Investment, Cash Flow, and Value: A Neoclassical Benchmark

Janice Eberly[*], Sergio Rebelo[†], and Nicolas Vincent[‡]

September 2006

Abstract

Simple specifications of the neoclassical investment model are resoundingly rejected by the data. While this may indicate the presence of other frictions, this hypothesis cannot be tested without an empirically-supported specification of the neoclassical benchmark. In this paper, we model a firm facing decreasing returns to scale (or imperfect competition), productivity shocks, and convex capital adjustment costs. We estimate the model from firm-level Compustat data using the simulated method of moments and demonstrate that this simple framework can replicate many well-known empirical characteristics, including cash flow sensitivity of investment. Morever, when we allow for regime-switching in the productivity shocks, the model broadly matches the first three moments and persistence of firm-level investment, cash flow, and Tobin's Q, providing a neoclassical benchmark against which frictions might be discerned.

[*]Northwestern University and NBER.
[†]Northwestern University, NBER and CEPR.
[‡]Northwestern University.

1. Introduction

In Hayashi's (1982) neoclassical investment model, firms are perfectly competitive, have constant-returns-to-scale production functions, and face convex investment adjustment costs, linearly homogeneous in investment and capital. Investment is a function of the shadow value of capital (Tobin's marginal Q) which is empirically difficult to measure. But under Hayashi's conditions, marginal Q happens to coincide with average Q, the ratio of the value of the firm to its capital stock, which can be measured empirically. A firm's optimal investment-capital ratio can be expressed solely as a function of average Q. When adjustment costs are quadratic this function is linear.

The implication that the investment-capital ratio depends only on average Q has been resoundedly rejected by the data. These rejections are often interpreted as suggesting that frictions ignored by Hayashi (1982) are important determinants of investment. Fazzari, Hubbard, and Petersen (1988) (henceforth FHP) pioneered the use of linear investment equations to test for frictions, focusing on financial constraints. Using firm-level data, they regress investment on various proxies for fundamentals, including average Q and sales, and then add cash flow to the regression. They find large cash-flow effects for small, non-dividend paying firms that they interpret as evidence of financial constraints.

In this paper we propose an alternative interpretation for the findings of FHP and the related empirical literature. Our interpretation is that Hayashi's (1982) frictionless benchmark is misspecified because the strong assumptions necessary to make marginal and average Q coincide do not hold in the data. One of FHP's results suggests this interpretation: they find a significant cash flow effect on investment for large firms.

To study the empirical plausibility of our interpretation we propose an alternative frictionless model which features decreasing returns to scale in production and a flow fixed production cost. Productivity or demand shocks are stochastic and follow a Markov chain. We find that a regime-switching process for these shocks is crucial for matching the higher moments in the data.

We use the simulated method of moments to estimate this model using data for the largest (top quartile) firms in Compustat data set. These firms are used by FHP as the frictionless benchmark, since they are less likely to be affected by financial and informational frictions.

We use our parameter estimates to simulate our model and create a simulated panel of firm-level data. We then compare regressions from the simulated data to our regressions from the Compustat sample. Regressions on simulated data feature cash-flow effects similar to those in the Compustat sample.

Why do cash-flow effects emerge in our simulated data even though firms face no financial or informational frictions? The firm's value function is not linearly homogeneous and the shadow (marginal) value of capital is not proportional to the average value of capital (average Q). Average Q is no longer a sufficient statistic for the shadow value of capital which is the sole determinant of investment. The shadow value of capital determines investment, but this shadow value is a function of all the state variables in the model: the capital stock and the shock (and the regime, if there is regime-switching). A single independent variable no longer summarizes all the state variables, so *any* additional independent variable that is correlated with the state variables has explanatory power in a regression equation.

Related research hypothesizes that various forms of misspecification can generate an empirical correlation between investment and cash flow, and thus masquerade as financial frictions. The role of decreasing returns to scale in the revenue function has been seriously considered as early as Schiantarelli and Georgoutsos (1990), who explore the implications of monopolistic competition in a Q model of investment. They show, using aggregate data, that this feature can explain an accelerator-type effect of output on investment, even controlling for current and lags of Tobin's Q . Gomes (2001) models firms facing costly external finance and no adjustment costs, as well as firms with no external financing friction. He shows that it is possible to find no cash flow effect at all, even when there is a financing friction, and conversely, to find a cash flow effect in simulated regressions even when external finance is not costly in the model. Cooper and Ejarque (2000) show that a model with decreasing returns in the revenue function and quadratic adjustment costs, like ours, can generate empirically plausible cash flow effects by calibrating their model to match the estimated cash flow effect in a panel of data, also using simulated method of moments. Their later (2003) paper, adds a financing constraint to the model and shows that the data cannot distinguish between a firm facing a simple financing constraint and one facing curvature in the revenue function, again calibrating to match the cash flow effect. Abel and Eberly (2001, 2005) show that when firm value contains growth options, these options drive firm value, but do not continuously affect investment. Continuous investment is instead highly correlated with cash flow, with occasional investment spikes associated with exercise of growth options.

This work is complementary to ours, in showing that cash flow effects can be generated by deviations from the linear homogeneity assumption. Our paper takes a more general approach, by specifying a simple Cobb-Douglas model with quadratic adjustment costs and asking which parameters best fit observed firm-level data. It is important to our approach that we do not calibrate to fit a friction – neither theoretical nor empirical. Instead, our null is a neoclassical specification, and we estimate to fit only the univariate empirical moments of the firm's investment, cash flow, and value (Tobin's Q). In this way we generate a neoclassical benchmark - against which frictions can be discerned.

Results in the large empirical literature on financial frictions are suggestive of our findings. Gilchrist and Himmelberg (1995) construct a measure of Tobin's Q using a VAR methodology. When they include cash flow as an explanatory variable for forecasting Tobin's Q , the power of cash flow to predict investment is diminished (disappearing in some subsamples) for large firms. This is consistent with our misspecification argument. Similarly, Erickson and Whited (2001) test for cash flow effects in firm level data and find that when they go beyond a classical measurement error specification and instead allow for higher (third) order moments and heteroskedasticity, evidence of a cash flow effect disappears for both large and small firms.

2. The Model of the Firm

The model of the firm that we consider features decreasing returns to scale in production, quadratic adjustment costs in investment, a fixed cost of production, and stochastic total factor productivity.

The firm's problem is given by the following Bellman equation, where we use x' to denote next period's value of x:

$$\begin{aligned} V(k,z) &= \max_{i,k'}[zk^{\alpha} - \phi - \xi\left(i/k - \delta\right)^2 k - i \\ &\qquad + \beta \int V(k',z')F(dz',z)], \end{aligned}$$

$$k' = i + (1-\delta)k.$$

Here $V(k,z)$ represents the value of a firm that has capital k and total factor productivity, z. The behavior of z is governed by the distribution $F(.)$. We denote the discount factor by β. The firm's output is zk^{α}, where α is the degree of returns to scale. The variable ϕ represents a fixed production cost paid in every period. Capital depreciates at rate δ. Investment is subject to quadratic adjustment costs, which are represented by the term $\xi(i/k-\delta)^2$. This formulation has the property that adjustment costs are zero whenever the capital stock remains constant. The parameter ξ controls the magnitude of the adjustment costs.

We consider two versions of the model. In the 'single-regime model' z follows a Markov chain with support:

$$z \in \{\mu-\sigma, \mu, \mu+\sigma\}.$$

In the 'regime-switching model' the support of z is given by:

$$z \in \left\{\mu^L-\sigma^L, \mu^L, \mu^L+\sigma^L, \mu^H-\sigma^H, \mu^H, \mu^H+\sigma^H\right\}$$

where $\mu^L < \mu^H$. Productivity alternates between two regimes, the low regime $(\mu^L-\sigma^L, \mu^L, \mu^L+\sigma^L)$ and the high regime $(\mu^H-\sigma^H, \mu^H, \mu^H+\sigma^H)$. The evolution of z is governed by a Markov chain. The transition matrix associated with this Markov chain has the property that both the high and low regimes are highly persistent and there is a small probability of switching between regimes.

The model is solved by value-function iteration. We assume that k can only take n_k discrete values. We start with a guess for the value function, $V^0(k,z)$ for each pair (k,z). We compute the policy function $k^{'}=h^0(k,z)$ by finding the value of k' that maximizes the value of the firm for each pair (k,z). The new value function, $V^1(k,z)$ is given by the following equation with $m=1$:

$$\begin{aligned} V^m(k,z) &= \max_{i,k'}[zk^\alpha-\phi-\xi\left\{[k'-(1-\delta)k]/k-\delta\right\}^2 k-[k'-(1-\delta)k] \\ &\qquad +\beta\int V^{m-1}(k',z')F(dz',z)] \\ k' &= i+(1-\delta)k \end{aligned}$$

We use the $V^1(k,z)$ to find a new policy function $k^{'}=h^1(k,z)$ and a new value function, $V^2(k,z)$. We continue to iterate until $V^{m-1}(k,z)$ and $V^m(k,z)$ converge for every pair (k,z).

3. Estimation

3.1. Data

To estimate the model we use a balanced panel of Compustat firms with annual data for the period 1981-2003. Using a balanced panel introduces a selection bias towards more stable firms which are the focus of our study. Our sample includes 776 firms and roughly 14,000 firm-year observations. We focus our analysis on the top quartile of firms sorted by size of the capital stock in 1981 but we provide comparisons across quartiles. The model leads us to look at four main variables: investment in property, plant, and equipment, the physical capital stock, Tobin's Q , and cash flow. We exclude from our sample firms that have made a major acquisition, to help ensure that investment measures purchases of new property, plant, and equipment. We estimate the physical capital stock using the perpetual inventory method, using book value as the starting value and four-digit industry-specific estimates of the depreciation rate. Tobin's Q is calculated as the market value of equity plus the book value of debt, divided by the capital stock estimate. Cash flow is measured using the Compustat item for Income before extraordinary items + depreciation and amortization + minor adjustments. We describe the data in more detail in the appendix.

In Table 1 we report summary statistics for the fourth quartile (largest) firms in our sample, both for the 1981-2003 period and for two subperiods, 1981-1992 and 1993-2003. The median value of the variables that we consider are similar to those reported in other studies using Compustat data. The median values are 1.3 for Q , 0.15 for the investment rate, and 0.17 for the ratio cash flow/capital stock. We report the standard deviations for logarithms as well as the levels of the main variables so that their volatility can be compared. Q is the most volatile variable, closed followed by the investment, I/K. Cash flow is much less volatile than investment. All variables except the logarithm of the stock of capital exhibit positive skewness. There is more skewness in the full sample than in each of the two subsamples. This observation, along with the significantly higher values of the mean and standard deviation of Q and cash flow in the second subsample, leads us to consider a regime switching model in our estimation strategy. We return to this point when we discuss estimation and the regime-switching model. Finally, the data exhibit serial correlation, especially in the (log) capital stock and Q , but to a lesser extent in investment and cash flow.

We use some simple regressions to summarize some features of the data. In a pooled, time-series-cross-section regression of investment on Tobin's Q , we obtain:

$$\frac{I}{K} = \underset{(0.002)}{0.14} + \underset{(0.0004)}{0.01} Q, \quad \text{adjusted } R^2 = 0.16.$$

These regression coefficients are similar to those obtained in other empirical studies. The coefficient on Tobin's Q is quantitatively small, but significant, with modest explanatory power for investment. When we add cash flow as an explanatory variable we find:

$$\frac{I}{K} = \underset{(0.002)}{0.12} + \underset{(0.0005)}{0.003} Q + \underset{(0.009)}{0.17} \frac{\text{cash flow}}{K}, \quad \text{adjusted } R^2 = 0.24.$$

Including cash flow increases significantly the explanatory power of the regression and reduces the size and significance of the coefficient on Tobin's Q . Cash flow itself has a large and statistically significant effect on investment.

Abel and Eberly (2003) point out that the presence of skewness in Q suggests that a simple linear regression is not the best fit. They suggest a semi-log specification, which in these data yields:

$$\frac{I}{K} = \underset{(0.0016)}{0.14} + \underset{(0.0016)}{0.06} \log(Q), \quad \text{adjusted } R^2 = 0.20.$$

When cash flow is added to this semi-log specification we find

$$\frac{I}{K} = \underset{(0.005)}{0.22} + \underset{(0.002)}{0.03} \log(Q) + \underset{(0.002)}{0.039} \log\left(\frac{\text{cash flow}}{K}\right), \quad \text{adjusted } R^2 = 0.33.$$

These results confirm the presence of a significant cash-flow effect in our sample of large firms, in both the linear and the semi-log regression specification. Scatter plots of investment rates relative to Tobin's Q and cash flow help clarify why the semi-log specification performs better. Figure 1a shows a linear investment regression fit to the data on investment versus Q , and in panel b we fit instead to log(Q). Figures 2a and 2b show the same regressions for investment versus cash flow, with similar results. The semi-log specification has a better fit, increasing the R^2 by more than 10 percentage points in both cases, with significant effects of both Tobin's Q and cash flow on firm-level investment. In what follows, we use the semi-log specification as our benchmark when we examine some of the implications of our estimated model relative to the data.

3.2. Estimation Procedure

We estimate the model using the simulated method of moments proposed by Lee and Ingram (1989). We first use our data to estimate the vector of moments Ψ_D. Next, we choose the vector of parameters, Φ, that we want to estimate. For each candidate parameter vector, Φ, we simulate our model and compute simulated moments, which we denote by $\Psi(\Phi)$. Our parameter estimates are obtained by minimizing the weighted distance, L, between actual and simulated moments:

$$L(\hat{\Phi}) = \min\left[\Psi(\Phi) - \Psi_D\right]' W \left[\Psi(\Phi) - \Psi_D\right]. \tag{3.1}$$

The weighting matrix, W, is obtained using the variance-covariance matrix of the empirical moments, Ω_D:

$$W = \frac{1}{\Omega_D(1 + 1/k)}, \tag{3.2}$$

where k = length of simulation/length of sample. The intuition underlying the minimization problem (3.1) is that it penalizes heavily the difference between simulated and empirical moments only when the empirical moments are precisely estimated.

We estimate the matrix Ω_D using a block-bootstrap method. This method works as follows. We form m samples. Each sample consists of data for n firms sampled with replacement from our data set. For each of the m samples we compute the vector of empirical moments. We use the m observations on the vector of moments to estimate the variance-covariance matrix of the empirical moments, Ω_D.

We solve the minimization problem (3.1) using an annealing algorithm. The first step in this algorithm is to choose initial values for the parameter vector, Φ, as well as admissible ranges for each of the parameters in the vector. In the second step, we set the "temperature" and the step size. As we discuss below, the temperature controls the probability that, given the best parameter vector so far, Φ^*, we accept a parameter vector Φ' that yields a higher value of L, $L(\Phi') > L(\Phi^*)$. This procedure is used to avoid convergence to a local minimum. We start with a high temperature value, so that the algorithm explores different regions of the state space. The third step is to generate a new parameter vector, Φ', by adding random shocks to the elements of Φ^* within their admissible range. The fourth step is to solve the model using value-function iteration for the parameter vector Φ'. In the fifth step we simulate 1940 representative firms. This number is 10 times the length of our sample, so the value of k in (3.2) is equal to 10. Step six is to simulate a panel of firms and calculate simulated moments In step seven we compute $L(\Phi')$. If $L(\Phi') < L(\Phi^*)$ we set $\Phi^* = \Phi'$. If $L(\Phi') > L(\Phi^*)$ we set $\Phi^* = \Phi'$ with probability $\exp\left[-\left(L(\Phi') - L(\Phi^*)\right)/\text{temperature}\right]$. In step nine we reduce the values of temperature and step size. We go back to step three. We continue to iterate until we can no longer lower the value of L. The vector of parameter estimates is the one that generates the lowest value of L. We denote this vector by $\hat{\Phi}$.

The standard errors of the estimated parameters are computed as

$$\hat{\Omega} = \frac{(\Gamma' W \Gamma)^{-1}}{n},$$

where Γ is the matrix of derivatives,

$$\Gamma = \frac{\partial \Psi(\hat{\Phi})}{\partial \hat{\Phi}},$$

which we compute numerically.

4. Results: single regime model

4.1. Parameter and moment estimates

We report the parameters estimates and standard errors in Table 2. Our estimate of the adjustment cost parameter, ξ, is 0.541 (with a standard error of 0.144). This estimate implies that the average investment adjustment cost is 0.7 percent of sales. Our estimate for the fixed cost of operating, ϕ, is 0.010 (with a standard error of 0.038). This estimate implies annual fixed operating costs that are 17 percent of annual sales. The mean value of the shock z is 0.166, and the spread is plus or minus 0.086. As we discuss below, these values match the mean and standard deviation of cash flow in the data.

Table 2 reports summary statistics for the Compustat data, compared to those for the simulated firms in the model, using the estimated model parameters. The left-hand column reports the statistics from the data, using the full 1981-2003 sample. The right-hand column reports the moments from the model, where those reported in bold (red) were used in the estimation procedure to match the model to the data (the moments in the Ψ_D vector). The algorithm matches all of these moments closely. The other moments in the column, however, are not "targeted" by the algorithm. These results indicate that the simple model, with modest decreasing returns to scale, matches well the serial correlation in sales, cash flow, and investment, has slightly lower serial correlation in Q than in the data, but has a much lower standard deviation and skewness of Q than is found in the data. We add i.i.d. measurement error to our estimate of Q to match the standard deviation of Q in the data which also increases the skewness of Q (now higher than in the data); however, the i.i.d. measurement error reduces the serial correlation of Q even further below that in the data. In order to further evaluate the performance of the model relative to the data, we regress investment on its determinants - both the state variables that are only observable in the model, as well Q and cash flow, as we did earlier in section 3.1 with the data.

4.2. Simulated regression results

Table 3 reports the results of estimating linear regressions on data from firms simulated from the estimated model. The first column demonstrates that using a semi-log specification on the true state variables of the model (k and the shock, z) yields an R^2 of 0.95, providing a very good approximation to the model. When we use the true measure of Tobin's Q , the model fits nearly as well ($R^2 = 0.93$), so Q provides a very good proxy for the two state variables. However, when we use the noisy measure of Q that better fits the properties of Q , the R^2 falls to 0.09 and the coefficient on Q is 0.039 (compared to 0.38 for the true Q). When cash flow is added to the regression with noisy Q , the coefficient on Q falls below 0.01, cash flow has a coefficient of 0.08, and the R^2 jumps up to 0.68. The final column substitutes the value of the shock, z, for cash flow in this regression; this substitution gives a coefficient estimate and R^2 that are nearly the identical to using cash flow as a dependent variable. Since there are no frictions in the model that would give a direct role to cash flow, these results strongly suggest that cash flow proxies for the impact of the shock.

While these results are instructive, it is clear from Table 3 that the model does not match the data along several dimensions, in particular the volatility of Q and the skewness of Q , cash flow, and investment. The former we address by adding i.i.d. measurement error. We also experimented with adding a behavioral bias to the model. Specifically, we assumed that managers forecast fundamentals using the correct Markov chain, but we allowed investors to use a distorted Markov chain. A Markov chain with higher persistence (larger values on the diagonal) generated enough volatility in Q , but did not address the skewness of Q found in the data.

We have already shown that a simple neoclassical model can generate a cash flow effect, however, our intent is to provide a neoclassical model that also fits the broad characteristics of the firm-level data. In order to address the skewness in Q (and in the other variables) we add a regime-switching component to the Markov chain.

5. Results: regime switching model

5.1. Parameter and moment estimates

The regime switching model allows for a second regime in the productivity shock z. In this case, we estimate the same structural parameters in the single regime model, as well an additional three parameters: the average value and range of the shock in the second regime, plus the regime-switching parameter. The estimated model parameters and standard errors are reported in Table 5. The estimated adjustment cost parameter ξ is 0.718 (with a standard error of 0.019), which implies that the average investment adjustment cost is 1.0% of sales. The estimated fixed cost of operating, ϕ is 0.022 (with a standard error of 0.009), which implies that annual fixed operating costs are 14% of annual sales. These values are very similar to those in the single regime model, and the standard errors are lower. Figure 3 plots the shocks in the two regimes. It is interesting to note that the data gravitate toward two overlapping regimes, where the high regime has a higher mean productivity, but also a higher standard deviation. In fact, the low shock in the high regime is lower than the low shock in the low regime. All of these parameters have low standard deviations and are precisely estimated. The estimated Markov chain (Table 6) exhibits substantial persistence. We calibrated the within-regime autocorrelation to 0.6 to match the serial correlation in the data. The regime switching parameter is precisely estimated at 0.057, which gives the total probability of switching from one regime to the other, but switches can occur from either the middle state or that state closest to the alternative regime (e.g., transiting from the highest low state to the high regime, or from the lowest high state to the low regime).

Table 7 reports summary statistics for the Compustat data, compared to those for the simulated firms in the model, using the estimated model parameters with regime switching. The left-hand column reports the statistics from the data, using the two subsamples, as well as the full 1981-2003 sample. The right-hand column reports the moments from the model for the low regime, the high regime, and the combined data, where those highlighted were used in the estimation procedure to match the model to the data (the moments in the Ψ_D vector). The algorithm matches all of these moments closely, as the algorithm is designed to do. The other moments in the column, however, are not "targeted" by the algorithm. These results indicate that the addition of regime switching improves the fit of the model, particularly for the higher moments in the data. The standard deviation and skewness of Q are substantially higher, and the model generates skewness in the investment and cash flow that are much more like the data (even a bit too much skewness). At the same time, the serial correlation properties are enhanced compared to the single regime model. Before running data regressions, we again match the standard deviation of Q by adding measurement error, but in the regime switching model this measurement error adjustment is much smaller, since the fundamental volatility of Q is higher in this case. We add 20 percent measurement error to Q , with serial correlation to preserve the serial correlation in Q , which already matches the data very well.[1]

In order to better understand the dynamics of the model, we calculated the elasticity of each moment in the Ψ_D vector with respect to the parameters of the model. This gives an indication of how changes in the parameter values affect the performance of the model. The matrix of elasticities is reported in Table 8. We highlight those which have the largest impact. In the first row of the table, for example, we see that average Q in the first (low) regime is largely determined by the fixed operating cost, which lowers average Q , and then the average shock in the low regime μ^L, as well as the average shock in the high regime μ^H and the probability of transiting to the high regime. Thus average Q is largely determined by expected cash flow, net of operating costs, as one would expect. Average Q in the high regime depends most on the average shock in the high regime. Importantly for our later results, average cash flow is largely determined by the average shock. Similarly, the standard deviation of cash flow in each regime has a unit elasticity with respect to the standard deviation of shocks in the regime. Finally, the standard deviation of investment (investment rate, i/k) is determined largely by the adjustment cost parameter ξ (negatively) and the mean shock in the low regime (also negatively), since this determines the volatility of investment

[1] To be precise, we generate $Q_{noise} = Q\exp(\varepsilon_t)$, where $\varepsilon_{t+1} = 0.73\varepsilon_t + \eta_{t+1}$ and $\eta_t \backsim N(0, 0.145)$ and the standard deviation of ε is 0.20.

across regimes.[2] The next two figures, Figure 4a and 4b, plot the value functions and policy functions for each state in the two regimes as a function of the firm's capital stock. Note that even though the shock regimes are overlapping, the value and policy functions are not. This owes to the estimated transition structure of the model. For example, even though the lowest individual shock is in the high regime, there is a higher probability of transiting from this state to the highest possible states, so the value of the firm is higher in this state than for a higher realization in the lower regime (where the firm's future prospects are bleaker).

5.2. Simulated regression results

In order to further evaluate the performance of the model relative to the data, we regress investment on its determinants - both the state variables that are only observable in the model, as well Q and cash flow, as we did earlier with the single regime model and in section 3.1 with the Compustat data. Table 9 reports these results for the two regime model. In the first column, we use the state variables from the model to explain investment in a semi-log specification. As we found previously, the semi-log specification does a good job of recovering the model, with an R^2 of 0.97. When we revert however, to standard observable variables, we find that a regression of investment on Tobin's Q has an R^2 of only 0.62, even when we control for the regime (which is an additional state variable in our model not typically included in Q models). If we ignore the regime switching and just regress investment on Tobin's Q , the coefficient on Q is only 0.091 and the R^2 falls to 0.38; if we use a noisy measure of Q that matches the empirical properties of Q , the coefficient on Q falls further to 0.068 and the R^2 falls to

[2] Notice that the elasticity of the investment standard deviation with respect to μ^H is 1.7, while the elasticity with respect to μ^L is -2.0. Since the two coefficients are of opposite sign and similar magnitude, this suggests that the difference between the two is what matters - or the average spread between the two regimes.

0.28. The importance of adding measurement error in Q falls significantly in the regime switching model; the regime-switch already disrupts the relationship between investment and the shocks, as we describe in more detail below. If we follow the empirical literature and add cash flow to this regression, the coefficient on Q falls to 0.05, cash flow enter significantly with a coefficient of 0.04 and the R^2 recovers to 0.39. As we demonstrated with the single regime model, this is nearly identical to the regression results if we include the shock realization instead of cash flow, as shown in the last column of the table, suggesting again that cash flow proxies for the shock.

The next table, Table 10, exploits the regime switching implications to further explore the fit of the model. This table reports regressions run on the simulated data in the first two columns, and then new regressions on the Compustat data in the two right-hand columns. Since we emphasize the importance of the two regimes in the model, we now use our two subsamples of data as empirical proxies for the two regimes. In the first column, we regress investment on a dummy for the high regime, as well as Tobin's Q. Q has a significant positive effect on investment, but the high regime dummy variable enters negatively, since Q is disproportionately high (compared to investment) in the high regime. These results are nearly identical to the results for the data reported in the third column. The coefficient on Q is 0.07 (versus 0.086 in the model), the high regime dummy is significantly negative, and the R^2 is 0.33. When we add cash flow to this regression in the model, it enters positively with a coefficient of 0.039, reduces the coefficient on Q and slightly reduces (in absolute value) the coefficient on the high regime dummy; while the R^2 increases modestly. In the data, the results are very similar. The coefficient on cash flow is 0.034, the coefficient on Q falls by about a third, the high regime dummy is smaller in absolute value, and the R^2 of the regression increases.

Thus, the regime switching model not only improves the fit of the model to the moments of the firm, it also matches the covariation and partial covariation among investment, cash flow, and valuation (Tobin's Q) observed in the data. Figure 5 demonstrates the role of regime switching in understanding the investment regressions. In the data and in the simulation, both the true Q and the "noisy" Q have relatively poor explanatory power for investment when there is regime switching (columns 3 and 4 of Table 9). Cash flow improves the fit of the regression, but not nearly as dramatically as it did in the single regime model, where using cash flow to proxy the shock brought the regression R^2 from 0.09 to 0.68 – whereas with regime switching, the addition of cash flow only increases the R^2 from 0.28 to 0.39.[3] Figure 5 plots the investment rate, I/K, as a function of the capital stock for each value of the shock, z. There is no longer a monotonic relationship between the current shock and current investment. The lowest investment rates occur on the lowest branch of the graph, when $z = 0.1188$ in the low regime. However, when z takes on its *lowest* value, $z = 0.0987$, in the high regime, investment is substantially *higher*. This occurs because of the transition probability within the regimes. Within the high regime, even when current z is very low, future prospects are bright because of the higher probability of transiting to the most favorable states. In the low regime, current z can be higher, but the prospects for the future are relatively bleak and thus investment remains low. The transition dynamics within and across regimes disrupt the monotonic relationship between investment and the shock, and hence between investment and cash flow. As Table 10 shows, however, a simple linear control for the regime does not restore the explanatory power of the regression, since the regimes do not differ by a constant shift term.

[3] When the dummy for the regime is included, the addition of cash flow to the Q regression increases the R^2 from 0.33 to 0.37 (see Table 10).

6. Extensions

In addition to the findings reported above, we explored a number of other model features and specifications in order to understand their implications. In particular, we looked at different specifications of the adjustment cost function, simple forms of borrowing constraints, and a non-constant interest rate.

The skewness in investment led us to consider asymmetric adjustment costs, both in the form of asymmetric quadratic adjustment costs and an irreversibility constraint. We examined an asymmetric quadratic adjustment cost, with a different parameters ξ for high and low investment. This generated skewness in investment, but not in cash flow. Once we allowed for skewness in cash flow through the shock process, the model naturally generated investment skewness so additional asymmetry from adjustment costs was not necessary quantitatively. Similarly, for these large firms, an irreversibility constraint was not quantitatively important. It rarely binds and did not have a noticeable impact on the results. Other authors, such as Doms and Dunne (1998), using smaller firms and disaggregated plant data, observe that aggregating to large firms tends to smooth out these non-convexities, so it is not surprising that this feature is not evident quantitatively in our sample in the moments we examine.

We also added a borrowing constraint to our model to examine whether it could identified if it was present. We used a simple specification that does not require adding an additional state variable to the model: we limiting borrowing to a fixed percentage of the capital stock. This constraint binds frequently, especially at the transition from the low regime to the high regime when the firm would like to grow quickly. However, even when only 10% of the capital stock can be used as collateral, there is no discernible effect on the moments of the model. This does not mean that there are no borrowing constraints; rather, it means that if they are present, they are difficult to detect in the data with the methods we have used here (which mirror much of the literature in the regression format).

Finally, we added interest rate variation by including the actual path of real interest rates to the model as we solved the optimization problem of the firm and the estimation algorithm. This generalization had no effect on the results, so we used a constant discount rate throughout our reported results.

7. Conclusions and future work

This research specifies and estimates a benchmark neoclassical model to fit the empirical moments of data on large firms. The important features of the model are decreasing returns to scale (or imperfect competition), quadratic adjustment costs on investment, a fixed operating cost, and regime switching productivity shocks. The estimated parameterization broadly matches the average, standard deviation, skewness, and persistence of investment, cash flow, and Tobin's Q . Moreover, simulated data from the model also closely matches the covariation and partial covariation of investment with Tobin's Q and cash flow as measured by regression analysis, even though the model was calibrated as a frictionless neoclassical model. This specification was intentional, since other frictions can be present in the economy and in the data, but they can only be identified relative to a well-specified neoclassical benchmark. A friction that is correlated with the misspecification of the model cannot be separately identified. The model estimated here is intended to provide this benchmark.

As we develop the model further, we intend to include a more sophisticated modelling of capital structure. This requires solving a more challenging numerical problem, but allows us to address quantitative questions in corporate and macro finance.

References

[1] Abel, Andrew, and Janice Eberly, "Q Theory Without Adjustment Costs & Cash Flow Effects Without Financing Constraints", working paper, 2001.

[2] Abel, Andrew, and Janice Eberly, "Investment and q with Fixed Costs: An Empirical Analysis", working paper, 2003.

[3] Abel, Andrew, and Janice Eberly, "Investment, Valuation, and Growth Options", working paper, 2005.

[4] Cooper, R. and Joao Ejarque, "Exhuming Q : Market Power vs. Capital Market Imperfections," working paper, June 2000.

[5] Cooper, R. and Joao Ejarque, "Financial Frictions and Investment: A Requiem in Q ", *Review of Economic Dynamics*, 6 (2003), 710-28.

[6] Doms, Mark and Timothy Dunne, "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics*, vol. 1(2), pp 409-429, April.

[7] Erickson, T. and Toni Whited, "Measurement Error and the Relationship between Investment and 'q'" *The Journal of Political Economy*, Vol. 108, No. 5. (Oct., 2000), pp. 1027-1057.

[8] Gilchrist, Simon and Charles Himmelberg, "Evidence on the role of cash flow for investment," *Journal of Monetary Economics* 36 (1995), pp. 541 572.

[9] Gomes, Joao, "Financing Investment," *The American Economic Review*, Vol. 91, No. 5 (Dec., 2001), pp. 1263-1285

[10] Hayashi, Fumio, "Tobin's marginal Q and average Q : A neoclassical interpretation," *Econometrica*, 50, pp 215-224.

[11] Lee, B. and Ingram B. (1991), “Simulation estimation of time series models”, *Journal of Econometrics*, 47, 197-205.

[12] Schiantarelli, Fabio, and D. Georgoutsos, "Monopolistic competition and the Q theory of investment, *European Economic Review*, Volume 34, Issue 5 , July 1990, pp 1061-1078.

Table 1.

Median Across Large Firms (4th quartile)	Full Sample 1981-2003	Sub-Samples 1981-1992	1993-2003
Time-Series Average			
Q	1.30	0.95	1.89
I/K	0.15	0.15	0.16
Cash Flow/K	0.17	0.15	0.20
Sales/K	1.54	1.38	1.97
log(K)	7.73	7.59	7.97
Time-Series Standard Deviations			
Q	0.66	0.26	0.59
log(Q)	0.43	0.28	0.28
I/K	0.06	0.05	0.05
log(I/K)	0.38	0.34	0.30
Cash Flow/K	0.08	0.05	0.09
log(Cash-Flow/K)	0.26	0.15	0.22
log(K)	0.23	0.15	0.12
Skewness			
Q	0.64	0.19	0.43
I/K	0.46	0.38	0.39
Cash Flow/K	0.26	-0.04	0.08
Sales/K	0.66	0.10	0.42
log(K)	-0.08	0.08	0.03
Serial Correlation			
Q	0.70	0.56	0.42
I/K	0.45	0.33	0.36
Cash Flow/K	0.43	0.33	0.21
Sales/K	0.71	0.56	0.52
log(K)	0.83	0.68	0.73

For each variable, we compute the time series average for each firm in the sample, and report the median across firms. “Q” is Tobin’s Q, I is investment in property, plant, and equipment, and K is the capital stock. Construction of the variables is described in the text and in the data appendix.

Table 2: parameter estimates for the single regime model

Parameter	**Estimate (standard error)**
Adjustment cost : ξ	0.541 (0.144)
Fixed cost: ϕ	0.010 (0.038)
Mean shock: μ	0.166 (0.132)
Shock range: σ	0.086 (0.074)
Calibrated parameters	
Discount factor: β	0.977
Returns to scale: α	0.80
Depreciation rate: δ	0.15
Persistence: ρ	0.60

Table 3. Summary statistics for Compustat data versus the single regime model

Median Across Large Firms	Data 1981-2003	Single Regime Model
Time-Series Average		
Q	1.2980	1.3160
I/K	0.1497	0.1514
Cash Flow/K	0.1689	0.1737
Time-Series Standard Deviations		
Q	0.6253	0.1793
Qnoise		0.6194
I/K	0.0548	0.0541
log(I/K)		
Cash Flow/K	0.0781	0.0791
log(Y/K)		
log(K)		
Skewness		
Q	0.64	0.1098
Qnoise	0.64	1.3915
I/K	0.46	0.1041
Cash Flow/K	0.26	-0.0706
Serial Correlation		
Q	0.84	0.5151
Qnoise	0.84	0.0328
I/K	0.53	0.4847
Cash Flow/K	0.60	0.61
Sales/K	0.82	0.7616
log(K)	0.95	

Table 4. Regression results from the single regime model

Dependent variable I/K

Regression	1	2	3	4	5
Constant	0.2101	0.0509	0.1411	0.2942	0.3514
	0.0011	0.0004	0.0008	0.0015	0.0019
ln(q)		0.3789			
		0.0014			
ln(qnoise)			0.0389	0.0089	0.0395
			0.0016	0.001	0.0016
ln(cash flow/k)				0.0769	
				0.0007	
ln(k)	-0.1499				
	0.0009				
ln(z)	0.136				0.1075
	0.0004				0.001
R2	0.952	0.9257	0.0897	0.6817	0.7099

Table 5: parameter estimates for the regime-switching model

Parameter	Estimate (standard error)
Adjustment cost : ξ	0.718 (0.019)
Fixed cost: ϕ	0.022 (0.0086)
Low regime mean shock: μ^L	0.168 (0.014)
Low regime shock range: σ^L	0.0496 (0.0037)
High regime mean shock: μ^H	0.2215 (0.017)
High regime shock range: σ^H	0.1228 (0.0086)
Switching parameter	0.057 (0.0009)
Calibrated parameters	
Discount factor: β	0.977
Returns to scale: α	0.80
Depreciation rate: δ	0.15
Persistence: ρ	0.60

Table 6: estimated Markov chain for the regime-switching model

Low Regime			High Regime		
0.1188	0.1684	0.2180	0.0987	0.2215	0.3443
0.64	0.32	0.04	0	0	0
0.157	0.6671	0.157	0.019	0	0
0.0385	0.3078	0.6157	0.038	0	0
0	0	0.038	0.6157	0.3078	0.0385
0	0	0.019	0.157	0.6671	0.157
0	0	0	0.04	0.32	0.64

Table 7: Summary statistics for Compustat data versus the regime-switching model

	Data					
Median Across Large Firms	1981-1992	1993-2003	1981-2003	Low	High	All
Time-Series Average						
Q	0.9515	1.8919	1.2980	0.9324	1.8083	1.3924
I/K	0.1457	0.1611	0.1497	0.1324	0.1686	0.1515
Cash Flow/K	0.1546	0.1994	0.1689	0.1433	0.1902	0.1679
Sales/K						
fcost/y						
Time Series Standard Deviations						
Q	0.2560	0.5891	0.6253	0.1507	0.3112	0.503
Qnoise				0.2592	0.5036	0.5987
I/K	0.0497	0.0457	0.0548	0.0398	0.0604	0.0547
log(I/K)						
Cash Flow/K	0.0458	0.0891	0.0781	0.0453	0.0885	0.0751
log(Y/K)						
log(K)						
Skewness						
Q	0.19	0.43	0.64			0.4354
Qnoise			0.64			0.8352
I/K	0.38	0.39	0.46			0.5452
Cash Flow/K	-0.04	0.08	0.26			0.4303
Sales/K						
log(K)						
Serial Correlation						
Q			0.84			0.8823
Qnoise			0.84			0.8298
I/K			0.53			0.685
Cash Flow/K			0.60			0.6329
Sales/K			0.82			0.555
log(K)			0.95			

Table 8: Elasticity of moments with respect to the parameters for the regime-switching model

	ξ	ϕ	μ^L	σ^L	μ^H	σ^H	γ
Average q_L	-0.1	-3.0	6.9	0.1	8.1	0.0	0.3
Average q_H	0.0	-0.6	0.4	0.0	2.6	0.1	-0.1
Average cash flow_L	0.0	-0.5	2.2	0.0	0.0	0.0	-0.1
Average cash flow_H	0.0	-0.1	0.0	0.0	0.7	-0.1	0.0
Std dev cash flow_L	0.0	0.1	-1.2	1.0	-0.3	0.0	0.1
Std dev cash flow_H	0.0	0.0	-0.2	0.0	-0.9	1.0	0.1
Std dev cash flow	0.0	0.1	-1.0	0.2	-0.5	0.7	-0.1
Std dev i/k	-0.4	0.0	-2.0	0.1	1.7	0.6	-0.2

Table 9. Regression results from the regime-switching model

Dependent variable I/K

Regression	1	2	3	4	5	6
Constant	0.1891	0.1509	0.1272	0.1330	0.2243	0.2427
	0.0003	0.0003	0.0003	0.0003	0.0100	0.0011
Dummy High Regime	0.1505	-0.1168				
	0.0002	0.0007				
ln(q)		0.2324	0.0908			
		0.0010	0.0006			
ln(qnoise)				0.0683	0.0495	0.0462
				0.0005	0.0005	0.0005
ln(cash flow/k)					0.0403	
					0.0005	
ln(k)	-0.1340					
	0.0002					
ln(z)	0.1104					0.0608
	0.0001					0.0006
R2	0.9672	0.6197	0.3836	0.2810	0.3942	0.4341

Table 10. Regression results from the regime-switching specification in the Compustat data versus the model

Dependent variable I/K

Regression	Model		Data	
Constant	0.1384	0.2146	0.1570	0.2238
	0.0003	0.0010	0.0020	0.0050
Dummy High Regime	-0.0193	-0.0096	-0.045	-0.039
	0.0007	0.0007	0.003	0.003
ln(q)				
ln(qnoise)	0.0860	0.0589	0.0700	0.0450
	0.0008	0.0008	0.0020	0.0020
ln(cash flow/k)		0.0391		0.034
		0.0005		0.002
ln(k)				
ln(z)				
R2	0.2941	0.3973	0.3320	0.3677

Figure 1a: Investment rate (I/K) versus Tobin's Q

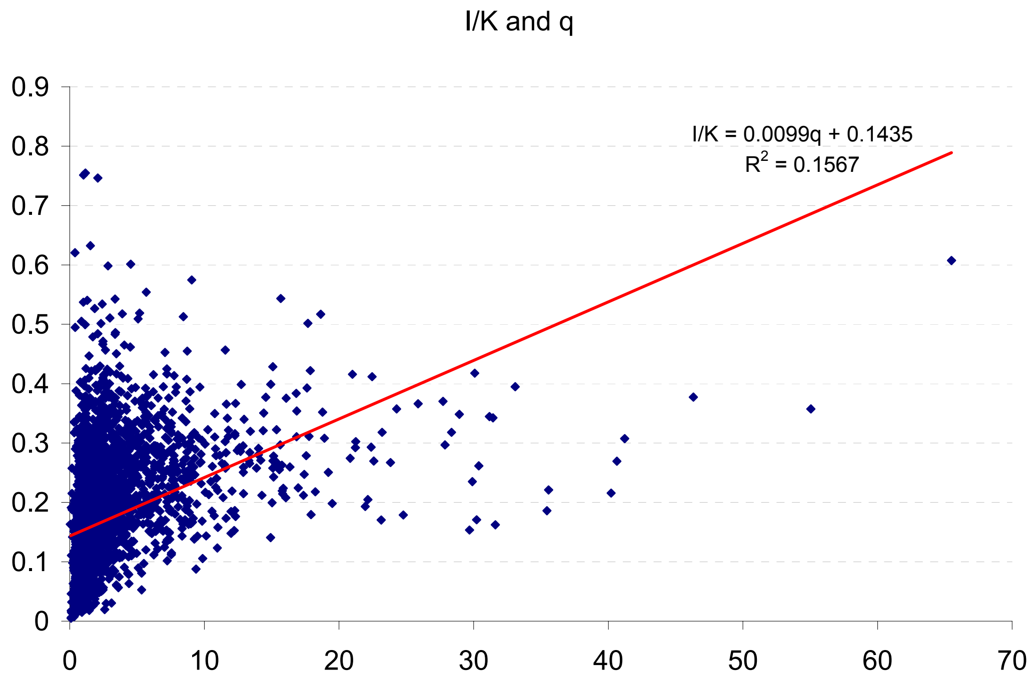

Figure 1b.

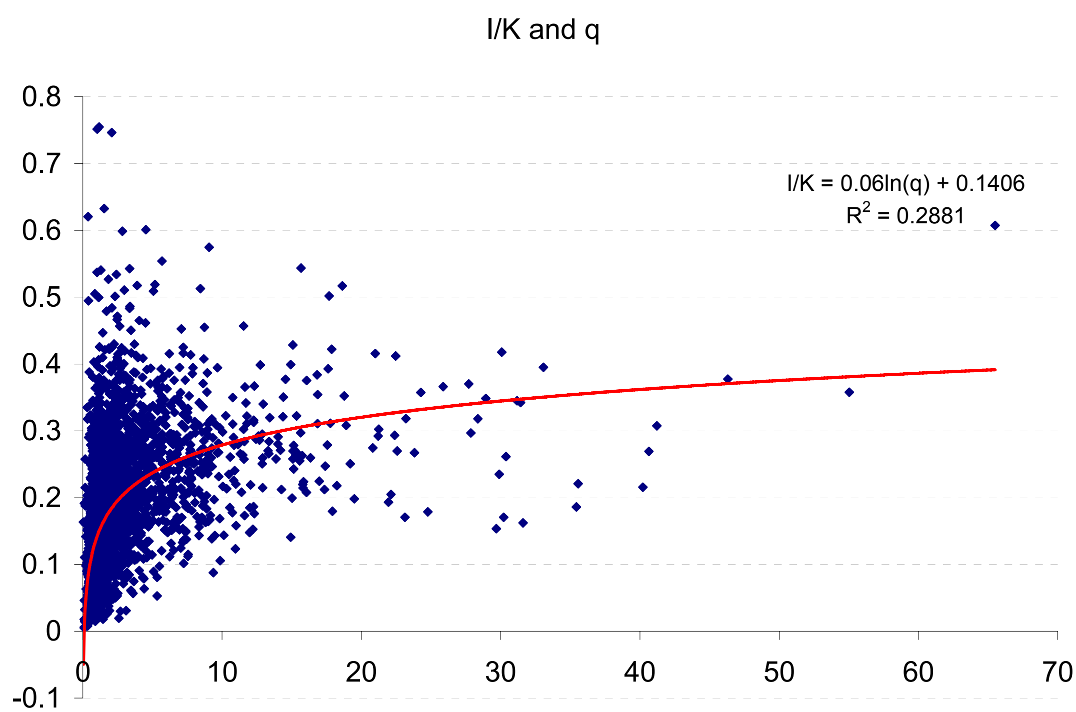

Figure 2a: Investment rate (I/K) versus cash flow (CF/K)

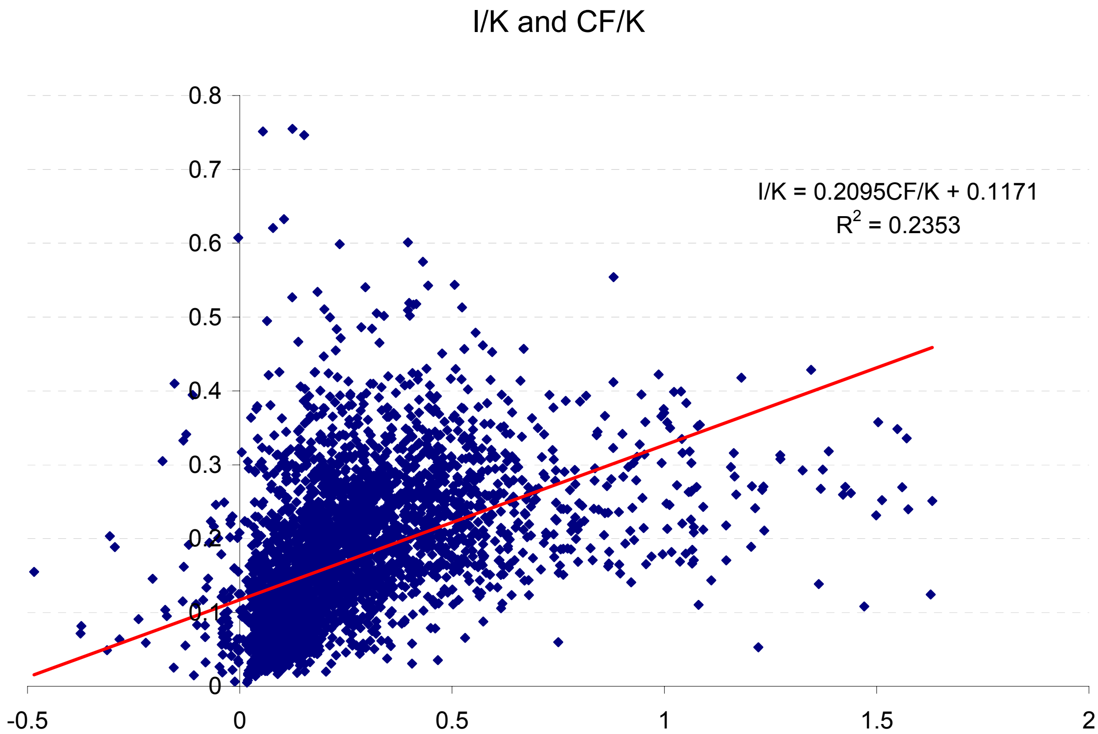

Figure 2b.

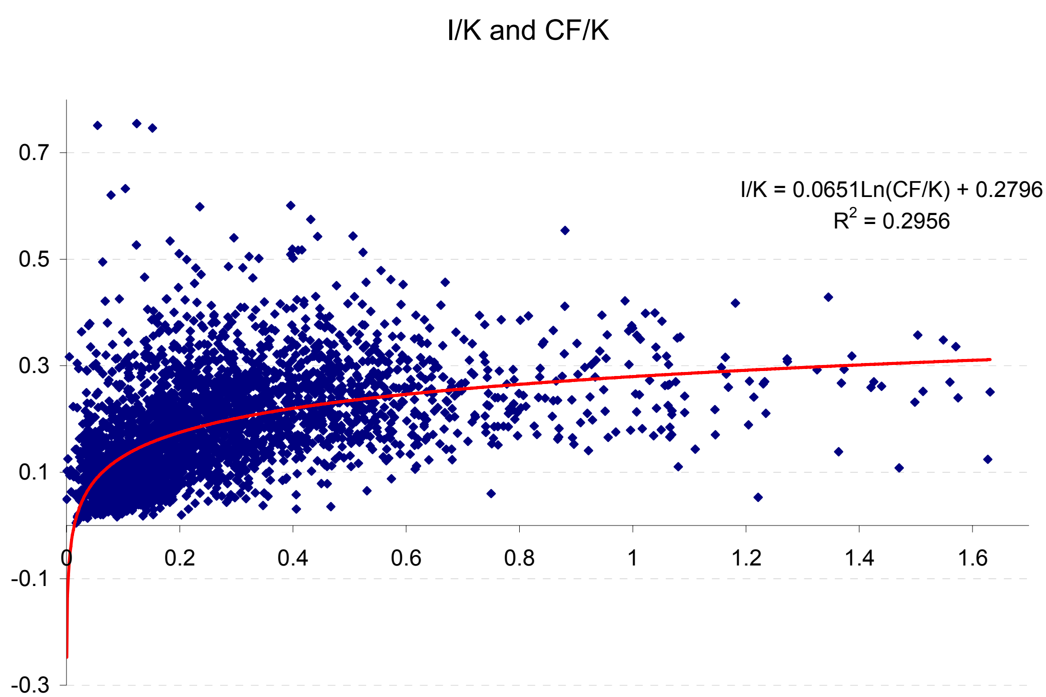

Figure 3a: Regime-switching model, value function by state in each regime

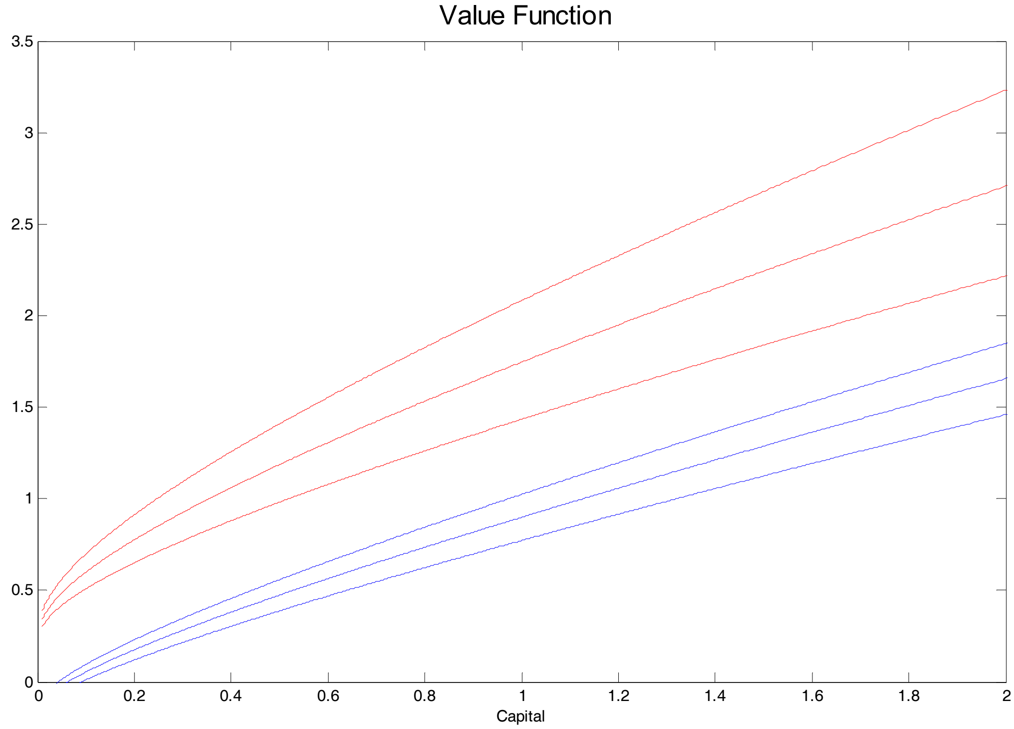

Figure 3b: Regime-switching model, policy function by state in each regime

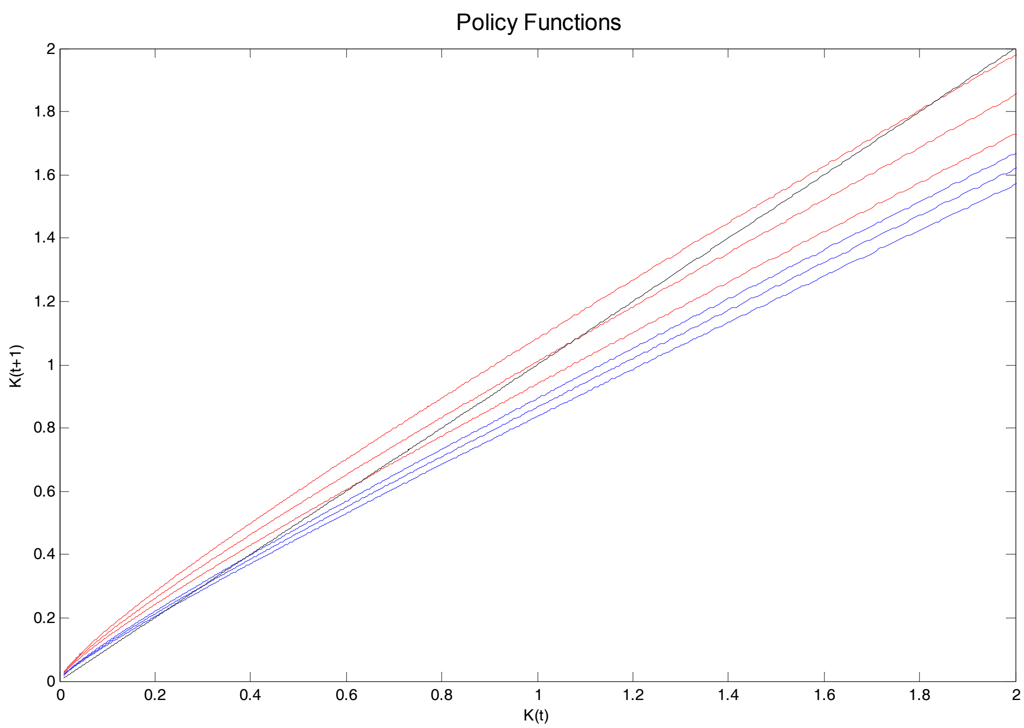

Figure 4: Regime-switching model, investment (I/K) by state in each regime