NBER WORKING PAPER SERIES

THE IMPACT OF UNCERTAINTY SHOCKS

Nicholas Bloom

Working Paper 13385
http://www.nber.org/papers/w13385

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2007

This was the main chapter of my PhD thesis, previously called "The Impact of Uncertainty Shocks: A Firm-Level Estimation and a 9/11 Simulation". I would like to thank my advisors Richard Blundell and John Van Reenen; Costas Meghir and my referees; my formal discussants Susantu Basu, Russell Cooper, Janice Eberly, Valerie Ramey and Chris Sims; and seminar audiences at the AEA, Bank of England, Bank of Portugal, Berkeley, Board of Governors, Boston College, Boston Fed, Chicago, Chicago Fed, Chicago GSB, Cowles conference, Hoover, Kansas City Fed, Kansas University, Kellogg, LSE, MIT, NBER EF&G, CM&E and Productivity groups, Northwestern, QMW, San Francisco Fed, Stanford, UCL, UCLA and Yale. The financial support of the ESRC (Grant R000223644) is gratefully acknowledged. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

The Impact of Uncertainty Shocks
Nicholas Bloom
NBER Working Paper No. 13385
September 2007
JEL No. C23,D8,D92,E22

ABSTRACT

Uncertainty appears to jump up after major shocks like the Cuban Missile crisis, the assassination of JFK, the OPEC I oil-price shock and the 9/11 terrorist attack. This paper offers a structural framework to analyze the impact of these uncertainty shocks. I build a model with a time varying second moment, which is numerically solved and estimated using firm level data. The parameterized model is then used to simulate a macro uncertainty shock, which produces a rapid drop and rebound in aggregate output and employment. This occurs because higher uncertainty causes firms to temporarily pause their investment and hiring. Productivity growth also falls because this pause in activity freezes reallocation across units. In the medium term the increased volatility from the shock induces an overshoot in output, employment and productivity. Thus, second moment shocks generate short sharp recessions and recoveries. This simulated impact of an uncertainty shock is compared to VAR estimations on actual data, showing a good match in both magnitude and timing. The paper also jointly estimates labor and capital convex and non-convex adjustment costs. Ignoring capital adjustment costs is shown to lead to substantial bias while ignoring labor adjustment costs does not.

Nicholas Bloom
Stanford University
Department of Economics
579 Serra Mall
Stanford, CA 94305-6072
and NBER
nbloom@stanford.edu

1. Introduction

Uncertainty appears to dramatically increase after major economic and political shocks like the Cuban Missile crisis, the assassination of JFK, the OPEC I oil-price shock and the 9/11 terrorist attacks. Figure 1 plots stock market volatility - one proxy for uncertainty - which displays large bursts of uncertainty after major shocks, temporarily increasing (implied) volatility by up to 200%.[1] These volatility shocks are strongly correlated with other measures of uncertainty, like the cross-sectional spread of firm and industry level earnings and productivity growth. Vector Auto Regression (VAR) estimations suggest that they also have a large real impact, generating a substantial drop and rebound in output and employment over the following six months.

Uncertainty is also a ubiquitous concern of policymakers - for example after 9/11 the Federal Open Market Committee (FOMC) worried about exactly the type of real-options effects analyzed in this paper, stating in October 2001 that "*the events of September 11 produced a marked increase in uncertainty...depressing investment by fostering an increasingly widespread wait-and-see attitude*".

But despite the size and regularity of these second moment (uncertainty) shocks there is no general structural model that analyzes their effects. This is surprising given the extensive literature on the impact of first moment (levels) shocks. This leaves open a wide variety of questions on the impact of major macroeconomic shocks, since these typically have both a first and second moment component.

The primary contribution of this paper is a structural framework to analyze these types of uncertainty shocks, building a model with a time varying second moment of the driving process and a mix of labor and capital adjustment costs. The model is numerically solved and estimated on firm level data using simulated method of moments. Firm-level data helps to overcomes the identification problem associated with the limited sample size of macro data. Cross-sectional and temporal aggregation are incorporated to enable the estimation of structural parameters.

With this parameterized model I then simulate the impact of a large temporary uncertainty shock and find that it generates a rapid drop, rebound and overshoot in employment, output and productivity growth. Hiring and investment rates fall dramatically in the four months after the shock because higher uncertainty increases the real option value to waiting, so firms scale back their plans. But once uncertainty has subsided, activity quickly bounces back as firms address their pent-up demand for labor and capital. Aggregate productivity growth also falls dramatically after the

[1] In financial markets implied share-returns volatility is the canonical measure for uncertainty. Bloom, Bond and Van Reenen (2007) show that firm-level share-returns volatility is significantly correlated with a range of alternative uncertainty proxies, including real sales growth volatility and the cross-sectional distribution of financial analysts' forecasts. While Shiller (1981) has argued that the *level* of stock price volatility is excessively high, Figure 1 suggests that *changes* in stock-price volatility are nevertheless linked with real and financial shocks.

Figure 1: Monthly US stock market volatility

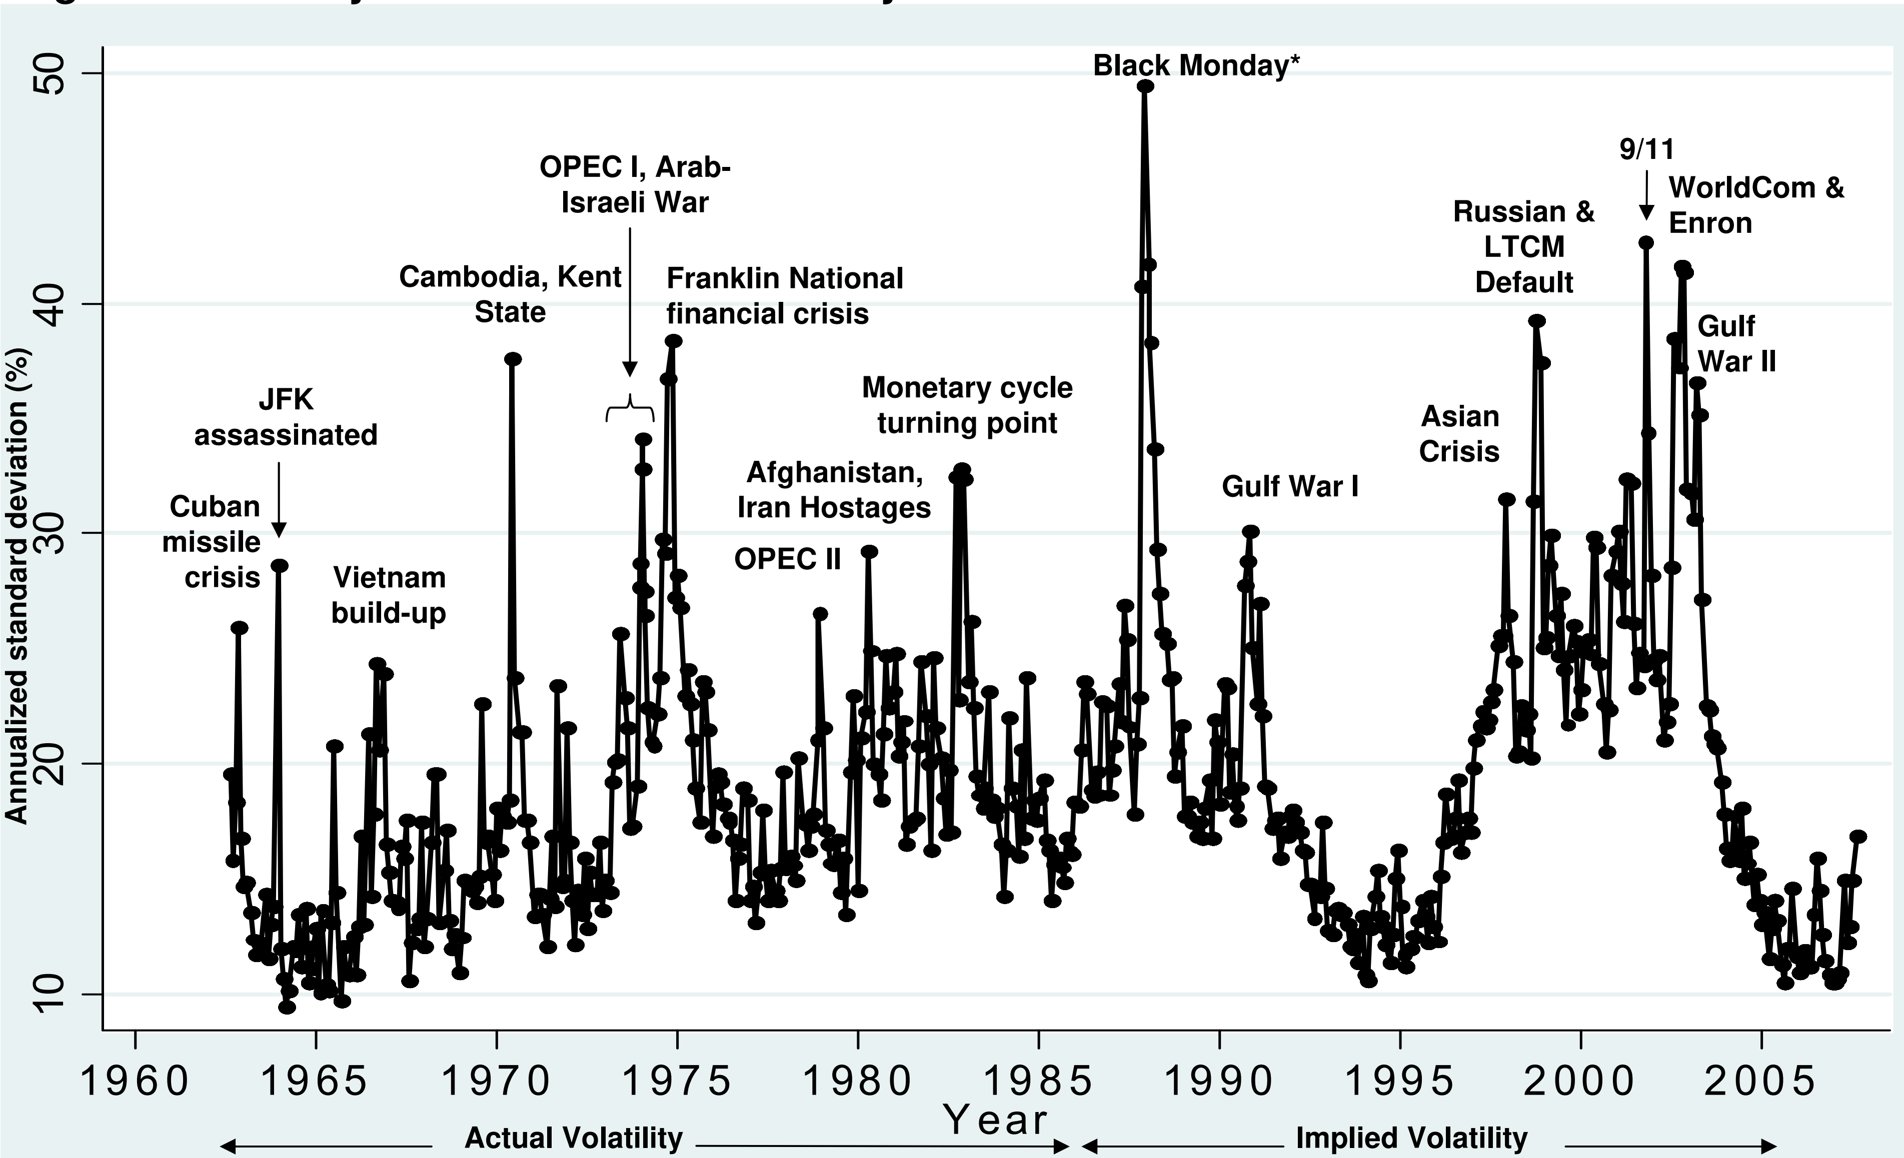

Notes: CBOE VXO index of % implied volatility, on a hypothetical at the money S&P100 option 30 days to expiration, from 1986 to 2007. Pre 1986 the VXO index is unavailable, so actual monthly returns volatilities calculated as the monthly standard-deviation of the daily S&P500 index normalized to the same mean and variance as the VXO index when they overlap (1986-2006). Actual and VXO are correlated at 0.874 over this period. The market was closed for 4 days after 9/11, with implied volatility levels for these 4 days interpolated using the European VX1 index, generating an average volatility of 58.2 for 9/11 until 9/14 inclusive. A brief description of the nature and exact timing of every shock is contained in Appendix A. Shocks defined as events 1.65 standard deviations about the Hodrick-Prescott detrended (λ=129,600) mean, with 1.65 chosen as the 5% significance level for a one-tailed test treating each month as an independent observation.

* For scaling purposes the monthly VXO was capped at 50 for the Black Monday month. Un-capped value for the Black Monday month is 58.2.

shock because the drop in hiring and investment reduces the rate of re-allocation from low to high productivity firms, which drives the majority of productivity growth in the model as in the real economy.[2] But again productivity growth rapidly bounces back as pent-up re-allocation occurs.

In the medium term the increased volatility arising from the uncertainty shock generates a 'volatility-overshoot'. The reason is that most firms are located near their hiring and investment thresholds, above which they hire/invest and below which they have a zone of inaction. So small positive shocks generate a hiring and investment response while small negative shocks generate no response. Hence, hiring and investment are locally convex in business conditions (demand and productivity). The increased volatility of business conditions growth after a second-moment shock therefore leads to a medium-term rise in labor and capital.

In sum, these second moment effects generate a rapid slow-down and bounce-back in economic activity, entirely consistent with the empirical evidence. This is very different from the much more persistent slowdown that typically occurs in response to the type of first moment productivity and/or demand shock that is usually modelled in the literature.[3] This highlights the importance to policymakers of distinguishing between the persistent first moment effects and the temporary second moment effects of major shocks.

I then evaluate the robustness of these predictions to general equilibrium effects, which for computational reasons are not included in my baseline model. To investigate this I build the falls in interest rates, prices and wages that occur after actual uncertainty shocks into the simulation. This has little short-run effect on the simulations, suggesting that the results are robust to general equilibrium effects. The reason is that the rise in uncertainty following a second moment shock not only generates a slowdown in activity, but it also makes firms temporarily extremely insensitive to price changes. This raises a second policy implication that the economy will be particularly unresponsive to monetary or fiscal policy immediately after an uncertainty shock, suggesting additional caution when thinking about the policy response to these types of events.

The analysis of uncertainty shocks links with the earlier work of Bernanke (1983) and Hassler (1996) who highlight the importance of variations in uncertainty.[4] In this paper I quantify and substantially extend their predictions through two major advances: first by introducing uncertainty as a stochastic process which is critical for evaluating the high frequency impact of major shocks;

[2] See Foster, Haltiwanger and Krizan (2000 and 2006).

[3] See, for example, Christiano, Eichenbaum and Evans (2005) and the references therein.

[4] Bernanke develops an example of uncertainty in an oil cartel for capital investment, while Hassler solves a model with time-varying uncertainty and fixed adjustment costs. There are of course many other linked recent strands of literature, including work on growth and volatility such as Ramey and Ramey (1995) and Aghion et al. (2005), on investment and uncertainty such as Leahy and Whited (1996) and Bloom, Bond and Van Reenen (2007), on the business-cycle and uncertainty such as Barlevy (2004) and Gilchrist and Williams (2005), on policy uncertainty such as Adda and Cooper (2000) and on income and consumption uncertainty such as Meghir and Pistaferri (2004).

and second by considering a joint mix of labor and capital adjustment costs which is critical for understanding the dynamics of employment, investment and productivity.

The secondary contribution of this paper is to analyze the importance of jointly modelling labor and capital adjustment costs. For analytical tractability and aggregation constraints the empirical literature has either estimated labor or capital adjustment costs individually assuming the other factor is flexible, or estimated them jointly assuming only convex adjustment costs.[5] I jointly estimate a mix of labor and capital adjustment costs by exploiting the properties of homogeneous functions to reduce the state space, and develop an approach to address cross-sectional and temporal aggregation. I find moderate non-convex labor adjustment costs and substantial non-convex capital adjustment costs. I also find that assuming capital adjustment costs only - as is standard in the investment literature - generates an acceptable overall fit, while assuming labor adjustment costs only - as is standard in the labor demand literature - produces a poor fit.

This framework also suggests a range of future research. Looking at individual events it could be used, for example, to analyze the uncertainty impact of major deregulations, tax changes, trade reforms or political elections. It also suggests there is a trade-off between policy "correctness" and "decisiveness" - it may be better to act decisively (but occasionally incorrectly) then to deliberate on policy, generating policy-induced uncertainty. More generally, the framework in this paper also provides one response to the "*where are the negative productivity shocks?*" critique of real business cycle theories.[6] In particular, since second moment shocks generate large falls in output, employment and productivity growth, it provides an alternative mechanism to first-moment shocks for generating recessions. Recessions could simply be periods of high uncertainty without negative productivity shocks. Encouragingly, recessions do indeed appear in periods of significantly higher uncertainty, suggesting an uncertainty approach to modelling business-cycles (see Bloom, Floetotto and Jaimovich, 2007). Taking a longer run perspective this paper also links to the volatility and growth literature, given the large negative impact of uncertainty on output and productivity growth.

The rest of the paper is organized as follows: in section (2) I empirically investigate the importance of jumps in stock-market volatility, in section (3) I set up and solve my model of the firm, in section (4) I characterize the solution of the model, in section (5) I outline my simulated method of moments estimation approach, in section (6) I report the parameters estimates using US firm data, in section (7) I take my parameterized model and simulate the high frequency effects of a large uncertainty

[5] See, for example; on capital Cooper and Haltiwanger (1993), Caballero, Engel and Haltiwanger (1995), Cooper, Haltiwanger and Power (1999) and Cooper and Haltiwanger (2003); on labor Hammermesh (1989), Bertola and Bentolila (1990), Davis and Haltiwanger (1992), Caballero and Engel (1993), Caballero, Engel and Haltiwanger (1997) and Cooper, Haltiwanger and Willis (2004); on joint estimation with convex adjustment costs Shapiro (1986), Hall (2004) and Merz and Yashiv (2005); and Bond and Van Reenen (2007) for a full survey of the literature.

[6] See the extensive discussion in King and Rebello (1999).

shock. Finally, section (8) offers some concluding remarks.

2. Do Jumps in Stock-Market Volatility Matter?

Two key questions to address before introducing any models of uncertainty shocks are: (i) do jumps in the volatility index in Figure 1 represents uncertainty shocks,[7] and (ii) do these have any impact on real economic outcomes? In section (2.1) I address the first question by presenting evidence showing that stock market volatility is strongly linked to other measures of productivity and demand uncertainty. In section (2.2) I address the second question by presenting Vector Auto Regression (VAR) estimations showing that volatility shocks generate a short-run drop of 1%, lasting about 6 months, with a longer run gradual overshooting. First moment shocks to the interest-rates and stock-market levels generate a much more gradual drop and rebound in activity lasting 2 to 3 years. A full data description for both sections is contained in Appendix A.[8]

2.1. Empirical Evidence on the Links Between Stock-Market Volatility and Uncertainty

The evidence presented in Table 1 shows that a number of *cross-sectional* measures of uncertainty are highly correlated with *time-series* stock-market volatility. Stock market volatility has also been previously used as a proxy for uncertainty at the firm level (e.g. Leahy and Whited (1996) and Bloom, Bond and Van Reenen. (2007)).

Columns (1) and (2) of Table 1 use the cross-sectional standard deviation of firms' pre-tax profit growth, taken from the quarterly accounts of public companies. As can be seen from column (1) stock-market time-series volatility is strongly correlated with the cross-sectional spread of firm-level profits growth. All variables in Table 1 have been normalized by their standard deviations (SD). The coefficient implies that the 2.47 SD rise in stock-market time-series volatility that occurred on average after the shocks highlighted in Figure 1 would be associated with a 1.31 SD rise in the cross-sectional spread of the growth rate of profits, a large increase. Column (2) re-estimates this including a full set of quarterly dummies and a time-trend, finding very similar results.[9]

Columns (3) and (4) use a monthly cross-sectional stock-returns measure and show that this is also strongly correlated with the stock-return volatility index. Columns (5) and (6) report the results from using the standard-deviation of annual 5-factor TFP growth within the NBER manufacturing industry database. There is also a large and significant correlation of the cross-sectional spread of industry productivity growth and stock-market volatility. Finally, Columns (7) and (8) use a

[7] I tested for jumps in the volatility series using the bipower variation test of Barndorff-Nielsen and Shephard (2006) and found statistically significance evidence for jumps. Full details in Appendix A1.

[8] All data and program files are also available at http://www.stanford.edu/~nbloom/

[9] This helps to control for any secular changes in volatility (see Davis et al. (2006)).

Table 1: The stock-market volatility index regressed on cross-sectional measures of uncertainty

Dependent variable is: **Stock market volatility**	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Explanatory variable is the period by period cross-sectional standard deviation of:								
Firm profit growth, Compustat quarterly	0.532	0.526						
	(0.064)	*(0.092)*						
Firm stock returns, CRSP monthly			0.537	0.528				
			(0.037)	*(0.038)*				
Industry TFP growth, 4-digit SIC annual					0.425	0.414		
					(0.118)	*(0.124)*		
GDP forecasts, Livingstone half-yearly							0.615	0.580
							(0.112)	*(0.121)*
Time trend	No	Yes	No	Yes	No	Yes	No	Yes
Month/quarter/half-year dummies	No	Yes	No	Yes	n/a	n/a	No	Yes
R^2	0.287	0.301	0.285	0.345	0.282	0.284	0.332	0.381
Time span	62Q1-05Q1		62M5-06M3		1962-1996		62H2-98H2	
Ave. units in cross-section	327		357		425		57.4	
Observations in regression	171		526		35		63	

Notes: Each column reports the coefficient from regressing the time-series of stock market volatility on the within period cross-sectional standard deviation (SD) of the explanatory variable calculated from an underlying panel. All variables are normalized to have a standard-deviation (SD) of one. Standard errors in italics in ( ) below the point estimate. So, for example, column (1) reports that the stock market volatility index is on average 0.532 SD higher in a quarter when the cross-sectional spread of firms' profit growth is 1 SD higher. The "Stock market volatility index" measures monthly volatility on the US stock market, and is plotted in Figure 1. The quarterly, half-yearly and annual values are calculated by averaging across the months within the period. The standard deviation of "Firm profits growth" measures the within-quarter cross-sectional spread of profit growth rates normalized by average sales, defined as $(\text{profits}_t - \text{profits}_{t-1})/(0.5\times\text{sales}_t + 0.5\times\text{sales}_{t-1})$. This comes from Compustat quarterly accounts using firms with 150+ quarters of accounts. The standard deviation of "Firm stock returns" measures the within month cross-sectional standard deviation of firm-level stock returns. This comes from the CRSP monthly stock-returns file using firms with 500+ months of accounts. The standard deviation of "Industry TFP growth" measures the within year cross-industry spread of SIC 4-digit manufacturing TFP growth rates. This is calculated using the 5-factor TFP growth figures from the NBER manufacturing industry database. The standard deviation of "GDP forecasts" comes from the Philadelphia Federal Reserve Bank's biannual Livingstone survey, calculated as the (standard-deviation/mean) of forecasts of nominal GDP one year ahead, using only half-years with 50+ forecasts. This series is linearly detrended to remove a long-run downward drift. "Ave. units in cross-section" refers to the *average* number of units (firms, industries or forecasters) used to measure the cross-sectional spread. "Month/quarter/half-year dummies" refers to quarter, month and half controls in columns (2), (4) and (8) respectively. A full description of the variables is contained in Appendix A.

measure of the dispersion across macro forecasters over their predictions for future GDP, calculated from the Livingstone half-yearly survey of professional forecasters. Once again, periods of high stock-market volatility are significantly correlated with cross-sectional dispersion, in this case in terms of disagreement across macro forecasters.

2.2. VAR Estimates on the Impact of Stock-Market Volatility Shocks

To evaluate the impact of uncertainty shocks on real economic outcomes I estimate a range of VARs on monthly data from July 1963 to July 2005.[10] The variables in the estimation order are log(industrial production), log(employment), hours, log(consumer price index), log(average hourly earnings), Federal Funds Rate, a stock-market volatility indicator (described below) and log(S&P500 stock market index). This ordering is based on the assumptions that shocks instantaneously influence the stock market (levels and volatility), then prices (wages, the CPI and interest rates) and finally quantities (hours, employment and output). Including the stock market levels as the first variable in the VAR ensures the impact of stock-market levels is already controlled for when looking at the impact of volatility shocks. All variables are Hodrick Prescott (HP) detrended ($\lambda = 129,600$) in the baseline estimations.

The main stock-market volatility indicator is constructed to take a value 1 for each of the shocks labelled in Figure 1 and a 0 otherwise. These sixteen shocks were explicitly chosen as those events when the peak of HP detrended volatility level rose significantly above the mean.[11] This indicator function is used to ensure identification comes only from these large, and arguably exogenous, volatility shocks rather than the smaller ongoing fluctuations.

Figure 2 plots the impulse response function of industrial production (the solid line with plus symbols) to a volatility shock. Industrial production displays a rapid fall of around 1% within four months, with a subsequent recovery and rebound from seven months after the shock. The one standard-error bands (dashed lines) are plotted around this, highlighting that this drop and rebound is statistically significant at the 5% level. For comparison to a first moment shock, the response to a 1% impulse to the Federal Funds Rate (FFR) is also plotted (solid line with circular symbols) displaying a much more persistent drop and recovery of up to 0.6% over the subsequent two years.[12] In Figure 3 the response of employment to a stock-market volatility shock is also plotted, displaying a similar large drop and recovery in activity. Figures A1, A2 and A3 in the Appendix confirm the

[10] I would like to thank Valerie Ramey and Chris Sims (my discussants at the NBER EF&G and Evora conferences) for their initial VAR estimations and subsequent discussions.

[11] The threshold was 1.65 standard deviations above the mean, selected as the 5% one-tailed significance level treating each month as an independent observation. The VAR estimation also uses the full volatility series (which does not require defining shocks) and finds very similar results, as shown in Figure A1.

[12] The response to a 5% fall in the level of the stock-market levels (not plotted) is very similar in size and magnitude to the response to a 1% rise in the FFR.

Figure 2: VAR estimation of the impact of a volatility shock on industrial production

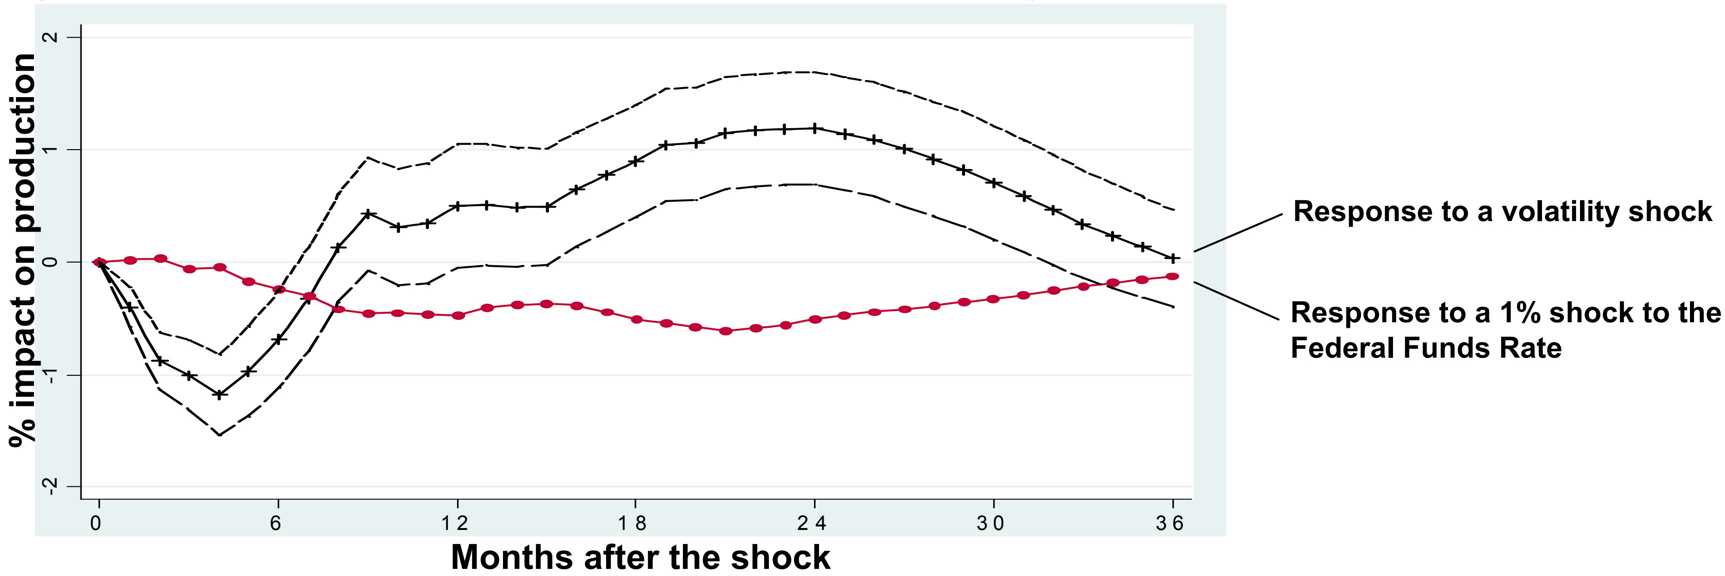

Figure 3: VAR estimation of the impact of a volatility shock on employment

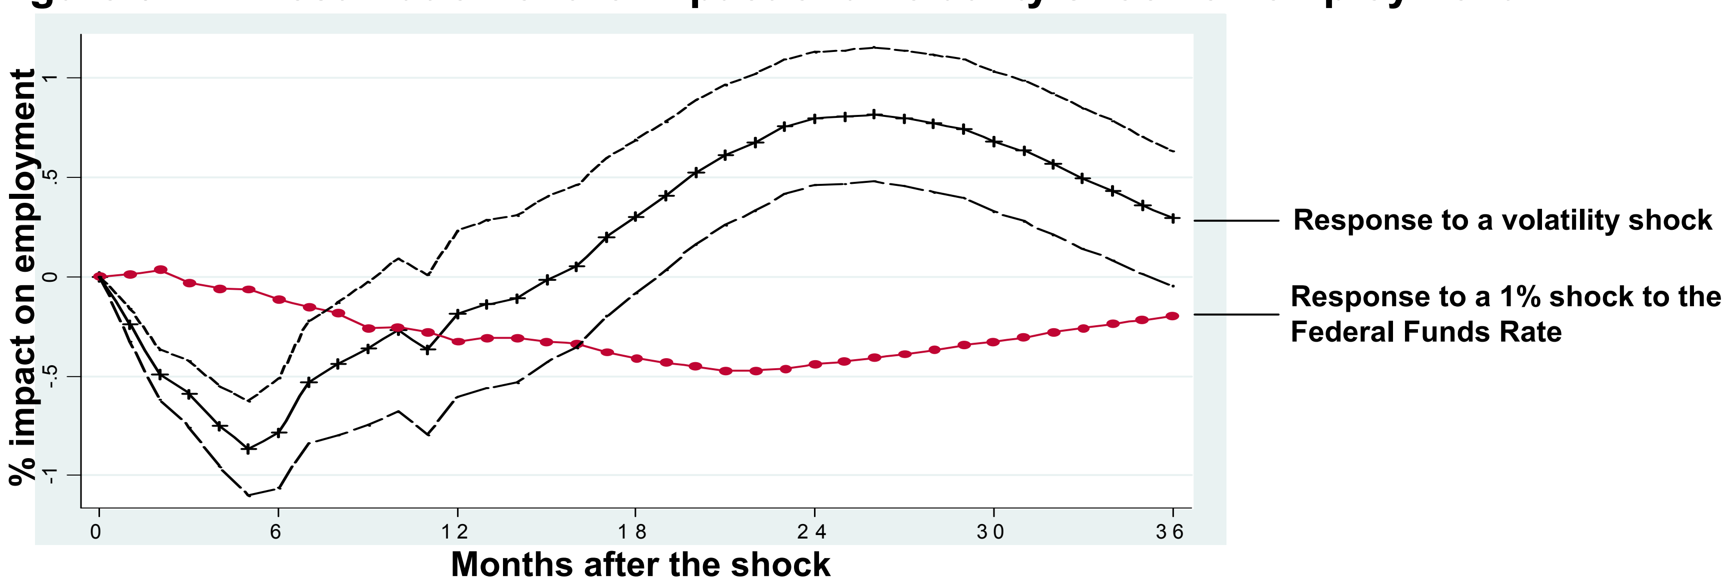

Notes: VAR Cholesky orthogonalized impulse response functions estimated on monthly data from July 1963 to July 2005 using 12 lags. Dotted lines in top and bottom figures are one standard error bands around the response to a volatility shock indicator, coded as a 1 for each of the 16 labelled shocks in Figure 1, and 0 otherwise. Variables (in order) are log industrial production, log employment, hours, log wages, log CPI, federal funds rate, the volatility shock indicator and log S&P500 levels. Detrending by Hodrick-Prescott filter with smoothing parameter of 129,600. The response to a 1% shock to the Federal Funds Rate (dotted line) is plotted to demonstrate the time profile in response to a typical first moment shock.

robustness of these VAR results to a range of alternative approaches over variable ordering, variable inclusion, shock definitions, shock timing and detrending. In particular, these results are robust to identification from uncertainty shocks defined by the 10 exogenous shocks arising from wars, OPEC shocks and terror events.[13]

3. Modelling the Impact of an Uncertainty Shock

In this section I model the impact of an uncertainty shock. I take a standard model of the firm[14] and extend this in two ways. First, I introduce uncertainty as a stochastic process to evaluate the impact of the uncertainty shocks shown in Figure 1. Second, I allow a joint mix of convex and non-convex adjustment costs for both labor and capital. The time varying uncertainty interacts with the non-convex adjustment costs to generate time-varying real-option effects, which drive fluctuations in hiring and investment. I also build in temporal and cross-sectional aggregation by assuming firms own large numbers of production units, which allows me to estimate the model's parameters on firm-level data.

3.1. The Production and Revenue Function

Each production unit has a Cobb-Douglas[15] production function

$$F(\widetilde{A}, K, L, H) = \widetilde{A} K^{\alpha} (LH)^{1-\alpha} \tag{3.1}$$

in productivity ($\widetilde{A}$), capital (K), labor (L) and hours (H). The firm faces an iso-elastic demand curve with elasticity (ϵ)

$$Q = BP^{-\epsilon}, \tag{3.2}$$

where B is a (potentially stochastic) demand shifter. These can be combined into a revenue function $R(\widetilde{A}, B, K, L, H) = \widetilde{A}^{1-1/\epsilon} B^{1/\epsilon} K^{\alpha(1-1/\epsilon)} (LH)^{(1-\alpha)(1-1/\epsilon)}$. For analytical tractability I define $a = \alpha(1 - 1/\epsilon)$, $b = (1 - \alpha)(1 - 1/\epsilon)$ and substitute $A^{1-a-b} = \widetilde{A}^{1-1/\epsilon} B^{1/\epsilon}$, where A combines the unit level productivity and demand terms into one index, which for expositional simplicity I will refer to as 'business conditions'. With these redefinitions we have[16]

$$S(A, K, L, H) = A^{1-a-b} K^{a} (LH)^{b}. \tag{3.3}$$

[13] In an earlier version of the paper (Bloom, 2006) I evaluated the impact of one particular uncertainty shock - the 9/11 terrorist attack - against consensus forecasts made two weeks before the attack. I showed that 9/11 appeared to generate a large drop and rapid rebound in hiring and investment lasting around 6 months.

[14] See, for example, Bertola and Caballero (1994), Abel and Eberly (1996) or Caballero and Engel (1999).

[15] While I assume a Cobb-Douglas production function other supermodular homogeneous unit revenue functions could be used. For example, by replacing (3.1) with a CES aggregator over capital and labor where $F(\widetilde{A}, K, L, H) = \widetilde{A}(\alpha_1 K^{\sigma} + \alpha_2 (LH)^{\sigma})^{\frac{1}{\sigma}}$ I obtained similar simulation results.

[16] This reformulation to A as the stochastic variable to yield a jointly homogeneous revenue function avoids any long-run Hartman (1972) or Abel (1983) effects of uncertainty reducing or increasing output because of convexity or concavity in the production function. See Caballero (1991) or Abel and Eberly (1996) for a more detailed discussion.

Wages are determined by undertime and overtime hours around the standard working week of 40 hours, which following the approach in Caballero and Engel (1993), is parameterized as $w(H) = w_1(1+w_2H^{\gamma})$, where w_1, w_2 and γ are parameters of the wage equation to be determined empirically.

3.2. The Stochastic Process for Demand and Productivity

I assume business conditions evolve as an augmented geometric random walk. Uncertainty shocks are modelled as time variations in the standard deviation of the driving process, consistent with the stochastic volatility measure of uncertainty in Figure 1.

Business conditions are in fact modelled as a multiplicative composite of three separate random-walks[17], a macro-level component (A_t^M), a firm-level component ($A_{i,t}^F$) and a unit-level component ($A_{i,j,t}^U$), where $A_{i,j,t} = A_t^M A_{i,t}^F A_{i,j,t}^U$ and i indexes firms, j indexes units and t indexes time. The macro level component is modelled as follows:

$$A_t^M = A_{t-1}^M(1+\sigma_{t-1}W_t^M) \qquad W_t^M \sim N(0,1), \tag{3.4}$$

where σ_t is the standard-deviation of business conditions and W_t^M is a macro-level i.i.d. normal shock. The firm level component is modelled as follows:

$$A_{i,t}^F = A_{i,t-1}^F(1+\mu_{i,t}+\sigma_{t-1}W_{i,t}^F) \qquad W_{i,t}^F \sim N(0,1), \tag{3.5}$$

where $\mu_{i,t}$ is a firm-level drift in business conditions and $W_{i,t}^F$ is a firm-level i.i.d. normal shock. The unit level component is modelled as follows:

$$A_{i,j,t}^U = A_{i,j,t-1}^U(1+\sigma_{t-1}W_{i,j,t}^U) \qquad W_{i,j,t}^U \sim N(0,1), \tag{3.6}$$

where $W_{i,j,t}^U$ is a unit-level i.i.d. normal shock. I assume W_t^M, W_t^F and $W_{i,t}^U$ are all independent of each other.

While this demand structure may seem complex, it is formulated to ensure that: (i) units within the same firm have linked investment behavior due to common firm-level business conditions and uncertainty shocks; and (ii) they display some independent behavior due to the idiosyncratic unit level shocks, which is essential for smoothing under aggregation. This demand structure also assumes that macro, firm and unit level uncertainty are the same. This is broadly consistent with the results from Table 1 for firm and macro uncertainty, which show these are highly correlated. For unit level

[17] A random-walk driving process is assumed for analytical tractability, in that it helps to deliver a homogenous value function (details in the next section). It is also consistent with Gibrat's law. An equally plausible alternative assumption would be a persistent AR(1) process, such as the following based on Cooper and Haltiwanger (2006): $\log(A_t) = \alpha + \rho\log(A_{t-1}) + v_t$ where $v_t \sim N(0,\sigma_{t-1})$, $\rho = 0.885$. To investigate this alternative I programmed up another monthly simulation with auto-regressive business conditions and no labor adjustment costs (so I could drop the labor state) and all other modelling assumptions the same. I found in this set-up there were still large real-options effects of uncertainty shocks on output, as plotted in Appendix Figure A4.

uncertainty there is no direct evidence on this. But to the extent this assumption does not hold - so that unit and macro uncertainty are imperfectly correlated - this will weaken the quantitative impact of macro uncertainty shocks (since total uncertainty will fluctuate less than one-for-one with macro uncertainty), but not the qualitative findings. The firm-level business conditions drift ($\mu_{i,t}$) is also assumed to be stochastic, to allow autocorrelated changes over time within firms. This is important for empirically identifying adjustment costs from persistent differences in growth rates across firms, as section (5) discusses in more detail.[18]

The stochastic volatility process (σ_t^2) and the demand conditions drift ($\mu_{i,t}$) are both assumed for simplicity to follow two point Markov Chains

$$\sigma_t \in \{\sigma_L, \sigma_H\} \quad \text{where } Pr(\sigma_{t+1} = \sigma_j | \sigma_t = \sigma_k) = \pi^\sigma_{k,j} \tag{3.7}$$

$$\mu_{i,t} \in \{\mu_L, \mu_H\} \quad \text{where } Pr(\mu_{i,t+1} = \mu_j | \mu_{i,t} = \mu_k) = \pi^\mu_{k,j}. \tag{3.8}$$

3.3. Adjustment Costs

The third piece of technology determining the firms' activities are the adjustment costs. There is a large literature on investment and employment adjustment costs which typically focuses on three terms, all of which I include in my specification:

Partial irreversibilities: Labor partial irreversibility, labelled C_L^P, derives from *per capita* hiring training and firing costs, and is denominated as a fraction of annual wages (at the standard working week). For simplicity I assume these costs apply equally to gross hiring and gross firing of workers.[19] Capital partial irreversibilities arise from resale losses due to transactions costs, the market for lemons phenomenon and the physical costs of resale. The resale loss of capital is labelled C_K^P and is denominated as a fraction of the relative purchase price of capital.

Fixed disruption costs: When new workers are added into the production process and new capital is installed some downtime may result, involving a fixed cost loss of output. For example, adding workers may require fixed costs of advertising, interviewing and training, or the factory may need to close for a few days while a capital refit is occurring. I model these fixed costs as C_L^F and C_K^F for hiring/firing and investment respectively, both denominated as fractions of annual sales.

[18] This formulation also generates 'business conditions' shocks at the unit-level that have a $\sqrt{3}$ times larger standard-deviation than at the macro level. This appears to be counter empirical given the much higher volatility of establishment data than macro data. However, because of the non-linearities in the investment and hiring response functions (due to non-convex adjustment costs) output and input growth is typically around 10 times more volatile at the *unit* level then at the smoothed (by aggregation) *macro* level in the simulation.

Furthermore, all that matters for the simulation results in section (7.1) is the change in the total variance of shocks to $A_{i,j,t}$, rather than the breakdown of this variance between macro, firm and unit level shocks.

[19] Microdata evidence, for example Davis and Haltiwanger (1992), suggests both *gross* and *net* hiring/firing costs may be present. For analytical simplicity I have restricted the model to *gross* costs, noting that *net* costs could also be introduced and estimated in future research through the addition of two *net* firing cost parameters.

Quadratic adjustment costs: The costs of hiring/firing and investment may also be related to the rate of adjustment due to higher costs for more rapid changes, where $C_L^Q L(\frac{E}{L})^2$ are the quadratic hiring/firing costs and E denotes gross hiring/firing, and $C_K^Q K(\frac{I}{K})^2$ are the quadratic investment costs and I denotes gross investment.

The combination of all adjustment costs is given by the adjustment cost function:

$$\begin{aligned} C(A,K,L,H,I,E,p_t^K) &= 52w(40)C_L^P(E^+ + E^-) + (I^+ - (1-C_K^P)I^-) + \\ &\quad \left(C_L^F 1_{\{E\neq 0\}} + C_K^F 1_{\{I\neq 0\}}\right) S(A,K,L,H) + C_L^Q L(\frac{E}{L})^2 + C_K^Q K(\frac{I}{K})^2 \end{aligned}$$

where E^+ (I^+) and E^- (I^-) are the absolute values of positive and negative hiring (investment) respectively, and $1_{\{E\neq 0\}}$ and $1_{\{I\neq 0\}}$ are indicator functions which equal 1 if true and 0 otherwise. New labor and capital take one period to enter production due to time to build. This assumption is made to allow me to pre-optimize hours (explained in section (3.5) below), but is unlikely to play a major role in the simulations given the monthly periodicity. At the end of each period there is labor attrition and capital depreciation proportionate to δ_L and δ_K respectively.

3.4. Dealing with Cross-Sectional and Time Aggregation

Gross hiring and investment is typically lumpy with frequent zeros in single-plant establishment level data but much smoother and continuous in multi-plant establishment and firm level data. This appears to be because of extensive aggregation across two dimensions: cross sectional aggregation across types of capital and production plants; and temporal aggregation across higher-frequency periods within each year (see Appendix section A4). I build this aggregation into the model by explicitly assuming that firms own a large number of production *units* and that these operate at a higher frequency than yearly. The units can be thought of as different production plants, different geographic or product markets, or different divisions within the same firm.

To solve this model I need to define the relationship between production units within the firm. This requires several simplifying assumptions to ensure analytical tractability. These are not attractive, but are necessary to enable me to derive numerical results and incorporate aggregation into the model. In doing this I follow the general stochastic aggregation approach of Bertola and Caballero (1994) and Caballero and Engel (1999) in modelling macro and industry investment respectively, and most specifically Abel and Eberly (2002) in modelling firm level investment.

The stochastic aggregation approach assumes firms own a sufficiently large number of production units that any single unit level shock has no significant impact on firm behavior. Units are assumed to independently optimize to determine investment and employment. Thus, all linkages across units within the same firm are modelled by the common shocks to demand, uncertainty or the price of capital. So, to the extent that units are linked over and above these common shocks the implicit assumption is that they independently optimize due to bounded rationality and/or localized incentive mechanisms (i.e. managers being assessed only on their own unit's Profit and Loss account).[20]

In the simulation the number of units per firm is set at 250, chosen by increasing the number of units until the results were no longer sensitive to this number.[21] This assumption will have a direct effect on the estimated adjustment costs (since aggregation and adjustment costs are both sources of smoothing) and thereby an indirect effect on the simulation. Hence, in section (5) I re-estimate the adjustment costs assuming instead the firm has 1 and 50 units to investigate this further.

The model also assumes no entry or exit for analytical tractability. This seems acceptable in the monthly time frame (entry/exit accounts for around 2% of employment on an annual basis), but is an important assumption to explore in future research. My intuition is that relaxing this assumption should increase the effect of uncertainty shocks since entry and exit decisions are extremely non-convex, although this may have some offsetting effects through the estimation of slightly "smoother" adjustment costs.

There is also the issue of time series aggregation. Shocks and decisions in a typical business-unit are likely to occur at a much higher frequency than annually, so annual data will be temporally aggregated, and I need to explicitly model this. There is little information on the frequency of decision making in firms, with the available evidence suggesting monthly frequencies are typical (due to the need for senior managers to schedule regular meetings), which I assume in my main results.

3.5. Optimal Investment and Employment

The firm's optimization problem is to maximize the present discounted flow of revenues less the wage bill and adjustment costs across its units. I assume that the *firm* is risk neutral to focus on the real options effects of uncertainty.[22]

Analytical methods can show that a unique solution to the firm's optimization problem exists, that is continuous and strictly increasing in (A, K, L) with an almost everywhere unique policy function.[23] The model is too complex, however, to be fully solved using analytical methods, so I use

[20] The semi-independent operation of plants may be theoretically optimal for incentive reasons (to motivate local managers) and technical reasons (the complexity of centralized information gathering and processing). The empirical evidence on decentralization in US firms suggests that plant-managers have substantial hiring and investment discretion (see for example Bloom and Van Reenen, 2007).

[21] The firms in my estimation sample have a mean (median) size of 13,540 (3,450) employees (see section 5.4) so this implies each unit has 54 (14) employees at the mean (median).

[22] In an earlier version of the paper, Bloom (2006), I provided (partial equilibrium) simulation results for *firm* risk-aversion. Including this reinforces the real-options effects because it induces a rise in the investment hurdle-rate after the uncertainty shock hits which then falls back as certainty returns. In general equilibrium these effects become ambiguous because of offsetting consumer risk-aversion effects.

[23] The application of Stokey and Lucas (1989) for the continuous, concave and almost surely bounded normalized returns and cost function in (3.9) for quadratic adjustment costs and partial irreversibilities, and Caballero and Leahy

numerical methods knowing that this solution is convergent with the unique analytical solution.

Given current computing power, however, I have too many state and control variables to solve the problem as stated. But the optimization problem can be substantially simplified in two steps. First, hours are a flexible factor of production and depend only on the variables (A, K, L), which are pre-determined in period t given the time to build assumption. Therefore, hours can be optimized out in a prior step, which reduces the *control* space by one dimension. Second, the revenue function, adjustment cost function, depreciation schedules and demand processes are all jointly homogenous of degree one in (A, K, L), allowing the whole problem to be normalized by one state variable, reducing the *state* space by one dimension.[24] I normalize by capital to operate on $\frac{A}{K}$ and $\frac{L}{K}$. These two steps dramatically speed up the numerical simulation, which is run on a state space of $(\frac{A}{K}, \frac{L}{K}, \sigma, \mu)$ making numerical estimation feasible. Appendix B contains a description of the numerical solution method.

The Bellman equation of the optimization problem before simplification (dropping the firm subscripts) can be stated as

$$V(A_t, K_t, L_t, \sigma_t, \mu_t) = \max_{I_t, E_t, H_t} \left\{ \begin{array}{c} S(A_t, K_t, L_t, H_t) - C(A_t, K_t, L_t, H_t, I_t, E_t) - w(H_t)L_t \\ +\frac{1}{1+r}E[V(A_{t+1}, K_t(1-\delta_K) + I_t, L_t(1-\delta_L) + E_t, \sigma_{t+1}, \mu_{t+1})] \end{array} \right\},$$

where r is the discount rate and $E[.]$ is the expectation operator. Optimizing over hours and exploiting the homogeneity in (A, K, L) to take out a factor of K_t enables this to be re-written as

$$Q(a_t, l_t, \sigma_t, \mu_t) = \max_{i_t, e_t} \left\{ \begin{array}{c} S^*(a_t, l_t) - C^*(a_t, l_t, i_t, l_t e_t) + \\ \frac{1-\delta_K+i_t}{1+r} E[Q(a_{t+1}, l_t, \sigma_{t+1}, \mu_{t+1})] \end{array} \right\}, \tag{3.9}$$

where the normalized variables are $l_t = \frac{L_t}{K_t}$, $a_t = \frac{A_t}{K_t}$, $i_t = \frac{I_t}{K_t}$ and $e_t = \frac{E_t}{L_t}$; $S^*(a_t, l_t)$ and $C^*(a_t, l_t, i_t, l_t e_t)$ are sales and costs after optimization over hours, and $Q(a_t, l_t, \sigma_t, \mu_t) = V(a_t, 1, l_t, \sigma_t, \mu_t)$, which is Tobin's Q.

4. An Example of the Model's Solution

The model yields a central region of inaction in $(\frac{A}{K}, \frac{A}{L})$ space, due to the non-convex costs of adjustment. Firms only hire and investment when business conditions are sufficiently good, and only fire and disinvest when they are sufficiently bad. When uncertainty is higher these thresholds move out - firms become more cautious in responding to business conditions.

To provide some graphical intuition Figure 4 plots in $(\frac{A}{K}, \frac{A}{L})$ space the values of the fire and hire thresholds (left and right lines) and the sell and buy capital thresholds (top and bottom lines) for low

(1996) for the extension to fixed costs.

[24] The key to this homogeneity result is the random-walk assumption on the demand process. With a random-walk driving process adjustment costs and depreciation are naturally scaled by unit size, since otherwise units would 'outgrow' adjustment costs and depreciation. The demand-function is homogeneous through the trivial re-normalization $A^{1-a-b} = \widetilde{A}^{1-1/\epsilon}B^{1/\epsilon}$.

Figure 4: Hiring/firing and investment/disinvestment thresholds

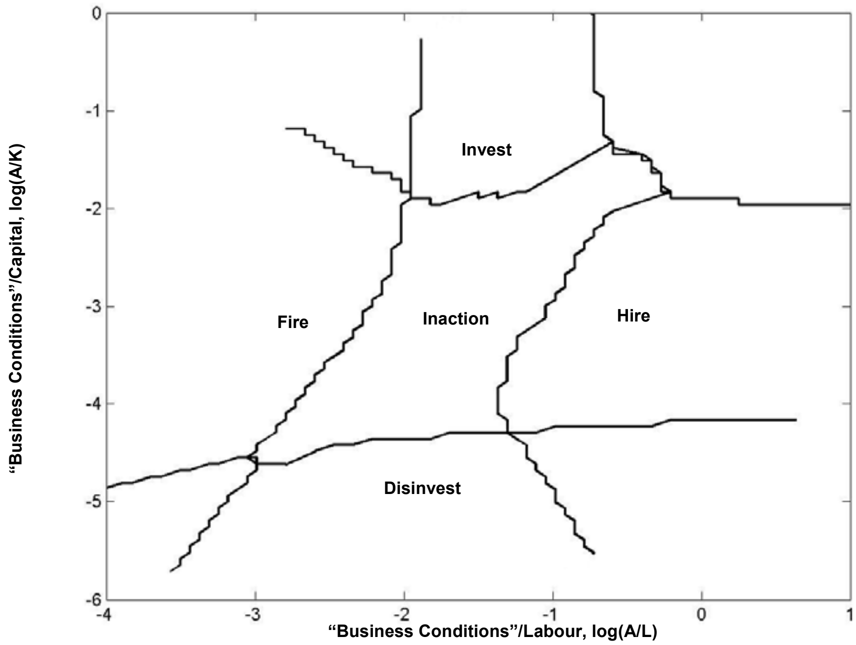

Notes: Simulated thresholds using the adjustment cost estimates "All" in table 3. All other parameters and assumptions as outlined in sections 3 and 4. Although the optimal policies are of the (s,S) type it can not be proven that this is always the case.

Figure 5: Thresholds at low and high uncertainty

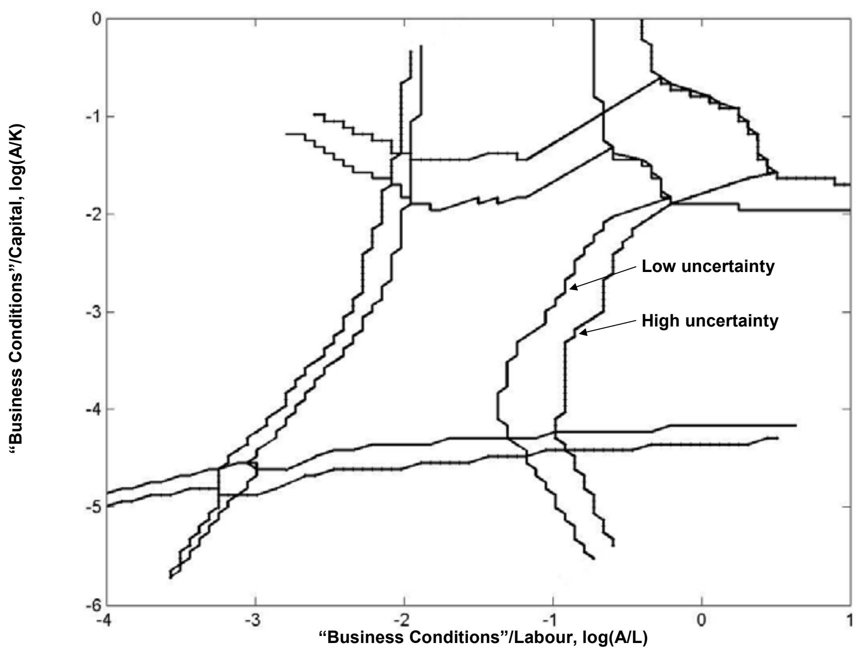

Notes: Simulated thresholds using the adjustment cost estimates "All" in Table 3. All other parameters and assumptions as outlined in sections 3 and 4. High uncertainty is twice the value of low uncertainty ($\sigma_H=2\times\sigma_L$).

uncertainty (σ_L) and the preferred parameter estimates in Table 3 column "All". The inner region is the region of inaction ($i = 0$ and $e = 0$), where the real option value of waiting is worth more than the returns to investment and/or hiring. Outside the region of inaction investment and hiring will be taking place according to the optimal values of i and e. This diagram is a two dimensional (two factor) version of the the investment models of Abel and Eberly (1996) and Caballero and Leahy (1996). The gap between the investment/disinvestment thresholds is higher than between the hire/fire thresholds due to the higher adjustment costs of capital.

Figure 5 displays the same lines for both low uncertainty (the inner box of lines), and also for high uncertainty (the outer box of lines). It can be seen that the comparative static intuition that higher uncertainty increases real options is confirmed here, suggesting that large changes in σ_t can have an important impact on investment and hiring behavior.

To quantify the impact of these real option values I run the thought experiment of calculating what *temporary* fall in wages and interest rates would be required to keep firms hiring and investment thresholds unchanged when uncertainty temporarily rises from σ_L to σ_H. The required wage and interest rate falls turn out to be quantitatively large - firms would need a 25% reduction in wages in periods of high uncertainty to leave their marginal hiring decisions unchanged, and a 7% (700 basis point) reduction in the interest rates in periods of high uncertainty to leave their marginal investment decisions unchanged. This can be graphically seen in Figure A5, which plots the low and high uncertainty thresholds, but with the change that when $\sigma_t = \sigma_H$ interest rates are 7 percentage points lower and wage rates 25% lower then when $\sigma_t = \sigma_L$.

Interestingly, re-computing these thresholds with permanent (time invariant) differences in uncertainty results in an even stronger impact on the investment and employment thresholds. So the standard comparative static result[25] on changes in uncertainty will tend to over predict the expected impact of time changing uncertainty. The reason is that firms evaluate the uncertainty of their discounted value of marginal returns over the lifetime of an investment or hire, so high current uncertainty only matters to the extent that it drives up long run uncertainty. When uncertainty is mean reverting high current values have a lower impact on expected long run values than if uncertainty were constant.

Figure 6 shows a one-dimensional cut of Figure 4 (using the same x-axis), with the level of

[25] See, for example, Dixit and Pindyck (1994). Hassler's (1996) model actually predicts that temporary shocks in uncertainty have a larger impact than permanent shocks. This arises in his model because to obtain analytical tractability he assumes fixed-costs only. With fixed costs the rise in uncertainty influences both the investment threshold and target, with these effects being smaller and larger respectively in response to a temporary versus permanent uncertainty shock. In his model the target effect dominates the threshold effect. In my model the addition of partial irreversible (and quadratic) adjustment costs reverses this so the threshold effects dominate, so permanent shocks have a larger impact than temporary shocks. This highlights the importance of estimating adjustment costs for determining the impact of uncertainty shocks.

Figure 6: The distribution of units between the hiring and firing thresholds

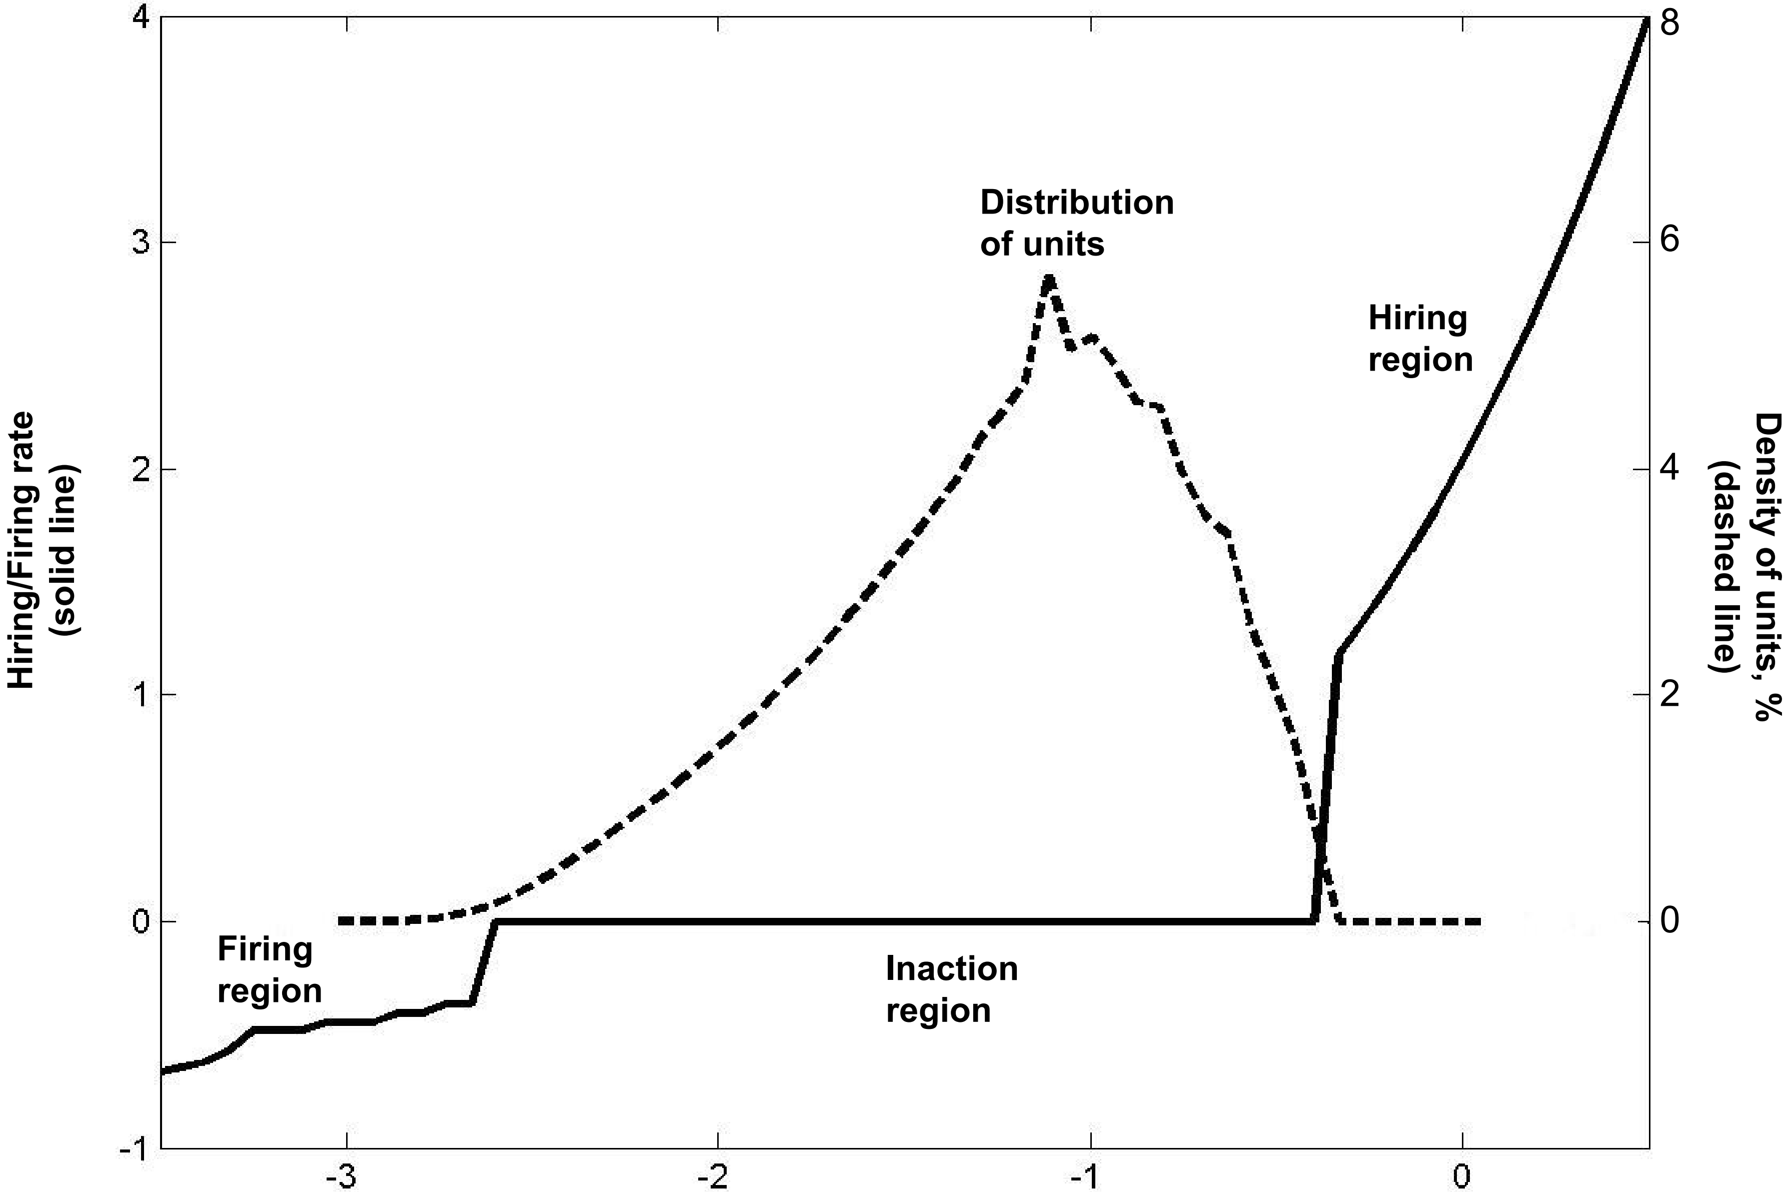

"Business Conditions"/Labor: Ln(A/L)

Notes: The hiring response (solid line) and unit-level density (dashed line) for low uncertainty (σ_L), high-drift (μ_H) and the most commonl capital/labor (K/L) ratio. All other parameters and assumptions as in sections 3 and 4. The distribution of units in (A/L) space is skewed to the right because productivity growth generates an upward drift in A and attrition generates a downward drift in L. The density peaks internally because of lumpy hiring due to fixed costs.

hiring/firing (solid line, left y-axis) and cross-sectional density of units (dashed line, right y-axis) plotted. These are drawn for one illustrative set of parameters: baseline uncertainty (σ_L), high demand growth (μ_H) and the modal value of capital/labor.[26] Three things stand out: first, the distribution is skewed to the right due to positive demand growth and labor attrition; second, the density just below the hiring threshold is low because whenever the unit hits the hiring threshold it undertakes a burst of activity (due to hiring fixed costs) that moves it to the interior of the space; and third, the density peaks at the interior which reflects the level of hiring that is optimally undertaken at the hiring threshold.

5. Estimating the Model

The econometric problem consists of estimating the parameter vector θ that characterizes the firm's revenue function, stochastic processes, adjustment costs and discount rate. Since the model has no analytical closed form these can not be estimated using standard regression techniques. Instead estimation of the parameters is achieved by simulated method of moments (SMM) which minimizes a distance criterion between key moments from the actual data and the simulated data. Because SMM is computationally intensive only 10 parameters can be estimated, with the remaining 13 predefined.

5.1. Simulated Method of Moments (SMM)

SMM proceeds as follows - a set of actual data moments Ψ^A is selected for the model to match. For an arbitrary value of θ the dynamic program is solved and the policy functions are generated. These policy functions are used to create a simulated data panel of size $(\kappa N, T+10)$, where κ is a strictly positive integer, N is the number of firms in the actual data and T is the time dimension of the actual data. The first ten years are discarded in order to start from the ergodic distribution. The simulated moments $\Psi^S(\theta)$ are then calculated on the remaining simulated data panel, along with an associated criterion function $\Gamma(\theta)$, where $\Gamma(\theta) = [\Psi^A - \Psi^S(\theta)]'W[\Psi^A - \Psi^S(\theta)]$, which is a weighted distance between the simulated moments $\Psi^S(\theta)$ and the actual moments Ψ^A.

The parameter estimate $\widehat{\theta}$ is then derived by searching over the parameter space to find the parameter vector which minimizes the criterion function:

$$\widehat{\theta} = \arg\min_{\theta\in\Theta}[\Psi^A - \Psi^S(\theta)]'W[\Psi^A - \Psi^S(\theta)] \tag{5.1}$$

Given the potential for discontinuities in the model and the discretization of the state space I use an annealing algorithm for the parameter search (see Appendix B). Different initial values of θ are selected to ensure the solution converges to the global minimum.

[26] Figure 6 is actually a 45° cut across Figure 4. The reason is Figure 6 holds K/L constant while allowing A to vary.

The efficient choice for W is the inverse of the variance-covariance matrix of $[\Psi^A - \Psi^S(\theta)]$. Defining Ω to be the variance-covariance matrix of the data moments, Lee and Ingram (1991) show that under the estimating null the variance-covariance of the simulated moments is equal to $\frac{1}{\kappa}\Omega$. Since Ψ^A and $\Psi^S(\theta)$ are independent by construction, $W = [(1+\frac{1}{\kappa})\Omega]^{-1}$, where the first term represents the randomness in the actual data and the second term the randomness in the simulated data. Ω is calculated by block bootstrap with replacement on the actual data. The asymptotic variance of the efficient estimator $\widehat{\theta}$ is proportional to $(1+\frac{1}{\kappa})$. I use $\kappa = 25$, with each of these 25 firm-panels having independent draws of macro shocks. This implies the standard error of $\widehat{\theta}$ is increased by 4% by using simulation estimation.

5.2. Predefined Parameters

In principle every parameter could be estimated, but in practice the size of the estimated parameter space is limited by computational constraints. I therefore focus on the parameters about which there is probably most uncertainty - the six adjustment cost parameters, the wage/hours trade-off slope, the baseline level of uncertainty and the two key parameters determining the firm-level demand drift, $\Theta = (PR_L, FC_L, QC_L, PR_K, FC_K, QC_K, \gamma, \sigma_L, \pi^{\mu}_{H,H}, \mu_L)$. The other thirteen parameters are based on values in the data and literature, and are displayed in Table 2 below.[27]

The predefined parameters outlined in Table 2 are mostly self-explanatory, although a few require further discussion. One of these is ϵ, which is the elasticity of demand. In a constant returns to scale production function set-up this translates directly into the returns to scale parameter on the revenue function, $a+b$. There are a wide range of estimates of the revenue returns to scale, with recent examples being 0.905 in Khan and Thomas (2003), 0.82 in Bachman, Caballero and Engel (2006) and 0.592 in Cooper and Haltiwanger (2006). I chose a parameter value of 0.75 which is: (i) roughly in the mid-point of this literature, and (ii) optimal for the speed of the numerical simulation since $a = 0.25$ and $b = 0.5$ so that capital and labor have integer fractions exponentials which compute much faster.[28] Given my assumption of constant-returns to scale and a constant-elasticity of demand this implies a markup of 33%, which is towards the upper-end of the range estimates for price-cost mark-ups. I also check the robustness of my results to a parameter value of $a+b = 0.83$, which is consistent with a 20% markup.

[27] This procedure could in principle be used iteratively to check my predefined parameters by using the estimated adjustment costs $\widehat{\theta}$ from the first round to estimate a subset of the predefined parameters in a second round of estimation and compare them to their predefined values.

[28] Integer fractional exponentials are more easily approximated in binary calculations (see Judd 1998, Chapter 2 for details). This is quantitatively important due to the intensity of exponential calculations in the simulation - for example moving from $a+b = 0.75$ to $a+b = 0.76$ slows down the simulation by around 15%. Choosing a lower value of $a+b$ also has the benefit of inducing more curvature in the value function so that less grid points are required to map any given space.

Table 2: Predefined parameters in the model

Parameter	**Value**	**Rationale** (also see the text).
α	1/3	Capital share in output is one third, labor share in output is two thirds.
ϵ	0.75	33% markup with constant returns to scale. Middle of the recent literature. I also try a+b=0.833 (20% markup).
w_1	0.8	Hourly wages minimized at a 40 hour week.
w_2	$2.4e-9$	Arbitrary scaling parameter. Set so the wage bill equals unity at 40 hours.
σ_H	$2 \times \sigma_L$	Uncertainty shocks 2× baseline uncertainty (Figure 1 data). σ_L estimated. I also try 1.5× and 3× baseline shocks.
$\pi^{\sigma}_{L,H}$	1/36	Uncertainty shocks expected every three years (16 shocks in 46 years in Figure 1).
$\pi^{\sigma}_{H,H}$	0.71	Average 2 month half-life of an uncertainty shock (Figure 1 data). I also try 1 and 6 month half-lives.
$(\mu_H + \mu_L)/2$	0.02	Average real growth rate equals 2% per year. The spread $\mu_H - \mu_L$ is estimated.
$\pi^{\mu}_{L,H}$	$\pi^{\mu}_{H,L}$	Firm-level demand growth transition matrix assumed symmetric. The parameter $\pi^{\mu}_{H,L}$ estimated.
δ_K	0.1	Capital depreciation rate assumed 10% per year.
δ_L	0.1	Labor attrition assumed 10% for numerical speed (since $\delta_L = \delta_K$). I also try $\delta_L = 0.2$.
r	6.5%	Long-run average value for US firm-level discount rate (King and Rebello, 1999).
N	250	Firms operate 250 units, chosen to achieve complete aggregation. I also try $N = 25$ and $N = 1$.

Notes: Reports the predetermined parameter values used in the estimations in section (6) and simulations in section (7).

The uncertainty process parameters are primarily taken from the macro volatility process in Figure 1, with the baseline level of uncertainty estimated in the simulation. The labor attrition rate is chosen at 10% per annum. This low figure is selected for two reasons: (i) to be conservative in the simulations of an uncertainty shock since attrition drives the fall in employment levels, so that lower levels reduces the impact of shocks; and (ii) for numerical speed as this matches the capital depreciation rate, so that the (L/K) dimension can be ignored if no investment and hiring/firing occurs. I also report a robustness test for using an annualized labor attrition rate of 20% which more closely matches the figures for annualized manufacturing quits in Davis, Faberman and Haltiwanger (2006).

5.3. Identification

Under the null any full-rank and sufficient order set of moments (Ψ^A) will identify *consistent* parameter estimates for the adjustment costs (Θ). However, the precise choice of moments is important for the *efficiency* of the estimator, suggesting moments which are "informative" about the underlying structural parameters should be included. The basic insights of plant and firm-level data on labor and capital is the presence of highly skewed cross-sectional growth rates and rich time-series dynamics, suggesting some combination of cross-sectional and time-series moments. Two additional issues help to guide the exact choice of moments.

5.3.1. Distinguishing the Driving Process from Adjustment Costs

A key challenge in estimating adjustment costs for factor inputs is distinguishing between the dynamics of the driving process and factor adjustment costs. Concentrating on the moments from only one factor - for example capital - makes it very hard to do this. To illustrate this first consider a very smooth driving process without adjustment costs, which would produce a smooth investment series. Alternatively consider a volatile driving process with convex capital adjustment costs, which would also produce a smooth investment series. Hence, without some additional moments (or assumptions) it would be very hard to estimate adjustment costs using just the investment series data.

So I focus on the joint (cross-sectional and dynamic) moments of the investment, employment and sales growth series. The difference in responses across the three series (investment, employment and sales growth) should identify the two sets of adjustment costs (for capital and labor).[29]

[29] An alternative is a two-step estimation process in which the driving process is estimated first and then the adjustment costs estimated given this driving process (see for example Cooper and Haltiwanger, 2006).

5.3.2. Distinguishing Persistent Differences from Adjustment Costs

A stylized fact from the estimation of firm and plant level investment and labor demand equations is the empirical importance of 'fixed-effects' - that is persistent differences across firms and plants in their levels of investment, employment and output growth rates. Without controls for these persistent differences the estimates of the adjustment costs could be biased. For example, persistent between-firm differences in investment, employment and sales growth rates due to different growth rates of demand would (in the absence of controls for this) lead to the estimation of large quadratic adjustment costs, necessary to induce the required firm-level autocorrelation.

To control for differential firm-level growth rates the estimator includes two parameters: the spread of firm-level business conditions growth, $\mu_H - \mu_L$, which determines the degree of firm-level heterogeneity in the average growth rates of business conditions as defined in (3.5); and the persistence of firm-level business conditions growth, $\pi^{\mu}_{H,H}$, as defined in (3.8). When $\mu_H - \mu_L$ is large there will be large differences in the growth rates of labor, capital and output across firms, and when $\pi^{\mu}_{H,H}$ is close to unity these will be highly persistent.[30] To identify these parameters separately from adjustment costs requires information on the time path of autocorrelation across the investment, employment and sales growth series. For example, persistent correlations between investment, sales and employment growth rates going back over many years would help to identify fixed differences in the growth rates of the driving process, while decaying correlations in the investment series *only* would suggest convex capital adjustment costs.

So I include moments for the second-order and fourth-order correlations of the investment, employment growth and sales growth series.[31] The second-order autocorrelation is chosen to avoid a negative bias in these moments from underlying levels measurement errors which would arise in a first-order autocorrelation measure, while the fourth-order autocorrelation is chosen to allow a sufficiently large time-period to pass (2 years) to identify the decay in the autocorrelation series. Shorter and longer lags, like the third-order, fifth-order and sixth-order order autocorrelations could also be used, but in experimentations did not make much difference.[32]

5.4. Firm-Level Data

There is too little data at the macroeconomic level to provide sufficient identification for the model. I therefore identify my parameters using a panel of firm-level data from US Compustat. I select the 20 years of data covering 1981 to 2000.

[30] Note that with $\pi^{\mu}_{H,H} = 1$ these will be truly 'fixed effect' differences.

[31] To note, a k^{th} order correlation for series $x_{i,t}$ and $y_{i,t}$ is defined as $Corr(x_{it}, y_{it-k})$

[32] Note that because the optimal weighting matrix takes into account the covariance across moments, adding extra moments that are highly correlated to included moments has very little impact on the parameters estimates.

The data were cleaned to remove major mergers and acquisitions by dropping the top and bottom 0.5% of employment growth, sales growth and investment rates. Only firms with an average of at least 500 employees and $10m sales (in 2000 prices) were kept to focus on larger more aggregated firms. This generated a sample of 2548 firms and 22,950 observations with mean (median) employees of 13,540 (3,450) and mean (median) sales of $2247m ($495m) in 2000 prices. In selecting all Compustat firms I am conflating the parameter estimates across a range of different industries, and a strong argument can be made for running this estimation on an industry by industry basis. However, in the interests of obtaining the "average" parameters for a macro simulation, and to ensure a reasonable sample size, I keep the full panel leaving industry specific estimation to future work.

Capital stocks for firm i in industry m in year t are constructed by the perpetual inventory method[33], labor figures come from company accounts, while sales figures come from accounts after deflation using the CPI. The investment rate is calculated as $(\frac{I}{K})_{i,t} = \frac{I_{i,t}}{0.5*(K_{i,t}+K_{i,t-1})}$, the employment growth rate as $(\frac{\Delta L}{L})_{i,t} = \frac{L_{i,t}-L_{i,t-1}}{0.5*(L_{i,t}+L_{i,t-1})}$ and the sales growth as $(\frac{\Delta S}{S})_{i,t} = \frac{S_{i,t}-S_{i,t-1}}{0.5*(S_{i,t}+S_{i,t-1})}$.[34]

The simulated data is constructed in exactly the same way as company accounts are built. First, firm value is created by adding up across the N units in each firm. It is then converted into annual figures using standard accounting techniques: simulated data for 'flow' figures from the accounting Profit & Loss and Cash-Flow statements (such as sales and capital expenditure) are added up across the 12 months of the year; simulated data for 'stock' figures from the accounting Balance Sheet statement (such as the capital stock and labor force) are taken from the year end values.

By constructing my simulation data in the same manner as company accounts I can estimate adjustment costs using firm-level datasets like Compustat. This has some advantages versus using census datasets like the LRD because firm-level data is: (i) easily available to all researchers in a range of different countries; (ii) is matched into firm level financial and cash-flow data; and (iii) is available as a yearly panel stretching back several decades (for example to the 1950s in the U.S.). Thus, this technique of explicitly building aggregation into estimators to match against aggregated quoted firm-level data should have a broader use in other applications.

5.5. Measurement Errors

Employment figures are often poorly measured in company accounts, typically including all part-time, seasonal and temporary workers in the total employment figures without any adjustment for

[33] $K_{i,t} = (1-\delta_K)K_{i,t-1}\frac{P_{m,t}}{P_{m,t-1}}+I_{i,t}$, initialized using the net book value of capital, where $I_{i,t}$ is net capital expenditure on plant, property and equipment, and $P_{m,t}$ are the industry level capital goods deflators from Bartelsman, Becker and Grey (2000).

[34] *Gross* investment rates and *net* employment growth rates are used since these are directly observed in the data. Under the null that the model is correctly specified the choice of net versus gross is not important for the consistency of parameter estimates so long as *the same* actual and simulated moments are matched.

hours, usually after heavy rounding. This problem is then made much worse by the differencing to generate growth rates.

As a first step towards reducing the sensitivity towards these measurement errors, the autocorrelations of growth rates are taken over longer periods (as noted above). As a second step, I explicitly introduce employment measurement error into the simulated moments to try and mimic the bias these impute into the actual data moments. To estimate the size of the measurement error I assume that firm wages (W_{it}) can be decomposed into $W_{it} = \eta_t \lambda_{j,t} \phi_i L_{it}$ where η_t is the absolute price level, $\lambda_{j,t}$ is the relative industry wage rate, ϕ_i is a firm specific salary rate (or skill/seniority mix) and L_{it} is the average annual firm labor force (hours adjusted). I then regress $\log W_{it}$ on a full set of year dummies, a log of the 4-digit SIC industry average wage from Bartelsman, Becker and Gray (2000), a full set of firm specific fixed effects and $\log L_{it}$. Under my null on the decomposition of W_{it} the coefficient on $\log L_{it}$ will be approximately $\frac{\sigma_L^2}{\sigma_L^2+\sigma_{ME}^2}$ where σ_L^2 is the variation in log employment and σ_{ME}^2 is the measurement error in log employment. I find a coefficient (s.e.) on $\log L_{it}$ of 0.882 (0.007), implying a measurement error of 13% in the logged labor force numbers.[35] This is reassuringly similar to the 8% estimate for measurement error in Compustat manufacturing firms' labor figures Hall (1987) calculates comparing OLS and IV estimates. I take the average of these figures and incorporate this into the simulation estimation by multiplying the aggregated annual firm labor force by $me_{i,t}$ where $me_{i,t} \sim i.i.d.\ LN(0, 0.105)$ before calculating simulated moments.

6. Adjustment Costs Estimates

In this section I present the estimates of the firms capital and labor adjustment costs. Starting with Table 3, the first column labelled "Data" in the bottom panel reports the actual moments from Compustat. These demonstrate that investment rates have a low spread but a heavy right skew due to the lack of disinvestment, and strong dynamic correlations. Labor growth rates are relatively variable but un-skewed, with weaker dynamic correlations. Sales growth rates have similar moments to those of labor, although slightly lower spread and higher degree of dynamics correlations.

The second column in Table 3 labelled "All" presents the results from estimating the preferred specification allowing for all types of adjustment costs. The estimated adjustment costs for capital imply a large resale loss of around 34% on capital, fixed investment costs of 1.5% of annual sales (about 4 working days) and no quadratic adjustment costs. The estimated labor adjustment costs

[35] Adding firm or industry specific wage trends reduces the coefficient on $\log W_{it}$ implying an even higher degree of measurement error. Running the reverse regression of log labour on log wages plus the same controls generates a coefficient (s.e.) of 0.990 (0.008), indicating that the proportional measurement error in wages (a typically much better recorded financial variable) is many times smaller than that of employment. The regressions are run using 2468 observations on 219 firms.

Table 3: Adjustment cost estimates

Adjustment Costs Specification:		All	Capital	Labor	Quad	None
Estimated Parameters:						
$\mathbf{C}_K^P$		33.9	42.7			
investment resale loss (%)		*(6.8)*	*(14.2)*			
$\mathbf{C}_K^F$		1.5	1.1			
investment fixed cost (% annual sales)		*(1.5)*	*(0.2)*			
$\mathbf{C}_K^Q$		0	0.996		4.844	
capital quadratic adjustment cost (parameter)		*(0.009)*	*(0.044)*		*(454.15)*	
$\mathbf{C}_L^P$		1.8		16.7		
per capita hiring/firing cost (% annual wages)		*(0.8)*		*(0.1)*		
$\mathbf{C}_L^F$		2.1		1.1		
fixed hiring/firing costs (% annual sales)		*(0.9)*		*(0.1)*		
$\mathbf{C}_L^Q$		0		1.010	0	
labor quadratic adjustment cost (parameter)		*(0.037)*		*(0.017)*	*(0.002)*	
$\boldsymbol{\sigma}_L$		0.443	0.413	0.216	0.171	0.100
baseline level of uncertainty		*(0.009)*	*(0.012)*	*(0.005)*	*(0.005)*	*(0.005)*
$\boldsymbol{\mu}_H - \boldsymbol{\mu}_L$		0.121	0.122	0.258	0.082	0.158
spread of firm business conditions growth		*(0.002)*	*(0.002)*	*(0.001)*	*(0.001)*	*(0.001)*
$\boldsymbol{\pi}_{H,L}^{\mu}$		0	0	0.016	0	0.011
transition of firm business conditions growth		*(0.001)*	*(0.001)*	*(0.001)*	*(0.001)*	*(0.001)*
$\boldsymbol{\gamma}$		2.093	2.221	3.421	2.000	2.013
curvature of the hours/wages function		*(0.272)*	*(0.146)*	*(0.052)*	*(0.009)*	*(14.71)*
Moments:	*Data*	Simulated moments - Data moments				
Correlation $(I/K)_{i,t}$ with $(I/K)_{i,t-2}$	*0.328*	0.060	-0.015	0.049	-0.043	0.148
Correlation $(I/K)_{i,t}$ with $(I/K)_{i,t-4}$	*0.258*	0.037	0.004	0.088	0.031	0.162
Correlation $(I/K)_{i,t}$ with $(\Delta L/L)_{i,t-2}$	*0.208*	0.003	-0.025	0.004	-0.056	0.078
Correlation $(I/K)_{i,t}$ with $(\Delta L/L)_{i,t-4}$	*0.158*	-0.015	-0.009	0.036	0.008	0.091
Correlation $(I/K)_{i,t}$ with $(\Delta S/S)_{i,t-2}$	*0.260*	-0.023	-0.062	-0.044	-0.102	0.024
Correlation $(I/K)_{i,t}$ with $(\Delta S/S)_{i,t-4}$	*0.201*	-0.010	-0.024	0.018	-0.036	0.087
Standard Deviation $(I/K)_{i,t}$	*0.139*	-0.010	0.010	-0.012	0.038	0.006
Coefficient of Skewness $(I/K)_{i,t}$	*1.789*	0.004	0.092	1.195	1.311	1.916
Correlation $(\Delta L/L)_{i,t}$ with $(I/K)_{i,t-2}$	*0.188*	-0.007	0.052	-0.075	0.055	0.053
Correlation $(\Delta L/L)_{i,t}$ with $(I/K)_{i,t-4}$	*0.133*	-0.021	0.024	-0.061	0.038	0.062
Correlation $(\Delta L/L)_{i,t}$ with $(\Delta L/L)_{i,t-2}$	*0.160*	0.011	0.083	-0.033	0.071	0.068
Correlation $(\Delta L/L)_{i,t}$ with $(\Delta L/L)_{i,t-4}$	*0.108*	-0.013	0.054	-0.026	0.045	0.060
Correlation $(\Delta L/L)_{i,t}$ with $(\Delta S/S)_{i,t-2}$	*0.193*	-0.019	0.063	-0.091	0.064	0.023
Correlation $(\Delta L/L)_{i,t}$ with $(\Delta S/S)_{i,t-4}$	*0.152*	0.003	0.056	-0.051	0.059	0.063
Standard Deviation $(\Delta L/L)_{i,t}$	*0.189*	-0.022	-0.039	0.001	-0.001	0.005
Coefficient of Skewness $(\Delta L/L)_{i,t}$	*0.445*	-0.136	0.294	-0.013	0.395	0.470
Correlation $(\Delta S/S)_{i,t}$ with $(I/K)_{i,t-2}$	*0.203*	-0.016	-0.015	-0.164	-0.063	-0.068
Correlation $(\Delta S/S)_{i,t}$ with $(I/K)_{i,t-4}$	*0.142*	-0.008	-0.010	-0.081	-0.030	-0.027
Correlation $(\Delta S/S)_{i,t}$ with $(\Delta L/L)_{i,t-2}$	*0.161*	-0.005	0.032	-0.105	-0.024	-0.037
Correlation $(\Delta S/S)_{i,t}$ with $(\Delta L/L)_{i,t-4}$	*0.103*	-0.015	0.011	-0.054	-0.005	-0.020
Correlation $(\Delta S/S)_{i,t}$ with $(\Delta S/S)_{i,t-2}$	*0.207*	-0.033	0.002	-0.188	-0.040	-0.158
Correlation $(\Delta S/S)_{i,t}$ with $(\Delta S/S)_{i,t-4}$	*0.156*	0.002	0.032	-0.071	-0.021	-0.027
Standard Deviation $(\Delta S/S)_{i,t}$	*0.165*	0.004	0.003	0.033	0.051	0.062
Coefficient of Skewness $(\Delta S/S)_{i,t}$	*0.342*	-0.407	-0.075	-0.365	0.178	0.370
Criterion, $\Gamma(\theta)$		404	625	3618	2798	6922

Notes to Table 3: The "Data" column (bottom panel only) contains the moments from 22,950 observations on 2548 firms. The other columns contain the adjustment costs estimates (top panel) and simulated moments minus the data moments (bottom panel) for: all adjustment costs ("All"), just capital adjustment costs ("Capital"), just labor adjustment costs ("Labor"), just quadratic adjustment costs with 1 unit per firm ("Quad") and no adjustment costs ("None"). So, for example, the number 0.328 at the top of the first column ("Data") reports that the second-lag of the autocorrelation of investment in the data is 0.328, and the number 0.060 to the right reports that in the "All" specification the simulated moment is 0.060 greater than the data moment (so is 0.388 in total). In the top panel standard-errors in italics in brackets below the point estimates. Parameters estimated using Simulated Method of Moments, and standard errors calculated using numerical derivatives. All adjustment-cost parameters constrained to be non-negative. Full simulation and estimation details in Appendix B.

imply limited hiring and firing costs of about 1.8% of annual wages (about 5 working days) and a high-fixed cost of around 2.1% of annual revenue (about 6 working days), with no quadratic adjustment costs. The standard errors suggest all of these point estimates are statistically significant except for the fixed cost of capital adjustment (C_K^F)

One question is how do these estimates compare to those previously estimated in the literature? Table 4 presents a comparison for some other estimates from the literature. Three factors stand

Table 4: A comparison with other capital and labor adjustment cost estimates

	Capital			Labor		
Source:	PI (%)	Fixed (%)	Quad	PI (%)	Fixed (%)	Quad
Column "All", Table 3, this paper	33.9	1.5	0	1.8	2.1	0
Ramey and Shapiro (2001)	40 to 80					
Caballero and Engel (1999)		16.5				
Hayashi (1982)			480			
Cooper and Haltiwanger (2006)	2.5	20.4	0.294			
Shapiro (1986)			36			16
Hall (2004)			0			0
Nickel (1986)				8 to 25		
Cooper, Haltiwanger & Willis (2004)					1.7	0

Note: 'PI' denotes partial irreversibilities, 'Fixed' denotes fixed costs, and 'Quad' denotes quadratic adjustment costs. Missing values indicate no parameter estimated in the main specification. Zeros indicate the parameter was not significantly different from zero. Nickel's (1986) lower[higher] value is for unskilled[skilled] workers. Shapiro's (1986) value is a weighted average of (2/3)×0 for production workers and (1/3)×48 for non-production workers. Quadratic adjustment costs defined monthly (12 times the yearly parameter). Comparability subject to variation in data sample, estimation technique and maintained assumptions.

out: first, there is tremendous variation is estimated adjustment costs, reflecting the variety of data, techniques and assumptions used in the different papers; second, my estimates of zero quadratic adjustment costs appear broadly consistent with recent papers using detailed industry or micro data; and third, studies which estimate non-convex adjustment costs report positive, and typically very substantial values.

For interpretation I also display results for four illustrative restricted models. First, a model with capital adjustment costs only, assuming labor is fully flexible, as is typical in the investment literature. In the "Capital" column we see that the fit of the model is worse, as shown by the significant rise in the criterion function from 404 to 625.[36] This reduction in fit is primarily due to the worse fit of the labor moments, suggesting ignoring labor adjustment costs is a reasonable approximation for modelling investment. Second, a model with labor adjustment costs only - as is typical in the dynamic labor demand literature - is estimated in the column "Labor", with the fit

[36] The χ^2 value for 3 degrees of freedom is 7.82, so column "Capital" can easily be rejected against the null of "All" given the difference in criterion values of 221. It is also true, however, that the preferred "All" specification can also be rejected as the true representation of the data given the χ^2 value for 10 degrees of freedom is 18.31.

substantially. This suggests that ignoring capital adjustment costs is problematic. Third, a model with quadratic costs only and no cross-sectional aggregation - as is typical in convex adjustment costs models - is estimated in the "Quad" column, leading to a moderate reduction in fit generated by excessive intertemporal correlation and an inadequate investment skew. Interestingly, industry and aggregate data are much more autocorrelated and less skewed due to extensive aggregation, suggesting quadratic adjustments costs could be a reasonable approximation at this level.[37] Finally, a model with no adjustment costs is estimated in column "None". Omitting adjustment costs clearly reduces the model fit. It also biases the estimation of the business-conditions process to have much larger firm-level growth fixed-effects and lower variance of the idiosyncratic shocks. This helps to highlight the importance of jointly estimating adjustment costs and the driving process.

In Table 3 there are also some estimates of the driving process parameters $\sigma_L, \mu_H - \mu_L$ and $\pi^{\mu}_{H,H}$, as well as the wage-hours curve parameter γ. What is clear is that changes in the adjustment cost parameters leads to changes in these parameters. For example, the lack of adjustment costs in column "Quad" generates an estimated uncertainty parameter of around 1/3 of the baseline "All" value and a spread in firm-level fixed costs of about 2/3 of the baseline "All" value. This provides support for the selection of moments that can separately identify the driving process and adjustment cost parameters.

6.1. Robustness Tests on Estimated Parameters

In Table 5 I run some robustness tests on the modelling assumptions. The first column "All" repeats the baseline results from Table 3 for ease of comparison.

The column "δ_L=0.2" reports the results from re-estimating the model with a 20% (rather then 10%) annual attrition rate for labor. This higher rate of attrition leads to higher quadratic adjustment costs for labor and capital, and lower fixed-costs for labor. This is because with higher labor attrition rates hiring and firing become more sensitive to current demand shocks (since higher attrition reduces the sensitivity to past shocks). To compensate the estimated quadratic adjustment costs estimates are higher and fixed costs lower. The column "a+b=0.83" reports the results for a specification with a 20% markup, in which the estimated adjustment costs look very similar to the baseline results.

In columns "N=25" and "N=1" the results are reported for simulations assuming the firm operates 25 units and 1 unit respectively.[38] These assumptions also lead to higher estimates for the quadratic adjustment costs and lower estimates for the non-convex adjustment costs to compensate

[37] Cooper and Haltiwanger (2006) also note this point.

[38] The specification with N=1 is included to provide guidance on the impact of simulated aggregation rather than for empirical realism. The evidence of aggregation in Appendix A4, and from the annual report of any large company with its typical multi-divisional, suggests aggregation is likely to be pervasive.

Table 5: Adjustment cost robustness tests

Adjustment Costs Specification:	All	δ_L=20%	a+b=0.83	N=25	N=1	Yearly
Estimated Parameters:						
$\mathbf{C}_K^P$	33.9	28.6	29.8	30.3	47.0	45.3
investment resale loss (%)	*(6.8)*	*(4.8)*	*(4.8)*	*(8.7)*	*(9.1)*	*(5.2)*
$\mathbf{C}_K^F$	1.5	2.1	2.1	0.9	1.3	2.1
investment fixed cost (% annual sales)	*(1.0)*	*(0.9)*	*(0.5)*	*(0.4)*	*(0.2)*	*(0.3)*
$\mathbf{C}_K^Q$	0	0.461	0	0.616	2.056	0.025
capital quadratic adjustment cost (parameter)	*(0.009)*	*(0.054)*	*(0.007)*	*(0.154)*	*(0.284)*	*(0.015)*
$\mathbf{C}_L^P$	1.8	1.0	0	0	0	2.0
per capita hiring/firing cost (% annual wages)	*(0.8)*	*(0.1)*	*(0.0)*	*(0.1)*	*(0.1)*	*(0.9)*
$\mathbf{C}_L^F$	2.1	0.3	1.7	1.3	0	2.0
fixed hiring/firing costs (% annual sales)	*(0.9)*	*(0.1)*	*(0.6)*	*(0.8)*	*(0.0)*	*(0.5)*
$\mathbf{C}_L^Q$	0	0.360	0	0.199	0.070	1.039
labor quadratic adjustment cost (parameter)	*(0.037)*	*(0.087)*	*(0.021)*	*(0.062)*	*(0.031)*	*(0.165)*
$\boldsymbol{\sigma}_L$	0.443	0.490	0.498	0.393	0.248	0.339
baseline level of uncertainty	*(0.009)*	*(0.019)*	*(0.012)*	*(0.013)*	*(0.008)*	*(0.011)*
$\boldsymbol{\mu}_H - \boldsymbol{\mu}_L$	0.121	0.137	0.123	0.163	0.126	0.228
spread of firm business conditions growth	*(0.002)*	*(0.002)*	*(0.001)*	*(0.002)*	*(0.002)*	*(0.005)*
$\boldsymbol{\pi}_{H,L}^{\mu}$	0	0	0	0	0	0.016
transition of firm business conditions growth	*(0.001)*	*(0.001)*	*(0.001)*	*(0.001)*	*(0.001)*	*(0.001)*
$\boldsymbol{\gamma}$	2.093	2.129	2.000	2.148	2.108	2.000
curvature of the hours/wages function	*(0.272)*	*(0.222)*	*(0.353)*	*(0.266)*	*(0.147)*	*(0.166)*
Moments:		Simulated moments - Data moments				
Correlation $(I/K)_{i,t}$ with $(I/K)_{i,t-2}$	0.060	0.021	0.065	0.002	0.078	-0.0289
Correlation $(I/K)_{i,t}$ with $(I/K)_{i,t-4}$	0.037	0.017	0.042	0.027	0.081	-0.088
Correlation $(I/K)_{i,t}$ with $(\Delta L/L)_{i,t-2}$	0.003	-0.020	-0.005	-0.038	0.014	0.025
Correlation $(I/K)_{i,t}$ with $(\Delta L/L)_{i,t-4}$	-0.015	-0.017	-0.008	-0.020	0.016	-0.023
Correlation $(I/K)_{i,t}$ with $(\Delta S/S)_{i,t-2}$	-0.023	-0.044	-0.023	-0.075	-0.003	0.033
Correlation $(I/K)_{i,t}$ with $(\Delta S/S)_{i,t-4}$	-0.010	-0.018	-0.008	-0.021	-0.001	-0.032
Standard Deviation $(I/K)_{i,t}$	-0.010	-0.009	-0.006	0.001	-0.008	0.014
Coefficient of Skewness $(I/K)_{i,t}$	0.004	-0.010	-0.022	-0.088	-0.188	0.043
Correlation $(\Delta L/L)_{i,t}$ with $(I/K)_{i,t-2}$	-0.007	0.054	0.007	-0.041	0.043	0.029
Correlation $(\Delta L/L)_{i,t}$ with $(I/K)_{i,t-4}$	-0.021	0.024	-0.002	-0.020	0.026	-0.027
Correlation $(\Delta L/L)_{i,t}$ with $(\Delta L/L)_{i,t-2}$	0.011	0.070	0.016	-0.018	0.052	0.014
Correlation $(\Delta L/L)_{i,t}$ with $(\Delta L/L)_{i,t-4}$	-0.013	0.040	-0.001	-0.009	0.024	-0.020
Correlation $(\Delta L/L)_{i,t}$ with $(\Delta S/S)_{i,t-2}$	-0.019	0.054	-0.008	-0.054	0.048	0.032
Correlation $(\Delta L/L)_{i,t}$ with $(\Delta S/S)_{i,t-4}$	0.003	0.058	0.017	0.003	0.032	-0.046
Standard Deviation $(\Delta L/L)_{i,t}$	-0.022	-0.044	-0.028	-0.021	-0.030	0.023
Coefficient of Skewness $(\Delta L/L)_{i,t}$	-0.136	0.207	-0.179	-0.082	0.036	-0.051
Correlation $(\Delta S/S)_{i,t}$ with $(I/K)_{i,t-2}$	-0.016	0.001	-0.024	-0.063	-0.023	0.048
Correlation $(\Delta S/S)_{i,t}$ with $(I/K)_{i,t-4}$	-0.008	-0.001	-0.005	-0.037	-0.016	-0.003
Correlation $(\Delta S/S)_{i,t}$ with $(\Delta L/L)_{i,t-2}$	-0.005	0.018	-0.021	-0.042	-0.007	0.040
Correlation $(\Delta S/S)_{i,t}$ with $(\Delta L/L)_{i,t-4}$	-0.015	0.003	-0.019	-0.033	-0.021	0.008
Correlation $(\Delta S/S)_{i,t}$ with $(\Delta S/S)_{i,t-2}$	-0.033	0.009	-0.050	-0.060	-0.015	0.087
Correlation $(\Delta S/S)_{i,t}$ with $(\Delta S/S)_{i,t-4}$	0.002	0.034	-0.009	-0.010	-0.024	-0.018
Standard Deviation $(\Delta S/S)_{i,t}$	0.004	-0.012	0.006	0.001	0.011	-0.009
Coefficient of Skewness $(\Delta S/S)_{i,t}$	-0.407	-0.132	-0.251	-0.484	-0.417	-0.176
Criterion, $\Gamma(\theta)$	404	496	379	556	593	656

Notes to Table 5: The columns contain the adjustment costs estimates (top panel) and simulated moments minus the data moments (bottom panel). The moments come from 22,950 observations on 2548 firms, and are reported in full in Table 3). The columns report results for: the baseline model with all adjustment costs ("All"), baseline model but with 20% annualized labor attrition ("δ_L=20%"), baseline model but with a 20% mark-up ("a+b=0.83"), baseline model but with only 25 units per firm ("N=25"), baseline model but with only 1 unit per firm("N=1"), and the baseline model but with the simulation run at a yearly frequency (rather than monthly and aggregated to the yearly level) ("$Yearly$"). So, for example, the number 0.060 at the top of the first column ("All") reports that the second-lag of the autocorrelation of investment in the data is 0.060 greater than the data moment (so is 0.388 in total). In the top panel standard-errors in italics in brackets below the point estimates. Parameters estimated using Simulated Method of Moments, and standard errors calculated using numerical derivatives. All adjustment-cost parameters constrained to be non-negative. Full simulation and estimation details contained in Appendix B.

for the reduction in smoothing by aggregation. Finally, the column "Yearly" reports the results for running the simulation at a yearly frequency without any time aggregation. Dropping time aggregation leads to higher estimated quadratic adjustment costs, again to compensate for the loss of smoothing by aggregation. Hence, modelling cross-sectional or time aggregation appears to matter for estimating adjustment costs since these play a role in smoothing data moments.

I also used the estimated parameters across all the columns to re-run the baseline simulation for the impact of an uncertainty shock (full details in section 7.6.2). The key result of a drop and rebound in activity was *qualitatively* robust for all the columns, although there was some variation in the magnitude of this.

7. Simulating an Uncertainty Shock

The simulation models the impact of a large, but temporary, rise in the variance of business conditions (productivity and demand) growth. This second-moment shock generates a rapid drop in hiring, investment and productivity growth as firms become much more cautious due to the rise in uncertainty. Once the uncertainty shock passes, however, activity bounces back as firms clear their pent-up demand for labor and capital. This also leads to a drop and rebound in productivity growth, since the temporary pause in activity slows down the reallocation of labor and capital from low to high productivity units. In the medium term this burst of volatility generates an overshoot in activity due to the convexity of hiring and investment in business conditions.

Of course this is a stylized simulation since other factors also typically change around major shocks. Some of these factors can and will be added to the simulation, for example allowing for a simultaneous negative shock to the first moment. I start by focusing on a second moment shock only, however, to isolate the pure uncertainty effects and demonstrate that these alone are capable of generating large short-run fluctuations. I then discuss the robustness of this analysis to price changes from general equilibrium effects, a combined first and second moment shock, different estimates for the adjustment costs, different predetermined parameters and different stochastic processes for the uncertainty shock.

7.1. The Baseline Simulation Outline

I simulate an economy of 1000 units (4 firms) for 15 years at a monthly frequency. This simulation is then repeated 25,000 times, with the values for labor, capital, output and productivity averaged over all these runs. In each simulation the model is hit with an uncertainty shock in month 1 of year 11, defined as $\sigma_t = \sigma_H$ in equation (3.7). All other micro and macro shocks are randomly drawn as per sections (3) and (5). The estimated values are taken from the "All" column in Table

3. This generates the average impact of an uncertainty shock, where the average is taken over the distribution of micro and macro shocks.

Before presenting the simulation results it is worth first showing the precise impulse that will drive the results. Figure 7a reports the average value of σ_t normalized to unity before the shock. It is plotted on a monthly basis, with the month normalized to zero on the date of the shock. Three things are clear from Figure 7a: first, the uncertainty shock generates a sharp spike in the average σ_t across the 25,000 simulations, second this dies off rapidly with a half-life of 2 months, and third the shock almost doubles average σ_t (the rise is less than 100% because some of the 25,000 simulations already had $\sigma_t = \sigma_H$ when the shock occurred). In Figure 7b I show the average time path of business conditions ($A_{j,t}$) showing that the uncertainty shock has no first moment effect.

In Figure 8 I plot aggregate detrended labor, again normalized to 1 at the month before the shock. This displays a substantial fall in the six months immediately after the uncertainty shock and then overshoots from months 8 onwards, eventually returning to level by around 3 years.

The initial drop occurs because the rise in uncertainty increases the real-option value of inaction, leading the majority of firms to temporarily freeze hiring. Because of the ongoing exogenous attrition of workers this generates a fall in net employment. Endogenizing quits would of course reduce the impact of these shocks since the quit rate would presumably fall after a shock. But in the model to offset this I have conservatively assumed a 10% annual quit rate - well below the 15% to 25% quit rate observed over the business cycle in recent JOLTS data (see Davis, Faberman and Haltiwanger. 2006). This low fixed quit rate could be thought of as the exogenous component due to retirement, maternity, sickness, family-relocation etc.

The rebound from months 4 onwards occurs because of the combination of falling uncertainty (since the shock is only temporary) and rising pent-up demand for hiring (because firms paused hiring over the previous three months). In order to make up the short-fall in labor firms begin to hire at a faster pace than usual so the labor force heads back towards it trend-level. This generates the rapid rebound in the total labor from month 3 until about month 6. From month 7 onwards this overshoot gradually returns to trend.

7.2. The Volatility Overshoot

One seemingly puzzling phenomenon, however, is the overshoot from month 7 onwards. Pure real-options effects of uncertainty should generate a drop and overshoot in the *growth rate* of labor (that is the hiring rate), but only a drop and convergence back to trend in the *level* of the labor force. So the question is what is causing this medium term overshooting in the level of the labor force?

This medium term overshoot arises because the increased volatility of business conditions leads

Figure 7a: The simulation has a large second moment shock

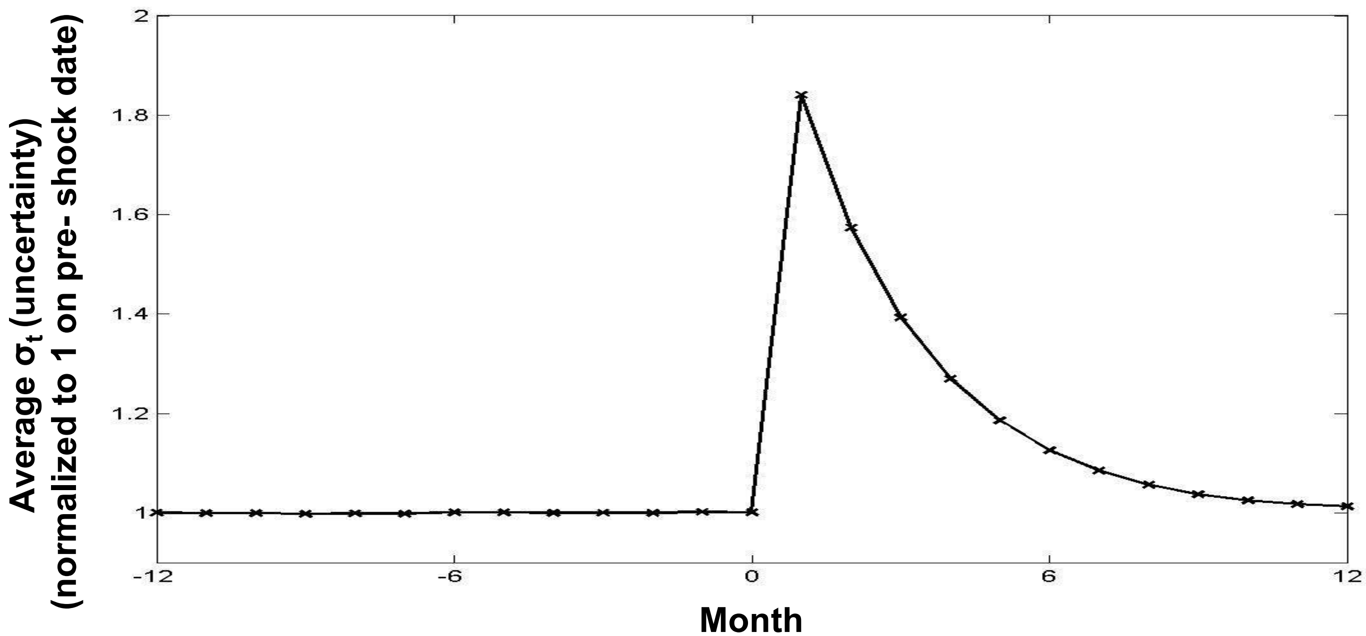

Figure 7b: The simulation has no first-moment shock

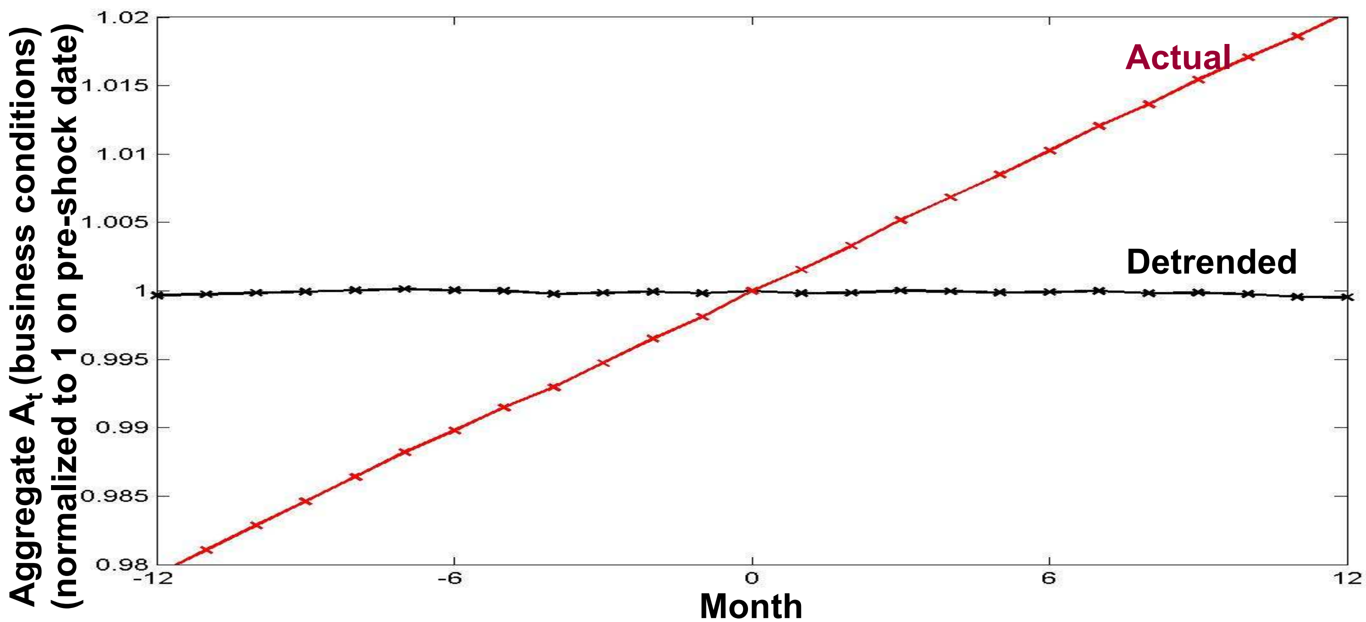

Notes: Simulations run on 1000 units. This is repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs are taken from the "All" values in table 3. All other parameters and assumptions as outlined in sections 3 and 4. The aggregate figure for A_t (business conditions) is calculated by summing up across all units within the simulation. This is detrended by removing its long run growth rate. The month is normalized to zero at the date of the uncertainty shock.

Figure 8: Aggregate (detrended) labor drops, rebounds and overshoots

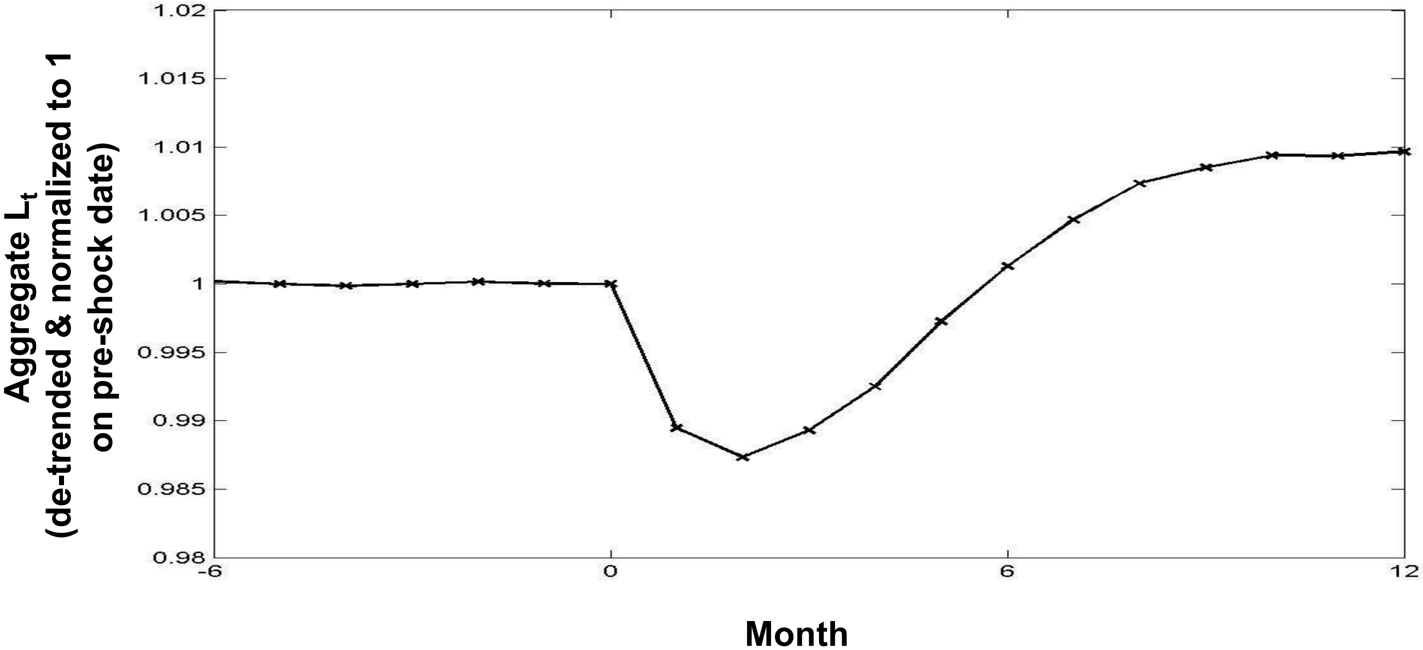

Notes: Simulation run on 1000 units. This is repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs are taken from the "All" values in table 3. All other parameters and assumptions as outlined in sections 3 and 4. The aggregate figures for L_t are calculated by summing up across all units within the simulation. They are detrended by removing their long-run growth rate. The month is normalized to zero at the date of the uncertainty shock.

Figure 9: Splitting out the uncertainty and volatility effects

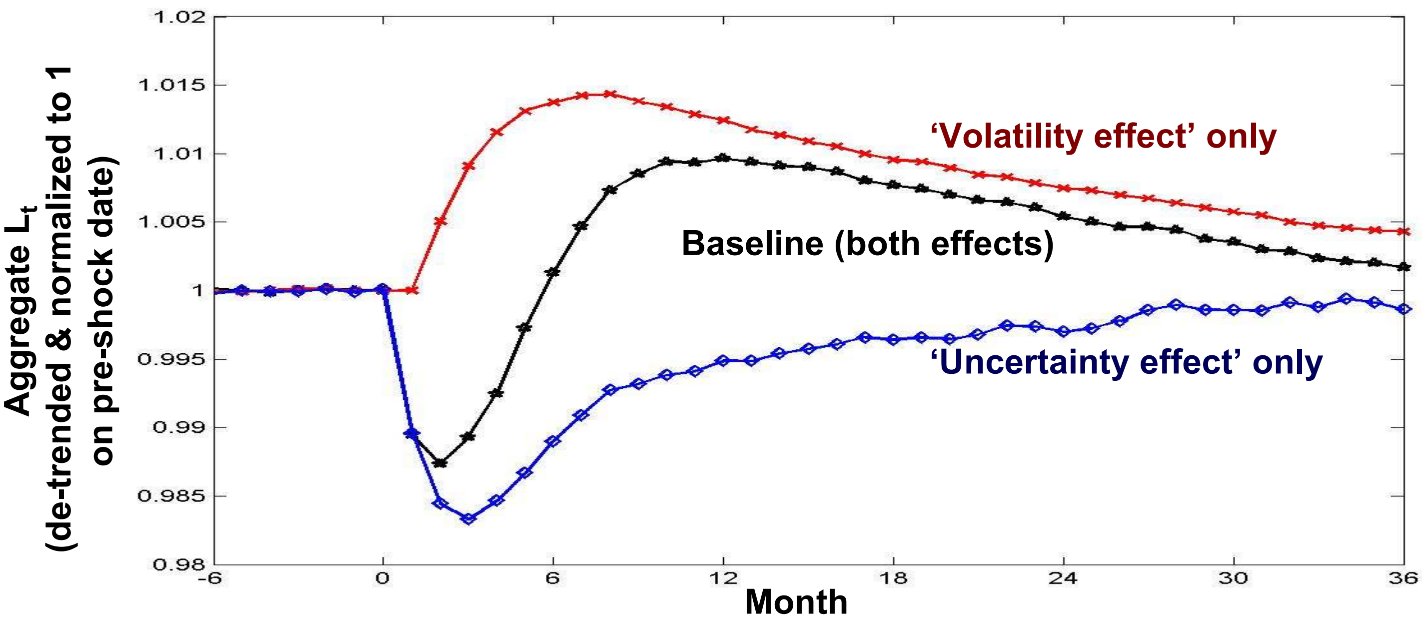

Notes: Simulation run on 1000 units. This is repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs are taken from the "All" values in table 3. All other parameters and assumptions as outlined in sections 3 and 4 for the **baseline** plot (which plots the same figure as in Figure 8 but extended out for 36 months). For the **volatility effect only** plot firms have expectations set to $\sigma_t=\sigma_L$ in all periods (i.e. uncertainty effects are turned off), while in the **uncertainty effect only** they have the actual shocks drawn from a distribution $\sigma_t=\sigma_L$ in all periods (i.e. the volatility effects are turned off).

more units to hit both the hiring and firing thresholds. Since more units are clustered around the hiring threshold then the firing threshold due to labor attrition and business conditions growth (see Figure 6) this leads to a medium term burst of net hiring. In effect hiring is convex in productivity just below the hiring threshold - firms that receive a small positive shock hire and firms that receive a small negative shock do not respond. So total hiring rises in the medium term with the increased volatility of productivity growth. Of course once firms have undertaken a burst of hiring they jump to the interior of the region of inaction and so do not hire again for some time. So in the long-run this results in labor falling back to its long-run trend path. I label this phenomenon the '*volatility overshoot*', since this medium-term hiring boom is induced by the higher unit-level volatility of business conditions shocks.[39]

Thus, the effect of a rise in σ_t is two fold. First, the real-options impact effect from increased uncertainty over future business conditions, which causes an initial drop in activity as firms pause investment and hiring. This happens rapidly since expectations change upon impact of the uncertainty shock, so that hiring and investment instantly freeze. Second, the effect from increased volatility of realized business conditions, which causes a medium term hiring-boom. This takes more time to occur because this is driven by the rise in the realized volatility of productivity growth. This rise is volatility accrues over several months. Because the higher uncertainty temporarily moves out the hiring and investment thresholds this further delays the volatility overshoot. Thus, the uncertainty drop always precedes the volatility overshoot.

These distinct uncertainty and volatility effects are shown in Figure 9. This splits out the expectational effects of higher σ_t from the realized volatility effects of higher σ_t. These simulations are shown for 36 months after the shock to highlight the long-run differences between these effects.[40] The 'uncertainty effect' is simulated by allowing firms expectations over σ_t to change after the shock (as in the baseline) but holds the variance of the actual draw of shocks constant. This generates a drop and rebound back to levels, but no volatility overshoot. The 'volatility effect' is simulated by holding firms expectations over σ_t constant but allowing the realized volatility of the business conditions to change after the shock (as in the baseline). This generates a volatility overshoot, but no initial drop in activity from a pause in hiring.[41] The baseline figure in the graph is simply the

[39]Another way to think about this is the cross-sectional distribution of firms changes in the medium term from the ergodic steady-state to one with a lower average A/L ratio. This is because units are more evenly distributed between the hiring and firing thresholds after the increased volatility (rather than clustered up around the hiring threshold). In the longer run this settles back down to its ergodic distribution, bringing the A/L ratio back up to its steady state. Interestingly, given the fixed costs in hiring this medium-term burst in activity also generates echo effects in future as L settles back down towards its long-run trend.

[40]In general I plot response for the first 12 months due to the partial equilibrium nature of the analysis, unless longer-run plots are expositionally helpful.

[41]In the figure the volatility effects also take 1 extra month to begin. This is simply because of the standard finance timing assumption in (3.4) that σ_{t-1} drives the volatility of $A_{j,t}$. Allowing volatility to be driven by σ_t delivers similar

aggregate detrended labor (as in Figure 8). This suggests that uncertainty and volatility have very different effects on economic activity, despite often being driven by the same underlying phenomena.

The response to aggregate capital to the uncertainty shock is similar to labor. Capital also displays a short-run fall as firms postpone investing, followed by a rebound as they address their pent-up demand for investment, with a subsequent medium-run overshoot as the additional volatility generates a burst of investment (see Appendix Figure A6).

7.3. Why Uncertainty Reduces Productivity Growth

Figure (10a) plots the time series for the growth of 'Aggregate productivity', defined as $\sum_j A_{j,t}L_{j,t}$ where the sum is taken over all j production units in the economy in month t. In this calculation the growth of business-conditions ($A_{j,t}$) can be used as a proxy for the growth of productivity under the assumption that shocks to demand are small in comparison to productivity (or that the shocks are independent). Following Baily, Hulten and Campbell (1992) I define three indices as follows[42]:

$$\underbrace{\frac{\sum A_{j,t}L_{j,t} - \sum A_{j,t-1}L_{j,t-1}}{\sum A_{j,t-1}L_{j,t-1}}}_{\text{Aggregate Productivity Growth}} = \underbrace{\frac{\sum (A_{j,t} - A_{j,t-1})L_{j,t-1}}{\sum A_{j,t-1}L_{j,t-1}}}_{\text{Within Productivity Growth}} + \underbrace{\frac{\sum A_{j,t}(L_{j,t} - L_{j,t-1})}{\sum A_{j,t-1}L_{j,t-1}}}_{\text{Reallocation Productivity Growth}}$$

The first term, "Aggregate Productivity Growth", is the increase in productivity weighted by employment across units. This can be broken down into two sub-terms: "Within Productivity Growth" which measures the productivity increase within each production unit (holding the employment of each unit constant), and "Reallocation Productivity Growth" which measures the reallocation of employment from low to high productivity units (holding the productivity of each unit constant).

In Figure 10a 'Aggregate Productivity Growth' shows a large fall after the uncertainty shock, dropping to around 15% of its value. The reason is that uncertainty reduces the shrinkage of low productivity firms and the expansion of high productivity firms, reducing the reallocation of resources towards more productive units.[43] This reallocation from low to high productivity units drives the majority of productivity growth in the model so that higher uncertainty has a first-order effect on productivity growth. This is clear from the decomposition which shows that the fall in Total is entirely driven by the fall in the "Reallocation" term. The "Within" term is constant since, by assumption, the first moment of the demand conditions shocks is unchanged.[44] In the bottom two panels (Figures

results because in the short-run the uncertainty effect of moving out the hiring and investment thresholds dominate.

[42] Strictly speaking Bailey, Hulten and Campbell (1992) defined four terms, but for simplicity I have combined the 'between' and 'cross' terms into a 'reallocation' term.

[43] Formally there is no reallocation in the model because it is partial equilibrium. However, with the large distribution of contracting and expanding units all experiencing independent shocks, gross changes in unit factor demand are far larger than net changes, with the difference equivalent to "reallocation".

[44] These plots are not completely smooth because the terms are summations of functions which are approximately squared in $A_{j,t}$. For example $A_{j,t}L_{j,t} \approx \lambda A_{j,t}^2$ for some scalar λ since $L_{i,t}$ is approximately linear in $A_{j,t}$. Combined

Figure 10a: 'Aggregate productivity' *growth* falls & rebounds after the uncertainty shock

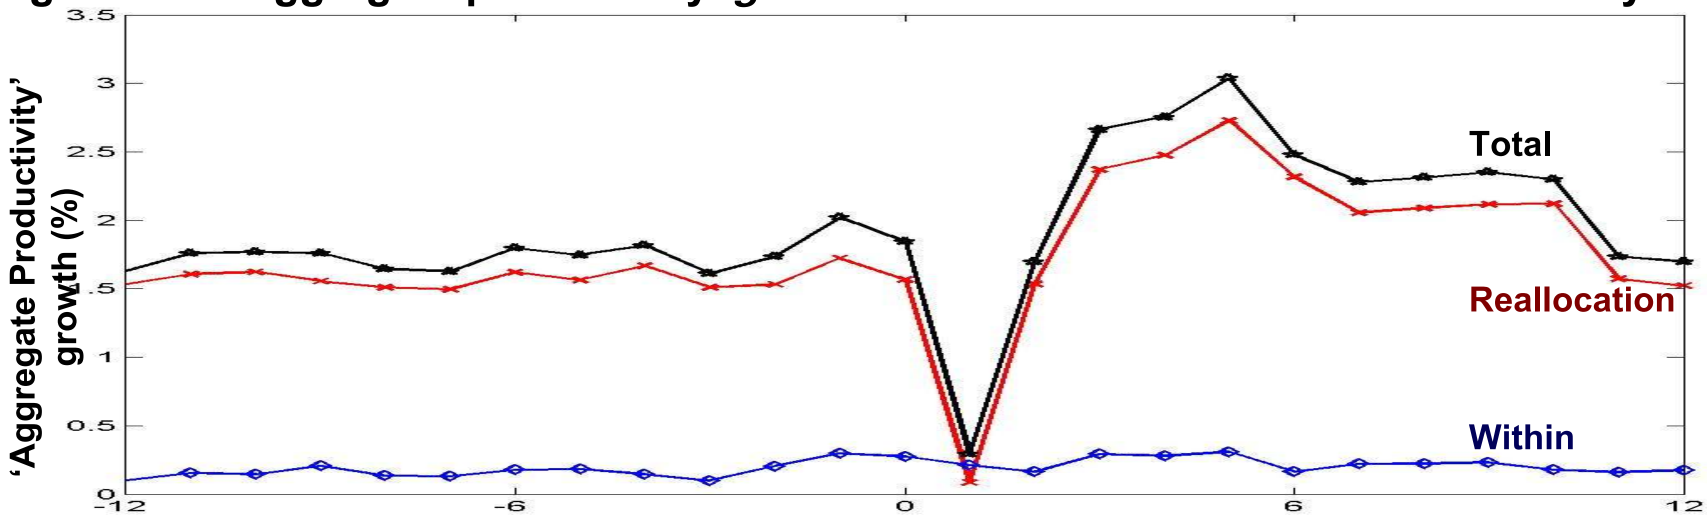

Figure 10b: 'Productivity' and hiring, period before the uncertainty shock

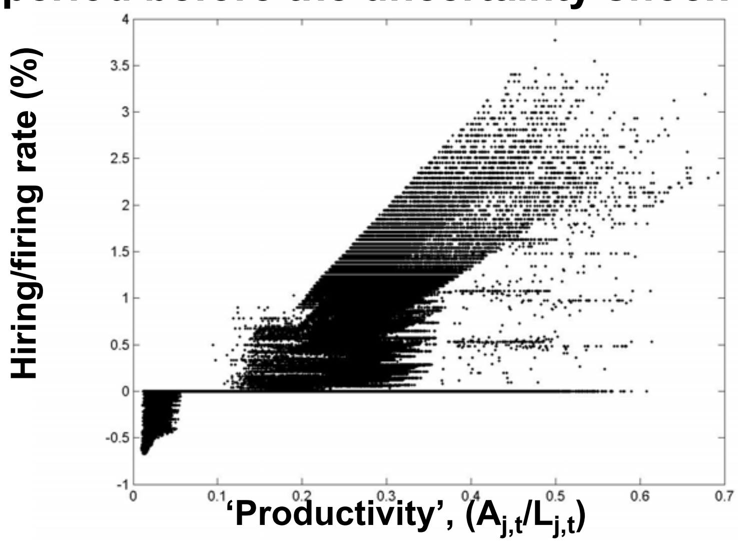

Figure 10c: 'Productivity' and hiring, period after the uncertainty shock

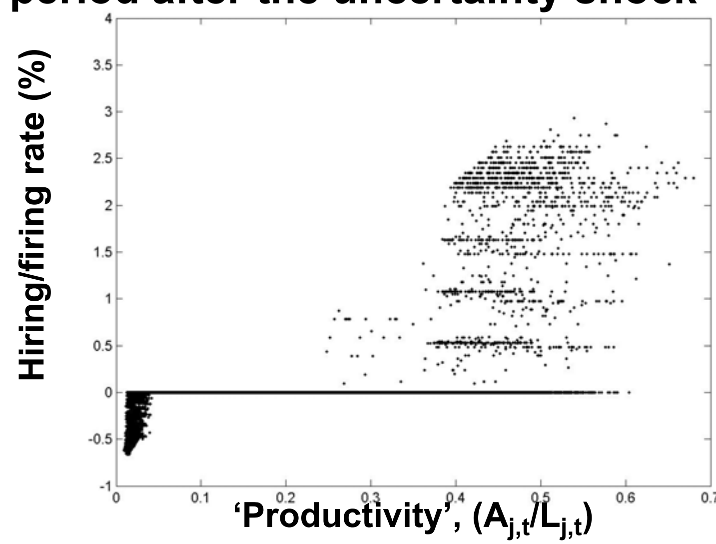

Notes: Simulations run on 1000 units. This is repeated 50000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs are taken from the "All" values in table 3. All other parameters and assumptions as outlined in sections 3 and 4. **'Aggregate Productivity' =** $\sum L_{j,t}A_{j,t}/\sum L_{j,t}$, where $A_{j,t}$ is unit level business conditions and $L_{j,t}$ is unit level employment. The summation is taken across all units in the simulation. 'Within' is defined as the productivity growth achieved holding unit size constant and 'Reallocation' is defined as the productivity growth requiring a change in unit size. In bottom panel unit-level business conditions ($A_{j,t}$) is used as proxy for productivity as discussed in section 6.1. The month is normalized to zero at the uncertainty shock.

10b and 10c) this reallocative effect is illustrated by two unit-level scatter plots of gross hiring against log productivity in the month before the shock (left-hand plot) and the month after the shock (right-hand plot). It can be seen that after the shock much less reallocative activity takes place with a substantially lower fraction of expanding productive units and shrinking unproductive units. Since actual US aggregate productivity growth appears to be 70% to 80% driven by reallocation[45] these uncertainty effects should play an important role in the real impact of large uncertainty shocks.

In Figure 11 plots the *level* of an alternative productivity measure, "Solow productivity". This is defined as aggregate output divided by factor share weighted aggregate inputs

$$\text{Solow productivity} = \frac{\sum_j A_{j,t}^{1/(\epsilon-1)} K_{j,t}^{\alpha}(L_{j,t} \times H_{j,t})^{1-\alpha}}{(\alpha \sum_j K_{j,t} + (1-\alpha) \sum L_{j,t} \times H_{j,t}}$$

I report this series because macro productivity measures are typically calculated in this way using only macro data (note the previous 'Aggregate Productivity' measure would require micro-data to calculate).[46] As can be seen in Figure 11 the detrended "Solow productivity" series also falls and rebounds after the uncertainty shock. Again, this initial drop and rebound is because of the initial pause and subsequent catch-up in the level of reallocation across units immediately after the uncertainty-shock. The medium-run overshoot is again due to the increased level of cross-sectional volatility, which increases the potential for reallocation, leading to higher aggregate medium-term productivity growth.

Finally, Figure 12 plots the effects on an uncertainty shock on output. This shows a clear drop, rebound and overshoot, very similar to the behavior of the labor, capital and productivity. What is striking about Figure 12 is the similarity of the size, duration and time-profile of the simulated response of output to an uncertainty shock to the VAR results on actual data shown in Figure 2. In particular both the simulated and actual data show a drop of detrended activity of around 1% to 2% after about three months, a return to trend at around 6 months, and a longer-run gradual overshoot.

7.4. Investigating Robustness to General Equilibrium

Ideally I would set up my model within a General Equilibrium (GE) framework, allowing prices to endogenously change. This could be done, for example, by assuming agents approximate the cross-sectional distribution of firms within the economy using a finite set of moments, and then using these

with the random walk nature of the driving process (which means some individual units grow very large) this results in lumpy aggregate productivity growth even in very large samples of units.

[45] Foster, Haltiwanger and Krizan (2000 and 2006) report that reallocation, broadly defined to include entry and exit, accounts for around 50% of manufacturing and 90% of retail productivity growth. These figures will in fact underestimate the full contribution of reallocation since they miss the within establishment reallocation, which Bernard, Redding and Schott's (2006) results on product switching suggests could be substantial.

[46] 'Aggregate productivity' (the series shown in growth rates in Figure 10a) looks very similar to the detrended level of 'Solow productivity'. Note output is approximated by $A^{1/(\epsilon-1)}K^{\alpha}(LH)^{1-\alpha}$ since $A^{1/(\epsilon-1)} = \widetilde{A}B^{1/(\epsilon-1)}$.

Figure 11: 'Solow productivity' (detrended) drops, rebounds and overshoots

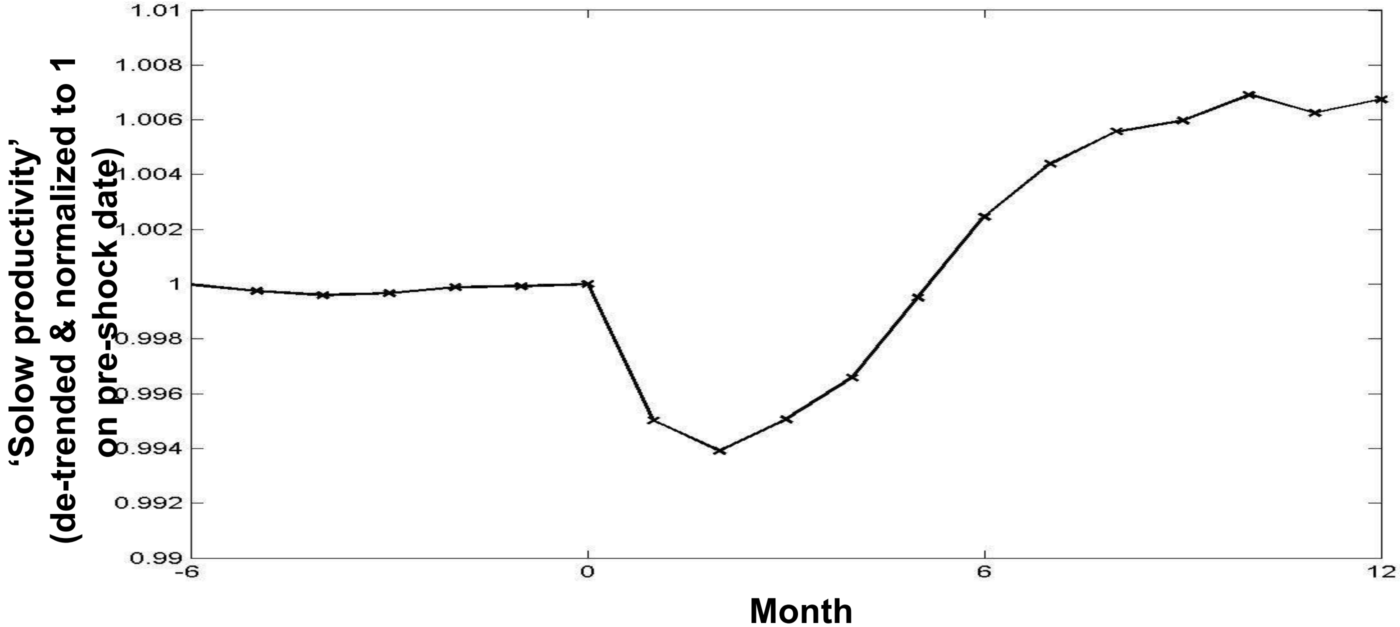

Figure 12: Aggregate (detrended) output drops, rebounds & overshoots

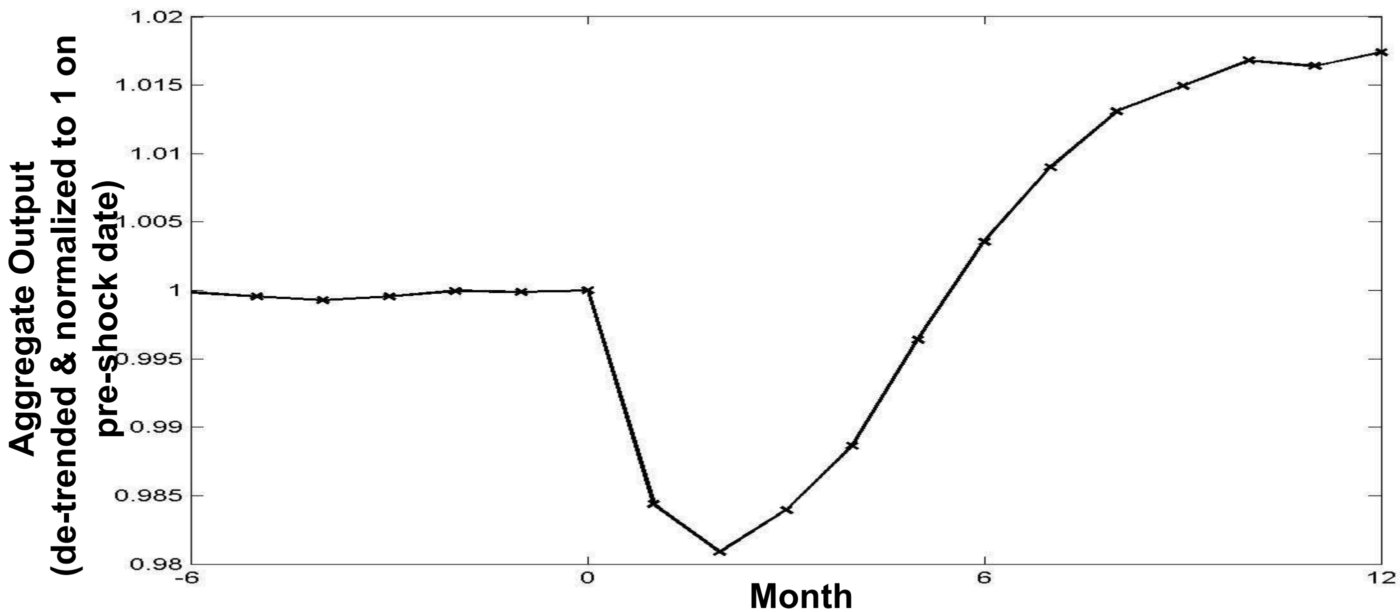

Notes: Simulations run on 1000 units. This is repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs are taken from the "All" values in table 3. All other parameters and assumptions as outlined in sections 3 and 4. **Solow productivity** defined as aggregate output divided by the factor share weighted aggregate inputs. Both series are detrended by removing their long-run growth rate. The month is normalized to zero at the uncertainty shock.

moments in a representative consumer framework to compute a recursive competitive equilibrium (see, for example, Krusell and Smith, 1998, Khan and Thomas, 2003, and Bachman, Caballero and Engel, 2006). However, this would involve another loop in the routine to match the labor, capital and output markets between firms and the consumer, making the program too slow to then loop in the Simulated Method of Moments estimation routine. Hence, there is a trade-off between two options: (1) a GE model with flexible prices but assumed adjustment costs[47], and (2) estimated adjustment costs but in a fixed price model. Since the effects of uncertainty are sensitive to the nature of adjustment costs I opted to take the second option and leave GE analysis to future work.

This means the results in this model could be compromised by GE effects if factor prices changed sufficiently to counteract factor demand changes.[48] One way to investigate this is to estimate the actual changes in wages, prices and interest rates arising after a stock-market volatility shock, and feed these into the model in an expectations consistent way. If these empirically plausible changes in factor prices radically changed these results this would suggest they are not robust to GE, while if they have only a small impact it is more reassuring on GE robustness.

To do this I use the estimated changes in factor prices from the VAR (see section 2.2), which are plotted in Figure 13. An uncertainty shock leads to a short-run drop and rebound of interest rates of up to 1.1% points (110 basis point), of prices of up to 0.5%, and of wages of up to 0.3%. I take these numbers and structurally build them into the model so that when $\sigma_t = \sigma_H$ interest rates are 1.1% lower, prices (of output and capital) are 0.5% lower and wages 0.3% lower. Firms expect this to occur, so expectations are rational.

In Figure 14 I plot the level of output after an uncertainty shock with and without these pseudo-GE prices changes. This reveals two surprising outcomes: first, the effects of these empirically reasonable changes in interest rates, prices and wages have very little impact on output in the *immediate* aftermath of an uncertainty shock; and second, the limited 'pseudo-GE' effects that do occur are greatest at around 3 to 5 months, when the level of uncertainty (and so the level of the

[47] Unfortunately there are no "off the shelf" adjustment cost estimates that can be used since no paper has previously jointly estimated convex and non-convex labor and capital adjustment costs. Furthermore, given the pervasive nature of temporal and cross-sectional aggregation in all firm and establishment level datasets, using one-factor estimates which also do not correct for aggregation may be problematic, especially for non-convex adjustment costs given the sensitivity of the lumpy behavior they imply to aggregation. This may explain the differences of up to 100 fold in the estimation of some of these parameters in the current literature.

[48] Kahn and Thomas (2003) find in their micro to macro investment model that GE effects cancel out most of the macro effects of non-convex adjustment costs on the *response to* shocks. With a slight abuse of notation this can be characterized as $\frac{\partial(\partial K_t/\partial A_t)}{\partial NC} \approx 0$ where K_t is aggregate capital, A_t is aggregate productivity and NC are non-convex adjustment costs. The focus of my paper on the direct impact of uncertainty on aggregate variables, is different and can be characterized instead as $\frac{\partial K_t}{\partial \sigma_t}$. Thus, their results are not necessarily inconsistent with mine.

More recent work by Bachman, Caballero and Engel (2006), however, finds the Kahn and Thomas (2003) results are sensitive to the choice of parameter values. Sim (2006) builds a GE model with capital adjustment costs and time varying uncertainty and finds that the impact of temporary increases in uncertainty on investment is robust to GE effects.

Figure 13: VAR estimation of the impact of a volatility shock on prices

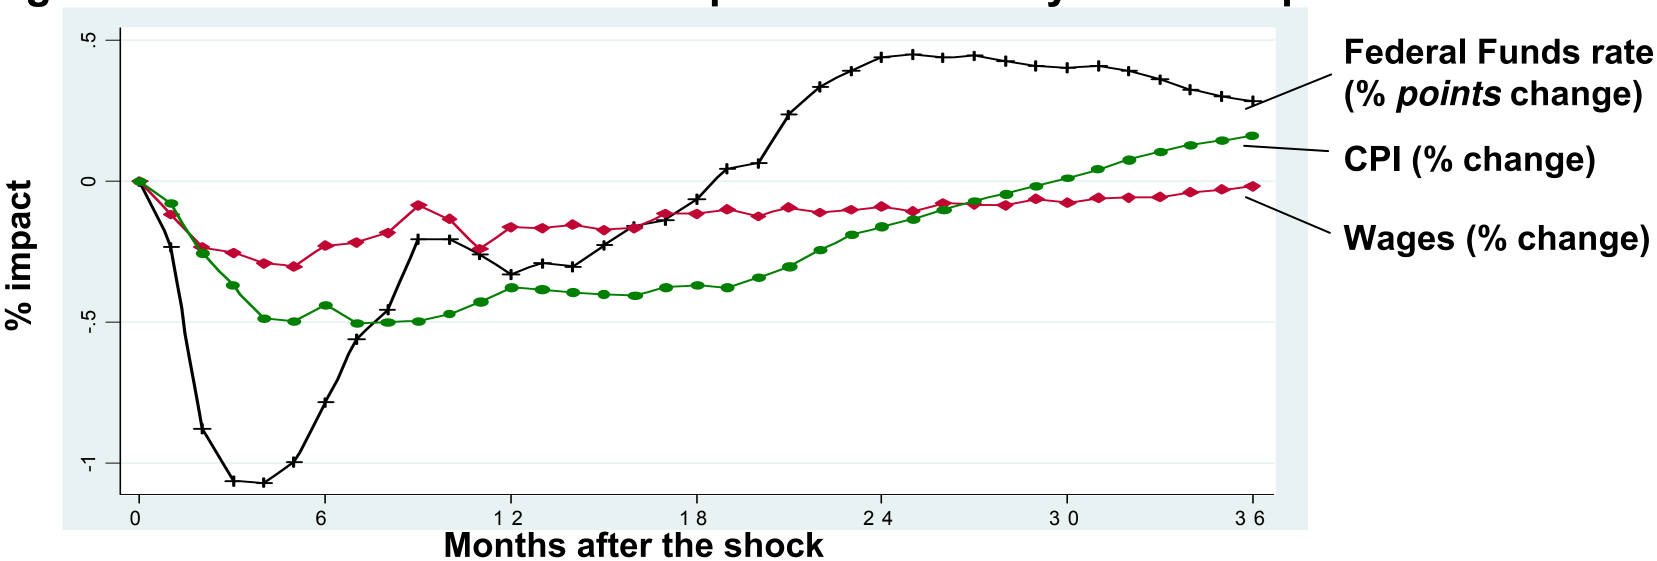

Notes: VAR Cholesky orthogonalized impulse response functions estimated on monthly data from July 1963 to July 2005 using 12 lags. Variables (in order) are log industrial production, log employment, hours, log wages, log CPI, federal funds rate, the volatility shock indicator and log S&P500 levels. Detrended using a Hodrick-Prescott filter with smoothing parameter of 129,600. Impact on the Federal Funds rate plotted as a percentage point change (so the shock reduces rates by up to 110 basis points) while the impact on the CPI and wages plotted as percentage changes.

Figure 14: Aggregate (detrended) output: partial-equilibrium and 'Pseudo-GE'

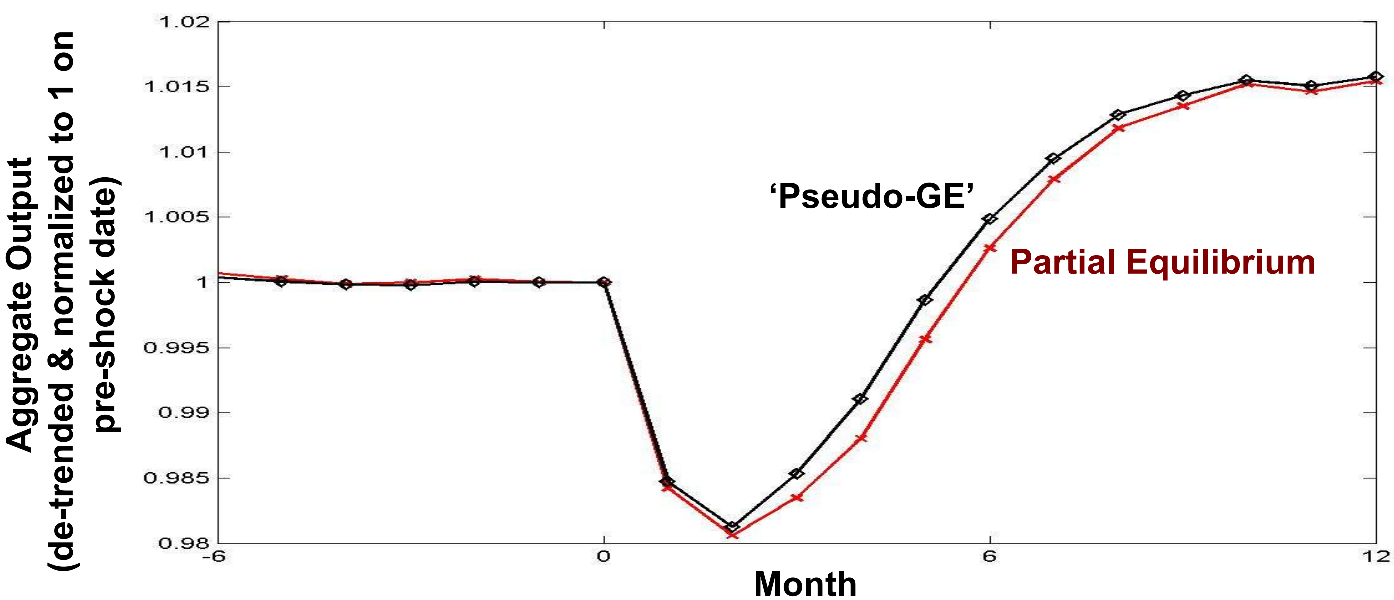

Notes: Simulations run on 1000 units. Repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs are taken from the "All" values in table 3. All other parameters and assumptions as outlined in sections 3 and 4. 'Pseudo-GE' allows interest rates, prices and wages to be 1.1% points, 0.5% and 0.3% lower during periods of high uncertainty. Series detrended by removing their long-run growth rate.

interest rate, price and wage reductions) are much smaller. To highlight the surprising nature of these two findings Figure A7 plots the impact of the 'pseudo-GE' price effects on capital, labor and output in a simulation *without* adjustment costs. In the absence of any adjustment costs these interest rate, prices and wages changes do have an extremely large effect. So the introduction of adjustment costs both dampens and delays the response of the economy to the 'pseudo-GE' price changes.

The reason for this limited impact of 'pseudo-GE' price changes is that after an uncertainty shock occurs the hiring/firing and investment/disinvestment thresholds jump out, as shown in Figure 5. As a result there are no units near any of the response thresholds. This makes the economy insensitive to changes in interest rates, prices or wages. The only way to get an impact would be to shift the thresholds back to the original low uncertainty position where the majority units are located. But as noted in section (4) the quantitative impact of these uncertainty shocks is equivalent to something like a 7% higher interest rate and a 25% higher wage rate, so these 'pseudo-GE' price reductions of 1.1% in interest rates, 0.5% in prices and 0.3% in wages are not sufficient to do this.

Of course once the level of uncertainty starts to fall back again the hiring/firing and investment/disinvestment thresholds begin to move back towards their low uncertainty values. This means they start to move back towards the region in (A/K) and (A/L) space where the units are located. So the economy becomes more sensitive to changes in interest rates, prices and wages. Thus, these 'pseudo-GE' price effects start to play a role. But this effect is limited by the fact that these prices effects are now reduced by the fall in uncertainty.

In summary, the rise in uncertainty not only reduces levels of labor, capital, productivity and output, but it also makes the economy temporarily extremely insensitive to changes in factor prices. This is the macro equivalent to the 'cautionary effects' of uncertainty demonstrated on firm-level panel data by Bloom, Bond and Van Reenen (2007).

For policymakers this is important since it suggests a monetary or fiscal response to an uncertainty shock is likely to have almost no impact in the immediate aftermath of a shock. But as uncertainty falls back down and the economy rebounds, it will become more responsive, so any response to policy will occur with a lag. Hence, a policymaker trying for example, to cut interest rates to counteract the fall in output after an uncertainty shock would find no immediate response, but a delayed response occurring when the economy was already starting to recover. This cautions against using first-moment policy levers to respond to the second-moment component of shocks, with policies aimed directly at reducing the underlying increased uncertainty likely to be far more effective.

7.5. A Combined First and Second Moment Shock

All the large macro shocks highlighted in Figure 1 comprise both a first and a second moment element, suggesting a more realistic simulation would analyze these together. This is undertaken in Figure 15, where the output response to a pure second moment shock (from Figure 12) is plotted alongside the output response to the same second moment shock with an additional first moment shock of -2% to business conditions.[49] Adding an additional first moment shock leaves the main character of the second moment shock unchanged - a large drop and rebound.

Interestingly, a first-moment shock on its own shows the type of slow response dynamics that the real data displays (see, for example, the response to a monetary shock in Figure 3). This is because the cross-sectional distribution of units generates a dynamic response to shocks.[50]

This rapid drop and rebound in response to a second moment shock is clearly very different to the persistent drop over several quarters in response to a more traditional first moment shock. Thus, to the extent a large shock is more a second moment phenomena - for example 9/ll - the response is likely to involve a rapid drop and rebound, while to the extent it is more a first moment phenomena - for example OPEC II - it is likely to generate a persistent slowdown. However, in the immediate aftermath of these shocks distinguishing them will be difficult, as both the first and second moment components will generate an immediate drop in employment, investment and productivity. The analysis in section (2.1) suggests, however, there are empirical proxies for uncertainty that are available real-time to aid policymakers, such as the VXO series for implied volatility (see notes to Figure 1), the cross-sectional spread of stock-market returns and the cross-sectional spread of professional forecasters.

Of course these first and second moment shocks differ both in terms of the moments they impact and also in terms of their duration, permanent and temporary respectively. The reason is that the second moment component of shocks is almost always temporary while the first moment component tends to be persistent. For completeness a persistent second moment shock would generate a similar effect on investment and employment as a persistent first moment shock, but would generate a slow-down in productivity *growth* through the "Reallocation" term rather than a one-time reduction in productivity *levels* through the "Within" term. Thus, the temporary/permanent distinction is important for the predicted time profile of the impact of the shocks on hiring and investment, and the first/second moment distinction is important for the route through which these shocks impact productivity.

[49] I choose 2% because this is equivalent to 1 years business conditions growth in the model. Larger or smaller shocks yield a proportionally larger or smaller impact.

[50] See the earlier work on this by, for example, Bertola and Caballero (1990, 1994) and Caballero and Engel (1993).

Figure 15: Combined first and second moment shocks

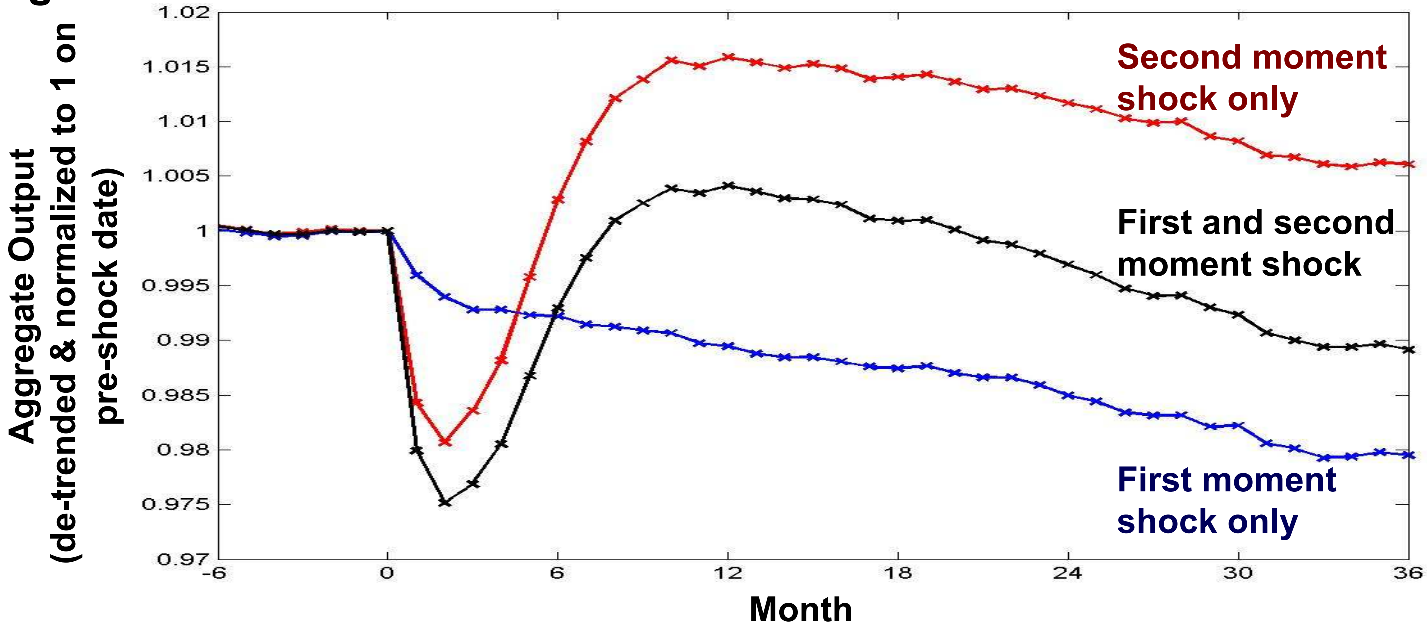

Figure 16: Different adjustment costs

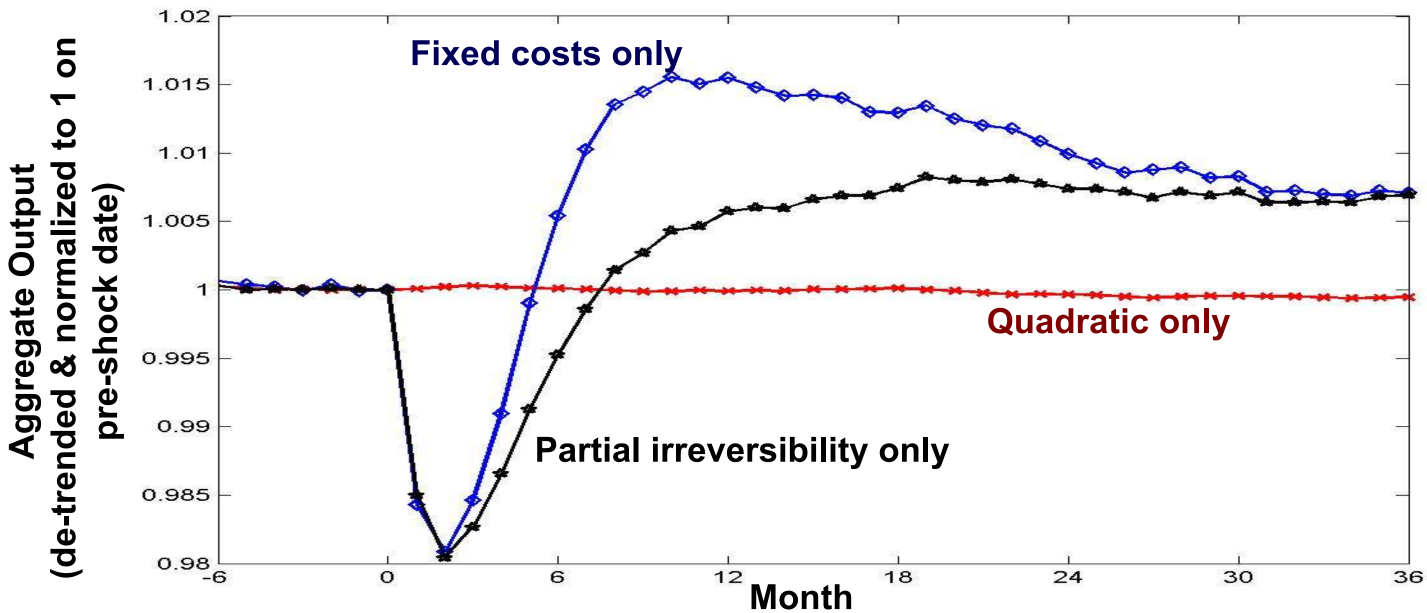

Notes: Simulations run on 1000 units. This is repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs in the top panel are taken from the "All" values in table 3. In the bottom panel the "fixed costs" specification has only the FC_K and FC_L adjustment costs from this estimation, the "partial irreversibility" has only the PR_K and PR_L from this specification, and the "Quadratic" has the adjustment costs from the "Quad" column in table 3. All other parameters and assumptions as outlined in sections 3 and 4. All series are detrended by removing their long-run growth rate. The month is normalized to zero at the uncertainty shock

The only historical example of a persistent second moment shock was the Great Depression, when uncertainty - as measured by share returns volatility - rose to an incredible 130% of 9/11 levels on average for the 4-years of 1929 to 1932. While this type of event is unsuitable for analysis using my model given the lack of general equilibrium effects and the range of other factors at work, the broad predictions do seem to match up with the evidence. Romer (1990) argues that uncertainty played an important real-options role in reducing output in the *onset* of the Great Depression, while Ohanian (2001) and Bresnahan and Raff (1991) report "inexplicably" low levels of productivity growth with an "odd" lack of output *reallocation* over this period.

7.6. Investigating Robustness to Different Parameter Values

7.6.1. Adjustment Costs

To evaluate the effects of different types of adjustment I ran three simulations: the first with fixed costs only, the second with partial irreversibilities only and the third with quadratic adjustment costs only.[51] The output from these three simulations is shown in Figure 16. As can be seen the two specifications with non-convex adjustment costs generate a distinct drop and rebound in economic activity. The rebound with fixed-costs is faster than with partial irreversibilities because of the bunching in hiring and investment, but otherwise they are remarkably similar in size, duration and profile. The quadratic adjustment cost specification appears to generate no response to an uncertainty shock. The reason is that there is no kink in adjustment costs around zero, so no option values associated with doing nothing.

In summary, this suggests the predictions are very sensitive to the inclusion of some degree of non-convex adjustment costs, but are much less sensitive to the type (or indeed level) of these non-convex adjustment costs. This highlights the importance of the prior step of estimating the size and nature of the underlying labor and capital adjustment costs.

7.6.2. Predefined Parameters

To investigate the robustness of the simulation results to the assumptions over the predefined parameters I re-ran the simulations using the different parameters from Table 5. The results, shown in Figure 17, highlight that the *qualitative* result of a drop and rebound in activity is robust to the different assumptions over the predetermined parameters. This is because of the presence of some non-convex component in all the sets of estimated adjustment costs in Table 5.

The size of this drop and rebound did vary across specifications, however. Running the simulation

[51] For fixed costs and partial irreversibilities the adjustment costs are the fixed-cost and partial irreversibility components of the parameter values from the "All" column in Table 3. For quadratic adjustment costs the values are from the "Quad" column in Table 3.

Figure 17: Simulation robustness to different parameter assumptions

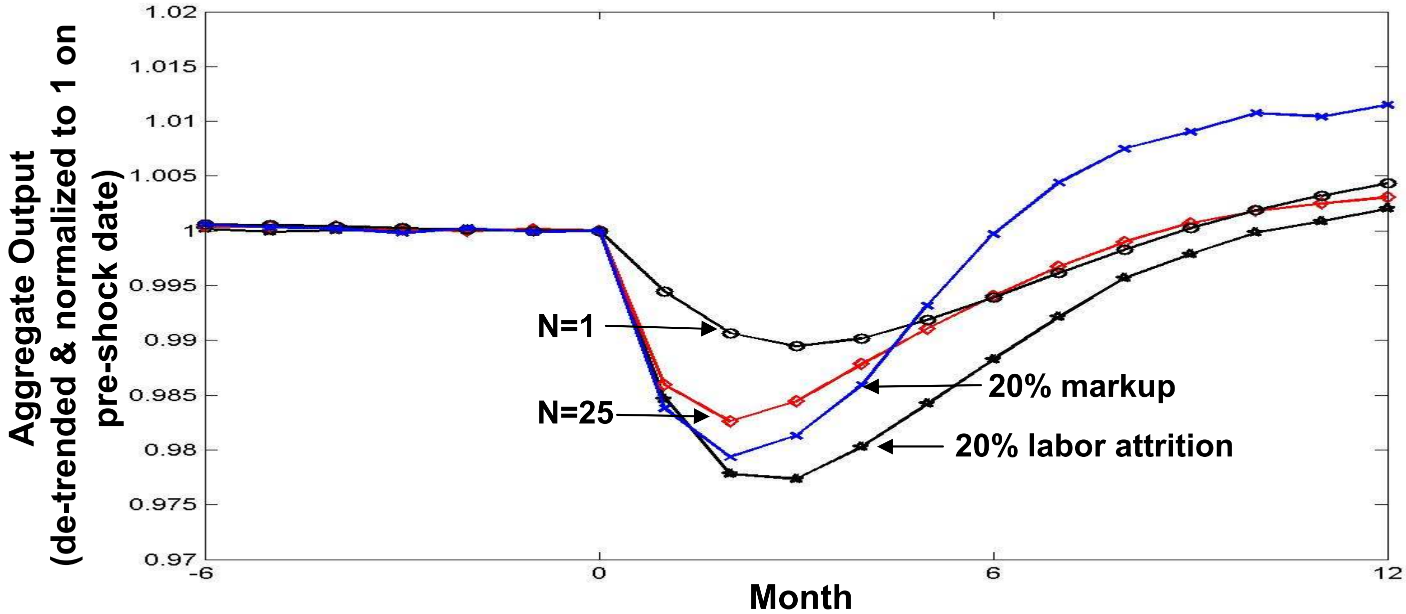

Notes: Simulations run on 1000 units. This is repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Parameter values are taken from the different columns of table 5, as indicated by the labels. All other parameters and assumptions as outlined in sections 3 and 4. The aggregate figures for output are calculated by summing up across all units within the simulation. They are detrended by removing their long-run growth rate. The month is normalized to zero at the date of the uncertainty shock.

with the "N=1" parameter estimates from Table 5 leads to drop of only 1%, about half the baseline drop of about 1.8%. This smaller drop was due to the very high levels of estimated quadratic adjustment costs that were required to smooth the investment and employment series in the absence of cross-sectional aggregation. Of course, the assumption of no cross-sectional aggregation ("N=1") is inconsistent with the aggregation evidence in Appendix A3 and the typical multi-divisional structure of large-firms. This simulation is presented simply to highlight the importance of building aggregation into estimation routines when it is also present in the data.

In the "δ_L=0.2" specification the drop was around 2.25%, about 30% above the baseline drop, due to the greater labor attrition after the shock. Hence, this more realistic assumption on 20% annual labor attrition (rather than 10% in the baseline) generates a larger drop and rebound in activity. The results for assuming partial cross-sectional aggregation ("N=25") and a 20% mark-up ("a+b=0.83") are both pretty similar to the baseline simulation (which has full cross-sectional aggregation and a 33% mark-up).

7.6.3. Durations and Sizes of Uncertainty Shocks

Finally, I also evaluate the effects of robustness of the simulation predictions to different durations and sizes of uncertainty shocks. In Figure 18 I plot the output response to a shorter-lived shock (a 1 month shock half-life) and a longer-lived shock (a 6 month shock half-life). Also plotted is the baseline (a 2.month shock half-life). It is clear that longer-lived shocks generate larger and more persistent falls in output. The reason is that the pause in hiring and investment lasts for longer if the rise in uncertainty is more persistent. Of course, because the rise in uncertainty is more persistent the cumulative increase in volatility is also larger so that the medium term 'volatility-overshoot' is also greater. Hence, more persistent uncertainty shocks generate a larger and more persistent drop, rebound and overshoot in activity. This is interesting in the context of the Great Depression, a period in which uncertainty rose to 260% of the baseline level for over 4-years, which in my (partial equilibrium) model would generate an extremely large and persistent drop in output and employment.

In Figure 19 I plot the output response to a smaller uncertainty shock ($\sigma_H = 1.5\times\ \sigma_L$), a larger uncertainty shock ($\sigma_H = 3\times\ \sigma_L$) and the baseline uncertainty shock ($\sigma_H = 2\times\ \sigma_L$). Surprisingly, the three different sizes of uncertainty shock lead to similar sized drops in activity. The reason is that real-option values are increasing, but concave, in the level of uncertainty.[52] So the impact of a 50% rise in uncertainty on the hiring and investment thresholds is about 2/3 of the size of the baseline 100% rise in uncertainty. Since the baseline impact on the hiring and investment thresholds is so large, even 2/3 of this pauses almost all hiring and investment. What is different across the different

[52] See Dixit (1993) and Abel and Eberly (1996) for an analytical derivation and discussion.

Figure 18: Different durations of uncertainty shocks

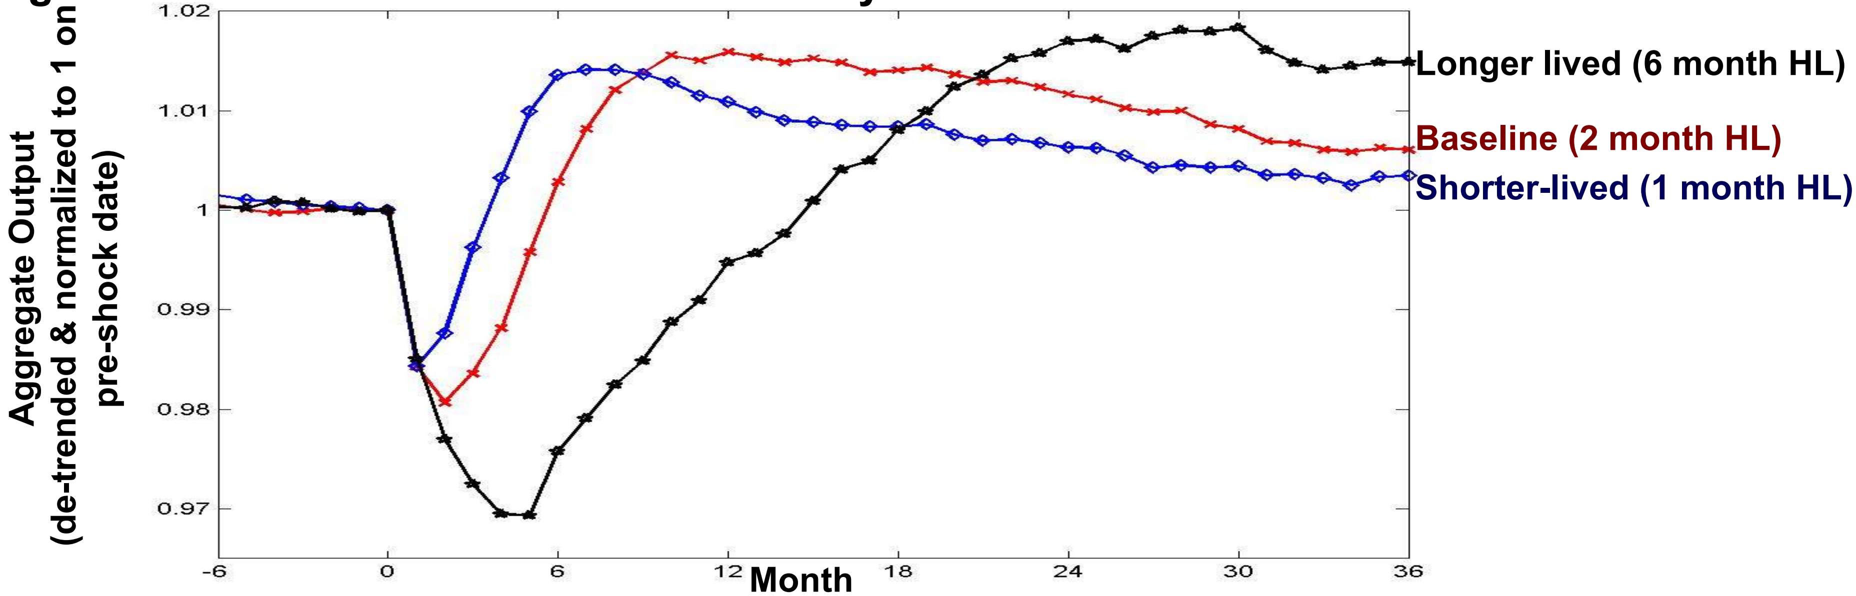

Figure 19: Different sizes of uncertainty shocks

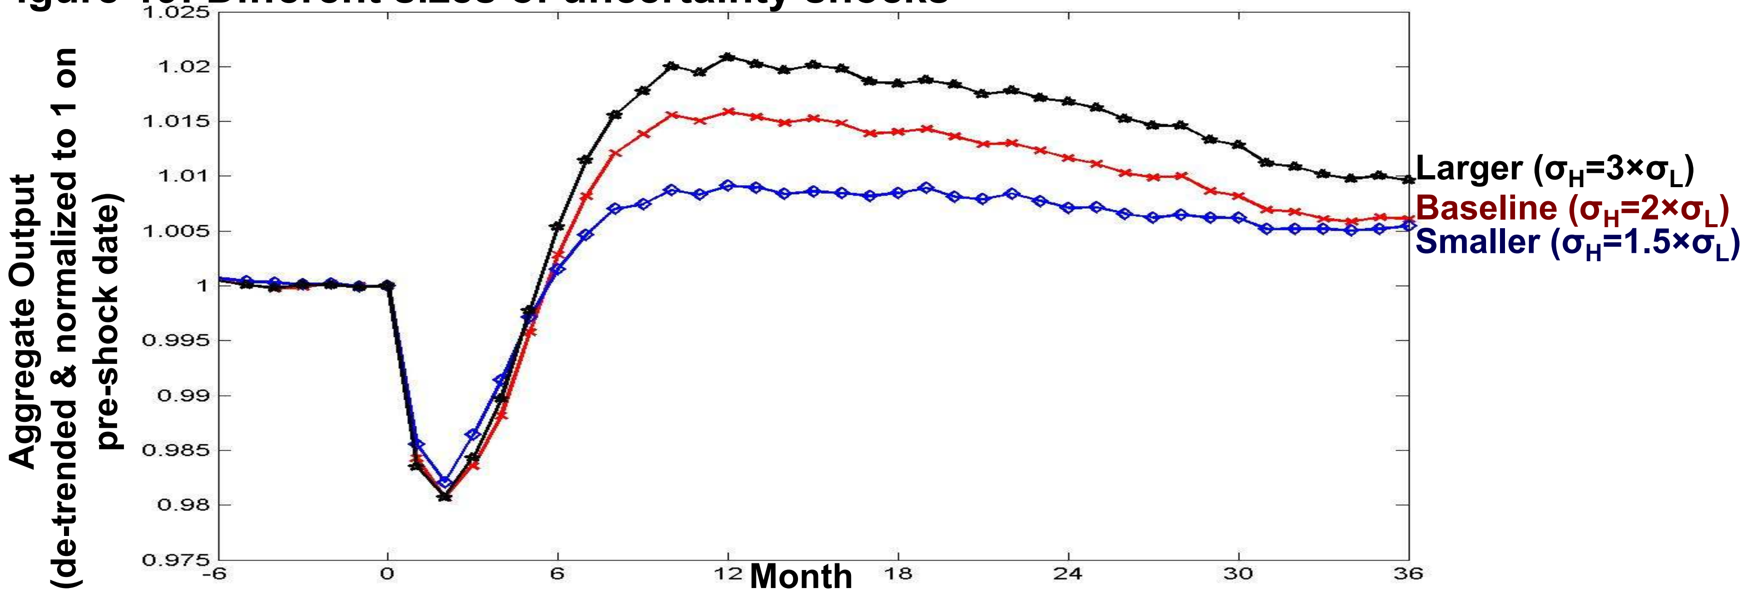

Notes for both figures: Simulations run on 1000 units, repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs taken from the "All" values in table 3. All other parameters and assumptions as outlined in sections 3 and 4. In the top panel the shorter and longer duration uncertainty shocks have **half-lives (HL)** of 1 month and 6 months respectively (baseline is 2 months). In the lower figure the larger and smaller uncertainty shocks have values of σ_H equal to 150% and 300% of σ_L level (baseline is 200%).

sizes of shocks, however, is that larger uncertainty shocks generate a larger medium-term 'volatility overshoot' because the cumulative increase in volatility is greater.

8. Conclusions

Uncertainty appears to dramatically increase after major economic and political shocks like the Cuban Missile crisis, the assassination of JFK, the OPEC I oil-price shock and the 9/11 terrorist attacks. If firms have non-convex adjustment costs these uncertainty shocks will generate powerful real-option effects, driving the dynamics of investment and hiring behavior. These shocks appear to have large real effects, with the uncertainty component alone generating a 1% drop and rebound in employment and output over the following six months, with a milder long-run overshoot.

This paper offers a structural framework to analyze these types of uncertainty shocks, building a model with a time varying second moment of the driving process and a mix of labor and capital adjustment costs. The model is numerically solved and estimated on firm level data using simulated method of moments. The parameterized model is then used to simulate a large macro *uncertainty* shock, which produces a rapid drop and rebound in output, employment and productivity growth. This is due to the effect of higher uncertainty making firms temporarily pause their hiring and investment behavior. In the medium term the increased volatility arising from the uncertainty shock generates a 'volatility-overshoot' as firms respond to the increased variance of productivity shocks, which drives a medium term overshoot and longer-run return to trend. Hence, the simulated response to uncertainty shocks generates a drop, rebound and longer-run overshoot, much the same as their actual empirical impact.

This temporary impact of a second moment shock is different from the typically persistent impact of a first moment shock. While the second moment effect has its biggest drop by month 3 and has rebounded by about month 6, persistent first moment shocks generate falls in activity lasting several quarters. Thus, for a policy-maker in the immediate aftermath of a shock it is critical to distinguish the relative contributions of the first and second moment component of shocks for predicting the future evolution of output.

The uncertainty shock also induces a strong insensitivity to other economic stimulus. At high levels of uncertainty the real-option value of inaction is very high, which makes firms extremely cautious. As a result the effects of empirically realistic General Equilibrium type interest rate, wage and price falls have a very limited short-run effect on reducing the drop and rebound in activity. This raises a second policy implication, that in the immediate aftermath of an uncertainty shock monetary or fiscal policy is likely to be particularly ineffective.

This framework also enables a range of future research. Looking at individual events it could be used, for example, to analyze the uncertainty impact of major deregulations, tax changes, trade reforms or political elections. It also suggests there is a trade-off between policy "correctness" and "decisiveness" - it may be better to act decisively (but occasionally incorrectly) then to deliberate on policy, generating policy-induced uncertainty. For example, when the Federal Open Markets Committee was discussing the negative impact of uncertainty after 9/11 it noted that "*A key uncertainty in the outlook for investment spending was the outcome of the ongoing Congressional debate relating to tax incentives for investment in equipment and software. Both the passage and the specific contents of such legislation remained in question*" (November 6th, 2001). Hence, in this case Congress's desire to revive the economy with tax incentives might have been counter-productive due to the increased uncertainty the policy process induced.

More generally these second moments effects contribute to the "*where are the negative productivity shocks?*" debate in the business cycle literature. It appears that second-moment shocks can generate short sharp drops and rebounds in output, employment, investment and productivity growth without the need for a first-moment productivity shock. Thus, recessions could potentially be driven by increases in uncertainty. Encouragingly, recessions do indeed appear in periods of significantly higher uncertainty, suggesting an uncertainty approach to modelling business-cycles (see Bloom, Floetotto and Jaimovich, 2007). Taking a longer run perspective this model also links to the volatility and growth literature given the evidence for the primary role of reallocation in productivity growth.

The paper also jointly estimates non-convex and convex labor and capital adjustment costs. I find substantial fixed costs of hiring/firing and investment, a large loss from capital resale and some moderate per-worker hiring/firing costs. I find no evidence for quadratic investment or hiring/firing adjustment costs. I also find that assuming capital adjustment costs only - as is standard in the investment literature - generates an acceptable overall fit, while assuming labor adjustment costs only - as is standard in the labor demand literature - produces a poor fit.

A. Appendix: Data

All data and Stata do files used to create the empirical Figures 1, 2, 3 and Table 1 are available on http://www.stanford.edu/~nbloom/. In this Appendix I describe the contents and construction of these datasets.

A.1. Stock Market Volatility Data

A.1.1. Testing for Jumps in Stock Market Volatility

To test for jumps in stock-market volatility I use the non-parametric bipower variation test of Barndorff-Nielsen and Shephard (2006). The test works for a time series $\{x_t,\ t = 1, 2, .N\}$ by comparing the squared variation, $SV = \sum_{t=3}^{N}(x_t - x_{t-1})^2$ with the bipower variation, $BPV = \sum_{t=3}^{N}(x_t - x_{t-1})(x_{t-1} - x_{t-2})$. In the limit as $dt \to 0$ if there are no jumps in the data then $E[SV] \to E[BPV]$ since the variation is driven by a continuous process. If there are jumps, however, then $E[SV] > E[BPV]$ since jumps have a squared impact on SV but only a linear impact on BPV. Barndorff-Nielsen and Shephard (2006) suggest two different test statistics - the *linear-jump* and *ratio-jump* test - which have the same asymptotic distribution but different finite-sample properties. Using the monthly data from Figure 1 I reject the null of no jumps at the 2.2% and 1.6% level using the linear and ratio tests, respectively. Using the daily VXO data underlying Figure 1 (available from January 1986 onwards, providing 5443 observations) I reject the null of no-jumps using both test at the 0.0% level.

A.1.2. Defining Stock Market Volatility Shocks

Given the evidence for the existence of stock-market volatility jumps I need to define what these are. The main measure is an indicator that takes a value of 1 for each of the sixteen events labelled in Figure 1, and 0 otherwise. These sixteen events are chosen as those with stock-market volatility more than 1.65 standard-deviations above the Hodrick Prescott detrended ($\lambda = 129,600$) mean of the stock market volatility series (the raw undetrended series is plotted in Figure 1). While some of these shocks occur in one month only, others span multiple months so there was a choice over the exact allocation of their timing. I tried two different approaches, the primary approach is to allocate each event to the month with the largest volatility spike for that event, with an alternative approach to allocate each event to the first month in which volatility went more than two standard-deviations above the HP detrended mean. The events can also be categorized in terms of terror, war, oil or economic shocks. So a third volatility indicator uses only the arguably most exogenous terror, war and oil shocks.

The volatility shock events, their dates under each timing scheme and their classification are shown in Table (A.1) below, while in section (A.1.3) below each event is described in more detail.It is noticeable from Table (A.1) that almost all the shocks are bad events. So one question for empirical identification is how distinct are stock-market volatility shocks from stock-market levels shocks? Fortunately, it turns-out these series do move reasonably independently because some events - like the Cuban Missile crisis - raise volatility without impacting stock-market levels while others - like Hurricane Katrina - generate falls in the stock-market without raising volatility. So, for example, the log detrended stock-market level has a correlation of -0.192 with the main 1/0 volatility shock indicator, a correlation of -0.136 with the 1/0 oil, terror and war shock indicator, and a -0.340 correlation with the log detrended volatility index itself. Thus, the impact of stock-market volatility can be separately identified from stock-market levels.

A.1.3. Details of the Volatility Events

I briefly describe each of the sixteen volatility shocks shown on Figure 1 to highlight the fact that these are typically linked to real shocks.

Cuban Missile Crisis: The crisis began on October 16, 1962 when U.S. reconnaissance planes discovered Soviet nuclear missile installations on Cuba. This led to a twelve day stand-off between U.S. President John F. Kennedy and Soviet premier Nikita Khrushchev. The crisis ended on October 28 when the Soviets announced that the installations would be dismantled. The Cuban Missile Crisis is often regarded as the moment when the Cold War came closest to escalating into a nuclear war.

Assassination of JFK: President Kennedy was assassinated in Dallas on November 22, 1963, while on a political trip through Texas. Lee Harvey Oswald was arrested 80 minutes later for the

Table A.1: Major Stock-Market Volatility Shocks.

Event	Max Volatility	First Volatility	Type
Cuban Missile Crisis	October 1962	October 1962	Terror
Assassination of JFK	November 1963	November 1963	Terror
Vietnam build-up	August 1966	August 1966	War
Cambodia and Kent State	May 1970	May 1970	War
OPEC I, Arab-Israeli War	December 1973	December 1973	Oil
Franklin National	October 1974	September 1974	Economic
OPEC II	November 1978	November 1978	Oil
Afghanistan, Iran Hostages	March 1980	March 1980	War
Monetary cycle turning point	October 1982	August 1982	Economic
Black Monday	November 1987	October 1987	Economic
Gulf War I	October 1990	September 1990	War
Asian Crisis	November 1997	November 1997	Economic
Russian, LTCM Default	September 1998	September 1998	Economic
9/11 Terrorist Attack	September 2001	September 2001	Terror
Worldcom and Enron	September 2002	July 2002	Economic
Gulf War II	February 2003	February 2003	War

assassination. Oswald denied shooting anyone and claimed that he was being set up. Oswald was fatally shot less than two days later in a Dallas police station by Jack Ruby.

Vietnam build-up: US troop numbers rose from 184,000 at the beginning of 1966 to almost 500,000 by the end of the year. This military build up caused considerable uncertainty over both the introduction of wage and price controls (following the precedent set in the Korean War) and the Budgetary implications of the escalating defense expenditure.[53]

Cambodia and Kent State: President Nixon, elected in 1968 promising to end the Vietnam war, announced on April 30, 1970, that the US had invaded Cambodia. This caused student protest across the US. On May 4th, 1970, four students were fatally shot by the Ohio National Guard during anti-war protest, followed by two more fatal shootings of student demonstrators at Jackson State by Mississippi State Police. This generated the largest national strike in US history, and considerable social unrest.

OPEC I and Arab-Israeli War: A coalition of Arab nations attacked Israel on October 6th, 1973, with the war lasting until October 26th. As a result of this conflict Arab members of OPEC plus Egypt and Syria stopped shipping petroleum to the US and Europe, and increased the price of Oil four-fold.

Franklin National Financial Crisis: On 8 October 1974, Franklin National bank was declared insolvent due to mismanagement and fraud, involving losses in currency speculation. It had been purchased in 1972 by Michele Sindona, a banker with close ties to the Italian Mafia. At the time the Franklin National was the 20th largest bank in the US, and its failure was the largest banking collapse in US history.

OPEC II: The Shah of Iran fled the country after the Iranian revolution brought the Ayatollah Khomeini to power. This severely damaged the Iranian oil sector, allowing OPEC to double oil-prices.

Afghanistan and Iran Hostages: On December 25th 1979 the Soviet Union invaded Afghanistan, generating uncertainty over whether the invasion would continue through into the oil-fields of neighboring Iran. The storming of the American embassy and capture of 66 diplomats and citizens, who were held hostage until January 1981 added to the political uncertainty.

Monetary cycle turning point: Market volatility appeared to stem from uncertainty over the timing of the recovery from the recession and the ability of the Reagan government to deliver its fiscal and supply-side policies.[54]

[53] For example Time Magazine (28/10/1966) reported "*Defense spending jumped by a startling $4.2 billion annual rate during the third quarter. So far, Defense Secretary Robert McNamara has been mum as to how much money he must have next year. Not until that fog lifts will the economy managers, or anybody else, be able to get a clear glimpse of the 1967 economy.*"

[54] Time Magazine (16/08/1982) reported "*On one point nearly everyone agreed: the chaotic trading and uncertainty were directly traceable to Washington's ongoing failure to slash the runaway federal deficits that triggered crippling interest rates in the first place.*"

Black Monday: A large stock-market crash on Monday October 19, 1987 in which the Dow Jones index fell by 22.8%. No major news or individual event was associated with the crash.

Gulf War I: On August 2nd 1990 Iraq invaded Kuwait. In response the US started deploying troops to Saudi Arabia. On January 12th, 1991, Congress authorized the use of military force in Kuwait by 52-47 in the Senate and by 259-183 in the House, the closest margin in authorizing force by Congress since the War of 1812. This knife-edge political support generated pre-invasion uncertainty around the US response.

Asian Crisis: On the 14 May 1997 the Thai Baht came under sustained speculative attack, leading to its devaluation on July 2 1997. This crisis spread (in varying degree) across Asia to the Philippines, Malaysia, South Korea, Indonesia, Singapore and Hong Kong.

Russian & LTCM Default: In August 17, 1998 Russia defaulted on its Ruble and domestic Dollar debt. This caused the Long Term Capital Management hedge fund to default on several billion dollars of financial contracts, threatening a major financial collapse.

9/11 Terrorist Attack: On the morning of September 11, 2001, al-Qaeda terrorists hijacked four planes, flying two into the towers of the World Trade Centre in Manhattan, one into the Pentagon and one into rural Pennsylvania (which crashed, presumed to be heading for the White House). This was followed by a wave of Anthrax letters which killed five people, initially also believed to be linked with al-Qaeda.

Worldcom and Enron: Enron, a major energy trading firm, filled for bankruptcy in December 2001 after admitting to the fabrication of its accounts. WorldCom, a large telecoms firm, announced in July 2002 that an internal audit had uncovered approximately $3.8 billion of overstated revenues. This was accompanied by a series of other accounting scandals involving major firms such as Tyco, AOL Time Warner, Bristol-Myers Squibb, Merck and Dynegy, casting doubt over the veracity of the accounts of many large firms.

Gulf War II: In October 2002 Congress gave the President Bush the authority to invade Iraq. The US worked to obtain UN approval for this, but by March 2003 it became clear this was not going to happen. On March 20, 2003, the US-led a small coalition force into Iraq. The period running up to this invasion generated substantial stock-market volatility over whether the UN would support the war, and if not whether President Bush would proceed without this support.

A.2. Cross-Sectional Uncertainty Measures

There are four key cross-sectional uncertainty measures:

Standard deviation of firm-level profits growth: This is measured on quarterly basis using Compustat Quarterly Accounts. It is the cross-sectional standard deviation of the growth rates of pre-tax profits (data item 23). Profit growth has a close fit to productivity and demand growth in homogeneous revenue functions[55], and is one of the few variables to have been continuously reported in quarterly accounts since the 1960s.[56] This is normalized by the firms average sales (data item 2), and defined as (profits$_t-$ profits$_{t-1}$)/(0.5×sales$_t$+ 0.5×sales$_{t-1}$). Only firms with 150 or more quarters of accounts with sales and pretax profits figures are used to minimize the effects of sample composition changes.[57] The growth rates are windsorized at the top and bottom 0.05% growth rates to prevent the series being driven by extreme outliers.

Standard deviation of firm-level stock returns: This is measured on a monthly basis using the CRPS data file. It is the cross-sectional standard deviation of the monthly stock returns. The sample is all firms with 500 or more months of stock-returns data. The returns are windsorized at the top and bottom 0.5% growth rates to prevent the series being driven by extreme outliers.

Standard deviation of industry-level TFP growth: This is measured on an annual basis using the NBER industry database (Bartelsman, Becker and Grey 2000). The cross-sectional spread is defined

[55] Consider, for example, a Cobb-Douglas *revenue* function, $AK^{\alpha}L^{\beta}$, where A is the productivity term, K is capital, and L is labor. Profit can be written as $\pi = pAK^{\alpha}L^{\beta} - rK - wL$ where p is the price, r is the cost-of-capital and w is the wage rate, initially assumed to be fixed. First, consider the situation where K and L are *costlessly adjustable.* Under profit maximization one can easily show that $K = \phi_1 A$ and $L = \phi_2 A$ where ϕ_1 and ϕ_2 are functions of α, β, p,r and w, so that the growth of profit/sales is a linear function of $\Delta A/A$. Alternatively consider the situation in which K and L are *totally fixed.* In this case the growth rate of profit/sales is also a linear function of the growth rate of A because $\Delta\pi/AK^{\alpha}L^{\beta} = \Delta AK^{\alpha}L^{\beta}/AK^{\alpha}L^{\beta} = \Delta A/A$. Of p, w and r will also fluctuate somewhat over time, but to the extent these fluctuations are common to all firms this will not effect the cross-sectional standard-deviation of profits growth.

[56] Note that employment is not reported quarterly, so no quarterly productivity figures are available.

[57] Limiting compositional change helps to address some of the issues raised by Davis, Faberman and Haltiwanger (2006), who find rising sales volatility of publicly-quoted firms but flat volatility of privately-held firms. I also include a time-trend in column (2) to directly control for this and focus on short-run movements.

as the standard deviation of the 5-factor TFP growth rates, taken across all SIC 4-digit manufacturing industries. The complete sample is a balanced panel for 422 of the 425 industries (results are robust to dropping these 3 industries).

Standard deviation of GDP forecasts. This is measured on a half-yearly basis using the Philadelphia Federal Reserve Bank's Livingstone survey of professional forecasters. It is defined as the cross-sectional standard deviation of the one-year ahead GDP forecasts normalized by the mean of the one-year ahead GDP forecasts. Only half-years with 50+ forecasts are used to ensure sufficient sample size for the calculations. This series is linearly detrended across the sample (1950 to 2006) to remove a long-run downward drift of forecaster variance.

A.3. VAR Data

The VAR estimations are run using monthly data from July 1962 until July 2005. The full set of VAR variables in the estimation are log industrial production in manufacturing (Federal Reserve Board of Governors, seasonally adjusted), employment in manufacturing (BLS, seasonally adjusted), average hours in manufacturing (BLS, seasonally adjusted), log consumer price index (all urban consumers, seasonally adjusted), log average hourly earnings production workers (manufacturing), Federal Funds Rate (effective rate, Federal Reserve Board of Governors), a monthly stock-market volatility indicator (described below) and the log of the S&P500 stock market index. All variables are HP detrended using a filter value of $\lambda = 129,600$.

In Figure A1 the industrial production impulse response function is shown for four different measures of volatility: the actual series in Figure 1 after HP detrending (square symbols), the 1/0 volatility indicator with the shocks scaled by the HP detrended series (dot symbols), an alternative volatility indicator which dates shocks by their first month (rather than their highest month) (triangle symbols), and a series which only uses the shocks linked to terror, war and oil (plus symbols). As can be seen each one of these shock measures generates a rapid drop and rebound in the predicted industrial production. In Figure A2 the VAR results are also shown to be robust to a variety of alternative variable sets and orderings. The VAR is re-estimated using a simple bivariate VAR with industrial production and the volatility indicator only (square symbols), also displaying a drop and rebound. The trivariate VAR (industrial production, log employment and the volatility indicator) also displays a similar drop and rebound (cross symbols), as does the trivariate VAR with the variable ordering reversed (circular symbols). Hence the response of industrial production to a volatility shock appears robust to both the basic selection and ordering of variables. In Figure A3 I plot the results using different HP detrending filter values: the linear detrended series ($\lambda = \infty$) is plotted (square symbols) alongside the baseline detrending ($\lambda = 129,600$) (cross-symbols) and the 'flexible' detrending ($\lambda = 1296$). As can be seen the results again appear robust. I also conducted a range of other experiments, such as adding controls for the oil price (spot price of West Texas), and found the results to be robust.

A.4. Evidence for Cross-Sectional and Temporal Aggregation

Table (A2) shows that as investment data is aggregated across units (going from the small establishments on the bottom row to firms on the top row) and across lines of capital (going from structures, equipment and vehicles columns on the left to the total column on the right) the investment zeros disappear. Table (A3) shows that going from quarterly to annual data generates a drop in the volatility of sales and investment data.

Table A.2: Cross-Sectional Aggregation and Zero Investment Episodes.

Annual zero investment episodes (%)	Structures	Equipment	Vehicles	Total
Firms	5.9	0.1	n.a.	0.1
Establishments (All)	46.8	3.2	21.2	1.8
Establishments (Single Plants)	53.0	4.3	23.6	2.4
Establishments (Single Plants, <250 employees)	57.6	5.6	24.4	3.2

Source: UK ARD plant-level data and UK Datastream firm level data

Table A.3: Temporal Aggregation and Time Series Volatility.

Standard deviation/mean of growth rates	Quarterly	Yearly
Sales	6.78	2.97
Investment	1.18	0.84

Source: Compustat firms with quarterly data 1993-2001

B. Appendix: Numerical Solution Method

This Appendix describes some of the key steps in the numerical techniques used to solve the firm's maximization problem. The full program, which runs on Matlab 64-bit, is on http://www.stanford.edu/~nbloom.

B.1. Value Function Iteration

The objective is to solve the value function (3.9). This value function solution procedure is used in two parts of the paper. The first is in the Simulated Method of Moments estimation of the unknown adjustment cost parameters, whereby the value function is repeatedly solved for a variety of different parameters, the data simulated and these moments used in the parameter search algorithm. The second is in the simulation where the value function is solved just once - using the estimated parameters choices - and then used to simulate a panel of 1000 units, repeated 25000 times. The numerical contraction mapping procedure used to solve the value function in both cases is the same. This proceeds following four steps:

(1) Choose a grid of points in (a, l, σ, μ) space. Given the log-linear structure of demand process I use a grid of points in $(\log(a), \log(l), \sigma, \mu)$ space. In the $\log(a)$ and $\log(l)$ dimensions this is equidistantly spaced, and in the σ and μ spacing this is determined by the estimated parameters. The normalization by capital in a and l - noting that $a = A/K$ and $l = L/K$ - also requires that the grid spacing in the $\log(a)$ and $\log(l)$ dimensions is the same (i.e. $a_{i+1}/a_i = l_{j+1}/l_j$ where $i, j = 1, 2, ..N$ index grid points) so that the set of investment rates $\{a_i/a_1, a_i/a_2, ...a_i/a_N\}$ maintains the state space on the grid.[58] This equivalency between the grid spaces in the $\log(a)$ and $\log(l)$ dimensions means that the solution is substantially simplified if the values of δ_K and δ_L are equal, so that depreciation leaves the $\log(l)$ dimension unchanged. When δ_K and δ_L are unequal the difference between them needs to be an integer of the grid spacing. For the $\log(a)$ dimension depreciation is added to the drift in the stochastic process, so there is no constraint on δ_K. Given the conversion to logs I need to apply the standard Jensen's correction to the driving process (3.4, 3.5 and 3.6), for example for (3.4) $\log(A_t^M) = \log(A_{t-1}^M) - (\sigma_{t-1}^2 - \sigma_L^2)/2 + \sigma_{t-1}W_t^M$. The uncertainty effect on the drift rate is second-order compared its real-options effect, so the simulations are virtually unchanged if this correction is omitted.

I used a grid of 40,000 points ($100 \times 100 \times 2 \times 2$). I also experimented with finer and coarser partitions and found that there was some changes in the value functions and policy choices as the partition changed, but the characteristics of the solution - i.e. a threshold response space as depicted in Figure (3) - was unchanged so long as about 60 or more grid points were used in the $\log(a)$ and $\log(l)$ dimensions. Hence, the qualitative nature of the simulation results were robust to moderate changes in the number of points in the state space partition.

(2) Define the value function on the grid of points. The is straightforward for most of the grid but towards the edge of the grid due to the random walk nature of the demand process this requires taking expectations of the value function off the edge of the state space. To address this an extrapolation procedure is used to approximate the value function off the edge of the state space. Under partial-irreversibilities and/or fixed-costs the value function is log linear outside the zone of inaction, so that so long as the state space is defined to include the region of inaction this approximation is exact. Under quadratic adjustment costs the value function, however, is concave so a log-linear approach is only approximately correct. With a sufficiently large state space, however, the probability of being at a point off the edge of the state space is very low so any approximation error will have little impact.

(3) Select a starting value for the value function in the first loop. I used the solution for the value function without any adjustment costs, which can be easily derived. In the SMM estimation routine I initially tried using the last solution in the next iteration, but found this could generate instability in the estimations loop. So I instead I always used the same initial value function.

[58] Note that some extreme choices of the investment rate will move the state off the l grid which induces an offsetting choice of employment growth rates e to ensure this does not occur.

(4) The value function iteration process. The speed of value function iteration depends on the modulus of contraction, which with a monthly frequency and a 6.5% annual discount rate is relatively slow. So I used value function acceleration (see Judd, 1998) in which the factor of acceleration λ was set to 0.33 as follows

$$Q_{i+1} = Q_i + \lambda(Q_i - Q_{i-1})$$

where Q_i is iteration number i for the value function in the numerical contraction mapping.[59] The number of loops was fixed at 250 which was chosen to ensure convergence in the *policy* functions. In practice, as Krusell and Smith (1998) note, value functions typically converge more slowly than the policy function rules associated with them. Thus, it is generally more efficient to stop the iterations when the policy functions have converged even if the value function has not yet fully converged.

B.2. Simulated Method of Moments Estimation (SMM)

To generate the simulated data for the SMM estimation (used to create $\Psi^S(\theta)$ in equation 5.1) I simulate an economy with 1000 firms, with 250 units each. This is run for 30 years, with the first 10 years discarded to eliminate the effects of any assumptions on initial conditions. I run this simulations 25 times to try to average out over the impact of any individual macro shocks. The same seed is always used in every simulation iteration. I also assume firms are initially distributed equally across μ_L and μ_H given the symmetry of the transition matrix for $\mu_{i,t}$. In order to ensure that first moment draws have a constant aggregate drift rate I numerically set $\sum_{i,j} A_{i,j,t} = \exp^{\frac{\mu_L+\mu_H}{2}t} \sum_{i,j} A_{i,j,0}$, consistent with (3.8) as $N \to \infty$, which in smaller samples stops extreme draws for individual units from driving macro averages

B.2.1. Estimation

I use a simulated-annealing algorithm for minimizing the criterion function in the estimation step in equation 5.1. This starts with a predefined first and second guess. For the third guess onwards it takes the best prior guess and randomizes from this to generate a new set of parameter guesses. That is, it takes the best-fit parameters and randomly 'jumps' off from this point for its next guess. Over time the algorithm 'cools', so that the variance of the parameter jumps falls, allowing the estimator to fine-tune its parameter estimates around the global best-fit. I restart the program with different initial conditions to ensure the estimator converges to the global minimum. The simulated annealing algorithm is extremely slow, which is an issue since it restricts the size of the parameter space which can be estimated. Nevertheless, I use this because it is robust to the presence of local-minima and discontinuities in the criterion function across the parameter space.

B.2.2. Numerical Standard Errors

To generate the standard errors for the parameter point estimates I generate numerical derivatives of the simulation moments with respect to the parameters and weight these using the optimal weighting matrix. One practical issue with this is the value of the numerical derivative, defined as $f'(x) = \frac{f(x+\epsilon)-f(x)}{\epsilon}$, is sensitive to exact value of ϵ chosen. This is a common problem with calculating numerical derivatives using simulated data with underlying discontinuities, arising for example from grid point defined value functions. To address this I calculate four values of the numerical derivative for an ϵ of +1%, +2.5%, +5% and -1% of the 'mid-point' of the parameter space[60] and then take the median value of these numerical derivative. This helps to ensure that the numerical derivative is robust to outliers arising from any discontinuities in the criterion function.

[59] I experimented with different values for λ and found 0.33 was a good trade off between speed (higher values are faster) and stability (higher values make the value function iteration less stable).

[60] For example, the mid-point of parameter space for C_K^F is taken as 0.01, so that ϵ is defined as 0.0001, 0.00025, 0.0005 and -0.0001.

Figure A1: VAR robustness to different shock definitions

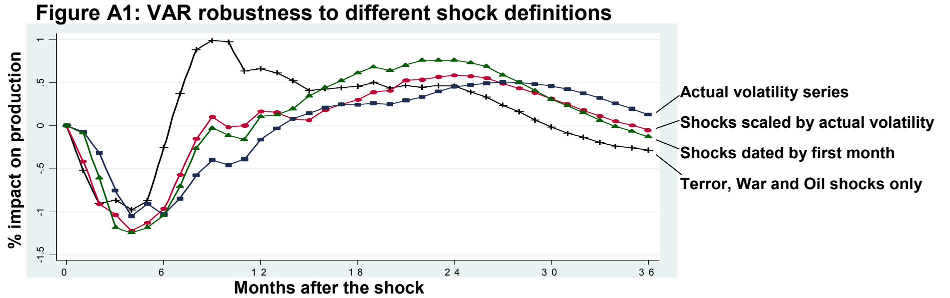

Figure A2: VAR robustness to different variable sets and ordering

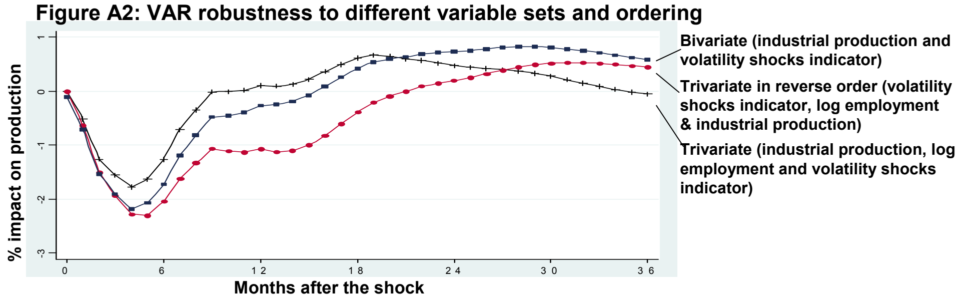

Notes for both figures: VAR Cholesky orthogonalized impulse response functions estimated on monthly data from July 1963 to July 2005 using 12 lags. All data detrended using a Hodrick-Prescott filter with smoothing parameter of 129600 (Stata's monthly default value). In top panel variables (in order) are log industrial production, log employment, hours, log wages, log CPI, federal funds rate, the volatility shock indicator and log S&P500 levels. The volatility indicator used is different for each plot as follows: "actual volatility" is the de-trended series itself, "shocks scaled by actual volatility" uses the 16 shocks but scales these by their actual de-trended level, "shocks dated by first month" uses the 16 events with the timing defined by their first month, and "terror, war and oil shocks only" uses a 1/0 indicator for just the 10 shocks defined as terror, war or oil related. In the bottom panel the standard volatility indicator is used (a 1/0 for each of the 16 shocks in Figure 1 timed by the peak volatility month) but the variable sets and ordering are different as noted.

Figure A3: VAR robustness to different variable detrending assumptions

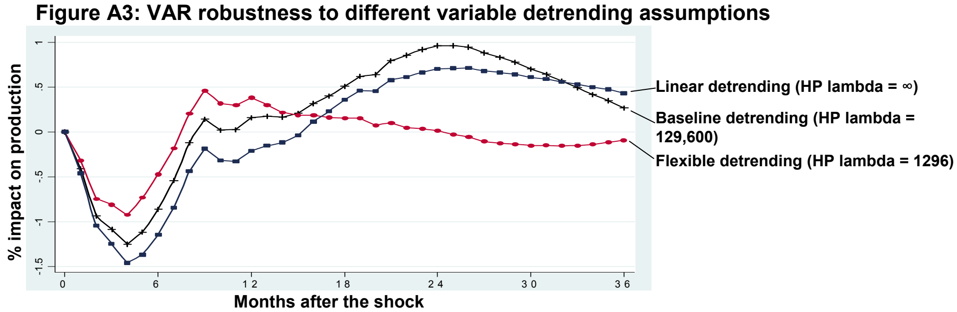

Notes for figures: VAR Cholesky orthogonalized impulse response functions estimated on monthly data from July 1963 to July 2005 using 12 lags. All data detrended using a Hodrick-Prescott filter with smoothing parameters noted above next to each series. Variables (in order) are log industrial production, log employment, hours, log wages, log CPI, federal funds rate, the baseline (1/0) volatility shock indicator and log S&P500 levels. Response to a one-unit change in the volatility shocks indicator is plotted (i.e. the Cholesky response scaled by 1/SD of the impulse variable).

Figure A4: Simulation with a stationary autoregressive demand process

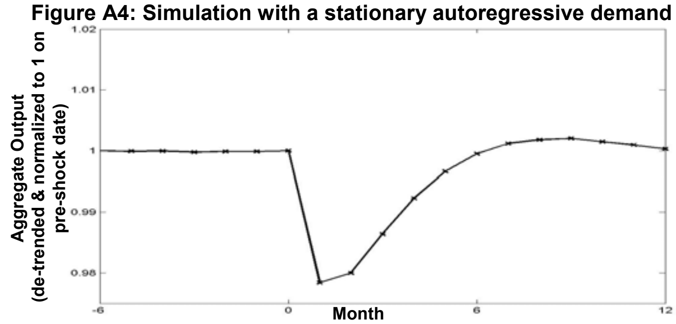

Notes for figures: Simulations run on 1000 units. This is repeated 10000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs for capital are taken from the "All" values in table 3. No adjustment costs for labor. Business conditions ($A_{i,j,t}$) follow a stationary autoregressive process, $A_{i,j,t} = \rho A_{i,j,t-1} + v_t$ where $v_t \sim N(0,\sigma_{t-1})$. Following Cooper and Haltiwanger (2006) I set monthly $\rho=0.885^{1/12}$. The month is normalized to zero at the date of the uncertainty shock. Full program available on http://www.stanford.edu/~nbloom/

Figure A5: Quantifying the size of the real-options effect

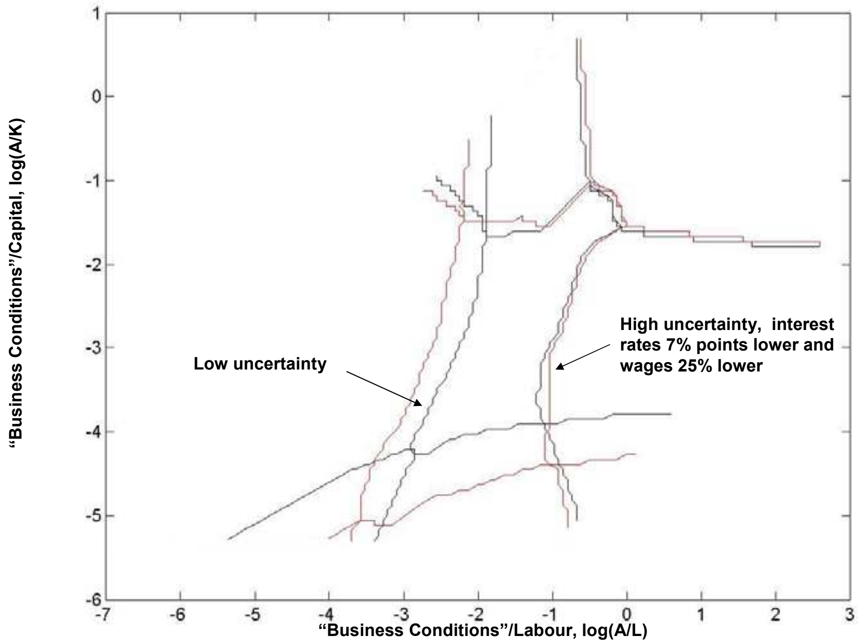

Notes: Simulated thresholds using the adjustment cost estimates "All" in Table 3. At σ_t=σ_H interest rates are 7% points (700 basis points) lower and wages 25% lower, to quantify the approximate size of the short-run rise in uncertainty. All other parameters and assumptions as outlined in sections 3 and 4.

Figure A6: Aggregate (detrended) capital drops, rebounds and overshoots

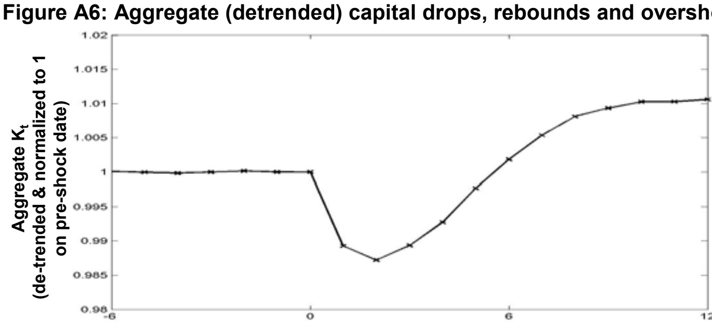

Notes: Simulations run on 1000 units. This is repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. Adjustment costs are taken from the "All" values in table 3. All other parameters and assumptions as outlined in sections 3 and 4. The aggregate figures for K_t are calculated by summing up across all units within the simulation. They are detrended by removing their long-run growth rate. The month is normalized to zero at the date of the uncertainty shock.

Figure A7: Impact of 'Pseudo-GE' price changes <u>without</u> adjustment costs

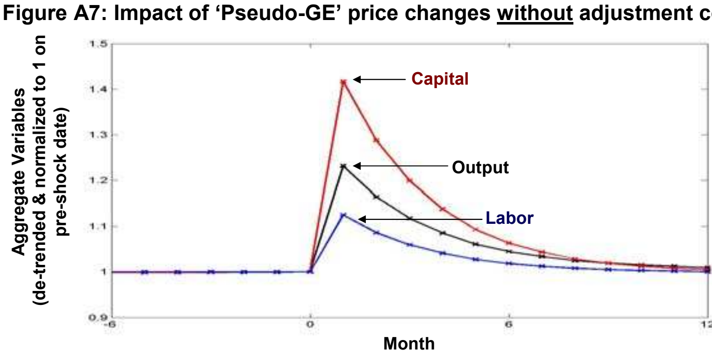

Notes: Simulation run on 1000 units. This is repeated 25000 times with the average plotted here. All micro and macro shocks drawn randomly except at month 0, when all simulations have σ_t set to σ_H. <u>All adjustment costs are set to zero</u>. All other parameters taken from the estimated "All" column in Table 3 and as outlined in sections 3 and 4. The simulation is 'Pseudo-GE', so interest rates, prices and wages are 1.1% points, 0.5% and 0.3% lower during periods of high uncertainty. All series are detrended by removing their long-run growth rate. The month is normalized to zero at the uncertainty shock.

C. References:

ABEL, A. (1983), "Optimal Investment under Uncertainty", *American Economic Review,* **73**, 228-233.

ABEL, A. and EBERLY, J. (1996), "Optimal investment with costly reversibility", *Review of Economic Studies*, **63**, 581-593.

ABEL, A. and EBERLY, J. (2002), "Investment and q with fixed costs: an empirical analysis", Wharton School, University of Pennsylvania, mimeo.

ADDA, J. and COOPER, R. (2000), "Balladurette and Juppette: a discrete analysis of scrapping subsidies", *Journal of Political Economy,* **108**, 778-806.

AGHION, P., ANGELETOS, G., BANERJEE, A. and MANOVA, K. (2005), "Volatility and growth: credit constraints and productivity-enhancing investment", NBER WP11349.

BACHMAN, R., CABALLERO, R. and ENGEL, E. (2006), "Lumpy Investment in Dynamic General Equilibrium", NBER WP 12336.

BAILY, M., HULTEN, C. AND CAMPBELL, D. (1992),"Productivity dynamics in manufacturing plants", *Brookings Papers on Economic Activity*, **1992**, 187-249.

BARNDORFF-NIELSEN, O. and SHEPHARD, N. (2006),"Econometrics of testing for jumps in financial economics using bipower variation", Journal of Financial Econometrics, **4**, 1-30.

BARTELSMAN, E., BECKER, R. and GRAY, W. (2000) NBER Productivity Database, available at www.nber.org.

BARLEVY, G. (2004), "The cost of business cycles under endogenous growth" *American Economic Review*, **94**, 964-990.

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

BERNARD, A., REDDING, S. and SCHOTT, P. (2006), "Multi-product firms and product switching", NBER WP 12293.

BERTOLA, G. and BENTOLILA, S. (1990), "Firing costs and labor demand: How bad is Eurosclerosis", *Review of Economic Studies*, **54**, 318-402.

BERTOLA, G. and CABALLERO, R. (1990), "Kinked Adjustment Costs and Aggregate Dynamics", *NBER Macroeconomics Annual,* **61**, 223-246.

BERTOLA, G. and CABALLERO, R. (1994), "Irreversibility and Aggregate Investment", *Review of Economic Studies,* **61**, 223-246.

BLOOM, N., (2006), "The impact of uncertainty shocks: firm-level estimation and a 9/11 simulation", Centre for Economic Performance discussion paper 718.

BLOOM, N., BOND, S. and VAN REENEN, J.(2007), "Uncertainty and Investment Dynamics", *Review of Economic Studies,* **74**, 391-415.

BLOOM, N., FLOETOTTO, M. and JAIMOVICH, N. (2007), "Really Uncertain Business Cycles", Stanford mimeo.

BLOOM, N. and VAN REENEN, J. (2007), "Measuring and Explaining Management Practices Across Firms and Countries", forthcoming *Quarterly Journal of Economics.*

BOND, S. and VAN REENEN, J. (2007), "Microeconometric models of investment and employment", forthcoming *Handbook of Econometrics.*

BRESNAHAN, T. and RAFF, D. (1991), "Intra-industry heterogeneity and the Great Depression: the American motor vehicles industry, 1929-1935", Journal of Economic History, vol. **51**, 317-331.

CABALLERO, R. (1991), "On the sign of the investment uncertainty relationship", *American Economic Review*, **81**, 279-288.

CABALLERO, R. and ENGEL, E. (1993), "Microeconomic adjustment hazards and aggregate dynamics", *Quarterly Journal of Economics*, **108**, 359-383.

CABALLERO, R. and ENGEL, E. (1999), "Explaining Investment Dynamics in U.S. Manufacturing: a Generalized (S,s) Approach", *Econometrica*, **67**, 783-826.

CABALLERO, R., ENGEL, E. and HALTIWANGER, J. (1995), "Plant-Level Adjustment and Aggregate Investment Dynamics", *Brookings Papers on Economic Activity*, 2, 1-54.

CABALLERO, R., ENGEL, E. and HALTIWANGER, J. (1997), "Aggregate Employment Dynamics: Building from Microeconomics Evidence", *American Economic Review,* **87**, 115-137.

CABALLERO,R. and LEAHY, J. (1996), "Fixed costs: the demise of marginal Q", NBER working paper 5508.

CHRISTIANO, L., EICHENBAUM, M. and EVANS, C. (2005), "Nominal rigidities and the dynamic effects of a shock to monetary policy", *Journal of Political Economy*, **113**, 1-45.

COOPER, R. and HALTIWANGER, J. (1993), "The Aggregate Implications of Machine Replacement: Theory and Evidence", *American Economic Review*, **83**, 360-82.

COOPER, R. and HALTIWANGER, J. (2006), "On the nature of the capital adjustment process", *Review of Economic Studies*, **73**, 611-633.

COOPER, R., HALTIWANGER, J. and WILLIS, J. (2004), "Dynamics of labor demand: evidence from plant-level observations and aggregate implications", NBER working paper 10297.

COOPER, R., HALTIWANGER, J. and POWER, L. (1999), "Machine Replacement and the Business Cycle: Lumps and Bumps" *American Economic Review*, **89**, 921-946

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

DAVIS, S., FABERMAN, R. and HALTIWANGER, J. (2006), "The Flow Approach to Labor Markets: New Data Sources and Micro-Macro Links", forthcoming *Journal of Economic Perspectives.*

DIXIT, A. (1993), *The Art of Smooth Pasting,* Fundamentals of Pure and Applied Economics, Harwood Academic Publishers, Switzerland.

DIXIT, A., and PINDYCK, R. (1994), *Investment Under Uncertainty*, Princeton University Press, Princeton, New Jersey.

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

GILCHRIST, S. and WILLIAMS, J. (2005), "Investment, capacity and uncertainty: a putty-clay approach", *Review of Economic Dynamics,* **8**, 1-27.

HALL, B. (1987), "The relationship between firm size and firm growth in the US manufacturing sector", *Journal of Industrial Economics*, **35**, 583-606

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

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

HARTMAN, R. (1972), "The effect of price and cost uncertainty on Investment", *Journal of Economic Theory*, **5**, 258-266.

HASSLER, J. (1996), "Variations in risk and fluctuations in demand - a theoretical model", Journal of Economic Dynamics and Control, **20**, 1115-1143.

JUDD, K. (1998), *Numerical Methods in Economics*, MIT Press, Cambridge, Massachusetts.

HAYASHI, F. (1982), "Tobin's marginal Q and average Q: a neoclassical interpretation", Econometrica, **50**, 213-224.

KHAN, A. and THOMAS, J. (2003), "Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?", *Journal of Monetary Economics*, **50**, 331-360.

KING, R. and REBELO, S. (1999), "Resuscitating real business cycles", in the *Handbook of Macroeconomics*, John Taylor and Michael Woodford (eds.), Elsevier, North-Holland.

KRUSELL, P and SMITH, A. (1998), "Income and wealth heterogeneity in the macroeconomy", *Journal of Political Economy*, **106**, 867-896.

LEAHY, J. and WHITED, T. (1996), "The Effects of Uncertainty on Investment: Some Stylized Facts", *Journal of Money Credit and Banking*, **28**, 64-83.

LEE, B. and INGRAM. B. (1991), "Simulation estimation of time series models", *Journal of Econometrics*, **47**, 197-205.

MEGHIR, C. and PISTAFERRI, L (2004), "Income variance dynamics and heterogeneity", *Econometrica*, **72**, 1-32,.

MERZ, M. and YASHIV, E. (2006), "Labor and the Market Value of the Firm", forthcoming American Economic Review.

NICKELL, S. (1986), "Dynamic models of labor demand", in: O. Ashenfelter and R. Layard, Eds., *Handbook of Labor Economics*, vol.1, North-Holland, Amsterdam.

OHANIAN, L. (2001), "Why did productivity fall so much during the Great Depression?", *American Economic Review,* **91**, 34-38.

RAMEY, G. and RAMEY, V, (1995) "Cross-country evidence on the link between volatility and growth", *American Economic Review*, **85**, 1138-1151.

RAMEY, V. and SHAPIRO, M, (2001) "Displaced Capital", *Journal of Political Economy*, 605-622.

ROMER, C. (1990) "The Great Crash and the Onset of the Great Depression.", *Quarterly Journal of Economics*, **105**, 597-624

SHAPIRO, M., (1986) "The Dynamic Demand for Labor and Capital", *Quarterly Journal of Economics*, **101**, 513-42.

SHILLER, R. (1981) "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?", *American Economic Review*, **71**, 421–436

SIM, J., (2006) "Uncertainty, Irreversible Investment and General Equilibrium", Boston University mimeo.

STOKEY, N. and LUCAS, R. (1989), *Recursive Methods in Economic Dynamics*, with PRESCOTT, E., Harvard University Press, Cambridge Mass.