

Investment and Value: A Neoclassical Benchmark*

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

November 2008

Abstract

Which investment model best fits firm-level data? To answer this question we estimate alternative models using Compustat data. We find that both a version of the Hayashi (1982) and a model with decreasing returns to scale in production fit firm-level data very well. Our estimates suggest that there is substantial measurement error in Q . This measurement error implies that the investment regressions that have received so much empirical attention are ineffectual to discriminate among alternative models. In fact, the models that we estimate generate plausible cash-flow and lagged-investment effects even though they were not designed to produce them.

*We thank Nick Bloom for helpful discussions on the estimation algorithm, Damba Lkhagvasuren for providing us with one of the algorithms used in the paper, and Caroline Sasseville and Niels Schuehle for research assistance. We also thank Bob King, Martin Eichenbaum, and seminar participants at Berkeley, Boston University, Chicago, Duke, Harvard, Indiana, Northwestern, Stanford, and Wharton for their comments.

[†]Northwestern University and NBER.

[‡]Northwestern University, NBER and CEPR.

[§]HEC Montréal.

1. Introduction

In Hayashi's (1982) neoclassical model of investment average Q (the value of existing capital divided by its replacement cost) is equal to marginal Q (the value of an additional unit of installed capital). This property famously implies that average Q is a sufficient statistic for determining a firm's investment decision. This implication has often been empirically rejected. A cash-flow effect is present in virtually every investment-regression specification and data sample. This findings has spawned a large literature that measures and interprets this effect. The lagged-investment effect has attracted less attention than the cash-flow effect but it is empirically much more important. Lagged investment is a much better predictor of current investment than either Q or cash flow. Both the cash-flow and the lagged-investment effect suggest that Hayashi's model is an inadequate description of the behavior of investment at the firm level.

In this paper we search for an empirically successful model of investment. We find that investment regression estimates are very sensitive to the presence of the measurement error in Q . So, instead of using investment regressions as our guide, we use a broad set of empirical moments (means, standard deviations, persistence, and skewness properties of cash flow, Q , and investment) to estimate three candidate models. Our estimates are based on the simulated method of moments implemented on firm-level data for the top quartile of Compustat firms sorted by the size of the capital stock in the beginning of the sample. These are the firms that Fazzari, Hubbard, and Petersen (1988) (henceforth FHP) use as their frictionless benchmark because they are unlikely to be affected by financial frictions.

We consider three models driven by stochastic shocks that can be interpreted as productivity or demand shocks. We assume that these shocks follow a regime-

switching process. This assumption is important, as it generates skewness in cash flows, as well as the low correlation between Q and current cash flow observed in the data.

The first model, which we call the “decreasing-returns model,” features decreasing returns to scale in production, a fixed operating cost, and quadratic capital adjustment costs. The conditions for Q to be a sufficient statistic for investment choice are not satisfied in this model because the production function is not homogeneous of degree one. However, in a single-regime version of the model, the decision rule for optimal investment can still be very closely approximated by a log-linear function of Q . The second model, which we call the “Hayashi model,” is a version of Hayashi’s (1982) model with quadratic investment adjustment costs. The third model, which we call the “CEE model” incorporates adjustment costs that penalize changes in the level of investment, as proposed by Christiano, Eichenbaum, and Evans (2005). This specification has gained currency in the macroeconomics literature because it generates impulse responses to monetary policy shocks that are consistent with those estimated using vector auto-regressions.

Surprisingly, we find that both the Hayashi model and the decreasing-returns model fit firm-level data very well. The CEE model also provides a reasonably good fit, but it generates excess persistence and insufficient skewness in investment. These properties result from the fact that the CEE model penalizes large changes in investment, generating a highly persistent investment series that exhibits very few investment spikes.

Our estimates suggest that there is substantial measurement error in Q . These findings are consistent with the results in Erikson and Whited (2000) who estimate the importance of measurement error in Q using the information contained in the third and higher-order moments of the joint distribution of the regression variables.

This measurement error can arise from any component of Q that is better observed by the firm than by the researcher, including the market value of debt or the replacement value of the capital stock.¹ More controversially, measurement error in Q can also arise from differences between the intrinsic value and the market value of equity, as suggested by Shiller (1989, 2000). Consistent with this view, measures of Q that do not rely on the market value of equity tend to be better predictors of investment than conventional measures of Q . Examples of these alternative Q measures, include estimates based on cash-flow forecasts (Gilchrist and Himmelberg (1995)), analyst forecasts of earnings growth (Cumins, Hassett, and Oliner (2006)), and bond prices (Philippon (2008)).

The cash-flow effect present in our data is likely to be caused by measurement error and/or model misspecification. We draw this inference because we find cash-flow effects in our sample, even though it only contains very large firms that are unlikely to face borrowing constraints.² To investigate this possibility we run investment regressions on data generated by simulating our three models. All three models generate cash-flow and lagged-investment effects. These results suggest that the investment regressions that have received so much empirical attention are ineffectual to discriminate between alternative models.

The decreasing-returns model generates effects that are remarkably similar to those we estimate in our data. In this model these effects emerge both from measurement error in Q and from misspecification in the investment regression, since average and marginal Q do not coincide. The optimal level of investment is

¹We studied the case in which measurement error arises from the use of book value as the seed in the perpetual inventory calculation of the capital stock. However, we found that this source of error alone decayed too quickly, owing to depreciation, to have a significant effect on our estimates.

²Several authors suggest that cash-flow effects can be generated by deviations from Hayashi's (1982) assumptions. For example, Schiantarelli and Georgoutsos (1990), Cooper and Ejarque (2003), Gomes (2001), Alti (2003), and Moyen (2004) study the implications of decreasing returns to scale, while Abel and Eberly (2001, 2005) analyze the effects of growth options.

a function of three state variables: the capital stock, the shock, and the regime. So *any* additional independent variable that is correlated with the state variables has explanatory power in a regression equation. As a result, cash-flow and lagged-investment effects emerge naturally, even though the model is *not* designed to produce them.

In this respect our analysis differs from that of Cooper and Ejarque (2003). These authors estimate a model similar to our decreasing-returns model to match five moments estimated by Gilchrist and Himmelberg (1995): the coefficients of cash-flow and Q in an investment regression, the serial correlation of investment rates, the standard deviation of profit rates, and the average value of average Q . Cooper and Ejarque (2003) show that it is possible to parameterize a model with no financial frictions so that it matches the cash-flow effects observed in the data.

Our paper is organized as follows. In Section 2 we present the decreasing-returns model. In Section 3 we discuss our data and estimation procedure. Section 4 presents the results for a version of the decreasing-returns model in which the demand or productivity shock has a single regime. We also discuss the effects of introducing asymmetric investment adjustment costs, investment irreversibility, a variable discount factor, as well as a behavioral bias. In Section 5 we discuss results for the regime-switching version of the model.³ Section 6 considers the Hayashi model. Section 7 contains results for the CEE model. Section 8 concludes.

³Since our estimates are based on firm-level data, this result does not imply that these features are not useful to understand investment in less aggregated data (e.g. at the plant level or in smaller firms).

2. The decreasing-returns model

The firm's problem is given by the following Bellman equation, where y' denotes next period's value of variable y :

$$\begin{aligned} V(K, X, z) &= \max_I [zK^\alpha X^{1-\alpha} - \phi X - \xi [I/K - \delta - (\gamma - 1)]^2 K - I \\ &\quad + \beta \int V(K', X', z') F(dz', z)], \\ K' &= I + (1 - \delta)K. \end{aligned}$$

The variable X denotes the level of exogenous technological progress. This variable grows at a constant rate $\gamma > 1$:

$$X' = \gamma X.$$

The production structure of this model can be given several interpretations. The stochastic variable z , which is governed by the distribution $F(\cdot)$, represents a shock to the revenue function, such as productivity or the price of the firm's output. This output is given by $zK^\alpha X^{1-\alpha}$. We can interpret the production function as requiring a single production factor, capital. Alternatively, we can think of output as being produced with capital, labor, and other variable factors, with labor and variable factors being adjustable without frictions. In this case $zK^\alpha X^{1-\alpha}$ represents output net of labor and other variable costs. Under this interpretation, which we adopt throughout the paper, the variable z can also represent movements in the real wage and the price of variable factors.

We assume that $\alpha < 1$. We can interpret this property as reflecting the presence of decreasing returns to scale in production. Alternatively, we can think of $\alpha < 1$ as resulting from a setting in which the production function exhibits constant-returns to scale but the firm has monopoly power and a faces constant-elasticity demand function.

The function $V(K, X, z)$ represents the value of a firm with capital stock K , technical progress, X , and total factor productivity, z . We denote the discount factor by β . Capital depreciates at rate δ . The variable ϕ represents a fixed operating cost paid in every period.

Investment, denoted by I , is subject to quadratic adjustment costs, which are represented by the term $\xi [I/K - \delta - (\gamma - 1)]^2 K$. This formulation has the property that adjustment costs are zero when the firm grows at its steady state growth rate, γ . The parameter ξ controls the size of the adjustment costs.

We define cash-flow (CF_t) as:

$$CF_t = zK^\alpha X^{1-\alpha} - \phi X - \xi [I/K - \delta - (\gamma - 1)]^2 K,$$

so we interpret investment adjustment costs as reducing output or revenue.

Investment, average, and marginal Q The optimal level of investment is given by:

$$1 + 2\xi [I/K - \delta - (\gamma - 1)] = \beta \int V_1(K', X', z') F(dz', z). \quad (2.1)$$

Equation (2.1) implies that investment is a function of marginal Q , defined as the value of an additional unit of installed capital ($\beta \int V_1(K', X', z') F(dz', z)$). This function is linear as a consequence of our assumption that adjustment costs are quadratic.

Average Q is defined as the ratio of firm value to the stock of capital:

$$Q = \frac{V(K, X, z)}{K}. \quad (2.2)$$

In this model marginal and average Q do not coincide, so investment cannot be written as a linear function of Q . The difference between average and marginal Q arises for three reasons: the presence of decreasing returns to production ($\alpha < 1$),

the presence of fixed costs (ϕ), and timing considerations that result from the discrete-time nature of our model.

To explain these timing considerations consider the case in which $\alpha = 1$ and $\phi = 0$. Then average Q , defined as (2.2) is still different from marginal Q . In order for marginal and average Q to coincide we must measure Q at the end of the period. We denote end-of-period Q by Q^* :

$$Q^* = \frac{V^*(K', X', z)}{K'},$$

Here $V^*(K', X', z)$ is the end-of-period value of the firm, that is the value of the firm after it receives its cash flow and incurs the cost of investment and the adjustment costs but before it learns z' . It is easy to show that, if $\alpha = 1$ and $\phi = 0$, marginal and end-of-period average Q coincide:

$$Q^* = \beta \int V_1(K', X', z') F(dz', z).$$

Using equation (2.1) we can write investment as a linear function of Q^* .

The fact that $V^*(K', X', z)$ is computed before the firm learns z' makes it difficult to compute empirically. For this reason we use the conventional definition of Q , given by equation (2.2), in our analysis, so that it more closely corresponds to empirical measures.

Single versus Regime-Switching Regime We consider two versions of the model. In the ‘single-regime model’, z follows a Markov chain where the mean shock is normalized to one and the support is given by:

$$z \in \{1 - \sigma, 1, 1 + \sigma\}.$$

We assume that the Markov chain for the single-regime model takes the form:

$$\pi = \begin{bmatrix} p^2 & 2p(1-p) & (1-p)^2 \\ p(1-p) & p^2 + (1-p)^2 & p(1-p) \\ (1-p)^2 & 2p(1-p) & p^2 \end{bmatrix}.$$

The first-order serial correlation of the shock implied by this matrix is: $\rho = 2p - 1$ (see Rouwenhorst (1995)).

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

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

where:

$$\begin{aligned}\mu^L &= 1 - \sigma^*, \\ \mu^H &= 1 + \sigma^*.\end{aligned}\tag{2.3}$$

The variable σ^* governs the distance between the means of the two regimes. 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.

It is useful to rewrite the firm’s problem in terms of detrended variables, $k = K/X$, $i = I/X$, and $v(k, z) = V(K, X, z)/X$:

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

We solve the model using the value-function iteration method (see Appendix 9.3).

3. Estimation

In this section we first describe the data used in our estimation and summarize some key features using simple regressions. We then describe our estimation procedure.

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 intended focus of our study. Our sample includes 776 firms and roughly 14,000 firm-year observations. We focus our analysis on the large firms in our data, defined as being those in the top quartile of firms sorted by size of the capital stock in 1981. In the beginning of the sample, the top quartile of firms represents 30 percent of aggregate private non-residential investment and 40 percent of corporate non-residential investment. We use data for the four variables present in our model: investment in property, plant, and equipment, the physical capital stock, 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. We use the book value of capital as the starting value for the capital stock and four-digit industry-specific estimates of the depreciation rate. 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 Appendix 9.1.

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 sub-periods, 1981-1992 and 1993-2003. Standard errors are indicated in parenthesis. We report the median across firms of selected time-series moments. An alternative would have been to compute moments for the average across firms of the variables of interest. However, this procedure would eliminate the idiosyncratic variability associated with individual firms. The median time-series averages are 1.3 for Q , 0.15 for

the investment-capital ratio, and 0.17 for the cash-flow-capital ratio. We report the standard deviations for both the logarithms and levels of the main variables. Q is the most volatile variable, closely followed by cash flow/capital and investment/capital. The estimates in Table 1 are similar to those reported in other studies that use Compustat data.

There are important differences across sub-samples. In particular, the mean and standard deviation of Q and cash flow in the second sub-sample are significantly higher than in the earlier period. All variables exhibit positive skewness, and there is more skewness in the full sample than in each of the two sub-samples. The systematic differences across sub-samples lead us to consider a regime switching model in our estimation strategy. Finally, Q exhibits strong serial correlation, while investment and cash flow exhibit moderate persistence.

In Table 2 we report pooled, time-series-cross-section regressions of the investment-capital ratio on $\log(Q)$, $\log(\text{cash flow/capital})$ and the lagged investment/capital ratio.⁴ The coefficient on $\log(Q)$ is quantitatively small (0.06), but significant, with modest explanatory power ($R^2 = 0.29$). Including cash flow increases significantly the explanatory power of the regression ($R^2 = 0.34$) and reduces the size (0.03) and significance of the coefficient on Q . Cash flow has a large and statistically significant effect on the investment-capital ratio. As discussed in the introduction, this cash-flow effect is surprising since we use data for the top quartile of Compustat firms, which a priori are unlikely to face borrowing constraints. We view this effect as stemming from measurement error and/or misspecification. We explore these possibilities in sections 4 and 5. Adding the lagged investment-

⁴We use a semi-log specification since, as discussed in Abel and Eberly (2002), the log specification fits the data better. The skewness in the independent variables, Q and the cash flow/capital ratio, favors the semi-log specification over a conventional linear regression. When we run linear regressions, the coefficient on Q is small but significant, and the coefficient on cash flow/capital is larger and also statistically significant. These results accord with the investment regression results reported in the literature.

capital ratio to the regression leads to a large improvement in the goodness of fit ($R^2 = 0.61$). Even though much of the investment literature focuses on the cash-flow effect, the lagged-investment effect is more important from an empirical standpoint.

Figures 1 through 3 provide scatter plots of pooled time-series-cross-section data that are useful to visualize the relation between different variables. Figure 1 shows a scatter plot of investment versus $\log(Q)$. Figure 2 shows a scatter plot of investment and $\log(\text{cash flow}/\text{capital})$. Figure 3 shows the close correlation between the investment-capital ratio and its lagged value.

Figure 4 displays the histograms for I/K , average Q , and $\text{cash-flow}/K$ for the firms in our sample. Figure 4 depicts three features of the data that are relevant for our analysis. The first feature is that only one percent of the observations for I/K are near zero. This finding suggests that investment irreversibilities do not play an important role in our data. In fact, when we introduce investment irreversibility into our model, in section 4.3, the irreversibility constraint never binds and the fit of our model does not improve. The second feature is that there are not large spikes in investment, suggesting that fixed costs of investment are not likely to be useful in matching the properties of the investment data. The third feature is that all three variables are skewed. This property suggests that asymmetric adjustment costs to investment may not be needed to match the skewness in investment, since the underlying cash flows and Q are already skewed. We investigate these possibilities formally in section 4.3.

3.2. Estimation Procedure

Our solution method does not yield an analytical representation for the population moments implied by the model. For this reason, we estimate the model using the simulated method of moments proposed by Lee and Ingram (1991). We first

use our data to estimate the vector of moments Ψ_D , as described in Section 3.1. We focus on the moments that are most directly related to the parameters of the model. The moment vector that we use to estimate the single regime model includes the mean, standard deviation, and serial correlation of cash flow (to identify the shock process), the variance of investment (to identify adjustment costs), and the mean of Q (to identify the fixed cost). We find the parameter vector $\hat{\Phi}$ that minimizes the distance between the empirical and simulated moments, $\Psi(\hat{\Phi})$ computed for the median firm,

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

The weighting matrix W is computed using a block-bootstrap method on our panel dataset (see 9.6 for a description). This estimation method gives a larger weight to moments that are more precisely estimated in the data

We solve the minimization problem (3.1) using an annealing algorithm. This procedure is used to avoid convergence to a local minimum. Finally, 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. The estimation method is discussed in more details in Appendix 9.6.

4. Results: decreasing-returns model, Single Regime

We choose the exogenous rate of technical progress to be $\gamma = 1.03$. This growth rate is chosen to match the real annual growth rate of corporate net cash flows

from January 1981 to January 2004. We set $\alpha = 0.8$. This value is consistent with the estimate of the average degree of returns to scale across industries by Burnside (1996). We fix α because we cannot separately identify α and ξ using the moments of the data that we consider. Both parameters control curvature, so when α changes, the value of ξ can be adjusted to restore the fit of the model.

4.1. Parameter and moment estimates

We report our parameter estimates and standard errors in Table 3. Our estimate of the adjustment cost parameter, ξ , is 0.4148 (with a standard error of 0.0035). This estimate implies that the average investment adjustment cost is 0.8 percent of revenue net of variable costs. Our estimate for the fixed operating cost, ϕ , is 87.07 (with a standard error of 2.23). This estimate implies annual fixed operating costs that are 22 percent of revenue net of variable costs. We normalize the average shock z to one. We estimate the spread between shocks to be 0.522. As we discuss below, these values allow the model to match the mean and standard deviation of the cash-flow to capital ratio in the data.

Table 1 reports summary statistics for a panel of firms constructed by simulating our model. The moments in bold are included in the Ψ_D vector, so our estimation algorithm seeks to make these moments as close as possible to those estimated from Compustat data. The algorithm matches all of these moments closely. The remaining moments are not “targeted” by the algorithm.

Table 1 shows that the single-regime model matches well the first-order serial correlations of sales, cash flow, and investment, although Q is significantly less persistent than in the data. Our main finding is that the model generates a much lower standard deviation and skewness of Q than those we find in the data. The volatility of Q generated by the model is one-fourth of the volatility of Q present in the data (0.157 versus 0.625).

We estimate the measurement error in Q so as to match the standard deviation, first-order serial correlation, and skewness of our empirical measure of Q . The estimated noise process generates $Q_t^{noise} = Q_t \exp(\varepsilon_t) + 0.7486\varepsilon_t$, where $\varepsilon_{t+1} = 0.8761\varepsilon_t + 0.1369\eta_{t+1}$ and $\eta_t \sim N(0, 1)$. Since the measurement error is serially correlated, it cannot be corrected in the investment regressions using instrumental variables.

4.2. Simulated regression results

To evaluate the performance of our model from a different angle we estimate investment regressions on a panel of firms constructed by simulating our model. We use as explanatory variables both the state variables, which are only observable in the model, as well as Q , cash flow, and lagged investment. We report our results in Table 2. The first column shows that regressing investment on the true state variables of the model (k and the shock, z) using a semi-log specification yields an R^2 of 0.95. This specification proves a very good description of how optimal investment depends on the state variables.

We discussed in Section 2 three reasons why average Q is different from marginal Q : the presence of decreasing returns to production, the presence of fixed costs, and timing issues that result from our discrete-time formulation. Still, we obtain a very good fit when we regress investment on $\log(Q)$ because, when Q is measured without error, average Q is an excellent proxy for marginal Q . In this sense, this model is not much different from the original Hayashi (1982) model.

When we use the noisy version of Q in our investment regressions the R^2 falls to 0.04 and the coefficient on Q is 0.018 (compared to 0.466 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 0.079, and the R^2 rises to 0.70. The final column reports the results of replacing cash flow with z in the investment

regression. This substitution yields a R^2 that is similar to that obtained using cash flow as a dependent variable. Since there are no frictions in the model, cash flow enters significantly in the regression because it is a proxy for the shock, z .

One shortcoming of the single-regime model is that it cannot explain the role of lagged investment in investment regressions. When we include lagged investment in the model-based regressions we obtain a very small coefficient (0.03, compared to 0.63 in the data) and no increase in explanatory power.

In summary, the decreasing-returns model can generate a cash-flow effect because when Q is measured with error, cash flow is a proxy for z . We also find that the model is inconsistent with the importance of lagged investment in investment regressions and with the skewness properties of Q , cash flow, and investment. In the next section we show that the performance of the model can be greatly improved by adding a regime-switching component to the Markov chain for z .

4.3. Other model specifications

We explored several alternative model specifications to identify the features that are important to replicate the key moments of our data. We considered different specifications of the adjustment cost function, a time-varying discount factor, as well as a behavioral bias.

The skewness in investment led us to consider asymmetric adjustment costs, both in the form of asymmetric quadratic adjustment costs and an irreversibility constraint. The asymmetric adjustment costs that we considered take the form:

$$\begin{aligned} \xi_1 (I/K - \delta)^2 K & \text{ if } I/K > \delta, \\ \xi_2 (I/K - \delta)^2 K & \text{ if } I/K < \delta. \end{aligned}$$

This formulation is similar to that considered in Zhang (2005). When $\xi_2 > \xi_1$, this formulation can match the skewness in investment. It does not, however, generate enough skewness and volatility in Q , and cannot explain the presence of

significant lagged-investment effects in empirical regressions.

We studied a version of the model that incorporates irreversibility in investment. This constraint is irrelevant because it never binds both in our data and in our model, simulated using the estimated parameter values. This result is not surprising. Other authors, such as Doms and Dunne (1998) show that aggregating data for smaller firms or for individual plants tends to smooth out non-convexities in investment.⁵

We found that introducing empirically plausible variability in the discount factor had almost no impact on the implications of our model for the moments of interest. For this reason, we computed our main results using a constant discount rate.

We introduced a behavioral bias into the model. Specifically, we assumed that managers forecast fundamentals using the correct Markov chain but investors forecast future shocks using a distorted Markov chain with higher persistence (larger diagonal values). This specification generated enough volatility in Q , but failed to replicate the skewness of Q found in the data.

Finally, we re-estimated the model using a more flexible specification for the shock distribution that allows for a skewed support. This model can match the skewness of investment in the data, but it requires skewness in cash flow that is four times greater than in the data. Moreover, even when we match the skewness of investment with this approach, there is still no lagged investment effect.

⁵Cooper and Haltiwanger (2006) use the Longitudinal Research Database to show that the properties of investment at the plant level are very different from those at the firm level. They estimate a model that captures key features of investment at the plant level. Since the plant-level data do not include Q and cash flow, these variables are not part of their analysis.

5. Results: decreasing-returns model with Regime Switching

The regime-switching specification allows for a second regime in the productivity shock z . The average shock is normalized to one. We separately estimate spreads across regimes (σ^* , see equation (2.3)) and within regimes (σ^L and σ^H). We also estimate the discount factor, the persistence of the shocks, as well as the switching parameters in the Markov chain. In all our regime-switching specifications, we use a moment vector that includes the mean and standard deviation of cash flow in both regimes, the overall standard deviation and serial correlation of cash flow, the mean of Q in both regimes and its overall serial correlation, and the standard deviation and skewness of investment. These moments are reported in bold in Table 1.

5.1. Parameter and moment estimates

We report the estimated model parameters and standard errors in Table 3. Our estimate for the adjustment cost parameter, ξ , is 0.9028 (with a standard error of 0.022). This estimate is much larger than the one we obtained for the single regime model (0.4148). This difference reflects the fact that the support of z is much wider in the regime-switching model. In the absence of adjustment costs, this wider support would generate higher volatility of investment than that of the single regime model. As a result, we need higher adjustment costs in the regime-switching model to match the empirical volatility of investment.

This value of ξ implies that the average investment adjustment cost is 1.3 percent of revenue net of variable costs. The estimated fixed operating cost, ϕ , is 87.81 (with a standard error of 1.74), which is similar to the value found for the single-regime model. These estimates imply that annual fixed operating costs are 25.1 percent of revenue minus variable costs.

Figure 5 plots the shocks in the two regimes. The high regime has a higher average productivity, but also a higher standard deviation. It is interesting to note that the support of the two regimes overlap. In fact, the low shock in the high regime is lower than the high shock in the low regime. All of these parameters are precisely estimated. The estimated Markov chain described in Table 4 exhibits strong persistence: the parameter ρ is 0.5289 (recall that our data has annual frequency). We also estimate the probabilities of switching regime from either the middle state or from the 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). These probabilities are 3.63 percent and 17.59 percent, respectively. These estimates imply that the (unconditional) probability of a regime switch is approximately 7 percent per year.

Table 1 reports summary statistics for the panel of firms simulated using the regime-switching model. The highlighted moments are included in the Ψ_D vector. The algorithm matches all of these moments quite closely. These results indicate that incorporating regime switching improves the fit of the model, particularly for the higher moments of the data. Compared to the single-regime model, the standard deviation of Q is substantially higher, and the model generates skewness in Q and investment that are much closer to the data. The serial correlation properties are also better than those of the single regime model. Before running investment regressions, we again add measurement error to Q . We estimate the measurement error process to match the standard deviation, persistence, and skewness of Q .⁶

In order to better understand the dynamics of the model, we calculate the elasticity of each moment in the Ψ_D vector with respect to the parameters of the model. This exercise shows how changes in parameter values affect the model's

⁶We generate $Q_t^{noise} = Q_t \exp(\varepsilon_t) + 0.1696\varepsilon_t$, where $\varepsilon_{t+1} = 0.8379\varepsilon_t + 1.0813\eta_{t+1}$ and $\eta_t \sim N(0, 1)$.

performance. We report this elasticity matrix in Table 5. In the first row of the table we see that average Q in the first (low) regime is heavily influenced by the fixed operating cost, ϕ , as well as by the discount factor β . The fixed operating cost, ϕ , influences only the moments of Q and cash flow but has no impact on the moments of investment.

Since we keep the average shock, z , constant in the model, the average cash flow for each of the two regimes is largely determined by the spread σ across regimes. This parameter establishes in turn the mean shocks μ^L and μ^H and affects the average cash flow in each regime. 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. The standard deviation of the investment-capital ratio is largely determined by the adjustment cost parameter ξ . The spread parameter σ is also an important determinant as it affects the volatility of investment across regimes. Finally, the skewness of investment is heavily influenced by the serial correlation of the shock.

Figure 6a and 6b plot the value functions and policy functions for each state in the two regimes as a function of the firm's capital stock. The lower bounds of the support of z in the two regimes ($\mu^L - \sigma^L$ and $\mu^H - \sigma^H$) are very similar. However, the value and policy functions evaluated at these two lower bounds take on very different values. The value of the firm is higher when the shock is $\mu^H - \sigma^H$ rather than when it is $\mu^L - \sigma^L$ even though $\mu^H - \sigma^H < \mu^L - \sigma^L$. This property reflects the fact that the probability of transiting to the highest value of the shock, $\mu^H + \sigma^H$, is higher when the current state is $\mu^H - \sigma^H$ than when the state is $\mu^L - \sigma^L$.

5.2. Simulated regression results

We now regress investment on its determinants using simulated data. We report our results in Table 2. In the first column, we use K , z , and a dummy variable

for the regime to explain investment using a semi-log specification. As in the single regime model, this specification provides a good approximation to the policy function for the investment-capital ratio, with a R^2 of 0.97.

A regression of investment on Q has a R^2 of only 0.56 (compared to 0.95 for the single-regime model) and the Q coefficient is equal to 0.1278. The difference between average and marginal Q is greater in this model, relative to the single-regime model, because the support of z is much wider.

If we use the noisy measure of Q the coefficient on Q falls to 0.0226 and the R^2 drops sharply to 0.12. When we control for the regime the R^2 rises from 0.12 to 0.35 while the coefficient on Q falls from 0.0226 to 0.0085.

When we add cash flow to this regression, the coefficient on Q falls to 0.0161. Cash flow enters significantly with a coefficient of 0.0364 and the R^2 rises from 0.12 to 0.21. As in the single regime model, we obtain similar results when we replace cash flow with z . Cash flow enters significantly in the investment regression because it is a proxy for z .

Finally, including lagged investment in the regression improves the fit considerably in both model and data regressions, lowering the coefficients on Q and cash flow. The parameter estimates are very similar in model and data regressions. Recall that this similarity is not present in the investment regressions for the single-regime model. In those regressions lagged investment is driven out by cash flow (see Table 2).

The presence of regime switching improves the ability of the model to fit the moments of the data. It also helps the model match the empirical covariation and partial covariation among investment, cash flow, and Q . These results suggest that the presence of regime switching is crucial to understanding investment regressions. In the data and in the simulation, both the true Q and noisy Q have relatively poor explanatory power for investment when there is regime switching

(Table 2). 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 raised the R^2 from 0.04 to 0.70. In the regime switching model, the addition of cash flow only increases the R^2 from 0.12 to 0.21. Figure 7 illustrates this property. It plots the investment rate, i/k , as a function of the capital stock for each value of the shock, z , in the regime switching model. The relation between the current shock and current investment is non-monotonic. The lowest investment rates occur on the lowest branch of the graph, when the shock is in the low regime and $z = 0.5957$. Investment rates are substantially higher when the shock is in the high regime and z takes on its lowest value: $z = 0.5701$. This property results from the fact that the probability of transiting between regimes is low. Within the high regime, even when current z is very low, future prospects are bright because there is a high probability of transiting to a high value of z . In the low regime, even when current z is high, the prospects for the future are relatively bleak and thus investment remains low. The transition dynamics within and across regimes break the monotonic relation between investment and z and between investment and cash flow.

A similar argument explains why the regime-switching model can replicate the lagged-investment effect present in the data. Since regime changes do not occur often, last period's level of investment is a good indicator of the current regime. In other words, lagged investment acts as a proxy for an aspect of the shock process (the regime) that is not embodied by cash flow. In contrast, in the single-regime model, the close relation between the shock and cash flow makes lagged investment redundant in explaining current investment.

6. Hayashi's Model

In this section we study a version of Hayashi's model by considering a special case of the decreasing-returns model in which returns to scale are constant ($\alpha = 1$) and the fixed cost of operating is zero ($\phi = 0$).

The firm's problem is given by the following Bellman equation:

$$V(K, z) = \max_{i, k'} [zK - \xi (I/K - v)^2 K - I + \beta \int V(K', z') F(dz', z)], \quad (6.1)$$

subject to:

$$K' = I + (1 - \delta)K. \quad (6.2)$$

We consider a regime-switching process and choose the Markov chain and the support of z so that the model matches the empirical volatility of the cash-flow-to-capital ratio. The support of z is given by:⁷

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

We solve the model taking advantage of the fact that the value function is homogeneous of degree one (see 9.4 for details).

One interesting finding is that if we set $v = \delta$, this model fails to match even the most basic moments of the data, such as the average value of Q and the volatility of I/K . The fact that the model generates infinite values for V and Q for many parameter combinations is at the heart of this failure. When the discount factor is high (i.e. the real interest rate is low) the average values of V and Q are often infinity. The value of the firm is finite only when the adjustment cost

⁷The performance of this regime-switching version of Hayashi's model is much better than that of a single-regime version. To conserve space we do not report results for the single-regime version.

parameter, ξ , is very high. However, high adjustment costs imply low investment volatility. When the discount factor is low (i.e. the real interest rate is high) it is possible to generate a finite firm value with low values of ξ . However, the low discount factor produces very low values for Q .

We now report results for a version of the model in which we estimate v . Table 3 reports the parameter estimates and standard errors for the Hayashi model with regime switching. The estimate of the adjustment cost parameter, ξ , is much higher than that obtained for the decreasing-returns model (3.986 versus 0.903). In the absence of adjustment costs investment would be finite in the decreasing-returns model because of the presence of decreasing returns to production. In contrast, without adjustment costs investment in the Hayashi model would alternate between $+\infty$ (when $z - \delta > 1/\beta - 1$) and $-\infty$ (when $z - \delta < 1/\beta - 1$). As a consequence, we need higher adjustment costs in the Hayashi model in order to match the volatility of investment observed in the data. Our parameter estimates imply that adjustment costs represent on average 4.6 percent of revenue net of variable costs.

Table 1 compares the implied data moments from the model to those in the data.⁸ The model matches closely the data moments, including the average level of Q in both regimes, and the overall volatility and skewness of Q . Since investment closely tracks Q in this model, overall investment volatility and skewness also match the data. However, the adjustment cost required to match the data reduces investment volatility within regimes (for example, the volatility of investment is 0.016 in the low regime, compared to 0.05 in the data). The model requires a large change in the average level of I_t/K_t across regimes (from 0.112 to 0.210 from the low to high regimes) that is not present in the data. Overall, the fit is comparable

⁸We added the average level of I_t/K_t to the moment vector used in the estimation of the Hayashi model. In the generalized-Hayashi model the ratio I_t/K_t is determined by the depreciation rate and the long run growth rate. This property is not present in the Hayashi model.

to that of the decreasing-returns model. In some dimensions the fit is superior (e.g., the dynamics of Q) in the Hayashi model, while in others (e.g., investment dynamics) the decreasing-returns model is a better fit.

6.1. Simulated regression results

Table 2 reports the results of estimating investment regressions on data simulated from the Hayashi model. The only reason why Q is not a sufficient statistic for investment is the timing issue that arises in discrete time, which we discuss in Section 2. So, it is not surprising that we find that Q is an excellent predictor of investment: the R^2 of the regression of investment on $\log(Q)$ is 0.98.

The second set of regressions use a version of the model where Q is measured with error. As with our previous model, this measurement error process is estimated so that the resulting Q matches the empirical standard deviation, persistence and skewness of Q .⁹ In this version of the model Q is no longer a sufficient statistic for the choice of investment, and cash-flow and lagged-investment effects emerge. However, these effects are much weaker than in the data. Regressing investment on noisy Q alone generates an R^2 of 0.37; adding only cash flow reduces the coefficient on Q from 0.0877 to 0.0713 with a coefficient on cash flow of 0.0339. Adding lagged investment raises the R^2 further to 0.83, with a lagged investment coefficient of 0.5871. In this specification the coefficient on Q is twice as large as it is in the data.

7. CEE Model

Many recent macroeconomic models incorporate a form of adjustment costs proposed by Christiano, Eichenbaum, and Evans (2005). In this formulation, adjust-

⁹We generate $Q_t^{noise} = Q_t \exp(\varepsilon_t) + 3.6559\varepsilon_t$, where $\varepsilon_{t+1} = 0.8189\varepsilon_t + 0.0419\eta_{t+1}$ and $\eta_t \sim N(0, 1)$.

ment costs depend on changes in the level of investment, so lagged investment effects are likely to arise naturally in investment regressions. In this section we study the properties of a version of our model that incorporates CEE-style adjustment costs.

The firm's problem, written in terms of detrended variables, is given by:

$$v(k, i_{-1}, z) = \max_{i, k'} \left[zk^\alpha - i - \phi + \beta\gamma \int V(k', i, z') F(dz', z) \right],$$

subject to:

$$\gamma k' = i [1 - \xi^*(\gamma i/i_{-1} - \gamma)^2] + (1 - \delta)k. \quad (7.1)$$

Here i_{-1} denotes the value of investment in the previous period. The presence of a third state variable in the value function requires us to adopt a different algorithm to solve the model. We describe this algorithm in the appendix.

There are four reasons why average and marginal Q do not coincide in this model. The first three reasons are common to the decreasing-returns model: there are decreasing returns to production, a fixed cost, and the timing issue that arises in discrete time. The fourth reason has to do with the fact that the value function depends, not only on k and z , but also on i_{-1} . If we set $\alpha = 1$ and $\phi = 0$ in the Hayashi model we obtain a value function that is homogeneous of degree one and so: $V(k, z) = V_1(k, z)k$, implying that $V(k, z)/k = V_1(k, z)$. If we set $\alpha = 1$ and $\phi = 0$ in the CEE model the value function is homogeneous in degree one in k and i_{-1} .¹⁰ This property implies that: $v(k, i_{-1}, z) = v_1(k, i_{-1}, z)k + v_2(k, i_{-1}, z)i_{-1}$. So, $v(k, i_{-1}, z)/k \neq v_1(k, i_{-1}, z)$.

We estimate that the adjustment cost parameter, ξ^* , is equal to 0.88, with a standard error of 0.022.¹¹ The other parameter estimates, shown in Table 3, are

¹⁰See Jaimovich and Rebelo (2008) for a proof.

¹¹The value of ξ^* estimated by CEE using macro data and a model with a constant returns to scale in production is 1.24. CEE estimate $\eta''(1) = 2.48$, where $\eta''(1)$ is the second derivative of the adjustment cost function evaluated at the steady state. In our case the adjustment cost function is quadratic, so $\xi^* = \eta''(1)/2$.

close to those for the decreasing-returns model. Average adjustment costs as a fraction of revenue net of variable costs are 0.8 percent. The fixed cost represents 25.4 percent of revenue net of variable costs.

Table 1 shows that the fit of the model with CEE adjustment costs is generally very good. This fit is comparable to that of the decreasing-returns model with two exceptions. First, the CEE formulation generates too much investment persistence. The first-order serial correlation of investment is 0.94 in the model and 0.60 in the data. The high degree of investment persistence generated by the model is not surprising since this specification penalizes changes in the level of investment. Second, the model does not generate enough skewness in investment (0.31 versus 0.42). This property is a direct consequence of the adjustment cost specification: an increase in ξ^* reduces both the standard deviation and skewness of investment, and the estimation procedure cannot find a set of parameter values which fits both moments.

Table 2 reports the results of estimating investment regressions on data simulated from the model with CEE adjustment costs.¹² This model generates a regression coefficient on Q that is very similar to the data. The cash-flow effect is weak and sometimes negative. The model generates a lagged investment effect that is much stronger than that found in the data (0.9275 versus 0.6253). This property reflects the fact that lagged investment is a state variable in this model.

8. Conclusions

We estimate three models of investment and examine their implications for the mean, standard deviation, skewness and persistence of investment, cash flow, and Q . While all three models can closely match the key data moments, the decreasing-

¹²We generate $Q_t^{noise} = Q_t \exp(\varepsilon_t) + 1.6006\varepsilon_t$, where $\varepsilon_{t+1} = 0.8275\varepsilon_t + 0.0865\eta_{t+1}$ and $\eta_t \sim N(0, 1)$.

returns model and the Hayashi model both replicate the salient empirical features of investment, cash flow and value in our sample of large firms. These models would nonetheless be rejected by tests based on investment regressions. We find empirically plausible cash-flow and lagged-investment effects in data simulated from these models when we incorporate our estimates of measurement error in the construction of Q . This result illustrates the importance of going beyond investment regressions when assessing investment models.

The estimated regime-switching process for the shocks is an important feature of the model. This process generates skewness in cash flow, Q , and investment, and also implies that Q and cash flow are not informationally redundant for the investment decision. When Q is mismeasured or misspecified (as in the decreasing-returns model), current cash flow does not perfectly predict expected investment opportunities. Instead there is a role for both cash flow and lagged investment to predict current investment, even when controlling for Tobin's Q . These findings show that a neoclassical model with quadratic adjustment costs can match the key data moments of large publicly-traded firms, while simultaneously generating empirically relevant cash flow and lagged investment effects.

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, 2002.
- [3] Abel, Andrew, and Janice Eberly, “Investment, Valuation, and Growth Options,” working paper, 2005.
- [4] Altı, Aydogan “How Sensitive Is Investment to Cash Flow When Financing Is Frictionless?” *The Journal of Finance*, 58: 707-722, 2003.
- [5] Burnside, Craig “Production Function Regression, Returns to Scale and Externalities,” *Journal of Monetary Economics*, 37: 177-201, 1996.
- [6] Christiano, Lawrence, Martin Eichenbaum, and Charles Evans “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113: 1–45, 2005.
- [7] Cooper, Russell and Joao Ejarque, “Financial Frictions and Investment: A Requiem in Q ,” *Review of Economic Dynamics*, 6: 710-28, 2003.
- [8] Cooper, Russell and John Haltiwanger, "On the Nature of Capital Adjustment Costs," *Review of Economic Studies* 73 (3) , 611–633, 2006.
- [9] Cummins, Jason G., Kevin A. Hassett, and Stephen D. Oliner “Investment Behavior, Observable Expectations, and Internal Funds,” *American Economic Review*, 96: 796-810, 2006.
- [10] Doms, Mark and Timothy Dunne, “Capital Adjustment Patterns in Manufacturing Plants,” *Review of Economic Dynamics*, 1: 409-429, 1998.

- [11] Erickson, T. and Toni Whited, "Measurement Error and the Relationship between Investment and 'q'," *The Journal of Political Economy*, 108: 5. 1027-1057, 2000.
- [12] Fazzari, Steven, R. Glenn Hubbard, and Bruce C. Petersen, "Financing Constraints and Corporate Investment," *Brookings Papers on Economic Activity*, 1:141-206, 1988.
- [13] Gilchrist, Simon and Charles Himmelberg, "Evidence on the role of cash flow for investment," *Journal of Monetary Economics* 36: 541-572, 1995.
- [14] Gomes, Joao, "Financing Investment," *The American Economic Review*, 91: 1263-1285, 2001.
- [15] Hayashi, Fumio, "Tobin's marginal Q and average Q : A neoclassical interpretation," *Econometrica*, 50, 215-224, 1982.
- [16] Jaimovich, Nir and Sergio Rebelo "Tobin's Marginal Q and Average Q Revisited," mimeo, Northwestern University, 2008.
- [17] Jorgenson, Dale and Ralph Landau, *Tax Reform and the Cost of Capital: an International Comparison*, Washington, DC: Brookings Institution, 1993.
- [18] King, Mervyn and Don Fullerton, *The Taxation of Income from Capital: a Comparative Study of the United States, the United Kingdom, Sweden, and West Germany*; University of Chicago Press, Chicago, 1984.
- [19] Lkhagvasuren, Damba "Big Locational Differences in Unemployment Despite High Labor Mobility," mimeo, Northwestern University, 2006.
- [20] Lee, B. and Ingram B. "Simulation estimation of time series models", *Journal of Econometrics*, 47, 197-205, 1991.

- [21] Moyen, Nathalie “Investment-Cash Flow Sensitivities: Constrained versus Unconstrained Firms,” *The Journal of Finance* 59: 2061–2092, 2004.
- [22] Philippon, Thomas “The bond market’s Q”, *Quarterly Journal of Economics*, forthcoming, 2008.
- [23] Rouwenhorst, K. Geert “Asset Pricing Implications of Equilibrium Business Cycle Models,” in Thomas Cooley, ed. *Frontiers of Business Cycle Research*, Princeton University Press, 1995.
- [24] Salinger, Michael A. and Lawrence Summers, “Tax Reform and Corporate Investment: A Microeconomic Simulation Study,” in Martin S. Feldstein, editor, *Behavioral Simulation Methods in Tax Policy Analysis*, University of Chicago Press, 1983, 247-287.
- [25] Schiantarelli, Fabio, and Dimitri Georgoutsos, “Monopolistic competition and the Q theory of investment, *European Economic Review*, 34: 1061-1078, 1990.
- [26] Shiller, Robert J. *Market volatility*, Cambridge, MA: MIT Press, 1989.
- [27] Shiller, Robert J. *Irrational exuberance*, Princeton: Princeton University Press, 2000.
- [28] Zhang, Lu “The Value Premium,” *The Journal of Finance*, 60: 67-103, 2005.

9. Appendix

9.1. Data Sources and Calculations

Annual data items from the dataset *cstsann* in the CRSP/Compustat Merged database, 1981-2003, are first listed, followed by the calculations underlying the constructed variables. Sources for non-Compustat items are given in parentheses.

- *I* : expenditures on property, plant, and equipment, data 30
- *CashFlow*: income before extraordinary items + depreciation and amortization + minor adjustments, calculated as follows (from the Compustat manual):

Income Before Extraordinary Items, 123
+ Depreciation and Amortization, 125
+ Extraordinary Items and Discontinued Operations, 124
+ Deferred Taxes, 126
+ Equity in Net Loss (Earnings), 106
+ Sale of Property, Plant, and Equipment and Sale of Investments – Loss(Gain),
213
+ Funds from Operations – Other, 217
+ Accounts Receivable – Decrease (Increase), 302
+ Inventory – Decrease (Increase), 303
+ Accounts Payable and Accrued Liabilities – Increase (Decrease), 304
+ Income Taxes – Accrued – Increase (Decrease), 305
+ Assets and Liabilities – Other (Net Change), 307
= Operating Activities – Net Cash Flow, 308

- *inventories*: total inventories (end of period), data 3

- *debt*: long-term debt (end of period), data 9
- *PPE*, book value of capital: property, plant, and equipment,
 - data 182: PPE - Beginning Balance – check if it is still reported after 1997;
 - data 187: PPE - Ending Balance (Schedule V);
 - data 184: PPE - Retirements (Schedule V) - not reported after 1997;
 - data 185: PPE - Other Changes (Schedule V) - not reported after 1997.
- P_k , price of capital: implicit price deflator for nonresidential investment, Economic Report of the President, Table B-3, various years.
- u , investment tax credit: obtained by year for 51 asset classes from Dale Jorgenson. These data are aggregated to the two-digit industry level using the BEA historical cost capital flow matrix (asset by industry by year). Specifically, the weight of asset type n in industry j in year t is calculated as $w_{n,j,t} \equiv I_{n,j,t} / \sum_n I_{n,j,t}$. The investment tax credit applied to industry j in year t , $u_{j,t}$, is then constructed as the weighted sum $u_{j,t} = \sum_n w_{j,n,t} u_{j,n,t}$.
- z , value of depreciation allowances: obtained by year for 51 asset classes from Dale Jorgenson. These data are aggregated to the two-digit industry level using the BEA historical cost capital flow matrix (asset by industry by year). Specifically, the weight of asset type n in industry j in year t is calculated as $w_{n,j,t} \equiv I_{n,j,t} / \sum_n I_{n,j,t}$. The value of depreciation allowances in industry j in year t , $z_{j,t}$, is then constructed as the weighted sum $z_{j,t} = \sum_n w_{j,n,t} z_{j,n,t}$.
- τ , corporate tax rate: obtained from King and Fullerton (1984), table 6.4, and Fullerton and Karayannis (in Jorgenson and Landau (1993)), p. 343, updated to 2003 by Dale Jorgenson.

- *market value of equity*: closing stock price times number of common shares outstanding (end of period) plus redemption value of preferred stock (end of period) = $\text{prc} * \text{shrout}/1000 + \text{data56}$, where,
 - prc: closing stock price from msf file (monthly stock - securities);
 - shrout: Common shares outstanding from msf file (monthly stock - securities);
 - data 56: Preferred Stock - Redemption Value.
- L , useful life of capital goods: by two-digit industry, the useful life of capital goods is calculated as $L_j \equiv \frac{1}{N_j} \sum_{i \in j} \frac{PPE_{i,t-1} + DEPR_{i,t-1} + I_{i,t}}{DEPR_{i,t}}$, where N_j is the number of firms, i , in industry j . Using the double-declining balance method, the implied depreciation rate for industry j , δ_j , is $2/L_j$.
- K , replacement value of capital stock: Using the method of Salinger and Summers (1983) the replacement value of the capital stock is constructed by firm from its book value using the recursion: $K_{i,t} = \left(K_{i,t-1} \frac{P_{K,t}}{P_{K,t-1}} + I_{i,t} \right) (1 - \delta_j)$, where the recursion is initialized using the book value of capital.
- Tobin's Q : $[(\text{market value of equity})_{t-1} + (\text{debt})_{t-1} - (\text{inventories})_{t-1}]/K_t$.

9.2. Sample Selection

Starting from the dataset cstsann in the CRSP/Compustat Merged database, the following filters were applied:

- If the firm was involved in a merger or acquisition, then delete (using aftnt35: = '01' as indication of a Merger & Acquisition)
- end-of-period capital (data 187) is not missing
- investment (data 30) is not missing
- operating profit (data 178) is not missing
- incorrect capital accumulation (only for data before 1994, due to data184 and data185 not being reported after 1997)
- if disinvestment $>$ end-of-period capital then delete
- if operating loss is greater than end-of-period capital then delete
- if operating profit is greater than 2.5 times end-of-period capital then delete
- if q is missing or $q < 0$ then delete
- if investment (data 30) < 0 then delete
- if dis-investment (data107) < 0 then delete

9.3. Solution Method, decreasing-returns model

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 $m = 1$:

$$V^m(k, z) = \max_{i, k'} [zk^\alpha - \phi - \xi \{[k' - (1 - \delta)k] / k - \delta\}^2 k - [\gamma k' - (1 - \delta)k] + \beta\gamma \int V^{m-1}(k', z') F(dz', z)].$$

We use $V^1(k, z)$ 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 (k, z) pair.

In practice, this method is slow to converge. To speed up the procedure in the context of our SMM estimation, which requires solving the model at every iteration, we instead adopt a hybrid method. We start with a policy function iteration approach: we iterate as above until $h^{m-1}(k, z)$ and $h^m(k, z)$ converge for every (k, z) pair. Once this is done, we iterate on the value function, keeping the policy function constant, until convergence. Not having to find a new policy function at that stage makes this hybrid procedure significantly faster.

9.4. Solution Method, Hayashi Model

The value function, $V(K, z)$, is homogeneous of degree one in the capital stock. This property follows from the fact that we can write the value function as a sum of functions that are homogeneous of degree one. The homogeneity property allows us to rewrite (6.1) as:

$$V(1, z) = \max_{i/k} [z - \xi (I/K - \delta)^2 - I/K + (I/K + 1 - \delta) \beta \int V(1, z') F(dz', z)], \quad (9.1)$$

Using the fact that $V_1(1, z) = V(1, z)$, we can write the optimal value of I/K as:

$$\frac{I}{K} = \frac{\beta \int V(1, z') F(dz', z) - 1}{2\xi} + \delta \quad (9.2)$$

We solve the model using value-function iteration. We start with a guess for the value function, $V^0(1, z)$ for each value of z . We use (9.2) to compute the optimal value of I/K associated with each value of z . We then compute the new value function, $V^1(k, z)$. This function is given by the following equation with $m = 1$:

$$V^m(1, z) = \max_{i/k} \left[z - \frac{1}{\xi} \left(\frac{\beta \int V^{m-1}(1, z') F(dz', z) - 1}{2} \right)^2 - \frac{\beta \int V^{m-1}(1, z') F(dz', z) - 1}{2\xi} + \delta + \left(\frac{\beta \int V^{m-1}(1, z') F(dz', z) - 1}{2\xi} + 1 \right) \beta \int V^{m-1}(1, z') F(dz', z) \right]. \quad (9.3)$$

9.5. Solution Method, CEE Model

We obtain numerical solutions to the model with CEE adjustment costs using the following algorithm developed in Lkhagvasuren (2006):

1. Define a coarse grid for (k, i_{-1}, z) ;
2. Choose a guess for $v(k, i_{-1}, z)$ and evaluate it on the coarse grid;
3. Choose a fine grid for i_{-1} ;
4. Generate a fine grid for k compatible with fine grid for i_{-1} using the resource constraint, (7.1);
5. Use bilinear interpolation to evaluate $v(k, i_{-1}, z)$ for every value of z on the fine grid for i_{-1} and z ;¹³
6. Find the optimal value of i for every (k, i_{-1}, z) combination;
7. Save the new value of $v(k, i_{-1}, z)$ evaluated on the coarse grid;
8. Save the policy function for i , $i(k, i_{-1}, z)$, evaluated on the fine grid;
9. Check whether the value function has converged;
10. If the value function has converged then stop; else go to step 5;

To simulate the model we can use a bilinear interpolation of $i(k, i_{-1}, z)$ evaluated for every z , for every pair (k, i_{-1}) evaluated on the fine grid. This interpolation procedure avoids k and i_{-1} having to take values on the real line.

¹³Bilinear interpolation is an extension of linear interpolation for bivariate functions. Suppose we know the values of the function $f(x, y)$ evaluated at four points: (x_1, y_1) , (x_2, y_1) , (x_1, y_2) , and (x_2, y_2) . Then $f(x, y) \simeq \frac{f(x_1, y_1)}{(x_2 - x_1)(y_2 - y_1)}(x_2 - x)(y_2 - y) + \frac{f(x_2, y_1)}{(x_2 - x_1)(y_2 - y_1)}(x - x_1)(y_2 - y) + \frac{f(x_1, y_2)}{(x_2 - x_1)(y_2 - y_1)}(x_2 - x)(y - y_1) + \frac{f(x_2, y_2)}{(x_2 - x_1)(y_2 - y_1)}(x - x_1)(y - y_1)$.

9.6. Estimation Method

The objective of the simulated method of moments is to find the parameter vector $\hat{\Phi}$ that minimizes the distance between empirical (Ψ_D) and simulated moments ($\Psi(\Phi)$):

$$L(\hat{\Phi}) = \min [\Psi(\Phi) - \Psi_D]' W [\Psi(\Phi) - \Psi_D]. \quad (9.4)$$

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

$$W = \frac{1}{\Omega_D(1 + 1/k)}, \quad (9.5)$$

where $k = \text{length of simulation} / \text{length of sample}$. We estimate the matrix Ω_D using a block-bootstrap method as follows: We form m samples. Each sample consists of data for n firms drawn 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 (9.4) using an annealing algorithm. The first step consists in choosing initial values for the parameter vector, Φ , admissible ranges for the parameters, as well as 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 worse fit ($L(\Phi') > L(\Phi^*)$). 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 parameter space.

The second step is to generate a new parameter vector, Φ' , by adding random shocks to the elements of Φ^* within their admissible range. Next we solve the model using value-function iteration for the parameter vector Φ' and simulate 1940 representative firms (each with 23 years of data). Since the number of firms

in our Compustat sample is equal to 194, this implies that k in (9.5) equals 10. The fourth step consists in computing the simulated moments and $L(\Phi')$. If $L(\Phi') < L(\Phi^*)$ we set $\Phi^* = \Phi'$. If $L(\Phi') > L(\Phi^*)$ we set $\Phi^* = \Phi'$ with probability $\exp[-(L(\Phi') - L(\Phi^*)) / \text{temperature}]$. Finally, we reduce the values of temperature and step size before going back to step two. The vector of parameter estimates is the one that generates the lowest value of L . We denote this vector by $\hat{\Phi}$.

To verify the convergence properties of our estimation procedure, we used a simple robustness check. Starting with a parameter vector $\tilde{\Phi}$, we simulate a panel of firms and compute the simulated moments, $\Psi(\tilde{\Phi})$. We then use the SMM procedure described above to fit these moments. Ideally, we would like the parameter estimates $\hat{\Phi}$ to be as close as possible to the true parameter values $\tilde{\Phi}$ (the ones that generated the data). Failure to do so may indicate that the estimation procedure is not adequate or that the model parameters are not identified. We find that our procedure can recover reasonably well the true parameter values. This is also confirmed by the fact that we obtain similar parameter estimates across SMM runs with different starting values.

Table 1: Summary statistics, data and model implications

	Median across large firms (4th quartile of Compustat firms)*			Decreasing returns model	Decreasing returns model			Hayashi model			CEE model		
	Full Sample	Subsamples		Single regime	Regime switching			Regime switching			Regime switching		
	1981-2003	1981-1992	1993-2003		All	Low	High	All	Low	High	All	Low	High
Time-series average													
Q	1.298	0.950	1.892	1.316	1.240	0.918	1.573	1.411	0.978	1.854	1.284	0.928	1.656
I/K	0.150	0.146	0.161	0.151	0.151	0.120	0.184	0.160	0.112	0.210	0.153	0.124	0.183
Cash Flow/K	0.169	0.155	0.199	0.172	0.161	0.137	0.186	0.187	0.165	0.210	0.164	0.137	0.192
Time-series standard deviations													
Q	0.625	0.256	0.589	0.157	0.383	0.151	0.239	0.542	0.169	0.418	0.433	0.162	0.294
Q + noise				0.625	0.625	0.440	0.600	0.625	0.331	0.539	0.625	0.424	0.587
ln(Q)	0.420	0.280	0.280										
I/K	0.055	0.050	0.046	0.055	0.054	0.036	0.050	0.056	0.016	0.034	0.057	0.044	0.054
ln(I/K)	0.370	0.330	0.300										
Cash Flow/K	0.078	0.046	0.089	0.079	0.073	0.043	0.087	0.082	0.050	0.101	0.074	0.044	0.087
Skewness													
Q	0.577	0.160	0.350	0.084	0.366			0.672			0.568		
Q + noise	0.577	0.160	0.350	0.577	0.577			0.577			0.577		
I/K	0.418	0.320	0.330	0.014	0.436			0.432			0.314		
Cash Flow/K	0.245	-0.040	0.050	-0.063	0.601			0.646			0.587		
Serial correlation													
Q	0.838	0.780	0.660	0.426	0.841			0.857			0.860		
Q + noise	0.838	0.780	0.660	0.838	0.838			0.838			0.838		
I/K	0.600	0.550	0.540	0.397	0.757			0.869			0.938		
Cash Flow/K	0.540	0.500	0.370	0.535	0.568			0.552			0.587		

*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 3: Parameter estimates (standard errors in parenthesis)

	Decreasing returns model		Hayashi model	CEE model
	Single regime	Regime switching	Regime switching	Regime switching
Estimated parameters				
Adjustment cost : ξ	0.4148 (0.0035)	0.9028 (0.0220)	3.986 (0.0686)	0.8793 (0.0453)
Adjustment cost : v			0.117 (0.001)	
Fixed cost: ϕ	87.07 (2.23)	87.8059 (1.738)		87.1262 (1.514)
Discount factor: β	0.9514 (0.0007)	0.9511 (0.0006)	0.9526 (0.001)	0.9508 (0.0005)
Shock range: σ	0.522 (0.0028)	0.1581 (0.0040)		0.1657 (0.0015)
Low regime center shock: μ_L			0.1557 (0.0011)	
High regime center shock: μ_H			0.2793 (0.002)	
Low regime shock range: σ_L		0.2462 (0.0021)	0.0657 (0.0006)	0.2427 (0.0022)
High regime shock range: σ_H		0.588 (0.0051)	0.1706 (0.0014)	0.5732 (0.0044)
Switching parameter 1		0.0363 (0.0021)	0.0589 (0.0203)	0.0236 (0.0018)
Switching parameter 2		0.1759 (0.0056)	0.2985 (0.0805)	0.1415 (0.0005)
Shock persistence: ρ	0.5345 (0.0013)	0.5289 (0.0051)	0.5583 (0.0226)	0.5340 (0.0023)
Calibrated parameters				
Mean shock: μ	1.00	1.00	1.00	
Returns to scale: α	0.80	0.80	1.00	
Depreciation rate: δ	0.12	0.12	0.15	
Growth: γ	1.03	1.03		

Table 4: Estimated Markov chains for regime-switching models

Decreasing returns model

Support of the distribution

Low Regime			High Regime		
$\mu^L - \sigma^L$	μ^L	$\mu^L + \sigma^L$	$\mu^H - \sigma^H$	μ^H	$\mu^H + \sigma^H$
0.5957	0.8419	1.0881	0.5701	1.1581	1.7461

Transition matrix

	$\mu^L - \sigma^L$	μ^L	$\mu^L + \sigma^L$	$\mu^H - \sigma^H$	μ^H	$\mu^H + \sigma^H$
$\mu^L - \sigma^L$	0.5844	0.3601	0.0555	0	0	0
μ^L	0.1735	0.6166	0.1735	0.0363	0	0
$\mu^L + \sigma^L$	0.0457	0.2968	0.4816	0.1759	0	0
$\mu^H - \sigma^H$	0	0	0.1759	0.4816	0.2968	0.0457
μ^H	0	0	0.0363	0.1735	0.6166	0.1735
$\mu^H + \sigma^H$	0	0	0	0.0555	0.3601	0.5844

CEE model

Support of the distribution

Low Regime			High Regime		
$\mu^L - \sigma^L$	μ^L	$\mu^L + \sigma^L$	$\mu^H - \sigma^H$	μ^H	$\mu^H + \sigma^H$
0.602	0.8542	1.1063	0.547	1.1458	1.7447

Transition matrix

	$\mu^L - \sigma^L$	μ^L	$\mu^L + \sigma^L$	$\mu^H - \sigma^H$	μ^H	$\mu^H + \sigma^H$
$\mu^L - \sigma^L$	0.5879	0.1744	0.0494	0	0	0
μ^L	0.3577	0.6265	0.325	0	0	0
$\mu^L + \sigma^L$	0.0544	0.1744	0.5341	0.0914	0.0246	0
$\mu^H - \sigma^H$	0	0.0246	0.0914	0.5341	0.1744	0.0544
μ^H	0	0	0	0.325	0.6265	0.3577
$\mu^H + \sigma^H$	0	0	0	0.0494	0.1744	0.5879

Hayashi model

Support of the distribution

Low Regime			High Regime		
$\mu^L - \sigma^L$	μ^L	$\mu^L + \sigma^L$	$\mu^H - \sigma^H$	μ^H	$\mu^H + \sigma^H$
0.09	0.1557	0.2214	0.1087	0.2793	0.45

Transition matrix

	$\mu^L - \sigma^L$	μ^L	$\mu^L + \sigma^L$	$\mu^H - \sigma^H$	μ^H	$\mu^H + \sigma^H$
$\mu^L - \sigma^L$	0.6071	0.3441	0.0488	0	0	0
μ^L	0.1619	0.6172	0.1619	0.0589	0	0
$\mu^L + \sigma^L$	0.0342	0.2414	0.4259	0.2985	0	0
$\mu^H - \sigma^H$	0	0	0.2985	0.4259	0.2414	0.0342
μ^H	0	0	0.0589	0.1619	0.6172	0.1619
$\mu^H + \sigma^H$	0	0	0	0.0488	0.3441	0.6071

Table 5: Elasticity of moments with respect to parameters, Decreasing returns model with regime-switching

	ξ	ϕ	σ^*	σ^L	σ^H	Switching parameter 1	Switching parameter 2	β	ρ
Average q, low regime	-0.2	-3.2	0.1	0.2	-0.3	0	-0.1	94.4	0
Average q, high regime	0	-1	0.3	0.1	0	-0.1	-0.1	32.2	0.2
Average cash-flow, low regime	0	-0.4	-0.3	0	-0.1	0	0	4.7	0.1
Average cash-flow, high regime	0	-0.2	0.1	0	-0.1	0	-0.2	-1.6	0
Standard deviation of cash flow, low regime	0	0.1	0.1	1	0.1	0	0.1	-9.5	-0.4
Standard deviation of cash flow, high regime	0.1	0.1	-0.1	0	1	0	0.3	-7.4	-0.1
Standard deviation, cash flow	0	0.1	0.1	0.2	0.7	-0.1	0.1	-9.1	-0.2
Standard deviation, I/K	-0.7	0	0.7	0	0.4	-0.1	-0.1	-2.1	0.5
Skewness I/K	-0.7	0	-0.2	-0.3	0.9	-1	1.6	1	4.8
Serial correlation, q	0.1	0.6	0	0	0	0	0	-16.6	-0.1
Serial correlation, CF/K	0	0.2	0.4	-0.1	-0.3	-0.2	-0.1	-4.5	0.6

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

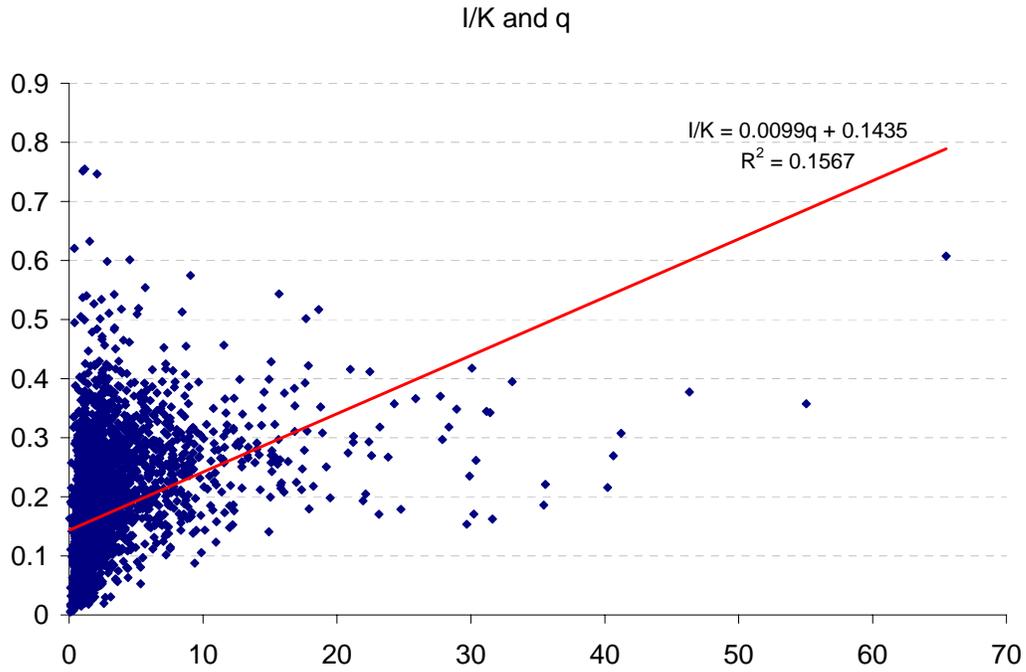


Figure 1b.

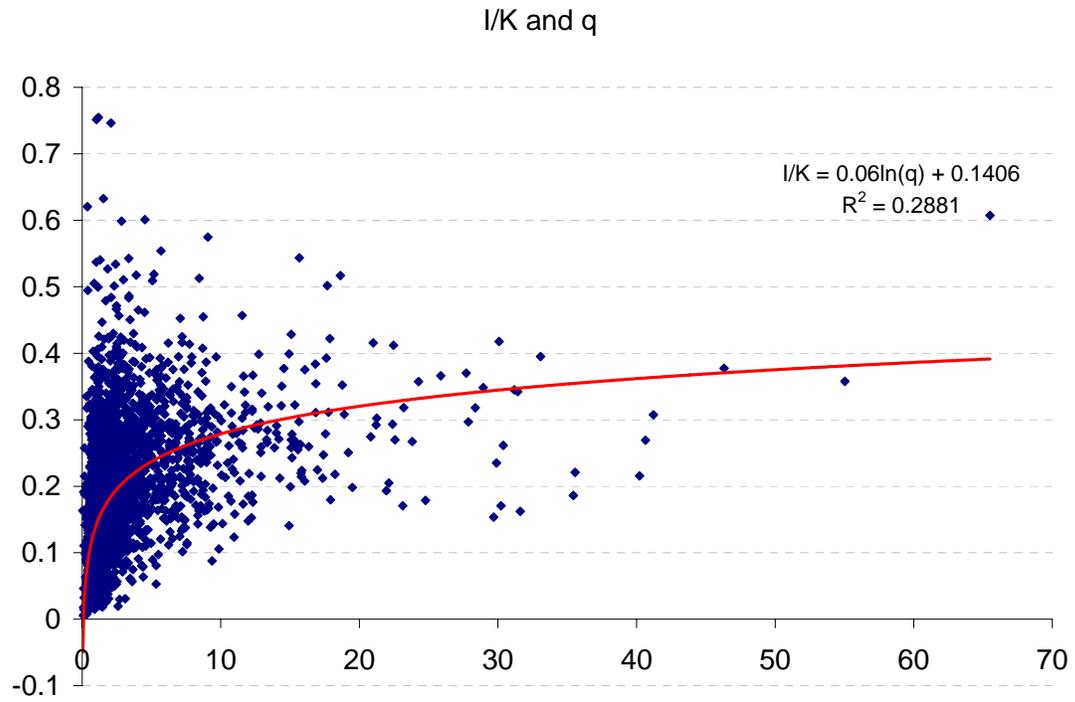


Figure 2a: Investment rate (I/K) versus cash flow (CF/K)

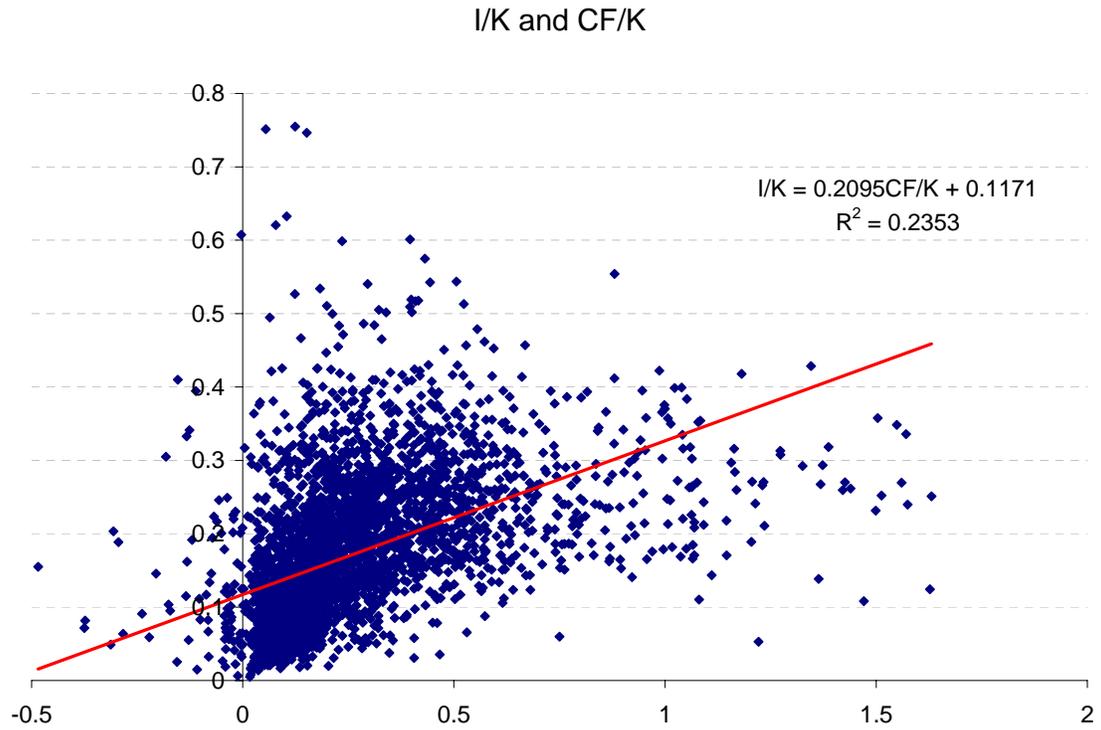


Figure 2b.

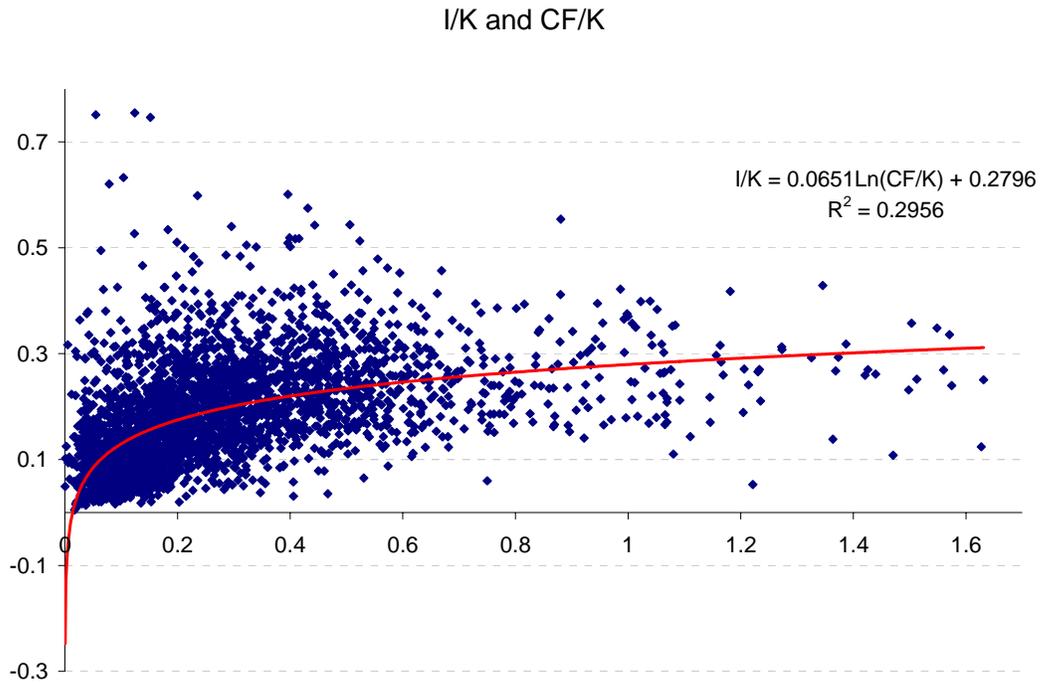


Figure 3: Investment rate $I(t)/K(t)$ versus lagged investment rate $I(t-1)/K(t-1)$

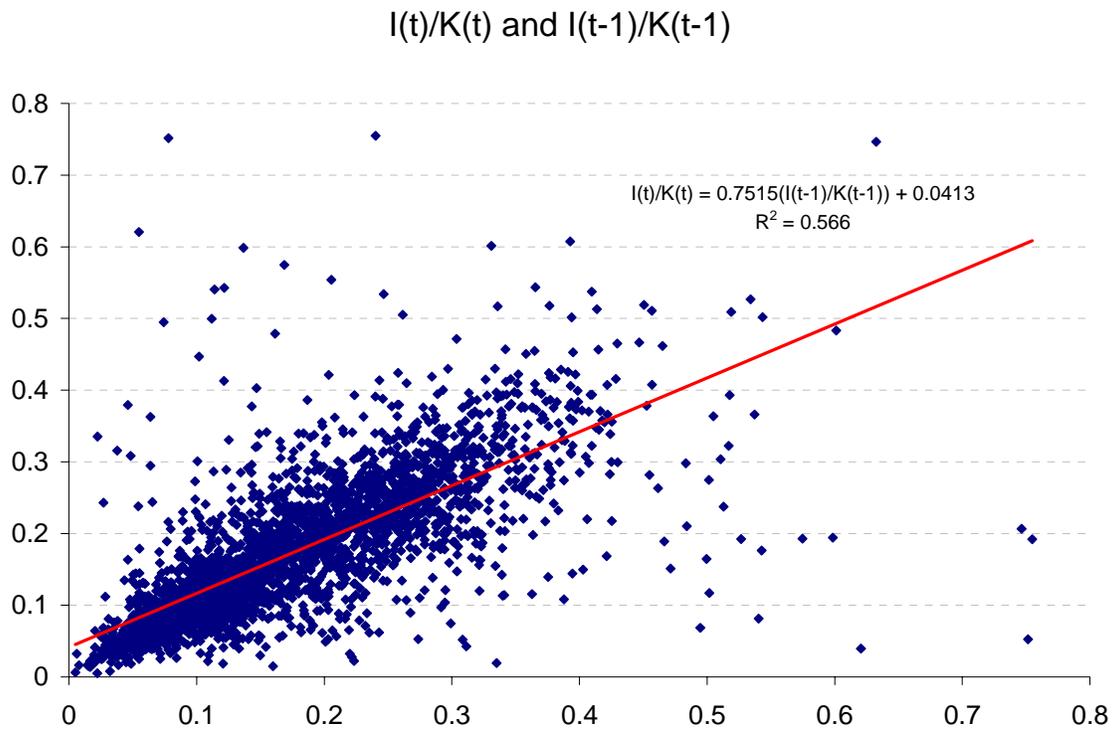


Figure 4: Histograms of I/K, Average Q and CashFlow/K

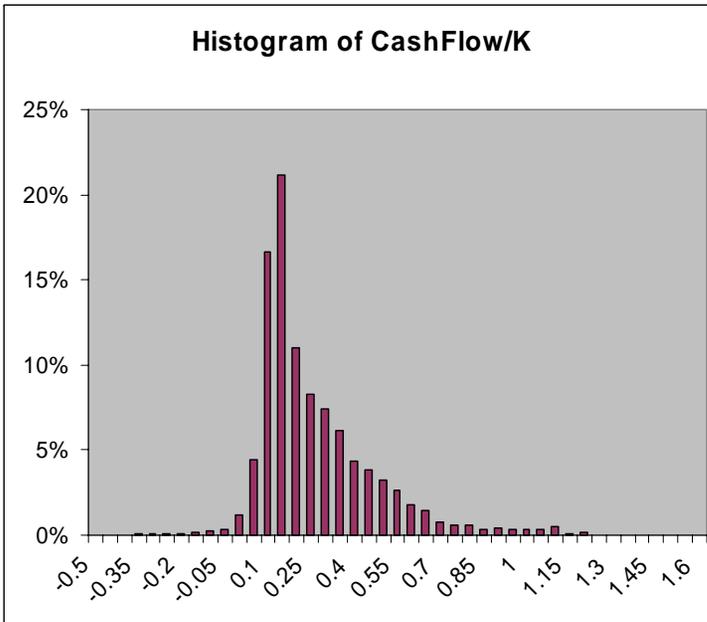
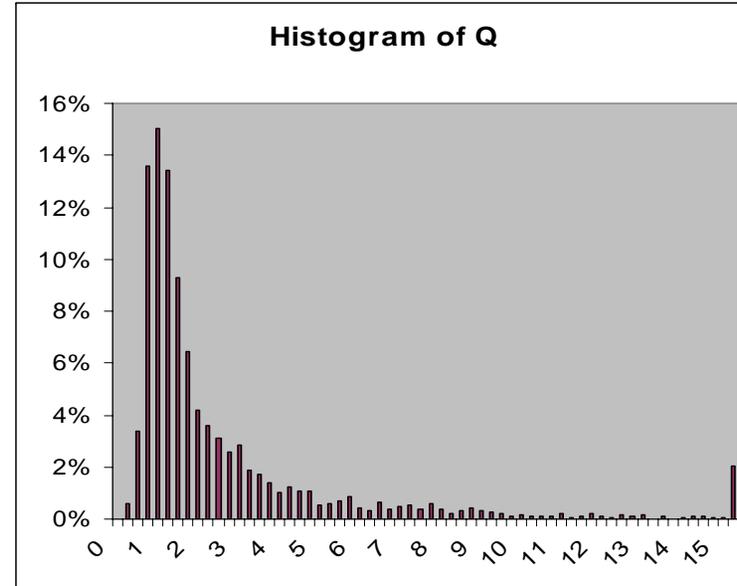
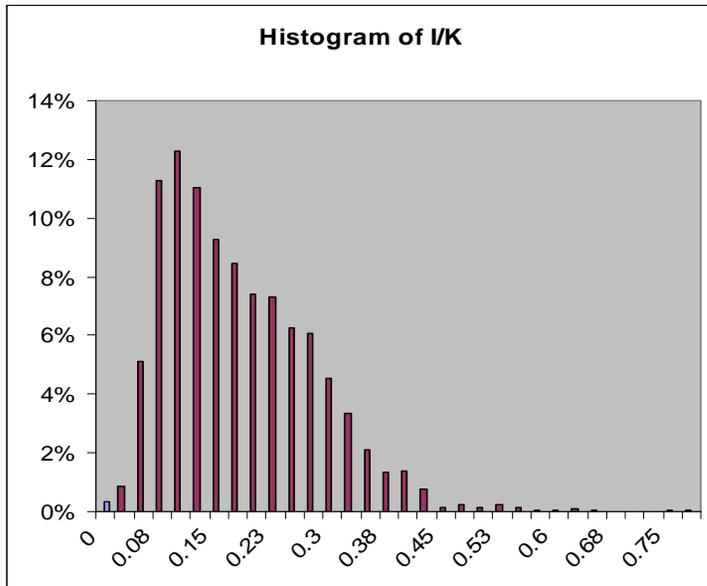


Figure 5: Regime-switching model, estimated distribution of Z shocks

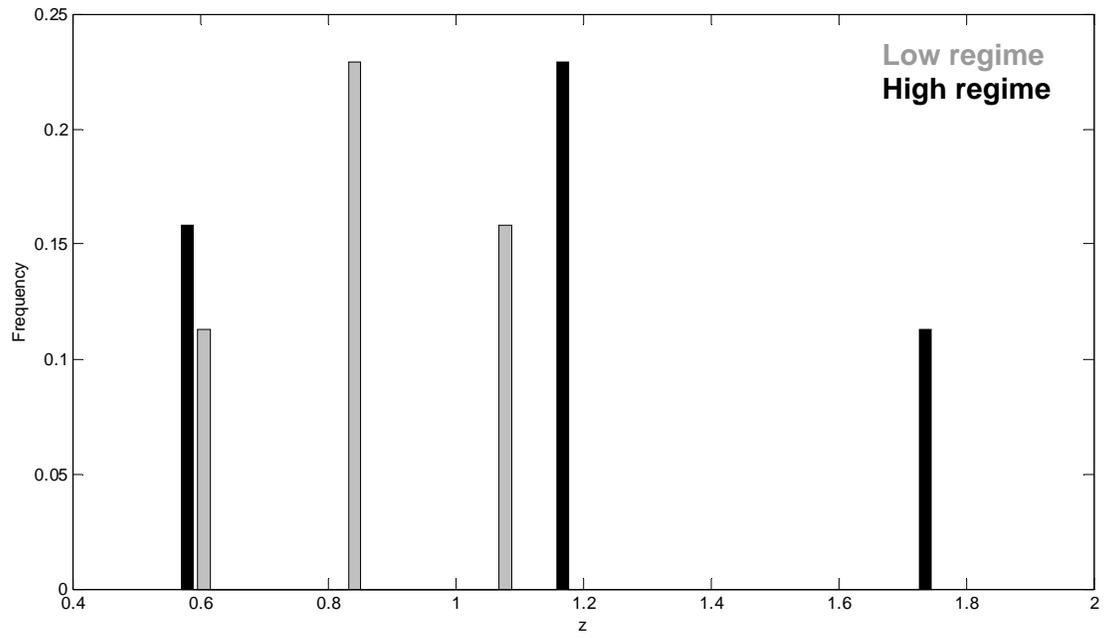


Figure 6a: Regime-switching model, value function by state in each regime

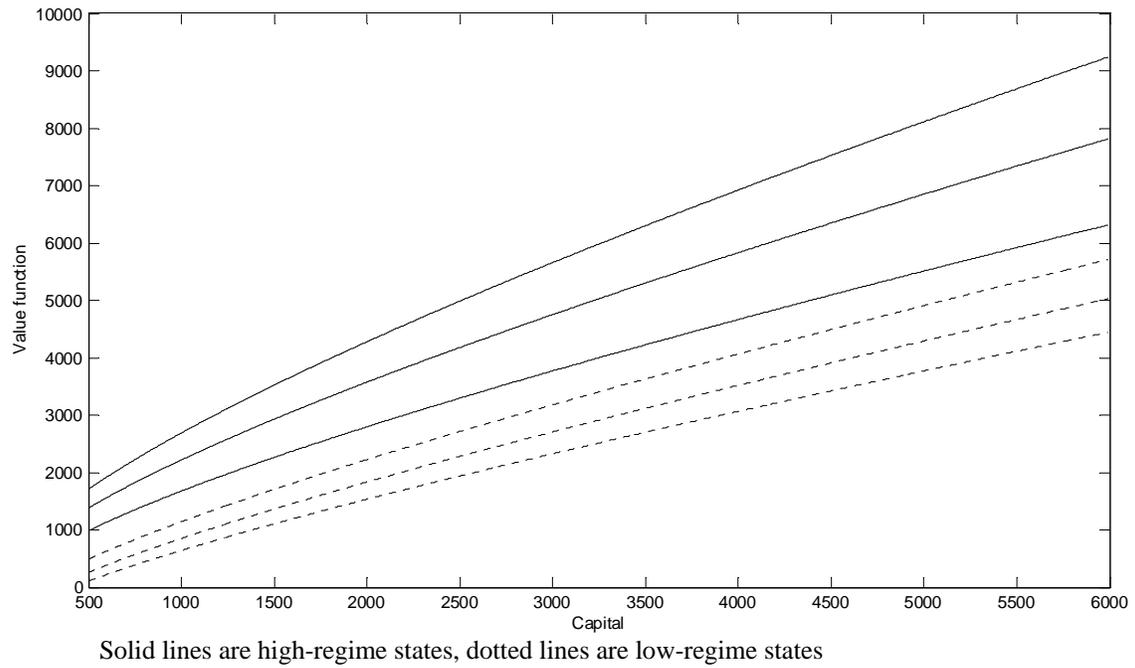


Figure 6b: Regime-switching model, policy function by state in each regime

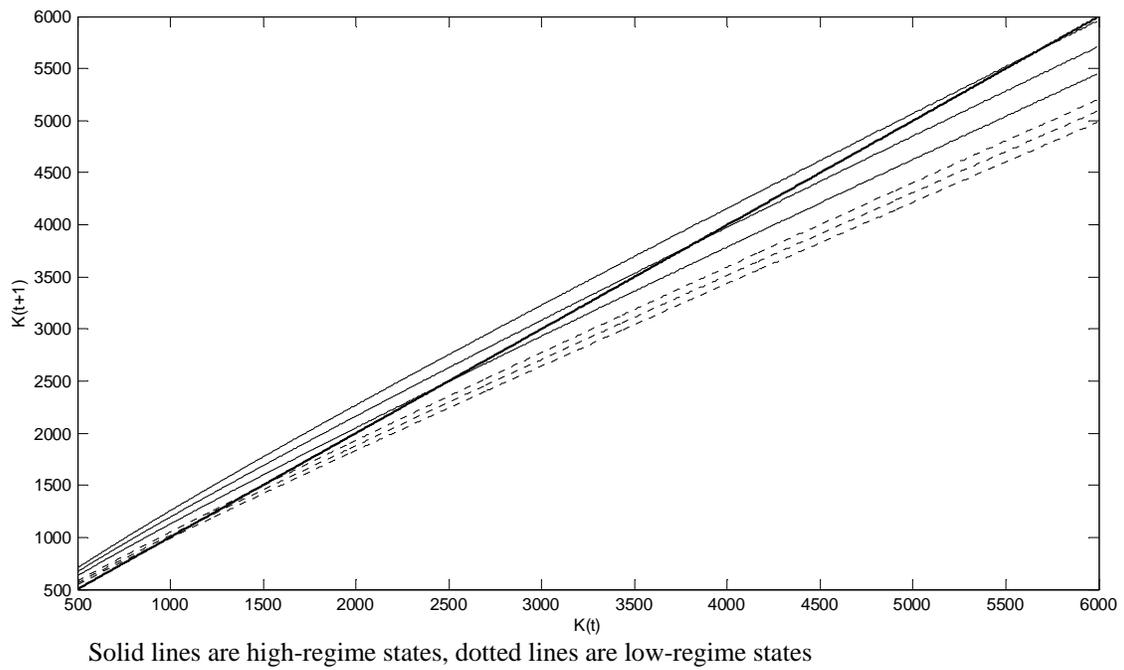
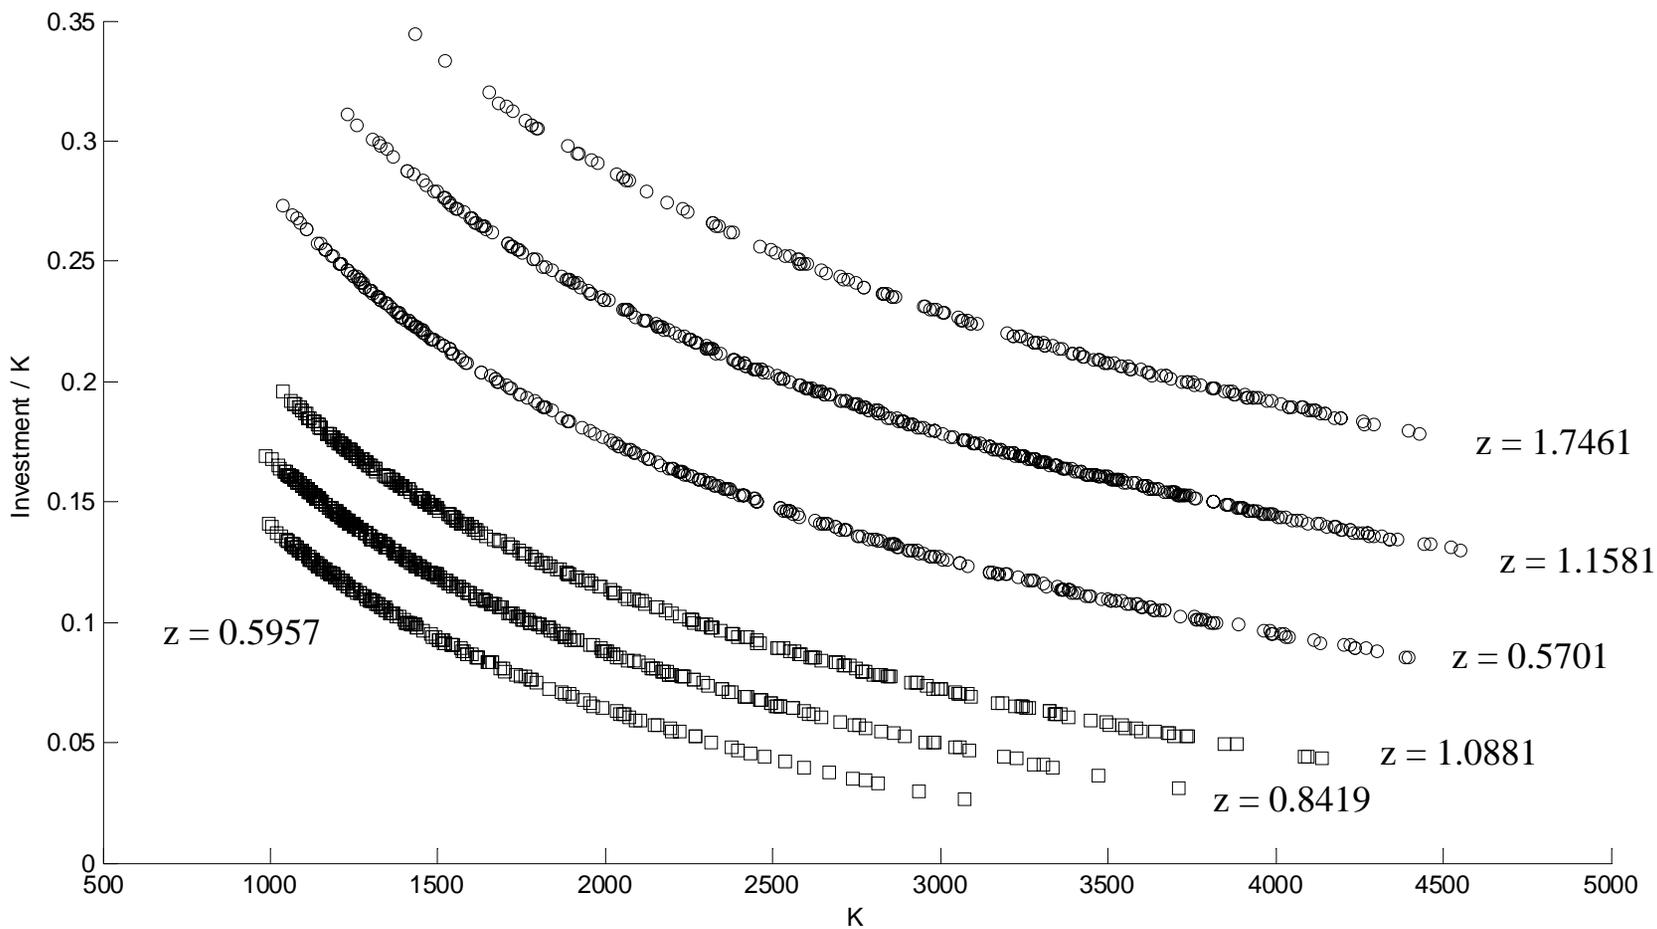


Figure 7: Regime-switching model, investment (I/K) by state in each regime



Circles are high-regime states, squares are low-regime states