



DP2011/06

Cyclical Changes in Firm Volatility

Emmanuel De Veirman and Andrew Levin

November 2011

JEL classification: C33, D22, E31, E32

www.rbnz.govt.nz/research/discusspapers/

Discussion Paper Series

ISSN 1177-7567

DP2011/06

Cyclical Changes in Firm Volatility*

Emmanuel De Veirman and Andrew Levin[†]

Abstract

We estimate changes in the volatility of firm-level sales, earnings and employment growth of US firms. Our method differs from existing measures for firm-level sales and employment volatility in that it not only captures longer-run changes in volatility, but also measures cyclical changes in firm volatility. We detect substantial cyclical variation in firm-specific volatility around trend. Firm-specific volatility was low in the early 1990s, rose in the mid- and late-1990s, and was high around 2000. Our results are consistent with the hypothesis, deduced from models with financial frictions, that rising idiosyncratic volatility before 2001 contributed to the coincident rise in the external finance premium and to the 2001 recession. Endogenous pricing models imply that price adjustment is less frequent, and disinflation more costly, when firm-specific volatility is low. Consistent with endogenous pricing models, we find that the output cost of disinflation was three times larger in the early 1990s than in the early 2000s.

* The Reserve Bank of New Zealand's discussion paper series is externally refereed. The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Reserve Bank of New Zealand. We thank Laurence Ball, David Baqaee, Andrew Coleman, Martin Ellison, Ippei Fujiwara, Francesco Furlanetto, Nicolas Goshenny, Dean Hyslop, Alejandro Justiniano, Matthias Kehrig, Ken Kuttner, David Lucca, Robert Moffitt, Fumio Ohtake, Pierre-Daniel Sarte, Christie Smith, Ayako Suzuki, Jo Van Biesebroeck, Alex Wolman, and Raf Wouters for feedback. We also thank participants at the 2011 Royal Economic Society meetings, the 2010 Australian Conference of Economists, the 2007 New Zealand Econometric Study Group, and at seminars at the National Bank of Belgium, the Bank of Finland, the Australian National University, Victoria University Wellington, the Federal Reserve Board, and the Reserve Bank of New Zealand for comments.

[†] Corresponding author: Emmanuel De Veirman, Reserve Bank of New Zealand, 2 The Terrace, PO Box 2498, Wellington. *email address*: emmanuel.deveirman@rbnz.govt.nz. Andrew Levin, Board of Governors of the Federal Reserve System, Washington DC. ISSN 1177-7567 ©Reserve Bank of New Zealand

1 Introduction

This paper develops an empirical procedure to jointly estimate firm-specific volatility and dispersion, and applies it to characterise cyclical changes in the volatility of US firms' sales, earnings and employment growth.

Our study relates to papers by Comin and Philippon (2005), Comin and Mulani (2006), and Davis, Haltiwanger, Jarmin, and Miranda (2006), which characterise trend changes in the volatility of US firm-level sales and employment growth. These papers compare trends in firm volatility with a long-term decline in aggregate volatility known as the Great Moderation.¹

While the above-mentioned papers focus on estimating longer-run changes in firm volatility, we also estimate cyclical changes in the volatility of firm-level sales, earnings and employment growth. We can do so because, unlike the above papers which compute rolling windows of volatility, we use a procedure that indicates changes in volatility at the same frequency as the available data. We document that in our data, firm volatility estimated using a rolling measure is similar to a smoothed version of our new volatility measure. On that basis, we argue that our volatility measure adds new information by uncovering cyclical deviations in firm volatility from trend.

This new information is economically important. In particular, knowing the level of firm volatility at any point of time is important in models where changes in firm volatility have an immediate, potentially short-lived dynamic macroeconomic effect. This applies to the following three model classes.

First, consider the effect of firm-specific volatility on the output-inflation trade-off and the degree of monetary policy non-neutrality. In models where firms endogenously decide whether to adjust prices, higher firm-specific volatility implies more frequent price adjustment.² Intuitively, firms will use these

¹ We do not take a stand in the debate about whether over a long horizon, US firm-level volatility has increased or decreased. See Davis and Kahn (2008) for an overview and joint interpretation of evidence on changes in aggregate and microeconomic volatility. Papers on the Great Moderation include Kim and Nelson (1999), McConnell and Perez-Quiros (2000), Blanchard and Simon (2001), Stock and Watson (2002), Ahmed, Levin, and Wilson (2004), Dynan, Elmendorf, and Sichel (2006), Cecchetti, Flores-Lagunes, and Krause (2006), Justiniano and Primiceri (2008), Gali and Gambetti (2009), and Den Haan and Sterk (2011).

² See the endogenous pricing model of Ball, Mankiw, and Romer (1988), and state-dependent pricing models with firm-specific productivity shocks such as Klenow and Willis (2006, 2007), Golosov and Lucas (2007), Gertler and Leahy (2008), Burstein and Hellwig (2008), Costain and Nakov (2011), Dotsey, King, and Wolman (2011), and Vavra (2011).

more frequent re-optimizations to also price in macroeconomic shocks that have not yet been priced in.³ Therefore, as Vavra (2011) explicitly documents, an increase in the volatility of firm-specific shocks implies that monetary policy and other aggregate demand shocks have larger short-run effects on the price level, and correspondingly smaller real effects. In other words, the Phillips curve is steeper when firm-specific volatility is higher, a result which is explicit in Gertler and Leahy (2008).

Second, consider the effect of a change in firm-specific volatility on the business cycle when financial frictions exist. In the model of Christiano, Motto, and Rostagno (2010), a rise in the volatility of firm-specific productivity shocks implies an increase in borrower default risk and the external finance premium, which in turn engenders a cyclical downturn in output growth.⁴

Third, consider the effect of firm-specific volatility on aggregate output when firms face irreversibilities in the form of non-convex adjustment costs to labour and capital. In Bloom (2009) and Bloom, Floetotto, and Jaimovich (2009), a temporary rise in aggregate and firm-specific volatility, interpreted as a rise in uncertainty, results in a wait-and-see effect implying that on net, firms temporarily reduce investment and hiring. As a result, aggregate output temporarily declines.⁵

Aside from the fact that we estimate cyclical changes in volatility, a second feature of our approach is that we separately estimate the firm-specific component of volatility in firm sales and employment growth by controlling for aggregate and sectoral factors behind variation in firm growth. In theoretical models, including those mentioned above, shocks are aggregate or firm-specific. Therefore, it is desirable to estimate firm-specific volatility separately rather than estimating a mixture of firm-specific and aggregate factors. To our knowledge, we are the first to estimate time-variation in the firm-specific component of firm sales and employment volatility. However, our approach is akin to a literature that estimates time-variation in the volatility of firm-level stock returns, decomposed into firm-specific, sectoral

³ Related to this point, Boivin, Giannoni and Mihov (2009) document that firms in sectors with more volatile sector-specific shocks adjust prices more rapidly in response to macroeconomic shocks than firms in sectors with less volatile sector-specific shocks.

⁴ See also Gilchrist, Sim, and Zakrajsek (2011).

⁵ See also Bachmann and Bayer (2011).

and aggregate volatility.⁶

A third feature of our approach is that we jointly estimate firm-specific volatility as well as the implied dispersion in the distribution of firm growth rates. This implies that our empirical volatility measure captures the intuition, incorporated in the theoretical models of Bloom, Floetotto, and Jaimovich (2009), Christiano, Motto, and Rostagno (2010) and the other forementioned theories, that firm-specific shocks engender dispersion in the firm growth distribution. Neither of the available empirical measures⁷ jointly estimate firm-specific volatility and dispersion.

We find that firm-specific volatility gradually increased throughout the boom period of the 1990s. On the other hand, firm-specific volatility gradually decreased after the Volcker disinflation and after the 2000 stock market crash and 2001 recession. Our interpretation is that firms tend to become gradually more stable after adverse aggregate events, but become gradually more volatile during extended boom periods. In line with that pattern, we find that firm-specific volatility is positively correlated with current and past output gaps in our sample, 1986-2005. We attribute these findings to the possibility that after a negative aggregate shock to growth, firms become less prone to invest in projects that are ex ante risky, such that measured ex post volatility gradually declines as the fraction of capital reflecting less risky investment projects increases.

Our intuition is that any such positive effect of aggregate output on firm volatility may co-exist with the negative effect of firm volatility on output implied by the above-mentioned financial frictions and irreversibility theories. To the extent that the financial frictions and irreversibility models primarily suggest that sharp increases in firm-specific volatility cause recessions, their empirical validation rests on a negative relation between volatility (at the firm and aggregate level) and aggregate output during recessions. Our sample

⁶ For evidence on changes in the volatility of stock returns in the United States, see Campbell, Lettau, Malkiel, and Xu (2001), Malkiel and Xu (2003), Fama and French (2004), Brown and Kapadia (2007), and Fink, Fink, Grullon, and Weston (2010). Hamao, Mei, and Xu (2007) provide evidence for Japan. For a similar methodology applied to differences across sectors in the volatility of sales and productivity, see Castro, Clementi, and Lee (2009).

⁷ Comin and Philippon (2005) and Comin and Mulani (2006) compute firm volatility, while Davis, Haltiwanger, Jarmin, and Miranda (2006) compute distinct volatility and dispersion measures. The empirical results of Bloom, Floetotto, and Jaimovich (2009) characterise changes in the dispersion of firm-level sales growth measured by the interquartile range of the sales growth distribution. Bachmann, Elstner, and Sims (2010) construct proxies for uncertainty from surveys on business expectations. Vavra (2011) characterises price dispersion by the interquartile range of price changes.

contains only two NBER recessions, and our findings suggest that the 1990-91 recession was not caused or exacerbated by an increase in firm-specific volatility. On the other hand, firm-specific volatility rose before the 2001 recession, in line with the hypothesis that this rise in firm-specific volatility increased borrowing costs and/or caused firms to postpone investment and hiring, and therefore contributed to the 2001 recession.

Finally, firm-specific volatility was low in the early 1990s and high in the early 2000s. In endogenous pricing models, this implies less frequent repricing in the earlier episode, which in turn implies that a larger contraction in aggregate demand is needed to reduce inflation by any given amount. In line with endogenous pricing models, we find that the output cost of reducing inflation by any given amount was three times higher in the early 1990s than it was in the early 2000s.

Our paper is structured as follows. Section 2 characterises the data and discusses sample selection. Section 3 details our approach for estimating cyclical changes in firm volatility. Section 4 presents results. Section 5 interprets our findings. Section 6 concludes.

2 Data

Using annual data from the Thomson Worldscope database, we compute firm-level growth rates in nominal net sales, nominal Earnings Before Interest and Taxes (EBIT), and the number of employees for the period 1986-2005.⁸ The present section characterises the levels and growth data and discusses our choices regarding sample selection and data treatment. Subsequent sections characterise volatility in these firm-level growth rates.

We work with two main samples. Our “full sample” or “unbalanced panel” includes 15,425 firms that are incorporated in the United States and are listed on at least one of the US stock exchanges. This sample excludes firms that trade on the US stock market only through American Depositary Receipts (ADRs), as well as other firms that indicate a physical address outside the

⁸ Results are virtually identical when we deflate sales and EBIT by the aggregate Producer Price Index (PPI). Net sales equals gross sales minus cash discounts, trade discounts, and returned sales and allowances for which credit is given to the customer. EBIT is sales and other income minus operating expenses, without subtracting net expenditure on interest and taxes. The number of employees accounts for full-time and part-time employees, but excludes seasonal and emergency employees.

US.⁹

We also construct a “balanced panel” containing the subsample of 588 firms for which we can compute all three growth rates (sales, earnings and employment) for each of the twenty years in the sample. This is an effort to control for changes in sample composition. To see the importance of this, consider the following two reasons why sample composition varies over time in the unbalanced panel.

Firstly, all firms in our dataset are publicly traded. Davis, Haltiwanger, Jarmin, and Miranda (2006) find that firms newly listed in the 1980s and 1990s exhibit much greater volatility than earlier cohorts. They also show that cohort effects explain two thirds of the observed trend increase in the volatility of publicly traded firms from the 1970s to the 2000s. Because our balanced panel contains the same firms at every point of time, cohort effects do not affect its results. For comparison, note that Comin and Philippon (2005) and Comin and Mulani (2006) consider unbalanced panels of publicly traded firms.

Secondly, Worldscope has expanded its coverage of listed firms over time.¹⁰ Therefore, when a new firm appears in the database, this does not necessarily reflect a new listing. On the other hand, firms in the sample tend to stay as long as they continue to exist as listed companies. By expanding its coverage of listed firms, Worldscope has implicitly relaxed the criteria for inclusion in the database. In particular, smaller firms tend to have been included at a later date than larger firms. This change in sample composition matters in particular because of the stylised fact that smaller firms tend to be more volatile.

Figure 1 illustrates that this gradual inclusion of small firms implies a trend decline in the level of net sales, operational profit and the number of employees, a decline which we do not take to be representative for the overall

⁹ To arrive at the sample of 15,425 firms, we excluded 17 entities in the SIC division public administration. Next, we excluded 1,302 firms only listed through ADRs and 53 additional firms that indicate an address outside of the US. Finally, we excluded 269 firms with empty primary SIC code and 199 firms that are listed under the heading ‘nonclassifiable’ in the SIC structure.

¹⁰ Coverage of listed firms expanded in particular during the 1990s. The number of US firms for which data on all three growth rates (sales, EBIT and employment) are available in Worldscope increases from 1,629 in 1986 to 6,025 in 1999. Due to delisting in the early 2000s, the number of firms for which all three growth rates are available decreases to 4,915 in 2005.

population of firms.¹¹ In Figures 1 through 4, large diamonds indicate the median in a given year, small diamonds indicate the 25th and 75th percentiles, and the dashed line indicates the mean. In the levels data, the mean substantially exceeds the 75th percentile, suggesting skewness in the distribution of firms.

A disadvantage of the balanced panel is that it contains a disproportionately large fraction of large firms. The fact that the balanced panel only includes firms that have never ceased to exist also means that, for firms of any given size, it tends to focus on stable firms with comparatively few observations with high volatility.

Furthermore, note that even if the balanced panel considers the same firms at any point of time, it does not quite hold sample composition constant in terms of firm size. The firms in the balanced panel tend to grow at a faster rate than aggregate output, implying that firm size gradually increases relative to the size of the economy. In that respect, we almost have the opposite issue as in the unbalanced panel. At any rate, Figure 2 indicates a stable upward trend in sales, profits and the number of employees in the balanced panel, with the exception of slowdowns during the recessions of the early 1990s and early 2000s.

Because of these disadvantages of the balanced panel, we investigate robustness on two counts. First, we find that our conclusions are robust to dividing the entire sample for continuously available firms into quartiles according to firm size, defined by the sales-to-GDP ratio. In that case, all observations in a subsample are within well-defined firm size bands, which is a way of holding firm size reasonably constant. To our knowledge, this is a new way of controlling for changes in sample composition in the literature on firm volatility.

Second, we estimate volatility in the unbalanced panel. In that exercise, we control for firm fixed effects, implying that our estimate of average volatility is less affected by the gradual extension of the sample to firms that are intrinsically more volatile. This is similar to the specifications in Comin and Philippon (2005) and Comin and Mulani (2006) that control for firm size and age. For every company in the unbalanced panel, we include all available observations, whether or not the company still exists at the end of

¹¹ There is a reversal in this trend from about 2000. This reflects the fact that Worldscope coverage of listed firms is virtually complete from that point on, in combination with actual positive economic growth. For the first few observations after 2000, the increase in average firm size plausibly also reflects the delisting that happened in the aftermath of the 2000 stock market crash, where smaller firms were more likely to delist.

the sample.

As will become apparent when we discuss results in Section 4, our main conclusions regarding cyclical swings in firm volatility hold for the balanced panel, the unbalanced panel, and apply across firm size quartiles. This robustness suggests that our conclusions are not driven by changes in sample composition.

Having discussed levels data, we now turn to growth rates. We compute annual growth rates in net sales and employment.¹² For Earnings Before Interest and Taxes (EBIT), we cannot compute the growth rate for any year t in the usual fashion, since doing so would yield meaningless results when EBIT is negative in t and/or $t - 1$. Therefore, we compute growth in operational profit based on the change in EBIT divided by lagged net sales:

$$\gamma_{it} = \frac{EBIT_{it} - EBIT_{i,t-1}}{SALES_{i,t-1}} * 100 \quad (1)$$

We do not account for firm entry or exit, in the sense that we require two consecutive observations on sales, employment, or EBIT in order to compute a growth rate.¹³

Even in the unbalanced panel, we only include observations for which data are available on all three growth rates. This ensures that the sample is the same for each of the three growth rates.

We windsorise the data in order to reduce the impact of outliers on our results. For every growth rate, we determine the 2.5th and 97.5th percentiles of all observations in the unbalanced panel. We replace any (negative) growth rate which falls below the 2.5th percentile by the value of the growth rate at the 2.5th percentile. Similarly, we replace any (positive) growth rate exceeding the 97.5th percentile by the value of the growth rate at the 97.5th percentile. We do not apply additional windsorizing to the balanced panel.

Figures 3 and 4 graph the distribution of sales, EBIT, and employment growth for the unbalanced and balanced panels, respectively. In both cases,

¹² We compute sales growth after dropping the twenty-three observations for which net sales is strictly negative. Reported results are based on growth rates computed in the regular fashion. For instance, the sales growth rate of firm i in year t is $\gamma_{it} = \frac{SALES_{it} - SALES_{i,t-1}}{SALES_{i,t-1}} * 100$. However, results are virtually indistinguishable when we use the measure in Davis, Haltiwanger, Jarmin, and Miranda (2006) : $\gamma_{it} = \frac{SALES_{it} - SALES_{i,t-1}}{(SALES_{i,t} + SALES_{i,t-1}) / 2} * 100$.

¹³ We do not account for entry and exit because, as we mentioned earlier in this section, new appearances in the sample do not necessarily reflect actual entries.

the firm-level data accurately indicate slowdowns during the recessions of the early 1990s and early 2000s. Taking into account the scale difference on those two graphs, dispersion in growth rates across firms appears to be larger in the unbalanced panel. We confirm this finding more formally in Section 4.

3 Estimation

In this section, we describe our procedure for estimating firm-specific volatility and dispersion. Before doing so, we discuss the existing approach for estimating firm-level sales and employment volatility.

3.1 Existing Approach

Comin and Philippon (2005), Comin and Mulani (2006) and Davis, Haltiwanger, Jarmin, and Miranda (2006) all measure firm-level volatility based on rolling ten-year standard deviations of firm-level growth rates.¹⁴ In particular, they compute the volatility of firm i in year t as:

$$\sigma_{it}^r = \sqrt{\frac{1}{10} \sum_{\tau=-4}^{\tau=5} (\gamma_{i,t+\tau} - \bar{\gamma}_{it})^2} \quad (2)$$

We chose the superscript for σ_{it}^r in order to differentiate this “rolling” volatility measure from our cyclical measure in equation (4) below. $\gamma_{i,t}$ represents net sales growth, EBIT growth, or employment growth between year $t - 1$ and t for firm i . $\bar{\gamma}_{it}$ captures the firm’s average growth in the years $t - 4$ to $t + 5$.

To translate the σ_{it}^r ’s into one estimate of firm-level volatility for any period t , the above-mentioned papers compute the cross-sectional median of σ_{it}^r and/or the (size-weighted or unweighted) average of σ_{it}^r across firms.

While equation (2) assigns the volatility estimate σ_{it}^r to a particular year t , measured volatility in fact depends on variation in growth rates between $t - 4$ and $t + 5$. In other words, σ_{it}^r does not yield a separate estimate of

¹⁴ Davis, Haltiwanger, Jarmin, and Miranda (2006) also consider a modified rolling standard deviation which accounts for firms with growth rates available for fewer than ten consecutive years.

volatility for any year. It is customary to assign the volatility estimate to a particular year, for instance to the mid-point as in equation (2). However, it is more precise to say that equation (2) captures average volatility for a ten-year period.

Since the cross-sectional median or average of σ_{it}^r does not constitute a separate estimate for firm volatility in any year, rolling it forward does not accurately capture year-on-year changes in volatility. Because rolling standard deviations in successive periods overlap, they tend to smooth out annual changes in firm volatility. In this respect, in Section 5.1 we will document that a rolling volatility estimate resembles a smoothed version of our paper's volatility estimate.

A volatility measure based on equation (2) is informative about longer-run changes in volatility, in the sense that it indicates how firm volatility changes from one ten-year period to another distinct ten-year period. Adopting the common definition of the business cycle as a cycle that takes between six quarters and eight years to complete,¹⁵ this implies that a rolling volatility measure indicates how volatility changes from one business cycle to the next. However, it smooths out shorter-run changes, and therefore does not accurately indicate how volatility changes at business cycle frequencies.

As a final point on the rolling volatility measure, note that empirical changes in firm-level variables reflect not only firm-specific events, but also aggregate conditions and sector-wide developments. Total firm-level volatility σ_{it}^r in equation (2) in principle reflects all these factors without disentangling one from the other.

3.2 Our Method

We now discuss our procedure for estimating cyclical changes in firm-specific volatility. As we will discuss in more detail at the end of this section, our method differs from the rolling approach in three ways. First of all, our measure captures volatility for every year separately, and therefore accounts for cyclical changes in volatility. Second, our measure separately estimates the firm-specific component of firm-level volatility. Third, our measure jointly captures firm-specific volatility and dispersion in the firm growth distribution.

Our volatility measure is based on the residual of the following regression:

¹⁵ See Stock and Watson (1999).

$$\gamma_{it} = c + a_i + b_t + d_{st} + e_{qt} + \varepsilon_{it} \quad (3)$$

As in equation (2), γ_{it} represents net sales growth, EBIT growth, or employment growth between year $t - 1$ and t for firm i . The constant is c .¹⁶

Due to the presence of firm fixed effects a_i , the residual ε_{it} depends on the deviation of a firm's growth rate from its average growth rate over the sample. This feature is similar to the fact that the rolling standard deviation in equation (2) depends on the deviation of firm growth from a time-varying mean growth rate.¹⁷

The effects of aggregate factors on firm growth are captured by time fixed effects b_t , by sector-time interactions d_{st} and size quartile-time interactions e_{qt} . By virtue of the two interaction terms, we allow aggregate factors to have a different effect on firm growth depending on a firm's size and the sector in which it operates.

To construct the firm size effects, we divide all observations for the relevant panel into quartiles according to the sales-to-GDP ratio. We will also use these quartiles when we consider volatility patterns for every firm size category separately. All of the results in this paper are virtually unchanged when we construct size quartiles based on the number of employees.

To construct sector effects, we classify firms into three broad sectors by means of data on each firm's primary Standard Industrial Classification (SIC) code. In particular, we classify firms into manufacturing firms, service firms (including firms in wholesale and retail trade), and firms operating in other sectors.¹⁸ As will become clear in Section 4.2, the results from the panels involving multiple sectors most closely reflect developments in the manufacturing sector.

We use heteroskedasticity-robust standard errors throughout this paper. As the results in Section 4 reveal, the typical standard deviation of ε_{it} changes

¹⁶ We do not weigh observations by firm size. As Section 4.2 documents, our main conclusions are robust across firm size quartiles. This suggests that assigning more weight to larger firms would not alter our main conclusions.

¹⁷ In our approach, the firm's mean growth rate does not need to be time-invariant. De Veirman and Levin (2009) model firm i 's mean growth rate as a Hodrick-Prescott trend of γ_{it} . This allows for gradual time-variation in the mean growth rate.

¹⁸ Manufacturing firms are from the SIC division manufacturing. Service firms are from the following three SIC divisions: services; wholesale trade; and retail trade. Firms in other sectors are in the five remaining private-sector SIC divisions: agriculture, forestry, fishing; mining; construction; transportation, communications, electric, gas, sanitary services; and finance, insurance, real estate.

substantially over time, indicating that accounting for heteroskedasticity is important.

We now focus on the residual from equation (3). As we mentioned before, the inclusion of firm fixed effects means that the residual depends on the deviation of firm i 's growth rate in year t from its average growth rate over the sample. Because equation (3) also controls for time, sector and size effects, the residual only captures firm-specific variation. In particular, it captures deviations of a firm's growth from its average growth rate for reasons other than economy-wide or sectoral developments or events shared with other firms of similar size.

We estimate the standard deviation of the residual by a term proportional to the absolute value of the estimated residual $\widehat{\varepsilon}_{it}$:

$$\widehat{\sigma}_{\varepsilon,it} = \sqrt{\frac{\pi}{2}} |\widehat{\varepsilon}_{it}| \quad (4)$$

Therefore, we estimate every firm i 's volatility at any point of time t from a single observation.¹⁹ We now discuss the bias and convergence properties of this estimator.

In Appendix 2, we prove that, if the true error term ε_{it} is normally distributed with mean 0, i.e. $\varepsilon_{it} \sim N(0, \sigma_{\varepsilon,it}^2)$, equation (4) yields an unbiased estimator $\widehat{\sigma}_{\varepsilon,it}$ of the true standard deviation $\sigma_{\varepsilon,it}$. In summary, the proof consists of taking expectations of equation (4) assuming that we know the true error $\varepsilon_{i,t}$, implying $E(\widehat{\sigma}_{\varepsilon,it}) = \sqrt{\pi/2} E(|\varepsilon_{i,t}|)$, and writing out the expectation $E(|\varepsilon_{i,t}|)$ explicitly in integral form. The latter step yields $E(\widehat{\sigma}_{\varepsilon,it}) = \sqrt{\pi/2} \int_{-\infty}^{+\infty} |\varepsilon_{i,t}| f(\varepsilon_{i,t}) d\varepsilon_{i,t}$. Substituting $f(\cdot)$ by the probability density function of the mean-zero normal distribution, Appendix 2 shows that $E(\widehat{\sigma}_{\varepsilon,it}) = \sigma_{\varepsilon,it}$.

We now turn to the speed of convergence. As we discuss later in this section, our focus is on the cross-sectional average of firm volatility $(1/N) \sum_{i=1}^N \sigma_{\varepsilon,it}$, with N the number of firms. We now assess whether the estimator for the cross-sectional average of firm volatility, $(1/N) \sum_{i=1}^N \widehat{\sigma}_{\varepsilon,it}$, converges to the truth quickly enough for us to be able to draw inference from our estimates given the empirical sample sizes we work with.

¹⁹ Kim and Nelson (1999) and McConnell and Perez-Quiros (2000) similarly estimate the variance and standard deviation, respectively, of aggregate output growth from a single observation. Equation (4) is the firm-level equivalent of McConnell and Perez-Quiros's (2000) formula for the volatility of aggregate output growth.

To do so, we run the following Monte Carlo experiment. For different values of N , we simulate a population of N firms at time t , with any firm i 's idiosyncratic shocks distributed $N(0, \sigma_{\varepsilon,it}^2)$. For every firm, we draw true volatility $\sigma_{\varepsilon,it}$ from a uniform distribution $U(0, 20)$, such that the true average standard deviation $(1/N) \sum_{i=1}^N \sigma_{\varepsilon,it} = 10$. This yields a simulated population of firms ranging from very stable to very volatile.²⁰

For every firm, we draw a single observation for ε_{it} , and compute its estimated standard deviation $\hat{\sigma}_{\varepsilon,it} = \sqrt{\pi/2} |\varepsilon_{it}|$. From the individual firms' estimated standard deviations, we compute the estimator $(1/N) \sum_{i=1}^N \hat{\sigma}_{\varepsilon,it}$. This estimator depends on the particular draws of the ε_{it} 's. To evaluate the probability that the estimator is close to the truth, we repeat the drawing of the ε_{it} 's and the estimation of $(1/N) \sum_{i=1}^N \hat{\sigma}_{\varepsilon,it}$ one million times, and capture the percentiles of the distribution of $(1/N) \sum_{i=1}^N \hat{\sigma}_{\varepsilon,it}$.

For different values of N , Table 1 shows the median, as well as the 2.5th and 97.5th percentiles, of the estimator $(1/N) \sum_{i=1}^N \hat{\sigma}_{\varepsilon,it}$. For all simulated sample sizes in the table, the median estimator is at or near the true value of 10.00. This also applies to the mean of the estimator (unreported), which follows from unbiasedness. Table 1 confirms convergence, in the sense that the 95 percent probability intervals around the median shrink as N grows large. When $N = 588$, which is the sample size of our empirical balanced panel, the estimator lies in the interval (9.30, 10.72) with 95 percent probability. In our unbalanced panel, the number of firms varies between 1,629 in 1986 and 6,025 in 1999. With $N = 1,629$, the estimator lies in (9.58, 10.43) with 95 percent probability. The 95 percent probability interval for $N = 6,025$ is (9.78, 10.22).

These probability intervals are narrow compared to the estimated changes in volatility that we show in Section 4. This indicates that our estimator converges sufficiently quickly in order to allow us to draw inference about changes in volatility from our empirical samples.

We now translate the estimated values for $\hat{\sigma}_{\varepsilon,it}$ into a single measure for every year t . In the balanced panel, we run the following regression:

$$\hat{\sigma}_{\varepsilon,it} = k + \delta_t + \nu_{it} \tag{5}$$

²⁰ The Monte Carlo results are conditional on our assumption that the cross-sectional distribution of true standard deviations is uniform. On the other hand, conditional on using a uniform distribution, the choice of 20 as the upper bound of the support is without loss of generality.

Equation (5) entails regressing firm-specific volatility on a constant k and time fixed effects δ_t . The error term ν_{it} is in principle independently and identically distributed with mean zero and variance σ_ν^2 .

The time fixed effect δ_t captures the cross-sectional average of firm-specific volatility $(1/N) \sum_{i=1}^N \hat{\sigma}_{\varepsilon,it}$ in year t . In Section 4, we characterise the evolution of firm-level volatility by plotting the time effects from equation (5) with respect to time.²¹

In the unbalanced panel, we instead run the following regression:

$$\hat{\sigma}_{\varepsilon,it} = k + \delta_t + \zeta_i + \nu_{it} \tag{6}$$

The inclusion of firm fixed effects ζ_i implies that, if an intrinsically more volatile firm is added at some point of time, this shows up as a relatively high estimate for that firm’s specific effect without necessarily implying an increase in measured average volatility.

As promised, we now discuss the differences between our estimator and rolling volatility measures in more detail. First, unlike σ_{it}^r in equation (2), our estimator $\hat{\sigma}_{\varepsilon,it}$ from equation (4) is based on a single observation, and therefore captures volatility in a given year. This fact allows us to capture changes in firm volatility over the course of the business cycle even with annual data.

Second, because we control for aggregate and sector-specific factors in equation (3), we estimate the firm-specific component of firm-level volatility. As discussed in the introduction, this makes our measure more directly comparable to theoretical concepts of firm-specific volatility.

Third, we jointly estimate firm-specific volatility and the dispersion it induces. To see this, note that when many firms have a high value for the absolute value $|\hat{\varepsilon}_{it}|$ in any year t , this means that firms tend to deviate substantially from the average growth rate that prevails at time t for firms in the same sector and of similar size. Saying that firms deviate substantially from the cross-sectional average is equivalent to saying that there is a high level of dispersion in the firm growth distribution. Therefore, $|\hat{\varepsilon}_{it}|$ captures dispersion.

On the other hand, when many firms have a high value for $|\hat{\varepsilon}_{it}|$, this implies high estimated firm-specific volatility through equation (4). Therefore, $|\hat{\varepsilon}_{it}|$

²¹ To avoid perfect collinearity, we omit the time dummy for 1986. Correspondingly, k indicates average firm-specific volatility in 1986, while $k + \delta_t$ indicates volatility for any other year t .

is proportional to firm-specific volatility. In sum, when many firms have a high value for $|\widehat{\varepsilon}_{it}|$, both dispersion and firm-specific volatility are high. In this way, our empirical measure mirrors the theories mentioned in the introduction by capturing the intuition that firm-specific shocks engender dispersion in the firm growth distribution.

4 Results

This section characterises the time-path of firm-specific volatility. We discuss the balanced panel before reporting results by firm size, by sector, and for the unbalanced panel. We will interpret the results in Section 5.

4.1 Balanced Panel

For the balanced panel of 588 US firms, Figure 5 characterises annual changes in firm-specific volatility, estimated from equations (3) through (5).

The upper left diagram graphs sales growth volatility. Net sales became gradually less volatile in the second half of the 1980s. Sales growth volatility remained comparatively low in the period 1990-1993. From 1994 onwards, volatility increased, peaking in 2000. After that, sales have gradually become more stable.

As will become clear shortly, this graph illustrates a general pattern: US firms became gradually more stable in the latter half of the 1980s, became gradually more volatile during the 1990s, and gradually stabilised again in the first half of the 2000s.

The top row of Table 2 reveals that on average over the subsample 1986-1995, the cross-sectional average of the estimated firm-specific standard deviation is 13.36 percent. The standard deviation increases to 16.22 in 1996-2005, an increase which is statistically significant at the one percent level.

The remaining two diagrams of Figure 5 document that earnings volatility and employment volatility track a similar pattern, in the sense that they also indicate that US firms became gradually more stable in the second half of the 1980s, became gradually more volatile in the 1990s, and stabilised after 2000.

Table 2 reveals that the average firm-specific standard deviation of EBIT growth increased from 5.93 to 7.07 percent from the first to the second half

of the sample, while employment volatility increased from 12.65 to 14.51 percent. Both increases are statistically significant at the one percent level.

Note that the estimated level of volatility depends on the way we scale the growth rates. As we are about to explain, this implies that we cannot interpret Figure 5 or Table 2 as saying that earnings are less volatile than sales and employment. Recall from equation (1) that we compute EBIT growth with lagged sales in the denominator. In absolute value, sales is often larger than EBIT, where the latter captures profits. The comparatively large absolute values of sales tend to imply comparatively small absolute values for the growth rate in equation (1), and therefore tend to imply low values for measured EBIT volatility. This scaling issue does not affect our interpretation of changes in volatility over time for any given variable.

4.2 Robustness

We consider robustness on three counts. First, we divide the balanced panel into firm size quartiles. Second, we discuss results by sector. Third, we report results for the unbalanced panel. At the end of the section, we also briefly compare this paper's evidence for the United States to results from a working paper by De Veirman and Levin (2009) for Japanese firms.

Figure 6 reports results by firm size quartile based on threshold sales-to-GDP ratios, as we anticipated in Section 3.2. The left column captures the quartile of smallest firm sizes, and firm size increases as we consider columns further to the right. As before, we estimate volatility from equations (3) through (5), but with one difference: we drop the firm size interaction term e_{qt} from equation (3).

In terms of average volatility over the sample, our results confirm the stylised fact that smaller firms are more volatile than larger firms. This applies for all three variables. In terms of changes in volatility over time, we find a similar pattern for each of the size quartiles as we did in the previous subsection. Irrespective of firm size, volatility tends to decrease in the second half of the 1980s, tends to increase in the 1990s, and tends to decrease in the first half of the 2000s. This suggests that the results from the overall balanced panel are not driven by changes in the firm size composition over the course of the sample.

Figure 7 presents results by sector. As we explain in Section 3.2, we divide firms into three broad categories: manufacturing, services, and other sectors.

As before, we use equations (3) through (5), but we now omit the sector interaction term d_{st} from equation (3).

The results from the previous subsection directly correspond to the volatility paths in the manufacturing sector. On the other hand, services and the ‘other sectors’ category display different patterns. In this respect, note that out of the 588 firms in the balanced panel, 353 are in manufacturing, while 114 are in services and 121 in other sectors. From this exercise, we conclude that our main results mostly speak for manufacturing firms.

We now consider the full Worldscope sample. In this case, we use equations (3), (4) and (6).

For the unbalanced panel, Figure 8 reveals a decrease in volatility in the second half of the 1980s, rising volatility in the 1990s, and declining volatility in the 2000s. Therefore, our conclusions regarding cyclical changes in volatility are robust to using the unbalanced panel. This further supports the notion that our main conclusions from the previous subsection are not driven by changes in the firm size composition.

A typical firm in the unbalanced panel is more volatile than a typical firm in the balanced panel. Table 3 reveals that volatility in the unbalanced panel is about twice as high in terms of sales and employment, and about three and a half times as high in terms of earnings, as in the balanced panel. This plausibly reflects the facts, mentioned in Section 2, that the balanced panel has a comparatively high concentration of large firms, and that it only includes firms that have never ceased to exist, and therefore tend to be more stable for any given firm size.

Comparing these results for US firms to the results from De Veirman and Levin (2009) for Japanese firms, we find that US firms in the balanced panel are about twice as volatile as Japanese firms in a similar panel. The difference is even more pronounced when comparing unbalanced panels.

We attribute this volatility difference in part to the likelihood that in the United States, it is easier for young and volatile firms to obtain a stock market listing than is the case in Japan.²² In addition, Japanese firms are plausibly less volatile because of features specific to the Japanese economy, such as the existence of structured business groups, and particularly tight links between government, banks and firms, that plausibly dampened the

²² At any rate, the findings of Davis, Haltiwanger, Jarmin, and Miranda (2006) suggest that young US firms could more easily obtain a listing in the 1980s and the 1990s than in earlier decades. See Brown and Kapadia (2007) for a similar argument.

effect of market forces on troubled firms. See Hamao, Mei, and Xu (2007) on the effect of business group protection in Japan on firm-level stock return volatility.

5 Interpretation

We first document that our measure adds new information to rolling measures by uncovering cyclical changes in volatility. Next, we discuss the relation between estimated firm volatility and the business cycle. Finally, we discuss economic implications of our findings.

5.1 Cyclical and Rolling Volatility

To document that our procedure yields new information, we compute firm volatility using a rolling volatility measure on our sample, and compare it to the previous section’s results. For firms in the balanced panel, we compute volatility for firm i in year t from a centered ten-year standard deviation:

$$\sigma_{it}^{roll} = \sqrt{\frac{1}{9} \sum_{\tau=-4}^{\tau=5} (\gamma_{i,t+\tau} - \bar{\gamma}_{it})^2} \quad (7)$$

The right-hand side is the same as in equation (2) except for the fact that we make a degrees-of-freedom correction. To mark that difference, we use a different superscript than in equation (2).

Next, we produce a single measure of firm volatility for year t by computing the unweighted cross-sectional average of the firm-level standard deviations:

$$\sigma_t^{roll} = \frac{1}{N} \sum_{i=1}^N \sigma_{it}^{roll} \quad (8)$$

where $N = 588$ for all years since we consider the balanced panel. The long dashed lines in Figure 9 graph σ_t^{roll} with respect to time for sales, earnings and employment growth. We refer to it as our “rolling volatility measure”. Since one value for σ_t^{roll} requires ten consecutive years of data, we can only compute eleven such values from our twenty-year sample. The fact that for each variable, the value for σ_t^{roll} assigned to 2000 exceeds the value assigned to 1990 corresponds to the finding reported in Section 4.1 that firms were

more volatile in the period 1996-2005 than in the period 1986-1995. The values of σ_t^{roll} graphed for years between 1990 and 2000 reflect overlapping windows.

The solid lines in Figure 9 repeat our estimates for firm-specific volatility from Figure 5. We refer to this as our “annual volatility measure”. The short dashed lines show the corresponding trends, computed by taking ten-year moving averages of annual volatility. We refer to this as our “smoothed volatility measure”.

Apart from a levels difference,²³ the rolling volatility measure is very similar to the trend component of our annual volatility measure. This motivates our claim that our annual volatility measure contains cyclical information that is not present in the rolling measure. The advantage of the annual measure is that it is instructive about the timing of year-on-year changes in volatility. This is not the case for the rolling and smoothed annual volatility measures. To illustrate this point, note that according to the rolling and smoothed measures, sales volatility increased faster around 1995 than it did at any other time in our sample. By looking at our annual volatility measure, which we argue captures the actual level of volatility in every year, we see that the fast increase in rolling and smoothed volatility around 1995 does not speak for actual changes in volatility in the mid-1990s. On the contrary, it reflects the fact that as the rolling window moves forward in time, it includes more high-volatility years around 2000, while it contains fewer low-volatility years around 1990.

5.2 Cyclical Pattern

Recall that Section 4’s main conclusion is that firm volatility gradually declined during the second half of the 1980s, gradually increased during the 1990s, and gradually decreased in the first half of the 2000s. We could only capture these turning points in volatility by implementing an approach that accurately indicates the timing of year-on-year changes in volatility.

Our interpretation of the path of firm-specific volatility over our sample is

²³ This levels difference cannot be explained by the fact that our annual volatility measure excludes aggregate and sectoral factors. When we add the aggregate and sector component to our annual volatility estimate, this only lifts the ten-year moving average up by a very small amount. One possible reason for there to be some difference is the fact that the degrees-of-freedom correction, while yielding an unbiased variance, leaves some bias in the rolling standard deviation. At any rate, the rolling and smoothed annual measures are of the same order of magnitude.

that firm volatility tends to gradually decline in the aftermath of adverse aggregate events, while it gradually increases during extended periods of stable growth. The Volcker disinflation, the 1987 stock market crash, and the 2000 stock market crash and 2001 recession all precede periods during which volatility gradually declines. Volatility does not decline, but stays low in the aftermath of the 1990-91 recession. On the other hand, volatility gradually increases during the productivity acceleration of the 1990s.

Our findings that adverse aggregate events tend to precede gradual declines in firm volatility, and that a long boom period coincides with a gradual increase in firm volatility, suggest that firm volatility is positively related to current and lagged aggregate output growth. We compute cross-correlations between firm volatility and the output gap to verify this intuition. To this end, we compute the output gap using the Congressional Budget Office's (CBO) estimate of potential output.²⁴ The contemporaneous correlations are all positive: 0.68 for sales volatility, 0.30 for earnings volatility, and 0.72 for employment volatility. For sales and employment volatility, the contemporaneous correlation is stronger, in absolute value, than the cross-correlation at any lead or lag. For earnings volatility, the strongest association is between firm-specific earnings volatility and the second lag of the output gap, with a correlation of 0.77. This confirms our intuition that firm volatility is positively related to current and lagged aggregate output.

At the end of this section, we will interpret the finding that firm volatility is pro-cyclical in our sample.

5.3 Macroeconomic Implications

The fact that our measure accurately indicates the timing of year-on-year changes in volatility means that it has clear economic implications in models where changes in firm-specific volatility have a dynamic macroeconomic effect. We first discuss endogenous pricing models before discussing financial frictions models and irreversibility models.

The endogenous pricing models we mentioned in the introduction imply that firms change their prices more frequently when firm-specific volatility is comparatively high. More frequent price adjustment in turn implies that any aggregate demand shock has larger nominal effects and smaller real effects.

²⁴ We use the 2011Q1 release for real Gross Domestic Product as well as for potential output.

In sum, the Phillips curve tends to be steep, and disinflating comes at a comparatively low cost, when firm-specific volatility is high.

In our sample period, 1986-2005, US inflation was low by historical standards. In addition, as the literature on the Great Moderation documents, aggregate volatility was low in our sample. Both facts imply that, as far as aggregate developments are concerned, firms in our sample have comparatively few reasons to reprice.²⁵ Golosov and Lucas (2007) show that in low-inflation economies, a model without firm-specific shocks substantially underpredicts actual repricing rates. They also show that accounting for firm-specific shocks helps their model to reproduce empirical repricing rates at low inflation rates. This suggests that firm-specific volatility is a particularly important determinant of repricing rates and the Phillips curve slope in a low-inflation sample such as ours.

We provide some evidence on this point by comparing sacrifice ratios, capturing output losses associated with disinflation, at different points of time. We compare the episode of the early 1990s, when firms were very stable, with the early 2000s, when firms were very volatile. Endogenous pricing implies that all other things equal, the Phillips curve is flatter in the earlier episode, which implies that a larger shortfall in aggregate demand is required in order to reduce inflation by any given amount. Therefore, the theoretical prediction is that the sacrifice ratio is larger in the earlier episode than in the later episode.

We compute the sacrifice ratio as the cumulative annual loss in real Gross Domestic Product (GDP) associated with a one percent decline in trend inflation. To this end, we apply the procedure of Ball (1994).²⁶ Table 4 documents the results for inflation in the Consumer Price Index (CPI). During the disinflation episode 1989Q4-1994Q3, trend CPI inflation declined 2.60 percentage points from peak to trough. The sacrifice ratio for this episode is 2.80, suggesting that a one percentage point decline in inflation required a loss of 2.80 percent of a year's output. In the period 2000Q2-2002Q2, infla-

²⁵ See DeFina (1991), De Veirman (2009) and Ball and Mazumder (2011) for empirical evidence on the relation between the Phillips curve slope on the one hand, and inflation and aggregate volatility on the other hand.

²⁶ In summary, we compute trend inflation as a centered nine-quarter moving average of inflation. We assume that output is at trend at the inflation peak and returns to trend four quarters after the inflation trough. We compute trend output for other quarters by log-linearly interpolating between those two points. The numerator of the sacrifice ratio is annualised cumulative deviations of output from trend. The denominator is the decline in trend inflation from peak to trough.

tion shrank by 1.46 percentage points.²⁷ The sacrifice ratio for that episode is 0.92.

The finding that the sacrifice ratio for the disinflation of the early 1990s is about three times larger than the ratio for the early 2000s suggests that firm volatility affected the output-inflation trade-off as predicted by endogenous pricing theories. Like Ball (1994) and Senda and Smith (2008), we focus on headline CPI inflation. However, the above finding is robust to using the CPI excluding food and energy. For comparison, Table 4 also shows that the sacrifice ratio for the Volcker disinflation, which took place in 1980Q1-1983Q4, is 1.94 for headline CPI inflation.

We now turn to financial frictions models and irreversibility models, both of which imply that a rise in firm-specific volatility causes a cyclical slowdown in aggregate growth. In financial frictions models such as Christiano, Motto, and Rostagno (2010), an increase in firm-specific volatility implies an increase in default risk. The resulting increase in the external finance premium tends to dampen aggregate output growth.

In irreversibility models such as Bloom (2009), an increase in firm-specific (and aggregate) volatility is interpreted as an increase in uncertainty. In combination with the assumption that firms face adjustment costs to labour and capital, this implies that on net, firms temporarily reduce hiring and investment after a rise in firm-specific volatility. The result is a temporary reduction in output.

There are two NBER recessions in our sample: 1990Q3-1991Q1 and 2001Q1-2001Q4. We find that firm-specific volatility is low before and during the recession of the early 1990s. Therefore, any wait-and-see effects on investment and hiring during the 1990-91 recession were not due to a rise in firm-specific uncertainty. Similarly, this result rules out that rising firm-specific volatility would have been a factor behind the increase, documented by De Graeve (2008), in the relative cost of external finance in the second half of the 1980s.²⁸

On the other hand, we find that firm-specific volatility was rising before the 2001 recession. This allows for the possibility that firm-specific uncertainty contributed to the 2001 recession because it caused firms to postpone investment and hiring. This increase in firm-specific volatility may also have

²⁷ Ball (1994), using data through the early 1990s, excludes episodes where inflation declined less than 2 percentage points. Our implementation differs in that respect.

²⁸ In De Graeve's (2008) model, changes in the external finance premium are largely driven by shocks to investment-specific technology, which affect entrepreneurs' capital formation and therefore their demand for external funds.

contributed to the fact that the external finance premium rose before the 2001 recession, a fact documented by Levin, Natalucci, and Zakrajsek (2004) and De Graeve (2008).²⁹

Finally, we turn back to the empirical correlation between firm volatility and the business cycle. Bloom, Floetotto, and Jaimovich (2009) document that in a sample that starts in the early 1960s, the interquartile range of sales growth rates is negatively correlated with real GDP growth. This negative correlation appears to be driven by the fact that sales growth dispersion tends to spike upwards during NBER recessions. Indeed, to the extent that the financial frictions and irreversibility models primarily speak for the possibility that sharp increases in firm-specific volatility cause recessions, their empirical validation primarily rests on a negative relation between volatility (at the firm and aggregate level) and aggregate output during recessions.

In the previous subsection, we discussed that in a sample starting in the mid-1980s, our measure for firm-specific volatility and dispersion relates positively to current and lagged output gaps. We do not interpret the pro-cyclicality of firm volatility in our sample as evidence against the financial frictions and irreversibility theories. To see why, recall that our sample only contains two NBER recessions, and that firm-specific volatility was comparatively low in the recession of the early 1990s.

Our intuition is that a negative effect of firm volatility on aggregate output, as predicted by financial frictions and irreversibility theories, may co-exist with a positive effect of aggregate output on firm-specific volatility, such that the empirical correlation could in principle be of either sign. In the remainder of this section, we provide intuition for such a positive effect of aggregate output on firm volatility.

Our finding that firm-specific volatility gradually rises during benign periods and decreases after an adverse aggregate event suggests that, while firm-specific shocks are by definition exogenous, firms and their contractual counterparties can influence the variance of these firm-specific shocks. After a large surprise decline in earnings, a firm may tend to assign a larger probability value to large negative shocks. Therefore, it may choose to invest in less risky projects, or spend more resources on hedging. According to this intuition, adverse aggregate events simultaneously make all firms less prone to take on projects that are ex ante risky. This may explain why ex post firm

²⁹ Levin, Natalucci, and Zakrajsek (2004) find that changes in the external finance premium during that episode are largely determined by changes in the expected cost of default. Recall from the previous footnote that in De Graeve (2008), changes in the external finance premium are largely driven by investment-specific technology shocks.

volatility gradually declines after adverse aggregate events as the fraction of capital corresponding to less risky investment gradually increases.

6 Conclusion

This paper develops and implements a method to estimate cyclical changes in firm-specific volatility and the associated dispersion in sales, earnings and employment growth.

Our new measure uncovers substantial cyclical swings in firm volatility around trend. These changes in firm-specific volatility have important macroeconomic implications. We find that US firms were comparatively stable in the early 1990s, but were volatile around 2000. In endogenous pricing models, this tends to imply that firms adjusted prices less frequently in the earlier episode than in the later episode. Consistent with that prediction, we find that reducing inflation by any given amount was three times more costly in the early 1990s than it was in the early 2000s.

Our findings suggest that firm-specific volatility did not rise in the early 1990s. Both from the perspective of financial accelerator models and from the point of view of irreversibility models, this implies that the 1990-91 recession resulted from aggregate rather than firm-specific factors. In particular, changes in firm-specific volatility did not contribute to the rise in the external finance premium in the second half of the 1980s. Similarly, there was no rise in firm-specific uncertainty in the early 1990s that would have tended to cause firms to postpone investment and hiring.

On the other hand, our evidence suggests that rising firm-specific volatility contributed to the 2001 recession. Firstly, our results are consistent with the hypothesis that rising firm-specific volatility in the late 1990s contributed to the coincident rise in the relative cost of external finance. Secondly, the same rise in firm-specific volatility would have tended to dampen investment and hiring because of a wait-and-see effect.

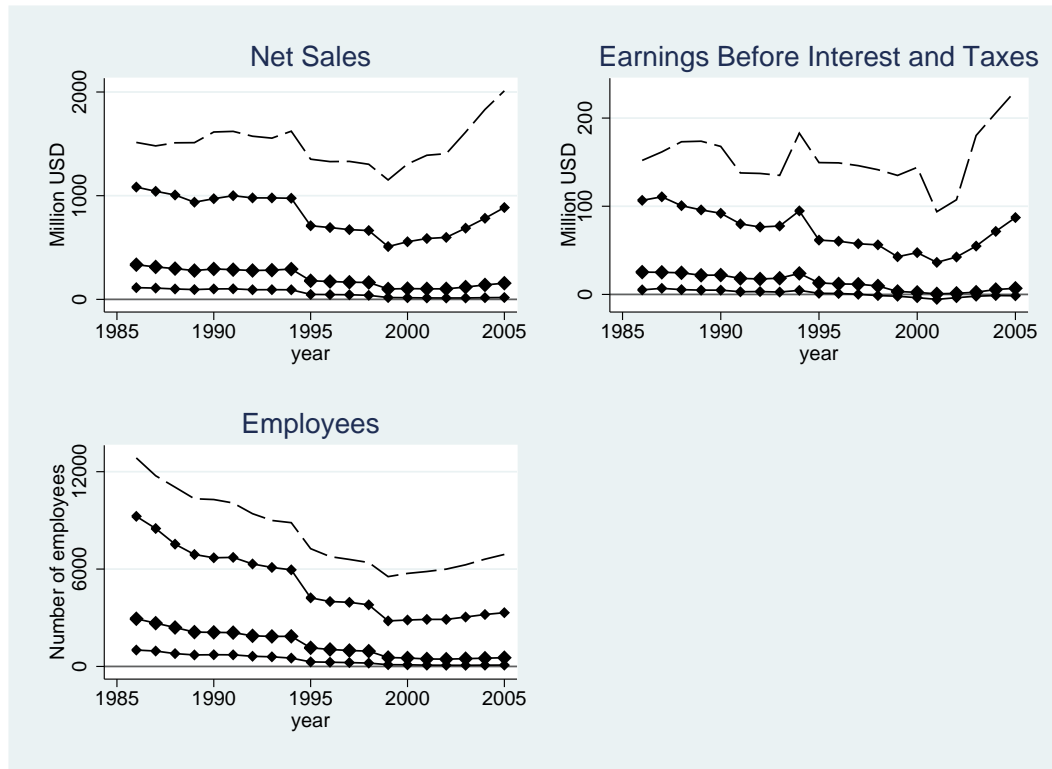
In our sample, firm volatility tends to gradually decrease in the aftermath of adverse aggregate events, and gradually increases during prolonged booms. Because of that pattern, we find that firm-specific volatility is positively correlated with current and lagged aggregate output. To the extent that financial accelerator and irreversibility models primarily suggest that sharp increases in firm-specific volatility cause recessions, the fact that we find a

positive correlation in between recession periods does not constitute evidence against those theories.

We conjecture that this positive correlation reflects the possibility that fluctuations in aggregate output affect firm-specific volatility. In particular, an adverse aggregate event may make firms less prone to undertake projects that are ex ante risky. On the other hand, a protracted boom may cause firms to believe that negative events are less likely, which could imply disproportionately more investment in risky projects.

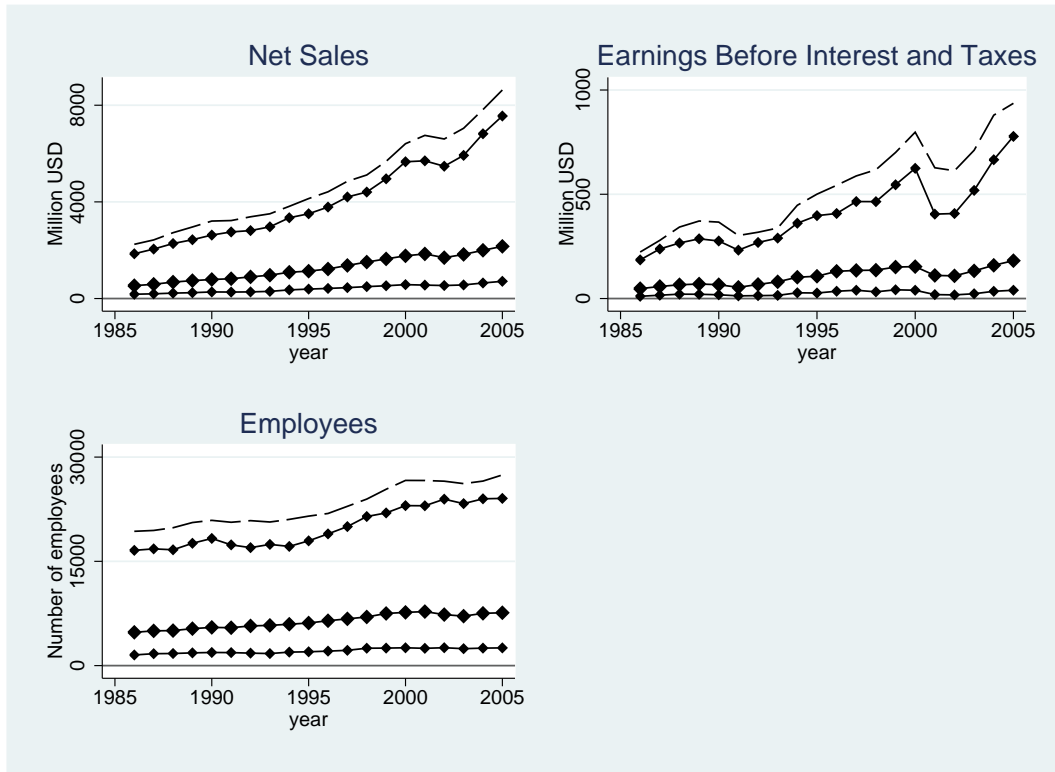
Appendix 1: Figures and Tables

Figure 1
Unbalanced Panel: Mean and Quartiles of Levels Data



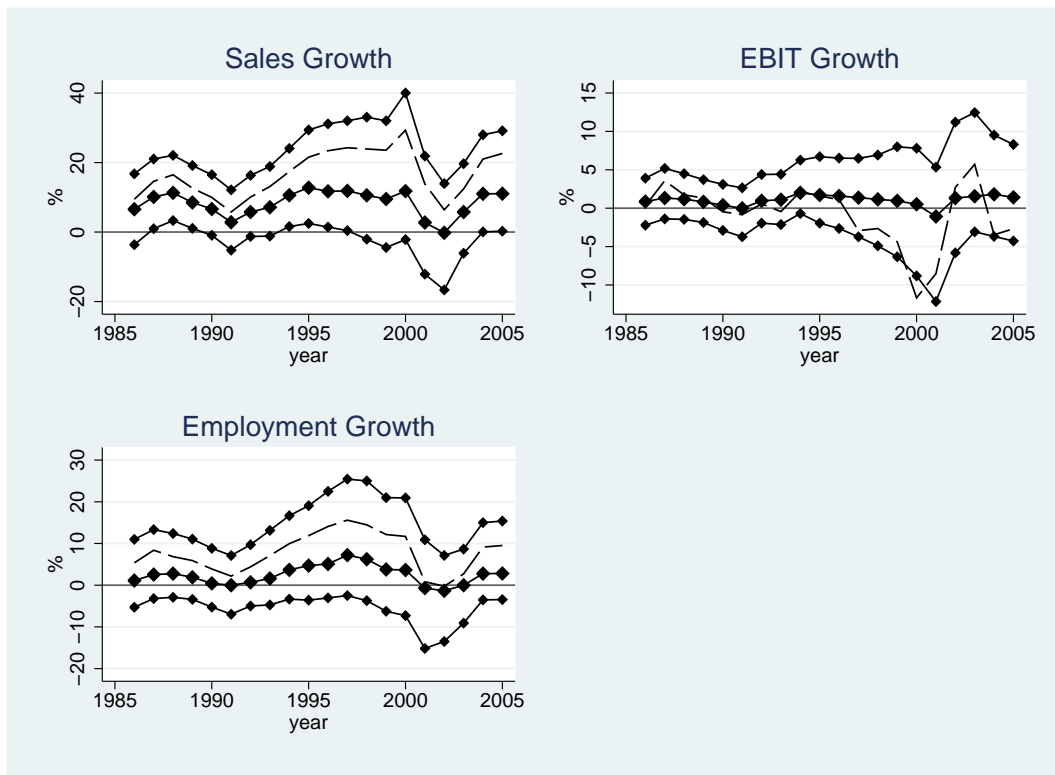
Note: This figure documents the distribution of net sales, earnings and the number of employees in every year for the period 1986-2005 for an unbalanced sample of 15,425 listed US firms from Worldscope. Large diamonds indicate the median, while small diamonds indicate the 25th and 75th percentiles. The dashed line is the mean. Net sales and Earnings Before Interest and Taxes (EBIT) are in nominal terms, and expressed in million USD. Employees stands for the number of employees. By any measure, firm size in the unbalanced panel gradually reduces through 2000. This reflects a change in sample composition: as time progresses, ever smaller firms enter the sample. See Section 2 for more details. That section also describes our approach to controlling for changes in sample composition.

Figure 2
Balanced Panel: Mean and Quartiles of Levels Data



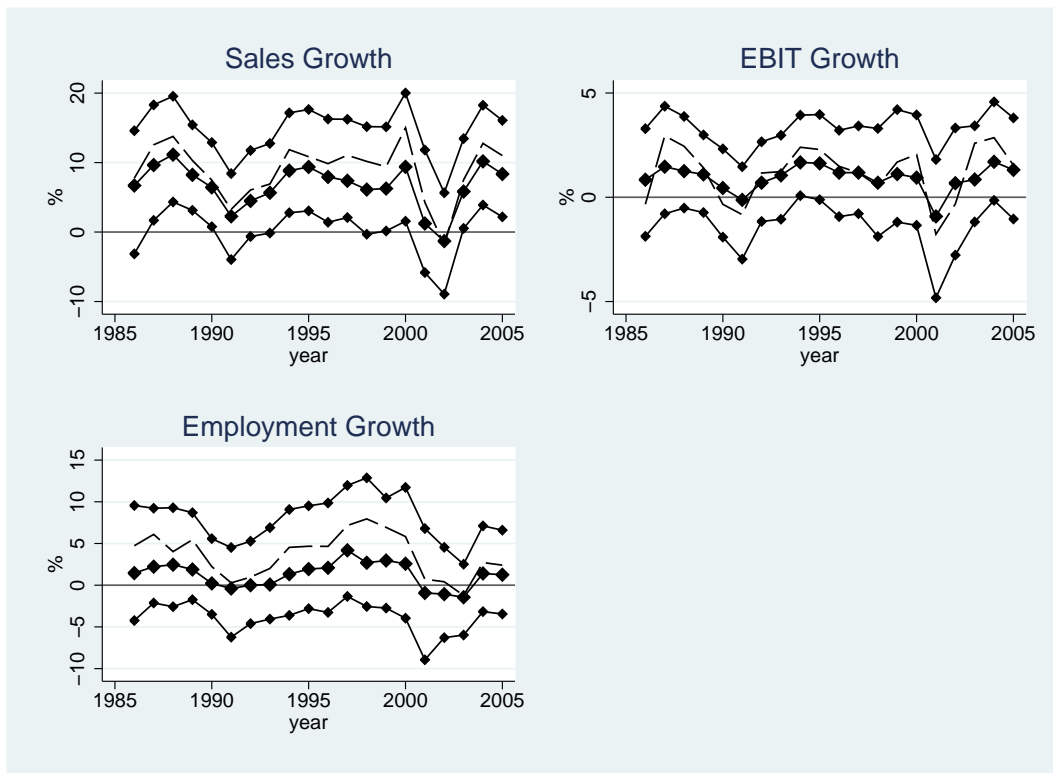
Note: This figure documents the distribution of net sales, earnings and the number of employees in every year for the 588 US firms in our Worldscope sample for which data are continuously available over the period 1986-2005. Large diamonds indicate the median, while small diamonds indicate the 25th and 75th percentiles. The dashed line is the mean. Units are as in Figure 1. Figure 2 indicates a stable upward trend in sales, earnings and the number of employees in the balanced panel, with the exception of slowdowns corresponding to the recessions of the early 1990s and early 2000s.

Figure 3
Unbalanced Panel: Mean and Quartiles of Growth Rates



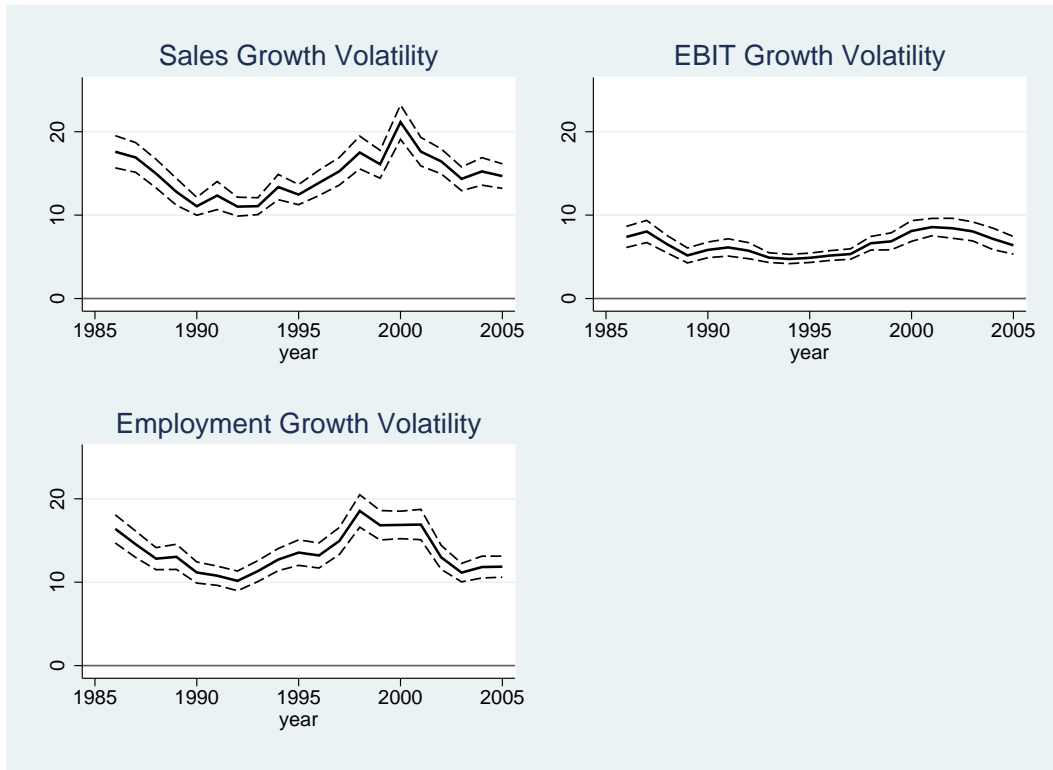
Note: This figure documents the distribution of the growth rates of net sales, earnings and the number of employees in every year over the period 1986-2005 for an unbalanced sample of 15,425 US firms from Worldscope. Large diamonds indicate the median, while small diamonds indicate the 25th and 75th percentiles. The dashed line is the mean. All growth rates are in percentage terms. The firm-level data accurately indicate slowdowns during the recessions of the early 1990s and early 2000s.

Figure 4
Balanced Panel: Mean and Quartiles of Growth Rates



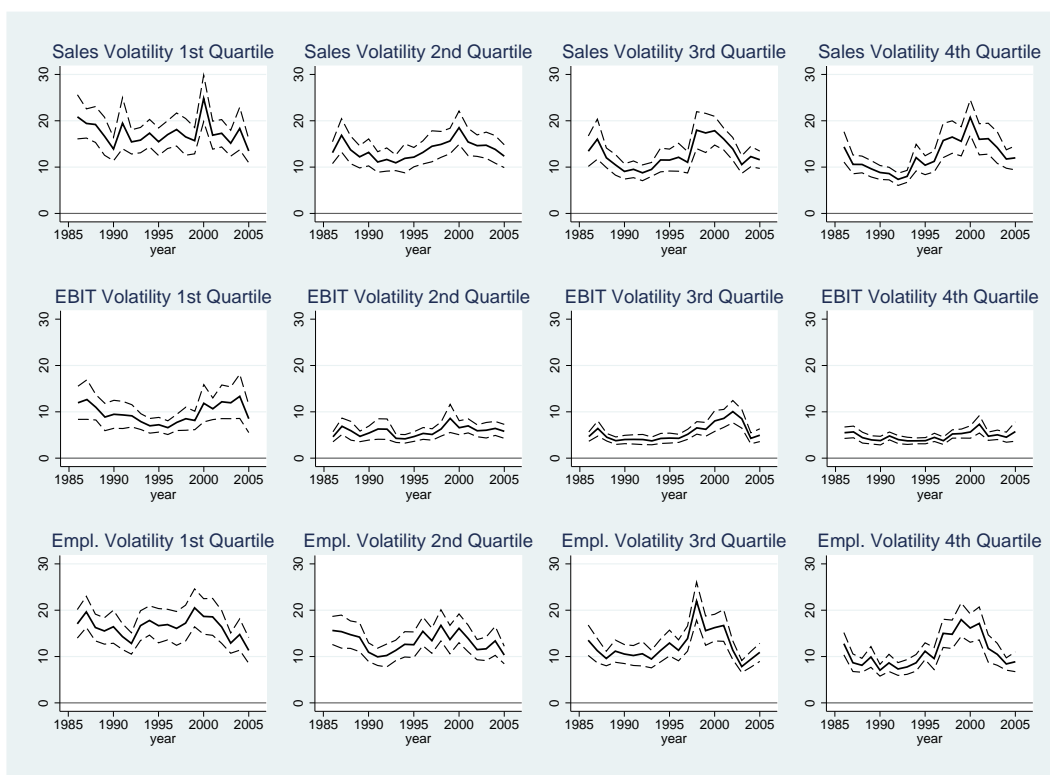
Note: This figure documents the distribution of the growth rates of net sales, earnings and the number of employees in every year for the 588 US firms in our Worldscope sample for which data are continuously available over the period 1986-2005. Large diamonds indicate the median, while small diamonds indicate the 25th and 75th percentiles. The dashed line is the mean. All growth rates are in percentage terms. Taking into account the scale difference with Figure 3, Figure 4 indicates that dispersion in the firm growth distribution is larger in the unbalanced panel. Comparing Table 2 with Table 3, or Figure 5 with Figure 8, we see that firm-specific volatility, and the dispersion it implies, are indeed substantially higher in the unbalanced panel.

Figure 5
Balanced Panel: Firm-Level Volatility



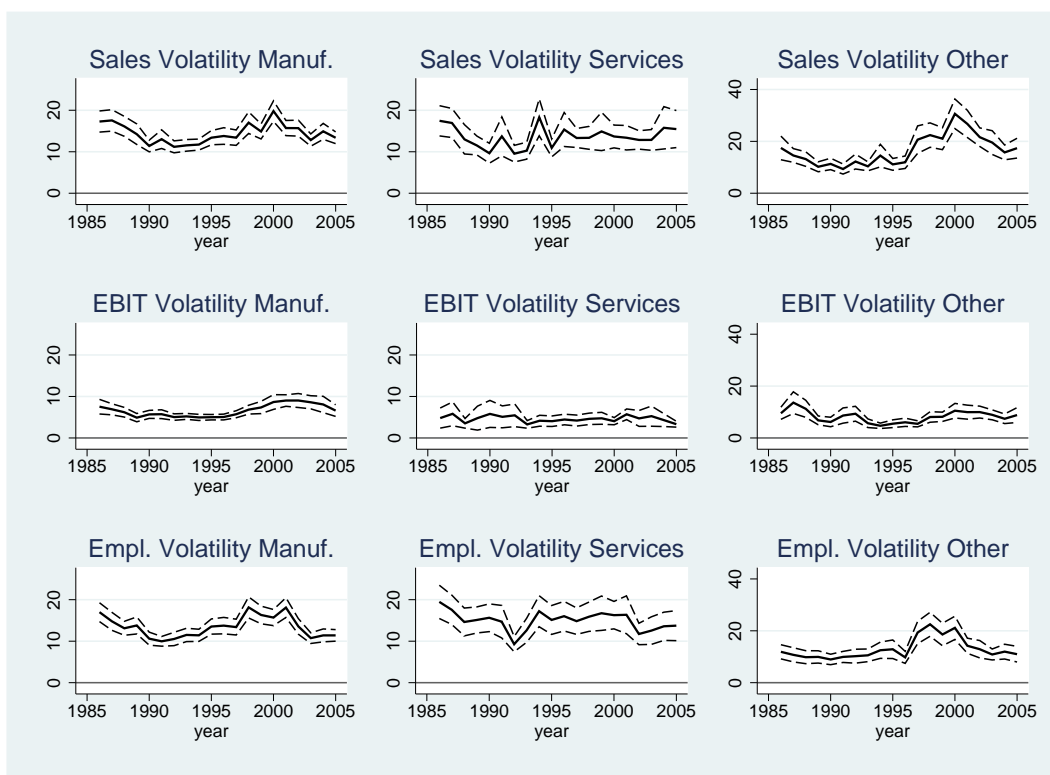
Note: This figure presents our main results. It graphs firm-level sales, earnings, and employment growth volatility for every year for the balanced panel of 588 US firms, along with a 95 percent confidence interval. Volatility is estimated from equations (3) through (5). The figure graphs the estimated time effects in the second-stage equation (5). The confidence intervals are computed from the corresponding, heteroskedasticity-robust, regression-based standard errors. US firms became gradually more stable in the latter half of the 1980s, became gradually more volatile during the 1990s, and gradually stabilised again in the first half of the 2000s. See Table 2 below for corresponding statistics.

Figure 6
Balanced Panel: Firm-Level Volatility by Firm Size Quartile



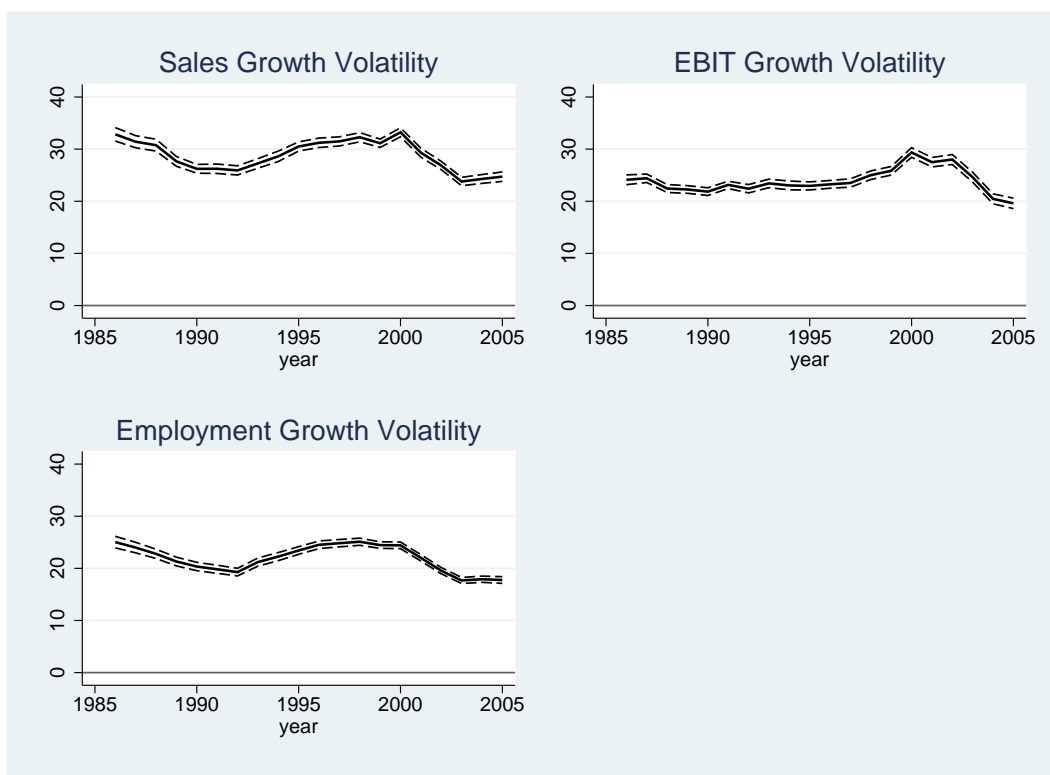
Note: This figure graphs firm-level sales, earnings, and employment volatility when subdividing the balanced panel's observations into firm size quartiles for 1986-2005 according to the sales-to-GDP ratio. The left column applies to the quartile of smallest firm sizes, and firm size increases as we consider columns further to the right. Volatility for every size quartile is estimated using equations (3) through (5), with the only difference that we omit the firm size interaction term from equation (3). Irrespective of firm size, volatility tends to decrease in the second half of the 1980s, tends to increase in the 1990s, and tends to decrease in the first half of the 2000s. This suggests that the results from the overall balanced panel are not driven by changes in the firm size composition over the course of the sample.

Figure 7
Balanced Panel: Firm-Level Volatility by Sector



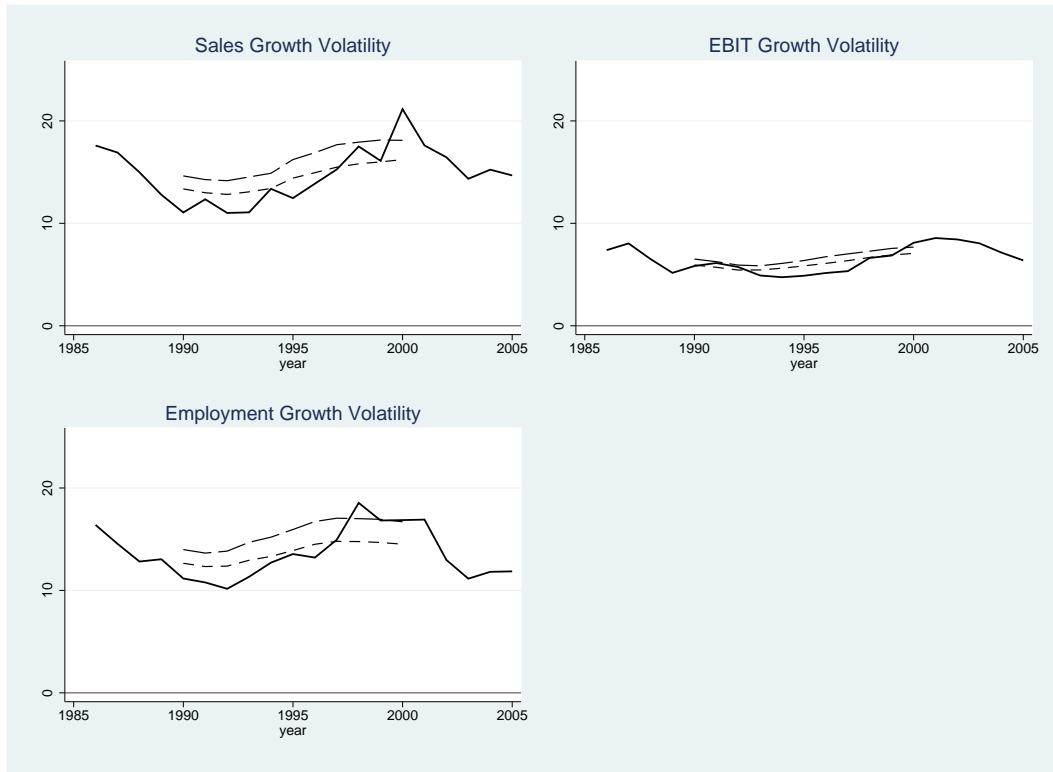
Note: This figure graphs firm-level sales, earnings, and employment volatility when dividing the 588 US firms in the balanced panel into three broad sectors. The left column pertains to manufacturing firms, the middle column to service providers (including wholesalers and retail traders), and the third to firms in other sectors. We define the three sectors in Section 3.2. We estimate volatility for every sector using equations (3) through (5), but we omit the sector interaction term from equation (3). This figure suggests that our main results, in Figure 5, mostly speak for manufacturing firms.

Figure 8
Unbalanced Panel: Firm-Level Volatility



Note: This figure graphs firm-level sales, earnings, and employment growth volatility for all 15,425 US firms in our Worldscope sample for 1986-2005, along with a 95 percent confidence interval. In this case, we estimate volatility using equations (3), (4) and (6). For the unbalanced panel, this figure reveals a decrease in volatility in the second half of the 1980s, rising volatility in the 1990s, and declining volatility in the 2000s. Therefore, our main conclusion, noted under Figure 5, is robust to using the unbalanced panel. This further supports the notion that our main conclusion is not driven by changes in the firm size composition. See Table 3 for corresponding statistics.

Figure 9
Balanced Panel: Annual versus Rolling Firm-Level Volatility



Note: In every panel, the solid line is our measure of annual volatility in the balanced panel of 588 US firms, as graphed previously in Figure 5. The short dashed lines stand for ten-year moving averages of annual volatility. The long dashes stand for rolling ten-year standard deviations computed using equations (7) and (8) on the same balanced panel. Apart from a levels difference, rolling ten-year standard deviations are very similar to the moving average of our annual volatility measure. As we discuss in Section 5.1, this suggests that our measure reveals new information by estimating the cycle as well as the trend in firm volatility.

Table 1
Convergence of Estimator for Firm Volatility

N =Number of firms	2.5th percentile	median	97.5th percentile
100	8.32	9.98	11.78
588	9.30	10.00	10.72
1,000	9.47	10.00	10.54
1,629	9.58	10.00	10.43
6,025	9.78	10.00	10.22
10,000	9.83	10.00	10.17

Note: This table documents convergence of our estimator for firm-specific volatility based on equation (4). For different sample sizes N , this table shows the median of the estimator $(1/N) \sum_{i=1}^N \hat{\sigma}_{\varepsilon,it}$ for the cross-sectional average of firm-specific volatility, along with the 2.5th and 97.5th percentiles. We obtain the percentiles by estimating $(1/N) \sum_{i=1}^N \hat{\sigma}_{\varepsilon,it}$ one million times from a simulated population of firms with true firm-specific standard deviations $\sigma_{\varepsilon,it}$ distributed uniformly over the support 0 to 20. In this case, true average firm-level volatility $(1/N) \sum_{i=1}^N \sigma_{\varepsilon,it} = 10$. The simulated sample sizes are chosen to match our empirical sample sizes. $N = 588$ in the balanced panel of firms. In the unbalanced panel of firms, sample size varies between $N = 1,629$ in 1986 and $N = 6,025$ in 1999. For all simulated sample sizes in the table, the median estimator is at or near the true value of 10.00. This table indicates convergence, in the sense that the 95 percent probability intervals around the median shrink as N grows large. These probability intervals are narrow compared to the estimated changes in volatility that we show in Figure 5. This indicates that our estimator converges sufficiently quickly in order to allow us to draw inference about changes in volatility from our empirical samples.

Table 2
Balanced Panel: Firm-Level Volatility

	Average volatility 1986-1995	Average volatility 1996-2005	Change in volatility	F-statistic for significant change
Sales growth	13.36 (0.24)	16.22 (0.27)	2.86** (0.36)	62.23 [0.00]
EBIT growth	5.93 (0.15)	7.07 (0.17)	1.13** (0.23)	25.03 [0.00]
Employment growth	12.65 (0.22)	14.51 (0.25)	1.86** (0.34)	30.39 [0.00]

Note: This table pertains to Figure 5. Standard errors of F-tests for significance are in round brackets, and p-values are in square brackets. The first column indicates average firm-specific volatility over a first subsample (1986-1995), and the second indicates average volatility over the subsample 1996-2005. The third column reports the change in volatility from the first to the second subsample. ** indicates significance at the 1 percent level. The fourth column reports the F-statistic and p-value for the null hypothesis of no change in volatility between the two subsamples.

Table 3
Unbalanced Panel: Firm-Level Volatility

	Average volatility 1986-1995	Average volatility 1996-2005	Change in volatility	F-statistic for significant change
Sales growth volatility	28.75 (0.22)	28.84 (0.09)	0.09 (0.31)	0.08 [0.78]
EBIT growth volatility	23.01 (0.22)	24.72 (0.09)	1.71** (0.31)	30.80 [0.00]
Employment growth volat.	21.95 (0.18)	21.83 (0.08)	-0.12 (0.26)	0.22 [0.64]

Note: This table pertains to Figure 8. Standard errors are in round brackets, and p-values in square brackets. Other notes are as under Table 2.

Table 4
Sacrifice Ratios

Disinflation episode	Inflation loss	Sacrifice ratio
1980Q1-1983Q4	-8.56	1.94
1989Q4-1994Q3	-2.60	2.80
2000Q2-2002Q2	-1.46	0.92

Note: This table shows sacrifice ratios for three disinflation episodes in the United States, computed as in Ball (1994). We use CPI inflation. For every disinflation episode, the leftmost column indicates the dates of the inflation peak and trough. The middle column indicates the total decline in trend inflation from peak to trough. The rightmost column indicates the sacrifice ratio, defined as the cumulative annual loss in real GDP required to reduce trend inflation by one percentage point. As we discuss in Section 5.3, our finding that the sacrifice ratio for the disinflation of the early 1990s is about three times larger than the ratio for the early 2000s suggests that firm-specific volatility affected the output-inflation trade-off as predicted by endogenous pricing theories.

Appendix 2: Unbiasedness Proof

This appendix proves unbiasedness of our estimator for the firm-level standard deviation, stated in equation (4). In particular, we prove that $E(\widehat{\sigma}_{\varepsilon,it}) = \sigma_{\varepsilon,it}$.

Taking expectations of equation (4), but assuming that we know the true error term $\varepsilon_{i,t}$, we obtain:

$$E(\widehat{\sigma}_{\varepsilon,it}) = \sqrt{\frac{\pi}{2}} E(|\varepsilon_{i,t}|) \quad (\text{A1})$$

Writing the expectation $E(|\varepsilon_{i,t}|)$ out as a function of the probability density function $f(\varepsilon_{i,t})$ yields:

$$E(\widehat{\sigma}_{\varepsilon,it}) = \sqrt{\frac{\pi}{2}} \int_{-\infty}^{+\infty} |\varepsilon_{i,t}| f(\varepsilon_{i,t}) d\varepsilon_{i,t} \quad (\text{A2})$$

Assuming that the error term is normally distributed with mean zero, i.e. $\varepsilon_{i,t} \sim N(0, \sigma_{\varepsilon,it}^2)$, this implies:

$$E(\widehat{\sigma}_{\varepsilon,it}) = \sqrt{\frac{\pi}{2}} \int_{-\infty}^{+\infty} |\varepsilon_{i,t}| \frac{1}{\sigma_{\varepsilon,it} \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,it}} \right)^2} d\varepsilon_{i,t} \quad (\text{A3})$$

Since $|\varepsilon_{i,t}| = |-\varepsilon_{i,t}|$, i.e. the absolute value is a function that is symmetric around the vertical axis, equation (A3) is equivalent to:

$$E(\widehat{\sigma}_{\varepsilon,it}) = \sqrt{\frac{\pi}{2}} 2 \int_0^{+\infty} \varepsilon_{i,t} \frac{1}{\sigma_{\varepsilon,it} \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,it}} \right)^2} d\varepsilon_{i,t} \quad (\text{A4})$$

Since we can rewrite $|\varepsilon_{i,t}| = \varepsilon_{i,t}$ for $\varepsilon_{i,t} \geq 0$, and after bringing the term $1/(\sigma_{\varepsilon,it} \sqrt{2\pi})$ outside the integral, we obtain:

$$E(\widehat{\sigma}_{\varepsilon,it}) = \frac{1}{\sigma_{\varepsilon,it}} \int_0^{+\infty} \varepsilon_{i,t} e^{-\frac{1}{2} \left(\frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,it}} \right)^2} d\varepsilon_{i,t} \quad (\text{A5})$$

Since the antiderivative of $\varepsilon_{i,t} e^{-\frac{1}{2} \left(\frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,it}} \right)^2}$ is $-\sigma_{\varepsilon,it}^2 e^{-\frac{1}{2} \left(\frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,it}} \right)^2}$, the fundamental theorem of calculus implies:

$$E(\widehat{\sigma}_{\varepsilon,it}) = \frac{1}{\sigma_{\varepsilon,it}} \left[-\sigma_{\varepsilon,it}^2 e^{-\frac{1}{2} \left(\frac{\varepsilon_{i,t}}{\sigma_{\varepsilon,it}} \right)^2} \right]_0^{+\infty} \quad (\text{A6})$$

Solving the right-hand side yields $(1/\sigma_{\varepsilon,it}) \left[0 - (-\sigma_{\varepsilon,it}^2) \right]$, which in turn equals $\sigma_{\varepsilon,it}$. Therefore,

$$E(\widehat{\sigma}_{\varepsilon,it}) = \sigma_{\varepsilon,it} \quad (\text{A7})$$

which proves unbiasedness.

References

- Ahmed, S, A Levin and B A Wilson (2004), “Recent U.S. Macroeconomic Stability: Good Policies, Good Practices, or Good Luck?”, *Review of Economics and Statistics*, 86, 824-832.
- Bachmann, R and C Bayer (2011), “Uncertainty Business Cycles - Really?”, NBER Working Paper 16862.
- Bachmann, R, S Elstner and E R Sims (2010), “Uncertainty and Economic Activity: Evidence from Business Survey Data”, NBER Working Paper 16143.
- Ball, L (1994), “What Determines the Sacrifice Ratio?” in *Monetary Policy*, ed. N G Mankiw, University of Chicago Press.
- Ball, L, N G Mankiw and D Romer (1988), “The New Keynesian Economics and the Output-Inflation Trade-Off”, *Brookings Papers on Economic Activity*, 1, 1-82.
- Ball, L and S Mazumder (2011), “The Evolution of Inflation Dynamics and the Great Recession”, *Brookings Papers on Economic Activity*, Spring 2011.
- Blanchard, O and J Simon (2001), “The Long and Large Decline in U.S. Output Volatility”, *Brookings Papers on Economic Activity*, 1, 135-174.
- Bloom, N (2009), “The Impact of Uncertainty Shocks”, *Econometrica*, 77, 623-685.
- Bloom, N, M Floetotto and N Jaimovich (2009), “Really Uncertain Business Cycles”, mimeo, Stanford University.
- Boivin, J, M P Giannoni and I Mihov (2009), “Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data”, *American Economic Review*, 99, 350-384.

- Brown, G and N Kapadia (2007), “Firm-Specific Risk and Equity Market Development”, *Journal of Financial Economics*, 84, 358-388.
- Burstein, A and C Hellwig (2008), “Welfare Costs of Inflation in a Menu Cost Model”, *American Economic Review: Papers and Proceedings*, 98, 438-443
- Campbell, J Y, M Lettau, B G Malkiel and Y Xu (2001), “Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk”, *Journal of Finance*, LVL, 1-43.
- Castro, R, G L Clementi and Y Lee (2009), “Cross-Sectoral Variation in Firm-Level Idiosyncratic Risk”, Federal Reserve Bank of Cleveland Working Paper 08-12.
- Cecchetti, S G, A Flores-Lagunes and S Krause (2006), “Assessing the Sources of Changes in the Volatility of Real Growth”, NBER Working Paper 11946.
- Christiano, L J, R Motto and M Rostagno (2010), “Financial Factors in Economic Fluctuations”, ECB Working Paper 1192.
- Comin, D and S Mulani (2006), “Diverging Trends in Aggregate and Firm Volatility”, *Review of Economics and Statistics*, 88, 374-383.
- Comin, D and T Philippon (2005), “The Rise in Firm-Level Volatility: Causes and Consequences”, *NBER Macroeconomics Annual* 2005.
- Costain, J and A Nakov (2011), “Price Adjustments in a General Model of State-Dependent Pricing”, *Journal of Money, Credit and Banking*, 43, 385-406.
- Davis, S J, J Haltiwanger, R Jarmin and J Miranda (2006), “Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms”, *NBER Macroeconomics Annual* 2006.
- Davis, S J and J A Kahn (2008), “Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels”, *Journal of Economic Perspectives*, 22, 155-180.
- DeFina, R H (1991), “International Evidence on a New Keynesian Theory of the Output-Inflation Trade-Off”, *Journal of Money, Credit and Banking*, 23, 410-422.
- De Graeve, F (2008), “The External Finance Premium and the Macroeconomy: US Post-WWII Evidence”, *Journal of Economic Dynamics and Control*, 32, 3415-3440.
- Den Haan, W J and V Sterk (2011), “The Myth of Financial Innovation and the Great Moderation”, *Economic Journal*, 121, 707-739.
- De Veirman, E (2009), “What Makes the Output-Inflation Trade-Off Change? The Absence of Accelerating Deflation in Japan”, *Journal of Money, Credit and Banking*, 41, 1117-1140.

- De Veirman, E and A Levin (2009), “Measuring Changes in Firm-Level Volatility: an Application to Japan”, Reserve Bank of New Zealand Discussion Paper 2009/20.
- Dotsey, M, R G King and A L Wolman (2011), “Inflation and Real Activity with Firm-Level Productivity Shocks”, mimeo, Federal Reserve Bank of Richmond.
- Dynan, K E, D W Elmendorf and D E Sichel (2006), “Can Financial Innovation Help to Explain the Reduced Volatility of Economic Activity?”, *Journal of Monetary Economics*, 53, 123-150.
- Fama, E F and K R French (2004), “New Lists: Fundamentals and Survival Rates”, *Journal of Financial Economics*, 73, 229-269.
- Fink, J, K E Fink, G Grullon and J P Weston (2010), “What Drove the Increase in Idiosyncratic Volatility During the Internet Boom?”, *Journal of Financial and Quantitative Analysis*, 45, 1253-1278.
- Gali, J and L Gambetti (2009), “On the Sources of the Great Moderation”, *American Economic Journal: Macroeconomics*, 1, 26-57.
- Gertler, M and J Leahy (2008), “A Phillips Curve with an Ss Foundation”, *Journal of Political Economy*, 116, 533-572.
- Gilchrist, S, J Sim and E Zakrajsek (2011), “Uncertainty, Financial Frictions, and Investment Dynamics”, mimeo, Boston University.
- Golosov, M and R E Lucas (2007), “Menu Costs and Phillips Curves”, *Journal of Political Economy*, 115, 171-199.
- Hamao, Y, J Mei and Y Xu (2007), “Unique Symptoms of Japanese Stagnation: an Equity Market Perspective”, *Journal of Money, Credit and Banking*, 39, 901-923.
- Justiniano, A and G E Primiceri (2008), “The Time-Varying Volatility of Macroeconomic Fluctuations”, *American Economic Review*, 98, 604-641.
- Klenow, P J and J L Willis (2006), “Real Rigidities and Nominal Price Changes”, Federal Reserve Bank of Kansas City Research Working Paper 06-03.
- Klenow, P J and J L Willis (2007), “Sticky Information and Sticky Prices”, *Journal of Monetary Economics*, 54, 79-99.
- Kim, C and C R Nelson (1999), “Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle”, *Review of Economics and Statistics*, 81, 608-616.
- Levin, A T, F M Natalucci and E Zakrajsek (2004), “The Magnitude and Cyclical Behavior of Financial Market Frictions”, Federal Reserve Board, Finance and Economics Discussion Series 2004-70.

Mankiel, B G and Y Xu (2003), “Investigating the Behavior of Idiosyncratic Volatility”, *Journal of Business*, 76, 613-644.

McConnell, M M and G Perez-Quiros (2000), “Output Fluctuations in the United States: What Has Changed Since the Early 1980’s?”, *American Economic Review*, 90, 1464-1476.

Senda, T and J K Smith (2008), “Inflation History and the Sacrifice Ratio: Episode-Specific Evidence”, *Contemporary Economic Policy*, 26, 409-419.

Stock, J H and M W Watson (1999), “Business Cycle Fluctuations in US Macroeconomic Time Series” in *Handbook of Macroeconomics*, eds. J B Taylor and M Woodford, Elsevier.

Stock, J H and M W Watson (2002), “Has the Business Cycle Changed and Why?”, *NBER Macroeconomics Annual* 2002.

Vavra, J (2011), “Inflation Dynamics and Time-Varying Uncertainty: New Evidence and an Ss Interpretation”, mimeo, Yale University.