequalities.

B A model with asymmetric information processing

Here we present a simple model based on information processing under ambiguity that delivers asymmetric hiring decision rules. There is 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 $\overline{w}$. 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

$$z_t^i = u_t^a + u_t^i - 0.5\left(\sigma_a^2 + \sigma_u^2\right)$$

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

$$\exp\left(z_t^i\right)\left(L_t^i\right)^{\alpha} - \overline{w}L_t^i$$

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. The key mechanism through which this simple signaling model generates asymmetries is to assume that firms are ambiguous about the information precision of their signals. In particular, similarly to the model used in Epstein and Schneider (2008), we assume that each firm is ambiguous about the value of σ_e and has a set of beliefs given by

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

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

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

where the minimization operator reflects the firm's ambiguity-aversion. Faced with uncertainty over the signal-to-noise ratio, the ambiguity-averse firm acts *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.

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

$$\max_{L_t^i} \min_{[\underline{\sigma}_{\varepsilon}, \overline{\sigma}_{\varepsilon}]} \exp\left[E^{\sigma_e}\left(z_t^i|s_t^i\right) + \frac{1}{2}var^{\sigma_e}\left(z_t^i|s_t^i\right)\right]\left(L_t^i\right)^{\alpha} - \overline{w}L_t^i. \tag{15}$$

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

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

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

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

The firm problem in (15) then simplifies to

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

The solution to this problem results in a hiring policy that is based on the worst case precision γ_t^* characterized by:

$$L_t^i = \left[\frac{\alpha}{\overline{w}} \exp\left(\gamma_t^* s_t^i\right)\right]^{\frac{1}{1-\alpha}}; \quad \gamma_t^* = \begin{cases} \overline{\gamma} & \text{if } s_t^i < 0 \\ \underline{\gamma} & \text{if } s_t^i \geq 0. \end{cases} \tag{16}$$

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 employment 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 that receives a negative signal $s_t^i = -x$ contracts its employment by more than it would expand it if the firm would have received a positive signal of the same magnitude, $s_t^i = x$.

C Robustness

C.1 Alternative TFP measures

Our preferred measure for TFP shocks – the one used in producing Figures 6 and 7 as well as Tables 5 and 7 – is obtained by estimating the fixed-effects panel model in equation (8). The fixed effect causes a well-known bias in the estimated parameters and TFP innovations. Though Monte Carlo studies suggest the bias in a long panel like ours should be quite small (see Nickell (1981)), we re-estimate equation (8) using the instrumenting proposed by Arellano and Bond (1991) and Blundell and Bond (2000). The resulting measure of TFP innovations is not only extremely similar to the one we obtain in the fixed effects panel regression, the resulting employment growth-TFP innovation relationship is equally asymmetric. If the previously observed asymmetric hiring rule were a mere artifact of biased estimates of TFP innovations, then using the bias-corrected TFP innovations should yield a symmetric estimate of the hiring rule. This re-estimated hiring rule[16] is presented in Table 8 which shows that the hiring rule is still asymmetric and similar to the results presented in Table 7.

C.2 Further analysis of composition effects

The asymmetric hiring rule we estimated previously could be an artifact of composition. We consider two of the most plausible ones: labor-saving capital investment and credit constraints in a firm's early life cycle.

As for the first, the logic runs as follows: firms that receive a positive technology shock might invest in new more productive machinery that requires less labor input thus muting employment growth after positive shocks. But after negative technology shocks firms might simply fire workers and divest machinery thus pronouncing an employment contraction after negative technology shocks. If that was true, then our observed asymmetry would be caused by firms that upgrade their technology. We therefore re-estimate our preferred linear-quadratic specification only for firms that do not simultaneously invest large amounts (which we identify as possibly labor-saving). Following the literature, we define a large investment project as exceeding an investment rate of 0.20. If our

[16]For simplicity we focus on our preferred specifications (I) and (II) from Table 7, but the asymmetry continues to hold in other other specifications.

Table 8: Alternative Measures of TFP Innovations

Regressions (I)/(III): $n_t^i = \beta_0 + \beta^i + \beta_1 t + \beta_2 Z_{t-1}^i + \beta_3 l_t^i + \beta_4 z + \beta_5 z^2$
Regressions (II)/(IV): $n_t^i = \beta_0 + \beta^i + \beta_1 t + \beta_2 Z_{t-1}^i + \beta_3 l_t^i + \beta_4 z + \beta_7 z \mathbb{I}\{z < 0\}$

	Arellano-Bond		Blundell-Bond	
	(I)	(II)	(III)	(IV)
Explanatory variables	quadr.	piece-wise linear	quadr.	piece-wise linear
z	0.0715***	0.0143***	0.0746***	0.0127***
	(0.0018)	(0.0029)	(0.0018)	(0.0029)
z^2	-0.0851***		-0.0896***	
	(0.0035)		(0.0046)	
$z \times \mathbb{I}\{z < 0\}$		0.1162***		0.1186***
		(0.0045)		(0.0046)
Trend	-0.0032***	-0.0032***	-0.0032***	-0.0032***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Z_{t-1}^i	0.0252**	0.0251**	0.0273**	0.0273**
	(0.0015)	(0.0015)	(0.0014)	(0.0014)
l_t^i	0.2972***	0.2972***	0.2963***	0.2963***
	(0.0020)	(0.0020)	(0.0020)	(0.0020)
Constant	-1.3019***	1.2961***	1.2999***	1.2941***
	(0.0389)	(0.0389)	(0.0389)	(0.0389)
Sample	ASM	ASM	ASM	ASM
N	1,536k	1,536k	1,536k	1,536k
R^2	0.1688	0.1688	0.1684	0.1684
Pos. Resp.	+1.42%	+0.39%	+1.34%	+0.33%
Neg. Resp.	−2.68%	−3.56%	−2.57%	−3.44%
Diff. sign.?	Yes***	Yes***	Yes***	Yes***

*, **, *** significantly different from 0 at the 10%, 5%, 1% level, respectively.

asymmetry was a mere artifact of this labor-saving capital upgrade, the asymmetry would disappear. Column (I) in Table 9, however, confirms that the asymmetry is still present, significant and quantitatively close to what we estimated before, even for firms that do not upgrade their capital stock by large amounts.

Second, we want to check if the observed asymmetry is driven by life-cycle dynamics. Young firms that receive positive technology shocks might not hire as many new employees as they want because they probably face credit constraints to dispose over working capital or finance necessary investments that go along an overall expansion. Thus constraints for firms at the early stages of their life cycle that tend to rely more on credit have a muted TFP shock-employment growth relationship. Naturally, credit constraints are one-sided, which means that firms that intend to fire after negative technology shocks need less working capital and capital, so they fire one-for-one. If all asymmetry was driven by such life0cycle patterns that are possibly linked to credit constraints, we should not see any asymmetric hiring for "mid-age" firms, that is firms three years or more after their creation and three years or more before their death.[17] Column (II) in Table 9 shows that the asymmetry does not disappear when focusing on mid-age firms.

We thus conclude that compositional effects driven by a firm's capital goods upgrade and a firm's life-cycle growth and decline are not responsible for the observed asymmetric hiring rule.

[17] We identify death in the data as exit in the Longitudinal Business Database rather than disappearance from our main manufacturing sample.

Table 9: Is Asymmetry Driven by Labor-Saving Capital Investment? By Life-Cycle Patterns?

Regressions (I)-(II): $n_t^i = \beta_0 + \beta^i + \beta_1 t + \beta_2 Z_{t-1}^i + \beta_3 l_t^i + \beta_4 z + \beta_5 z^2$

Explanatory variables	(I) no large investments	(II) mid-age establishments
z	0.0750***	0.0759***
	(0.0019)	(0.0018)
z^2	-0.1415***	-0.1459***
	(0.0064)	(0.006)
Trend	-0.0025***	-0.0022***
	(0.0001)	(0.0001)
Z_{t-1}^i	0.0247***	0.0311***
	(0.0015)	(0.0014)
l_t^i	0.3010***	0.2983***
	(0.0020)	(0.0020)
Constant	-1.3628***	-1.4048***
	(0.0391)	(0.0387)
Sample	ASM $\frac{i}{k} < 0.2$	ASM mid-age
N	1,406k	1,421k
R^2	0.178	0.190
Pos. Response	+0.89%	+0.89%
Neg. Response	−1.79%	−1.82%
Diff. sign.?	Yes***	Yes***

Note: *, **, *** significantly different from 0 at the 10%, 5%, 1% level, respectively.

C.3 Industry-level asymmetry and employment patterns

We redo regression (10) at the four-digit NAICS industry level to recover the asymmetry as in the body of the paper. We then compute the moments of the employment growth distribution for that same industry: long-run dispersion, dispersion difference between boom and recession years and long-run skewness. As displayed in Figure 7, industries with a more asymmetric response to news display higher long-run dispersion, more countercyclical dispersion and more negative skewness.

Table 10: Industry Asymmetry and Industry Employment Patterns

NAICS-4	Asymmetry ϕ_g	Rel. XS Skewness	Employment Share
3111	0.986	–0.881	0.003
3112	0.274	–0.723	0.004
3113	0.067	–0.458	0.004
3114	0.264	–0.333	0.012
3115	0.996	–0.741	0.011
3116	0.830	–1.177	0.015
3117	0.533	–0.545	0.002
3118	0.747	–0.288	0.015
3119	0.614	–0.028	0.007
3121	0.930	–0.556	0.012
3122	1.000	–0.582	0.004
3131	0.321	–1.367	0.008
3132	0.437	–1.186	0.024
3133	0.380	–0.769	0.004
3141	0.693	–0.579	0.006
3149	0.383	–0.079	0.005
3151	0.966	–1.450	0.011
3152	0.826	–0.617	0.066
3159	0.000	–0.100	0.002
3161	0.509	–0.795	0.002
3162	0.790	–1.033	0.012
3169	0.039	–0.732	0.004
3211	0.359	–0.760	0.009
3212	0.305	–0.542	0.006
3219	0.016	–0.310	0.018
3221	0.901	–0.796	0.013
3222	0.547	–0.056	0.022
3231	0.905	–0.154	0.031
3241	0.462	–0.380	0.008
3251	0.651	–0.443	0.015
3252	0.399	–0.447	0.009
3253	0.232	–0.456	0.003
3254	0.289	–0.063	0.008
3255	0.567	–0.400	0.005
3256	0.433	–0.178	0.007
3259	0.547	–0.417	0.008

continued . . .

NAICS-4	Asymmetry ϕ_g	Rel. XS Skewness	Employment Share
3261	0.374	–0.081	0.023
3262	0.617	–0.099	0.014
3271	0.505	–0.724	0.006
3272	0.294	0.264	0.010
3273	0.487	–0.572	0.012
3274	0.513	–0.127	0.001
3279	0.399	–0.159	0.006
3311	0.072	–0.506	0.028
3312	0.515	–0.353	0.004
3313	0.745	–0.582	0.006
3314	0.178	–0.970	0.007
3315	0.355	–0.513	0.018
3321	0.267	–0.219	0.009
3322	0.405	–0.185	0.004
3323	0.437	–0.358	0.019
3324	0.028	–0.458	0.009
3325	0.298	–0.108	0.006
3326	0.326	–0.331	0.004
3327	0.000	–0.006	0.014
3328	0.099	–0.187	0.006
3329	0.267	–0.234	0.022
3331	0.537	–0.345	0.017
3332	0.483	–0.374	0.011
3333	0.000	–0.421	0.007
3334	0.505	–0.302	0.011
3335	0.432	–0.053	0.014
3336	0.141	–0.086	0.010
3339	0.497	–0.324	0.018
3341	0.970	–0.314	0.009
3342	0.408	–0.275	0.013
3343	0.421	–0.605	0.005
3344	0.010	–0.338	0.019
3345	0.600	–0.427	0.026
3346	0.870	–0.722	0.001
3351	0.614	–0.407	0.005
3352	0.577	–0.291	0.009
3353	0.418	–0.270	0.015
3359	0.398	–0.499	0.013
3361	0.347	–0.377	0.020
3362	0.327	–0.365	0.007
3363	0.204	0.078	0.039
3364	0.399	–0.291	0.035
3365	0.390	–0.199	0.003
3366	0.386	–0.482	0.010

continued . . .

NAICS-4	Asymmetry ϕ_g	Rel. XS Skewness	Employment Share
3369	0.643	–0.132	0.002
3371	0.142	–0.281	0.018
3372	0.468	0.082	0.006
3379	0.152	–0.221	0.002
3391	0.883	0.197	0.008
3399	0.175	–0.180	0.028
Average	0.451	–0.409	0.012
Std. Deviation	0.265	0.323	0.010
Raw corr. with ϕ_g	1	–0.283** (0.110)	0.039 (0.120)
Weighted corr. with ϕ_g	1	–0.326**	0.204

C.4 Alternative specifications for quantitative significance

In this appendix we report results based on alternative specifications to the benchmark signal structure and parametric function of equations (11) and (12), respectively. In Table 11 we collect the moments defined in section 4.6. The first two columns give their informal and formal description. Column (M1) reports values for the benchmark specification. The rest of the remaining columns are based on alternative model specifications to be introduced below.

One extension is to allow for noise in the idiosyncratic signal. More specifically, as in equation (4), we write the general representation of the signal as

$$s_t^i = u_t^a + v_t^a + u_t^i + v_t^i$$

where v_t^a and v_t^i are the common and idiosyncratic components of the noise distributed as $v_t^a \sim \mathcal{N}\left(0, \sigma_{v,a}^2\right)$ and $v_t^i \sim \mathcal{N}\left(0, \sigma_{v,b}^2\right)$, respectively. We parametrize each of the variances of the noise through the implied signal to noise ratios

$$\kappa_x = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_{v,x}^2}$$

for both the aggregate and idiosyncratic components, i.e. $x = a, b$. Allowing for noise in the signals means that the econometrician recovers the conditional expectation $g\left(z_t^i\right) = E\left[f(s_t^i)|z_t^i\right]$, as defined in equation (5). We discuss some of the conclusions that can be drawn about $f(s_t^i)$ based on the observed $g\left(z_t^i\right)$ in Section 2.3. In this context however, we maintain for expositional purposes the assumption that $f(s_t^i)$ is still given by the linear-quadratic function in (12), with the same coefficients β^1 and β^2 as those estimated for the estimated $g\left(z_t^i\right) = \beta^1 z_t^i + \beta^2 \left(z_t^i\right)^2$. Thus, in the experiments reported below, allowing for noise amounts to maintaining the same decision rule while increasing the variance of either the idiosyncratic component, $u_t^i + v_t^i$, or the aggregate one, given by $u_t^a + v_t^a$.

In column (M2), we set $\kappa^a = \kappa^b = 0.5$, so that the variance $\sigma_{v,a}^2 = \sigma_a^2$ and $\sigma_{v,b}^2 = \sigma_b^2$. We observe that all statistics increase as we have essentially made both types of shocks more volatile. The larger volatility will have larger effects through the same curvature of the hiring rule. In column (M3), we keep $\kappa^a = 1$, so that there is no aggregate noise, but activate $\kappa^b = 0.5$. We see that the properties of cross-sectional and time series skewness are the same as in column (M2) since the amount of idiosyncratic variance is the same. The cyclicality of the cross-sectional dispersion is very similar to the benchmark case, with a very small difference caused by the larger variance of idiosyncratic noise. Lastly, in column (M4) we revert to $\kappa^b = 1$ and keep $\kappa^a = 0.5$. Similarly to the above logic, the cross-sectional and time series skewness are the same as the benchmark case while the cyclicality of the cross-sectional dispersion is larger.

A final alternative specification that we investigate is to change the hiring decision rule to a piece-wise linear function in which employment growth continues to respond stronger to bad than to good signals. We maintain the benchmark assumption of no additional noise in the signal as in equation (11) and consider the decision rule

$$f(s_t^i) = \beta_{pw}^1 s_t^i + \beta_{pw}^2 s_t^i \mathbb{I}\left\{s_t^i < 0\right\}$$

where the coefficients β_{pw}^1 and β_{pw}^2 equal 0.0347 and 0.0947, respectively, as estimated for specification III in Table 7. Column (M5) reports the results for this case. We see that the main qualitative features are maintained and that the piece-wise linear function in general produces stronger quantitative effects.

Table 11: Quantitative significance of estimated concavity of hiring decision rule

Distribution Moment Description	Statistic	(M1)	(M2)	(M3)	(M4)	(M5)
XS dispersion cyclicality	$\ln\left[\frac{IQR(n_t^i \mid u_t^a=-\sigma_a)}{IQR(n_t^i \mid u_t^a=\sigma_a)}\right]$	22%	32%	23%	31%	36%
XS skewness	$\gamma(n_t^i \mid u_t^a = 0)$	−1.17	−1.55	−1.55	−1.17	−1.01
Average firing and hiring	$\frac{E(n_t^i \mid n_t^i<0,u_t^a=0)}{E(n_t^i \mid n_t^i\geq 0,u_t^a=0)}$	−1.73	−2.15	−2.15	−1.73	−2.72
TS skewness	$\frac{1}{N}\sum_{i=1}^{N}\frac{1}{T}\sum_{t=1}^{T}\frac{(n_t^i-\overline{n}^i)^3}{(Vol^i)^{3/2}}$	−1.12	−1.49	−1.49	−1.12	−0.99

Note Column 1 and 2 give the informal and formal description of the distribution moments, where XS refers to cross-section and TS to time series. Column (M1) the corresponding values for the benchmark specification. Column (M2) refers to the alternative specification with noise, where $\kappa^a = \kappa^b = 0.5$; (M3) with $\kappa^a = 1, \kappa^b = 0.5$ (M4) with $\kappa^a = 0.5, \kappa^b = 1$ and (M5) is the piece-wise linear model without noise.