Buried by aggregation: excavating the dynamics of investment, employment and uncertainty

Nicholas Bloom*

June 2004

Abstract

Micro evidence on investment and employment has uncovered stark evidence of lumpy intermittent behaviour. This is consistent with non-convex adjustment cost models in which uncertainty plays a central role in driving firm behaviour. This lumpy micro-behaviour, however, is buried by cross-sectional and time series aggregation at the firm level, obscuring any underlying structural model. To excavate beneath this aggregation we construct a model with capital, labour, and time varying uncertainty, which we explicitly aggregate to yearly firm values and structurally estimate against Compustat using indirect inference. This model then provides a laboratory to examine the impact of policy changes and shocks on investment and employment. Three experiments are considered: (i) temporarily raising uncertainty as in OPEC I and II, and 9/11/2001; (ii) permanently increasing labour adjustment costs from US to EU levels; and (iii) cutting capital taxes by 1%.

Acknowledgement: I am grateful for comments on earlier versions of this paper from Jerome Adda, Steve Bond, Thomas Buettner, Russel Cooper, Peter Davis, Jan Eberly, Mike Elsby, Mark Schankerman, and John Van Reenen, and seminar participants at London School of Economics, London Business School and UCL.

Correspondence: n.bloom@lse.ac.uk

*London School of Economics CEP and CEPR (n.bloom@lse.ac.uk)

1. Introduction

There is now a large body of evidence demonstrating that investment and employment at the plant level is lumpy and intermitted rather than smooth and continuous. This is starkly documented by, for example, Doms and Dunne (1993) for investment and Davis and Haltiwanger (1992) for employment[1]. This type of lumpy behaviour is consistent with models of non-convex adjustment costs, where adjustment rules take on a threshold form, with infrequent rates of adjustment and rich time series dynamics[2].

In these models uncertainty plays a major role in determining the gap between thresholds and the responses of investment and employment, and so is central to the impact of these non-convexities on firm behaviour. The aim of this paper is to understand the impact of uncertainty on the adjustment dynamics of investment and employment. We do this by constructing a rich model of firm behaviour which we use to try and estimate the underlying demand and adjustment cost parameters. This approach is innovative in that we explicitly model and address three central issues which play an important role in shaping observed behaviour under uncertainty:

Time Varying Uncertainty

The literature on investment and uncertainty has, for analytical tractability, focused on cross-sectional changes in uncertainty rather than time series variation. While this is essential for delivering analytical results it misses out on important time series variation. For example, uncertainty varies over time at the macro level due to shocks like OPEC I and II, and 9/11, at the industry level due to shocks like regulatory changes and new technologies, and at firm level due to shocks from demand and factor supplies. In this paper we model an underlying demand process with time varying uncertainty.

Labour and Capital Adjustment Costs

The empirical literature on investment and labour demand has, again for tractability, focused on estimating either investment or labour demand by making

[1] Other work on investment includes Coooper and Haltiwanger (1993), Caballero, Engle and Haltiwanger (1995), Cooper, Haltiwanger and Power (1999) and Cooper and Haltiwanger (2003), and on employment Hammermesh (1989), Caballero and Engel (1993), Caballero, Engel and Haltiwanger (1997) and Cooper, Haltiwanger and Willis (2004)..

[2] See, Bertola (1988), Pindyck (1988), Dixit (1989), Abel and Eberly (1996), and Eberly and Van Mieghem (1997) on models of (partial) irreversibility, Caballero and Leahy (1996) on fixed cost models of investment and and Abel and Eberly (1994) who combine these together in a Tobin's Q framework.

the assumption that the other factor of production are perfectly flexible[3]. But both capital and labour are typically costly to adjust, and potentially becoming increasingly so as workers become more skilled and ICT goods become an increasing share of physical capital. Since capital and labour are linked in production, modelling their join dynamics will be important. In this paper we jointly solve for the optimal investment and employment behaviour with adjustment costs for each factor.

Aggregation across time and units

While micro-level data appears to be lumpy and intermittent firm data is smooth and continuous, suggesting that cross-sectional aggregation obscures the underlying non-convexities. Furthermore, even at the truly micro-level there is the additional problem of temporal aggregation, because the frequency of data collection is generally much lower than that of decision making. Firms often make investment and employment decisions on a weekly or monthly basis but data is typically collected quarterly or annually[4]. We explicitly build aggregation across units and across time into our model and estimation routine.

The paper starts in section 2 by building a model of firm behaviour, section 3 discusses our data and estimation routine. Section 4 discusses these results, while section 5 provides some policy simulations while section 6 concludes.

2. The Model

2.1. Overview

We model a firm as a collection of a very large, but fixed, number of production units. Each unit faces an iso-elastic demand curve for its product which is produced with Cobb-Douglas technology. Both demand and technology are affected by multiplicative shocks described by a geometric random walk with time varying drift and uncertainty. These shocks have a unit specific idiosyncratic component and a common firm component. We work in discrete time.

[3] One of the exceptions was Shapiro (1986), who estimated a joint investment and employment demand equation with quadratic adjustment costs for each factor.

[4] For evidence of cross–sectional and time series aggregation see tables (A.1) and (A.2) in the Appendix. From this it appears that even within lumpy micro-level data there is evidence of cross-sectional aggregation with zeros occuring much more frequently in smaller single production units, reflecting the fact that even establishment level data may conceal several different plants each with multiple production lines.

Firms can adjust their capital stock, labour force and average hours. Capital and labour adjustment entails adjustment costs, while hours can be freely raised or lowered, but at the penalty of a higher hourly wage rate. Since our aim is to model infrequent and lumpy investment and employment at the micro-unit level these adjustments costs include a fixed cost and partial irreversibility component, as well as a more traditional convex cost component. As is standard in non-convex models the resulting microeconomic policy is one of inaction interspersed with periods of (potentially large) investment and employment bursts. Aggregation across units within the firm and across time up to yearly values obscures this lumpiness but nevertheless plays an important role.

2.2. The Production unit

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

$$R(Y, K, L, H) = Y^{1-\alpha-\beta} K^{\beta} (L \times H)^{\beta} \tag{2.1}$$

which nests a Cobb-Douglas production function in capital (K), labour (L) and hours (H) and an iso-elastic demand curve. Demand and technology conditions are combined into an index (Y) - henceforth called demand conditions[5]. Wages are determined by undertime and overtime hours around the standard weekly norm of 40 hours, following the form and parameters used in Cooper, Haltiwanger and Willis (2004)

$$w(H) = w_0 + w_1 \times (H - 40) + w_2 \times (H - 40)^2 \tag{2.2}$$

Working hours can be costlessly adjusted (at the cost of higher hourly wages rates), so overtime acts as a short-run pressure valve in response to demand changes.

The demand conditions parameter Y evolves as a geometric random walk with a time-varying auto-regressive drift and variance

$$\begin{aligned}
Y_t &= Y_{t-1} \times (1 + \mu_t + \sigma_t V_t) & V_t &\sim N(0,1) && (2.3) \\
\mu_t &= \mu_{t-1} + \rho_\mu(\mu^* - \mu_{t-1}) + \sigma_\mu X_t & X_t &\sim N(0,1) && (2.4) \\
\sigma_t &= \sigma_{t-1} + \rho_\sigma(\sigma^* - \sigma_{t-1}) + \sigma_\sigma Z_t & Z_t &\sim N(0,1) && (2.5)
\end{aligned}$$

This generates permanent demand shocks with short run cycles of faster and slower growth, and short-run cycles of lower and higher uncertainty. Rapid growth and

[5]Note that this unitary homogeneity in (Y,K,L) can be achieved from any iso-elastic demand curve and Cobb-Douglas production function by an arbitrary redefinition of the demand and technology shocks.

high uncertainty this period implies rapid expected growth and higher uncertainty next period, but in the long run all rates converge in expectation to the average level. This is a generalisation (in discrete time) of the Geometric Brownian Motion typically used in modelling investment and uncertainty, and is consistent with the typical assumptions over the short and long run profile of firm growth in the finance valuation literature (i.e. Brealey and Myers, 2003, Copeland et al. 2003). It also implies a log-normal cross-sectional distribution of firm sizes, and also Gribrat's law that size and growth rates are independent, which while statistically rejected on large panels of data is a reasonable economic first approximation[6].

The third piece of technology determining the firms activities are the investment and employment adjustment costs. There is a long literature investigating and explaining investment and employment adjustment costs[7] which typically focuses on three terms, all of which we include in our specification:

1) Partial irreversibilities

The resale price of capital is usually less than the purchase price due to a combination of transactions costs, the market for lemons phenomena and the physical costs of resale. Labour partial irreversibility derives from hiring, training and firing costs. So we model

$$\begin{aligned} \text{Purchase price capital} = \quad & B > S \; = \text{Resale price capital} && (2.6) \\ \text{PDV hiring 1 worker} = \quad & H > F \; = \text{PDV saved by firing 1 worker} && (2.7) \end{aligned}$$

2) 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 the factory may need to close for a few days while a refit is occurring, or hiring new workers will require fixed costs of advertising, interviewing and training. We model this in terms of units of lost output

$$\begin{aligned} \text{Investment disruption costs} &= FCI \times R(Y, K, L, H) \times |I| > 0 && (2.8) \\ \text{Hiring disruption costs} &= FCE \times R(Y, K, L, H) \times |E| > 0 && (2.9) \end{aligned}$$

3) Quadratic adjustment costs

[6] See Scherer (1980) for evidence on the widespread phenomenon of log-normal distributions of firm sizes. Evans (1987) and Dunne et al. (1989) test Gibrat's law on a large sample of firms and statistically reject it, finding a typical coefficient of around -0.03 on the regression of growth in logged sales on lagged logged sales. This implies an auto-regressive logged sales coefficient of 0.97, which is economically close to our coefficient of 1 in the geometric random walk.

[7] See, for example, Nickell (1977) and Hammermesh (1996).

The costs of investment and employment may also be related to the rate of adjustment due to higher costs for more rapid changes

$$\begin{aligned}\text{Quadratic investment cost} &= \lambda_I K(\frac{I}{K})^2 && (2.10)\\ \text{Quadratic hiring costs} &= \lambda_E L(\frac{E}{L})^2 && (2.11)\end{aligned}$$

The combination of all adjustment costs is defined by the firm's adjustment cost function, $C(Y, L, I, E)$.

2.2.1. The Value Function

Investment and employment enters production immediately. At the end of each period capital and labour depreciates by δ_K and δ_L respectively so that, for example, $K_{t+1} = (K_t + I_t)(1 - \delta_K)$. The firms optimization problem can then be simplified by noting that the revenue function, adjustment cost function $C(Y, L, I, E)$, depreciation schedules and expectations operators are all jointly homogenous of degree one in (K, L, Y). This allows us to normalize by one state variable. Thus, after dividing (K, L, Y) by K the firms maximization can be stated as

$$\begin{aligned}Q(y_t, l_t, \mu_t, \sigma_t) &= \max_{i,e,H} R(\frac{y_t}{1+i}, \frac{l_t(1+e)}{1+i}, H) - C(y_t, l_t, i, e) - \frac{l_t(1+e)}{1+i} w(H)\\ &\quad + \frac{(1-\delta_K)(1+i)}{1+r} E[Q(y_{t+1}, l_{t+1}, \mu_{t+1}, \sigma_{t+1})] && (2.12)\end{aligned}$$

where lower case variables are $l = \frac{L}{K}$, $y = \frac{Y}{K}$, $i = \frac{I}{K}$ and $e = \frac{E}{L}$, and $Q(y, l, \mu, \sigma)$ is Tobin's Q. Normalizing by K is an essential step, which enables us to remove one major state variable from our problem, rendering our numerical solution routine computationally feasible.

2.3. Optimal investment and employment

The model is too complex to solve analytically and requires numerical simulation. However analytical results can be used to show the solution has a unique valued continuous solution and an (almost everywhere) unique policy function[8]. This

[8] Application of Stokey and Lucas (1989) for quadratic and partially irreversible investment costs, Caballero and Leahy (1996) for the extension to fixed costs, and Alvarez and Stokey (1998) for unbounded homogeneous functions.

means any numerical results we derive will be convergent to the unique solution. Results for an illustrative set of parameters are displayed in figure (2.1) overleaf to illustrate the resultant policy functions. The plot contains 4 lines plotted in $(\frac{Y}{K}, \frac{Y}{L})$ space for values of the fire and hire thresholds (left and right lines) and the sell and buy capital thresholds (top and bottom lines). The inner region is the region of inaction where ($E = 0$ and $I = 0$) while outside that region investment and hiring will be taking place according to the optimal values of E and I. The gap between the investment/disinvestment thresholds is higher than between the hire/fire thresholds due to the higher irreversibilities of capital.

Figure (2.2) displays the same lines for two different values of current uncertainty - $\sigma = 0.1$ in the "inner box" of lines and $\sigma = 0.3$ for the "outer box" of lines. It can be seen that the comparative static intuition that higher time varying uncertainty also increases real options is confirmed here, suggesting that large changes in σ can have quantitatively important impacts on investment and labour demand behaviour even with time varying uncertainty.

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 results 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 the project, 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.

2.4. Firm level investment

2.4.1. Cross-Sectional Aggregation

Aggregation up to the firm level is a difficult problem and several simplifying assumptions need to be taken. We assume each firm owns a number of production units. These units can be thought of as different production plants, different geographic markets, or different functions and divisions within the same firm. Production units are subject to both the common firm level demand process outlined above and a plant level idiosyncratic demand process with mean 0 and variance σ_U^2.

Firms own a sufficiently large number of these production units that any single unit level shock has no significant impact on firm behaviour - that is full

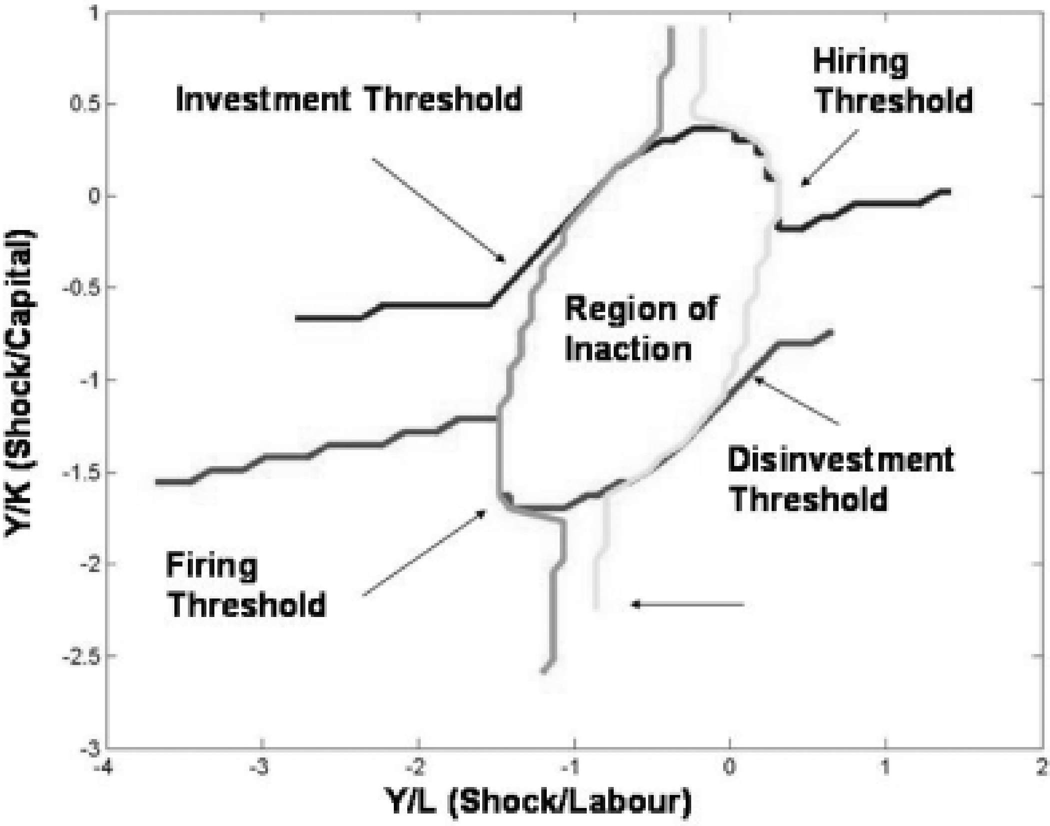

Investment Threshold
Hiring
Threshold
Region of
Inaction
Disinvestment
Threshold
Firing
Threshold
Y/K (Shock/Capital)
Y/L (Shock/Labour)
1
0.5
0
-0.5
-1
-1.5
-2
-2.5
-3
-4
-3
-2
-1
0
1
2

Figure 2.1:

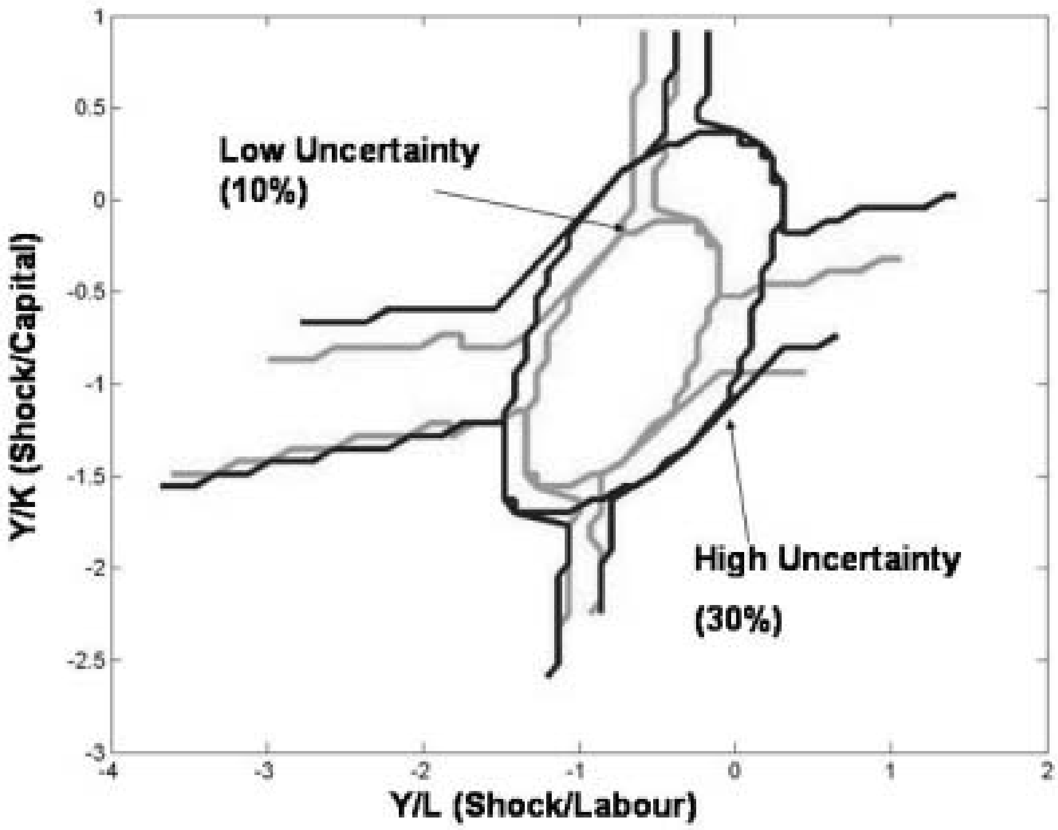
Low Uncertainty
(10%)
High Uncertainty
(30%)
Y/K (Shock/Capital)
Y/L (Shock/Labour)
1
0.5
0
-0.5
-1
-1.5
-2
-2.5
-3
-4
-3
-2
-1
0
1
2

Figure 2.2:

aggregation occurs. Units independently optimized to determine investment and employment, following the optimal behaviour outline above, with firm values being stochastic aggregates[9]. Thus all linkages across units within the same firm are modelled by common shocks to the demand process in levels, growth or uncertainty. Thus, to the extent that units are linked over and above these common shocks we are assuming they independently optimize due to bounded rationality and/or localized incentive mechanisms (i.e. managers being assessed only on their own units P&L).

Our intuition is that a relaxation of this assumption to allow more complex interaction between units within firms - for example through internal capital markets or sales to common output markets - would reduce the smoothing of aggregation. This will be similar in fashion to a lower variance in the unit level shock σ_U^2. Thus, while our estimated value for σ_U^2 may be downwardly biased, since this is estimated as a free parameter it will help to capture the first-order impact of linkages across units. Thus, the behaviour of the model should hopefully approximate a more complex model with some interactions between units. This is something, however, we also investigate for robustness in section (3.5.3) below.

2.4.2. Time series aggregation

We model demand and technology shocks, investment decisions and employment decisions on a monthly basis. Management meetings in a typical business-unit generally happen at quarterly, monthly or even weekly frequencies, with this depending on market conditions, technology and management practices. The main point for us is that they almost invariably happen at a higher frequency than that of data collection, and that we need to explicitly model this. Since one of our main variables of interest - uncertainty - is measured in Compustat using monthly returns variance we selected a monthly frequency for modelling purposes, although in section X we experiment with estimation on quarterly data to investigate the role of time-series aggregation.

To generate a modelled counterpart to Compustat "flow" figures found in the accounting P&L, such as sales and capex, are added up across the year while "stock" figures found in the balance sheet, like the capital stock and labour force,

[9] This follows the approach of Bertola and Caballero (1994), Caballero and Engel (1999) and Abel and Eberly (1999) who driven by the need for analytical tractability also assume a large number of indepedently maximising units or lines of capital to model macro, industry and firm level investment respectively.

are taken as year end values. Thus we aim to match up our modelled data as closely as possible to the Compustat empirical data.

3. Empirical Evidence

3.1. Data

The data is a panel of firm from the US Compustat dataset from 1983-2001 inclusive which was extracted and initially constructed using the same procedure as Bond and Cummins (2004). The data was cleaned to remove mergers and acquisitions by dropping observations with large jumps in the sales, employment and capital stock figures. The sample was also trimmed to focus on large continuing firms by keeping only firms with at least 15 years of data and over 500 employees. This was done to reduce the impact of entry and exit and focus on larger more aggregated firms[10]. The summary statistics are detailed in the table (3.1) below.

Table 3.1: Descriptive Statistics for the Compustat Data Set.

	mean	median	stan. dev.	min.	max.
$(I_{i,t}/K_{i,t-1})$ (investment rate)	0.191	0.144	0.189	-0.595	1.997
$\Delta Y_{i,t}/Y_{i,t-1}$ (real sales growth rate)	0.073	0.048	0.215	-0.658	1.971
$dL_{i,t}/L_{i,t-1}$ (employment rate)	0.052	0.020	0.223	-0.659	2.000
$Y_{i,t}$ (real sales, 1990 $M)	3855	652	12360	23.887	168,920
$L_{i,t}$ (employment, 1000's)	20.1	4.788	52.02	0.500	876.8
$\sigma_{i,t}$ (variance monthly returns,%)	33.13	32.32	11.66	6.581	80.779
observations per firm	17.3	18	1.04	15	18

3.1.1. The measurement of uncertainty

Measuring uncertainty is a tricky issue. What we need is a proxy variable which is well correlated with true uncertainty - σ_t^2 - which can be generated in simulated

[10] While this focus on larger continuing firms reduces the need to model entry and exit decisions it does undoubtedly introduce a selection bias here. This selection halved the sample size from 1073 to 521 firms, although it only reduced the total number of employees and sales covered by 22.4% and 21.6% respectively.

data and also in the Compustat data, and is measurable at an annual frequency. Fortunately, there is one variable that meets these exacting requirements, which is the variance of monthly share returns within each (accounting) year. Monthly share returns are calculated in the simulation as net cash flow plus capital gains per month per $ of value, and its yearly standard deviation has a correlation with true uncertainty of 0.788. Monthly share returns are calculated in Compustat (data item RIXM) as net cash-flow (dividends, buy-backs and rights issues) plus capital gains per $ of value[11]. Thus our uncertainty proxy is both well correlated with true uncertainty and also measured on the same basis and at the same frequency in Compustat and in the simulated data.

3.2. Estimation Approach

This model has no analytical closed form so can not be written down in a form which would enable standard regression estimation. Instead estimation of the underlying parameters is achieved by a methodology called indirect inference (see Gourierioux, Monfort and Renault 1996, and Smith, 1993). The basis of this methodology is a regression (hereafter termed the reduced form regression) which is run on both actual and simulated data. The simulated data set is created by solving the dynamic programming problem given a vector of parameters. The resulting policy functions are then used to create a panel data set comparable to Compustat. The structural parameters are chosen so that the coefficients of the reduced form regression from the simulated data are "close" to the estimates from the actual data.

The choice of a reduced form regression is thus a crucial piece of the analysis. For our purposes, the reduced form regression needs to satisfy two criteria. First, the parameters of the regression should be "informative" about our underlying structural parameters. That is, as the structural parameters are varied the regression coefficients should be responsive[12]. Second, the reduced form regression should summarize relevant aspects of the investment and employment decisions.

[11] This share returns variance measure has been previously used by Leahy and Whited (1996) and Bloom, Bond and VanReenen (2003). Share returns volatility are also significantly correlated at between 0.3 to 0.4 (depending on the calculation) with the dispersion of IBES analysts forecasts for each firm-year, which is the other typical firm-level uncertainty proxy used in the literature (see Bond et al, 2004).

[12] The formal condition is that there exists a one-to-one mapping between the structural parameters and the moments calculated from the data (the conceptual parallel to the rank condition in standard regression).

As emphasized above, the basic insights of the recent literature on non-convexities is that they imply rich cross-sectional non-linearities and time-series dynamics in the relationship between investment, employment and fundamentals, and uncertainty plays a central role in this. This is used to focus on two sets of reduced form regression equations: the first estimates the dynamics and non-linearities of investment and employment responses, and the second estimates the impact of uncertainty levels and changes on investment and employment.

Focusing first on the on the dynamics and non-linearities we estimate

$$i_{i,t} = \alpha_1 i_{i,t-1} + \alpha_2 s_{i,t-1} + \alpha_3 s^+_{i,t-1} + f_i + \eta_t \tag{3.1}$$
$$e_{i,t} = \beta_1 e_{i,t-1} + \beta_2 s_{i,t-1} + \beta_3 s^+_{i,t-1} + f_i + \eta_t \tag{3.2}$$

where $i_{i,t}$, $e_{i,t}$ and $s_{i,t}$ are the investment, employment and sales growth rate of firm i in period t. The term $s^+_{i,t-1}$ is the sales growth rate multiplied by a indicator which is 1 for positive sales growth rate and zero otherwise. This is included to pick up any non-linearities in response to positive and negative sales changes, reflecting the importance of asymmetric responses in Compustat[13]. Lagged values are used to prevent the feedback from current investment and employment decisions into current sales, which would otherwise drive a strong positive correlation as is clear from equation (2.1). Finally the f_i terms are included to control for unobserved heterogeneity[14], while the η_t are included to control for macro movements in factor prices like interest rates and the role of the business cycle. The ability to condition for this in an indirect inference framework motivates our choice of this as an estimating strategy over matching unconditional moments.

Focusing secondly on the impact of uncertainty levels and changes we estimate

$$i_{i,t} = \gamma_1 sd_{i,t} + \gamma_2 \Delta sd_{i,t} + f_i + \eta_t \tag{3.3}$$
$$e_{i,t} = \delta_1 sd_{i,t} + \delta_2 \Delta sd_{i,t} + f_i + \eta_t \tag{3.4}$$

[13] An alternative specification which we experimented with was instead to include a sales growth squared term following, for example, Cooper and Haltiwanger (2002). We found in our Compustat data, however, that the coefficient on this squared term was extremely sensitive to the data cleaning routine used to remove mergers and acquisitions (and very large measurement errors) while the co-efficient on $s^+_{i,t-1}$ was not.

[14] Clearly there is unobserved heterogeneity in Compustat. To match our model with data requires us either to build the unobserved differences across plants into our analysis or to purge Compustat of these. We have chosen the latter approach, although in principle an exercise explicitly trying to build in heterogeneity through parameter variation across firms could be undertaken.

where $sd_{i,t}$ and $\Delta sd_{i,t}$ are the level and change in logged returns uncertainty. This specification is chosen because of the central role that uncertainty plays in determining micro unit level thresholds, with γ_1 and γ_2 (δ_1 and δ_2) the elasticity of investment (employment) to uncertainty. Again the f_i terms are included to control for unobserved heterogeneity[15] while the η_t terms take out common macro changes. Current values Combining all four reduced form regression equation together yields ten parameters in total.

3.3. Estimation Procedure

The econometric problem consists of estimating the parameters that characterize (a) firms production function and mark-up, (b) the distribution of adjustment costs, (c) the wage curve, (d) the firm demand process, and (e) the unit level demand process. For tractability we need to limit the number of parameters being estimated and so take the starting values from the literature or Compustat wherever possible and use iteration to check their impact on the results.

These starting values (and their sources) are: (i) hiring cost of 1 month and firing cost of 2 months wages (Nickell, 1986 and XXXX); (ii) capital resale loss of 25% (Ramey and Shapiro, 2001); (iii) mean demand drift rate 5% (Compustat, average growth of real sales); (iv) mean-reversion in demand drift of 0.65 (Compustat, calculated from 5 year correlation of sales growth[16]); and (v) mean, variance and mean-reversion of demand uncertainty 50%, 30% and 0.30 (Compustat, estimated from a Garch regression on monthly returns index data).

This leaves six parameters for estimation $\Theta = (FCI, FCE, \lambda_I, \lambda_E, \sigma_\mu^2, \sigma_U^2)$, the fixed and quadratic adjustment cost parameters for investment and employment, the variance of the demand drift and the unit level uncertainty. With ten reduced form regression parameters and six estimation parameters the estimation is over-identified, and these parameters can be estimated as follows. For an arbitrary value of Θ the dynamic program is solved and policy functions are generated. using these policy functions the decision rule is simulated and given arbitrary initial conditions to create a simulated Compustat. The reduced form regressions (3.1),

[15] Clearly there is unobserved heterogeneity in Compustat. To match our model with data requires us either to build the unobserved differences across plants into our analysis or to purge Compustat of these. We have chosen the latter approach, although in principle an exercise explicitly trying to build in heterogeneity through parameter variation across firms could be undertaken.

[16] A 5 year gap was taken to allow the impact of adjustment costs on the dynamics of sales to dissipate.

(3.2), (3.3) and (3.4) are then estimated on the simulated Compustat. Let $\Psi^S(\Theta)$ denote this vector of stacked regression coefficients ignoring the fixed effects, and further let Ψ^C represent the same stacked vector of coefficients from Compustat.

The estimated value $\widehat{\Theta}$ then minimizes the weighted distance between the actual and simulated regression coefficients. Formally, we solve

$$\Gamma(\Theta) = \min_{\Theta}[\Psi^C - \Psi^S(\Theta)]' W [\Psi^C - \Psi^S(\Theta)] \tag{3.5}$$

where W is a weighting matrix. We use the optimal weighting matrix given by the inverse variance-covariance of the regression coefficients from our true Compustat regressions (3.1), (3.2), (3.3) and (3.4), calculated through bootstrapping[17].

Of course, the $\Psi^S(\Theta)$ function is not analytically tractable. Thus, the minimization is performed using numerical techniques. Given the potential for discontinuities in the model and the discretization of the state space, we used a simulated annealing algorithm to perform the optimization.

3.4. Results

The model is simulated on a grid of $(\frac{Y}{K}, \frac{L}{K}, \mu, \sigma)$ of dimension (40,40,3,3) with an hours grid of (36, 40, 44). The number of units per firm is set at 100 as a computationally feasible approximation to full aggregation[18]. This currently takes about 2 hours to converge, which is too slow for simulated annealing so **RESULTS ARE CURRENTLY "CALIBRATED" RATHER THAN ESTIMATED** at present while we speed up the program and wait for a more powerful PC (presently 2003 vintage)**.** Calibration results are based on experimentation with around 40 different parameter choices, thus are not fully optimized with no standard-errors or over-identification tests.

For interpretation alongside the results for the calibrated parameter set we also display results from two illustrative parameters sets:

Investment Costs Only

This model assumes labour is fully flexible and only capital incurs adjustment costs. This is included to understand the impact of only modelling the adjustment costs for one factor of demand (capital in this example) on the behaviour of both factors. That is, to what extent will making the common assumption in the

[17] The bootstrap estimates these regressions 1000 times on Compustat data (drawn with replacement) to generate the VCV between these 1000 coefficient vector estimates.

[18] It appears that adding any more units beyond 100 has no significant impact on the reduce form regression parameters, and so mimics full aggregation. We are currently looking into this.

literature that labour is fully flexible and only capital incurs adjustment costs generate sufficient smoothing and non-linearities to fit both the investment and employment series? In this draft version of the paper we present results for same parameters set as our calibrated model except without any employment costs. While making direct comparisons easier, this analysis of the investment costs only case is not a "fair" test of the model since the parameters are not optimally chosen. So the next steps when the speeded up simulation is running will be to estimate the optimal parameters for an investment only adjustment cost function and compare this directly to our results for investment and employment adjustments costs.

Quadratic Costs, no aggregation

This model assumes that only quadratic costs for investment and employment occur (so no fixed-costs or partial irreversibilities) and that the firm only operates one unit. This is included to understand to what extent would a typical convex adjustment costs model of firm fit our investment and employment data. By assuming quadratic adjustment costs at the firm level we can generate smooth firm level data without resorting to aggregation, so this tests to need for aggregation to generate smoothing. In addition with quadratic-costs only uncertainty has no real-options impact so a second test is on the impact of uncertainty in the model. That is if we remove any real options effects of uncertainty what is its role in this model in explaining investment and employment.

The parameter vectors for the calibrated model and the two comparison models are outlined in table (3.2) below. This includes both the free calibrated parameters and the values for prior fixed parameters if these are changed across models.

Table 3.2: Key parameters in the simulated model

Parameters	Free (Calibrated)						Prior		
	$\boldsymbol{\sigma}_{\mu}^2$	$\boldsymbol{\sigma}_U^2$	$\boldsymbol{\lambda}_I$	$\boldsymbol{\lambda}_E$	**FCI**	**FCE**	$S-B$	$H+F$	units
Calibrated	5%	30%	0.75	0	0	2/52	0.25	3/12	100
Investment costs only	5%	30%	0.75	0	0	0	0.25	0	100
Quadratic costs only	5%	30%	0.75	0.5	0	0	0	0	1

3.4.1. Investment and Employment dynamics

Table (3.3) reports the estimated reduced form regression coefficients for Compustat and our three versions of the model. Turning first to Compustat there is clear evidence of strong dynamics in investment, which even after removing fixed effects, shows a large coefficient on lagged investment as well as lagged sales. Employment appears to show less dynamics with a low negative coefficient on the lag, although this is still significant. One issue for estimating dynamic employment equations is measurement error[19]. Since the employment rate is a differenced transformation of employment levels, measurement error in the levels will generate a negative (and potentially large) bias in the autoregressive employment rate coefficient, which provides one potential explanation for the seemingly low degree of auto-correlation which we explore below.

Looking next at asymmetries the large and significant coefficient on the positive sales growth term $s^+_{i,t-1}$for both investment and employment provides evidence of a strong non-linearities.

Table 3.3: Aggregation Obscuring Zero Investment Episodes.

Equation	Investment rate, $i_{i,t-1}$			Employment rate, $e_{i,t}$		
Coefficient	$i_{i,t-1}$	$s_{i,t-1}$	$s^+_{i,t-1}$	$e_{i,t-1}$	$s_{i,t-1}$	$s^+_{i,t-1}$
Compustat	0.191	0.034	0.121	-0.019	-0.093	0.204
	(0.021)	(0.014)	(0.016)	(0.023)	(0.022)	(0.024)
Calibrated	0.212	-0.021	0.139	0.034	-0.011	0.154
	(0.023)	(0.024)	(0.027)	(0.018)	(0.025)	(0.029)
Investment Costs Only	0.100	-0.016	0.249	0.150	-0.185	0.213
	(0.020)	(0.019)	(0.022)	(0.017)	(0.025)	(0.028)
Quadratic Costs, 1 unit	0.576	-0.045	-0.115	0.610	-0.121	-0.159
	(0.021)	(0.024)	(0.023)	(0.022)	(0.026)	(0.024)

Note: standard errors displayed in brackets below

For the calibrated model comparing to the investment results to Compustat the match is reasonably good, with a similarly strong dynamics (indicated by the coefficient on lagged investment) and a high level of asymmetries (indicated by the coefficient on $s^+_{i,t-1}$). The fit of the employment equation is far less satisfactory displaying a low but positive autoregressive coefficient and less asymmetries than implied by the Compustat data. One potential rationale is insufficiently high

[19] Of course by introducing fixed-effects was are also biasing down the coefficient on the lagged dependent variable, but this is of order $\frac{1}{T}$ (Nickell, 1983) and we have around 18 years per firm on average.

values for the fixed costs, which tends to generate negative auto-correlation. Although we could raise the value of this parameter the current value of 2 weeks of global revenue seems high so this is unappealing, and so instead in section (3.4.3) we investigate an alternative explanation based on measurement error.

In the third row we report the results for a model with investment costs only. Interestingly both the investment and the employment coefficients change, with investment becoming less autoregressive and more asymmetric because employment adjustment costs play an important role in driving investment responses. The reason is the auto-regressive investment coefficient is driven partly by lagged investment proxying for the distribution of Y/L (demand conditions/labour) across units between their hiring and firing thresholds. Since this distribution is autocorrelated, indirectly drives investment and is not proxied by sales it will be picked up partly by lagged investment. Without these employment adjustment costs the hiring and firing thresholds are the same, so there is no distributional dynamics, reducing the size of the lagged investment coefficient and so increasing the size of the direct sales coefficient. Turning to the employment equation, despite the absence of any adjustment costs, this still displays substantial dynamics with a large autoregressive coefficient and a strongly asymmetric response. This suggests that investment costs can lead through to significant employment dynamics.

Finally in the fourth row the quadratic adjustment cost model with only one unit per firm displays extremely strong autocorrelated dynamics for both investment and employment. It also appears to display a negative response to lagged shocks with negative asymmetries, although this is driven by the high-coefficient on lagged investment and employment. In the absence of any lagged investment and employment terms the coefficients on $s_{i,t-1}$ and $s^{+}_{i,t-1}$ rise to 0.425 (0.015) and 0.05 (0.02) for investment and 0.410 (0.016) and 0.021 (0.021) for employment, which is positive with small or insignificant asymmetry. Hence, while the quadratic model can generate smooth firm level investment behaviour without the need for cross-sectional aggregation is generates (at least for this parameter choice) investment and employment dynamics which do not match those observed in Compustat

3.4.2. Uncertainty levels and changes

Table (3.4) reports the estimates from the reduced form uncertainty levels and changes regression. Looking at row one for the Compustat regressions we can see a positive correlation with the level of uncertainty and a negative correlation with the change in uncertainty, and this is closely replicated in the second row for the calibrated model. The reason for the negative coefficient on the change in uncertainty is that higher uncertainty moves the thresholds apart and reduces the responsiveness of investment and employment. Since the median value of both of investment and employment is positive the short-run net effect of an increase (decrease) in uncertainty is positive (negative). This is supported by a sample split around median investment and employment in Compustat and the simulated data, where in all four cases the response to changes in uncertainty for the bottom 50% was much lower than for the top 50% and generally insignificant or positive[20]. The positive coefficient on the uncertainty levels term appears to derive from a correlation between large positive shocks and high values of current uncertainty.

In the third row the results for the model with investment costs shows a similar level of positive levels and negative coefficients for investment, but with a lower levels and insignificant change coefficients for employment. For investment this arises from the real-options effects on investment discussed above, but for employment since there are no adjustment costs this only arises from a weaker indirect effect from investment. Finally in the fourth row the signs are reversed for the quadratic model, which is due to the absence of any real options effects and concavity on the value function leading to a net negative levels effect.

Table 3.4: Investment and Employment Response to Uncertainty

Equation	Investment rate, $i_{i,t}$		Employment rate, $e_{i,t}$	
Coefficient	$sd_{i,t-1}$	$\Delta sd_{i,t-1}$	$sd_{i,t-1}$	$\Delta sd_{i,t-1}$
Compustat	0.026	-0.015	0.020	-0.007
	(0.007)	(0.005)	(0.010)	(0.007)
Calibrated	0.025	-0.009	0.026	-0.004
	(0.004)	(0.004)	(0.004)	0.004)
Investment Only	0.029	-0.009	0.004	0.007
	(0.003)	(0.005)	(0.004)	0.004)
Quadratic Costs, 1 unit	-0.011	0.036	-0.001	0.045
	(0.006)	(0.006)	(0.007)	0.006)

Note: standard errors displayed in brackets below

[20] In Compustat the coefficient (standard error) on lagged change in uncertainty for investment and employment below the median is 0.012 (0.0018) and 0.017 (0.004) respectively. In the callibrated data the same coefficients are -0.001 (0.001) and -0.001 (0.001).

3.4.3. Measurement error in employment levels

An issue for understanding the dynamics of employment is the role of measurement error. As a non-financial measure it is typically provided in company accounts to provide a guide to firm size and so is often heavily rounded (31% of employment figures end in 00 versus only 1% of sales figures). Since employment growth rates are generated from changes in employment levels errors in levels will generate negative correlations in first differences[21].

To understand the potential impact of this we added normally distributed iid measurement error with a mean absolute deviation of 2.5% (benchmarked for the median firm of 5,000 that rounds its employees numbers to the nearest against 50) to our simulated data. From table (3.5) it is clear that even this moderate level of measurement error generates negative auto-correlation, with the lagged coefficients in all three models falling substantially. There is also a much closer overall correspondence between Compustat and the calibrated model, suggesting a model with reasonable employment adjustment costs and small iid measurement error in employment levels could be consistent with observed data.

Table 3.5: Impact of measurement error on estimated employment dynamics

Equation	Employment		
Coefficient	$e_{i,t-1}$	$s_{i,t-1}$	$s^{+}_{i,t-1}$
Compustat	-0.019	-0.093	0.204
	(0.023)	(0.022)	(0.024)
Calibrated + measurement error	-0.037	0.075	0.193
	(0.016)	(0.025)	(0.029)
Investment Only + measurement error	0.025	-0.055	0.244
	(0.015)	(0.025)	(0.028)
Quadratic Only + measurement error	0.354	0.084	-0.056
	(0.020)	(0.025)	(0.025)

Note: standard errors displayed in brackets below

[21] For capital this is less of an issue since iid errors in the measured capex and the initial capital stock will not lead to any strong bias in the estimated dynamics.

3.5. Robustness Checks

3.5.1. Estimation on quarterly Compustat data

Compustat also provides quarterly data on sales and investment (but not employment), and the monthly returns index which can be used to generate a quarterly returns variance. Thus, to investigate the time aggregation issue further we simulate quarterly sales, investment and returns variance data and compare the investment regressions coefficients (3.1) and (3.3) to those from estimation on similar quarterly Compustat data.

Table 3.6: Quarterly estimation on actual and simulated investment

Equation	Investment rate, $i_{i,t}$			Investment rate, $i_{i,t}$	
Coefficient	$i_{i,t-1}$	$s_{i,t-1}$	$s^{+}_{i,t-1}$	$sd_{i,t-1}$	$\Delta sd_{i,t-1}$
Compustat					
Calibrated + measurement error					
Investment Only + measurement error					
Quadratic Only + measurement error					

3.5.2. Estimation by industry sector

3.5.3. Changing the magnitude of unit level uncertainty

4. Policy Simulations

This model can be used to study the micro distributional and macro aggregated investment and employment impact of shocks and policy changes[22]. The three experiments we run here are: firstly the effect of a large once-off temporary rise in uncertainty; secondly a permanent rise in the level of employment adjustment

[22] While is it important to model general equilibrium effects (see i.e. Thomas (2001), Veraciertio (2001) and Gilchrist and Williams (2000)) a detailed micro-macro can provide important insights from its ability to model a much richer range of adjustment costs and shocks.

costs, and finally a rise in taxation modelled by introducing taxes into the value function (**??**)..

4.1. Rise in uncertainty following a large macro-shock

We investigate the impact of a large temporary increase in uncertainty, induced by for example OPEC I and II and 9/11 type shocks, on the mean and distribution of investment and employment rates. To do this we take our preferred parameter choices outlined in section (3.4) and simulate a steady state economy for 1000 firms operating 100 units each. The model is initial run for 20 years (at a monthly frequency) to ensure it is in its ergodic steady state and then observed monthly for a further ten years during which the policy simulation is undertaken.

The steady state distribution of uncertainty, σ_t, takes three potential values in the steady state, which are 30%, 50% and 70% annual standard deviation in demand conditions. Firms face a monthly transition matrix between these states, which yields a steady-state distribution of uncertainty values of (25%,50%,25%) across the three values[23]. Hence, mean uncertainty is 50% while the bottom and top quartiles are 30% and 70% respectively. To simulate a high common uncertainty shock we shift all firms to the top value of 70% and hold them there for 4 months. To ensure a clean comparison to normal conditions no changes in the actual distribution of shocks occurs (which continues to be governed an underlying σ_t process). That is only firms perceptions of uncertainty rises but the actual variation of the demand process they face. This is undertaken to isolate the pure behavioural "real options" effects of uncertainty. That is any change in investment and employment outcomes are purely due to the uncertainty impact on behaviour rather than effects through demand conditions[24].

In figure (4.1) we plot the investment outcomes following the shock to uncertainty during the first four months of Year 5. In the top panel are displayed the 95th percentile of investment (left hand scale) and 5th percentile of investment (right hand scale). When uncertainty rises firms' investment and disinvestment thresholds move apart (as displayed in figure (2.2)) which reduces both levels of

[23]These uncertainty values and transition matrixes have been callibrated to ensure a similar mean, distribution and auto-correlation of Compustat standard-deviation of share returns to simulated standard deviation of share returns.

[24]We could have changed the value of the actual σ_t parameter but this makes interpretation hard as both uncertainty and the distribution of realized shocks are changing. By leaving the process for true σ_t constant, but changing firms perceptions, allows us to disentangle these two effects.

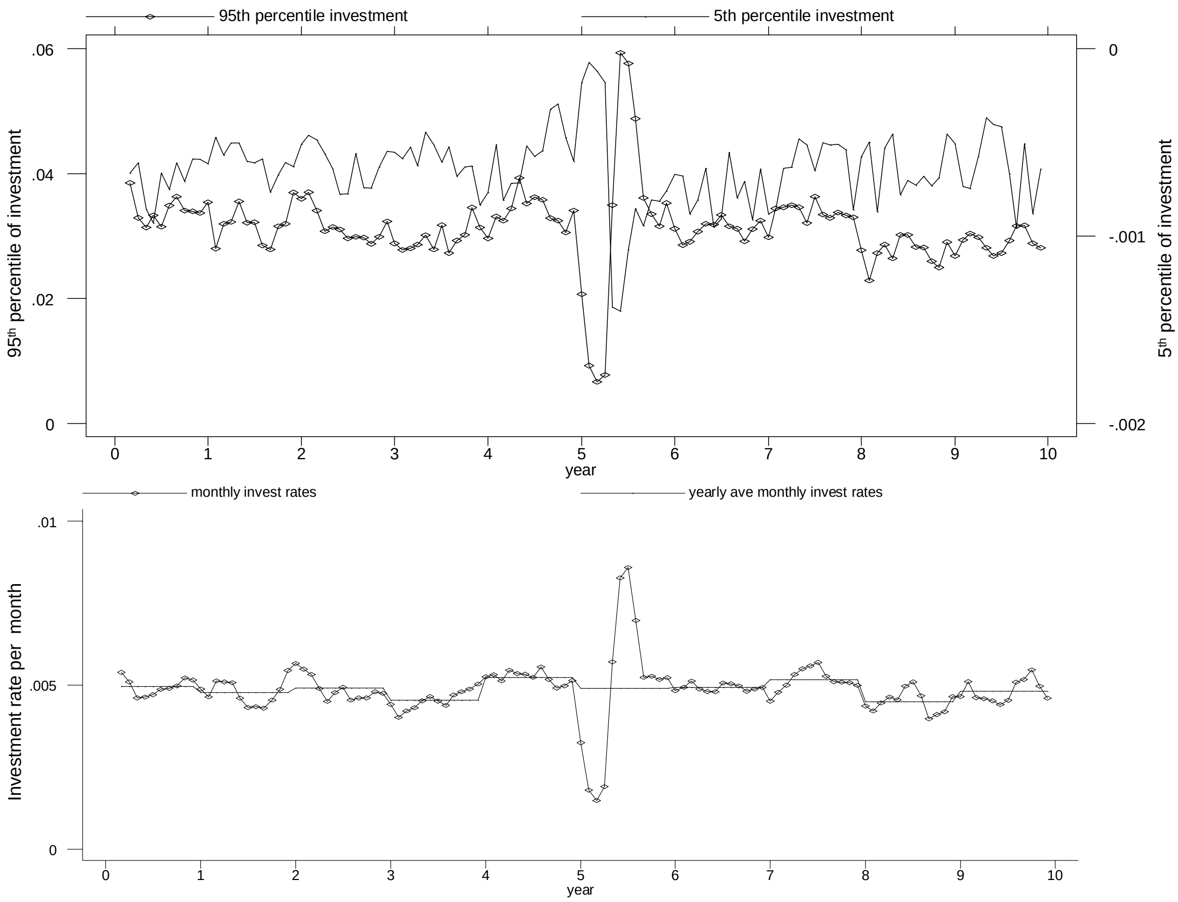

Figure 4.1: Investment rates before, during and after an uncertainty shock

both positive investment and negative disinvestment This can be seen to clearly reduce the level of investment at the 95th percentile during the shock (down from around 0.04% per month on average to around 0.01% per month) and also to reduce the absolute level of disinvestment at the 5th percentile (down from -0.01% per month on average to around -0.001% per month). But, after the shock has passed these investment percentiles rebound as the firms react to the pent-up backlog of demand accumulated during the period of inaction. This pause and then catch-up generates these s-shaped dynamic investment profiles.

This suggests that uncertainty changes can play a significant role in driving the distribution of investment rates across firms and industries. So for example, a declining industry may show a net positive response to higher uncertainty (as this reduces contractionary disinvestment) while an expanding industry will show a net negative impact (as this delays expansionary investment). For a policy-maker studying the data during or just after this shock this distributional analysis will be important in helping to distinguish between a longer-run slowdown or a short-run pause with a forthcoming rebound.

In the bottom panel of figure (4.1) we plot the average investment rates on a monthly basis and aggregated up to a yearly basis. Looking first at monthly investment we see the net impact of the rise in uncertainty is to reduce average investment. This is because the combined effects of depreciation and positive growth lead to positive average investment rates, so that a pause in investment reduces average overall levels of investment. The impact of this at a monthly level is pretty significant - during the uncertainty shock period investment falls to around 40% its average value and then rises to around 160% of its average value in the next 3 months during the rebound. Hence, these shocks can play a potentially powerful short-run role. But as can be seen from the average yearly values at lower frequencies this uncertainty shock appears to have little impact[25] as yearly data obscure much of the higher-frequency dynamics. However, this will nevertheless be important for policy making at higher frequencies where decisions need to be made on shorter-run data trends.

In figure (4.2) we plot the employment outcomes. Again in the top panel are displayed the 95th and 5th percentiles of net employment growth and in the bottom panel the monthly and yearly aggregate average values. These plots look similar to those for investment and the comments above apply in equal measure. The

[25] This is, in part, due to the positioning of the shock at the beginning of the year so that the investment fall and rebound are offset within the same year. But even if this occurred midway or towards the end of a year the observed impact would still be significantly reduced.

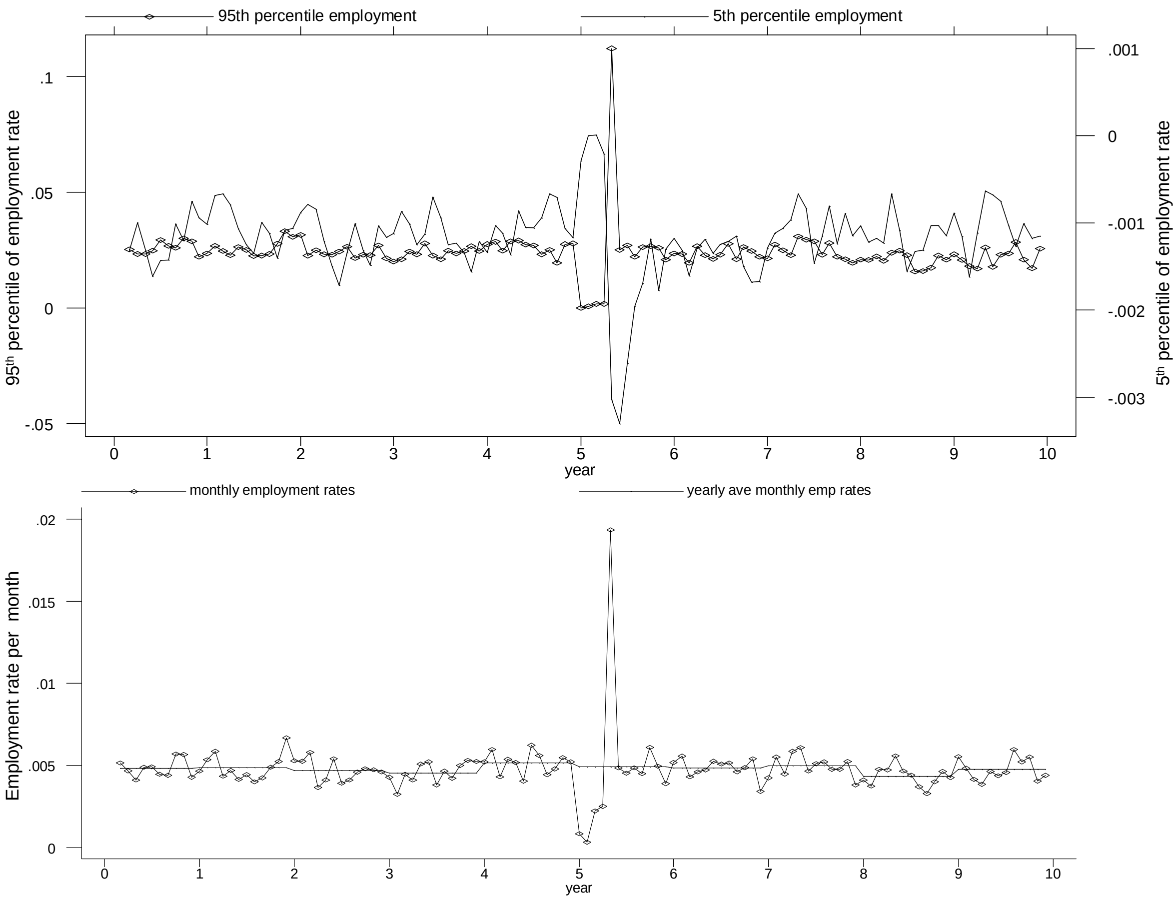

Figure 4.2: Employment rates before, during and after an uncertainty shock

one notable difference is the monthly time profile of the responses within the shock period and immediately afterwards. Whereas the investment response is slow and gradual the employment response is more sudden, particularly the rebound which occurs almost entirely in the month following the end of the uncertainty shock. This difference is driven by the differential adjustment cost parameters for capital and labour outlined in table (3.2) above.

Employment adjustment costs are mainly fixed costs with some partial irreversibilities, so the response tends to be "bang-bang" with lumpy quick adjustments, while the estimated investment adjustment costs contain important quadratic component which generates smoother slower responses. Finally the interaction of these two factors is also important, in that the investment response is also shaped by the employment response. In particular, investment falls more sharply in month 2 of the shock after employment has adjusted than month 1, and similarly its rebounds almost as strongly in the second month after employment has rebounded in the first month. Hence, both the differential individual factor adjustment cost parameters and their combined interactions appear to play an important role in shaping adjustment profiles for both factors.

Finally in figure (4.3) we plot the time series levels and changes for "revenue productivity". Revenue productivity is the empirical counterpart of the demand process Y, and is measured as $\log(sales_t) - \alpha \log(capital_t) - \beta \log(labour_t)$. Aggregate values for sales, capital and labour are generated from summing across all 1000 firms. This measure conflates the impact of demand shocks and productivity changes, but since both operate through same channel, the demand conditions Y, the revenue productivity measure will simulate the impact of uncertainty on true "total factor" productivity.

In the top panel of figure (4.3) we plot the level (in dark feint with diamonds on left-hand scale) which can be seen to vary significantly due to natural stochastic variation. However, there is there is a notable fall at the beginning of year 5 when the uncertainty shock occurred. While this is not enormous - equivalent to about 6 months trend growth - it is still clearly visible. This fall in revenue productivity is also notable when looking at the plot of the change in revenue productivity (in light feint on right-hand scale) where the single biggest fall occurred in the first month of the uncertainty shock and the biggest single rise the month after the shock ended.

The reason for this impact of uncertainty on aggregate productivity is that uncertainty impedes the resizing of firms, in that high uncertainty reduces the shrinkage of low productivity firms and the expansion of high productivity firms.

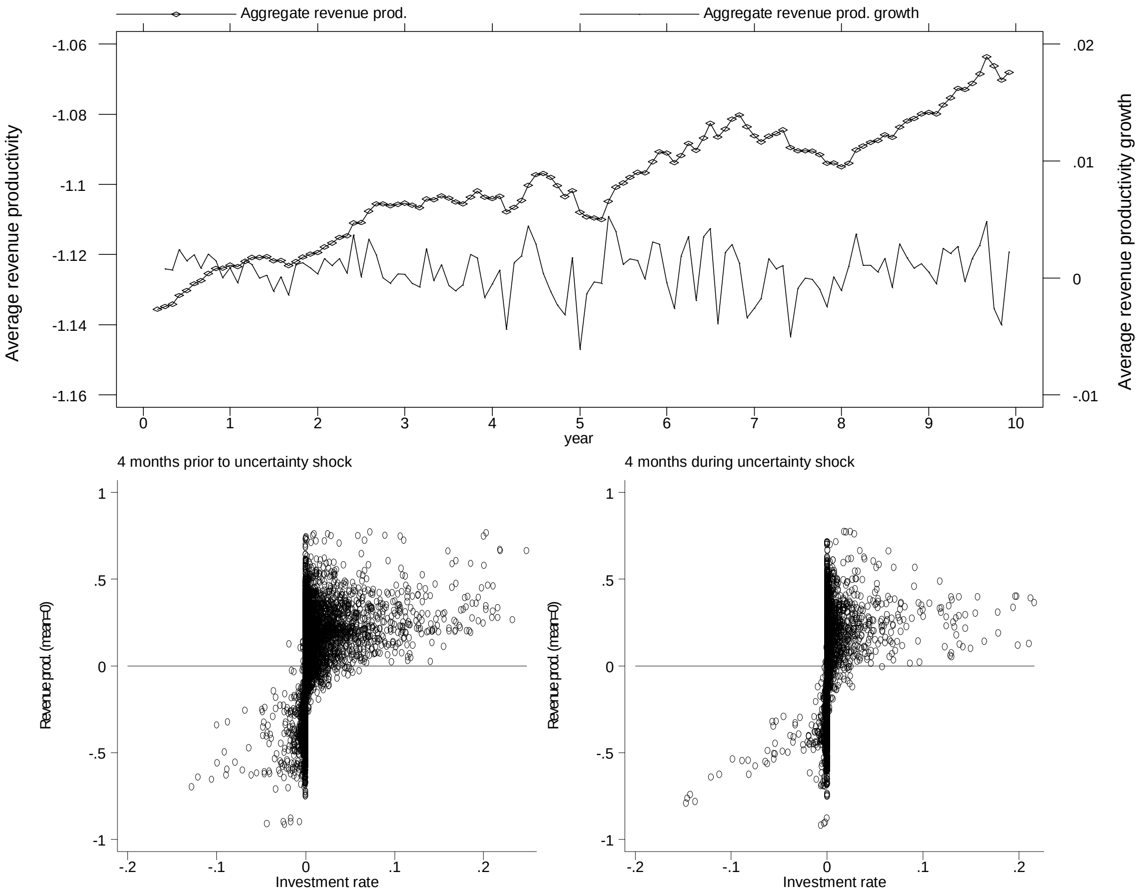

Figure 4.3: Revenue productivity levels and growth before, during and after the uncertainty shock

Hence, aggregate productivity levels and growth rates are reduced due to lower levels of "re-allocation" of capital and labour[26]. This re-allocative effect is large enough to have a moderate, albeit short-run, impact of aggregate productivity. This is illustrated in the bottom two panels of figure (4.3) where we plot the firm-level revenue productivity against rates of investment. On the left hand side we plot the distribution for the four months prior to the uncertainty shock (as a comparison period) while the right panel is the four months during the shock period. It can be seen that during the period of the uncertainty shock less adjustment happened - with less expansionary by productive firms and less contractionary investment by unproductive firms. Hence, by effectively reducing the rate of adjustment uncertainty will also play a role in driving productivity dynamics.

4.2. Bringing US Labour Adjustment Costs to European Levels

4.3. Reducing corporation tax by 1%

5. Conclusions

[26]While this is not a general equilibrium model so that there is strictly no concept of re-allocation, given the large number of firms the aggregate capital and labour stocks grow approximately constantly so that changes in one firm are on average off-set by changes in all other firms.

A. Appendix

Table A.1: Aggregation Obscuring Zero Investment Episodes.

% zero investment	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 (Small, Single Plants)	57.6	5.6	24.4	3.2

Note: Source UK ARD and Datastream

Table A.2: Aggregation Reduces Time Series Volatility.

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

Note: Source Compustat common support of firms with yearly and full quarterly data 1993-2001.

B. References - Incomplete

ABEL, A. and EBERLY, J. (1994), "A Unified Model of Investment under Uncertainty", American Economic Review, **84**, 1369-1384.

ABEL, A. and EBERLY, J. (1996), "Optimal Investment with Costly Reversibility", Review of Economic Studies, **63**, 581-593.

ABEL, A. and EBERLY, J. (2001), "Investment and q with Fixed Costs: An Empirical Analysis", Wharton School, University of Pennsylvania, mimeo.

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

BERTOLA, G. (1988), "Irreversible Investment", Massachusetts Institute of Technology doctoral thesis, re-printed in Research in Economics (1998), **52**, 3-37.

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

CABALLERO, R. and ENGEL, E. (1993), "Microeconomic Adjustment Hazards and Aggregate Dynamics", Quarterly Journal of Economics, 108, 313-358.

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.

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

COOPER, R., HALTIWANGER, J. and POWER, L. (1999), "Machine Replacement and the Business Cycle: Lumps and Bumps", American Economic Review, **86**, 921-946.

COOPER, R. and HALTIWANGER, J. (2000), "On the Nature of the Capital Ajustment Process", NBER working paper 7925.

COOPER, R., HALTIWANGER, J. and WILLIS, J. (2004), "Dynamics of Labor Demand: Evidence from Plant-Level Observations and Aggregate Implications", NBER working paper 10297.

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

DIXIT, A. (1989), "Entry and Exit Decisions under Uncertainty". Journal of Political Economy, **97**, 620-638.

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

DOMS, M. and DUNNE, T. (1993), "Capital Adjustment Patterns in Manufacturing Plants", Center for Economic Studies, U.S. Bureau of the Census.

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

GOURIEROUX, C, MONFORT, A. and RENAULT, E. (1993), "Indirect Inference", Journal of Applied Econometrics, 8, 85-118.

HAMERMESH, D. (1989), "Labour Demand and the Structure of Adjustment Costs", American Economic Review, 79, 674-89.

HAMERMESH, D. (1993), "Labour Demand", Princeton, NJ: Princeton University Press

HASSLER, J. (1996), "Variations in Risk and Fluctuations in Demand", Journal of Economic Dynamics and Control, **20**, 1115-1143.

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

NICKELL, S. (1985), The Investment Decision of Firms, Cambridge Economic Handbooks, Cambridge University Press, Cambridge.

PINDYCK, R. (1988), "Irreversible Investment, Capacity Choice, and the Valuation of the Firm", American Economic Review, **78**, 969-985.

RAMEY, V. and SHAPIRO, M, "Displace Capital", Journal of Political Economy, 605-622.

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

SMITH, A (1993), "Estimating nonlinear time-series models using simulated vector autoregressions", Journal of Applied Econometrics, S63-S84.