What Drives Business Cycle Fluctuations: Aggregate or Idiosyncratic Uncertainty Shocks?

Bijapur, Mohan

London School of Economics and Political Science

1 November 2014

What Drives Business Cycle Fluctuations: Aggregate or Idiosyncratic Uncertainty Shocks?

Mohan Bijapur [a]

London School of Economics

November 2014

Abstract

We study jointly the roles of aggregate and idiosyncratic uncertainty shocks in driving business cycle fluctuations. By decomposing total stock return volatility of over 20,000 publicly-listed US firms from 1962 to 2012, we construct separate indices for aggregate and idiosyncratic uncertainty, and run a horse race between them in an otherwise standard macroeconomic VAR. We find that idiosyncratic uncertainty shocks account for a large fraction of fluctuations in economic activity at business cycle frequencies, whereas the impacts of aggregate uncertainty are negligible. Idiosyncratic uncertainty, and not aggregate uncertainty, shocks produce the "sharp drop and rapid rebound" response in activity characterized in Bloom (2009). Idiosyncratic uncertainty shocks to large firms have more powerful macroeconomic impacts than small firms, suggesting "Granular" origins to the role of uncertainty in the macroeconomy. We also find evidence of an economy-wide "buffering effect", in which the effects of large and small firms' shocks exhibit a negative covariance which dampens down the aggregate effects of idiosyncratic uncertainty shocks on economic fluctuations.

Keywords: Idiosyncratic uncertainty shocks; aggregate uncertainty shocks; business cycles; Granular origins.

JEL classification: E32.

---

[a] Department of Economics, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK. E-mail address: M.Bijapur@lse.ac.uk. Tel: +44 207 955 6488.

I would like to thank Marios Angeletos, Gianluca Benigno, Xavier Gabaix, Seb Graves, Wouter den Haan, Ricardo Reis, Roman Sustek, Silvana Tenreyro and Shengxing Zhang for extremely helpful discussions and feedback.

1. Introduction

In recent times, much attention has been paid to the role played by uncertainty in the real economy. Of particular concern has been the sluggish nature of the US economy's recent recovery from the Great Recession, which many commentators have attributed in large part to heightened levels of uncertainty, with, for example, the September 2012 Federal Open Market Committee (FOMC) meeting noting that *"a high level of uncertainty regarding the European fiscal and banking crisis and the outlook for US fiscal and regulatory policies was weighing on confidence, thereby restraining household and business spending,"* and former Chairmen Ben Bernanke[1] indicating that *"increases in risk aversion and uncertainty, together with tight credit conditions, may have impeded the commercial application of new technologies and slowed the pace of business formation."*

Much of the existing debate has centred on the impacts of aggregate (or macroeconomic) uncertainty. However, it has long been known that idiosyncratic uncertainty also exhibits significant stochastic time variation (notably from the well-known paper of Campbell, Lettau, Malkiel and Xu (2001), henceforth referred to as CLMX (2001)). In this paper, we study jointly the roles of aggregate and idiosyncratic uncertainty shocks in driving business cycle fluctuations. We construct separate indices for aggregate and idiosyncratic uncertainty, and run a horse race between them in an otherwise standard macroeconomic VAR model. Our results are quite striking – whereas idiosyncratic uncertainty shocks are found to be a key driver of the business cycle, accounting for a large fraction of fluctuations in economic activity, the effects of aggregate uncertainty shocks are negligible.

Following the seminal paper of Bloom (2009), a burgeoning recent literature has emerged examining the effects of shocks to uncertainty in driving macroeconomic fluctuations[2]. The novel contributions of our paper can be outlined as follows. First, the existing literature has yielded somewhat inconclusive findings, with numerous papers finding aggregate

---

[1] Former Chairman of the Federal Reserve Ben Bernanke's speech at the Economic Club of New York, November 20, 2012.

[2] Other innovations to the conventional approaches of modelling technology, preferences and policy shocks have been proposed recently. In particular, in a ground-breaking paper, Angeletos et al. (2014) explore the role of variation in "confidence" about the state of the economy, finding it accounts for about one half of GDP uncertainty at business-cycle frequencies.

uncertainty to be a dominant force in business cycle fluctuations (Alexopoulos and Cohen (2009); Baker et al. (2013); Baker and Bloom (2012); Benigno et al. (2012); Bloom (2009), Caggiano et al. (2014); Gilchrist et al. (2013); Mumtaz and Surico (2013); Mumtaz and Zanetti (2013); while others find it has little impact at the macroeconomic level (Bachman and Bayer (2013); Bachmann et al. (2013); Bekaert et al. (2013); Born and Pfeifer (2014)) and yet others more attribute these effects to idiosyncratic, not aggregate, uncertainty (Arellano, Bai and Kehoe (2012); Christiano, Motto and Rostagno (2012)). But the majority of these papers study the effects of either aggregate uncertainty or idiosyncratic uncertainty in isolation. Our paper is different in that we combine both sources of uncertainty in a single model and analyze jointly their effects, in order to shed light on their relative importance in explaining macroeconomic fluctuations. Bloom et al. (2012)) appears to be the only other paper which studies both shocks in tandem. They quantify the impact of time-varying uncertainty on the economy in a dynamic stochastic general equilibrium (DSGE) model with heterogeneous firms, but they find that both types of uncertainty have significant impacts on economic activity. Their approach differs in three respects, first they estimate their uncertainty indices by measuring the dispersion of plant-level shocks to total factor productivity (TFP), whereas we use stock return data to measure uncertainty. Second, our modelling choices differ, they formulate a DSGE model, whereas we estimate VAR models. Third, other than the variance of shocks, aggregate and idiosyncratic uncertainty are observationally equivalent in their DSGE model, hence they do not allow for potential differences in the propagation mechanisms and dynamic responses of the two types of uncertainty shocks.

Our second contribution is to investigate the transmission mechanisms of uncertainty shocks to the real economy. The existing literature identifies three key channels – first, in Bloom (2009), greater uncertainty increases the real option value to waiting, hence firms scale back their investment and hiring plans. This induces a characteristic “sharp drop and rapid rebound” profile of responses. There is a sharp drop in activity, as firms scale back their plans rapidly (since expectations change with the onset of the uncertainty shock), followed by a rapid rebound, as firms address their pent-up demand for labour and capital, after the uncertainty shock subsides. Second, in Christiano, Motto and Rostagno (2012) (henceforth referred to as CMR), higher idiosyncratic risk induces tighter credit conditions, hence borrowing firms are constrained to scale back their investment plans. Third, in Arellano, Bai and Kehoe (2012) (henceforth referred to as ABK), hiring inputs is risky because financial frictions limit firms' ability to insure against shocks that occur between the time of production and the receipt of revenues. Hence, an increase in idiosyncratic uncertainty induces firms to reduce their inputs to reduce exposure to such risk. But it is plausible that smaller firms would be more significantly affected than larger firms under the second and third channels. Large firms tend to be more immune to bank credit rationing, given that they are better able to draw on internal equity, due to stronger balance sheets and liquidity positions, and also able to access the public equity and bond markets[3], which are closed to small firms[4].

On both counts, our evidence favours the "Bloom channel". We find evidence of the characteristic "sharp drop and rapid rebound" profile of responses in activity following idiosyncratic uncertainty shocks, but interestingly, different to Bloom, not with aggregate uncertainty shocks. The effects of these shocks, in contrast, are shallower and more protracted in nature. Furthermore, our evidence contradicts the "CMR" and "ABK" channels, in that we find that idiosyncratic uncertainty shocks to large firms have a significantly more powerful effect on output fluctuations than small firms.

We build a proxy for idiosyncratic uncertainty using the methodology developed in CLMX (2001). Specifically, we measure idiosyncratic returns as the difference between firm-level and market returns, using daily data for over 20,000 publicly-listed US firms from 1962 to 2012. We measure within-month volatility of returns by computing the variance of daily returns, which yields a sample of over three million firm-month observations on volatility data. We use these to construct the idiosyncratic uncertainty index by aggregating across all firms, on a month-by-month basis, by taking the value-weighted average of individual firm volatilities. Similarly, we construct an index of aggregate uncertainty by estimating the

[3] Adrian, Colla and Shin (2012) find evidence of such a substitution effect, in that although bank lending to firms declined during the recent financial crisis, bond financing actually increased to make up much of the gap.

[4] It is also plausible that the "sharp drop and rapid rebound" profile of responses would not be observed under this channel. Tighter credit conditions tend to bite only with a lag, since firms generally maintain a "precautionary" buffer stock of unused credit, and hence can draw down existing, pre-arranged credit lines. A rapid recovery is also unlikely, as tighter credit rationing tends to persist until long after the episode of higher cashflow uncertainty resulting from the uncertainty shock has subsided.

within-month volatility of aggregate returns. We then append both indices to a parsimonious macroeconomic VAR, consisting of real industrial production, inflation, and the Federal Funds rate, which is estimated on US monthly data from 1962 to 2012, with identification imposed via the Cholesky decomposition.

Next, we decompose the idiosyncratic uncertainty index into two sub-indices, one for large and one for small firms. We rank firms on a year-by-year basis according to annual average market capitalization, and include only the largest 100 firms in the large firm index (which account for approximately 40-50% of the total market over our sample period), and put all remaining firms in the small firms index. We then replicate the original VAR model, but substitute the whole economy idiosyncratic uncertainty index with the disaggregated counterpart for large firms, and then repeat again using the small firms index[5].

Our main results are outlined as follows. We find that idiosyncratic uncertainty shocks account for approximately 30% of volatility in industrial production at business cycle frequencies, larger than the share attributed to monetary policy shocks. Moreover, breaking the sample at 1990 to exclude the possibility of structural breaks, this share increases to around 45%. They lead to a significant decline in growth in production, which falls rapidly over the course of a year, and then rebounds sharply, almost back to trend within the next 6 months. In contrast, the share attributed to aggregate uncertainty shocks is little more than 5%. Whereas idiosyncratic uncertainty is a primitive force driving macroeconomic fluctuations, fluctuations in aggregate uncertainty are mostly an endogenous response to other disturbances in the macroeconomy. Our results are robust to a variety of checks, including different orderings of variables in the VAR, existence of structural breaks, de-trending of volatility series, problems of over-parameterization, and exclusion of the Dotcom bubble and Lehman's collapse.

Regarding the sectoral decomposition of idiosyncratic uncertainty shocks, we find that for large firms, the proportion of fluctuations in growth accounted for is actually higher at

[5] Although it might seem more straightforward to include both the large and small firms uncertainty indices simultaneously in the VAR, rather than studying their effects sequentially as we do, this leads to a serious problem of over-parameterization, resulting in significant loss of estimation efficiency (with the IRFs having large confidence bands, and the FEVDS displaying implausible results. However, for robustness, we include both the large and small firms uncertainty series simultaneously in a more parsimonious version of the VAR, with some endogenous variables omitted (see section 6 below).

approximately 35% after 2 years, whereas for small firms it is lower at approximately 25% at a similar forecast horizon. A puzzling finding is that the sum of the contributions of large and small firms exceeds significantly that of idiosyncratic uncertainty in the whole economy - around 30% of output fluctuations is missing. Our resolution of this puzzle is down to the existence of an economy-wide "buffering effect" – volatility fluctuations in the large and small firms sector appear to co-move with output in opposite directions over certain stages of the business cycle, which dampens down the aggregate effects of idiosyncratic uncertainty shocks on growth.

Our results also relate to the well-known "Granular Hypothesis" of Gabaix (2011), which shows that if the size distribution of firms is sufficiently fat-tailed, as is documented empirically, then even if shocks are independently distributed, the conditions under which the Central Limit Theorem hold are no longer satisfied, so that the well-known diversification argument no longer holds, and hence idiosyncratic shocks, in particular shocks to large firms, can survive in the aggregate. Hence, our results that idiosyncratic uncertainty shocks to large firms have more significant aggregate effects than small firms suggest there may be granular origins to the role of uncertainty in business cycle fluctuations (although the key difference is that whereas Gabaix focusses on time variation in the first moment, we focus on time variation in the second moment).

The remainder of the paper is structured as follows: section 2 details the measurement of the aggregate and idiosyncratic uncertainty indices, and describes the volatility data. Section 3 outlines the model specification, and discusses data and estimation issues. Section 4 presents the results, whilst section 5 goes into further detail on the disaggregated effects of large and small firms' uncertainty shocks. Section 6 checks the robustness of our results and section 7 concludes.

2. Measurement of aggregate and idiosyncratic uncertainty

We follow the methodology of CLMX (2001), who decompose the total return volatility of an individual firm into its constituent components, which has the appealing property that it allows us to extract the idiosyncratic component without the need to estimate covariances or betas for individual firms. But whereas they disaggregate into three components:

aggregate volatility, industry-specific volatility and firm-specific (idiosyncratic) volatility, we adhere to the traditional CAPM paradigm and decompose return volatility into only two components – aggregate and idiosyncratic volatility[6]. We start off by presenting the traditional CAPM decomposition (which requires estimation of betas), then apply the more tractable decomposition of CLMX (2001) (which does not).

Let R_{it} denote the excess return of firm *i* in period *t*. Let w_{it} denote the weight of firm *i* in the total market, where weights are defined according to market capitalization. Then the excess market return is simply the value-weighted average of the individual firm excess returns:

$$R_{mt} = \sum_{i} w_{it} R_{it} \tag{1}$$

Next, from CAPM we can write:

$$R_{it} = \beta_i R_{mt} + \tilde{\epsilon}_{it} \tag{2}$$

where β_i denotes the beta for firm *i* with respect to the market return and $\tilde{\epsilon}_{it}$ denotes the firm-specific residual. As is standard in CAPM, $\tilde{\epsilon}_{it}$ is orthogonal to the market return, which enables the following variance decomposition:

$$Var(R_{it}) = \beta_i^2 Var(R_{mt}) + Var(\tilde{\epsilon}_{it}) \tag{3}$$

Hence, it is clear to see that the CAPM decomposition requires estimation of betas. Next, we turn to the CLMX (2001) decomposition. Let ϵ_{it} be defined as the difference between the firm-level return R_{it} and the market return R_{mt}, i.e.

$$R_{it} = R_{mt} + \epsilon_{it} \tag{4}$$

Substituting in from (2) and re-arranging yields:

$$\epsilon_{it} = \tilde{\epsilon}_{it} + (\beta_i - 1) R_{mt} \tag{5}$$

Hence, computing the variance decomposition of (4) and taking into account (5), yields:

---

[6] Relating our approach to theirs, it is easily shown that our measure of idiosyncratic volatility simply bundles together their measures of firm-specific and industry-specific volatility.

$$Var(R_{it}) = Var(R_{mt}) + Var(\epsilon_{it}) + 2Cov(R_{mt}, \epsilon_{it}) \tag{6}$$

$$= Var(R_{mt}) + Var(\epsilon_{it}) + 2(\beta_i - 1)\, Var(R_{mt})$$

Note that although the variance of individual firm returns contains covariance terms, if we aggregate across all firms by taking the weighted average of individual variances:

$$\sum_i w_{it}\ Var(R_{it}) = \ Var(R_{mt}) + \sum_i w_{it}\ Var(\epsilon_{it}) = \sigma^2_{mt} + \sigma^2_{\epsilon t} \tag{7}$$

we find that the individual betas drop out given the standard property of CAPM that $\sum_i w_{it}\beta_{im} = 1$. Thus, we can use the firm-specific residuals ϵ_{it} to construct an index of idiosyncratic volatility, without the need to estimate betas. Finally, we show how our measure of idiosyncratic volatility relates to the "true" measure based on CAPM.

Using (5) and computing the variance decomposition of ϵ_{it}, we get:

$$Var(\epsilon_{it}) = (\beta_i - 1)^2\ Var(R_{mt}) + Var(\tilde{\epsilon}_{it}) \tag{8}$$

Aggregating (8) across all firms yields:

$$\sigma^2_{\epsilon t} = \tilde{\sigma}^2_{\epsilon t} + CSV_t(\beta_{im})\sigma^2_{mt} \tag{9}$$

where $\tilde{\sigma}^2_{\epsilon t} = \sum_i w_{it} Var(\tilde{\epsilon}_{it})$ is the average variance of the CAPM firm-specific residual $\tilde{\epsilon}_{it}$, and $CSV_t(\beta_{im}) = \sum_i w_{it}(\beta_i - 1)^2$ is the cross-sectional variance of betas across all firms in the market. Equation (9) shows that our measure of idiosyncratic volatility contains two components, the first term represents the intrinsic measure of idiosyncratic volatility a la CAPM, and the second term depends on the cross-sectional variance of individual firm betas and market volatility. However, according to CLMX (2001), plausible estimates of cross-sectional variance in betas are sufficiently small, such that the second component accounts for only a small fraction of time variation in the volatility series. Thus, our measure of idiosyncratic volatility can be viewed as a decent proxy for the intrinsic measure derived from CAPM.

Estimation

We use firm-level return data from the Center for Research in Security Prices (CRSP) US Stock Database for all firms trading on the NYSE, Amex and Nasdaq exchanges. Our sample period runs from July 1962 (daily data from the Amex exchange only goes back this far) to

December 2012 (the latest available data point). The total number of firms increases from 3256 to 5567 over this period, with the coverage peaking at 9837 in 1999.[7]

We replicate the methodology (with some modifications) of CLMX (2001) to estimate the volatility series $\sigma_{\epsilon t}^2$ and σ_{mt}^2 (see their paper for a more detailed description). Excess returns are defined as the spread over the 30 day Treasury bill rate (computed at the equivalent daily rate). We estimate the within-month volatility of excess market returns by computing:

$$MKT_t = \sum_{s \in t} (R_{ms} - \mu_m)^2 \tag{10}$$

where μ_m is defined as the within-month sample mean of the market return R_{ms}, and estimate the within-month volatility of idiosyncratic residuals by computing:

$$\hat{\sigma}_{\epsilon it}^2 = \sum_{s \in t} \epsilon_{is}^2 \tag{11}$$

As explained above, we aggregate across all firms in the sample by computing the value-weighted average of individual firm volatilities:

$$FIRM_t = \sum_i w_{it} \hat{\sigma}_{\epsilon it}^2 \tag{12}$$

to give us our idiosyncratic volatility index.

Descriptive analysis

It is useful to begin with a descriptive analysis of our volatility time series. Figures 1 to 4 plot the four monthly volatility series, which have been annualized (by multiplying by 12). The top panels plot the raw data, whereas the bottom panels plot a smoothed series (estimated as a 12 month backward moving average). To begin with, we update the results on trends in volatility in CLMX (2001)[8]. Consistent with their paper, we find evidence of a clear, upward trend in idiosyncratic volatility during the 1960-2000 period (this is more evident in the smoothed graph, which strips out the high frequency noise and focusses on the slow-

[7] A large number of firms de-list following the shakeouts of the 2000 Dotcom bust and 2008 financial crisis.

[8] Their time series end in 1997 (the latest available data point at their time of writing), where ours continue to 2012.

moving component). However, from around 2000 onwards, this gives way to a downward trend[9]. Furthermore, looking at the disaggregated idiosyncratic volatility series (Figures 3

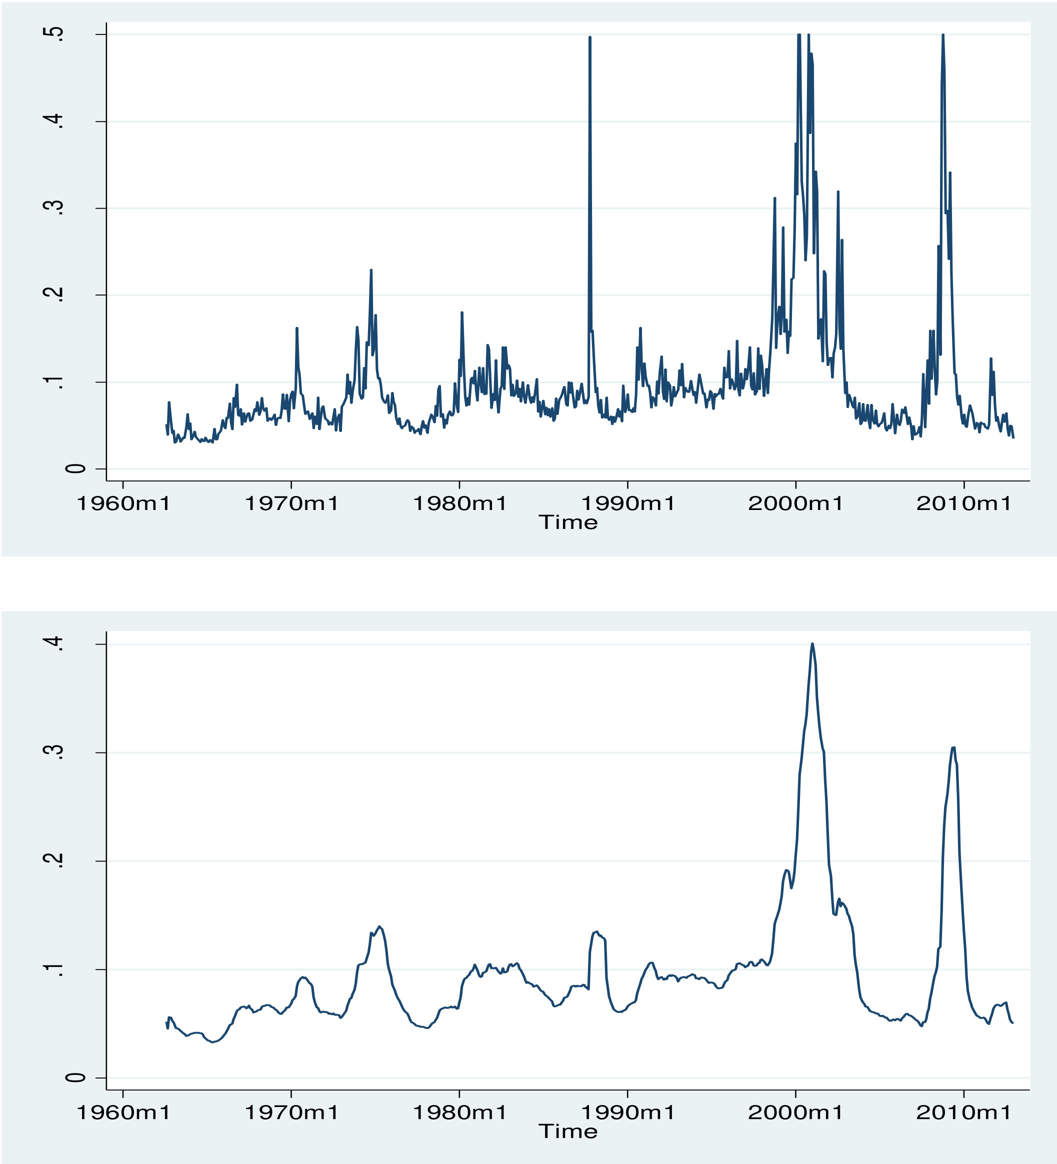

FIGURE 1: ANNUALIZED IDIOSYNCRATIC VOLATILITY

The top panel shows the annualized variance within each month of daily firm-level returns relative to the market, calculated using equations (11) and (12), for the period July 1962 to December 2012. The bottom panel shows a backwards 12-month moving average of FIRM. The series have been truncated where volatility exceeds 0.5.

[9] Zhang (2010) also finds evidence of upward trend in idiosyncratic uncertainty from 1980 until 2000, and a downward trend after 2000.

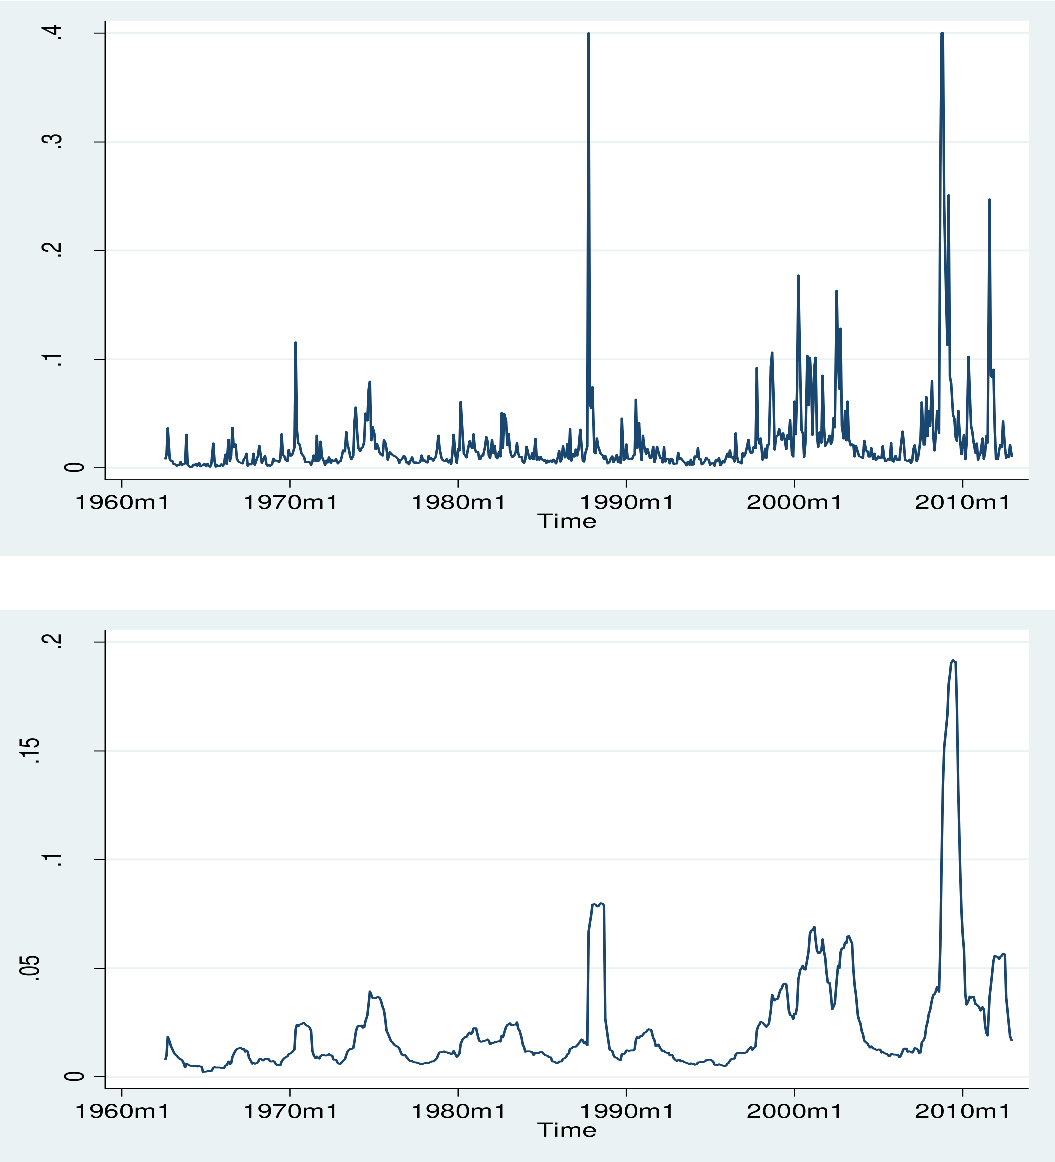

FIGURE 2: ANNUALIZED AGGREGATE VOLATILITY

The top panel shows the annualized variance within each month of daily market returns, calculated using equation (10), for the period July 1962 to December 2012. The bottom panel shows a backwards 12-month moving average of MKT. The series have been truncated where volatility exceeds 0.4.

and 4), it is clear that the trend component we observe is almost entirely driven by the small firms sector, with large firms exhibiting no obvious trend in volatility. In contrast, aggregate volatility exhibits no discernible trend over the entire time series.

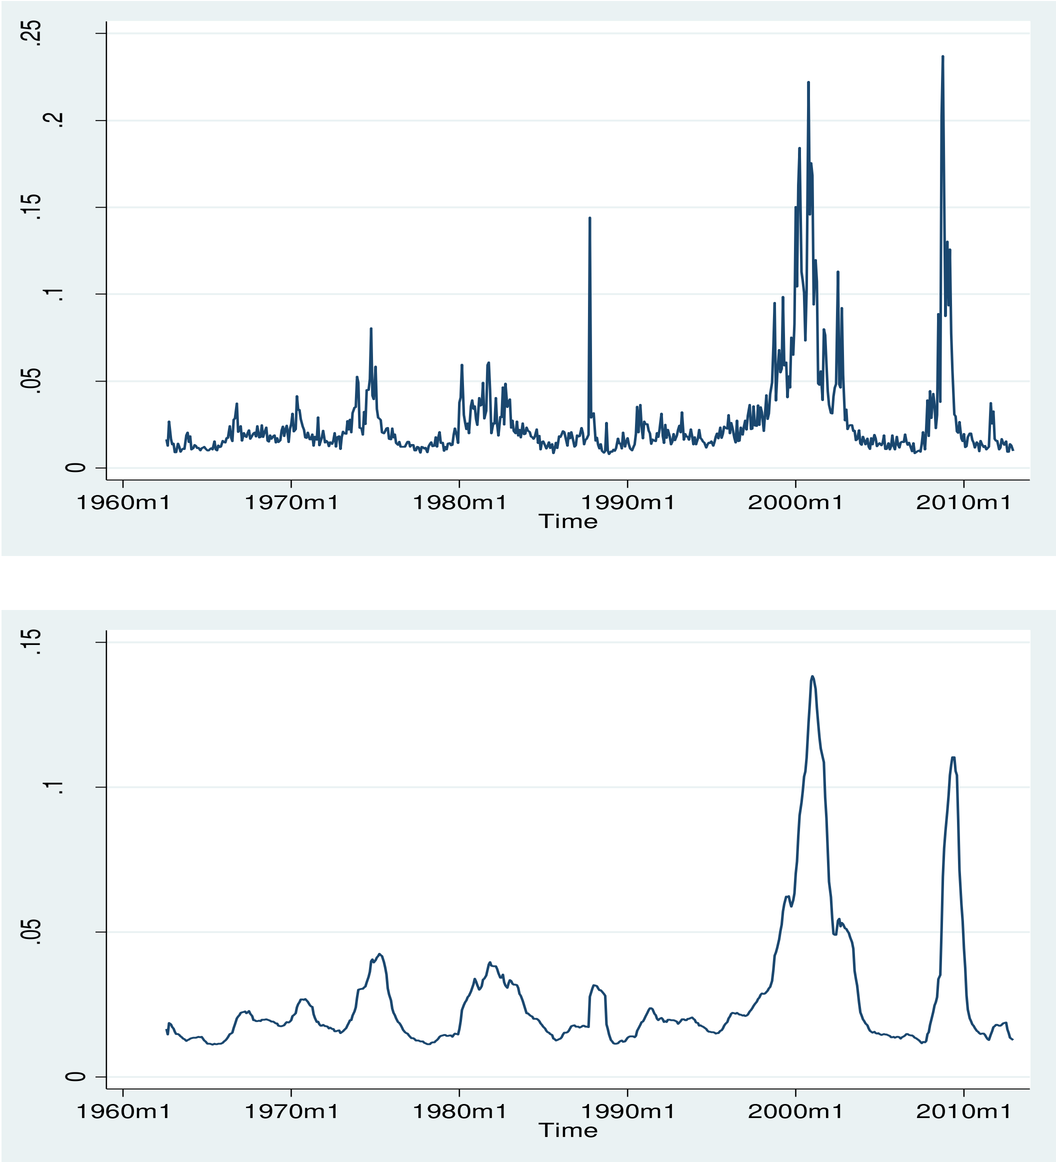

FIGURE 3: ANNUALIZED IDIOSYNCRATIC VOLATILITY (LARGE FIRMS)

The top panel shows the annualized variance within each month of daily firm-level returns relative to the market, calculated using equations (11) and (12), for the period July 1962 to December 2012, for the largest 100 firms only (according to market capitalization). The bottom panel shows a backwards 12-month moving average of the same series.

A casual inspection (in particular of the smoothed graphs) reveals a high degree of comovement between aggregate volatility and idiosyncratic volatility. Both series show spikes in the mid-1970s (coinciding with the first oil shock), and also the financial crisis of

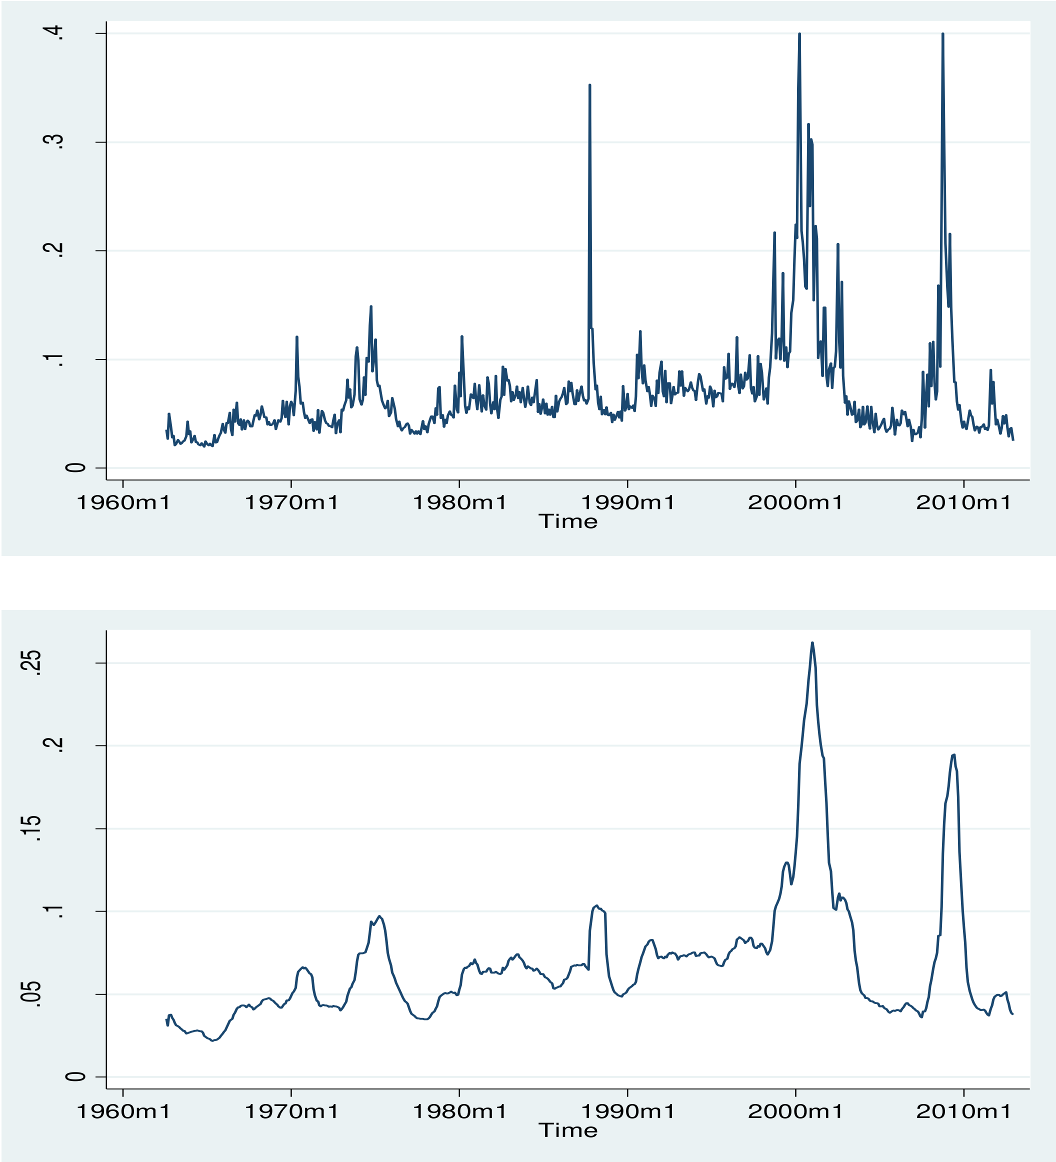

FIGURE 4: ANNUALIZED IDIOSYNCRATIC VOLATILITY (SMALL FIRMS)

The top panel shows the annualized variance within each month of daily firm-level returns relative to the market, calculated using equations (11) and (12), for the period July 1962 to December 2012, excluding the largest 100 firms (according to market capitalization). The bottom panel shows a backwards 12-month moving average of the same series. The series have been truncated where volatility exceeds 0.4.

2008. However, the comovement is less pronounced for some episodes, for example, there is a large spike in 1987 for aggregate volatility (due to the stock market crash), but a less significant increase for idiosyncratic volatility. Interestingly, the reverse situation occurs during the Dotcom Bust of 2000.

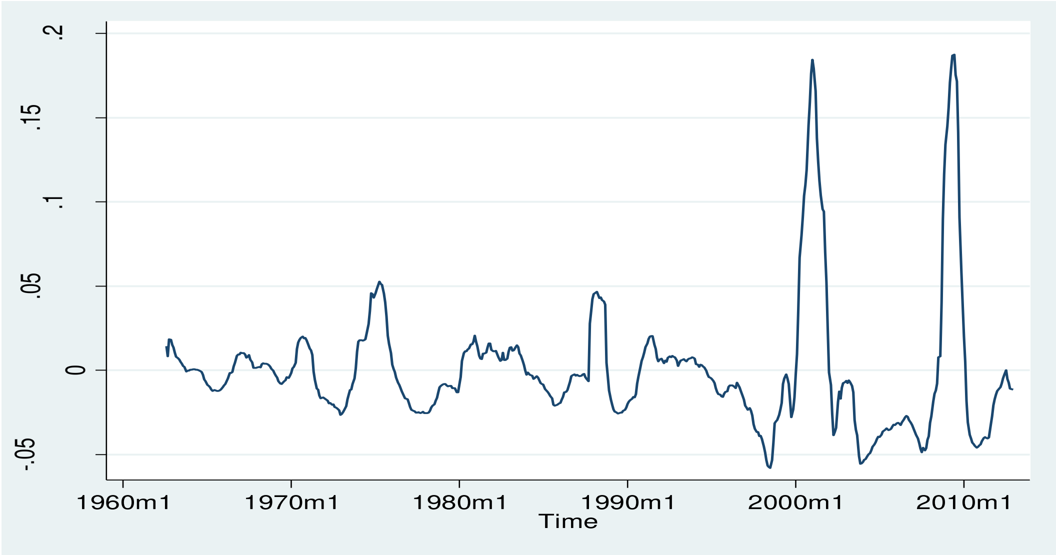

FIGURE 5: IDIOSYNCRATIC VOLATILITY (DE-TRENDED)

This shows the annualized variance within each month of daily firm-level returns relative to the market, calculated using equations (11) and (12) for the period July 1962 to December 2012, which is then de-trended using the Hodrick-Prescott filter (setting the appropriate Ravn-Uhlig value for monthly data of λ = 129,600), and smoothed by taking a backwards 12-month moving average.

Moreover, there are some episodes where the series seem to move quite differently. This becomes more distinct when we de-trend and smooth the volatility series[10]. For instance, in the early 1980s, idiosyncratic volatility increases above trend, but aggregate volatility stays muted[11] (see Figures 5 and 6). We also see the opposite effects occurring around 2011, with aggregate volatility ticking up, but idiosyncratic volatility remaining relatively subdued.

We also observe a very high degree of comovement between the disaggregated volatility series (again this is more visible in the de-trended and smoothed graphs, Figures 7 and 8). The key difference is that idiosyncratic volatility of small firms is significantly higher than that of large firms.

[10] We detrend using the Hodrick-Prescott filter (setting the appropriate Ravn-Uhlig value for monthly data of λ = 129,600), and smooth as before by taking a backwards 12 month moving average.

[11] This episode is particularly unusual as it is one of the few recessionary episodes in which aggregate uncertainty does not show a sharp increase. But this might be due to the unusual nature of this recession, in that unlike other severe recessions, credit conditions did not tighten significantly (see Bijapur (2010)). Thus, small firms, which tend to be more affected by external financing constraints, may have been spared the worst effects of the recession, and thus uncertainty was less of a market-wide phenomenon. Indeed, looking at the disaggregated series, we find that uncertainty in the large firms sector increases more significantly than the small firms sector.

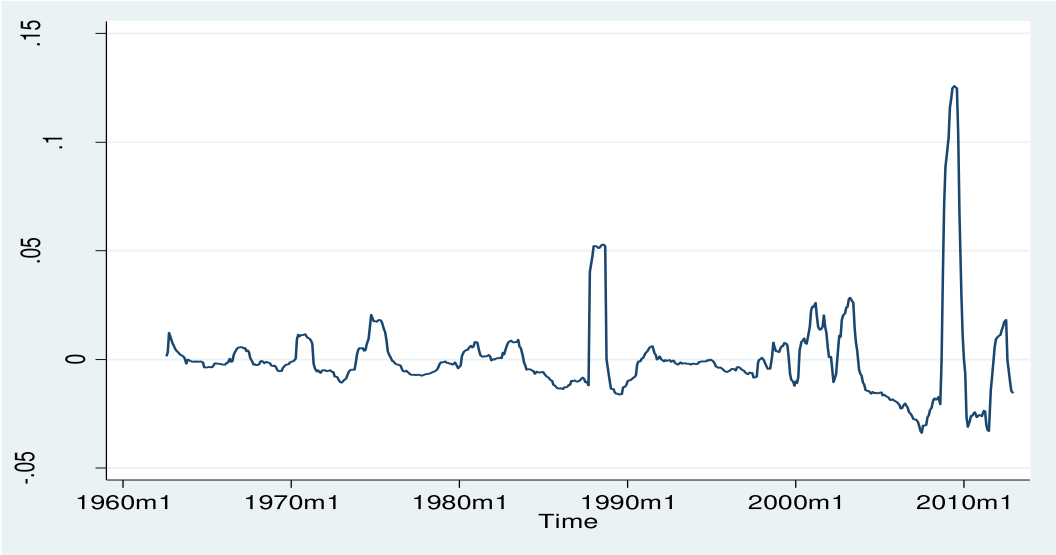

FIGURE 6: AGGREGATE VOLATILITY (DE-TRENDED)

This shows the annualized variance within each month of daily market returns, calculated using equation (10), for the period July 1962 to December 2012, which is then de-trended using the Hodrick-Prescott filter (setting the appropriate Ravn-Uhlig value for monthly data of λ = 129,600), and smoothed by taking a backwards 12-month moving average.

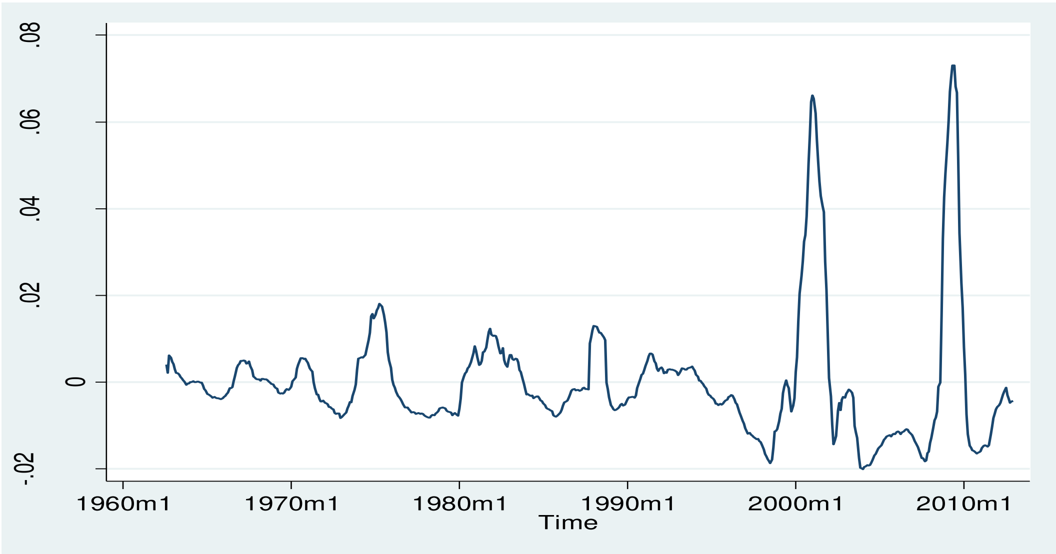

FIGURE 7: IDIOSYNCRATIC VOLATILITY – LARGE FIRMS (DE-TRENDED)

This shows the annualized variance within each month of daily firm-level returns relative to the market, calculated using equations (11) and (12), for the period July 1962 to December 2012, for the largest 100 firms only (according to market capitalization), which is then de-trended using the Hodrick-Prescott filter (setting the appropriate Ravn-Uhlig value for monthly data of λ = 129,600), and smoothed by taking a backwards 12-month moving average.

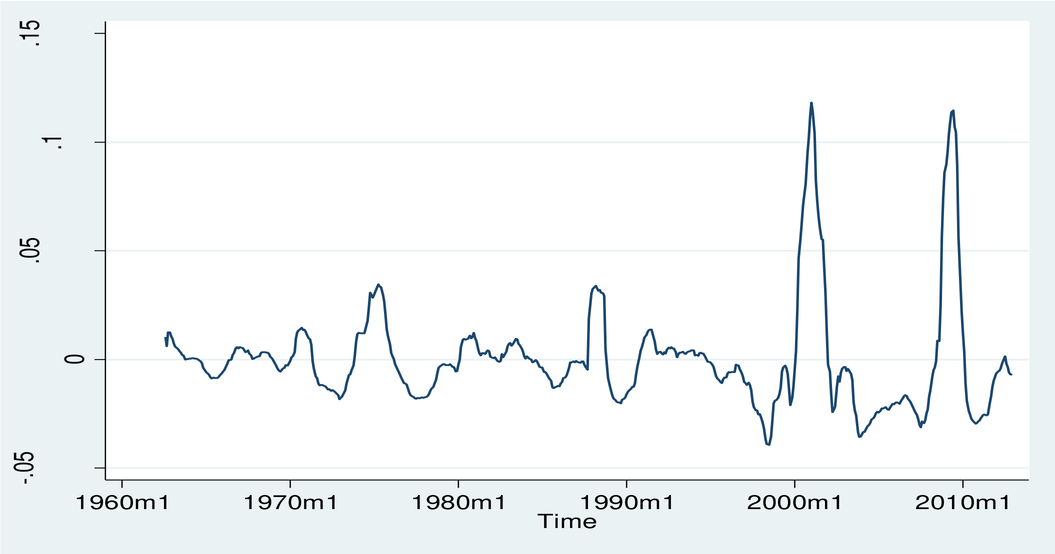

FIGURE 8: IDIOSYNCRATIC VOLATILITY – SMALL FIRMS (DE-TRENDED)

This shows the annualized variance within each month of daily firm-level returns relative to the market, calculated using equations (11) and (12), for the period July 1962 to December 2012, excluding the largest 100 firms (according to market capitalization), which is then de-trended using the Hodrick-Prescott filter (setting the appropriate Ravn-Uhlig value for monthly data of λ = 129,600), and smoothed by taking a backwards 12-month moving average.

3. Model specification, data and estimation

We model the macroeconomy using a stylized vector autoregression (VAR) consisting of real industrial production, inflation and the policy rate, which we expand to include our idiosyncratic and aggregate uncertainty indices. Although admittedly simple, it avoids the problem of over-parameterization, and also overcomes the need to make ad hoc identifying restrictions which would arise in a more extensive model. Furthermore, given that our main objective here is not to undertake a detailed analysis of the transmission mechanisms to an exhaustive set of macroeconomic aggregates, but instead to quantify the relative contributions of two different sources of uncertainty shocks, we believe our simpler approach suffices.

We impose identification through the Cholesky decomposition, which boils down to specifying the order in which shocks enter the system of equations in the VAR[12]. We specify the ordering of endogenous variables in the VAR as:

$$X_t = [IU_t, AU_t, MP_t, \pi_t, y_t]' \tag{13}$$

where *X* denotes our vector of endogenous variables, *IU* denotes idiosyncratic uncertainty, *AU* denotes aggregate uncertainty, *MP* denotes the policy rate, π denotes inflation and *y* denotes industrial production. The assumptions underpinning this particular choice of ordering are based on Bloom (2009): (i) financial variables should be affected before real economy variables due to the efficient nature of financial markets, and (ii) prices tend to adjust to the impact of shocks before quantities. To these we add a third assumption: (iii) shocks affect idiosyncratic uncertainty first, and then propagate to aggregate uncertainty through the general equilibrium effects. We envisage this mechanism arising through the Granular Hypothesis of Gabaix (2011), namely shocks to individual, large firms could permeate throughout the economy as a whole due to the interconnectedness of firms[13]. We check the robustness of our results to a range of different orderings in section 6 below.

Data and estimation

Our estimations are based on US monthly data. Variables are specified as follows: real industrial production[14] (seasonally adjusted, annual growth rates), inflation (Consumer Price

---

[12] The complete structural VAR model can be characterized by three components: (1) the matrix of structural parameters (denoted A), (2) the matrices of lag coefficients (denoted Φ), and (3) the variance-covariance matrix of the structural innovations (denoted Ω), giving:

$$AX_t = \Phi(L)X_t + \varepsilon_t, \qquad E(\varepsilon\varepsilon') = \Omega$$

where X_t is the vector of endogenous variables and ε_t is the vector of structural innovations. The exclusion restrictions are applied to matrix A, which through the Cholesky decomposition has lower triangular form.

[13] This choice of ordering is further justified by the results of our Granger causality tests, which indicate strongly that idiosyncratic uncertainty leads production, where there is no such unambiguous lead-lag relationship between aggregate uncertainty and production. Thus, ordering idiosyncratic uncertainty first avoids the possibility that the contribution of these shocks to fluctuations in production are incorrectly attributed to aggregate uncertainty shocks.

[14] given that GDP data is unavailable at a monthly frequency.

Index for all urban consumers: all items (seasonally adjusted, annual growth rates), policy rate (Federal Funds target rate). All macroeconomic data are taken from the Global Financial Data database, and all stock market data from the CRSP database. We estimate the VAR over sample period July 1962 – Dec 2012.

Preliminary inspection of the data reveals spikes in both uncertainty series due to the stock market crash of October 1987 and the Dotcom Bust of April 2000. The consensus view in the academic and policy literature seems to be that these episodes of turbulence were confined to financial markets and did not propagate to the real economy in any significant way. Hence, in order to prevent these one-off events from overpowering and distorting the estimation results, we exclude these outliers with time dummies.

We undertake Augmented Dickey Fuller (ADF) tests for unit roots on all variables. Given the trade-off between too few lags (giving rise to potential serial correlation and bias of the test statistic), and too many lags (decreasing the power of the test), we test using both 1 and 9 lags on the differences. In both cases, we reject the null of non-stationarity at the 1 per cent level for all variables, with the exception of the inflation variable and Federal Funds Rate, which are therefore de-trended using the Hodrick-Prescott filter (setting the appropriate Ravn-Uhlig value for monthly data of $\lambda = 129,600$). ADF tests on the de-trended series indicate the previously detected non-stationarity is eliminated.

A potential complication in the analysis is that our results might reflect longer term comovements in the data, as oppose to the cyclical fluctuations which are the focus of our interest. Indeed, CLMX (2001) document strong evidence of an upward trend in idiosyncratic uncertainty within their sample, hence in order to isolate short run movements around these trends, we de-trend both our volatility time series before embarking on the VAR analysis.

We include 25 lags of each variable in the VAR - the Lagrange Multiplier (LM) test for optimal lag structure indicates the optimal number of lags is 25, whereas the Akaike information criterion (AIC) puts it at 27, but we opt for 25 in order to gain from parsimony.

4. Results

Correlations and lead-lag relationships

In order to gain some preliminary insights into the sources of comovements in the data, it is useful to examine the bivariate correlations (at various leads and lags of up to one year) between our set of endogenous variables, before characterizing the relationships more formally with the VAR analysis. Both uncertainty time series[15] show strong evidence of serial correlation, which is more persistent for idiosyncratic uncertainty. There is also strong evidence of comovement between idiosyncratic and aggregate uncertainty, with the contemporaneous correlation being the strongest, and each series exhibiting positive and statistically significant correlations with the lagged value of the other up to 6 months. Although the correlations between lagged values of idiosyncratic uncertainty with contemporaneous aggregate uncertainty are slightly higher than the converse, the lead-lag relationship between the two uncertainty indices remains ambiguous.

Table 1 below reports the pairwise correlations of our two uncertainty indices with industrial production growth. There is strong evidence that both uncertainty measures lead production, with all lagged (and contemporaneous) values exhibiting negative[16] and highly statistically significant correlations with production up to one year. But notably, the correlations are stronger for idiosyncratic uncertainty, peaking at -0.39 after 6 months, and for aggregate uncertainty peaking at -0.26. For lead values, the evidence is weak - although both uncertainty series exhibit negative correlations with production, they are statistically significant only up to 3 months, and the magnitudes are considerably smaller. Thus the preliminary evidence on lead-lag relationships points to both uncertainty series leading production, rather than the converse.

[15] We report all correlations using the de-trended uncertainty measures only.

[16] i.e. uncertainty is higher during recessions and lower during booms.

Lead/lag (months)	Idiosyncratic uncertainty	Aggregate uncertainty
+12	-0.28**	-0.13**
+6	-0.40**	-0.26**
+3	-0.35**	-0.24**
0	-0.22**	-0.15**
-3	-0.12**	-0.08*
-6	-0.04	-0.03
-12	0.07	0.04

TABLE 1: CORRELATIONS WITH REAL INDUSTRIAL PRODUCTION GROWTH

NOTE: This table reports the lead, lagged and contemporaneous correlation structure of the aggregate and idiosyncratic uncertainty series (both de-trended) with real industrial production growth. * and ** indicate significance at the 5% and 1% levels respectively.

Granger causality tests

Next, we gain further insights by investigating to what extent the different uncertainty measures can forecast each other, and also forecast production growth and other endogenous variables. We do this by conducting Granger causality tests on each of the VAR equations.

Starting with the full VAR system which includes all five endogenous variables, unsurprisingly, we find strong evidence that each uncertainty measure Granger causes the other. However, in the production equation, although tests for both inflation and the policy rate are statistically significant, neither of the uncertainty measures is. However, although these variables are not independently significant in terms of their forecasting power, this does not preclude their joint significance. Indeed, the lack of individual significance can be explained given that the two uncertainty measures are highly positively correlated. This, combined with the problem of our inclusion of numerous lags, means that the Granger causality test, which tests whether each of the uncertainty series is independently significant in forecasting production, is less likely to yield a positive result. Hence, to get round this, we repeat the exercise but include each individual uncertainty series sequentially in the VAR system. In both cases, the results are quite different. First, our results shed light on the lead-lag relationship between idiosyncratic uncertainty and production. We find that idiosyncratic uncertainty Granger causes production with a very high significance level (whereas the tests for inflation and policy rate are no longer significant at the 5% level). In contrast, we find that production has very little predictive power for idiosyncratic uncertainty.

Second, we also find aggregate uncertainty has forecasting power for production, but this is significantly weaker than with the idiosyncratic uncertainty measure. Furthermore, production has strong predictive power for aggregate uncertainty, hence the lead-lag relationship is more ambiguous for this uncertainty measure.

Third, the results also lend support to the view that movements in aggregate uncertainty are primarily an endogenous response to macroeconomic fluctuations, with all three macro variables having significant forecasting power for our aggregate uncertainty measure. However, one cannot say the same for idiosyncratic uncertainty - only inflation has significant predictive ability for this variable. But analysis of the reduced form system is incapable of disentangling the underlying causal relationships, which requires that we move to a more structural model.

Idiosyncratic uncertainty

	y_t	π_t	MP_t	IU_t
y_{t-i}	-	0.014	0.000	0.234
π_{t-i}	0.069	-	0.746	0.000
MP_{t-i}	0.060	0.000	-	0.203
IU_{t-i}	0.002	0.000	0.063	-

Aggregate uncertainty

	y_t	π_t	MP_t	AU_t
y_{t-i}	-	0.017	0.000	0.010
π_{t-i}	0.024	-	0.756	0.005
MP_{t-i}	0.010	0.000	-	0.030
AU_{t-i}	0.014	0.000	1.000	-

TABLE 2: GRANGER CAUSALITY

NOTE: This table reports the *p*-values of Granger-causality VAR tests. The null hypothesis is that lags 1 through i of the series indicated in the row do not help to forecast the series indicated in the column, conditional on the other variables in the VAR.

VAR analysis

In this section, we present further insights into the effects of shocks to our two different uncertainty measures on the macroeconomy, by undertaking Cholesky-orthogonalized impulse response functions (IRFs) and forecast error variance decompositions (FEVDs) on our set of endogenous variables. We present results for the whole economy idiosyncratic uncertainty index, and also compare these with the results for the disaggregated measures for large and small firms.

IRFs

Starting with the IRFs, an unanticipated increase in idiosyncratic uncertainty of one standard deviation causes a significant decline in real industrial production (in level terms), which falls by approximately ¾ of a percentage point below trend 18 months after impact. Growth falls rapidly over the course of a year, but then subsequently rebounds sharply, almost 2/3 of the way back to trend within the next 5 months. The sharp drop and rapid rebound in growth is even more noticeable for idiosyncratic uncertainty shocks originating in the large firms sector, with the trough in growth rates slightly deeper, and the pace of recovery faster with growth almost back to trend within the following year. Their effects are of a similar order of magnitude to monetary policy shocks.

The effects of an aggregate uncertainty shock are, however, more subdued, and also more protracted. Following impact, growth dips gently, bottoming out after a year, and also displays a much more gradual pace of recovery, with growth still below trend after another year. In level terms, production is only around ¼ of a percentage point below trend after 18 months.

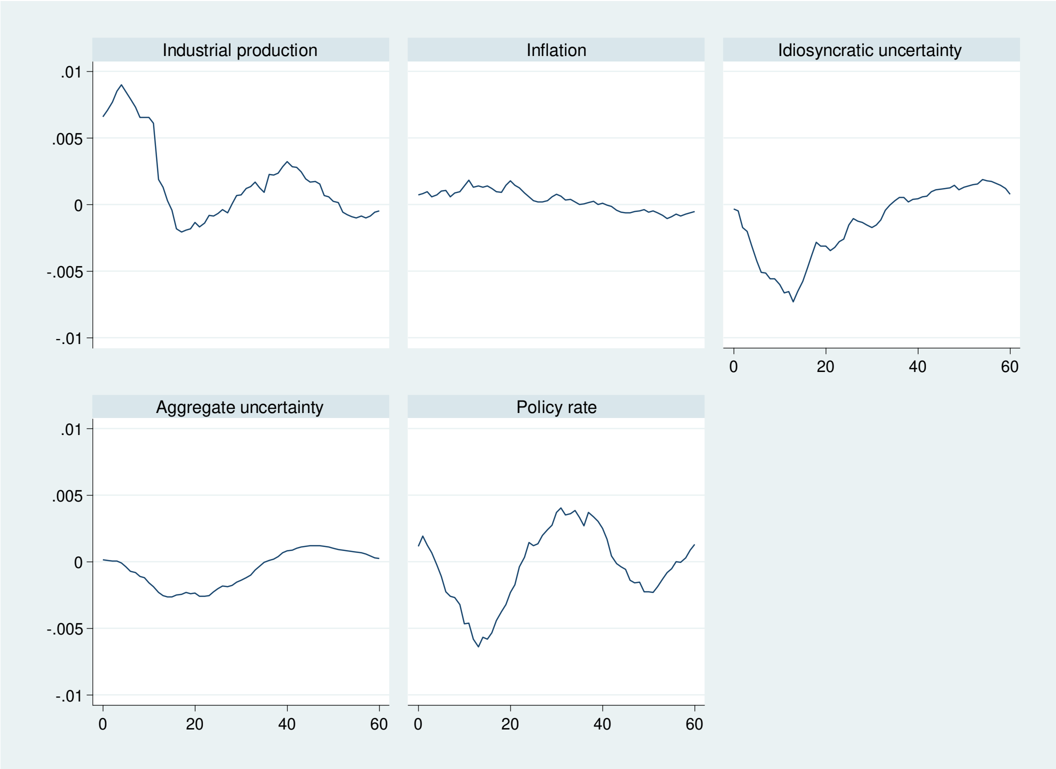

FIGURE 9: IMPULSE RESPONSE FUNCTIONS (EFFECTS OF SHOCKS ON INDUSTRIAL PRODUCTION)

The figure depicts the impulse responses of real industrial production growth to a one standard deviation orthogonalized shock to all variables.

It is interesting to note that the profile of responses to an idiosyncratic uncertainty shock displays the characteristic rapid drop and subsequent rebound in activity characterized in Bloom (2009)[17]. His intuition for these effects is that firms respond to a surge in uncertainty by delaying major investment decisions and hiring decisions, due to the "real option" value of waiting. This happens rapidly because the uncertainty shock affects firms' expectations about future business conditions, hence investment and hiring freezes are implemented with immediate effect. Once the uncertainty has dissipated, activity rebounds rapidly as firms address their pent-up demand for investment and hiring. However, the key difference is that Bloom associates this profile of responses with aggregate uncertainty shocks, whereas we find this for idiosyncratic, but not aggregate, uncertainty shocks.

We also investigate long-run effects, and find that a typical idiosyncratic uncertainty shock results in a permanent loss in production of around ¾ of a percentage point after 5 years

[17] Although Bloom uses levels of production (relative to trend), whereas we use growth rates.

(regarding the sectoral decomposition, shocks to large firms result in a permanent loss of around the same magnitude, but for small firms it is smaller at around ½ a percentage point). We interpret this as due to the scaling down of investment plans and postponement of projects by firms during the period of heightened uncertainty. This results in an erosion of future productive potential. Although Bloom notes that firms attempt to "catch up" on foregone investment opportunities during the recovery phase, our findings suggest that they never fully do so[18]. In contrast, although monetary policy shocks have an equally severe effect on production in the short run, in the long run, the economy almost fully recoups these losses such that there is no permanent effect on the level of production, consistent with the behaviour of a traditional demand shock. But, as expected, supply shocks (i.e. shocks to production) have a significant, persistent effect on output, resulting in a loss of around ¾ of a percentage point after 5 years. Hence, the long-run effects of an idiosyncratic uncertainty shock are of the same order of magnitude as a supply shock.

It is also of interest to examine the profile of responses of the uncertainty series themselves. We find that both uncertainty shocks decay quite rapidly, and are almost completely dissipated after a year[19], suggesting that periods of heightened uncertainty tend to be quite short-lived. It is also noteworthy that idiosyncratic uncertainty shocks lead to a sizeable increase in aggregate uncertainty, but again this decays rapidly and then overshoots for quite a protracted period (in contrast, aggregate uncertainty shocks have almost no effect on idiosyncratic uncertainty). These results suggests that in addition to a direct transmission mechanism from idiosyncratic uncertainty to real activity, there may also be an indirect channel through which an idiosyncratic uncertainty shock leads to propagation of aggregate uncertainty, which in turn drives business cycle fluctuations.

Finally, we examine the monetary policy response to uncertainty shocks. The IRFs indicate there is very little response to either type of uncertainty shock, suggesting that central banks tend not to ease policy significantly in the wake of spikes in uncertainty. However, the

[18] Alternatively, it might be that output takes longer to recover than our 5 year horizon (due to "time to build" investment lags etc.). However, running the IRFs up to 10 years does not affect the results, there is still a permanent loss in production of around ¾ of a percentage point.

[19] although aggregate uncertainty actually overshoots slightly for the following year before returning to steady state.

standard error bands are quite wide, hence it is difficult to pin down with confidence the true policy response.

FEVDs

To gain insights into how significant are uncertainty shocks in driving business cycle fluctuations, relative to the standard macroeconomic shocks modelled here, we conduct FEVDs, running out to a five year horizon (in order to match average business cycle frequencies). We find that idiosyncratic uncertainty shocks account for a large fraction of output fluctuations over most of the forecast range, dominating all other shocks with the exception of supply shocks. In contrast, the fraction attributed to aggregate uncertainty shocks is negligible throughout.

In the short-term (less than 6 months after impact), the contribution of idiosyncratic uncertainty is relatively small, with supply shocks dominating and accounting for over 80% of fluctuations in production. However, the effects of idiosyncratic uncertainty shocks propagate rapidly, peaking at more than 30% after 2 years, and stay around this level for the duration. The proportion attributed to aggregate uncertainty shocks starts off small, and stays low at around 5% throughout. Looking at the sectoral decomposition of idiosyncratic uncertainty shocks, we find that for large firms, the proportion of fluctuations in production accounted for is actually higher at around 35% after 2 years, whereas for small firms is lower at around 25% at a similar forecast horizon.

Interestingly, the contribution of monetary policy shocks is significantly lower than idiosyncratic uncertainty shocks, at less than 20% after 2 years, albeit increasing slightly to 25%[20] by the end of the 5 year forecast horizon. Inflation shocks, however, have a negligible impact throughout.

[20] consistent with several estimates reported in the existing literature.

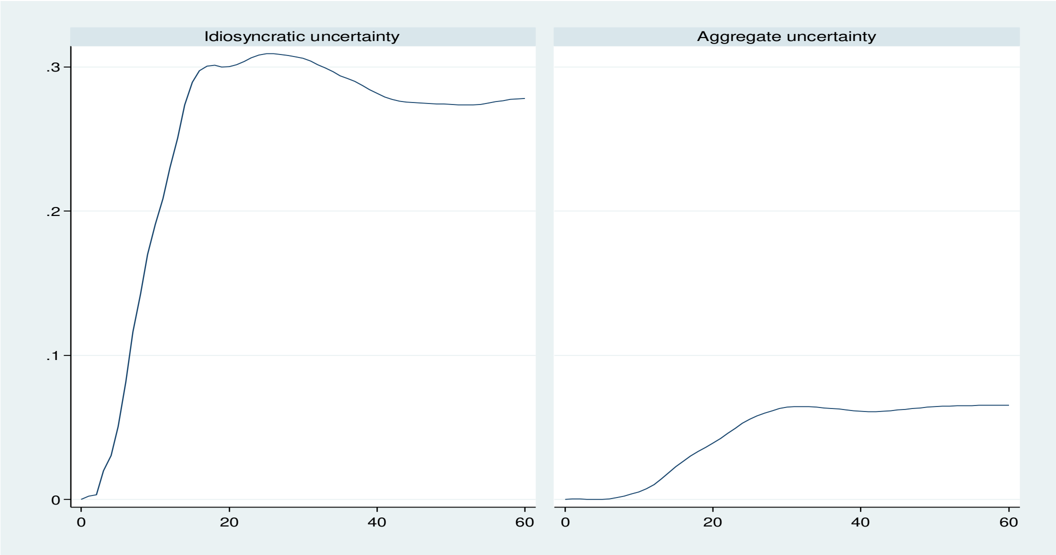

FIGURE 10: FORECAST ERROR VARIANCE DECOMPOSITION

The figure depicts the forecast error variance decomposition of real industrial production growth from one standard deviation orthogonalized shocks to idiosyncratic and aggregate uncertainty.

It is also useful to analyze the factors driving movements in the uncertainty measures themselves. Regarding idiosyncratic uncertainty, this is mainly self-driven by idiosyncratic uncertainty shocks, although monetary policy also plays a significant role. This suggests that idiosyncratic uncertainty is a primitive force driving fluctuations, as opposed to an endogenous response reflecting other macroeconomic factors. In contrast, aggregate uncertainty is driven significantly by all our candidate shocks (with the exception of inflation), with idiosyncratic uncertainty shocks having the largest effect. Thus, it appears that fluctuations in aggregate uncertainty are mostly an endogenous response to other disturbances in the macroeconomy.

Shocks	Industrial production	Inflation	Policy rate
Idiosyncratic	31	19	1
Aggregate	5	2	2

TABLE 3: FORECAST ERROR VARIANCE DECOMPOSITIONS

NOTE: This table reports the percentage contributions of one standard deviation orthogonalized shocks to idiosyncratic and aggregate uncertainty to other endogenous variables, at a 24 month forecast horizon.

5. Comparing the effects of large and small firms' uncertainty shocks

A somewhat puzzling finding is that when we aggregate up the individual shares of fluctuations in production attributed to large and small firms, the numbers do not seem to add up to what we get for the whole economy idiosyncratic uncertainty shocks. Specifically, at a 2 year forecast horizon, the share of output fluctuations attributed to uncertainty shocks for large firms is around 35%, whereas for small firms it is around 25%. But for the aggregated equivalent of the idiosyncratic uncertainty series, the figure is around 30%. Hence, around 30% of output fluctuations seems to be missing.

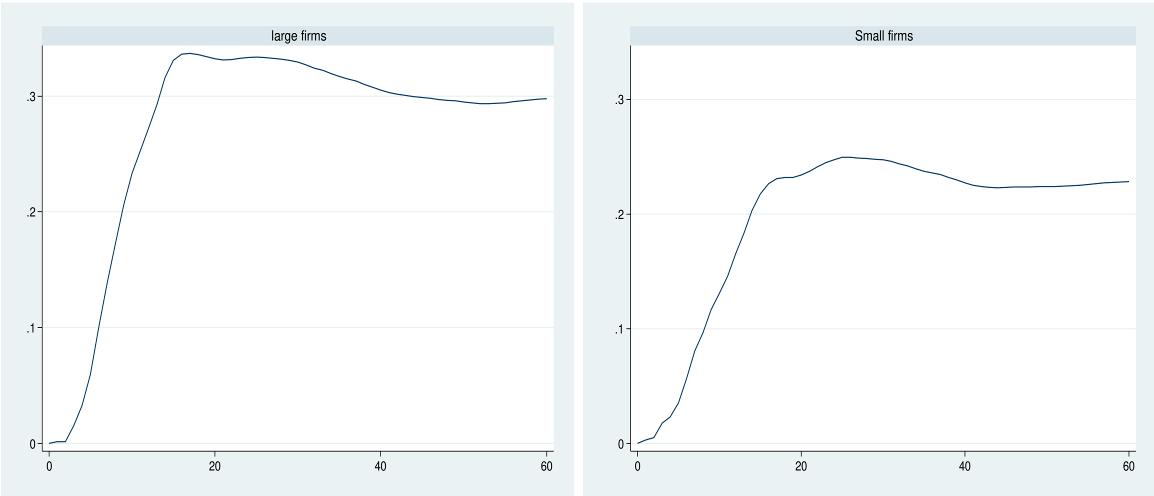

FIGURE 11: COMPARING LARGE AND SMALL FIRMS – FEVDS

The figure depicts the forecast error variance decomposition of real industrial production growth from one standard deviation orthogonalized shocks to idiosyncratic uncertainty of large firms and small firms.

This puzzle might be explained by omitted variable bias – the idiosyncratic uncertainty series for large and small firms are highly correlated, hence given that we include the large (small) firms index in isolation in the VAR, it is plausible that its contribution to economic fluctuations is biased upwards, as it captures the effects of small (large) firms also. This would effectively lead to a problem of "double-counting" when summing together the individual contributions of large and small firms' uncertainty shocks. However, this explanation is unsatisfactory in that it does not explain why the contribution of large firms' shocks (35%) actually exceeds that of the aggregated equivalent (30%).

To investigate further, we reformulate the VAR to include both large and small firm idiosyncratic uncertainty indices simultaneously, but omit inflation and interest rates in order to avoid the problem of over-parameterization from including too many endogenous variables. We then replicate this VAR model, but substitute the two disaggregated uncertainty series with the single aggregated equivalent, and compare the results. Our results are consistent with the previous formulation. From a two year forecast horizon onwards, the contribution of large and small firms shocks sums to around 25%, whereas the aggregated equivalent only accounts for around 15%[21].

Our proposed resolution of this puzzle is based on the pattern of comovements of large and small firms' shocks with production. Although there is a strong positive contemporaneous correlation between the innovations of large and small firm uncertainty measures, they appear to co-move with production in opposite directions over certain stages of the business cycle, with a large (small) firms' shock having a positive (negative) impact on production. Hence, there appears to be an economy-wide "buffering effect" in operation, which dampens down the aggregate effects of idiosyncratic uncertainty shocks on economic fluctuations. Inspecting the IRFs for the effects of shocks to large and small firms on activity, we find supportive evidence of this. Although both sectors experience steep initial drops in production, large firms recover more rapidly and more vigorously than small firms, with growth turning sharply positive for large firms around 18 months after impact, whereas as for small firms it remains negative.

[21] We also check whether this puzzling phenomenon is an artefact of the unique episode in the late 1990s - early 2000s which saw a significant increase in idiosyncratic uncertainty due to the incidence of the Dotcom bubble and bust. However, re-estimating the VARs but excluding the time period 1998Q4-2001Q4 from the sample, we find that our results are robust.

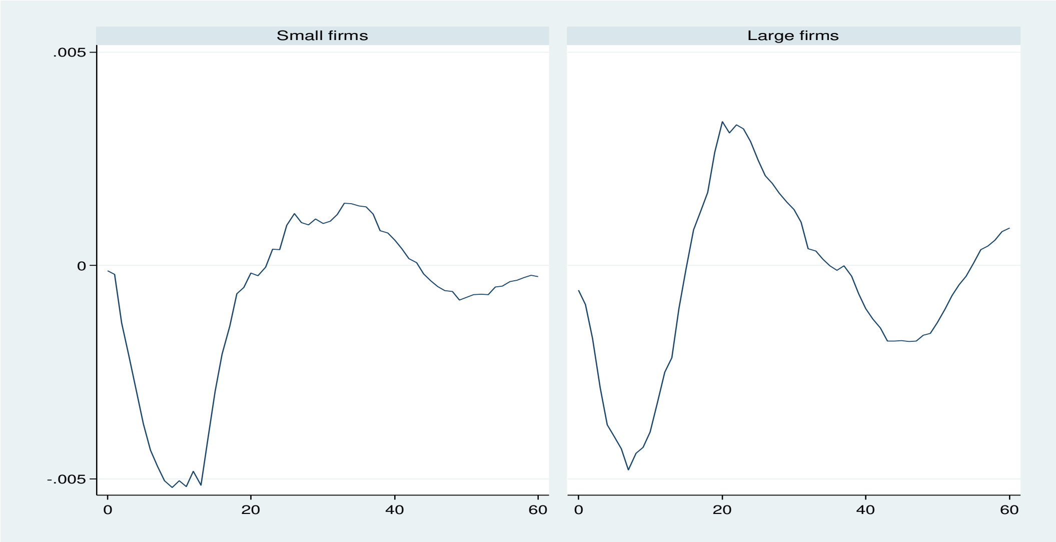

FIGURE 12: COMPARING LARGE AND SMALL FIRMS – IRFS

The figure depicts the impulse responses of real industrial production growth to one standard deviation orthogonalized shocks to idiosyncratic uncertainty of large firms and small firms.

Our intuition for these observed effects is as follows. An idiosyncratic uncertainty shock would likely result in increased cash-flow volatility, thus exacerbating financial frictions and making it harder for firms to fund investment through bank lending. However, large firms would be less constrained by tighter bank credit rationing than small firms, given that they are more able to draw on internal equity due to stronger balance sheets and liquidity positions, and also able to access external equity and bond markets, which are generally closed to small firms. Thus, large firms would have a better opportunity to expand and grab market share from their (shrinking) smaller rivals. We formalize this intuition as follows.

We start with the benchmark case in which we use the aggregated idiosyncratic uncertainty series in the VAR. Assuming the VAR is stable, we can re-write it in vector moving-average form:

$$X_t = \mu + \sum_{s=0}^{\infty} \varphi_s u_{t-s} \tag{14}$$

where X is an M-variate stochastic process, μ is a vector of time-invariant means, and u is a vector of innovations with $u_t \sim N(0, \Sigma)$. It is more useful to look at the MA representation with orthogonal innovations. Factorization of Σ yields:

$$FF' = \Sigma \quad \text{and} \quad G\Sigma G' = \mathrm{I} \tag{15}$$

where $G = F^{-1}$, which allows us to re-write (14) as:

$$X_t = \mu + \sum_{s=0}^{\infty} \varphi_s F v_{t-s} \tag{16}$$

where $u_t = Fv_t$ and $Ev_t v_t' = I$. Denoting $\varphi_s^* = \varphi_s F$, the K-step ahead forecast error is:

$$X_t - E_{t-k}[X_t] = \sum_{s=0}^{K-1} \varphi_s^* v_{t-s} \tag{17}$$

Given the standard SVAR assumption of mutual orthogonality of v, the variance-covariance matrix of the K-step ahead forecast errors is given by $\sum_{s=0}^{K-1} \varphi_s^* \, \varphi_s^{*\prime}$. Isolating our specific variable of interest, production (denoted *y*), the k-step ahead forecast error is given by:

$$y_t - E_{t-k}[y_t] = \sum_{i=0}^{k-1} a_i^{11}\, \varepsilon_{1t-i} + \sum_{i=0}^{k-1} a_i^{12}\, \varepsilon_{2t-i} + \cdots + \sum_{i=0}^{k-1} a_i^{1m}\, \varepsilon_{mt-i} \tag{18}$$

where a_i^{pq} denotes the pth row and qth column of matrix φ_i^*. Hence, the k-step ahead forecast error variance for variable *y* is given by:

$$E(y_t - E_{t-k}[y_t])^2 = \sigma_1^2 \sum_{i=0}^{k-1} (a_i^{11})^2 + \sigma_2^2 \sum_{i=0}^{k-1} (a_i^{12})^2 + \cdots + \sigma_m^2 \sum_{i=0}^{k-1} (a_i^{1m})^2 \tag{19}$$

where the mth variable denotes the IV series.

Next, we compare with the case in which we replace the aggregated idiosyncratic uncertainty series with the disaggregated counterparts for large and small firms, denoted by variables m and m+1 respectively. As previously, we assume:

$$E[v_t v_s{}'] = 0_K \quad for\ all\ \ s \neq t \tag{20}$$

However, unlike previously, we have:

$$E[v_t v_t{}'] \neq I_K \tag{21}$$

given there exists a non-zero contemporaneous correlation between the two sectoral idiosyncratic uncertainty series. Thus, the k-step ahead forecast error variance for variable *y* becomes:

$$
\begin{aligned}
&E(y_t - E_{t-k}[y_t])^2 \\
&\qquad = \sigma_1^2 \sum_{i=0}^{k-1} (a_i^{11})^2 + \cdots + \sigma_{m+1}^2 \sum_{i=0}^{k-1} (a_i^{1m+1})^2 \\
&\qquad + 2 \sum_{i=0}^{k-1} a_i^{1m} a_i^{1m+1} E(\varepsilon_{mt-i}\, \varepsilon_{m+1t-i})
\end{aligned}
\tag{22}
$$

We assume the sectoral idiosyncratic uncertainty series are covariance stationary, i.e.

$$E(\varepsilon_{mt-i}\varepsilon_{m+1t-i}) = \sigma_{m,m+1} \quad for\ all\ i \tag{23}$$

Substituting (23) into (22) yields:

$$
\begin{aligned}
&E(y_t - E_{t-k}[y_t])^2 \\
&\qquad = \sigma_1^2 \sum_{i=0}^{k-1} (a_i^{11})^2 + \cdots + \sigma_{m+1}^2 \sum_{i=0}^{k-1} (a_i^{1m+1})^2 \\
&\qquad + 2\sigma_{m,m+1} \sum_{i=0}^{k-1} a_i^{1m} a_i^{1m+1}
\end{aligned}
\tag{24}
$$

Next, let $\tilde{E}(y_t - E_{t-k}[y_t])^2$ denote the mis-specified forecast error variance in which mutual orthogonality is assumed between *all* innovations, hence:

$$\tilde{E}(y_t - E_{t-k}[y_t])^2 = \tilde{\sigma}_1^2 \sum_{i=0}^{k-1} (\tilde{a}_i^{11})^2 + \cdots + \tilde{\sigma}_{m+1}^2 \sum_{i=0}^{k-1} (\tilde{a}_i^{1m+1})^2 \tag{25}$$

where $\tilde{a}_i^{pq}$ denotes the pth row and qth column of matrix $\tilde{\varphi}_i^*$, and $\tilde{\varphi}_i^*$ is the equivalent of φ_i^* in the previous case above.

Thus, given $\sigma_{m,m+1} > 0$,

$$\tilde{E}(y_t - E_{t-k}[y_t])^2 > E(y_t - E_{t-k}[y_t])^2 \Leftrightarrow \sum_{i=0}^{k-1} a_i^{1m} a_i^{1m+1} < 0 \tag{26}$$

which requires a_i^{1m}, a_i^{1m+1} to have opposite signs over some portion of the forecast horizon.

6. Robustness checks

We begin by investigating the robustness of our results to different orderings of variables in the VAR. We find that the key flavour of our results continues to hold for a variety of different orderings. In particular, we check whether our results are sensitive to the decision to order idiosyncratic before aggregate uncertainty in the VAR. Ordering the aggregate uncertainty variable first, and again repeating the exercise with this variable taking all other possible positions in the estimation order, we find that the share of fluctuations in activity attributed to aggregate uncertainty shocks remains negligible, and never exceeds the share reached in our original formulation.[22]

Next, we completely reverse the order of variables compared to the original formulation, with production entering first and idiosyncratic uncertainty last. Here, we find that both types of uncertainty shocks become relatively insignificant in explaining fluctuations in activity, however repeating this exercise for large firms only reveals a different result. The contribution of idiosyncratic uncertainty shocks remains significant (despite being last in the ordering), peaking at around 15%, although only half its magnitude in the original formulation. However, aggregate uncertainty remains insignificant, despite being positioned before idiosyncratic uncertainty in the estimation order.

We also investigate whether our results are affected by a problem of over-parameterization, due to inclusion of five endogenous variables and a large number of lags. We drop the inflation variable from the VAR, as this consistently has the smallest forecasting power for fluctuations in activity in our previous estimations, but otherwise we retain the original estimation order. We find that the fraction explained by idiosyncratic uncertainty shocks peaks at over 20%, whereas aggregate uncertainty never exceeds 2%. Reversing the estimation order, the results remain robust with idiosyncratic uncertainty shocks remaining significant, and double that of aggregate uncertainty shocks which remains at approximately only 5%, despite being ordered before idiosyncratic uncertainty. Again, the effects are stronger for large firms, in the reverse ordering the impact of idiosyncratic uncertainty shocks peaks at over 20%, approximately 4 times that for aggregate uncertainty.

[22] Indeed, we try 20 different orderings and the contribution of aggregate uncertainty remains negligible for all of these, even when it is placed first in the ordering.

Robustness to structural breaks

Another concern is that the possibility of structural breaks within our estimation period might affect the results. Our first concern is with the well-documented "Great Moderation", i.e. the period of significantly diminished output volatility, occurring from the mid-1980s (McConnell and Perez-Quiros (2000)) until the outbreak of the financial crisis of 2008, and also the incidence of more efficient monetary policy-making and reduced occurrence of monetary policy shocks, occurring from around 1990 (Cecchetti et al. (2006)). Our interest here is whether the effects of idiosyncratic uncertainty might also suffer from a "moderated" impact during this time period. We investigate whether our results are robust to these effects by re-estimating the VAR, FEVDs and IRFs but with the sample period truncated in 1990.

Our key results appear to be unaffected by the possibility of structural breaks. On the contrary, breaking the sample at 1990[23], we find the share of fluctuations in activity attributed to idiosyncratic uncertainty shocks increases significantly to approximately 45% at a two year forecast horizon (compared with the whole sample estimate of around 30%), whilst aggregate uncertainty shocks remain insignificant at around 5%. The time profile of responses to uncertainty shocks also remains similar – an idiosyncratic uncertainty shock leads to a sharp fall and rapid rebound in output, whereas an aggregate uncertainty shock induces a more gradual and persistent decline. Also consistent with previous results, idiosyncratic uncertainty shocks to large firms account for a larger share of fluctuations in activity than to small firms (45% as opposed to 35% for the latter).

Our second concern is that our results might be distorted by the effects of a unique episode in the late 1990s - early 2000s which saw a significant increase in idiosyncratic uncertainty due to the incidence of the Dotcom bubble and bust. However, re-estimating the VAR but excluding the relevant time period (1998Q4-2001Q4)[24] from the sample, we find that our results continue to hold, and if anything are actually reinforced. Specifically, the share of fluctuations in activity attributed to idiosyncratic uncertainty shocks increases significantly

[23] Breaking the sample at different points during the mid-1980s to 1990 yields similar results.

[24] Bekaert et al. (2012) test for structural breaks for these effects and find the ends of 1997/1998 and 2001/2002 are consistently identified as break points.

to approximately 40% after two years, whereas aggregate uncertainty shocks become even more insignificant, with their share falling to 0% (to a first order approximation).

7. Concluding Remarks

In this paper, we study jointly the roles of aggregate and idiosyncratic uncertainty shocks in driving business cycle fluctuations. By decomposing total stock return uncertainty of over 20,000 publicly-listed US firms, we construct separate indices for aggregate and idiosyncratic uncertainty, and run a horse race between them in an otherwise standard macroeconomic VAR. We find that idiosyncratic uncertainty shocks account for approximately 30% of volatility in industrial production at business cycle frequencies, larger than the share attributed to monetary policy shocks. They generate a significant decline in real economic activity, and exhibit the characteristic “sharp drop and rapid rebound” response identified in Bloom (2009). In contrast, aggregate uncertainty shocks have a much more muted impact, with their contribution to volatility in industrial production little more than 5%. Idiosyncratic uncertainty shocks to large firms also have more powerful macroeconomic impacts than small firms, suggesting “Granular” origins to the role of uncertainty in the macroeconomy.

We also shed light on the transmission mechanisms of uncertainty shocks to the real economy. Our findings are supportive of the channel identified in Bloom (2009), in which greater uncertainty increases the real option value to waiting, hence firms scale back their investment and hiring plans. In contrast, our evidence casts doubt on the mechanism formulated in CMR (2012), which predicts that smaller firms would be more significantly affected by uncertainty shocks than larger firms.

One possible direction for future research would be to investigate the propagation mechanisms through which such Granular effects might arise. In particular, it would be interesting to explore whether there are general equilibrium effects through which an idiosyncratic uncertainty shock to a large firm might permeate throughout the economy as a whole, and how this depends on the interconnectedness of firms.

Another promising extension would be to investigate more deeply the causal relationships between idiosyncratic uncertainty and key macroeconomic variables. Although our research shows that these shocks generate significant impacts on the real economy, a fully structural model is required to tackle the identification problem.

Finally, it is also of interest to investigate the optimal policy response to idiosyncratic uncertainty shocks, and the differences between policy responses to aggregate and idiosyncratic uncertainty shocks. But these, and other interesting questions, are left for future research.

References

[1] Adrian, T., Colla, P., Shin, H. S., 2013. Which Financial Frictions? Parsing the Evidence from the Financial Crisis of 2007 to 2009. NBER Macroeconomics Annual, University of Chicago Press, vol. 27(1), 159-214.

[2] Alexopoulos, M., Cohen, J., 2009. Uncertain Times, Uncertain Measures. University of Toronto, Department of Economics. Working Paper No. 325.

[3] Angeletos, G. M., Collard, F., Dellas, H., 2014. Quantifying Confidence. MIT and University of Bern. Mimeo.

[4] Arellano, C., Bai, Y., Kehoe, P. J., 2012. Financial frictions and fluctuations in volatility. Staff Report 466, Federal Reserve Bank of Minneapolis.

[5] Bachmann, R., Bayer, C. ,2013. "Wait-and-see" business cycles. J. Monet. Econ. 60 (6) , 704-719.

[6] Bachmann, R.,Elstner, S., Sims, E., 2013. Uncertainty and economic activity: evidence from business survey data. Am. Econ. J.: Macroecon. 5 (2), 217-249.

[7] Baker, S., Bloom, N., 2012. Does Uncertainty Reduce Growth? Using Disasters as Natural Experiments. Stanford University. Mimeo.

[8] Baker, S., Bloom, N. ,Davis, S. J., 2013. Measuring Economic Policy Uncertainty. Stanford University and the University of Chicago Booth School of Business. Mimeo.

[9] Bekaert, G., Hodrick, R. J., Zhang, X., 2012. Aggregate Idiosyncratic Volatility. Journal of Financial and Quantitative Analysis 47, 1155-1185.

[10] Bekaert, G., Hoerova, M., Duca, M. L., 2013. Risk, uncertainty and monetary policy. J. Monet. Econ. 60, 771–788.

[11] Benigno, G., Benigno, P., Nisticò, S., 2012. Risk, monetary policy and the exchange rate. NBER Macroecon. Annu. 26, 247-309.

[12] Bloom, N., 2009. The impact of uncertainty shocks. Econometrica 77, 623-685.

[13] Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., Terry, S.J., 2012. Really Uncertain Business Cycles. NBER Working Paper series, Working Paper No. 18245.

[14] Bijapur, M., 2010. Does monetary policy lose effectiveness during a credit crunch?. Economics Letters 106, 42-44.

[15] Born, B., Pfeifer J., 2014. Policy risk and the business cycle. Journal of Monetary Economics 68, 68-85.

[16] Campbell, J., M. Lettau, B. G. Malkiel, and Y. Xu. 2001. Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk. Journal of Finance 56: 1–43.

[17] Caggiano, G., Castelnuovo, E., Groshenny, N., 2014. Uncertainty Shocks and Unemployment Dynamics. Journal of Monetary Economics 67, 78-92.

[18] Cecchetti, S. G., Flores-Lagunes, A., Krause, S., 2006. Has monetary policy become more efficient? A cross-country analysis. Economic Journal 116, 408-433.

[19] Christiano, L., Motto, R., Rostagno, M., 2014. Risk shocks. Am. Econ. Rev. 104 (1), 27-65.

[20] Gabaix, X., 2011. The Granular Origins of Aggregate Fluctuations. Econometrica 79, 733-772.

[21] Gilchrist, S., Sim, J.W., Zakrajsek, E., 2013. Uncertainty, Financial Frictions, and Irreversible Investment, Boston University and Federal Reserve Board. Mimeo.

[22] McConnell, M., Perez-Quiros, G., 2000. Output Fluctuations in the United States: What Has Changed Since the Early 1980's?. American Economic Review 90, 1464-1476.

[23] Mumtaz, H., Surico, P., 2013. Policy uncertainty and aggregate fluctuations. Queen Mary University of London and London Business School. Mimeo.

[24] Mumtaz, H., Zanetti, F., 2013. The Impact of the Volatility of Monetary Policy Shocks. Journal of Money, Credit, and Banking 45, 535-558.

[25] Zhang, C., 2010. A re-examination of the causes of time-varying stock return volatilities. Journal of Financial and Quantitative Analysis 45, 663-684.