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

Cosmin Ilut† Duke & NBER

Matthias Kehrig‡ UT Austin

Martin Schneider§ Stanford & NBER

August 1

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. We use establishment level Census data to document skewness properties of employment growth as well as firms' responses to anticipated TFP shocks that are consistent with the mechanism.

*Preliminary and incomplete – comments welcome. We thank 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.

Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

†cosmin.ilut@duke.edu.

‡matthias.kehrig@austin.utexas.edu.

§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 emphasized in the recent literature, its causes have not been discussed much. In fact, existing studies often rely on exogenous volatility shocks and derive comovement of macro and micro volatility from correlated shocks to macro and micro volatilities in fundamentals such as 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 future profitability. We use establishment level Census data to show that asymmetric adjustment not only helps understand the behavior of volatilities, but is also consistent with the skewness properties of employment growth as well as the joint distribution of employment growth and Solow residuals.

To see how our 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 (but not the variance) of all firms' signals. In particular, 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 is thus stronger than during a spell of good aggregate shocks.

It then follows that both macro and micro volatility of employment growth are countercyclical. Indeed, firms' stronger responses to bad signals generates 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 derive only from firms' endogenous nonlinear responses, and not from exogenous changes in volatility.

Asymmetric responses to dispersed signals lead to a number of predictions beyond the comovement of macro and micro volatility. In particular, employment growth should be negatively skewed in both the cross section and the time series. We find that both properties hold in the micro data: in particular, the cross section is negatively skewed across years while time series skewness obtains both at the aggregate level and at the level of the individual firm.

Moreover, asymmetric responses should be apparent in the relationship between employment growth and observable fundamentals. As an example of an observable fundamental, we consider the Solow residual, or measured TFP. If firms respond more (less) to bad (goods) signals about the Solow residual, then bad (good) realizations of the Solow residual should be preceded, on average, by larger drops (smaller increases) in hiring. In line with this prediction, both parametric and nonparametric regressions reveal a robust concave relationship between employment growth and the Solow residual.

The mechanism we study relies on two key elements. First, hiring decisions respond to dispersed signals with a common cyclical component. For example, firms may receive noisy private signals about an aggregate component of profitability (for example future demand conditions). They may care about the signals because they are trying to forecast future profitability. The dispersion of signals then corresponds to the dispersion of noise. Alternatively, firms might respond to firm-specific productivity innovations that have both aggregate and idiosyncratic components, so firm productivity thus serves as a noisy signal of aggregate productivity.[1]

The second element of our mechanism is that firms respond more to bad news than to good news. It does not matter exactly why such asymmetric adjustment occurs. One possibility is that the logistics of the hiring process directly makes hiring more costly than firing: for example, hiring new workers entails costly search, whereas firing is free. A second candidate for asymmetric adjustment comes from information processing: if firm decision makers are averse to Knightian uncertainty (ambiguity) and are uncertain about the quality of signals, then it is also optimal to respond more to bad news.[2]

For our empirical work we rely on confidential Census data on U.S. manufacturing establishments 1972-2009 to measure moments that speak to the presence of our mechanism. We first verify the behavior of micro- and macro-level volatility over the business cycle.

We define recessions using NBER dates. To understand the behavior of micro volatility, we compute the interquartile range (IQR) of annual employment growth rates. The countercyclicality of both aggregate as well as cross-sectional volatility can be seen in Figure 1: the IQR spikes up by roughly 20% in recessions when aggregate employment growth deviates typically by 5.7% from the long-run mean compared to only 2.9% in boom years.

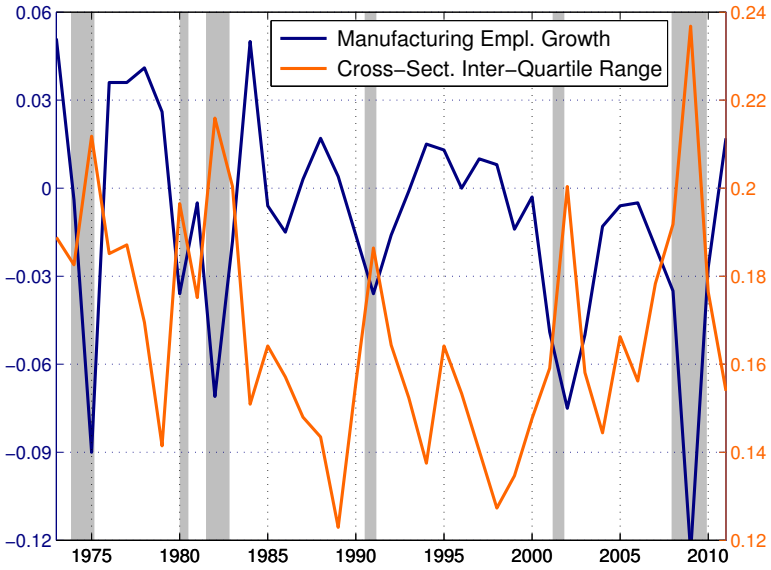
We further consider negative skewness of employment growth both in the cross section and the time series. We show that the cross section of employment growth rates is negatively skewed in all years: net firing by the typical firm that contracts employment relative to the mean is 1.8% compared to net hiring by the typical firm that expands employment relative to the mean which is about 1.16%. At the same time the growth rate of aggregate employment is also negatively skewed. This is again consistent with the idea that firms that receive bad signals lay off lots of workers while firms who receive good signals hire few workers. We emphasize that all aggregate implications are derived in a micro-founded framework which accounts for the composition of firms at the micro level. There are no asymmetries at the macro level that are assumed analogously to what we find at the micro level, so our conclusions about asymmetric aggregates are immune to the cautions of Caballero (1992).

While the above facts on volatility and negative skewness are consistent with asymmetric adjust-

[1] In this case, firms might care about these shocks for technological reasons, even if they do not need to infer aggregate productivity for decision making. The difference is not important for our mechanism.

[2] The information processing argument illustrates that our mechanism does not require adjustment costs to labor. Indeed, ambiguity averse firms evaluate hiring decisions *as if* taking a worst case assessment of future profits. With ambiguity about signal quality, the worst case then depends on what the signal says: for a good signal, the worst case interpretation is that it is noisy, whereas for a bad signal the worst case is that it is very precise. Updating from ambiguous signals thus endogenously generates asymmetric actions.

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.

ment, they could also derive from properties of the underlying shock distributions. To investigate asymmetric adjustment more directly, we thus consider firm responses to movements in the Solow residual, a source of shocks to profitability that has been prominently featured in the literature. We first estimate production functions to measure Solow residuals and then obtain innovations by fitting an autoregressive process with firm fixed effects.

The typical relationship between employment growth and innovations to the Solow residual is concave. We establish this result with both nonparametric and nonlinear parametric regressions, controlling in various ways for other state variables that are possibly relevant to the firm. We further show that employment growth is substantially more negatively skewed than innovations to the Solow residual. In fact, the latter are actually mildly positively skewed. We conclude that at least one important source of shocks to firms is propagated through asymmetric adjustment in firm behavior. Our parametric and non-parametric estimations agree on the quantitative results of typical positive versus negative shocks: firms that are faced with a typical positive TFP innovation, that corresponds to 18% higher output holding inputs fixed, will increase employment by 1.2%. Firms that are faced with a negative TFP innovation of the same size respond much more strongly and reduce employment by 1.8%.

We investigate the main facts relevant for our mechanism not only at the aggregate, but also at the industry level. It is then interesting to compare cross sectional moments across industries.

We find that skewness is correlated with a measure of asymmetric adjustment derived from our nonparametric regressions. These results again suggest that moments of employment growth and their movement is at least in part generated by endogenous asymmetric adjustment.

The paper is structured as follows. Section 2 introduces a general class of signal structures and concave decision rules that translate symmetric signals with time-invariant higher moments in actions for which the aggregate and cross-sectional higher moments can be time-varying. In Section 3 we use confidential Census data on U.S. manufacturing establishments to measure moments for employment growth that speak to the presence of our mechanism.

Related literature Following McConnell and Perez-Quiros (2000) and Stock and Watson (2002), many macroeconomists have asked why the volatility of US macroeconomic aggregates has been changing. Studies that aim to explain medium term changes in volatility such as the "great moderation" typically try to identify underlying changes in the volatility of exogenous shocks. For example, Justiniano and Primiceri (2008); Justiniano et al. (2011) find evidence for changing volatility in the shocks to savings and investment technology while Sims and Zha (2006); Fernández-Villaverde et al. (2012) consider long-run changes in monetary or fiscal policies. Carvalho and Gabaix (2013) point to changes in the sectoral composition of the U.S. economy, holding fixed the distribution of shocks within sectors. Our paper focuses on the business cycle frequency and does not address the great moderation directly; our results suggest, however, that the trends identified by those authors could generate joint changes in macro and micro volatilities.

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; instead, we use minimal structure (firms' asymmetric response) to exhibit a mechanism that connects cross sectional and aggregate moments. It would be interesting to exploit this mechanism also in the context of business cycle models.

Several papers have recently provided explanations for countercyclical movements in the dispersion of measured productivity. 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.

A number of papers have recently looked at the macroeconomic effects of asymmetric firm de-

cision 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 "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). 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 a continuum of firms making employment decisions. Every firm receives a signal s about future profitability. Firms's 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

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

where Δe 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. At this point, we do not take a stand on the reason for asymmetric adjustment. We propose one possibility based on information processing in Section B.

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 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 $G_{\Delta e}(\Delta e|a)$ 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

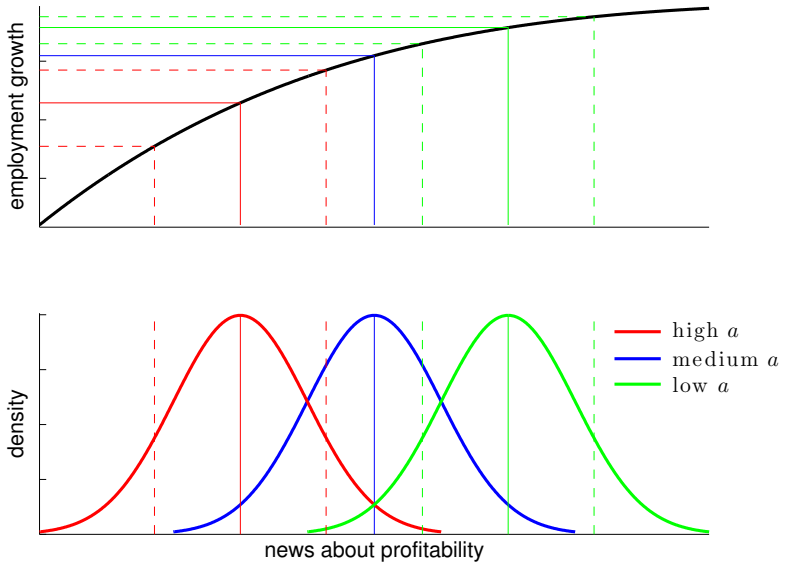
The main intuition for countercyclical volatilities can be seen in Figure 2. The top panel plots the concave response function f, with the signal realization on the horizontal axis and employment growth on the vertical axis. The bottom panel shows three densities of signal realizations $g_s(s|a)$, distinguished 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 volatility require only our signal structure and concavity of the response function. We summarize this result in the following proposition, with proofs detailed in the Appendix:

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.

Proposition 1 (Macro & micro volatilities)

For any two aggregate shock realization $a < a'$,

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

$$\frac{d}{d\tilde{a}} E\left[\Delta e|\tilde{a}\right]\Big|_{\tilde{a}=a} > \frac{d}{d\tilde{a}} E\left[\Delta e|\tilde{a}\right]\Big|_{\tilde{a}=a'},$$

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

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

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

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

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 then 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 $G_{\Delta e}^{-1}\left(.75|a\right) - G_{\Delta e}^{-1}\left(.25|a\right)$, 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 heteroskadastic 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.

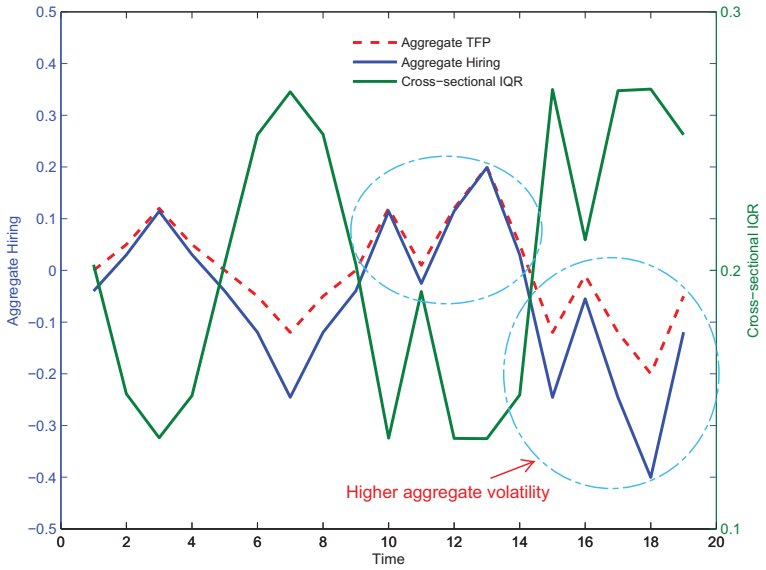
2.2 Implications for micro and macro skewness of employment

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

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

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

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.

Proposition 2 (Micro and macro skewness)

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

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

It is important to note that 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 $\gamma(\Delta e|a')$ of the cross sectional distribution of employment growth at a' is higher than (lower than, equal to) the skewness $\gamma(\Delta e|a)$ of cross sectional distribution of employment growth at 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 simply changes in the slope of the response function. It is therefore possible to have countercyclical volatility combined with procyclical skewness: this is true for example if the response function is quadratic and thus exhibit decreasing 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 is 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 thus 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

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

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 Δe_{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 $u_t^a + u_t^i$:

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

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 of a nonlinear and concave decision rule under plausible assumptions.

3 Employment growth distributions

The decision rule in the previous section showed how asymmetric responses to news lead to countercyclcial volatility in the cross-section and in the aggregate over time as well as negative skewness

in the cross section and over time. We will now check for the empirical counterparts of what we derived above. First, we study the employment dynamics both in the aggregate and across firms. We distillate empirical moments from the data and check for consistency with our asymmetry setup above. Of course, other forces could have caused these moments as well, notably countercyclical and negatively skewed technology shocks. We thus estimate in a second step productivity at the micro level and study its dynamics in a similar way as we did with employment. Third, we show that the relationship between productivity innovations and employment growth is nonlinear and suggests an asymmetric responses on the part of firms. We propose a variety of robustness checks to eliminate the possibility that the TFP-employment asymmetry is a measurement artifact caused by other forces. Finally, we show how asymmetric responses vary across industries and how the predictions about employment skewness vary with it.

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), the Plant Capacity Utilization Survey (PCU) 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.[3]

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

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

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)[4] and confirm our results using their growth rate. Note that this method is 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, eliminates any "fixed effects" specific to the the firm or the industry. 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 is also 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. This makes employment a good object to study possible non-linearities in a firm's policy.

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.

3.2.1 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 deviate 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 and it rose by almost half its normal value in the Great Recession 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 industries (results not disclosed) in order to avoid our results being driven by cyclical composition

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

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)
$Corr(\#\text{boom qtrs/year}_t, ...)$	-0.517^{***}	-0.336^{***}	−0.135
	(0.097)	(0.127)	(0.112)
$Corr(\#\text{boom qtrs/year}_{t-1}, ...)$	-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).

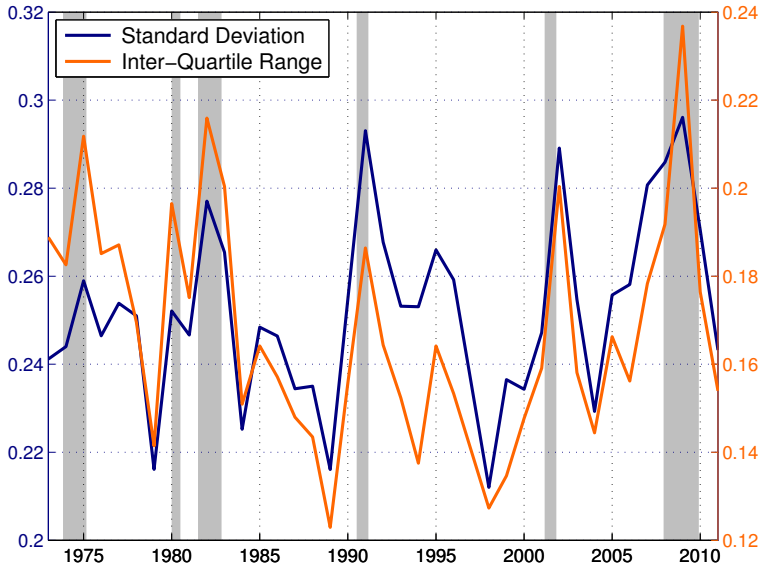
changes between industries with different but constant long-run dispersion. We also confirm that employment-weighted dispersion follows a similar pattern, so we know that outlier firms with very few employees are not driving our results.

3.2.2 Cross-sectional employment skewness

In the previous section we demonstrated how the asymmetric decision rule causes the cross-sectional employment dispersion to be countercyclical. Though this fact is robustly borne out in the data, this and other countercyclical dispersion features have attracted a lot of attention recently and, consequently, alternative explanations. those explanations, however, have no particular prediction about the cross-sectional skewness. Our mechanism, in contrast, 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. As Proposition 2 and 3 mention, skewness should be negative on average, but the cyclicality of skewness is ambiguous and depends on the exact curvature of the concave hiring rule. 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 displays the historical evolution of the third centered moment and the coefficient of

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.

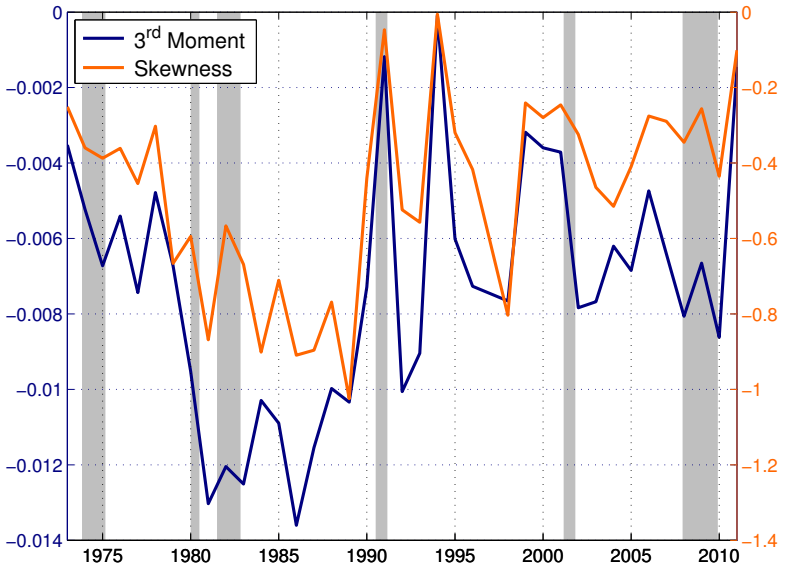
skewness:

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

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

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 3. Using these moments an average skewness of -0.477 is consistent with the top quartile firm growing employment by 6.7% above the median while the firm at the bottoms quartile decreases employment by 9.9% below the median. Both measures are acyclical, with a correlation coefficient with aggregate employment not significantly different than zero. The lack of definite cyclical patterns in the cross-sectional skewness is consistent with the predictions of Proposition 3.

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=1973-2009: $3^{rd}Moment_t = \frac{1}{N}\sum_{i=1}^{N}(n_t^i - \overline{n}_t)^3$ and $Skewness_t = 3^{rd}Moment_t/StD_t^3$. Further details as described in the notes to Figure 4.

As we did above with the dispersion measures, we confirm that no outlier industry or firms at the tails drive the negative employment skewness.

3.2.3 A firm's employment dynamics over time

The previous subsections presented empirical evidence for the implications of an asymmetric hiring function for the cross-sectional employment growth distribution. But the asymmetric decision rule makes additional predictions about a firm's time-series employment growth distribution. In particular, the employment pattern of an individual firm over time should be negatively skewed. As above, this employment pattern is subject to the qualification that any possible time series skewness of a firm could be driven by negatively skewed TFP innovations (this will be checked in the next subsection). In order to check the time series predictions of the asymmetric hiring rule, we construct each firm's time-series skewness and then average over all firms to obtain the time

series skewness of a "typical firm:"

$$FirmSkewness = \frac{1}{N}\sum_{i=1}^{N} Skew^i$$
$$= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\tau_T^i}\sum_{i=\tau_1^i}^{\tau_T^i}\frac{(n_t^i - \overline{n}^i)^3}{(Vol^i)^{3/2}} \quad (4)$$

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. 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; vice versa for the opposite patterns. Panel B shows that 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 tables displays the average time-series employment skewness of a firm as described in (4). 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.

4 Technology dynamics

Of course the employment data facts by themselves are consistent with our model mechanism, but they are far from proving it. We will proceed to present more direct evidence of the asymmetric mechanism. We still chose to present the empirical cross-sectional and firm-level employment dynamics because no theory limits this analysis. Furthermore, the empirical researcher might find these facts interesting in their own right and useful source to match heterogeneous firm models. Further statements will require making more restrictive assumptions about the factors driving a firm's employment choice. We will now proceed along that route in order to establish a more tight link between cause and effect. Since our analysis is at the firm level, we consider profitability – or revenue productivity – as the main driver of a firm's choices such as hiring and firing. This will tell us if firms respond asymmetrically to symmetric profitability shocks or if the shocks themselves are asymmetric. This means we have to identify productivity shocks at the level

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.

4.1.1 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.

4.1.2 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:

$$sr_t^i = gt + \overline{A} + \alpha^i + Z_t^i \tag{5}$$

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

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

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.

4.1.3 Identifying TFP innovations

To identify technology shocks, we assume that Z_t^i follows an AR(1) over time: $Z_t^i = \rho Z_{t-1}^i + z_t^i$.[5] After detrending the Solow residual in equation (5) and imposing the assumption of an AR(1) of Z_t^i, we get:

$$\begin{aligned} X_t^i \equiv sr_t^i - gt - \overline{A} = \alpha^i + Z_t^i &= \alpha^i + \rho Z_{t-1}^i + z_t^i \\ X_t^i &= \alpha^i(1-\rho) + \rho X_{t-1}^i + z_t^i \\ &= \alpha^i(1-\rho) + \rho X_{t-1}^i + u_t^a + u_t^i. \end{aligned} \tag{6}$$

Then we can estimate the objects α^i, ρ and z_t^i in equation (6) with a panel fixed effects regression.[6]

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

[5] 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.

[6] The fixed effect causes a well-known bias in the estimates of (6). Monte Carlo studies suggest the bias with 37 periods should be rather small. To check for robustness of our results, we also estimate (6) 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; see Table 7.

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?

Equipped with the estimated technology shocks, we relate the employment dynamics across firms and over time to the movements in respective technology. We start by looking at the summary statistics of TFP innovations displayed in Table 3. Evidently, the distribution of TFP innovations is not skewed negatively; the coefficient of skewness near zero reflects a fairly symmetric technology distribution – a feature reflected in Figure 6 below.

Table 4: Time-series skewness of a typical firm

Firm Skewness of...	Variable			No. Firms
	dsr_t^i	z_t^i	n_t^i	
A. Full Sample				
Unweighted	-0.042	$+0.047$	-0.187	149,800
Employment-weighted	-0.105	$+0.027$	-0.379	149,800
B. Balanced Panel				
Unweighted	-0.067	$+0.037$	-0.386	1,900
Employment-weighted	-0.115	$+0.041$	-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. dsr_t^i is the rate of change of the Solow residual as defined in (5) and z_t^i the TFP innovation as defined in (6).

Next, we want to relate firm-level moments of employment and technology over time. Firm employment can be negatively skewed over time if employment is negatively skewed as well. In that case, firms would symmetrically follow the exogenous shocks. As reported in Table 4, neither the growth rate of Solow residuals, dsr, nor TFP innovations, z, exhibit negative skewness which

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

In the previous sections, we have confirmed the cross-sectional and firm-level implications of the asymmetric decision rule for moments of the employment distribution. We have further shown that the moments of primitive technology shocks do not exhibit the same patterns as moments from the employment distribution.[7] We now proceed from these indirect implications to more direct evidence for our conjectured asymmetric hiring rule. To do so, we estimate a firm's employment growth on its technology. Since we do not want to impose any parametric assumption on this analysis, we choose to estimate this relationship non-parametrically. That way, we do not impose any restrictions on the particular shape of the relationship between technology shocks and employment changes and let the data speak for themselves. The goal of this exercise is to study whether the employment-asymmetry relationship exists and how it looks like in particular. A concave hiring rule could be piece-wise linear, logarithmic or quadratic to name some example, and we are interested in the particular shape of asymmetry because it matters for the quantitative analysis we plan to carry out in Section XXX.

The non-parametric approach comes with a significant disadvantage: There are possibly many factors influencing employment growth other than technology shocks. One may think that a similar technology shock leads to a different firm hiring response in different industries. Young firms might hesitate more to expand their employment stock compared to older firms with more experience in the marketplace. Ideally, one would partition the data into many homogeneous subgroups and estimate this relationship for each group. Naturally, the curse of dimensionality would make this endeavor barely feasible (the number of regressions grows exponentially in controls) and impossible to disclose (some of these homogeneous subgroups would just have too few observations to avoid disclosure, esp. in the tails). After studying many variations and subgroups, we chose to disclose the non-parametric regression of the mid-sized firms (union of the fourth and fifth octile) in the machinery industry (NAICS 333) in the years 1985-1990. This subsample accounts for differences across industries, size and time, yet it leaves enough observations to reliably estimate the employment growth-TFP innovation relationship.

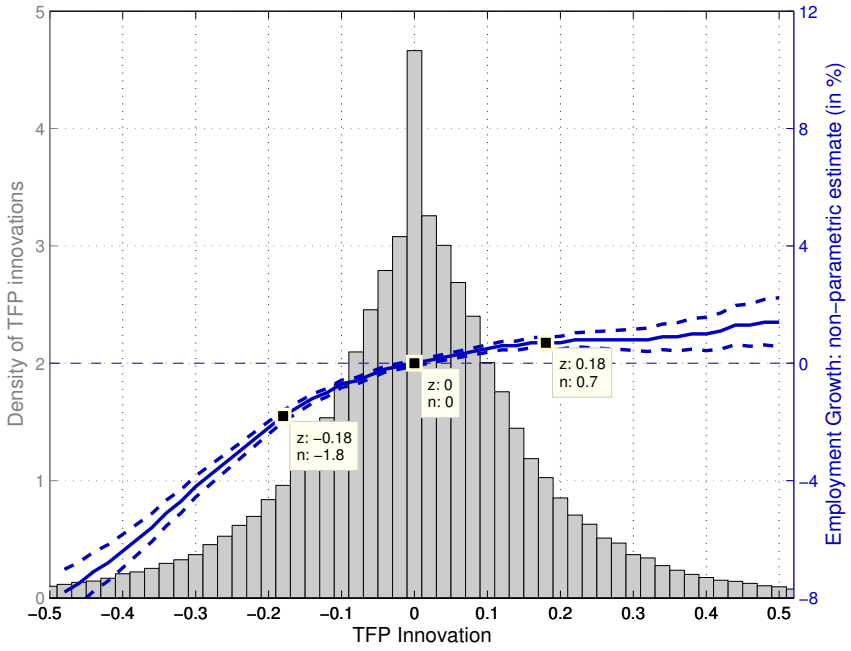
Implementing our information-theoretic framework empirically requires us to think about the timing of variables: The decision rule $f(\cdot)$ posits that employment growth n is chosen based on contemporaneous signals s which carry information about *future* productivity, so employment precedes TFP innovations by one period. Unfortunately, mapping this setup into Census data is ambiguous. TFP innovation 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. So a large portion of employment

[7] That the dispersion of technology shocks is countercyclical is the only exception.

stock in t is actually decided in the previous period $t-1$ based on the signal in $t-1$ which reflects information about the TFP innovation in period t. Under this assumption employment growth in the data should respond to TFP innovations measured in the *same period*.[8]

4.3.1 The asymmetric 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.

Figure 6 displays the non-parametrically estimated employment response given technology shocks in the above described subsample. This 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 displays the mean employment change. We can clearly see that employment

[8]We are fully aware that this mapping is not perfect. 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.

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 rather than piece-wise linear.

To quantify the strength of the asymmetry, we will study three metrics: First, the average employment response at $+1/-1$ standard deviation. Second the sensitivity of the employment response at $+1/-1$ standard deviation. Third, the share of predicted employment changes explained by non-linear terms in $g(\cdot)$.

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. Thus, the average employment response is stronger for negative shocks. Second, employment changes are more sensitive for negative TFP innovations. The marginal employment response at a positive standard deviation technology shock is 2.4% while it amounts to 17.5% at a negative standard deviation technology shock. Lastly, we introduce a metric that tries to capture the asymmetry in the estimated employment response function $g(\cdot)$. We represent g by a first-order Taylor expansion

$$g(z) = g(0) + g'(0)\, z + H$$

where H collects higher-order terms. Therefore, we can derive the share of fluctuations in $g(z)$ that is not explained by linear terms as ϕ_g:

$$var\left(g(z)\right) = \left[g'(0)\right]^2 var(z) + var(H) + 2g'(0) cov\left(z, H\right)$$
$$\phi_g \equiv 1 - \frac{\left[g'(0)\right]^2 var(z)}{var\left(g(z)\right)} = \frac{var(H) + 2g'(0) cov\left(z, H\right)}{var\left(g(z)\right)}$$

From the estimated function g we can recover $var\left(g(z)\right)$, $g'(0)$ as well as $var\left(z\right)$ and therefore $var\left(H\right)$ and $cov\left(z, H\right)$. If the function g is linear, then H is zero and the asymmetry measure is $\phi_g = 1$. So by looking at ϕ_g we are getting a measure of the "excess" volatility that is due to the presence of higher-order effects.[9] In the above estimation $\phi_g = 0.489$ which means that almost half the variation in $g(z)$ comes from asymmetries in the estimated decision rule.

4.3.2 Asymmetry is pervasive across time, industry and size

Figure 6 is meant to be a representative slice of the data which we cannot study in all its fullness using non-parametric methods. It is still instructive in revealing the nature of relationship between employment growth and TFP innovations. We found this relationship to be concave, but many factors could cause the observed asymmetry in Figure 6. Ideally, we would want to look at this figure for a set of establishments that share the same characteristics along any dimension that

[9] It is true that this excess volatility could also come from convexity; it is really only saying "what is contributed by higher order terms in decision rule."

may matter for employment growth; for example, size, time, age, industry, overall TFP level (as opposed to TFP innovation), etc. But such controlling for various states of the firms would require a number of subsamples that grows exponentially in the number of controls. To avoid the curse of dimensionality we redo the same non-parametric regression by conditioning on time only (in five-year intervals), then by condition on size only (grouping establishments into eight equally numerous size groups), then by conditioning on industry only. Table 5 summarizes our findings.

Table 5: Pervasive asymmetry: Non-parametric evidence across subsamples

	(I)	(IIa)	(IIb)	(IIIa)	(IIIb)
	Std. Dev. of TFP	Empl. growth elasticity wrt TFP of firm at...		Empl. growth diff. (in %) betw. mean and firm at...	
Sample	Innovation	–1 StDev	+1 StDev	–1 StDev	+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
311 Food	0.176	0.035	0.009	–0.4	0.3
312 Beverages and Tobacco	0.261	0.046	0.015	–0.2	0.2
313 Textile Mills	0.162	0.075	0.033	–1.0	0.6
314 Textile Product Mills	0.174	0.102	0.056	–0.7	0.5
315 Apparel	0.157	0.023	–0.010	–1.3	–0.1
316 Leather Products	0.161	0.094	0.016	–0.8	0.6
321 Wood Products	0.167	0.087	0.040	–1.3	1.0
322 Paper	0.130	0.046	0.027	–0.2	0.1
323 Printing and Related	0.139	0.033	–0.029	–0.2	–0.1

continued . . .

	(I)	(IIa)	(IIb)	(IIIa)	(IIIb)
	Std. Dev.	Empl. growth elasticity		Empl. growth diff. (in %)	
	of TFP	wrt TFP of firm at...		betw. mean and firm at...	
Sample	Innovation	–1 StDev	+1 StDev	–1 StDev	+1 StDev
324 Petroleum and Coal	0.173	0.079	0.045	–0.6	0.5
325 Chemicals	0.193	0.045	0.015	–0.4	0.2
326 Plastics and Rubber	0.163	0.183	0.097	–0.8	0.7
327 Nonmetallic Minerals	0.171	0.070	0.015	–0.9	0.6
331 Primary Metals	0.196	0.178	0.037	–1.0	0.7
332 Fabricated Metals	0.164	0.165	0.050	–1.5	1.0
333 Machinery	0.170	0.140	0.078	–1.1	0.9
334 Computers	0.225	0.043	0.013	–0.6	0.4
335 Electrical Equipment	0.163	0.097	0.050	–0.8	0.6
336 Vehicles	0.167	0.118	0.076	–2.3	1.6
337 Furniture	0.133	0.123	0.096	–1.0	0.8
339 Miscellaneous	0.181	0.076	0.049	–0.6	0.3

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. The industry "Printing and Related" (NAICS 323) is in smaller font as it is no longer considered a part of manufacturing.

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 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. 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 parametric analysis allows up to flexibly add these controls without loosing tractability. Once we understand which controls are relevant in the employment growth-TFP innovation relationship, we

In the following, we want to study the shape of the employment growth-asymmetry relationship and run the following regression

$$n_t^i = \beta^\zeta \zeta_t^i + \beta^{time} t + \beta^Z Z_{t-1}^i + \beta^L \log(L_t^i) + c \tag{7}$$

where t is a time trend, Z_{t-1}^i the lagged TFP level, $\log(L_t^i)$ the logged 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

$$\beta^\zeta \zeta_t^i = \begin{cases} \beta^1 z_t^i + \beta^2 (z_t^i)^2 & \text{Specification (I)} \\ \beta^1 z_t^i + \beta^2 (z_t^i)^2 + \beta^3 (z_t^i)^3 & \text{Specification (II)} \\ \beta^1 z_t^i + \beta^2 z_t^i \mathbf{I}\left\{z_t^i < 0\right\} & \text{Specification (III)} \\ \beta^1 z_t^i + \beta^2 (z_t^i)^2 \mathbf{I}\left\{z_t^i \geq 0\right\} & \text{Specification (IV).} \end{cases}$$

Specification (I) implies that employment growth is concave in TFP innovations (assuming $\beta^1 > 0, \beta^2 < 0$), specification (II) allows for more flexibility fits, specification (III) assumes a linear relationship, but has potentially different slopes for positive and negative innovations, specification (IV) is linear for negative innovations and concave for positive ones.

The rows "Pos. Resp." and "Neg. Resp." at the bottom of Table 6 evaluate the typical employment response at +1/–1 standard deviation TFP innovation: $\beta^\zeta \times \pm StD(\zeta_t^i)$. Say the unweighted estimates for specification (I) imply an employment response after a positive TFP innovation of $(0.0827 - 0.097 \times 0.18) \times 0.18 = 1.17\%$.

Across all specifications we observe that the employment does in fact respond asymmetrically to TFP innovations: the typical negative response is at least 1.5 times as strong in absolute value as the typical positive response. The regressions also demonstrate that our assumption of a concave function are supported by the data.

Table 6: Employment Asymmetry – parametric specifications

Coeff.	A. Unweighted Regression (I) quadr.	(II) cubic	(III) piece-wise linear	(IV) neg. linear pos. quadr.	B. Employment-weighted Regression (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)		
$z \times \mathbf{I}\{z < 0\}$			0.0947***				0.0795***	
			(0.0057)				(0.0217)	
$z^2 \times \mathbf{I}\{z \geq 0\}$				-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^i_{t-1}	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^i_t	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.

4.5 Industry analysis

If the asymmetric decision rules at the micro level is one of the factors responsible for employment dispersion and employment skewness, we can exploit the rich data to test more implications. Industries that exhibit a particularly strong asymmetry should exhibit a cross-sectional employment distribution that – on average – is more negatively skewed. This industry-level information is distilled into Figure 7 where we produce a scatter plot of skewness against asymmetry.

4.6 Robustness

4.6.1 Alternative TFP measures

We check for robustness. First, we want to understand if TFP innovations that are measured in a different way make a difference. Our preferred measure – used in producing Figures 6 and 7 as well as Tables 5 and 6 – are obtained by estimating the fixed-effects panel model in equation (6). Since this suffers from a well-known bias, we also estimate the same equation using the estimator by Arellano and Bond (1991); Blundell and Bond (2000). We these TFP innovations we redo our asymmetry analysis displayed in Table 6. If our preferred TFP innovation estimates suffered from a severe bias, the estimated asymmetry might vanish. As the results in Table 7 show, this is not the case. Using both Arrelano-Bond TFP innovations as well as Blundell-Bond TFP innovations delivers similar results to those in Table 6. The asymmetry certainly does not vanish.

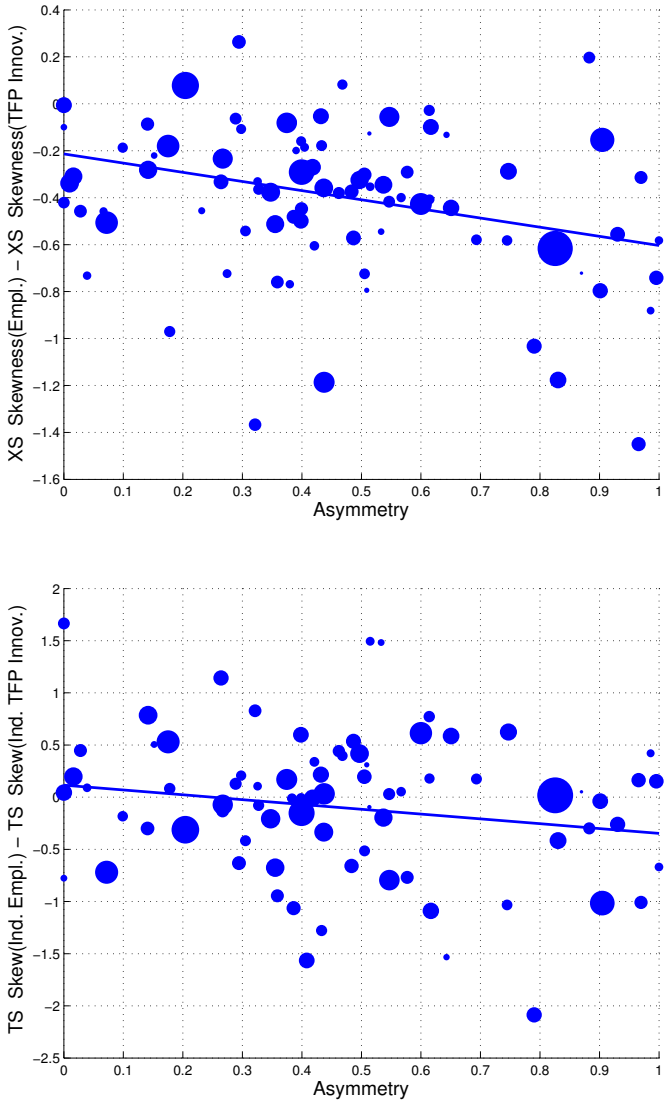
4.6.2 Asymmetric responses to news or asymmetric hiring costs?

4.6.3 Endogenous Quits

When workers quit firms that did not get a profitability shock and firms do not replace these workers (say, because of labor market frictions), then we would measure a declining employment growth while there were no shocks. So we might think that dispersion is larger than it actually is and skewness to be more negative than it actually is. Since quit rates are procyclical, this bias would be procyclical too. Our measured dispersion in booms would be hence too wide and skewness too negative which means that our measurements are a lower bound on actual dispersion and skewness caused by productivity shocks.

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

Figure 7: Industry Evidence: Hiring Asymmetry and Employment Patterns



Note: The figure plots the cross-sectional skewness of employment relative to that of TFP shocks against the asymmetry of the hiring rule within that NAICS-4 digit industry. The asymmetry measure equals the share of non-parametrically estimated employment growth variation predicted by non-linear terms: $1 - var(z)[g'^2/var(g(z))$ where $g(z)$ denotes the non-parametrically estimated employment growth given z. A value of 0 means that all variability in $g(z)$ is explained by linear terms in $g(\cdot)$, while a value of 1 means all variation is fully explained by the non-linear terms. The correlation coefficient between asymmetry and cross-sectional skewness is -0.283 with a standard error of 0.XXX.

Table 7: Robustness: Alternative Measures of TFP Innovations

Coeff.	Arellano-Bond (I) quadr.	Arellano-Bond (II) piece-wise linear	Blundell-Bond (III) quadr.	Blundell-Bond (IV) 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)	
$z \times \mathbf{I}\{z < 0\}$		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^i_{t-1}	0.0252**	0.0251**	0.0273**	0.0273**
	(0.0015)	(0.0015)	(0.0014)	(0.0014)
l^i_t	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***

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

Table 8: Robustness: Hiring Constraints? Labour-Saving technical change?

Coefficient	(I) baseline on PCU sample	(II) w/ hiring constraints	(III) no large investments	(IV) mid-age establishments
z	0.0679***	0.0673***	0.0750***	0.0759***
	(0.0043)	(0.0044)	(0.0019)	(0.0018)
z^2	-0.1200***	-0.1156***	-0.1415***	-0.1459***
	(0.0139)	(0.0142)	(0.0064)	(0.006)
$\mathbf{I}\{n\text{-constr.}\}$		0.0249***		
		(0.0031)		
$z \times \mathbf{I}\{n\text{-constr.}\}$		0.0063		
		(0.0156)		
$z^2 \times \mathbf{I}\{n\text{-constr.}\}$		-0.0723		
		(0.0554)		
Sample	PCU	PCU	ASM $\frac{i}{k} < 0.2$	ASM mid age
N	214k	214k	1,406k	1,421k
R^2	0.207	0.207	0.178	0.190
Pos. Response	+0.83%	+0.83%	+0.89%	+0.89%
Neg. Response	−1.60%	−1.57%	−1.79%	−1.82%
Diff. sign.?	Yes***	Yes***	Yes***	Yes***

Note: *, **, *** significantly different from 0 at the 10%, 5%, 1% level, respectively.
We want to know if the observed asymmetry in employment is specific to some labor-specific frictions such as adjustment costs, sticky wages etc. If employment responses are asymmetric because wages are downwardly rigid, then we would expect positive shocks to raise wages more than negative shocks decrease them. Columns (I) and (II) shod that this is not the case: even wages decline more in downturns than they increase in upturns, thus suggesting that the asymmetric labor response of firms takes place along both the wage and employment margins. Excluding observations with large investment projects (investment rate exceeding 20%) basically does not change the result.
After positive signals about future profitability, firms start increasing their inventories (materials and goods) today. But they reduce inventories even more strongly after a negative signal.

5 Conclusion

— to be done —

References

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

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 Guiseppe Moscarini. Business cycles and endogenous uncertainty. *Working Paper*, 2011.

Rüdiger Bachmann, Steffen Elstner, and Eric Sims. Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics*, 5(2):217–249, April 2013.

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, 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.

Ricardo J. Caballero. A fallacy of composition. *American Economic Review*, 82(5):1279–1292, December 1992.

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.

Jésus Fernández-Villaverde, Pablo Guerrón-Quintana, Keith Kuester, and Juan F. Rubio-Ramírez. Fiscal volatility shocks and economic activity. *Working Paper*, 2012.

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

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.

Cosmin Ilut. Ambiguity aversion: Implications for the uncovered interest rate parity puzzle. *American Economic Journal: Macroeconomics*, 4(3):33–65, July 2012.

Alejandro Justiniano and Giorgio E. Primiceri. The time varying volatility of macroeconomic fluctuations. *American Economic Review*, 98(3):604–641, June 2008.

Alejandro Justiniano, Giorgio E. Primiceri, and Andrea Tambalotti. Investment shocks and the relative price of investment. *Review of Economic Dynamics*, 14(1):102–121, January 2011.

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

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

Margaret M. McConnell and Gabriel Perez-Quiros. Output fluctuations in the United States: What has changed since the early 1980s? *American Economic Review*, 90(5):1464–1476, December 2000.

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.

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.

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.

Christopher A. Sims and Tao Zha. Were there regime switches in U.S. monetary policy? *American Economic Review*, 96(1):54–81, March 2006.

James H. Stock and Mark W. Watson. Has the business cycle changed and why? In Mark Gertler and Kenneth Rogoff, editors, *NBER Macroeconomics Annual*, volume 17, chapter 4, pages 159–218. MIT Press, January 2002.

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.

Laura L. Veldkamp. Slow boom, sudden crash. *Journal of Economic Theory*, 124(2):230–257, 2005. Learning and Bounded Rationality.

A Appendix

Proof of Proposition 1. *Part 1.* Write the derivative at 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

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

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

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

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

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

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

$$\max_{L_t^i} \min_{[\underline{\sigma}_\varepsilon, \overline{\sigma}_\varepsilon]} \exp\left[E^{\sigma_e}\left(z_t^i|s_t^i\right) + \frac{1}{2}var^{\sigma_e}\left(z_t^i|s_t^i\right)\right]\left(L_t^i\right)^\alpha - \overline{w}L_t^i. \tag{8}$$

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

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

The firm problem in (8) 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:

$$L_t^i = \left[\frac{\alpha}{\overline{w}} \exp\left(\gamma_t^* s_t^i\right)\right]^{\frac{1}{1-\alpha}}; \quad \gamma_t^* = \begin{cases} \overline{\gamma} & \text{if } s_t^i < 0 \\ \underline{\gamma} & \text{if } s_t^i \geq 0. \end{cases} \tag{9}$$

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 receives a negative signal $s_t^i = -x$, then employment contracts by more than it would expands if the firm receives a positive signal of the same magnitude $s_t^i = x$.

C Additional empirics

C.1 Industry-level asymmetry and employment patterns

We redo regression (7) 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 9: 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

continued ...

NAICS-4	Asymmetry ϕ_g	Rel. XS Skewness	Employment Share
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
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

continued ...

NAICS-4	Asymmetry ϕ_g	Rel. XS Skewness	Employment Share
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
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.2 Are Other Controls Asymmetric too?

Table 10: Robustness: Other controls

Specification Variable	(I) Real Wage	(II) Real Wage no $i/k > 0.2$	(III) Goods Inventories	(IV) Materials Inventories
Firms w/ pos. shock	-0.01%	-0.01%	$+4.2\%^{***}$	$+7.3\%^{***}$
Firms w/ neg. shock	$-0.6\%^{***}$	$-0.6\%^{***}$	$-9.9\%^{***}$	$-8.7\%^{***}$
N obs.	1,216k	1,216k	1,216k	1,216k

Note: We want to know if the observed asymmetry in employment is specific to some labor-specific frictions such as adjustment costs, sticky wages etc. If employment responses are asymmetric because wages are downwardly rigid, then we would expect positive shocks to raise wages more than negative shocks decrease them. Columns (I) and (II) shod that this is not the case: even wages decline more in downturns than they increase in upturns, thus suggesting that the asymmetric labor response of firms takes place along both the wage and employment margins. Excluding observations with large investment projects (denoted by and investment rate exceeding 20%) basically does not change the result.

Our information story is also corroborated by the behavior of inventories. After positive signals about future profitability, firms start increasing their inventories (materials and goods) today. But they reduce inventories even more strongly after a negative signal.