

The Impact of Uncertainty Shocks: A Firm-Level Estimation and a 9/11 Simulation

Nick Bloom*

January 2007

Abstract

Uncertainty appears to vary strongly over time, temporarily rising by up to 200% around major shocks like the Cuban Missile crisis, the assassination of JFK and 9/11. This paper offers the first structural framework to analyze 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 employment, investment and productivity growth, and a moderate loss in GDP. This temporary impact of a second moment shock is different from the typically persistent impact of a first moment shock, highlighting the importance for policymakers of identifying their relative magnitudes in major shocks. The simulation of an uncertainty shock is then compared to VAR estimations on monthly data and a 9/11 event-study, both of which display a surprisingly good match.

Keywords: Labor, investment, uncertainty, real options

JEL Classification: D92, E22, D8, C23.

Acknowledgement: This was the main chapter of my PhD thesis. I would like to thank my formal discussants Russel Cooper, Janice Eberly, Valerie Ramey and Chris Sims, and participants at seminars at the American Economics Association, Bank of England, Bank of Portugal, Berkeley, the Board of Governors, Boston College, Boston Fed, Chicago Fed, Chicago GSB, Kellogg, LSE, MIT, NBER EF&G and CM&E groups, Northwestern, San Francisco Fed, Stanford, UCL and Yale. The financial support of the ESRC (Grant R000223644) is gratefully acknowledged.

Correspondence: nbloom@stanford.edu

*Dept. of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305; the Centre for Economic Performance, the NBER and SIEPR.

1. Introduction

Major shocks to the economic and political system appear to cause large variations in macro uncertainty over time. Figure 1 presents a stockmarket volatility proxy for uncertainty¹, plotted monthly from 1962 to 2005. This varies dramatically over time, driven by major events like the Cuban missile crisis, the assassination of JFK, the OPEC I oil-price shock, and 9/11. These shocks generate large but short-lived bursts of uncertainty, increasing (implied) volatility by up to 200%. Uncertainty is also a ubiquitous concern of policymakers - for example Figure 2 plots the frequency of the word “uncertain” appearing in the Federal Open Market Committees (FOMC) minutes, which displays a clear jump and decay around 9/11.

But despite the size and regularity of these second moment (uncertainty) shocks there is still no general structural model of 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 to model the second moment effects of major shocks on employment, investment and productivity growth. This links with the earlier work of Bernanke (1983) and Hassler (1996) who highlight the importance of variations in uncertainty.² In this paper I quantify and substantially extend their predictions through two major advances: first by modelling uncertainty as a stochastic process which is critical for evaluating the high frequency impact of major shocks; and second by modelling a joint mix of labor and capital adjustment costs which is critical for understanding the dynamics of employment, investment and productivity. I then build in temporal and cross-sectional aggregation and estimate this model on firm level data using simulated method of moments

¹In 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) and others have 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.

²Bernanke does this in 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 the business-cycle and uncertainty such as Justiniano and Primiceri (2005) and Gilchrist and Williams (2005), on policy uncertainty such as Adda and Cooper (2000) and on income and consumption uncertainty such as Attanasio (2000) and Meghir and Pistaferri (2004).

to identify the structural parameters. Using firm-level data overcomes the identification problem of limited macro data.

With this parameterized model I then simulate the impact of a large temporary uncertainty shock and find that this generates a rapid drop and rebound in hiring, investment 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 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. But again productivity growth rapidly bounces back as pent-up re-allocation occurs. In sum, these second moment effects generate a rapid slow-down and bounce-back in economic activity, generating a short-run loss of GDP, but with little longer run impact. 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 a range of issues. One is general equilibrium effects, which are not included in my model, for which I conclude that the predictions are likely to be *qualitatively* robust. This is for two reasons: first prices are relatively inflexible over the monthly time frame analyzed, with stickiness in wages and prices and a zero nominal interest rate floor. This prevents prices and wages adjusting fast enough to fully address the very short-run impact of an uncertainty shock, and interest rates from falling far enough to offset the large (temporary) rise in firm's hurdle rates. Second, even with fully flexible prices delaying the reallocation of some factors of production at higher uncertainty will be optimal due to adjustment costs. High uncertainty makes the appropriate allocation of factors unclear, and if it is expensive to get this wrong due to adjustment costs, this will induce an optimal pause in activity until uncertainty returns to normal levels. I also examine the impact of firm-level risk aversion

³See, for example, Cooley (1995), King and Rebelo (1999), and Christiano, Eichenbaum and Evans (2005) and the references therein.

and find that this amplifies the real-options effects, increasing the immediate cut-back in investment and hiring, and thereby generating a stronger re-bound. In addition I consider a combined first and second moment shock - which is typical of the major shocks shown in Figure 1 - and find this generates a rapid drop and partial rebound. Finally, I re-run the simulations for different adjustment costs and find the predictions are sensitive to the inclusion of non-convex adjustments costs but not their magnitude.⁴

A comparison of these predictions to a range of empirical evidence is then undertaken. First, I estimate a VAR using the monthly volatility data in Figure 1 and find a sharp drop and rebound in the response of industrial production and employment to volatility shocks, consistent with the model. This is estimated using a binary indicator for the sixteen large shocks labelled in Figure 1, so should identify the impact of these arguably exogenous uncertainty shocks. Second, I also undertake a comparison to one recent uncertainty shock - the 9/11 attack. Compared to the consensus economic forecasts made just before 9/11 the attack appears to have caused a rapid drop and rebound in activity. I also examine cross-sectional investment activity around 9/11 and find a compression and rebound in quarterly investment rates after 9/11. Finally, I also look at Central Banks minutes and find supportive evidence for a real-options effect. For example, the October 2001 minutes for the FOMC report “*the events of September 11 produced a marked increase in uncertainty....depressing investment by fostering an increasingly widespread wait-and-see attitude*”.

The secondary contribution of this paper is to analyze the importance of jointly modelling labor and capital adjustment costs. The empirical literature has for analytical tractability and aggregation constraints either estimated labor or capital adjustment costs individually assuming the other factor is flexible, or estimated them jointly assuming only convex adjustment costs. These alternative approaches, however, have produced a range of different results.⁵ I estimate a joint mix of labor and capital adjustment costs by exploiting the properties of homogeneous functions to reduce the state space, and develop an

⁴ Non-convex adjustment costs include any lump-sum investment or hiring/firing costs (like closing a plant for a capital refit or employee unrest for a labor layoff) or any degree of irreversibility in investment or hiring (like capital resale losses or labor recruitment, induction, training or firing costs).

⁵ See, for example on capital Doms and Dunne (1993), Cooper and Haltiwanger (1993), Caballero, Engle 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 & Engel (1993), Caballero, Engel and Haltiwanger (1997) and Cooper, Haltiwanger and Willis (2004); and on joint estimation with convex adjustment costs Shapiro (1986), Hall (2004) and Merz & Yashiv (2005).

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 an acceptable fit for the labor moments but a poor fit for investment and output moments.

The framework in this paper also contributes to many of the debates in the business cycle literature⁶, including the lack of negative technology shocks which a second moment shock can substitute for; the explanation for why a positive total-factor productivity (TFP) shock has a negative *impact* effect on hiring and investment, which a 2nd moment shock arising from the 1st moment shock can explain⁷; the instability of VAR estimates without controls for volatility which second moment shocks can rationalize; and the role of non-convexities in aggregation, which second moment shocks bring center stage. Taking a longer run perspective this model also links to the volatility and growth literature given the negative impact of uncertainty on productivity growth in the model.

The rest of the paper is organized as follows: in section (2) I set up and solve my model of the firm, in section (3) I outline my simulated method of moments estimation approach, in section (4) I report the parameters estimates using US firm data, in section (5) I take my parameterized model and simulate the high frequency effects of a large uncertainty shock, and in section (6) I compare this to the 9/11 shock. Finally, section (7) offers some concluding remarks.

2. The Model

2.1. Overview

I model a firm as a collection of a very large number of production *units*. Each unit faces an iso-elastic demand curve for its product which is produced with a Cobb-Douglas technology in capital, labour and hours. Both demand and productivity are affected by

⁶See King and Rebelo (1999) and the references therein on the first point; Gali (1999) and Basu, Fernald and Kimball (2006) on the second point; Cogley and Sargent (2002) and Sims (2001 and 2005) on the third point; and Thomas (2002) and Veracierto (2002) on the last point.

⁷For example, if agents know the underlying process driving TFP can change over time then a large TFP shock will signal a shift in the underlying process, temporarily increasing uncertainty over the future evolution of TFP.

multiplicative shocks described by a geometric random walk with time varying drift and uncertainty. These shocks have a unit specific idiosyncratic component, a common firm component and an economy wide macro component. There is also a stochastic capital price. I work in discrete time.

Firms can adjust their capital stock and labor force, but this entails adjustment costs, while hours can be freely raised or lowered but at the penalty of a higher hourly wage rate outside the normal 40 hour week. These adjustments costs allow for a fixed cost and partial irreversibility component, as well as a more traditional convex cost component.

2.2. The Production Unit

Each production unit has a revenue function $R(X, K, L, H)$

$$R(X, K, L, H) = X^\varphi K^{\alpha(1-\epsilon)} (L \times H)^{(1-\alpha)(1-\epsilon)} \quad (2.1)$$

which nests a Cobb-Douglas production function in capital (K), labor (L) and hours (H) and an iso-elastic demand curve with elasticity (ϵ).⁸ Demand and productivity conditions are combined into an index (X) - henceforth called “demand conditions”. For analytical tractability I define $a = \alpha(1 - \epsilon)$, $b = (1 - \alpha)(1 - \epsilon)$, and normalize the demand conditions parameter by the substitution $Y^{1-a-b} = X^\varphi$, so that the revenue function $\tilde{R}(Y, K, L, H)$ is now homogeneous of degree 1 in (Y, K, L) ⁹

$$\begin{aligned} R(X, K, L, H) &= \tilde{R}(Y, K, L, H) \\ &= Y^{1-a-b} K^a (L \times H)^b \end{aligned} \quad (2.2)$$

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

⁸While I assume a Cobb-Douglas production function any supermodular homogeneous unit revenue function could be used. As an experiment I replaced (2.1) with a CES aggregator over capital and labor where $R(X, K, L, H) = X^\alpha (K^\sigma + (L \times H)^\sigma)^{\frac{1}{\sigma}}$ so that $\tilde{R}(X, K, L, H) = Y^{1-\gamma} (K^\sigma + (L \times H)^\sigma)^{\frac{\gamma}{\sigma}}$, where $Y = X^{\frac{\alpha}{1-\gamma}}$. This substitution generated similar simulation results.

⁹This reformulation to Y as the stochastic variable also avoids any Hartman (1972) or Abel (1983) effects of uncertainty which reduces (increases) output because of convexity (concavity) in the production function arising from homogeneity of degree less than (greater than) 1. See Caballero (1991) or Abel and Eberly (1996) for a more detailed discussion.

I assume demand conditions evolve as an augmented geometric random walk, consistent with Gibrat's law that long-run firm growth rates are independent of firm size. 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.

One question is whether shocks to stock-market volatility are a good proxy for shocks to uncertainty? While stock market volatility clearly contains noise from financial bubbles - such as the 1987 crash - there is a variety of evidence suggesting it is also strongly linked to uncertainty. First, stock market volatility is strongly linked with political uncertainty - for example Voth (2002) shows that variations in the probability of political unrest and revolution explained 10% to 50% of stock-market volatility in inter-war Europe and the US. Mei and Guo (2002) show electoral political uncertainty is highly significant in explaining stock-market volatility in 22 developing countries, and Wolfers and Zitewitz (2006) show that over 30% of the variation in the S&P500 between September 2002 and February 2003 can be explained by the changing probability of ousting Saddam Hussein. In fact many of the large spikes in stock-market volatility in Figure 1 are clearly linked to political, war and terror induced uncertainty. Second, stock market volatility is also significantly correlated with factor price and output price volatility. Monthly measures of the daily volatility of interest rates, exchange rates and oil-prices are all highly significant in explaining monthly stock-market volatility, accounting for 22% of its variation (see Appendix E). Third, professional forecaster uncertainty over GDP as measured by the cross-sectional dispersion of future forecasts is also highly correlated (0.426) with stock market volatility (Ramey, 2006).

Demand conditions are in fact modelled as a multiplicative composite of three separate sub random-walks, a macro-level component (Y_t^M), a firm-level component ($Y_{i,t}^F$) and a unit-level component ($Y_{i,j,t}^U$), where $Y_{i,j,t} = Y_t^M \times Y_{i,t}^F \times Y_{i,j,t}^U$ and i indexes firms, j indexes units and t indexes time. The macro level component is modelled as follows:

$$Y_t^M = Y_{t-1}^M \times (1 + \mu + \sigma_t W_t^M) \quad W_t^M \sim N(0, 1) \quad (2.3)$$

where μ is the mean drift in demand conditions, σ_t^2 is the variance of demand conditions and W_t^M is a macro-level iid normal shock. The firm level component is modelled as follows:

$$Y_{i,t}^F = Y_{i,t-1}^F \times (1 + \theta^F \sigma_t W_{i,t}^F) \quad W_{i,t}^F \sim N(0, 1) \quad (2.4)$$

where θ^F is the relative uncertainty of the *firm* level shock and $W_{i,t}^F$ is a firm-level iid normal shock. The unit level component is modelled as follows:

$$Y_{i,j,t}^U = Y_{i,j,t-1}^U \times (1 + \theta^U \sigma_t W_{i,j,t}^U) \quad W_{i,j,t}^U \sim N(0, 1) \quad (2.5)$$

where θ^U is the relative uncertainty of the *unit* level shock and $W_{i,j,t}^U$ is a unit-level iid normal shock, independent of W_t^M and $W_{i,t}^F$. While this demand structure may seem complex it is formulated to ensure that units within the same firm have linked investment behavior due to common firm-level demand and uncertainty shocks, but that they also display some independent behavior due to the idiosyncratic unit level shocks, which is essential for smoothing under aggregation.

The variance of demand conditions (σ_t^2) common to the macro, firm and unit level process is also stochastic and follows an auto-regressive process

$$\sigma_t = \sigma_{t-1} + \rho_\sigma (\sigma^* - \sigma_{t-1}) + \sigma_\sigma S_t + \sigma_Z Z_t \quad S_t \sim \{0, 1\} \quad Z_t \sim N(0, 1) \quad (2.6)$$

where ρ_σ is the rate of convergence of σ_t to its long run mean σ^* , σ_σ is the size of macro uncertainty shocks S_t which are drawn from a $\{0,1\}$ process where $P(S_t = 0) = 1 - \lambda$ and $P(S_t = 1) = \lambda$, and σ_Z is the standard deviation of the smaller ongoing macro fluctuations drawn from a Gaussian process. This process is based on the fact that macro volatility is approximately AR(1) with infrequent large jumps and ongoing smaller fluctuations (Figure 1). Rises in uncertainty will also increase firm (and unit) level uncertainty, consistent with the fact that cross-sectional firm-level share returns spreads also increase significantly around major shocks, for example rising by 75% after 9/11¹⁰, as these events typically affect firms differently.

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

Partial irreversibilities: Labor partial irreversibility derives from hiring, training and firing costs, is labelled PR_L , and is denominated as a fraction of annual wages (at the

¹⁰The monthly average of the daily cross-sectional standard deviation of returns for 1721 CRSP firms rose from 2.86% in August 2001 to 4.95% in September 2001. More generally Campbell et al. (2001) calculate macro time series and firm-level cross-sectional volatility measures and find a very similar pattern of jumps after major shocks like OPEC I and II.

standard working week). For simplicity I assume these costs apply equally to gross hiring and gross firing of workers. Capital partial irreversibilities arise from resale losses due to transactions costs, the market for lemons phenomena and the physical costs of resale. The resale loss of capital is labelled PR_K and is denominated as a fraction of the relative purchase price of capital, labelled p_t^K . This price of capital is stochastic and is assumed to follow a mean reverting process.

$$p_t^K = p_{t-1}^K + \rho_{PK}(p^{K*} - p_{t-1}^K) + \sigma_{PK}T_t \quad T_t \sim N(0, 1) \quad (2.7)$$

where p^{K*} is the mean price of capital (normalized to unity), ρ_{PK} is the rate of reversion to this mean, σ_{PK} is the relative variance in the price of capital and T_t is an iid normal shock. This stochastic capital price is introduced to generate some separation between the capital and labor processes.

Fixed disruption costs: When new workers are added into the production process and new capital 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 FC_L and FC_K for hiring/firing and investment respectively, both denominated as fractions of annual revenue.

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 $QC_{LL}(\frac{E}{L})^2$ are the quadratic hiring/firing costs and E denotes gross hiring/firing, and $QC_{KK}(\frac{I}{K})^2$ are the quadratic investment costs

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

$$C(Y, K, L, H, I, E, p_t^K) = 52 \times w(40) \times PR_L(E^+ + E^-) + p_t^K(I^+ - (1 - PR_K)I^-) + FC_L(E \neq 0) + FC_K(I \neq 0) + QC_{LL}(\frac{E}{L})^2 + QC_{KK}(\frac{I}{K})^2$$

where E^+ (I^+) and E^- (I^-) are the absolute values of positive and negative hiring (investment) respectively, and $(E \neq 0)$ and $(I \neq 0)$ 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. At the end of each period there is labor attrition and capital depreciation proportionate to δ_L and δ_K respectively.

2.3. The Firm

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 (see appendix table A1); and temporal aggregation across higher-frequency periods within each year (see appendix table A2). I build this aggregation into the model by explicitly assuming firms own a large number of production *units* and these operate at a higher frequency than yearly. These 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 easy or palatable, 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. In the simulation this is set at 250 units per firm, chosen by increasing the number of units until the results were no longer sensitive to this number.¹¹ Units are assumed to independently optimized 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).

Of course in practice these assumptions are unlikely to hold and units will be linked within the firm, so the question is how sensitive these results are to this assumption. I test

¹¹In the UK ARD census microdata - which is very similar to the US LRD - the average size of a manufacturing production local unit is 20.8 employees. The median size of firms in my estimating data is 4,500 employees (see section 3.3), suggesting a median of around 220 local units per firm, similar in magnitude to my assumption of 250 units per firm.

this by estimating a specification (column 6 table 3) in which the number of units is 25 rather than 250, approximating a firm with very strong links within sub-sets of 10 units, for example if these served common markets. I find the results are reasonably similar despite this large reduction in the degree of aggregation.¹² The model also assumes no entry or exit for analytical tractability, which seems acceptable in this monthly time frame.

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 is typical, which I assume in my main results.

2.4. 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. In the main results I assume the firm is risk neutral to focus on the real options effects of uncertainty, but I also provide a simulation result for a risk-averse firm showing risk-aversion actually reinforces the real-options effects.

Analytical results can be used to show a unique solution to the firm's optimization problem exists which is continuous and strictly increasing in (Y, K, L) with an almost everywhere unique policy function.¹³ The model is too complex, however, to fully solve using analytical methods, so I use numerical methods knowing this solution is convergent with the unique analytical solution.

Given current computing power, however, I have too many state and control variables to solve this even using numerical methods. 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 (Y, K, L) , which are pre-determined in period t given time to build, so can be optimized out in a prior step. This reduces the *control* space by one dimension. Second, the revenue function, adjustment cost function, depreciation schedules

¹²The results of Bloom et al. (2005) are also re-assuring on this point as they find the qualitative real-options effects of uncertainty are robust to aggregation across types of capital within the same unit.

¹³The application of Stokey and Lucas (1989) for the continuous, concave and almost surely bounded normalized returns and cost function in (2.8) for quadratic adjustment costs and partial irreversibilities, and Caballero and Leahy (1996) for the extension to fixed costs.

and demand processes are all jointly homogenous of degree one in (Y, K, L) , allowing the whole problem to be normalized by one state variable, reducing the *state* space by one dimension. I normalize by capital to estimate on $\frac{Y}{K}$ and $\frac{L}{K}$.¹⁴ These two steps dramatically speed up the numerical simulation, which is run on a state space of (y, l, σ, p^k) of dimension $(120, 120, 5, 2)$, making numerical estimation feasible.¹⁵ 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:

$$\begin{aligned} V(Y_t, K_t, L_t, \sigma_t, p_t^k) &= \max_{I_t, E_t, H_t} \tilde{R}(Y_t, K_t, L_t, H_t) - C(Y_t, K_t, L_t, H_t, I_t, E_t, p_t^K) - w(H_t)L_t \\ &\quad + \frac{1}{1+r} E[V(Y_{t+1}, K_t(1-\delta_K) + I_t, L_t(1-\delta_L) + E_t, \sigma_{t+1}, p_{t+1}^k)] \end{aligned}$$

where r is the discount rate and $E[\cdot]$ is the expectations operator. Optimizing over hours to define $H_t^* = h(Y_t/K_t, L_t/K_t)$, and exploiting the homogeneity in (Y, K, L) to take out factors of K_t or K_{t+1} enables this to be re-written as:

$$\begin{aligned} K_t V(y_t, 1, l_t, \sigma_t, p_t^k) &= \max_{i_t, e_t} K_t R_H(y_t, 1, l_t) - K_t C_H(y_t, 1, l_t, i_t, l_t e_t, p_t^K) \\ &\quad + K_{t+1} \frac{1}{1+r} E[V(y_{t+1}, 1, l_t, \sigma_{t+1}, p_{t+1}^k)] \end{aligned}$$

where $R_H(Y_t, K_t, L) = \tilde{R}(Y_t, K_t, L, h(Y_t/K_t, L_t/K_t)) - w(h(Y_t/K_t, L_t/K_t))L_t$, $C_H(Y_t, K, L_t, I_t, E_t, p_t^K) = C(Y_t, K_t, L_t, h(Y_t/K_t, L_t/K_t), I_t, E_t, p_t^K)$ and the normalized variables are $l = \frac{L}{K}$, $y = \frac{Y}{K}$, $i = \frac{I}{K}$ and $e = \frac{E}{L}$. Finally, by dividing through by K_t we obtain

$$\begin{aligned} Q(y_t, l_t, \sigma_t, p_t^k) &= \max_{i_t, e_t} R^*(y_t, l_t) - C^*(y_t, l_t, i_t, l_t e_t, p_t^K) \quad (2.8) \\ &\quad + \frac{1 - \delta_K + i_t}{1+r} E[Q(y_{t+1}, l_t, \sigma_{t+1}, p_{t+1}^k)] \end{aligned}$$

where $Q(y_t, l_t, \sigma_t, p_t^k) = V(y_t, 1, l_t, \sigma_t, p_t^k)$ which is in fact Tobin's Q, $R^*(y_t, l_t) = \tilde{R}_H(y_t, 1, l_t)$, and $C^*(y_t, l_t, i_t, l_t e_t, p_t^K) = C_H(y_t, 1, l_t, i_t, l_t e_t, p_t^K)$.

¹⁴An alternative normalization by labor (L) is equally feasible, while the normalization by the demand process (Y) is mathematically feasible but (after initial experimentation) turned out to numerically difficult due to Jensen's inequality effects from taking reciprocals of stochastic variables.

¹⁵Of course I also need an optimal control space (i, e) of dimension $(120, 120)$, so that the full returns function in the Bellman equation has dimensionality $(120, 120, 120, 120, 5, 2)$.

2.5. A Numerical Example

As an example of the predictions of the model Figure 3 plots in $(\frac{Y}{K}, \frac{Y}{L})$ 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 the preferred parameter estimates in section (4).¹⁶ The inner region is the region of inaction ($i = 0$ and $e = 0$). 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 4 displays the same lines for two different values of current uncertainty, $\sigma_t = 19\%$ in the inner box of lines (low uncertainty) and $\sigma = 37\%$ for the outer box of lines (high uncertainty). 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 a quantitatively important impact on investment and hiring behavior. In this example doubling uncertainty increases the additional real-options premium on the investment hurdle rate¹⁷ from 6% at $\sigma_t = 19\%$ to 10% at $\sigma = 37\%$, increasing firms (risk-neutral) discount rate by 4%.

Interestingly, re-computing these thresholds with permanent (time invariant) uncertainty results in a stronger impact on the investment and employment thresholds. So the standard comparative static result¹⁸ 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. This is why adopting this more complex stochastic volatility approach is important for analyzing the impact of high frequency uncertainty shocks.

¹⁶See table 2 column (2).

¹⁷Following Abel and Eberly (1996) we can define the investment hurdle rate (c) as $c = r + \delta_K + \phi(\sigma)$, where r is the real interest rate, δ_K the depreciation rate and $\phi(\sigma)$ the additional real options premia.

¹⁸See, for example, Dixit and Pindyck (1994).

3. Estimating the Model

The econometric problem consists of estimating the parameters that characterize the firm's revenue function, stochastic processes, adjustment costs and discount rate, denoted Θ . 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.

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 then solved and policy functions 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 W weighted distance between the simulated moments $\Psi^S(\Theta)$ and the actual moments Ψ^A .

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

$$\hat{\Theta} = \min_{\Theta} [\Psi^A - \Psi^S(\Theta)]'W[\Psi^A - \Psi^S(\Theta)] \quad (3.1)$$

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

The optimal 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 Ψ^A , Lee and Ingram (1989) show that under the estimating null the variance-covariance of the simulated moments, $\Psi^S(\Theta)$, 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. A value for Ω is calculated by block bootstrap with replacement on the actual data.

The asymptotic distribution of the efficient W weighted estimator can be shown to be

$$\sqrt{N}(\hat{\Theta} - \Theta) \xrightarrow{D} N(0, [E[\partial\Psi(\Theta)/\partial\Theta]]'[(1 + \frac{1}{\kappa})\Omega]E[\partial\Psi(\Theta)/\partial\Theta]]) \quad \text{as } N \rightarrow \infty \quad (3.2)$$

where $E[\partial\Psi(\Theta)/\partial\Theta]$ is taken at $\hat{\Theta}$. Since I use $\kappa = 10$ this implies the standard error of $\hat{\Theta}$ is increased by only 5% by using simulation estimation, plus any additional imprecision from using a discretized state space.

3.1. 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 probably least known six adjustment cost parameters, $\Theta = (PR_L, FC_L, QC_L, PR_K, FC_K, QC_K)'$, and predefined all the other parameters based on values in the literature and the raw data.¹⁹

The predefined parameters are as follows: (i) capital (α) and labor ($1 - \alpha$) parameters of 1/3 and 2/3 and an elasticity (ϵ) of -3 (from a 50% mark-up); (ii) a capital depreciation rate (δ_K) of 10%, an exogenous labor quit rate (δ_L) of 10% and a discount rate (r) of 6%; (iii) a wage level parameter (w_1) set to 1/3 (to generate about 20 employees per unit), an hours parameter (w_2) set to 7e-06 (to generate an optimal week of 40 hours) and a wage curvature parameter (γ) of 2.5 (to generate an overtime share of 27% (Trejo, 1993)); (iv) an annual real demand drift (μ) of 5% (Compustat sample average real sales growth); (v) mean uncertainty (σ^*) of 29.0%, to ensure the simulated annual standard deviations of monthly share returns matches the actual mean Compustat annual standard deviations of monthly share returns; (vi) a macro uncertainty shock size σ_σ of size σ^* , probability λ of 1/60 and mean-reversion of ρ_σ of 0.42 based on macro shocks doubling uncertainty, occurring twice a decade and having about a 1.5 month half-life (Figure 1), and macro Gaussian volatility standard deviation of $\sigma_z = \sigma^*/10$ reflecting the fact that the (non-shock) variation in macro volatility is about 10% of the average level (Figure 1); (vii) the relative variance of firm-level shocks of (θ^F) of 1.132 (to match the firm level share returns standard deviation of 29.0% with the average S&P100 macro returns standard deviation

¹⁹This procedure could, of course, be used iteratively to check my predefined parameters by using the estimated adjustment costs $\hat{\Theta}$ 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.

of 19.2%) and plant-level shocks (θ^U) of 0.34 (from UK plant-level data²⁰); and (viii) the mean price of capital (p^{K*}) normalized to 1, a price of capital mean-reversion (ρ_{PK}) 0.27 and standard-deviation (σ_{PK}) of 0.12 from the NBER 4-digit industry data set (see Becker et al. 2000).

Given these values for $\sigma^*, \sigma_\sigma, \rho_\sigma, \sigma_z$ and λ the simulated uncertainty process can be modelled. This is achieved using a five-point grid, where the optimal grid points and transition matrices are calculated using the quadrature procedures in Tauchen (1986) and Tauchen and Hussey (1991).²¹

3.2. 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 chosen. 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. This is used to focus on four cross sectional moments, the standard deviation and skewness coefficients of investment and employment growth rates, and six dynamic moments, the intertemporal correlations of investment, employment growth and sales growth rates.

To demonstrate these moments provide identification Table 1 presents their values for each of the adjustment cost parameters in turn, and then for sets of combinations of these. Columns (2) and (5) present the moments for partial irreversibility in labor then capital respectively, which compared to the no adjustment cost benchmark (column 1), display much stronger dynamics, a lower standard-deviations and a heavy skew in each factor. Columns (3) and (6) present the moments for fixed costs for labor then capital, which display moderate dynamics, little reduction in the standard-deviation and a heavy skew in each factor. While columns (4) and (7) present the moments for quadratic adjustment costs in labor then capital, which display strong dynamics, a lower standard-deviation but little skew in each factor. Thus, all six adjustment costs generate distinct patterns across

²⁰Calculated from the decomposition of the variance of employment growth rates in local *unit* data (single postal address production sites) within and between *firms* for 1996 to 2002 in the UK ARD (which is similar to the US LRD).

²¹See also the discussion and Matlab routines in Adda and Cooper (2003).

Table 1: Identification of adjustment costs

Adjustment costs	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
PR_L : hire/fire per head	0	0.250	0	0	0	0	0	0.250	0	0
FC_L : hire/fire fixed	0	0	0.050	0	0	0	0	0	0.050	0
QC_L : rapid hiring/firing	0	0	0	2.000	0	0	0	0	0	2.000
PR_K : disinvestment loss	0	0	0	0	0.250	0	0	0.250	0	0
FC_K : investment fixed	0	0	0	0	0	0.050	0	0	0.050	0
QC_K : rapid investment	0	0	0	0	0	0	2.000	0	0	2.000
Labor growth $(\Delta L/L)_{i,t}$ moments										
Standard. Deviation	0.352	0.233	0.290	0.231	0.282	0.320	0.272	0.226	0.275	0.190
Coefficient Skewness	0.021	0.699	0.463	0.090	0.283	0.225	0.041	0.707	0.678	0.156
Corr. with $(I/K)_{i,t-2}$	-0.030	0.099	0.042	0.143	0.034	-0.015	0.064	0.127	0.030	0.197
Corr. with $(\Delta L/L)_{i,t}$	-0.026	0.115	0.005	0.152	0.041	0.007	0.061	0.119	0.011	0.172
Corr. with $(\Delta S/S)_{i,t-2}$	-0.014	0.140	0.068	0.187	0.054	0.014	0.071	0.167	0.095	0.220
Investment $(I/K)_{i,t}$ moments										
Standard Deviation	0.436	0.368	0.383	0.365	0.202	0.320	0.147	0.197	0.300	0.137
Coefficient Skewness	0.059	0.172	0.175	0.055	1.459	0.903	0.119	1.612	0.969	0.232
Corr. with $(I/K)_{i,t-2}$	-0.023	-0.013	-0.019	0.001	0.141	-0.017	0.285	0.149	0.018	0.291
Corr. with $(\Delta L/L)_{i,t}$	-0.031	0.009	-0.011	0.023	0.121	0.071	0.241	0.145	0.049	0.293
Corr. with $(\Delta S/S)_{i,t-2}$	-0.024	0.019	0.002	0.027	0.199	0.075	0.301	0.200	0.110	0.323

The top part contains selected values for adjustment costs while the bottom panel contains the associated simulated data moments. The simulation data sample is created from a balanced panel of 5000 firms over 10 years. Each firm year is aggregated across 250 units per firm and 12 months per year.

the ten moments, providing identification for the adjustment cost parameters.

Comparing across the columns in Table 1 it is also clear that while the adjustment costs have the largest impact on the factor they apply to, the other factor's moments are also affected. For example, in column (2) the introduction of partial irreversibility in labor makes the labor growth moments smoother and much more skewed, but also has a similar (but weaker) effect on the investment moments. Columns (8) to (10) suggest this cross-factor impact is weaker if both factors have adjustment costs, but is nevertheless still important. For example, comparing columns (8) to (2) we see that the addition of partial irreversibility for capital has a noticeable effect on the labor moments despite labor already being subject to its own partial irreversibility. This highlights the importance of allowing for a full set of labor and capital adjustment costs when estimating these.

3.3. 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 10 years of data covering 1991 to 2000.²²

The data was cleaned to remove major mergers and acquisitions by dropping observations with jumps of +200% or -66% in the employment and capital stocks. Only Manufacturing firms with 500+ employees and a full 10 years of data were kept to focus on a larger more aggregated firms and reduce the impact of entry and exit.²³ This generated a sample of 579 firms and 5790 observations with median employees of 4500 and median sales of \$850m (2000 prices). In selecting all manufacturing 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.

²²This data spans two macro uncertainty shocks - the Asian and Russian crises - and so these are also included when generating simulated data by introducing macro uncertainty shocks in the equivalent months. This data also spans several macro demand shocks, but since the simulation and estimation both use firm level data without any cross-firm linkages first moment macro shocks are modelled as part of the firm level shock process.

²³While this focus on larger continuing firms reduces the need to model entry and exit decisions it does undoubtedly introduce a selection bias. In terms of coverage the total number of employees in the Compustat panel averages 8.9 million per year (including foreign employees) while the average domestic manufacturing employment reported by the Bureau of Labor Statistics for the same period is 16.7 million.

Capital stocks for firm i in industry m in year t are constructed by the perpetual inventory method²⁴, 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 $(\frac{\Delta L}{L})_{i,t} = \frac{\Delta L_{i,t}}{0.5*(L_{i,t}+L_{i,t-1})}$ and the sales growth as $(\frac{\Delta S}{S})_{i,t} = \frac{\Delta S_{i,t}}{0.5*(S_{i,t}+S_{i,t-1})}$.²⁵ Yearly firm uncertainty, $sd_{i,t}$, is calculated as the yearly standard deviations of monthly share returns (net cash flow plus capital gains per \$ of equity).²⁶

The simulated data is constructed in exactly the same manner as actual company accounts data, enabling the moments of the actual and simulated data to be directly matched. So simulated data for flow figures from the accounting Profit & Loss and Cash-Flow statements (such as sales and capital expenditure) values are added up across units across the year, while data for stock figures from the accounting Balance Sheet statement (such as the capital stock and labor force) are added up across units at the year end.

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 a number of advantages versus using census datasets like the LRD because firm-level data is: (i) easily available to all researchers, (ii) covers all sectors of the economy in a range of different countries; (iii) is matched into firm level financial and cash-flow data; and (iv) is available as a yearly panel stretching back several decades (for example to the 1950s in US). Thus, this technique should have a much broader use in firm-level estimation.

3.4. 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 hours, usually after heavy rounding. This problem is then made much

²⁴ $K_{i,t} = (1 - \delta_K)K_{i,t-1} \frac{P_{m,t}}{P_{m,t-1}} + I_{i,t}$, 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 et al. (2000).

²⁵ 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.

²⁶ A similar share returns variance measure has been previously used by Leahy and Whited (1996). The leverage adjustment normalizes the standard deviation of firms returns by $\frac{E+D}{E}$ where E is the market value of common plus preferred stock and D is the book value of long-term debt.

worse by the differencing to generate growth rates.

As a first step towards addressing these measurement errors intertemporal correlations of growth rates are taken between periods t and $t - 2$ to reduce the sensitivity to levels measurement error. 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 SIC-4 digit industry average wage from Becker et al. (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 $\frac{\sigma_L^2}{\sigma_L^2 + \sigma_{ME}^2}$ 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 0.898 (0.010), implying a measurement error of 11% in the logged labor force numbers.²⁷ This is reassuringly similar to the 8% estimate for measurement error in Compustat manufacturing firms' labor figures IIall (1987) calculates comparing OLS and IV estimates. This 11% measurement error is incorporated into the simulation estimation by multiplying the aggregated annual firm labor force by $me_{i,t}$ where $me_{i,t} \sim iid LN(0, 0.11)$ before calculating simulated moments.

4. Adjustment Costs Estimates

Turning to Table 2 the first column reports the actual moments for Compustat. These demonstrate that labor growth rates are relatively variable but un-skewed, with weak dynamic correlations. Investment is less variable but has a heavy right skew due to the lack of disinvestment, and much stronger dynamic correlations.

The second column in Table 2 presents the results from estimating the preferred specification. The estimated adjustment costs for labor imply limited hiring and firing costs of 9.5% of annual wages (about five-weeks of wages), a high-fixed cost of around 4.6% of

²⁷ 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 0.967 (0.010), indicating that the proportional measurement error in wages is less than one third that of employment. The regressions are run on 194 firms (those who report wage data) with 1603 observations.

Table 2: Adjustment cost estimates

	(1) Data	(2)	(3)	(4)	(5) Simulated	(6)	(7)
Adjustment costs specification		All	Capital	Labor	Quadratic	All	All
Number of units per firm	250	250	250	1		25	250
Share noise adjustment	No	No	No	No		No	Yes
<hr/>							
Estimated Adjustment Costs							
PR_L : hire/fire per head		0.095		0.000		0.102	0.109
FC_L : hire/fire fixed		0.046		0.059		0.040	0.047
QC_L : rapid hiring/firing		0.000		0.650	1.444	0.000	0.000
PR_K : disinvestment loss	0.421	0.145				0.435	0.461
FC_K : investment fixed	0.006	0.021				0.002	0.007
QC_K : rapid investment	4.743	6.052		4.634		5.323	4.635
<hr/>							
Labor growth $(\Delta L/L)_{i,t}$ moments							
Standard Deviation	0.197	0.234	0.265	0.234	0.234	0.251	0.223
Coefficient of Skewness	0.213	0.437	0.033	0.680	0.225	0.682	0.424
Correlation with $(I/K)_{i,t-2}$	0.102	0.152	0.117	0.097	0.162	0.137	0.130
Correlation with $(\Delta L/L)_{i,t}$	0.111	0.106	0.057	0.119	0.163	0.095	0.110
Correlation with $(\Delta S/S)_{i,t-2}$	0.137	0.174	0.120	0.153	0.189	0.156	0.165
<hr/>							
Investment $(I/K)_{i,t}$ moments							
Standard Deviation	0.141	0.146	0.164	0.359	0.177	0.162	0.140
Coefficient of Skewness	1.404	1.031	0.991	0.199	0.295	1.251	0.923
Correlation with $(I/K)_{i,t-2}$	0.305	0.318	0.345	-0.003	0.244	0.259	0.325
Correlation with $(\Delta L/L)_{i,t}$	0.139	0.207	0.170	0.022	0.244	0.158	0.199
Correlation with $(\Delta S/S)_{i,t-2}$	0.201	0.325	0.267	0.032	0.271	0.247	0.303
Criterion, $\Gamma(\Theta)$		229	313	2357	351	239	206

Notes: “Data” in column (1) is balanced panel of 579 firms over 10 years from Compustat. Simulation data in columns (2) to (7) is a balanced panel of 5000 firms over 10 years. Each firm year is aggregated across 250 units per firm and 12 months per year. The criterion function, $\Gamma(\Theta)$, is minimized in the parameter search and provides a goodness of fit measure with lower values signifying a better fit.

annual revenue (about two weeks sales), and no quadratic adjustment costs. The estimated capital adjustment costs imply heavy resale costs of about 42%, a fixed resale cost of about 0.6% of annual revenue (about 1/2 a weeks sales), and a moderate quadratic adjustment coefficient of 4.743.²⁸

One question is how do these estimates compare to those previously estimated in the literature. The available evidence is as follows: for labour partial adjustment costs (PR_L) Nickell (1986) reports about 1 months wages for unskilled workers consistent with my estimates, but several months for skilled workers which is higher than my estimates although my fixed cost term may proxy for the additional costs; for labour quadratic adjustment costs (QC_L) Hall (2004) suggests a value of 0 consistent with my estimate; and on fixed labor disruptions costs (FC_L) Cooper et al (2004) suggest a value of around 1.2% which is lower than my 4.6% estimate, although their figure is estimated from annual plant-level data without any provision for aggregation or *capital* adjustment costs. The available evidence on investment adjustment partial irreversibilities (PR_K) appears roughly consistent with my values, with Ramey and Shapiro (2001) estimating resale losses of between 40% to 80% based on aerospace plant closure data. For fixed disruption costs (FC_K) there is an extremely wide span with Caballero and Engel (1999) and Cooper and Haltiwanger (2004) estimating higher costs of 16.5% and 20.4% while Thomas (2002) estimates costs of around 0.1%, although all these estimates are on an annualized basis, without provision for temporal aggregation or *labour* adjustment costs, and use a variety of different methodologies. For quadratic adjustment costs (QC_K), my estimate lies within an even larger span of estimates (normalized to a monthly basis), ranging from 0 for industry data (Hall, 2004), to 0.294 on establishment level data (Cooper and Haltiwanger, 2004) to 480 on firm level (Hayashi, 1982), with again these based on a variety of differing assumptions over timing, other factors adjustment costs and temporal aggregation.

For interpretation I also display results in columns 3 to 5 for three 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 column 3 we see that the fit of the “Capital” adjustment costs only model is worse than the “All” adjustment costs model (column 2), as shown by the rise in the criterion function from 229 to 313. However,

²⁸Note the quadratic adjustment cost coefficient is on a monthly basis so should be normalized by 12 for comparison to values estimated on annual data.

this reduction in fit mainly arises from the labor moments, suggesting that ignoring labor adjustment costs is a reasonable approximation for investment modelling.²⁹ Second, a model with “Labor” adjustment costs only - as is typical in the dynamic labor demand literature - is estimated in column 4, with the fit substantially reduced by an extremely poor fit on the capital moments, with the labor moments themselves looking reasonable. This suggests that ignoring capital adjustment costs is a reasonable approximation for narrowly modelling labor demand, although this would be unsuitable for modelling any functions of capital such as output or productivity. Finally, a model with quadratic costs only and no cross-sectional aggregation - as is typical in convex adjustment costs models - is estimated in column 5, 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.³⁰

In columns 6 and 7 I run a couple of robustness tests on the modelling assumptions. In column 6 I estimate the model with a smaller number of units to examine the impact of cross-unit links within firms. The estimated adjustment costs parameters are somewhat higher. This is because less aggregation induces less smoothing and lower intertemporal correlations, requiring higher adjustment costs to compensate.

Since the parameters of the uncertainty process are determined to match the moments of actual share-returns variance column 7 checks the extent to which the estimated adjustment costs are sensitive to the potentially excessive volatility of share-returns. Jung and Shiller (2002) provide evidence that excess volatility is more likely to be a phenomena of overall stock-market returns than relative firm-level share returns.³¹ Vuolteenaho (2002) undertakes a variance decomposition of firm-level share returns *relative* to the S&P500 and finds around 5/6 of this can be attributed to cash-flow volatility (equivalent to “demand conditions” volatility in the model). In column 7 I estimate a specification in which the demand process is set at 5/6 of the variance of firm share returns *relative* to the S&P500. To do this I re-calculate from Compustat the leverage adjusted annual stan-

²⁹Thus, the labor adjustment costs are principally identified from the labor moments. This is also apparent in Table 2 where each factor is more sensitive to its own adjustment costs.

³⁰Cooper and Haltiwanger (2003) also note this point.

³¹This appears due to the “Samuelson Dictat” that individual agents will find it easier to arbitrage away relative mispricing for individual shares than absolute mispricing for the whole stock-market.

dard deviation of monthly share returns *relative* to the S&P500, take 5/6 of these values and use their mean (24.1%) to re-calculate the parameter σ^* and re-estimate the model. This provides alternative adjustment costs estimates using what is more of a lower bound for the true cash-flow returns variance. The re-estimated non-convex adjustment costs are moderately higher and quadratic costs moderately lower to offset the lower skew and standard-deviation in the demand process.

5. Simulating an uncertainty shock

5.1. Overview

I start by running the *thought experiment* of analyzing a second moment uncertainty shock in isolation. Of course this is only a very stylized simulation since many 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 in the context of risk-aversion, different estimates for the adjustment costs, general equilibrium effects, and a combined first and second moment shock.

5.2. The Simulation

An *uncertainty* shock is defined in the model as a positive draw ($S_t = 1$) for the uncertainty jump process in equation (2.6). Given the choice of $\sigma_\sigma = \sigma^*$ this will double average firm uncertainty that period³², approximating the doubling of micro and macro uncertainty after major shocks displayed in Figure 1.

I simulate an economy of 250,000 firms for 10 years (at a monthly frequency) without any first or second moment shocks ($W_t^M = 0$, $Z_t = 0$ and $S_t = 0$ in every month) to generate a steady-state ergodic distribution, using the preferred “All” parameter specification from Column (2). The model is then hit with an uncertainty shock ($S_t = 1$) in month 1 of year 6, with no other macro shocks for the next 2 years, ensuring the only macro impulse is the uncertainty shock.

³²This is only true on average because there are also minor fluctuations due to the Gaussian shocks (Z_t). For simplicity there are set to zero in the simulation ($Z_t = 0$).

In figure 5a (the top panel) I plot the total monthly net-employment growth, which displays a substantial fall in the four months immediately after the uncertainty shock and a bounce-back in months 5 to 9. This occurs because the rise in average uncertainty generates valuable real options, making firms much more cautious so that they pause their employment behavior. Once the uncertainty begins to dissipate firms increase net-hiring to address their pent-up demand from the proceeding period of inaction. The impact of this at a monthly level is large - during the five months after the uncertainty shock aggregate net-hiring falls and becomes negative as hiring freezes while exogenous quits continue. Endogenizing quits would reduce the impact of those shocks of course. But in the model I have very conservatively assumed a 10% annual quit rate - well below the typical 20% floor for the quit rate throughout the business cycle so that a 10% rate can reasonably be assumed to be exogenous due to retirement, maternity, incapacity, relocation etc. Hiring then rebounds and mildly overshoots trend for the next few months as firms address their labor shortages from the period of prior inaction.

In figures 5b (bottom panel) I plot the 99th, 95th, 5th and 1st percentiles of hiring to demonstrate the distributional impact of the uncertainty shock on hiring. After the shock the hiring and firing thresholds move apart (as illustrated in figure 4), and this reduces both hiring and firing activity, which compresses the distribution of activity across firms. Again, after the shock has passed these hiring percentiles rebound as the firms react to pent-up demand accumulated during the period of inaction.

In figures 6a I plot the macro investment outcome. This looks similar to hiring, with again a rise in uncertainty causing a temporary pause in firms activities, with a subsequent bounce-back to clear pent-up demand. Gross investment falls to around 50% of its long-run value in the 5 months after the shock, and then mildly overshoots trend for the next few months. Figure 6b demonstrates the cross-sectional compression of investment rates that occurs after the shock.

Figure (7a) plots the time series for aggregate productivity growth, defined in terms of the demand conditions growth, $\Delta \log(Y)$.³³ Following Baily, Hulten and Campbell (1992)

³³While Y combines demand and productivity effects, since both operate through the same channel, this will also simulate the response of productivity.

I define four indices as follows:

$$\begin{aligned}
& \Delta \sum_i \sum_j \log(Y_{i,j,t}) \frac{L_{i,j,t}}{\sum_i \sum_j L_{i,j,t}} && \text{Total productivity growth} \\
&= \sum_i \sum_j \Delta \log(Y_{i,j,t}) \frac{L_{i,j,t-1}}{\sum_i \sum_j L_{i,j,t-1}} && \text{Within productivity growth} \\
&+ \sum_i \sum_j \log(Y_{i,j,t-1}) \Delta \frac{L_{i,j,t}}{\sum_i \sum_j L_{i,j,t}} && \text{Between productivity growth} \\
&+ \sum_i \sum_j \Delta \log(Y_{i,j,t}) \Delta \frac{L_{i,j,t}}{\sum_i \sum_j L_{i,j,t}} && \text{Cross productivity growth}
\end{aligned}$$

where $L_{i,j,t}$ is employment, and Δ is the difference operator. The first term, “Total” growth, is the increase in employment weighted by productivity. This can be broken down into three sub-terms: “Within” growth which measures the productivity increase within each production unit, “Between” growth which measures the reallocation of employment from low to high productivity units, and “Cross” productivity growth which measures the correlation between productivity growth and employment growth.

In figure 7a “Total” productivity shows a large fall after the uncertainty shock, dropping to around 35% of its value immediately after the shock. 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.³⁴ 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” productivity growth is entirely driven by the fall in the reallocation “Between” term. The “Within” term is constant since, by assumption, the *mean* draw for demand conditions shocks is unchanged, while the “Cross” term is zero because of the random walk nature of Y and the 1 period time to build. In the bottom two panels this reallocation 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 reallocation activity takes place with a substantially lower fraction of expanding productive units and shrinking unproductive units. Since actual US

³⁴Formally 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”.

aggregate productivity growth is probably about 70% or 80% driven by reallocation³⁵ these uncertainty effects should play an important role in the real impact of large uncertainty shocks.

As another way to quantify the impact of a second moment shock Table 3 reports the lost output during the first 2, 4 and 6 months after the uncertainty shock. This is broken down into the lost output due to the temporary fall in the *level* of factor inputs as result of the fall in employment growth and investment (figures 5 and 6), and the lost output due to the temporary fall in productivity as a result of the fall in reallocation (figure 7). As can be seen the uncertainty impact potentially reduces GDP by around 1.3% within the first 6 months.

Table 3: GDP loss from an uncertainty shock (% annual)

	First 2 months	First 4 months	First 6 months
Input factors	0.30	0.74	1.16
TFP (reallocation)	0.07	0.11	0.14
Total (input factors and TFP)	0.37	0.85	1.30

Notes: Simulations run with 250,000 firms, each with 250 plants at a monthly frequency. Adjustment costs include capital and labor convex and non-convex terms as in “All” in column (2) of Table 2. Input factor losses are due to lower levels of capital and labor while TFP losses are due to lower levels of factor reallocation.

5.2.1. Comparing first and second moment shocks

This rapid drop and rebound in response to a second moment shock is very different to the typically persistent drop over several quarters from a more traditional first moment shock.³⁶ Thus, to the extent a large shock is more a second moment phenomena - for example 9/11 - 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

³⁵Foster, Haltiwanger and Krizan (2000 and 2004) 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 (2005) results on product switching suggests could be substantial.

³⁶See, for example, Cooley (1995), King and Rebelo (1999), or Christiano, Eichenbaum and Evans (2005) and the references therein.

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 suggests that there are two pieces of information available to help policymakers with this, however. First daily stock-market volatility proxies, for example the VVO series³⁷, will provide a direct and immediate indicator of the financial markets view of the uncertainty component of any shock. Second the distribution of responses across firms should assist in identification since the second moment element of a shock will generate a compression of the distribution while the first moment element should generate a downward shift in all percentiles.³⁸

Of course these first and second moment components of shocks differ both in terms of the moments they impact - first or second moment - and in terms of their duration - permanent or temporary. This co-distinction is driven by the fact that the second moment component 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 “Between” 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.³⁹

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 and risk-aversion role in reducing output in the *onset* of the Great Depression, while Ohanian

³⁷The VVO is an index of the financial market’s expectation of near term volatility of the S&P100 equity index. It is provided on a daily basis by the Chicago Board Options Exchange and is calculated from a basket of call and put options.

³⁸Of course, this will be hard to distinguish if the shock differentially impacts sectors, generating a cross-sectoral spread. In this case it will be important to look at the within-sector spread of activity.

³⁹There would also be notable cross-sectional differences at the plant/firm level between a first and second moment shock. A second moment shock would generate a bigger spread of size weighted TFP and a narrower spread of investment and employment growth rates than a first moment shock.

(2001) and Bresnahan and Raff (1991) report “inexplicably” low levels of productivity growth with an “odd” lack of output *reallocation* over this period.

5.3. Risk aversion

In the model in section (2) I assume firms behave as if risk-neutral. Including risk-aversion effects on the part of firms into the model actually amplifies the impact of an uncertainty shock since firms will cut back investment and hiring immediately after the shock when their discount rate rises, which will then generate a stronger re-bound due to a larger pent-up demand when the discount rate falls again. To illustrate this Figure 8a re-plots the investment response under risk-neutrality alongside an example investment response based on a new numerical solution and simulation using the same parameters as in the “All” specification, but additionally incorporating an approximate risk-adjustment which is linear in uncertainty and takes the value of 3% at the average level of uncertainty ($\sigma_t = 29\%$).⁴⁰ It can be seen that including this risk-aversion effect increases the size of the post shock contraction and the subsequent bounceback.

5.4. Adjustment cost robustness

I also evaluate the robustness of the simulation predictions against different estimates for the non-convex adjustment costs which drive the real-options effects of uncertainty. The results from Dixit (1993) and Abel and Eberly (1996) demonstrate that (in a continuous time model) the non-response thresholds depicted in figures (3) and (4) have an infinite derivative with respect to non-convex adjustment costs around their zero. Thus, even small values of non-convex adjustment costs should generate real-options threshold behavior.

To evaluate this figures 9a, 9b and 9c plot aggregate hiring, the hiring percentiles and productivity growth for a simulation assuming only moderate partial irreversibilities, with $PR_L = 0.1$, $PR_K = 0.1$ and all other adjustment costs set to zero. These figures demonstrate a clear drop and rebound in aggregate activity, with a compression of the cross-sectional hiring distribution and a corresponding fall in “Between” and “Total” productivity growth. Figures 9d, 9e and 9f plot the aggregate hiring, the hiring percentiles and productivity growth for a simulation assuming only moderate fixed costs, with $FC_L = 0.01$

⁴⁰This example is based on a conservative 3% estimate of the average equity risk premia (Kocherlakota, 1996). Smaller or larger risk premiums generate proportionally smaller or larger risk-effects.

and $FC_K = 0.01$ and all other adjustment costs set to zero. Again these figures demonstrate a smaller, but nevertheless distinct, drop and rebound in activity, a compression of cross-sectional activity and a fall in “Total” productivity growth driven by a fall in “Between” productivity growth. However, running a simulation with only *quadratic* adjustment costs introduces no compression of cross-sectional activity and almost no change in aggregate activity. Hence, 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 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.

5.5. General equilibrium

Ideally I would set up my model within a General Equilibrium (GE) framework, allowing prices to 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 moments in a representative consumer framework to compute a recursive competitive equilibrium (see, for example, Krusell and Smith, 1998, and Khan and Thomas, 2003). 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 costs⁴¹, and (2) estimated adjustment costs but in a fixed price model. The results on the first-order sensitivity of the results to the presence of non-convex adjustment costs in section (5.4) and the arguments suggesting a limited sensitivity to GE effects over the *monthly* time frame I discuss (see below) suggests taking the second option and leaving a GE analysis to future work.

This, of course, means the results in this model could be compromised by GE effects if factor prices changed sufficiently to counteract factor demand changes. There are two

⁴¹Unfortunately 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 will be problematic, especially for non-convex adjustment costs given the sensitivity of the lumpy behaviour they imply to aggregation. These problems may explain the differences of up to 100 fold in the estimation of some of these parameters in the current literature (see section 4).

reasons to doubt this would substantially occur, however.⁴²

First, prices are not completely flexible over the monthly time-frame analyzed in the simulation. For labor it appears unlikely that wages could change sufficiently rapidly to offset large monthly employment changes in the first 3 months (Bewley, 1999). For example, post 9/11 despite the largest monthly drop in employment since 1980 wages did not fall. The same is also likely to be true for the price of capital goods, which appear to have a multi-month (state-independent) reset period.⁴³ Interest rates falls will occur, but nominal rates are bounded at zero and so are unlikely to be able to fall enough to offset the large real-options and risk-aversion effects of a major uncertainty shock. The average increase in investment hurdle rates in the simulation for a doubling of uncertainty was 4% for the increased real options premia and a further 2% to 5% for the additional risk premia. This will be substantially greater than any interest rates cuts. Again, as an example, in the 3 months after 9/11 the FOMC cut rates by only 1.75%, about 1/4 of the simulated impact of the shock on firms hurdle rates. Furthermore, the simulated productivity effects highlighted in section (5.2) are entirely redistributional, and so should also be robust to GE effects.

Second, even with price flexibility the costs of adjusting capital and labor make it welfare optimal to delay the reallocation of some factors of production while uncertainty is high. High uncertainty makes the appropriate allocation of factors unclear, and if it is expensive to get this wrong due to adjustment costs, this will induce an optimal pause for a few months until uncertainty returns to normal levels. Thus, even a fully flexible general equilibrium model would display a marked slowdown and rebound in activity.

5.6. 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 8b, where the investment response to a second moment

⁴²The recent papers by Thomas (2002) and Veracierto (2002) are also linked with this issue. In their models 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 $\frac{\partial^2 M_t}{\partial Y_t \partial NC} \approx 0$ where M_t is some macro-variable like capital or employment, Y_t is a macro shock variable and NC is a non-convex adjustment cost. The focus of my paper on the direct impact of uncertainty on macro variables, is different and can be characterized instead as $\frac{\partial M_t}{\partial \sigma_t}$. Thus, their results are not necessarily inconsistent with mine.

⁴³See Bils and Klenow (2004) and Klenow and Kryvtsov (2005).

shock (from Figure 6a) is plotted alongside the investment response to the same second moment shock with an additional first moment shock of -5%.⁴⁴ Adding an additional first moment shock leaves the main character of the second moment shock unchanged - a large drop and rebound - but eliminates the overshoot due to the persistent impact of the first moment shock

6. Evaluating the simulation against empirical evidence

6.1. VAR evidence

One way to evaluate the plausibility of the simulation is to compare the shape, size and duration of an output response to a stock-market volatility shock. I do this by estimating a range of VARs on industrial production, employment, hours, inflation, interest rates, stock-market levels and volatility on monthly data from 1962 to 2006.⁴⁵ The main findings are show in Figures 10 and 11 and discussed below, with fuller details in Appendix C.

Figure 10a plots the impulse response function of industrial production to an indicator function for a volatility shock which takes a value 1 for each of sixteen shocks displayed in Figure 1 and a 0 otherwise. These sixteen shocks are chosen as those whose Hodrick Prescott (HP) detrended ($\lambda = 1600$) stock-market volatility series is more than two standard-deviations above the mean. This indicator function is used to ensure identification comes only from the largest, and arguably most exogenous, volatility shocks. The response of industrial production (the solid line with plus symbols) displays a rapid fall of around 1.4% within four months, with a subsequent recovery and rebound from months eight onwards. Both the magnitude and the timing of this are consistent with the predictions from the simulation model. 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 the response to a 1% impulse to the federal funds rate is also plotted (solid line with circular symbols) displaying a much more persistent drop and recovery over the subsequent two years. In Figure 10b the response of employment using this stock-market volatility shock indicator is also plotted, displaying a similar drop and

⁴⁴I choose 5% because this is equivalent to 1 years demand conditions growth in the model. Larger or smaller shocks yield a proportionally larger or smaller impact.

⁴⁵I would like to thank Valerie Ramey and Chris Sims (my discussants at the NBER and Bank of Portugal conferences) for their initial VAR estimations and subsequent discussions.

recovery in activity.⁴⁶

Figure 11a confirms the robustness of these VAR results to some alternative choices of the definition of the volatility variable. In the top panel the industrial production impulse response function is shown for four different measures of volatility: the actual series in Figure 1 after IIP detrending (square symbols), the 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 impulse response function. In Figure 11b 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 (square symbols), but nevertheless displays a sharp 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. Other tests I have run include changing the detrending assumptions - for example HP filter values of ($\lambda = 160$, $\lambda = 16,000$ and $\lambda = 1,000,000$)- and the results again appear robust.⁴⁷

6.2. A 9/11 event study

Another way to evaluate the plausibility of the simulation is to compare this against actual data from a large uncertainty shock. While this is not a test in any sense it does provide a basic sense check for one large uncertainty shock. I choose 9/11 because high frequency consensus forecasts are available for this period providing a forecast baseline. In addition Central Bank minutes are also available from the late 1990s providing a richer background contextual picture.

Looking first at figure (12a), which plots actual quarterly changes in net-employment, there is evidence of a sharp-drop in net employment growth in the quarter after 9/11, with a rapid rebound in 2002 Q1. The size of the immediate fall is large, with 2001

⁴⁶ Interestingly the recovery in employment is slower than output, generating a mildly 'job-less' recovery.

⁴⁷ These experiments, the data and estimation programs will be put on <http://www.stanford.edu/~nbloom/>

Q4 representing the largest quarterly fall in employment growth since 1980. Compared to predicted employment changes from the August 15th 2001 consensus forecasts, 9/11 appears to have generated a net job-loss of around 1 million jobs in the subsequent four months, but with little longer run fall in employment growth. Turning to investment, figure (12b) plots quarterly investment as % contribution to real GDP growth, which demonstrates a similar sharp fall after 9/11, with 2001 Q4 representing the lowest quarterly figure since 1982. Again compared to the prior 9/11 predictions the short-run effects are large - with the drop in investment cutting annual GDP growth by about 3% over the subsequent 4 months - but with a rapid bouncceback in 2002 Q1 and no apparent longer run effects. Thus, macro employment growth and investment are consistent with the predictions from the model, particularly after allowing for risk aversion and/or a simultaneous 1st moment shock.⁴⁸

Another prediction of the model is that cross-sectional investment activity should fall rapidly after the 9/11 uncertainty shock as firms pause. To evaluate this I build a quarterly panel of Compustat investment rates (details in Appendix D). Figure (13a) plots the cross sectional spread of investment activity with a noticeable drop after 9/11. In figures (13b) and (13c) this cross-sectional spread is plotted for the quarters before and after 9/11 with the histogram showing the contraction in the cross-section of investment rates, mainly driven by a drop in larger positive rates reflecting the interpretation of 9/11 as a combination of a negative first moment shock alongside a second moment shock. While these figures are not conclusive they are supportive of the prediction that 9/11 reduced hiring and investment activity.

Because high frequency macro data can be noisy I also look at Central Bank minutes. While the Central Banks did not structurally model the uncertainty impact of 9/11, they did have a strong sense that the real-options effects of uncertainty were important. For example, the FOMC minutes from October 2nd state “*The events of September 11 produced a marked increase in uncertainty and anxiety among contacts in the business sector....depressing investment by fostering an increasingly widespread wait-and-see attitude about undertaking new investment expenditures*”. This view appears to have been widespread with Michael Moskow (President of the Chicago Federal Reserve Board) stating

⁴⁸ After 9/11 GDP growth did pick up, however, driven by the rapid growth in consumer durable expenditure from higher car sales. This was due to a combination of generous zero-rate financing deals introduced by GM, Ford and Chrysler plus an apparent patriotic move to “Buy American”.

almost two-months later on November 27th “*Because the attack significantly heightened uncertainty...it appeared that some households and some business would enter a wait-and-see mode....They are putting capital spending plans on hold*”. Other Central Banks also discussed this phenomena, for example the Bank of England stated in its October 17th minutes “*A general increase in uncertainty could lead to a greater reluctance to make commitments....Labour hiring and discretionary spending decisions are likely to be deferred for a while, to allow time for the situation to clarify*”.

7. Conclusions

Uncertainty appears to dramatically increase after major economic and political shocks like the Cuban Missile crisis, the assassination of JFK and 9/11. If firms have non-convex adjustment costs these uncertainty shocks will generate powerful real-options, driving the dynamics of investment and hiring behavior. This paper offers the first structural framework to analyze these types of uncertainty shocks, building a model with a time varying second moment of the driving process and a rich mix of labor and capital adjustment costs. This 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 employment, investment and productivity growth - due to the effect of higher uncertainty making firms temporarily pause - but with limited longer run 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 in month 1 and has completely rebounded by month 5, a persistent first moment shock will generate a drop in activity lasting several quarters. Thus, for a policy maker in the immediate aftermath of a major shock trying to forecast ahead it is critical to distinguish between persistent first moment effects and temporary second moment effects of the shock. I suggest two pieces of information which could help: first measures of financial uncertainty from implied volatility indices, and second the spread of activity across firms as a first moment shock will generate a fall in activity across all percentiles while a second moment shock will generate a compression of the percentiles.

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 and political elections. More generally these second moments effects contribute to many of the debates in the business cycle literature including the lack of negative technology shocks which a second moment shock can substitute for; the explanation for why a positive total-factor productivity (TFP) shock has a negative impact effect on hiring and investment, which a 2nd moment shock arising from the uncertainty induced by a 1st moment shock can help explain; the potential instability of VAR estimates without controls for volatility which second moment shocks rationalize; and the role of non-convexities in aggregation which second moment shocks bring center stage. Finally, 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 and its sensitivity to uncertainty.

A. Appendix: Data

Table A.1: 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.2: 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 in Matlab for 64-bit Linux, is provided on <http://www.stanford.edu/~nbloom>.

The objective is to solve the value function (2.8). 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 parameter choices in the moment 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 large panel of 250,000 firms subject to a variety of first and second moment shocks. 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 (y, l, σ, p^k) space. Given the log-linear structure of demand process I use a grid of points in $(\log(y), \log(l), \sigma, p^k)$ space. In the $\log(y)$ and $\log(l)$ dimensions this is equidistantly spaced, and in the σ and p^k dimensions the spacing is determined by Tauchen and Hussey's (1991) quadrature method. The normalization by capital in y and l - noting that $y = Y/K$ and $l = L/K$ - also requires that the grid spacing in the $\log(y)$ and $\log(l)$ dimensions is the same (i.e. $y_{i+1}/y_i = l_{j+1}/l_j$ where $i, j = 1, 2, \dots, N$ index grid points) so that the set of investment rates $\{y_i/y_1, y_i/y_2, \dots, y_i/y_N\}$ maintains the

state space on the grid.⁴⁹ This equivalency between the grid spaces in the $\log(y)$ 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. For the $\log(y)$ dimension depreciation is added to the drift in the stochastic process.

I used a grid of 144,00 points ($120 \times 120 \times 5 \times 2$). 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 to 80 grid points were used in the $\log(y)$ 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. This 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. To confirm this I tested a log-quadratic approximation and found this induced no change in the solution.⁵⁰

(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 after the first iteration I used the value function from the last set of parameters for the starting value.

(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% 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 0.5 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.⁵¹ The number of loops was fixed at 500 which was chosen to ensure convergence in

⁴⁹ 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.

⁵⁰ A log-quadratic approximation was considerably slower however. This is because every combination of points outside the state space with a non-zero probability of occurrence requires interpolation, involving an additional loop within the value function loop. Hence, this approximation is called on extremely frequently during the program making the total running time very sensitive to its speed. There are other potentially more accurate approximations that could be used - such as cubic splines - but these will be computationally even slower.

⁵¹ I experimented with different values for λ and found 0.5 was a good trade off between speed (higher values are faster) and stability (higher values dampen errors in the value function less).

the *policy* functions. In practice, as Krusell and Smith (1998) note, value functions typically converge more slowly than the policy functions rule 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.

C. Appendix: VAR estimations

C.1. Standard VAR variables

The VAR estimations are run using monthly data from July 1962 until January 2007. 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), average monthly inflation (All Urban consumers, BLS), Federal Funds Rate (Federal Reserve Board of Governors), a monthly stock-market volatility indicator (described below) and log of the S&P500 stock market index.

This ordering is based on the assumptions that economic and political shocks instantaneously influence stock market levels and volatility, then the Federal Funds rate on the assumption interest rates (can potentially) respond quickly to shocks, then hours and prices as flexible firm-level variables, and finally employment and output as costly to adjust firm-level variables. Including

C.2. Volatility variables

The main measure for volatility shocks is an indicator that takes a value of 1 for the each of the sixteen shocks labelled in Figure 1, and 0 otherwise. These sixteen shocks are chosen as those with stock-market volatility more than two standard-deviations above the Hodrick Prescott detrended ($\lambda = 1600$) 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 (C.1) below, while in section (C.3) below each event is described in more detail.

It is noticeable from table (C.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

Table C.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
OPFC 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	War

level has a correlation of -0.233 with the main 1/0 volatility shock indicator, a correlation of -0.152 with the 1/0 oil, terror and war shock indicator, and a -0.424 correlation with the log detrended volatility index itself.

C.3. Details of the Volatility Events

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 over the continued existence of the installations. The crisis ended on October 28 when the Soviet's 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 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 the introduction of wage and price controls (following the precedent set in the Korean War) and the Budgetary implications of the War.⁵²

⁵²For 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*

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 sector 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 (who was subsequently poisoned in an Italian jail). 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. An additional factor in the volatility of March 1980 was the build-up and collapse of the billionaire Hunt Brothers attempt to corner the world silver market. Together with two Arab investors the Hunt Brothers drove prices up by over 500% in the proceeding three months, only for these to collapse by 50% on “Silver Thursday”, March 27th when they declared bankrupt.

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.⁵³

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 events were associated with the event, without any clear evidence over its exact cause.

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 uses of military force to liberate Kuwait by 52-47 in the Senate and by 259-183 in the House, the closest margins in authorizing force by Congress since the War of 1812. This marginal political support generated pre-invasion uncertainty around the US response.

Asian Crisis: On the 14 May 1997 the Thai Baht came under sustained speculative

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 1987 economy.”

⁵³Time 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.”*

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 headed 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, filed for bankruptcy in December 2001 after admitting to 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 would President Bush proceed anyway.

D. Appendix: Cross-sectional investment

This appendix describes in detail the cross-sectional investment figures used to construct Figures 14a, 14b and 14c in section(6.2).

The data is Standard & Poor's COMPUSTAT dataset. Quarterly investment rates are constructed from quarterly capital expenditure (Data item 90) less quarterly capital resale (Data item 83) normalized by the average of the current and lagged quarters net property, plant and equipment (Data item 42). Observations with zero or missing sales, capital stocks or inventories were dropped. Missing values for capital resale were set to zero if all other variables (including capital purchases) were present. Only US manufacturing firms with \$20m or more average sales and ten years of complete quarterly data from 1990 to 2005 were kept, with this timing chosen to provide 10 years of quarterly data prior to 2001. Only firms with the financial year ends falling in March, June, September and December were kept, dropping 19% of the sample. This was done to enable the construction of quarterly periods corresponding to the same three calendar months. Quarterly investment rates were windsorized at -100% and +100%, affecting 0.12% of the observations, to minimize the impact of large outlier values. The final sample covers 23,625 quarterly observations from 375 firms.

The standard deviation of quarterly investment was normalized at the quarterly level based on the pre-2001 SD of investment. This is undertaken because quarterly accounts

have a quarterly standard-deviation trend induced by auditors balancing quarters within each accounting year. Without this normalization the drop after 9/11 is actually even higher, rising from 0.035 to 0.042. This is because the 4th quarter accounts have the highest standard-deviation on average because most accounting years start on October 1st, so the fourth calendar quarter is the first accounting quarter. Quarterly accounts data also lag macro data by about half a quarter. This is because quarterly accounts are published within 2 months of the end of the quarter, which given the lags in both internal data collection and the external lags in accounts production and verification, means they are skewed towards the beginning of the quarter.

E. Appendix: Volatility measures

This appendix describes the construction of monthly volatility series on interest rates, exchange rates and oil prices, and their comparison to stock-market volatility measures.

Monthly interest rate volatility data is calculated as the daily standard deviation of the Moody's Aaa interest rate series (source Federal Reserve, Board of Governors). Monthly exchange rate volatility data is calculated as the average of the monthly standard deviations of the daily US\$ exchange rate with the French Franc, German Mark, Japanese Yen, Swiss Franc and UK Pound (source Federal Reserve Bank, New York). The average across five different monthly exchange rate volatilities is used to create a weighted average foreign exchange rate volatility measure. Monthly oil price volatility data is calculated using daily European Brent Spot Price (source Energy Information Administration). A regression is then run of log (monthly stock market volatility) on log (monthly volatility of interest rates), log (average monthly volatility of exchange rates) and log (monthly volatility of oil prices), spanning 207 months from March 1987. The regression coefficients (standard errors) are 0.117 (0.043), 0.152 (0.044) and 0.176 (0.030) respectively, with an R-squared of 0.224

F. 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 Juprette: a discrete analysis of scrapping subsidies", *Journal of Political Economy*, 108, 778-806.

ADDA, J. and COOPER, R. (2003), "Dynamic Economics", MIT Press, Cambridge.

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

ATTANASIO, O. (2000), "Consumer Durables and Inertial Behaviour: Estimation and Aggregation of (S,s) Rules for Automobile Purchases", *Review of Economic Studies*, **67**, 667-696.

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

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**, pp 964-990.

BASU, S., FERNALD, J. and KIMBALL, M. (2006), "Are technology improvements contractionary" forthcoming *American Economic Review*.

BECKER, R., HALTIWANGER, J., JARMIN, R., KLIMEK, S. and WILSON, D. (2004). "Micro and macro data integration: the case of capital", Bureau of the Census.

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

BERNARD, A., REDDING, S. and SCHOTT, P. (2005), "Multi-product firms and the dynamics of product mix", Dartmouth mimeo.

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

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

BEWLEY, T., (1999) *Why Wages Don't Fall During a Recession*, Cambridge and London: Harvard University Press.

BILS, M. and KLENOW, P (2004), "Some evidence on the importance of sticky prices", *Journal of Political Economy*, **112**, 947-985.

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

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, pp 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. (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.

CABALLERO, R. and ENGEL, E. (2004), "A comment on the economics of labor adjustment: mind the gap: reply", *American Economic Review*, 94, 1238-1246.

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.

COGLEY, T and SARGENT, T. (2002), "Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII US", *Review of Economic Dynamics*, **8**.

COOLEY, T (eds) (1995), *Frontiers of Business Cycle Research*, Princeton University Press, Princeton, New Jersey.

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. (2003), "On the nature of the capital adjustment process", NBER WP7925.

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., and WILLIS, J. (2004), "A comment on the economics of labor adjustment: mind the gap", *American Economic Review*, 94, 1223-1237.

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

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.

DOMS, M. and DUNNE, T. (1998), "Capital Adjustment Patterns in Manufacturing Plants", *Review of Economic Dynamics*, **67**, 117-145.

EBERLY, J. and VAN MIEGHEM, J. (1997), "Multi-Factor Dynamic Investment under Uncertainty", *Journal of Economic Theory*, **75**, 345-387.

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.

GALI, J. (1999), "Technology, employment and the business cycle: Do technology shocks explain aggregate fluctuations?", *American Economic Review*, **89**, 249-271.

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

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

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

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

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

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

HOWARD, R (1964), "Dynamic Inference", *Research in the Control of Complex Systems*, **10**, 89-104.

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

JUNG, J. and SHILLER, R. (2002), "One simple test of Samuelson's dictum for the stock market", NBER working paper 9348.

JUSTINIANO, A and PRIMICERI, G. (2005), "The time varying volatility of macroeconomic fluctuations", Northwestern mimeo.

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.

KLENOW, P and KRYVTSOV, O. (2005), "State-dependent or time-dependent pricing: does it matter for recent U.S. inflation?", Stanford mimeo

KOCHERLAKOTA, N. (1996), "The equity premium: it's still a puzzle", *Journal of Economic Literature*, **34**, 42-71.

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), 1-32.

MEI, J. and GUO, L. (2002), "Political Uncertainty, Financial Crisis and Market Volatility", NYU mimeo.

MERZ, M. and YASHIV, E. (2005), "Labor and the Market Value of the Firm", LSE mimeo.

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*, vol. 91(2), pp. 34-38.

POTERBA, J. and SUMMERS, L. (1986), "The persistence of volatility and stock market fluctuations", *American Economic Review*, **76**, 1142-1151.

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

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

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

SIMS, C, (2001) "Stability and instability in US monetary policy behavior", Princeton University mimeo.

SIMS, C, (2005) "Were there regime switches in US monetary policy", Princeton University mimeo

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

TAUCHEN, G. (1986), "Finite state markov-chain approximations to univariate and vector autoregressions", *Economic Letters*, **20**, pp177-181.

TAUCHEN, G. and HUSSEY, R. (1991), "Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models", *Econometrica*, **59**, pp371-396.

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

TREJO, S. (1993), "Overtime pay, overtime hours and labor unions", *Journal of Labor Economics*, **11**(2), 253-278.


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

VOTH, H. (2002), "Stock Price Volatility and Political Uncertainty: Evidence from the Interwar Period", MIT mimeo.

VUOLTEENAHRO, T. (2002), "What drives firm-level stock returns?", *Journal of Finance* LVII(1), Applied Econometrics, S63-S84.

WOLFERS, J. and ZITEWITZ, E. (2006), "Using Markets to Inform Policy: The Case of the Iraq War", Wharton mimeo.


Figure 1: Monthly US stock market volatility



Notes: CBOE VXX index of % implied volatility, on a hypothetical at the money S&P100 option 30 days to expiry, from 1986 to 2004. Pre 1986 the VXX 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 VXX index when they overlap (1986-2004). Actual and VXX 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 VXX index, generating an average volatility of 58.2 for 9/11 until 9/14 inclusive.


* For scaling purposes the monthly VXX was capped at 50 for the Black Monday month. Un-capped value for the Black Monday month is 58.2. A brief description of the nature and exact timing of every shock is contained in Appendix

Figure 2: Frequency of the word “uncertain” in the FOMC minutes (% of all words)




Source: [count of “uncertain”/count all words] in minutes posted on <http://www.federalreserve.gov/fomc/previouscalendars.htm#2001>

Figure 3: Hiring/firing and investment/disinvestment thresholds



Notes: Simulated thresholds using the adjustment cost estimates "All" in column (2) of table 2. All other parameters and assumptions as outlined in Sections 2 and 3.

Figure 4: Thresholds at different levels of uncertainty



Notes: Simulated thresholds using the adjustment cost estimates “All” in column (2) of table 2. “Low uncertainty” defined as $\sigma=19\%$ and “high uncertainty” defined as $\sigma=37\%$. All other parameters and assumptions as outlined in sections 2 and 3.

Figure 5a: Aggregate net hiring rate (%)





Figure 5b: Percentiles of firm net hiring rates (%)



Notes: Simulations run with 250,000 firms, each with 250 plants, at a monthly frequency. Adjustment costs are set at the estimated “All” values in column (2) of table 2. All other parameters and assumptions as outlined in sections 2 and 3. Macro uncertainty shock ($S_{t=0}=1$) in period 0, otherwise ($S_{t\neq 0}=0$).

Figure 6a: Aggregate investment rate (%)





Figure 6b: Percentiles of firm investment rates (%)



Notes: Simulations run with 250,000 firms, each with 250 plants, at a monthly frequency. Adjustment costs are set at the estimated "All" values in column (2) of table 2. All other parameters and assumptions as outlined in sections 2 and 3. Macro uncertainty shock ($S_{t=0}=1$) in period 0, otherwise ($S_{t\neq 0}=0$).

Figure 7a: Productivity growth rates (%)




Figure 7b: Productivity and hiring, period before the uncertainty shock






Figure 7c: Productivity and hiring, period after the uncertainty shock



Notes: Simulations run on 250,000 firms, each with 250 plants, at a monthly frequency. Adjustment costs are set at the estimated "All" values in column (2) of table 2. All other parameters and assumptions as outlined in sections 2 and 3. Macro uncertainty shock ($S_{t=0}=1$) in period 0, otherwise ($S_{t\neq 0}=0$). "Total", "Between", "Within" and "Cross" productivity growth defined as in section 5.2 following Foster et al. (2000).

Figure 8a: Risk-aversion effects




9/11 uncertainty shock

risk-neutral

risk-averse

Figure 8b: Combined 1st and 2nd moment shocks



Notes: Simulations run with 250,000 firms, each with 250 plants, at a monthly frequency. Adjustment costs are set at the estimated "All" values in column (2) of table 2. Macro uncertainty shock ($S_{t=0}=1$) in period 0, otherwise ($S_{t\neq 0}=0$). All other parameters and assumptions as outlined in sections 2 and 3, except: in top panel (8a) discount rate set at $0.06 + 0.03 \times \sigma_{it}/26\%$; in bottom panel (8b) -5% demand conditions shock also occurs in period 0 for the "1st and 2nd moment shock" case.

Figure 9a: Aggregate hiring, Partial Irreversibilities only




Figure 9b: Hiring percentiles, Partial Irreversibilities only




Figure 9d: Aggregate hiring, Fixed Costs only




Figure 9e: Hiring percentiles, Fixed Costs only





Figure 9f: Productivity growth, Fixed Costs only



Notes: Hiring and productivity growth rates are % at annualized values. Simulations run on 250,000 firms, with 250 plants at a monthly frequency. In top panels (a, b and c) assume 10% partial irreversibilities only ($PR_L=0.1$ and $PR_K=0.1$ and all other adjustment costs zero), and in bottom panels (d, e and f) 1% fixed costs only ($FC_L=0.01$ and $FC_K=0.01$ and all other adjustment costs zero). All other parameters as in sections 2 and 3. Macro uncertainty shock ($S_{t=0}=1$) in period 0, otherwise ($S_{t=0}=0$). “Total”, “Between”, “Within” and “Cross” productivity growth defined in section 5.2 following Foster et al. (2000).

Figure 10a: VAR estimation of the impact of a volatility shock on industrial production





Figure 10b: 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 1 for the 15 labelled shocks in Figure 1, and 0 otherwise. Variables (in order) are log industrial production, log employment, hours, inflation, federal funds rate, log stock market levels and the volatility shock indicator. All data detrended using a Hodrick-Prescott filter with smoothing parameter of 14400.

Figure 11b: VAR robustness to different shock definitions





Figure 11a: VAR robustness to different variable sets and ordering



Notes: 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 14400. In top panel variables (in order) are log industrial production, log employment, hours, inflation, federal funds rate, log stock market levels and the volatility indicator. 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 var as noted.

Figure 12a: Quarterly total private employment, thousands net-change¹





Figure 12b: Quarterly gross private investment, % contribution to real GDP growth²



¹ BLS Current Employment Statistics survey, total private employees (1000s), seasonally adjusted, quarterly net change. Series CES0500000001.

² BEA NIPA, gross private domestic investment contribution to real GDP growth, seasonally adjusted, quarterly at annualized values. Table 1.1.2.

³ Federal Reserve Bank of Philadelphia's "Survey of Professional Forecasters", taken quarterly from 33 economic forecasters, www.phil.frb.org

Figure 13a: Cross Sectional standard deviation of investment rates





Figure 13b: Investment rate histogram, 2001 Q3 (before 9/11)



Figure 13c: Investment rate histogram, 2002 Q1 (after 9/11)



Source: Compustat quarterly investment rates (%). Numerator equals plant, property and equipment purchases less resales, plus net change in inventories; denominator equals total stock of net fixed assets plus inventories averaged over the current and prior quarter. Balanced panel of all 375 publicly quoted manufacturing firms with at least \$20m average sales and complete quarterly data from 1990 to 2005. The standard deviation (SD) of quarterly investment has been normalized at the quarterly level based on the pre-2001 SD of investment. This is undertaken because quarterly accounts have a quarterly SD trend induced by auditors balancing quarters within each accounting year, generating a falling SD trend through the accounting year. Without this normalization the drop after 9/11 rises from 0.035 to 0.042, as the 4th quarter accounts have the highest SD on average (most accounting years start on October 1st). Quarterly accounts data also lags macro data by about half a quarter. Quarterly accounts are published within 2 months of the end of the quarter, which given the lags in reporting and accounts preparation, means they are skewed towards the beginning of the quarter.