Really Uncertain Business Cycles

Nicholas Bloom, Max Floetotto and Nir Jaimovich*
Stanford University

January 2009

Abstract

We propose uncertainty shocks as an additional impulse driving business cycles. First, we demonstrate that uncertainty, measured by a number of proxies, appears to be strongly countercyclical. When uncertainty is included in a standard vector-auto-regression, increases in uncertainty lead to a large drop and rebound in economic activity. Second, we build a dynamic stochastic general equilibrium model that extends the benchmark neoclassical growth model along two dimensions. It allows for the existence of heterogeneous firms with non-convex adjustment costs in both capital and labor and time variation in uncertainty, modelled as a changing variance of innovations to productivity. We find that increases in uncertainty lead to large drops in employment and investment. This occurs because uncertainty makes firms cautious, leading them to pause hiring and investment. This freezing in activity also reduces the reallocation of capital and labor across firms, leading to a large fall in productivity growth. Taken together, the freeze in hiring, investment and productivity growth lead to a business cycle sized drop and rebound in output following a rise in uncertainty.

Keywords: uncertainty, adjustment costs and business cycles..
JEL Classification: D92, E22, D8, C23.

*We would like to thank our formal discussant Eduardo Engel and participants at the AEA for comments and suggestions. Correspondence: Nick Bloom, Department of Economics, Stanford University, Stanford, CA 94305, nbloom@stanford.edu.

1 Introduction

What are the primary drivers of the business cycle and, in particular, what drives recessions? In this paper, we tackle this question by proposing a new mechanism that generates cyclical fluctuations: variations in the level of uncertainty, or second-moment shocks. The idea to link uncertainty to the business cycle is not new. John Maynard Keynes (1936) argued that changes in investor sentiments, the so-called *animal spirits*, could lead to economic downturns. While this can be interpreted as an argument for the role of uncertainty, it has not traditionally played a large role in modern studies of business cycle fluctuations.

This paper attempts to fill this gap by taking three distinct steps. First, we address the empirical behavior of uncertainty over the business cycle. Evidence on the time series variation in uncertainty is scarce as no measure of uncertainty is directly observable. To circumvent this difficulty, we use various proxies for uncertainty. These include measures of cross firm and cross industry dispersion, time series variation of aggregate data as well as measures of forecaster disagreement. We find that all of our uncertainty proxies are strongly countercyclical — uncertainty appears to rise sharply during recessions, increasing by between 20% to 70%. We use a variety of methods to combine all these proxies in a single aggregate uncertainty index. One possibility is to calculating the index as a simple average of our seven proxies after re-scaling each of them to an average of one in non-recession periods. Measured by this index, uncertainty increases by 42.5% during recessions. Taking this one step further, we investigate the conditional association of uncertainty and the business cycle by including the uncertainty index in a standard vector-auto-regression (VAR). Here, we try to control for the first moment by including the standard Basu, Fernald and Kimball (2006) series for total factor productivity (TFP) and find that an increase in uncertainty is associated with a large drop and rebound in economic activity.

Second, in order to theoretically investigate the role of fluctuations in uncertainty we study a dynamic stochastic general equilibrium model that allows for shocks to both the level of technology (the first moment) as well as uncertainty (the second moment). More specifically, we model shocks to uncertainty as variation in the standard deviation of the innovations to productivity. Various features of the model are specified to conform as closely as possible to the standard frictionless neoclassical growth model. The main deviation from this benchmark model lies in assuming that firms incur convex and non-convex adjustment costs in both capital and labor. The non-convexities imply that firms become more cautious in investing and hiring when uncertainty increases. This is because the option value of waiting increases when uncertainty is high – it is expensive for firms to invest and then disinvest or to hire and then fire. The time varying option value implies that in the presence of higher aggregate macro uncertainty, aggregate investment and employment levels fall as it becomes optimal for each individual firm to wait. In addition, time varying uncertainty reduces productivity growth as it lowers the extent of reallocation in the economy. When uncertainty rises productive firms expand less and unproductive firms contract less. In the model, as in actual US economy, reallocation is a key factor driving aggregate productivity[1]. When aggregate uncertainty falls back down firms start to invest and hire heavily to address pent-up demand from their previous period of caution. Reallocation also picks up and there is a rapid rebound of productivity.[2]

Third, we will evaluate the empirical performance of the model using detailed US Census micro data. In particular, we will look at response of capital and labor to productivity shocks at the establishment level. This will allow us to distinguish between first and second moment shocks in the data. First moment shocks shift the Ss bands to the left and thus imply more disinvestment and less investment. Second moment shocks on the other hand lead to an outward shift of the Ss bands and to a freeze in both investment and disinvestment.

Our work is related to several strands in the literature. First, we add to the extensive literature building on the Real-Business Cycle (RBC) framework by studying the role of productivity shocks in causing business cycles. In the standard RBC literature recessions are caused by large negative technology shocks.[3] The reliance on negative technology shocks has proven to be controversial as it suggests that recessions are times of technological regress. As King and Rebelo (1999) ask "if these shocks are large and important, why cannot we read about them in the Wall Street Journal?"[4] As discussed above, our work provides a rationale for falling measured productivity. Counter-cyclical increases in uncertainty lead to a freeze in economic activity, substantially lowering productivity growth during recessions. In our model, however, the drop in TFP is not causing the recession, but rather an artifact of a recession which is in turn caused by an increase in uncertainty.

The paper also relates to the literature on investment under uncertainty. A growing body of work has shown that uncertainty can directly influence firm-level investment and

[1] See, for example, Foster, Haltiwanger and Krizan (2000, 2006), who report that reallocation, broadly defined to include entry and exit, accounts for around 50% of manufacturing and 80% of retail productivity growth in the US. These figures will in fact understate the full contribution of reallocation as they miss the within establishment reallocation, which Bernard, Redding and Schott's (2006) results on product switching suggest could be important. Hence, this suggests that reallocation accounts for the large majority of US productivity growth.

[2] This is not to claim that other mechanisms, in which fluctuations in uncertainty can affect economic activity are irrelevant. Instead, this specification allows us to isolate the role of cyclical movements in uncertainty levels due to the specific mechanism we analyze within a framework that has the benchmark frictionless neoclassical growth model as a special limiting case.

[3] See for example the discussion in Rebelo (2005): "Most RBC models require declines in TFP in order to replicate the declines in output observed in the data."

[4] This reasoning has lead many researchers to study models with other disturbances, which still mostly focus on first moment (level) shocks. A partial list of these alternative shocks includes oil shocks, investment specific shocks, monetary shocks, government expenditure shocks, news shocks, and terms-of-trade shocks. Yet, in most models, negative technology shocks continue to be an important driver of economic downturns.

employment in the presence of adjustment costs.[5] Probably, of most relevance is Bloom (2009) who solves a partial equilibrium model with stochastic volatility and shows how uncertainty shocks lead to a drop in investment and hiring by firms. Our paper is different in two important dimensions: first, we are looking at the business cycle in a general equilibrium model while he uses a partial equilibrium framework; and second, he does not test the impact of increases in uncertainty on firms using micro-data.

Finally, the paper is related to the literature studying macroeconomic models with micro-rigidities.[6] The paper contributes to this literature by finding that micro-rigidities of the type considered here (i.e. non-convex adjustment costs) coupled with time varying uncertainty have important GE effects.

The remainder of this paper is organized as follows. Section 2 discusses the behavior of uncertainty over the business cycle and presents evidence that variations in uncertainty affect aggregate output. In section 3 we formally present the model and define the recursive equilibrium. Since most of the business cycle literature has concentrated exclusively on first-moment shocks where the level of uncertainty is held constant, we depart from the standard log-linearization techniques and use non-linear methods instead. Specifically, section 3.4 presents our solution algorithm which builds on the work of Krusell and Smith (1998), Kahn and Thomas (2008) and Bachman, Caballero and Engel (2008). The model is calibrated and simulated in Section 4 where we study the role of uncertainty shocks in driving the business cycle. Section 6 tests the predictions of the model with the Census micro data, while section 5 studies the impact of policy shocks in the presence of time-varying uncertainty. Section 7 concludes.

2 The rise in uncertainty during recessions

This section consists of two parts. The first part presents a range of proxies for uncertainty and their movement over the business cycle. The different measures are also aggregated to form an uncertainty index. The second part provides evidence that time variation in the uncertainty index are associated with a large drops and rebound in economic activity in a VAR that also controls for the first moment.

[5] See, for example; Bernanke (1983), Pindyck (1988), Dixit (1990), Bertola and Bentolila (1990), Bertola and Caballero (1994), Dixit and Pindyck (1994), Abel and Eberly (1996), Hassler (1996), and Caballero and Engel (1999).

[6] See for example, Thomas (2002), Veraciertio (2002), Kahn and Thomas (2006 and 2008), Bachman, Caballero and Engel (2008), and House (2008).

2.1 Measuring uncertainty over the business cycle

There is no single good measure for uncertainty. However, we compile a wide range of sensible proxies for uncertainty and provide evidence that they are all strongly countercyclical. Each of these uncertainty proxies has obvious drawbacks, taken together, however, we believe they paint a robust picture of counter cyclical uncertainty. Details of the construction of each measure are contained in Appendix A.

2.1.1 Cross-firm and industry evidence

One approach to modelling time-series uncertainty is to examine the cross sectional spread of firm and industry level growth rates. The rationale for this is that uncertainty implies fatter tails in the distribution of productivity in the context of our model. Row (1) of Table 1 reveals that the cross-sectional spread of firm-level sales growth rates is 23.6% higher during quarters defined as recessionary by the NBER business cycle dating committee. This rise in the cross-sectional spread is also negatively correlated with real GDP growth at -0.455.

One interpretation of this result is that it simply reflects differential responses of firms in different industries to a common macro shock. For example, an oil shock will have a large negative impact on firms in energy intensive industries like aluminium production, but a positive effect on firms in energy producing industries like coal and gas production.[7] To address this issue we calculate the cross-sectional dispersion both within and across industries. We find the results to be practically identical. For example, recomputing this change in firm-level sales growth within each 2-digit SIC industry results in a rise of 20.4% (standard error of 3.3%), very similar to the rise of 23.6% in the pooled data.[8] Hence, there is a clear within industry spread of sales-growth during recessions, suggesting these results can not be driven simply by differential industry responses to common shocks.

This increase in the cross-sectional dispersion of sales growth is shown graphically in Figure 1, which plots the raw cross-sectional increase in sales dispersion (solid line) and the within SIC2 industry increase in sales dispersion (dashed line). Shown on the same graph as grey bars are the recessions according to the NBER. Two things stand out: the two lines are extremely similar, indicating that the increase in cross-sectional variance over the cyclical is very similar within and across broad industry groups. Second, cross-sectional sales spreads appears to rise strongly during recessions, most notably during the 1974/75 and the 2001 recession.

[7] This is essentially the critique Abrahams and Katz (1986) raised against Lillien (1982) who showed a strong correlation between cross-industry variation in unemployment and overall unemployment levels, arguing for a large role for structural unemployment during periods of rapid cross-industry movements of employment. Abrahams and Katz (1986) argued that it could instead be interpreted as a differential industry level response to common negative macro shocks.

[8] All caculations undertaken using 2-digit SIC cells with 25 or more firms, covering 29 2-digit SIC industries.

The result in row (2) shows that the cross-sectional spread of firm-level stock returns is also counter-cyclical, rising by 19.6% during recessions. Again, this is also true both within and across industries. For example, within 2-digit SIC industries the rise in the spread of stock-returns is 18.8% (standard error of 3.7%), again almost identical to the overall rise of 19.6%.[9] This is shown graphically in Figure 2, which again plots the raw cross-sectional spread of stock returns (solid line) and the within SIC2 spread (dashed line).

The results in row (3) of Table 1 reveal that the cross-sectional spread of industry level output growth for manufacturing industries is also substantially higher during recessions, rising by 61.5% as compared to non-recessionary periods. This is shown in Figure 3, which reveals that cross-industry spread rose, especially during the recessions of the 1970s and 1980s. Row (4) reveals that the cross-sectional spread of industry level productivity growth in manufacturing also appears to be significantly higher during recessions. Since industry productivity growth is only measured annually, we regress the annual spread on the share of quarters in a recession in that year to evaluate its association with the business cycle.[10] This coefficient suggests that cross-sectional industry productivity growth is 32.5% higher on average during years with all four quarters in a recession (and proportionally less for one, two or three quarters in a recession).

2.1.2 Macro measures of uncertainty

An alternative approach uses high-frequency aggregate data to infer the underlying process of stochastic volatility. This captures the common macroeconomic component of uncertainty, in contrast to the idiosyncratic firm and industry level components discussed above. In row (5) we see that the conditional heteroskedasticity of output growth is 57.0% higher in recessions. This is estimated from an regression of monthly industrial production on up to 12 lags and an GARCH(1,1) error process.[11] In row (6) we look at an index stock-market volatility, and find that recessionary quarters are associated with a 39.1% higher volatility of stock-market returns. Hence, output and stock-market data both suggest macro uncertainty is substantially higher during recessions.[12] These two aggregate measures of uncertainty are plotted in Figures 4 and 5, and again these measures appear to be reasonably

[9] All caculations undertaken using 2-digit SIC cells with 25 or more firms, covering 27 2-digit SIC industries. Campbell et al. (2001) also report this result, that cross-sectional stock returns variance is counter-cyclical.

[10] This comes from the NBER industry database, and is only measured annually as industry level capital-stock data is only available annually.

[11] Longer lags in the GARGH process were not significant.

[12] Hamilton and Lin (1996) provide evidence similar to Schwert (1989) that stock-market volatility is much higher during recessions. Engle and Rangel (2006) look at data from 48 countries (developed and developing) and find similar results for this panel (stock-market volatility is significantly higher when GDP growth is low). Officer (1973) compiles stock-market volatility data from about 1899 to 1960 and finds that the volatility of industrial productivity is correlated with stock-market volatility (especially during the Great Depression).

counter-cyclical.

2.1.3 Cross-forecaster evidence

Finally, we can interpret the extent of disagreement between forecasters over future macroeconomic variables as a proxy for uncertainty. It is important to note that theoretically cross-sectional forecaster disagreement and uncertainty are not necessarily correlated.[13] Moreover, our model does not induce forecaster disagreements. However, there is an extensive empirical literature that argues in favor of using disagreement among macro forecasts - as measured by mean forecast error - as a proxy for uncertainty.[14] In rows (7) and (8) we see that the dispersion of professional forecasts over future GDP and unemployment increases substantially during recessions, rising by 58.8% and 71.7% respectively. This supports the result from section (2.1.2) above, that macro uncertainty rises substantially during recessions. These two measures of forecaster disagreement are shown in figures 6 and 7.

2.1.4 The uncertainty index over the business cycle

In the final row of Table 1 we report an aggregated index of uncertainty. There are different options to combine the proxies we compiled. Here, we chose the simplest option and calculated the average of the seven quarterly uncertainty proxies after normalizing all of them to an average of one during non-recessionnary periods. The uncertainty index rises by about 42.5% during recessions, and is strongly negatively correlated with real industrial production growth at -0.52.[15] Figure 8 plots this series over time and it is obvious that recessions are periods of higher uncertainty. Interestingly, this relationship is stronger for the four recessions in the 1970s and 1980s, and weaker for the last two recessions in the early 1990s and 2000s. The recession in 2008 looks likely to generate another large uncertainty spike, since the data we plot only extends until 2008Q2, before the financial turmoil associated with the collapse of Lehman Brother. Two other notable spikes in the graph occur in 1987Q4 and 1998Q3, which are the Black Monday and the Russia/LTCM stock-market crashes respectively. Both episodes generated large increases in both firm and aggregate stock-returns uncertainty measures, but did not coincide with recessions.

[13] See for example Amador (2008).

[14] See, for example, Zarnowitz and Lambros (1987), Bomberger (1996), and Giordani and Soderlind (2004).

[15] One alternative is to use the principal component factor method. Here, it loads relatively evenly on all the indicators, with weights ranging from 0.17 to 0.27. This index is also strongly negatively correlated with real industrial production growth at -0.56.

2.2 A VAR analysis

The previous section has documented extensively that there is a significant unconditional correlation between uncertainty — as measured by an array of proxies — and the business cycle. The remainder of section 2 will provide some evidence that points to the existence of a conditional association. That is we find a large negative association between uncertainty and GDP growth even after controlling for TFP, wages, prices and interest rates.

At the macro-level it is obviously hard to identify the causal impact of any variable because of the absence of exogenous variation in the variable of interest. Nevertheless, we provide a VAR impulse response functions for a realistically calibrated increase in uncertainty to demonstrate that: (i) uncertainty can lead to a sizeable drop and subsequent rebound of economic activity; and (ii) this holds even with controls for TFP shocks. We will later see, that this empirical finding is consistent with the predictions of our model.

The uncertainty index is added to an otherwise standard VAR with the following variables (in the estimation order): the uncertainty shock, log(TFP) from Basu, Fernald and Kimball (2006)[16], the Federal Funds Rate, log(average hourly earnings), log(consumer price index), hours, log(employment) and log(GDP).[17] This ordering is based on the assumptions that the second-moment shock (uncertainty) can influence the first-moment shock (productivity), that first and second moment shocks influence prices (wages, the CPI and interest rates) and then quantities (hours, employment and output). Including TFP before the prices and quantities in the VAR ensures that the impact of first-moment productivity shocks is already controlled for when looking at the impact of the uncertainty index. All figures are scaled to show the impact of a 40% increase in uncertainty. This is meant to capture the empirical behavior of uncertainty in recession as illustrated in the previous section. A set of robustness checks in figures A1 and A2 in the appendix shows that our results are robust to different choices of variables, variable ordering and different methods of detrending.

In figure 9 we see the response of GDP to a 40% rise in the uncertainty index. There is an immediate drop in GDP which keeps following for two more quarters to reach a low of minus of 1.7% after two quarters falling by a rebound to trend within four quarters. In figure 10 we see the response of employment to an increase in uncertainty. Here, the maximum drop amounts to almost 3% after two quarters and a rebound to level within five quarters.

[16]We wish to thank John Fernald for providing this quarterly data.
[17]All variables are Hodrick Prescott (HP) detrended ($\lambda = 1600$) in the baseline estimations.

3 Model

The existing business cycle literature has focused almost completely on first moment shocks, such as productivity shocks, investment specific shocks, monetary shocks, government expenditure shocks, news shocks etc. It is an open question whether fluctuations in uncertainty matter for business cycles. By building an otherwise standard model that allows for time variation in uncertainty, firm level heterogeneity and non-convex adjustment costs, we can evaluate the quantitative significance of the mechanism proposed here.

Directly measuring time-variation in uncertainty is difficult. Moreover, while all the uncertainty proxies inroduced above appear to be strongly countercyclical, each could potentially be interpreted as the result of a specific first-moment shock. A structural model with shocks to both the first and second moment can provide us with testable implications that can help in distinguishing first and second moment shocks.

We thus proceed to analyze the potential role of variation in uncertainty within a dynamic stochastic general equilibrium model where heterogeneous firms are subject to both first and second moment shocks. In the model each firm uses capital and labor to produce a final good. Firms that adjust their capital stock and employment incur non-convex adjustment costs. As is standard in the RBC literature, firms are subject to an exogenous process of productivity. We assume that the productivity process has an aggregate and an idiosyncratic component. In addition to these first moment shocks, we allow the second moment of the innovations to productivity to vary over time. That is, shocks to productivity can be fairly small in normal times, but become potentially large when uncertainty is high. We now formally present the model.

3.1 Technology

The economy is populated by a large number of heterogeneous firms that employ capital and labor to produce a single final good. We assume that each firm operates a diminishing returns to scale production function with capital and labor as the variable inputs.[18] Specifically, a firm indexed by j produces output according to

$$y_{j,t} = A_t z_{j,t} k_{j,t}^{\alpha} n_{j,t}^{\nu} \ , \ \ \alpha + \nu < 1. \tag{1}$$

Each firm's productivity is a product of two separate processes: aggregate productivity, A_t, and an idiosyncratic component, $z_{j,t}$. Both the macro and the firm level component of

[18] An alternative model has a setup of monopolisticly competitive firms in which each firm produces a differentiated good. Note that the assumption of decreasing returns to scale implies that there is a fixed factor of production that pins down firm size.

productivity follow an autoregressive process.

$$\log(A_t) = \rho^A \log(A_{t-1}) + \sigma_t^A \epsilon_t \tag{2}$$
$$\log(z_{j,t}) = \rho^Z \log(z_{j,t-1}) + \sigma_t^Z \epsilon_{j,t} \tag{3}$$

We depart from the benchmark RBC model in that we allow the variance of innovations to the productivity processes, σ_t^A and σ_t^Z, to vary over time, following a two point Markov process.

$$\sigma_t^A \in \{\sigma_L^A, \sigma_H^A\} \quad \text{where } Pr(\sigma_{t+1}^A = \sigma_j^A | \sigma_t^A = \sigma_k^A) = \pi_{k,j}^{\sigma A} \tag{4}$$
$$\sigma_t^Z \in \{\sigma_L^Z, \sigma_H^Z\} \quad \text{where } Pr(\sigma_{t+1}^Z = \sigma_j^Z | \sigma_t^Z = \sigma_k^Z) = \pi_{k,j}^{\sigma Z} \tag{5}$$

Capital and labor adjustment costs. A firm's capital stock evolves according to the standard law of motion

$$\gamma k_{j,t+1} = (1 - \delta_k) k_{j,t} + i_{j,t} \tag{6}$$

where the $\gamma - 1$ is the trend growth rate of output. We assume that adjusting the capital stock involves adjustment costs. Based on prior empirical evidence we consider two types of adjustment costs.[19] The first one involves a non-convexity; conditional on undertaking an investment a fixed cost F^K is incurred independently of the scale of investment. The second capital adjustment cost we consider is a partial irreversibility; that is, the resale value of $1 of capital is S, which is below the purchase price of capital, $1 > S > 0$.

Similarly, we assume that whenever the firm changes the number of hours it employes it faces an adjustment cost. Specifically, we assume that the low of motion for hours worked is governed by

$$n_{t,t} = (1 - \delta_n) n_{j,t-1} + s_{j,t}. \tag{7}$$

That is, at each period a constant fraction, δ_n, of hours worked is exogenously destroyed due to retirement, illness, maternity leave, exogenous quits etc. Whenever the firm chooses to increase/reduce its stock of hours relative to $(1 - \delta_n) n_{j,t-1}$, a fixed cost F^L is incurred independently of the size of the change in hours. The firm also has to pay a cost per worker hired, H, or fired, F, representing for example variable interviewing and training costs or severance packages.

[19] The adjustment cost functions are a reduced form representations of a host of frictions that exit in practice. For instance, one could think of the effort that needs to be put into installing the capital, the disruption of current production when adding new machines or the relatively low resale value of machines that are specific to a certain plant. There is an extremely large literature focused on estimating these labor and capital adjustment costs, including, Nickell (1986), Caballero and Engel (1999), Ramey and Shapiro (2002), Hall (2004), Cooper and Haltiwanger (2006) and Bloom (2009).

The firm's value function. We denote by $V(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$ the value function of a firm. The seven state variables are given by (1) a firm's capital stock k, (2) a firm's hours stock from the previous period n_{-1}, (3) the firm idiosyncratic productivity $z_{j,t}$, (4) aggregate productivity A_t, (5) macro uncertainty σ_t^A, (6) micro uncertainty σ_t^Z and (7) the joint distribution of idiosyncratic productivity and firm-level capital stocks and hours worked in the last period μ_t which is defined for the product space $S = Z \times R_+ \times R_+$.

The dynamic problem of the firm consists of choosing investment and hours to maximize the present discounted value of future profit streams.

$$
\begin{aligned}
& V(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu) = \qquad (8) \\
& \max_{i,n} \left\{ \begin{array}{c} y_{j,t} - w(A, \sigma^A, \sigma^Z, \mu)n - i \\ -AC^k(k, k') - AC^n(n_{-1}, n) \\ +E\left[m\left(A, \sigma^A, \sigma^Z, \mu, A', \sigma^{A\prime}, \sigma^{Z\prime}, \mu'\right) V(k', n, z'; A', \sigma^{A\prime}, \sigma^{Z\prime}, \mu')\right] \end{array} \right\},
\end{aligned}
$$

given a law of motion for the joint distribution of idiosyncratic productivity, capital and hours,

$$
\mu' = \Gamma(A, \sigma^A, \sigma^Z, \mu), \qquad (9)
$$

and the stochastic discount factor m.

We denote by $K(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$ and $N^d(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$ the policy rules associated with the firm's choice of capital for the next period and current demand for hours worked.

3.2 Households

The economy is populated by a large number of identical households that we normalize to a measure one. Households choose a path of consumption, labor supply and investments in firm shares to maximize lifetime utility. We use the measure ϕ to denote those one period shares. The dynamic problem of the household is given by

$$
W(\phi, A, \mu) = \max_{\{C, N, \phi'\}} \left\{U(C, N) + \beta E\left[W(\phi', A', \mu')\right]\right\} \qquad (10)
$$

subject to the law of motion for μ and a sequential budget constraint

$$
\begin{aligned}
& C + \int q(k', n, z; A, \sigma^A, \sigma^Z, \mu)\phi'(dkdndz) \qquad (11) \\
\leq \; & w(A, \sigma^A, \sigma^Z, \mu)N + \int \rho(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)\phi(dkdndz).
\end{aligned}
$$

The households receive labor income as well as the sum of dividends and the resale value of their investments, $V(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$. With these resources the household consumes and buys new shares at a price $\rho(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$ per share of the different firms in the economy.

We denote by $C(\phi, A, \mu)$, $N^s(\phi, A, \mu)$, $\Psi(k', n, z; A, \sigma^A, \sigma^Z, \mu)$ the policy rules determining current consumption, time worked and quantities of shares purchased in firms that begin the next period with a capital stock that equals k' and who currently employ n hours. We are now ready to define the recursive competitive equilibrium.

3.3 Recursive competitive equilibrium

A recursive competitive equilibrium in this economy is defined by a set of quantity functions $\{C, N^s, \Psi, K, N^d\}$, pricing functions $\{w, q, \rho, m\}$, and lifetime utility functions $\{W, V\}$, such that:

1. V and $\{K, N^d\}$, are the value function and policy functions, respectively, solving (8).

2. W and $\{C, N^s, \Psi\}$ are the value function and policy functions, respectively, solving (10).

3. Asset markets clear

$$\Psi(k', n, z; A, \sigma^A, \sigma^Z, \mu) = \mu\left(z, k', n\right) \text{ for every triplet } \left(z, k', n\right) \in S$$

4. Goods markets clear

$$\begin{aligned} & C(\phi, A, \mu) \\ = & \int_S \left[\begin{array}{l} Azk^\alpha N^\nu(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)^\nu - \left(K(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu) - (1-\delta)k\right) \\ -AC^k(k, K(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)) - AC^n(n_{-1}, N(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)) \end{array} \right] \\ & \mu\left(dkdndz\right) \end{aligned}$$

5. Labor markets clear

$$N^s(\phi, A, \mu) = \int_S \left[N^d(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)\right] \mu\left(dkdndz\right)$$

6. The evolution of the joint distribution of z, k and n is consistent. That is, $\Gamma(A, \sigma^A, \sigma^Z, \mu)$ is generated by $K(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$, $N^d(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu)$ and the exogenous stochastic evolution of A, z, σ^Z and σ^A with the appropriate summation of firms' optimal choices of capital and hours worked given current state variables.

3.4 Sketch of the numerical solution

The model can be simplified substantially. We combine the firm and household problems into a single dynamic optimization problem as in Kahn and Thomas (2008). Starting with the household problem one can derive

$$w = -\frac{U_N(C,N)}{U_C(C,N)} \quad (12)$$

$$m = \beta\frac{U_C(C',N')}{U_C(C,N)} \quad (13)$$

where equation (12) is the standard optimality condition for labor supply and equation (13) is the standard expression for the stochastic discount factor. To ease the burden of computation it is useful to assume that the momentary utility function for the household can be specified as follows

$$U(C,N) = \frac{C_t^{1-\gamma}}{1-\gamma} - \theta N_t, \quad (14)$$

implying that the wage rate is a function of the marginal utility of consumption,

$$w_t = \frac{\theta}{C_t^{-\eta}}. \quad (15)$$

Kahn and Thomas (2003, 2008) and Bachmann, Caballero and Engel (2008) define the intertemporal price of consumption goods as $p(A, \sigma_z, \sigma^A, \mu) \equiv U_C(C,N)$. Using this approach, we can redefine the firm problem in terms of marginal utility, denoting the new value function as $\tilde{V} \equiv pV$. The firm problem can be expressed as

$$\tilde{V}(k, n_{-1}, z; A, \sigma^A, \sigma^Z, \mu) = \max_{\{i,n\}} \left\{ \begin{array}{c} p(A,\sigma^A,\sigma^Z,\mu)\left(y - w(A,\sigma^A,\sigma^Z,\mu)n - i - AC^k(k,k') - AC^n(n_{-1},n)\right) \\ +\beta E\left[\tilde{V}(k',n,z';A',\sigma^{A'},\sigma^{Z\prime},\mu')\right] \end{array} \right\} \quad (16)$$

We employ nonlinear techniques to solve for the optimal policy functions instead of log linearizing the model as is standard in the RBC literature.[20] Our solution uses a variation of the algorithm proposed by Krusell and Smith (1998). We reduce the large state vector of the model to include only the aggregate states (A, σ^A, σ^Z) and a small set of moments of the firm distribution which we will denote by Ω. The simplest example would be to use the average capital stock employed by all firms, but the set of moments can easily be

[20]In a log linearized model there is by construction no role for time varying second moments.

extended to include, for instance, the standard deviation of firm specific capital holdings. The solution algorithm then works as follows (full details in Appendix B). In iteration l, perform the following steps.

1. Forecast the intertemporal price $\hat{p}$ and next period's moments $\hat{\Omega}'$ as functions of the current aggregate state:

$$\begin{aligned} \hat{p} &= f_1^{(l)}(z, A, \sigma^A, \sigma^Z, \Omega) \\ \hat{\Omega}' &= f_2^{(l)}(z, A, \sigma^A, \sigma^Z, \Omega) \end{aligned}$$

2. For a given forecast of $\hat{p}$ we know the current period wage w from (**??**). We can then find the value function $\tilde{V}^l$ associated with those forecasting functions by solving (**??**) substituting the approximated state Ω for the joint distribution μ and $f_2^{(l)}$ for the law of motion Γ.

3. Simulate the economy for T periods. Here the forecasting rule for the intertemporal price is not used. Instead, in each period the market clearing price p_t is calculated as the price that combines firm optimization and goods market clearing. For a given price, the simplified firm optimization problem becomes

$$\max_{\{i,n\}} \left\{ p\left(y - wn - i - AC^k(k, k') - AC^n(n_{-1}, n)\right) + \beta E\left[\tilde{V}^l(k', n, z'; A', \sigma^{A\prime}, \sigma^{Z\prime}, \hat{\Omega}')\right]\right\}$$

   which uses the value function calculated in step II and the moment forecasting function from step I. Market clearing is achieved when aggregation of the optimal policies from this problem yield market clearing in the goods market

$$C = \int \left(y + i - AC^k - AC^n\right) \mu(dkdndz).$$

   This simulation yields sequences of prices $\{p_t\}$ and moments $\{\Omega_t\}$.

4. Update the forecasting functions $f_1^{(l+1)}$ and $f_2^{(l+1)}$ from the observed moments and equilibrium prices. Restart the algorithm at step 1 and iterate until the forecasting functions converge.[21]

[21] For the forecasting functions, we use the first moments of the distribution over capital and labor as well as the aggregate uncertainty state. Interestingly, this provides a very good fitand R2 of above 0.99.

4 Simulation

This section motivates the choice of parameter values used in the simulations and also presents simulation results for our preferred specification.

4.1 Calibration

First, we will discuss our parameter choices. At this stage, we do not provide an extensive set of robustness checks. Instead, we have tried to stay as close as possible to parameters that are widely used. Most importantly, we will the discuss our assumptions regarding the adjustment cost specification and the calibrated process of uncertainty – two important drivers of our results. Table 2 summarizes all the calibrated parameter values.

Frequency and preferences. We set the time period to equal a quarter. The household's discount rate is calibrated to yield an annual interest rate of 4%. Moreover, setting η to equal one implies that the momentary utility function features an elasticity of intertemporal substitution of one. Our utility function (14) assumes an infinite Frisch labor supply elasticity, or indivisible labor, which simplifies our solution algorithm substantially. We set the parameter θ such that households spend a third of their time working in the non stochastic steady state. The trend growth rate of per capita consumption is set to equal 1.6% annually.

Production function, depreciation and adjustment costs. The capital depreciation rate is set to match the average annual depreciation rate of 10%, leading us to set $\delta_k = 0.025$.

The annual exogenous quit rate of labor is a key parameter, and is set to 15%. This estimate is based on the quit rate reporeted in the Bureau of Labor Statistic JOLTS data.[22] In JOLTS the monthly quit figure varies between 1.6% and 2.4%, with the lowest value occuring in November 2008. Since this November 2008 quit rate occured during the depths of the 2008 recession, we believe this represents the level of exogenous quits that occur during a recession. Annualizing this figure we get 19.2%. We assume a lower 15% annualized figure to be conservative, but note that the 19.2% quit rate would make the impact of uncertainty shocks even more dramatic.

We set the firm's production function elasticity with respect to its capital stock, $\alpha = 0.25$ and $\nu = 0.5$, consistent with a capital cost share of 1/3 and a 25% markup when the firm faces an iso-elastic demand curve.

[22] JOLTS standards for Job Openings and Labor Turnover Data, which the BLS has been collecting since January 2001. Hence, this data spans two NBER defined recessions. It distinguished between quits, layoffs and other seperations. Our figures are seasonally adjusted for total private employment.

The existing literature provides a wide range of estimates for capital and labor adjustment costs.[23] We expect future versions of this paper to provide more guidance on the sensible level of adjustment costs for a model like this despite the fact that a full structural estimation from the model is going to be a computational challenge. We thus set our adjustment cost parameters to match Bloom (2009), which is the only paper (we are aware of) that jointly estimates capital and labor convex and non-convex adjustment costs. Fixed costs of capital adjustment are set to 1.5% of annual sales, the resale loss of capital amounts to 40%. The fixed cost of adjusting hours, is set to 2.1% of annual wages, the hiring and firing costs equals 1.8% again in terms of wages.

Aggregate and idiosyncratic TFP process. We approximate the autoregressive processes (2) and (3) with Markov chains. The support for the process is set to include three standard deviations on either side of the mean. For idiosyncratic TFP, we increase this range to three times the standard deviation on either side. The parameters of the processes are taken from Khan and Thomas (2008) and adjusted to the quarterly frequency. Hence, ρ^A and ρ^Z are set to yield an annual persistence parameter of 0.859, whereas the standard deviation of innovations to the aggregate and idiosyncratic productivity process are set to yield an annual equivalent of 0.014 and 0.022, respectively.

The calibrated process for uncertainty. When the economy enters the uncertain state, we assume that the standard deviation of innovations to both idiosyncratic and aggregate productivity doubles. This is calibrated so that the average increase in our simulated data of (i) cross-firm sales growth, (ii) cross-firms stock-returns, (iii) macro output volatility, and (iv) macro stock market volatility are 40% on average, matching the increases in these uncertainty proxies in the real data shown in section 2. That is, we calibrate our model to achieve the same increase in the simultaed uncertainty proxies as we observe in actual data during recessions. In doing so we mimic every step to measure uncertainty in the simulations as in the actual data. In the absence of a business cycle dating comittee in the model, we define recessions as 2+ quarters of declining output in the simulation.

We assume for simplicity that the stochastic volatility processes σ_t^A and σ_t^Z follow 2-state Markov chains. In fact, the current version of the paper assumes that the values of aggregate and idiosyncratic uncertainty are driven by just one exogenous process that determines whether the economy is in the normal or uncertain regime. We estimate the transition probabilities from the uncertainty index we assemble in section 2 as follows. We find that the first order serial coloration is 0.75. Under the assumption that the process of

[23] See, for example, Hayashi (1982), Nickel (1986), Shapiro (1986), Caballero and Engel (1999), Ramey and Shapiro (2001), Hall (2004), Cooper, Haltiwanger and Willis (2004), Cooper and Haltiwanger (2006) as well as Mertz and Yashiv (2007).

the uncertainty index can be approximated an AR1 process, this implies that the diagonal entries of the Markov chain have to sum to 1.75. Informal introspection of the uncertainty index leads us to set the diagonal entries such that in the ergodic distribution the economy is 80% in normal times and 20% in volatile times. These two condition imply that the Markov chain is given by,

	normal	*uncertain*
normal	0.95	0.05
uncertain	0.2	0.8

4.2 The effect of an uncertainty shock

We first study the effects of an isolated increase in uncertainty. We simulate 1000 economies until each reaches its ergodic distribution with respect to z and the distribution over k and n while forcing all economies to be in the low uncertainty state. At period 0 all economies are subject to an uncertainty shock. That is, agents learn that starting in period 1 the distribution of productivities fans out. The magnitude of the shock is such that the standard deviation of the innovation to idiosyncratic and aggregate productivity doubles. The expected duration of the process is governed by the transition matrix. This leads to an increase in cross-sectional stock-return variance of about 40% during the recession period, matching the real data discussed in section 2.

Figure 13 plots the fraction of economies that are in the high uncertainty state at each quarter, while figures 14 and 15 show that the aggregate and firm-level productivity does not change in this experiment. Thus, the baseline simulation results are driven entirely by changes in uncertainty. Figure 16 plots the cross sectional dispersion of idiosycratic productivities. In what follows we report the behavior of the variables of interest as the average over these economies.

The time profile of hours worked is shown in figure 17. In period zero, once uncertainty rises, firms defer most hiring decisions, leading to a fall of about 2% in hours worked. Hours continue to fall until they reach a level of about 3% below the steady state. By quarter 3 uncertainty falls enough and productivity fans out suffciently that firms with high productivity draws start to increase hours. By quarter 5 the eceonomy have reverted back to the initial trend.

What is potentially surprising is the overshoot in labor that arises from quarter 6 onwards. Labor use actually rises above its long-run value for about ten quarters before returning to this long-run level. The reason is that in the economy many firms are bunched near the hiring threshold due to labor attrition. Then, small increases in productivity will lead them to hire more workers, while small decreases will simply move them towards the interior of the Ss bands. As a result the increased variance of the productivity shock induced by higher uncertainty increases aggregate medium-run hiring. Firms that have received positive productivity shocks hit their Ss band and hire. Firms the receive a negative shock move to the interior of the Ss space and do nothing.

Since at period 0 each firm's capital stock is given, the fall in aggregate hours worked manifest itself in a drop in aggregate output as can be seen in figure 18. Interestingly, both the magnitude and evolution of output after the uncertainty shock is very similar to the VAR findings from section 2. The magnitude of the drop is 1.7% in the simulation and 1.8% in the VAR, while the duration of the drop and rebound is 6 months in both the simulation and VAR. Moreover, both simulation and VAR also show an overshoot, although this is much more pronounced in the data than in the simulation. This provides empirical support for the notion that the model – which is an extension of a standard neoclassical business cycle model with standard parameter values – is capturing something important in the determination of business cycles. The uncertainty shock induces a drop and rebound in investment as shown in Figure 19. In period 0 once uncertainty rises investment expenditure instantly falls by 15% as firms postpone investment decisions. Substantial capital adjustment costs lead many firms to defer new investment expenditure until after uncertainty has subsided. From quarters 2 to 4 investment rises, but it is still below trend. By quarter 5, however, uncertainty has fallen sufficiently to lead firms start to addressing their pent-up demand for investment, which continues for about 10 quarters.

Figure 20 plots the time profile of consumption. Upon impact, i.e. in quarter zero, consumption jumps up and then falls below trend for about 12 quarters. The reason for the initial spike in consumption is that the freeze in investment and hiring reduces the resources spent on capital and labor adjustment. Since the interest rate drops upon impact, consumers are signaled that consumption is cheap which leads to an increase of consumption in period 0. In other words, even though consumers know they face higher uncertainty in the future and they would like to save more, they do not increase savings in the first period because the returns to saving have become (temporarily) low and very risky. This logic suggests that if we extended the model to allow for some alternative savings technology – for example inventories or savings abroad – this initial spike in consumption would disappear as the representative consumer would just increase savings through this channel when the uncertainty shock hit to reduce the subsequent drop in consumption.

Finally, figure 21 plots the value for the aggregate Solow residual, defined as $Y_t/K_t^{\alpha}N_t^{\nu}$. This also shows a clear drop and rebound after the uncertainty shock, despite the fact that the average micro and macro productivity shock were unchanged as shown in figures 14 and 15. The reason is that uncertainty freezes the reallocation of capital and labor from low to high productivity firms. In normal times unproductive firms contract and productive firms expand, helping to maintain high productivity levels. When uncertainty is high, firms reduce expansion and contraction, shutting of much of this productivity enhancing reallocation leading to a fall in productivity growth rates. When uncertainty falls back down again firms rapidly address their pent-up demand for reallocation so that productivity returns towards its long-run trend. After the drop and rebound aggregate TFP overshoots from quarter 5 until about quarter 15. The reason is that the greater dispersion in firm-level productivity provides more opportunity for cross-sectional reallocation of output, increasing aggregate productivity.

Thus, to conclude, a second moment shock induces a fall in investment, hours, and output. Intriguingly, the Solow residual also falls even though this is an effect rather than a cause of the drop in economic activity. Consumption also exhibits a fall quarters 2 onwards, although there is an inital one period jump in the basic model. Finally, figure 22 shows that detrended output falls fast enough in reaction to the uncertainty shock, that even when the 1.6% trend growth is added back in output will still drop.

5 Policy in the presence of uncertainy

To be added.

6 Model's implication versus the Micro data

Our application for Census data access has just been approved. We expect to add results soon.

7 Conclusions

This paper proposes time variation in uncertainty as an additional impulse driving business cycles. The argument proceeds in three steps.

First, we demonstrate that uncertainty, measured by a number of proxies, appears to be strongly countercyclical. When added to a standard vector-autoregression, increases in uncertainty lead to a large drop and rebound in economic activity. This result holds even after controlling for first moment effects.

Second, we study a dynamic stochastic general equilibrium model that allows for shocks to both the level of technology (the rst moment) as well as uncertainty (the second moment). More specically, we model shocks to uncertainty as variation in the standard deviation of the innovations to productivity. We nd that increases in uncertainty lead to large drops in employment and investment. This occurs because uncertainty makes rms cautious, leading them to pause hiring and investment. This freezing in activity also reduces the reallocation of capital and labor across rms, leading to a large fall in productivity growth. Taken together, the freeze in hiring, investment and productivity growth lead to a business cycle sized drop and rebound in output following a rise in uncertainty.

Third, we will evaluate the empirical performance of the model using detailed US Census micro data. In particular, we will look at the response of capital and labor to productivity shocks at the establishment level. This will allow us to distinguish between rst and second moment shocks in the data. First moment shocks shift the Ss bands to the left and thus imply more disinvestment and less investment. Second moment shocks on the other hand lead to an outward shift of the Ss bands and to a freeze in both investment and disinvestment. Our application to gain access to Census data has just been approved. We plan to add results to this part of the project soon.

References

Abel, A.B. and Eberly, J.C. (1996), "Optimal Investment with Costly Reversibility", Review of Economic Studies, 63(4), 581-593.

Abraham, K.G. and Katz, L.F. (1986), "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?", Journal of Political Economy, 9(3), 507-522.

Aghion, P., Angeletos, G., Banerjee, A. and Manova, K. (2005), "Volatility and Growth: Credit Constraints and Productivity-Enhancing Investment", NBER WP 11349.

Amador, M. (2008), "XXX".

Bachman, R., Caballero, R.J. and Engel, E.M.R.A. (2008), "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model", Yale mimeo.

Bartelsman, E., Becker, R. and Gray, W. (2000), "NBER Productivity Database", www.nber.org.

Barlevy, G. (2004), "The Cost of Business Cycles Under Endogenous Growth", American Economic Review, 94(4), 964-990.

Basu, S., Fernald, J. and Kimball, M.S. (2006), "Are Technology Improvements Contractionary?", American Economic Review, 96(5), 1418-1448.

Bernanke, B. (1983), "Irreversibility, Uncertainty and Cyclical Investment", Quarterly Journal of Economics, 98, 85-106.

Bentolila, S. and Bertola, G. (1990), "Firing Costs and Labor Demand: How Bad is Eurosclerosis", Review of Economic Studies, 57(3), 381-402.

Bernard, A., Redding, S. and Schott, P. (2006), "Multi-Product Firms and Product Switching", NBER WP 12293.

Bertola, G. and Caballero, R.J. (1994), "Irreversibility and Aggregate Investment", Review of Economic Studies, 61, 223-246.

Bloom, N. (2009), "The Impact of Uncertainty Shocks", forthcoming Econometrica.

Bloom, N. (2007), "Uncertainty and the Dynamics of R&D", American Economic Review Papers and Proceedings, 97(2), 250-255.

Bomberger, W. A. (1996), "Disagreement as a Measure of Uncertainty", Journal of Money, Credit and Banking, 28(3), 381-392.

Caballero, R.J. and Engel, E.M.R.A (1999),"Explaining Investment Dynamics in U.S. Manufacturing: a Generalized (S,s) Approach", Econometrica, 67(4), 783-826.

Caballero, R.J. and Leahy, J. (1996), "Fixed costs: the demise of marginal Q", NBER WP 5508.

Campbell, J., Lettau, M., Malkiel B. and Xu, Y. (2001), "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk", Journal of Finance, 56(1), 1-43.

Davis, S.J. and Haltiwanger, J. (1992), "Gross Job Creation, Gross Job Destruction, and Employment Reallocation", Quarterly Journal of Economics, 107, 819-863.

Davis, S.J., Haltiwanger, J., Jarmin, R. and Miranda, J. (2006), "Volatility and Dispersion in Business Growth Rates: Publicly Traded Versus Privately Held Firms", NBER WP 12354.

Dixit, A.K. and Pindyck, R.S. (1994), "Investment Under Uncertainty", Princeton University Press, Princeton.

Engle, R.F. and Rangel, J.G. (2008), "The Spline-GARCH Model for Low-Frequency Volatility and Its Global Macroeconomic Causes", Review of Financial Studies, 21(3), 1187-1222.

Faggio, G., Salvanes K.G. and Van Reenen, J. (2007), “Understanding Wage and Productivity Dispersion in the United Kingdom”, mimeo.

Foster, L., Haltiwanger, J. and Krizan, C.J. (2000), "Aggregate Productivity Growth: Lessons from Microeconomic Evidence", New Developments in Productivity Analysis, NBER, University of Chicago Press.

Foster, L., Haltiwanger, J. and Krizan, C.J. (2006), "Market Selection, Reallocation and Restructuring in the U.S. Retail: Trade Sector in the 1990s", Review of Economics and Statistics, 88, 748-758.

Koren, M. and Tenreyro, S. (2007), "Volatility and Development", Quarterly Journal of Economics, 122(1), 243-287.

Gilchrist, S. and Williams, J.C. (2005), "Investment, Capacity, and Uncertainty: A Putty-Clay Approach", Review of Economic Dynamics, 8(1), 1-27.

Giordani,P. and Soderlind, P (2004), "Inflation Forecast Uncertainty", European Economic Review, 47(6), 1037-1059.

Hall, R.E. (2004), "Measuring Factor Adjustment Costs", Quarterly Journal of Economics, 119, 899-927.

Hamilton, J.D. and Lin, G. (1996), "Stock Market Volatility and the Business Cycle", Journal of Applied Econometrics, 11(5), 573-593.

Hammermesh, D. (1989), "Labor demand and the structure of adjustment costs", American Economic Review, 79, 674-89.

Hassler, J. (1996), "Variations in Risk and Fluctuations in Demand - a Theoretical Model", Journal of Economic Dynamics and Control, 20, 1115-1143.

House, C.L. (2008), "Fixed Costs and Long-Lived Investments", NBER WP 14402.

Jaimovich, N. and Siu, H. (2008), "The Young, the Old, and the Restless: Demographics and Business Cycle Volatility", forthcoming American Economic Review.

Keynes, J.M. (1936), "The General Theory of Empployment, Interest and Money", Macmillan, London.

Lillien, D.M. (1982), "Sectoral Shifts and Cyclical Unemployment", Journal of Political Economy, 90(4), 777-793.

Khan, A. and Thomas, J.K. (2003), "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?", Journal of Monetary Economics, 50, 331-360.

Khan, A. and Thomas, J.K. (2008), "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics, Econometrica, 76(2), 395-436.

King, R.G. and Rebelo, S.T. (1999), "Resuscitating Real Business Cycles", in Handbook of Macroeconomics, John B. Taylor and Michael Woodford (eds.), Elsevier.

Krusell, P and Smith, A.A. (1998), "Income and Wealth Heterogeneity in the Macroeconomy", Journal of Political Economy, 106(5), 867-896.

Mele, A. (2007), "Asymmetric Stock Market Volatility and the Cyclical Behavior of Expected Returns", Journal of Financial Economics, 86, 446-478.

Merz, M. and Yashiv, E. (2007), "Labor and the Market Value of the Firm", American Economic Review, 97(4), 1419-1431.

Nickell, S.J. (1986), "Dynamic Models of Labor Demand", in O. Ashenfelter and R. Layard (eds.), Handbook of Labor Economics, North-Holland, Amsterdam.

Officer, R. (1973), "The Variability of the Market Factor of the New York Stock Exchange", Journal of Business, 46, 434-453.

Pindyck, R.S. (1988), "Irreversible Investment, Capacity Choice, and the Value of the Firm", American Economic Review, 78(5), 969-985.

Ramey, V.A. and Ramey, G. (1994),"Cross-Country Evidence on the Link between Volatility and Growth", American Economic Review, 85, 1138-1151.

Ramey, V.A. and Shapiro, M.D. (2001) "Displaced Capital: A Study of Aerospace Plant Closings", Journal of Political Economy, 109, 605-622.

Rebelo, S.T. (2005), "Real Business Cycle Models: Past, Present, and Future", Scandinavian Journal of Economics, 107(2), 217-238.

Schwert, G.W. (1989), "Why Does Stock Market Volatility Change Over Time?", Journal of Finance, 44, 1115–1153.

Storesletten, K., Telmer, C.I. and Yaron, A. (2004), "Cyclical Dynamics in Idiosyncratic Labor-Market Risk", Journal of Political Economy, 112(3), 695-717.

Thomas, J.K. (2002), "Is Lumpy Investment Relevant for the Business Cycle?", Journal of Political Economy, 110(3), 508-534.

Veracierto, M.L. (2002),"Plant Level Irreversible Investment and Equilibrium Business Cycles", American Economic Review, 92(1), 181-197.

Zarnowitz, V. and Lambros, L.A. (1987), "Consensus and Uncertainty in Economic Prediction", Journal of Political Economy, 95(3), 591-621.

A Appendix: Uncertainty Data

To be added.

B Appendix: Numerical Solution Method

To be added.

Table 1: The Increase in Measures of Uncertainty During Recessions

		% increase during recessions, mean *(standard deviation)*	correlation with quart. ind. production growth	period covered
(1)	**Firm sales growth spread** (quarterly cross-sectional interquartile range)	23.6 *(3.8)*	-0.455	65Q3 to 08Q2
(2)	**Firm stock returns spread** (quarterly cross-sectional interquartile range)	19.6 *(3.6)*	-0.306	62Q2 to 08Q2
(3)	**Industry output growth spread** (quarterly cross-sectional interquartile range)	61.5 *(5.3)*	-0.481	72Q1 to 08Q3
(4)	**Industry productivity growth spread** (annual cross-sectional interquartile range)	32.5 (9.4)	-0.533	1962-1996
(5)	**Macro output growth volatility** (quarterly average ARCH errors on monthly AR(12))	57.0 *(13.3)*	-0.249	63Q1 to 08Q3
(6)	**Macro stock returns volatility** (quarterly standard deviation of daily stock returns)	39.1 *(6.8)*	-0.367	62Q2 to 08Q4
(7)	**Forecaster predicted industrial production spread** (quarterly standard deviation/mean)	58.8 *(10.1)*	-0.249	68Q4 to 08Q3
(8)	**Forecaster predicted unemployment spread** (quarterly standard deviation/mean)	71.7 *(7.1)*	-0.502	62Q4 to 08Q3
(9)	**Uncertainty index** (average of normalized individual measures)	42.5 *(3.4)*	-0.524	68Q4 to 08Q2

Notes: The first column reports the precentage increase of the variable during a recession, with standard errors in ( ). For example, row (1) reports that the interquartile range of quarterly firm-level sales growth is 23.6% higher during a recession, with a standard error of 3.8%. Recessions defined using the NBER dates. The second column shows the correlation with quarterly real GDP growth (all correlations significant at the 1% level). Row (1) contains the interquartile range of firm sales growth, defined as $(\text{sales}_{t+2} - \text{sales}_{t-2})/(0.5\times\text{sales}_{t+2} + 0.5\times\text{sales}_{t-2})$, where the 4-quarter window is used to remove quarterly accounting effects. Data from quarterly Compustat. Row (2) contains the interquartile range of firm stock returns, using CRSP monthly data. Row (3) reports the quarterly average of the monthly interquartile range of growth rates of industrial production, covering the manufacturing sector broken down into 196 NAICS 3 to 6 digit sectors, from the Federal Reserve Board. Row (4) reports the annual cross-sectional interquartile range of 5-factor TFP (from the NBER manufacturing database). The 32.5% figure denotes the % increase in the spread for a year with all four quarters in recession, estimated by regressing the annual spread against an indicator for the fraction of quarters in a recession in that year. Row (5) contains the average quarterly conditional standard deviation of monthly industrial production estimated from an AR(12) regression with GARCH(1,1) errors. Row (6) contains the quarterly average of the CBOE's VXO index (implied volatility on the S&P 100) from 1987 onwards, and the standard deviation of daily returns on the S&P 500 prior to 1987. Rows (7) and (8) contain standard deviation/mean of the 4-quarter ahead forecasts for unemployment and nominal GDP, obtained from the Survey of Professional Forecasters. Row (9) contains the average of the seven quarterly indicators in columns (1) to (3) and (5) to (8) after they have all be normalized to have a value of unity during non-recessions. Full data details in Appendix A.

Table 2: Parameters in the Model

Parameter	**(Annual) Value**	**Rationale** (also see the text)
β	0.977	Annual interest rate of 4%.
η	1	Elasticity of intertemporal substitution set to 1.
θ	1.813	Households spend one third of their time working in steady state.
γ	0.016	Trend growth rate in per capita consumption set to 1.6%.
α	0.25	Constant-returns-to-scale production function, 25% markup with iso-elastic
ν	0.50	demand curve, capital share of one third and labor share of two thirds.
δ_k	0.10	Annual depreciation of 10%.
δ_n	0.15	Annual exogenous quit rate of 15% (JOLTS).
F^K	1.5%	Fixed cost of changing the capital stock is 1.5% of annual sales (Bloom (2009)).
S	40%	Resale loss of capital of 40% (Bloom (2009)).
F^L	2.1%	Fixed cost of changing the number of hours is 2.1% of annual wage bill (Bloom (2009)).
H	1.8%	Per hour hiring cost of 1.8% of annual wage (Bloom (2009)).
F	1.8%	Per hour firing cost of 1.8% of annual wage (Bloom (2009)).
ρ^A	0.859	Persistence of aggregate productivity (Khan and Thomas (2008)).
ρ^Z	0.859	Persistence of idiosyncratic productivity (Khan and Thomas (2008)).
σ^A_{normal}	0.014	Std. of innovations to aggregate productivity in normal state.
σ^Z_{normal}	0.022	Std. of innovations to idiosyncratic productivity in normal state.
$\sigma_{uncertain}$	$2 * \sigma_{normal}$	Matches increase of uncertainty proxies in simulated data by 40%.
$\pi^\sigma_{L,H}$	0.05	Serial correlation of uncertainty index is 0.75, hence diagonal elements have to sum
$\pi^\sigma_{H,H}$	0.80	to 1.75. In the ergodic distribution the economy spends 80% in normal state.

Figure 1 Cross-firms sales growth spread

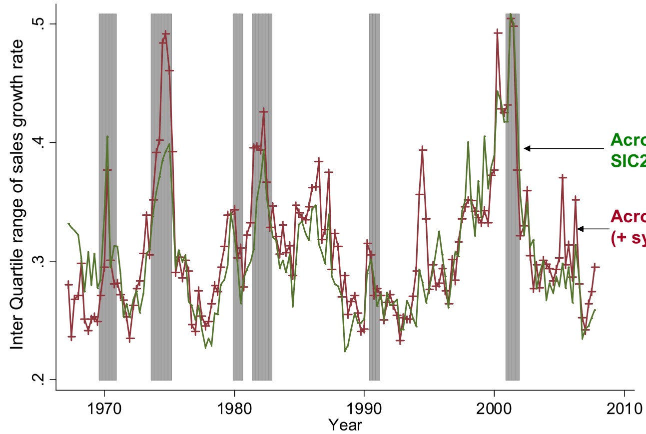

Notes: All Compustat firms with at least 25 years of quarterly accounts used (to reduce the impact of compositional changes). Only quarters with at least 500 firms kept to ensure sufficient sample size. Sales growth spread is defined as the inter-quartile range (IQR) across firms within each quarter. Sales growth is defined over a four quarter period to remove the effects of the quarterly accounting cycle, with this centered around the current quarter so that (Sales Growth)$_t$ = (Sales$_{t+2}$ - Sales$_{t-2}$)/(0.5xSales$_{t+2}$ + 0.5xSales$_{t-2}$). Spread within SIC2 industry uses only SIC2 year cells with at least 25 observations to calculate the IQR. Quarterly values are averages across the 24 SIC2 classes with 25 or more firms per quarter. The grey shaded columns are recessionary quarters defined according to the NBER.

Figure 2: Cross-firm stock-returns spread

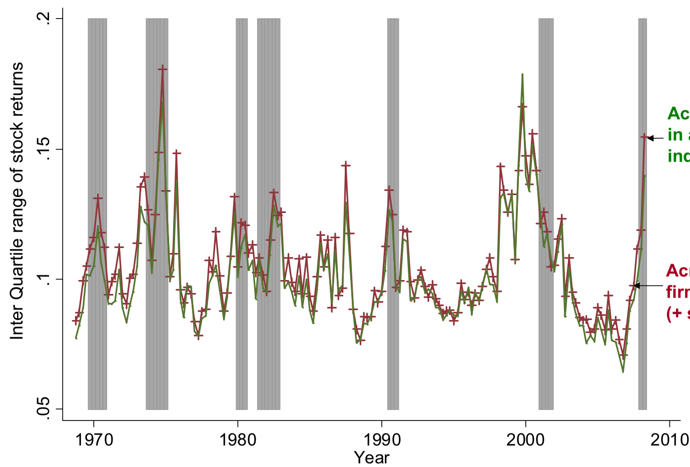

Notes: All firms with at least 25 years of quarterly returns used to reduce the impact of compositional changes. Only quarters with at least1 1000 firms kept to ensure sufficient sample size. Stock returns spread calculation using the inter-quartile range (IQR) across firms within each quarter. Spread within SIC2 industry uses only SIC2 year cells with at least 25 observations to calculate the IQR. Quarterly values are averages across the 25 SIC2 classes with 25 or more firms per quarter. The grey shaded columns are recessionary quarters defined according to the NBER.

Figure 3: Cross-industry output spread

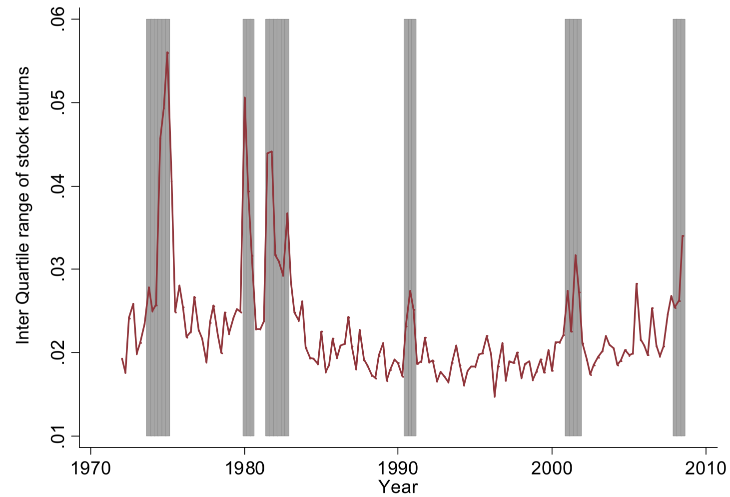

Note: Plots the inter-quartile range of the 3-month growth rates of industrial production, 1972 to 2008. Sample covers the manufacturing sector broken into 196 NAICS sectors, taken from the Federal Reserve Board. The grey shaded columns are recessionary quarters defined according to the NBER.

Figure 4: Industrial-production growth volatility

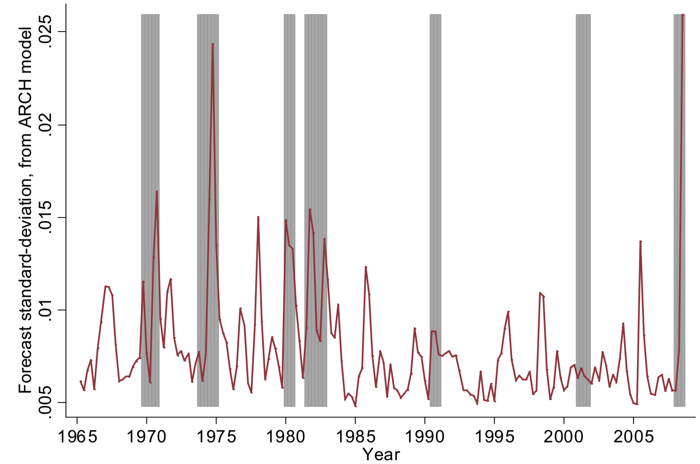

Notes: Industrial production standard-deviation (SD) defined as predicted SD from a regression of monthly log industrial production (Final products and non-industrial supplies seasonally adjusted, FRB Statistics) on its own 12 lags, with an GARCH(1,1) error term. Longer lags in industrial productions or the GARCH error term were not significant. Quarterly values of SD averaged over the monthly values within the quarter. The grey shaded columns are recessionary quarters defined according to the NBER.

Figure 5: Stock-market volatility

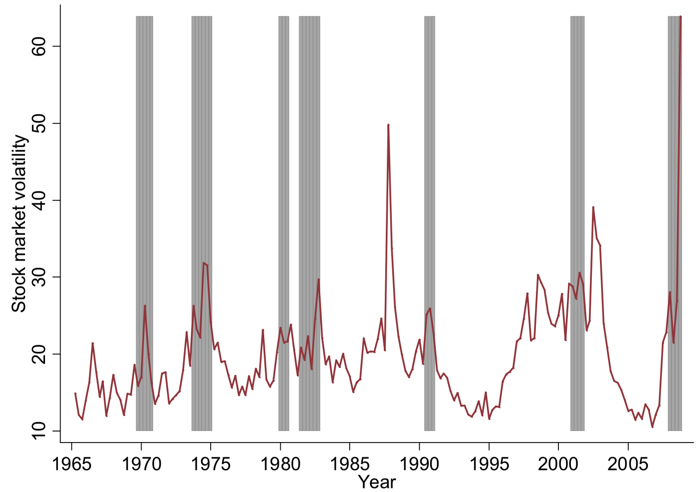

Notes: Stock market volatility used actual quarterly standard deviation of daily returns until 1987, and average quarterly implied volatility from 1987 onwards (see Bloom, 2007 for details). The grey shaded columns are recessionary quarters defined according to the NBER.

Figure 7: Forecaster industrial production dispersion

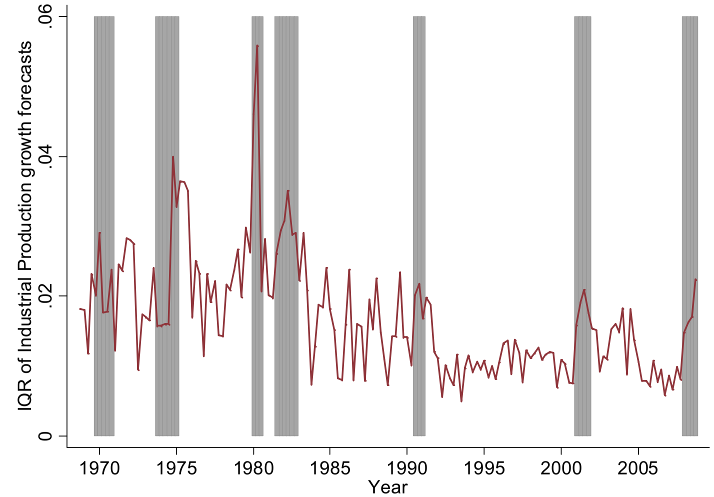

Notes: Standard deviation of cross-sectional forecasts of 4 quarters ahead GDP growth (nominal) from the Survey of Professional Forecasters. Forecasts collected quarterly with an average of 41 forecasters per period. The grey shaded columns are recessionary quarters defined according to the NBER.

Figure 6: Forecaster unemployment dispersion

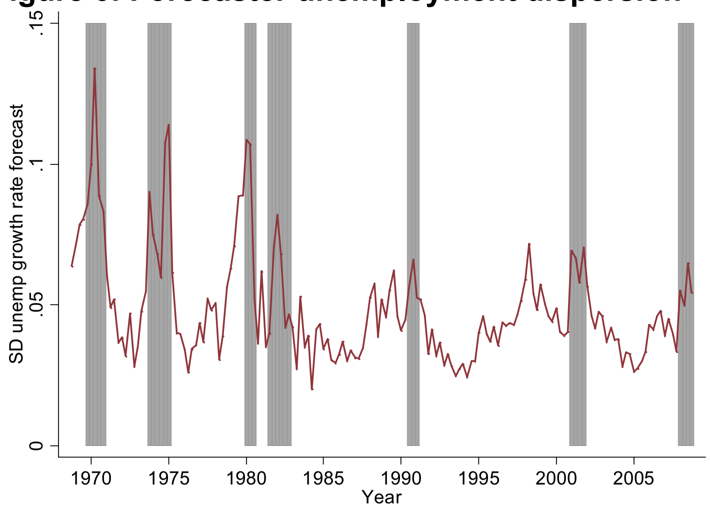

Notes: Standard deviation of cross-sectional forecasts divided by average of cross-sectional forecasts, 4 quarters ahead unemployment rates from the Survey of Professional Forecasters. Forecasts collected quarterly with an average of 41 forecasters per period. The grey shaded columns are recessionary quarters defined according to the NBER.

Figure 8: The aggregate uncertainty proxy

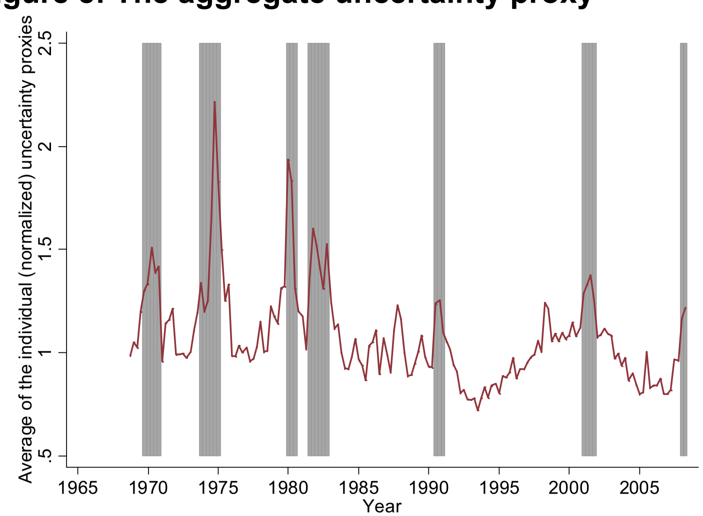

Notes: Uncertainty proxy defined as the average of the seven quarterly measures of uncertainty, after these have been normalized to have a value of unity on average during non-recessions. During 1968Q4-1971Q4 and 2008 only 6/7 indicators are available, so the average value is taken over those 6. Gray-shading denotes NBER defined recession quarters.

Figure 9: VAR estimation of the 'impact' of a recessionary increase in uncertainty on GDP

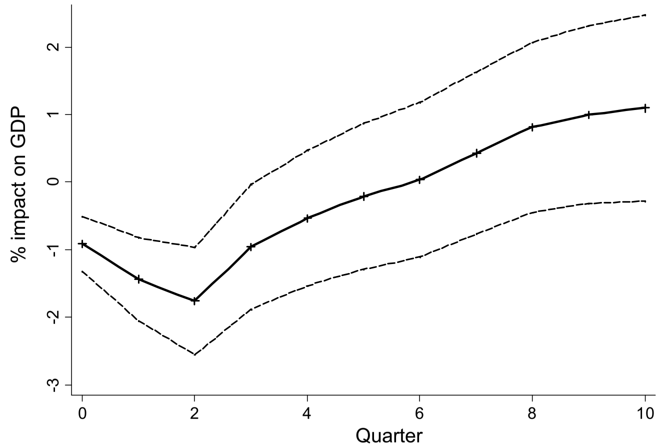

Notes: VAR Cholesky orthogonalized impulse response to a 40% increase in the uncertainty index, calibrated to the average increase on the uncertainty index during a recession. Estimated on quarterly data from 1968:4 to 2006:5 using 4 lags. Dotted lines in top and bottom figures are the 95% confidence intervals. Variables (in order) are the uncertainty index, log TFP (from Basu et al. 2006), the federal funds rate, log wages, log CPI, hours, log employment and log real GDP.

Figure 10: VAR estimate of the 'impact' of a recessionary increase in uncertainty on employment

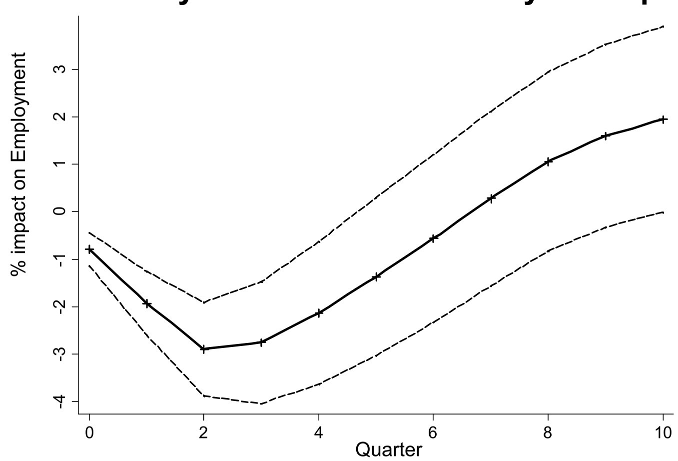

Notes: VAR Cholesky orthogonalized impulse response to a 40% increase in the uncertainty index, calibrated to the average increase on the uncertainty index during a recession. Estimated on quarterly data from 1968:4 to 2006:2 using 4 lags. Dotted lines in top and bottom figures are the 95% confidence intervals. Variables (in order) are the uncertainty index, log TFP (from Basu et al. 2006), the federal funds rate, log wages, log CPI, hours, log employment and log real GDP.

Figure A1: VAR estimates of the impact of a recessionary increase in uncertainty on GDP

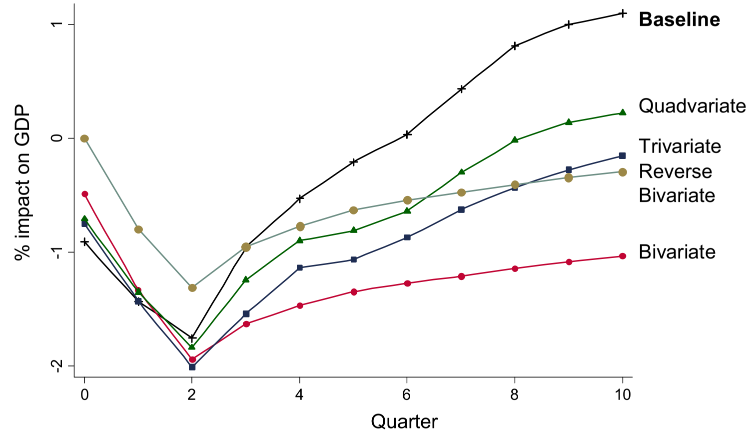

Notes: VAR Cholesky orthogonalized impulse responses to a 40% increase in the uncertainty index, calibrated to the average increase on the uncertainty index during a recession. Estimated on quarterly data from 1968:4 to 2006:2 using 4 lags. Variables (in order) are: (i) "Bivariate": Variables (in order) are the uncertainty index and log real GDP; (ii) "Trivariate" uncertainty index, FFR and log real GDP; (iii) "Quadvariate" uncertainty index, FFR, log employment log GDP and (iv) "Reverse Bivariate" log real GDP and the uncertainty index. "Baseline" as in previous graphs.

Figure A2: VAR estimates of the impact of a recessionary increase in uncertainty on GDP

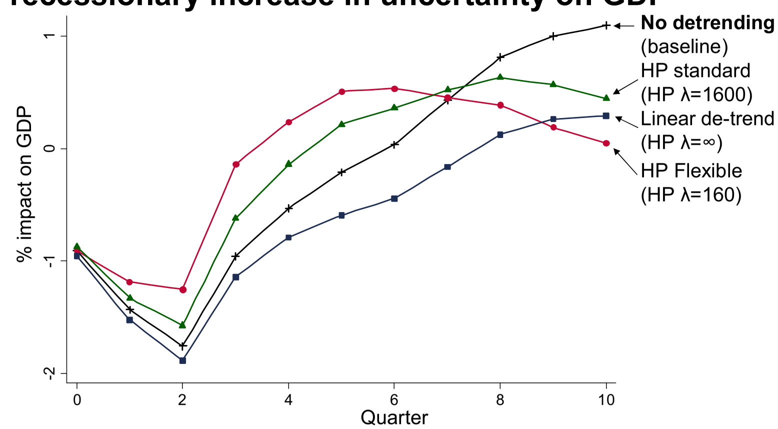

Notes: VAR Cholesky orthogonalized impulse responses to a 40% increase in the uncertainty index, calibrated to the average increase on the uncertainty index during a recession. Estimated on quarterly data from 1968:4 to 2006:2 using 4 lags. Variables (in order) are Variables (in order) are the uncertainty-index, log TFP (from Basu et al. 2006), the federal funds rate, log wages, log CPI, hours, log employment and log real GDP. Detrending as follows: (i) "HP Standard" by Hodrick-Prescott filter with a smoothing parameter λ of 1600; (ii) "HP λ=160" has an HP smoothing parameter of 160; (iii) "Linear de-trended" has an HP smoothing parameter of $\lambda=\infty$, and (iv) "No detrending" as in baseline

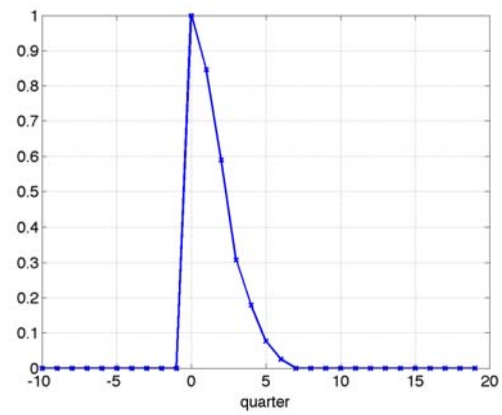

Figure 13: Simulation of an Uncertainty Shock. This figure shows the share of simulated economies that are in the high uncertainty state. In period zero, the uncertainty shock hits. The persistence of the uncertainty shock is governed by the Markov process presented in section 4.

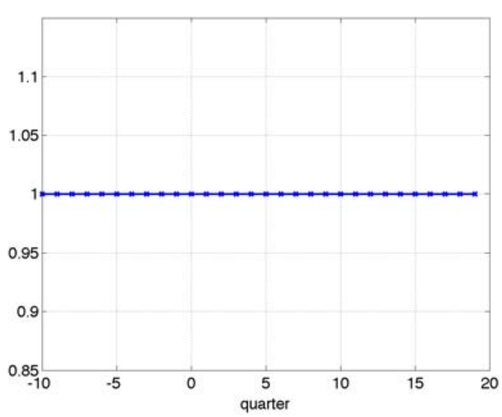

Figure 14: Aggregate Productivity During Uncertainty Shock. The figure shows aggregate productivity during the uncertainty shock. There is no first moment shock.

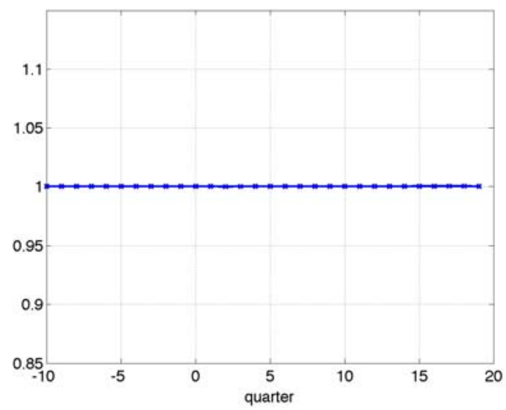

Figure 15: Average Idiosyncratic Productivity During Uncertainty Shock. Each firms idiosyncratic productivity follows the AR1 process with time-variation in the innovations variance as described in the main text. There is no first moment shock.

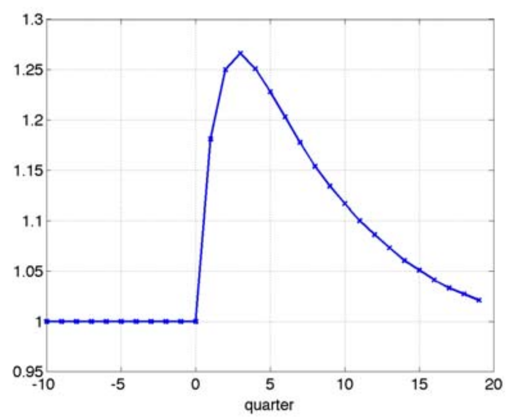

Figure 16: Standard Deviation of Firms' Idiosyncratic Productivities During Uncertainty Shock. This figure shows the standard deviation of the firms' idiosyncratic productivities. In period zero, the uncertainty shock hits. The persistence of the uncertainty shock is governed by the Markov process presented in section 4. The data is normalized to one in the period before impact.

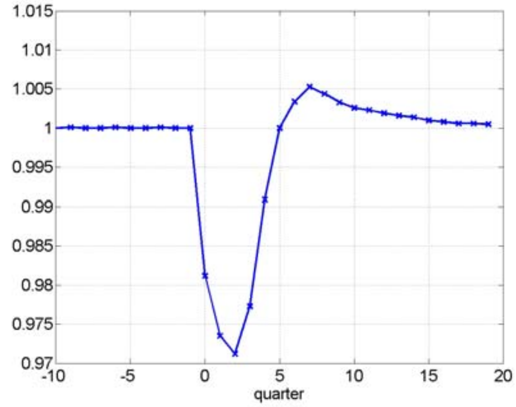

Figure 17: Aggregate Labor During Uncertainty Shock. This figure shows the average response of aggregate labor to an uncertainty shock that hits in quarter 0. The underlying data is taken from a large number of simulations where the length of the uncertainty shock follows the Markov process presented in section 4. Aggregate productivity is fixed at its mean. The series is normalized to 1 in the period before impact.

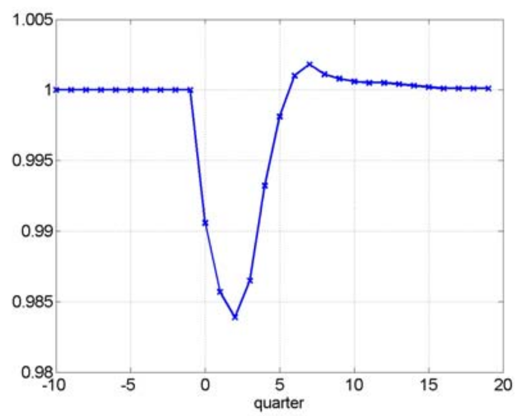

Figure 18: Aggregate Output During Uncertainty Shock. This figure shows the average response of aggregate output to an uncertainty shock that hits in quarter 0. The underlying data is taken from a large number of simulations where the length of the uncertainty shock follows the Markov process presented in section 4. Aggregate productivity is fixed at its mean. The series is normalized to 1 in the period before impact.

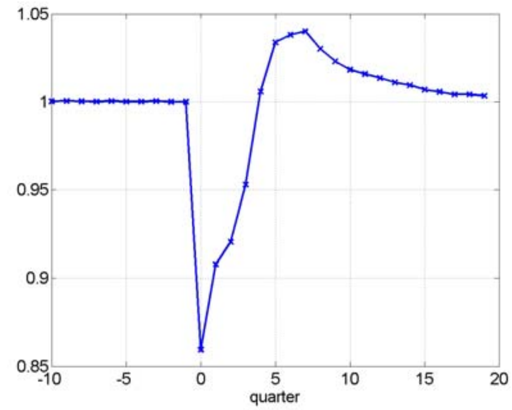

Figure 19: Aggregate Investment During Uncertainty Shock. This figure shows the average response of aggregate investment to an uncertainty shock that hits in quarter 0. The underlying data is taken from a large number of simulations where the length of the uncertainty shock follows the Markov process presented in section 4. Aggregate productivity is fixed at its mean. The series is normalized to 1 in the period before impact.

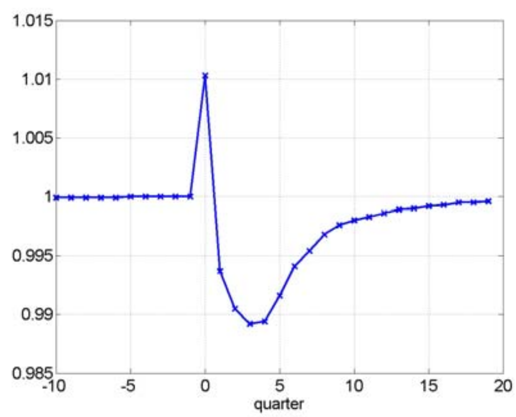

Figure 20: Consumption During Uncertainty Shock. This figure shows the average response of consumption to an uncertainty shock that hits in quarter 0. The underlying data is taken from a large number of simulations where the length of the uncertainty shock follows the Markov process presented in section 4. Aggregate productivity is fixed at its mean. The series is normalized to 1 in the period before impact.

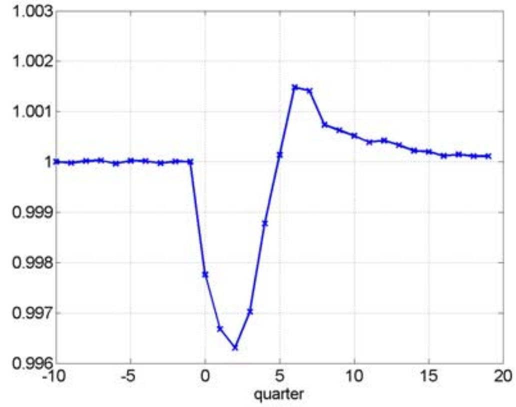

Figure 21: Aggregate Measured TFP During Uncertainty Shock. This figure shows the average response of the Solow Residual to an uncertainty shock that hits in quarter 0. The Solow Residual is calculated as $\frac{Y}{K^{\alpha}L^{\nu}}$. The underlying data is taken from a large number of simulations where the length of the uncertainty shock follows the Markov process presented in section 4. Aggregate productivity is fixed at its mean. The series is normalized to 1 in the period before impact.

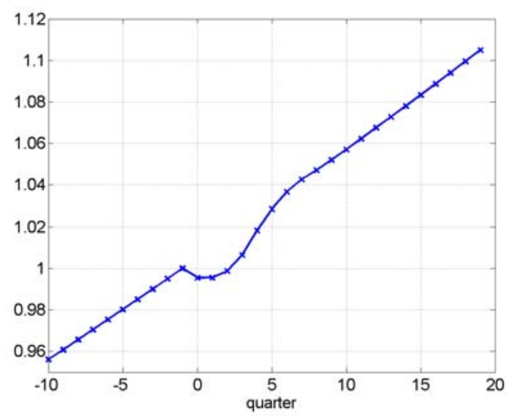

Figure 22: Uncertainty Shocks Can Create Recessions Without Negative Shocks to Aggregate TFP. This figure shows the average response of aggregate output to an uncertainty shock that hits in quarter 0. This graph adds the trend back to the IRF of aggregate output above. The underlying data is taken from a large number of simulations where the length of the uncertainty shock follows the Markov process presented in section 4. Aggregate productivity is fixed at its mean. The series is normalized to 1 in the period before impact.