6-2012

Technological Heterogeneity and Corporate Investment

Theodosios Dimopoulos
University of Lausanne

Stefano Sacchetto
Carnegie Mellon University, sacchetto@cmu.edu

Technological Heterogeneity and Corporate Investment[*]

Theodosios Dimopoulos[†] **Stefano Sacchetto**[††]

Abstract

We study a dynamic model of corporate investment with fixed and convex capital adjustment costs, and estimate the parameters of the model separately for each firm in a sample of U.S. companies. We evaluate empirically the degree of parameter heterogeneity among firms; quantify the cross-sectional distribution of capital adjustment costs; and assess the magnitude of the estimation bias when one assumes that firms are characterized by a homogeneous set of parameter values. The results show that a considerable amount of parameter heterogeneity exists across firms. Average fixed adjustment costs are 1.15% of the firm's capital, they account for the majority of total adjustment costs, and they are underestimated when assuming parameter homogeneity across firms. Adjustment costs decline with firm size, and convex adjustment costs are positively related to a firm's average merger and acquisition expenditure.

JEL Classifications: D21, D92, E22, L11

Keywords: Firm Heterogeneity, Capital Adjustment Costs

[*†] Swiss Finance Institute, and University of Lausanne. Address: Institut de banque et finance (IBF), HEC Lausanne, Université de Lausanne, Extranef 237, 1015, Lausanne, Switzerland. Phone: +1-41-21-692-3398. Email: Theodosios.Dimopoulos@unil.ch.

[††] Tepper School of Business, Carnegie Mellon University. Address: Tepper School of Business, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA. Phone:+1-412-268-8832. E-mail: sacchetto@cmu.edu.

First draft, please do not cite without authors' permission. All remaining errors are our responsibility.

1. Introduction

There is a growing body of research that studies and estimates dynamic models of corporate investment and financing choices.[1] The aim is to infer from the data the value of economic variables that are not directly observable by researchers, such as the costs of external financing (Hennessy and Whited, 2007), or capital adjustment costs (Cooper and Haltiwanger, 2006).

The approach followed by most papers in this literature is to parameterize an intertemporal model of firm-level investment and financing decisions, and estimate the parameters by matching a set of simulated moments from the model to their empirical counterparts. This estimation procedure neglects firm heterogeneity, because it implicitly assumes that all firms in the sample are described by the same set of parameters. This can lead to biased estimates because, typically, the investment models considered are non-linear, and the parameters that characterize an average, representative firm are different from the average parameters across the firms in the sample.[2]

Two sources of firm heterogeneity can give rise to this bias. The first, which is the main focus of our paper, is represented by differences in production technologies that determine how invested capital translates into cash flows. A second source is the presence of financing frictions that determine the choice between internal and external financing, the capital structure mix, and the allocation of the cash flows among the firm's security holders (Glover, 2011, and Morellec, Nikolov, and Schurhoff, 2012).

In this paper, we study a dynamic model of investment with capital adjustment costs, and estimate the parameters of the model separately for each firm in a sample of U.S. companies. This allows us to evaluate empirically the degree of technological heterogeneity among firms; quantify the cross-sectional distribution of capital adjustment costs; and assess the magnitude of the bias when one assumes homogeneity across firms.

As a basis for estimation, we employ a neoclassical model of investment that has become standard in the literature (e.g., Cooper and Haltiwanger, 2006, and Riddick and Whited,

[1] See Strebulaev and Whited (2012) for a recent survey of this literature.

[2] The issue of neglected persistent firm heterogeneity is well recognized in the literature. For example, Welch (2011) argues that "common problems that plague empirical research in corporate finance are strong firm-size effects that are not fully understood, [and] residual heterogeneity across firms and industries."

2009). The model is in discrete time, and the horizon is infinite. Firms are risk neutral and maximize the present value of future cash flows. In each period, cash flows are determined by operating profits, which are affected by firm-specific persistent productivity shocks, the value of investment, and capital adjustment costs. The latter are incurred whenever a firm acquires or sells capital, which depreciates over time. The specification of the adjustment cost function includes both a fixed and a convex component in the firm's investment-to-capital ratio.

In the model, firms are heterogeneous with respect to six parameters: the depreciation rate of capital, the persistence and volatility of the productivity shock, a profit curvature parameter that affects the marginal returns to capital, and two parameters that describe the cost adjustment function and measure, respectively, fixed and convex adjustment costs.

We estimate a separate set of parameters for each firm in a sample of 1,068 public U.S. companies in the period 1972 to 2006. Our source of data is the Compustat database. To facilitate estimation, we restrict attention to firms with at least 20 years of consecutive observations. Estimation is based on the Simulated Method of Moments (SMM) (McFadden, 1989, and Pakes and Pollard, 1995). This procedure requires to minimize a weighted distance between a set of simulated moments that are derived from the solution of the model, and their empirical counterparts in the data.

For each firm, the time-series moments that we choose to match are the average Tobin's Q, the persistence of the ratio of operating profits to capital, the variance and skewness of the investment to capital ratio, and the coefficients of an OLS regression of the investment ratio on a constant, Tobin's Q, and the firm profitability ratio. These moments are chosen on the basis of their informativeness about the structural parameters. For example, the profit curvature parameter affects negatively the value of average Tobin's Q, and positively the correlation of investment with Tobin's Q and profitability. The persistence and volatility of the productivity process are positively related to the autocorrelation of profitability and to the variance of investment, respectively. A high value of the parameter governing convex adjustment costs implies less variable investment, whereas a higher fixed adjustment cost parameter means less frequent but more aggressive investment bursts and, therefore, higher investment skewness.

The estimation results show that a considerable amount of parameter heterogeneity exists across firms. To quantify potential biases, we compare the cross-sectional averages of our estimates to the values obtained by assuming that all firms are characterized by a homogeneous set of parameters. The curvature of the profit function and the persistence of idiosyncratic productivity are estimated with a relatively low amount of bias. The distribution of the curvature parameter is skewed to the left, and its minimum estimated value, 0.37, is well below the cross-sectional average of 0.83. The distribution of the variance of the productivity shocks exhibits fat tails and its average lies at a considerably lower level than the estimate obtained by neglecting heterogeneity. This bias obtains because of the non-linear way in which the variance of investment and the variance of productivity shocks are related.

Using the parameter estimates, for each firm we compute by Monte Carlo simulation the expected total adjustment costs, and evaluate the relative importance of the fixed and convex components. Average fixed adjustment costs are 1.15% of the firm's capital, and they account for the majority of total adjustment costs (1.4%). Moreover, the cross section of firms' fixed adjustment costs is characterized by a bimodal shape: for 31% of firms, these costs assume a value close to zero, whereas for approximately 42% of firms, they range between 1.5% and 3%. Finally, we find that fixed adjustment costs are underestimated when assuming parameter homogeneity across firms, in which case they are found to be 0.22% of capital.

Next, we analyze the determinants of the cross-sectional dispersion of adjustment costs. We find that, although these costs exhibit clustering at the sectoral level, 85% of their variation is driven by idiosyncratic firm characteristics. In other words, biases due to heterogeneity are bound to persist when one estimates the model assuming that firms within the same sector have homogeneous parameters. To analyze within sector variation of adjustment costs, we regress their average simulated value on firm-specific characteristics. This analysis shows that adjustment costs decline as firm size, measured by sales, increases. To put this relationship into perspective, fixed adjustment costs are on average 1.2% of capital value for firms in the bottom size decile and 0.8% for firms in the top decile. We argue that aggregation of asset-specific fixed adjustment costs within firm is a plausible explanation for this effect. In addition, we find that convex adjustment costs are positively related to the average merger and acquisition expenditure that a firm undertakes. This implies that merger integration costs are an important source of investment frictions at the corporate level. Finally, firms characterized by intensive R&D activity exhibit higher convex and lower fixed adjustment costs, after controlling for sector effects. This can be attributed to high adjustment costs in human capital, in which R&D intensive firms heavily invest.

Our paper is related to several recent contributions in the literature on corporate investment and financing. Cooper and Haltiwanger (2006) develop a model of investment with both fixed and convex adjustment costs and calibrate its parameters using plant-level data from the Longitudinal Research Database.[3] Building on this model, DeAngelo, DeAngelo, and Whited (2011) propose a dynamic model of capital structure and investment to explain the slow speed of adjustment of leverage ratios and the frequent use of debt to finance investment spikes. The same model of corporate investment also serves as a building block in Riddick and Whited (2009) and Eisfeldt and Muir (2012), who study the optimal firm retention policy in an intertemporal setting with costly external finance; and in Nikolov and Whited (2009), who study cash accumulation and corporate investment when the management-investor relationship is affected by agency problems.[4]

A common assumption that these papers make for estimation is that firms are characterized by the same set of parameter values.[5] In this sense, firms are ex-ante homogeneous, and heterogeneity is generated over time only by the realization of firm-specific productivity shocks. We find that, by neglecting persistent firm heterogeneity, not only are parameters estimated with a bias, due to non-linearities in the optimal investment policy, but this bias can vary systematically across firms. For instance, we find that fixed adjustment costs are underestimated more severely for smaller rather than larger firms.

Notable exceptions that account for heterogeneity in firm parameters are the papers by Morellec, Nikolov, and Schurhoff (2012), who use Simulated Maximum Likelihood to

[3]Notice that estimates of investment policy obtained by plant level data are not necessarily representative of the parameters characterizing total corporate investment, since the latter involves a degree of aggregation within firm.

[4]The aggregate effects capital adjustment costs have also been the focus of recent macroeconomic models (e.g., Khan and Thomas, 2008).

[5]For example, whereas DeAngelo, DeAngelo, and Whited (2011) use a panel data approach to estimate the profit curvature parameter, in their simulations all firms are assumed to be characterized by the same parameter values.

estimate a dynamic capital structure model with agency costs, and Glover (2011), who estimates the cross-sectional distribution of financial distress costs in a trade-off model of capital structure. Whereas these papers focus on financing policies, to the best of our knowledge, ours is the first paper that estimates at the firm level a structural model of corporate investment with adjustment frictions.

The paper is structured as follows. Section 2 studies a dynamic model of firm investment with capital adjustment costs, describes the data, and provides details of the estimation procedure. In section 3, we discuss the structural parameter estimates and the simulated cross-sectional distribution of adjustment costs. Section 4 concludes.

2. Model and Estimation

In this section, we present the model of corporate investment with fixed and convex capital adjustment costs that forms the basis for estimation. Furthermore, we discuss sample construction and provide definitions to the empirical variables. Finally, we describe our estimation strategy and the choice of the moments to match.

2.1. Model

We study an economy populated by risk-neutral firms that discount future cash flows at rate r. Time is discrete, and the horizon is infinite. Firm j's operating profits are $\pi_j(z_j, k_j) = z_j k_j^{\alpha_j}$, where k is capital, $\alpha_j \in (0,1)$ characterizes the firm-specific decreasing returns to scale, and $z_j \in [\underline{z}_j, \overline{z}_j]$ is a random profitability shock with law of motion given by

$$ln(z_j') = \rho_{z,j} ln(z_j) + \sigma_{\varepsilon,j}\varepsilon_j, \tag{1}$$

where $\rho_{z,j} \in (0,1)$, $\sigma_{\varepsilon,j} \geq 0$ and ε_j follows a truncated standard normal distribution. The prime symbol denotes values one period in the future. The dynamics of capital are determined by the firm's investment choices $i_j = k_j' - (1-\delta_j)k_j$, where $\delta_j \in (0,1)$ denotes the depreciation rate. Investment results in adjustment costs

$$A(k_j, k_j') = \left[c_{0,j} \mathbf{I}(k_j' - (1-\delta_j)k_j \neq 0) + \frac{c_{1,j}}{2} \left(\frac{k_j' - (1-\delta_j)k_j}{k_j} \right)^2 \right] k_j, \tag{2}$$

where $\mathbf{I}(.)$ is an indicator function. The specification of adjustment costs in (2) is standard in the literature, as it accounts for investment irreversibility in a parsimonious way (Cooper and Haltiwanger, 2006; Riddick and Whited, 2009; DeAngelo, DeAngelo, and Whited, 2001; and Eisfeldt and Muir, 2012). The adjustment cost function contains two terms. The first represents a fixed cost that the firm incurs whenever investment is non-zero, e.g., costs of organizational restructuring within the firm, or integration costs in the case of a merger. Such fixed costs give rise in equilibrium to lumpy investment behavior: periods of high investment are followed by periods of inactivity. The second term in (2)is a convex adjustment cost, which generates investment smoothing over time. Examples are inventory adjustment costs, machine set-up costs, or overtime costs (Holt, Modigliani, Muth, and Simon, 1960; Hamermesh and Pfann, 1996). Scaling the adjustment cost function by the capital level implies that lumpy investment and investment smoothing can take place irrespective of firm size.

In each period, firm j's cash flows are

$$e\left(z_j, k_j, k_j'\right) \equiv \pi_j(z_j, k_j) - i_j - A(k_j, k_j'), \tag{3}$$

with a negative value of e corresponding to external financing proceed. The firm's technology is characterized by the structural parameter vector $\theta_j = (\alpha_j, \delta_j, \rho_{z,j}, \sigma_{\varepsilon,j}, c_{0,j}, c_{1,j})$. We assume that θ_j is time invariant. Therefore, unlike most prior papers in the literature, firm heterogeneity in the model results not only from transitory productivity shocks, but also from permanent differences in investment policies across firms.

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

$$V(z_j, k_j; \theta_j) = \max_{k_j' \geq 0} e_j(z_j, k_j, k_j') + \frac{1}{1+r} \int_{\underline{z}_j}^{\overline{z}_j} V(z_j', k_j'; \theta_j) dF(z_j' | z_j; \rho_{z,j}, \sigma_{\varepsilon,j}), \tag{4}$$

where $F(.|z_j;\rho_{z,j},\sigma_{\varepsilon,j})$ is the productivity shock's transition c.d.f. implied by equation (1).[6] We solve the Bellman equation using value function iteration. To that end, we discretize the state space for both the idiosyncratic productivity shock and the level of capital. Since firms' technologies are heterogeneous, the ergodic sets for z and k differ depending on the parameters. To compute expectations over future productivity shocks, we follow Tauchen (1986). We construct an equally spaced grid with $n_z = 25$ points for $log(z_j)$ that spans eight standard deviations of the ergodic distribution, that is $log(z_j) \in \left[-\frac{4\sigma_{\varepsilon,j}}{\sqrt{1-\rho_{z,j}^2}}, \frac{4\sigma_{\varepsilon,j}}{\sqrt{1-\rho_{z,j}^2}}\right]$. It can be shown that the level of capital never exceeds a firm-specific upper bound $\overline{k}_j = \left(\frac{\overline{z}_j}{\delta_j}\right)^{\frac{1}{1-\alpha_j}}$, where $\overline{z}_j = \exp\left(\frac{4\sigma_{\varepsilon,j}}{\sqrt{1-\rho_{z,j}^2}}\right)$. The grid of capital is then constructed as $k_j \in \left[\overline{k}_j(1-\delta_j)^{\frac{n_k(j)-1}{2}}, ..., \overline{k}_j(1-\delta_j)^{1/2}, \overline{k}_j\right]$. We set the number of points of the capital grid for firm j, $n_k(j)$, so that the left endpoint of the grid lies below 0.01, with a minimum $n_k(j)$ equal to 100. We do so in order to allow the firm to choose a sufficiently low capital level to implement almost complete disinvestment.

2.2. Data

Our source of data is the Compustat Industrial Annual database. To select our sample, we choose a set of filters that are common in the corporate finance literature.[7] We start from the full data sample between 1972 and 2006, and we delete all firms that have a primary Standard Industrial Classification (SIC) code between 4900 and 4999 (regulated firms), 6000 and 6999 (financial firms), or greater than 9000 (quasi-public firms). We also delete firm-year observations with missing values and observations with negative values for total assets, gross property plant and equipment, or sales.[8] We then choose, for each firm, the longest consecutive time series of data and drop firms with less than twenty observations. Our final sample consists of 29,895 yearly observations for 1,068 firms, with an average of about 28 observations per firm.

[6] Our model satisfies assumptions 9.4-9.7 in Stokey, Lucas, and Prescott (1989) and therefore the value function exists and it is unique.

[7] See, for example, Riddick and Whited (2009) and DeAngelo, DeAngelo, and Whited (2011).

[8] We replace with zeros the missing values for deferred taxes (item 74), sale of property, plant and equipment (item 107), and acquisitions (item 129). According to Frank and Goyal (2003), Compustat records these variables "as missing when a firm does not report a particular item or combines it with other data items."

Variable definitions are as follows: total assets are item 6 in Compustat; sales are item 12; gross property, plant and equipment (PPE) is item 7; the investment ratio is the difference between PPE expenditure (item 30) and sale of PPE (item 107), divided by gross PPE; profitability is operating income before depreciation (item 13) divided by total assets; the depreciation rate is item 14 divided by PPE; the numerator of Tobin's Q is the sum of the market value of equity (computed as the product of item 199 – the stock price at the fiscal-year end – and the number of common shares outstanding, item 25) and total assets minus the book value of equity (item 60 plus item 74, deferred taxes), the denominator is total assets; the R&D ratio is item 46 divided by sales; acquisition expenditure is item 129. Variables are deflated to constant 2005 dollars using the GDP deflator.[9] Finally, to reduce the effect of outliers when computing moments, we winsorize the variables at the top and bottom 1%.

2.3. Estimation Procedure

To assess the degree of technological heterogeneity in the cross-section of firms, we estimate the set of structural parameters θ_j for each company in our dataset, $j = 1, \ldots, 1068$, using the SMM procedure introduced by McFadden (1989) and Pakes and Pollard (1989). Denote by N_j the number of observations in the time-series data sample for firm j, by S the number of simulated samples, by η_N the vector of moments estimated using the real data, and by $\eta_{NS}(\theta)$ the vector of simulated moments, which is a function of the vector of structural parameters, θ_j. The structural parameter estimator is given by

$$\widehat{\theta_j} = \underset{\theta \in \mathbf{\Theta}}{\arg\min} \left(\eta_{NS}\left(\theta\right) - \eta_N\right)' \widehat{\mathbf{W}}_N \left(\eta_{NS}\left(\theta\right) - \eta_N\right), \tag{5}$$

where $\widehat{\mathbf{W}}_N$ is an optimally chosen weighting matrix equal to $\frac{1}{N_j}\widehat{\mathbf{\Sigma}}^{-1}$, with $\mathbf{\Sigma} = Var\left(\eta_N\right)$. The estimate $\widehat{\mathbf{\Sigma}}$ is based on actual data and is obtained by bootstrapping the moments in η_N.

At a firm-specific value of the structural parameters, our algorithm simulates a sample of $S = 1000$ firms for $T = 200$ years. This procedure involves drawing S independent

[9]Source: Bureau of Economic Analysis (www.bea.gov), NIPA Table 1.1.9.

time series of z_j in $[\underline{z}_j, \overline{z}_j]$, according to the transition law (equation (1)). We use linear interpolation to find the value function and the optimal next-period capital values between grid points. For each of the S simulated firms, we compute a set of moments and then average their value. The moments are computed only on the basis of the last 20 periods of the simulation, to allow the states of simulated firms to converge to their steady state distribution. We then minimize the SMM criterion function in equation (5), and repeat the procedure across every firm in our dataset. Following Strebulaev and Whited (2012), we use simulated annealing to avoid local optima, followed by the Nelder-Meade algorithm to accelerate convergence to the global optimum.

2.4. Selection of Moments

We set the depreciation rate δ_j to its average empirical value for each firm, and the interest rate to 0.04, a standard value in the literature. We estimate the remaining parameters using the SMM procedure. To do so, we select moments on the basis of their informativeness about the structural parameters. We choose seven moments to match: average Tobin's Q, defined in the model as V/k; the autocorrelation of firm profitability, defined as π/k; the variance and skewness of the investment-to-assets ratio, i/k; and OLS coefficients of a regression of the investment ratio on Q, the firm profitability ratio and a constant term. Using the regression coefficients as moments serves a dual purpose: it facilitates matching the average investment rate within firm, and it helps identify the effects of time varying productivity shocks on investment.

Each structural parameter affects essentially all moments in a non-linear way. To gain intuition, however, about what guides estimation, we highlight some key relationships between parameters and moments. A higher curvature parameter α for the profit function leads to higher investment rate and a decline in Tobin's Q. At the same time, investment responds more strongly to shocks, and its correlation with Tobin's Q and profitability increases. The autocorrelation ρ_z and the variance σ_ε^2 of the productivity shock are positively related to the autocorrelation of operating profits and the variance of investment, respectively. Fixed adjustment costs make investment less responsive to productivity shocks, leading to periods of investment inactivity followed by investment bursts. In that sense, adjustment costs generate skewness in the investment ratio, and reduce Tobin's Q. Unlike the fixed cost parameter c_0, the quadratic cost parameter c_1, reduces the variance of investment and drives down its skewness.

To illustrate the degree of firm heterogeneity in our sample, we provide kernel density estimates of firm-specific empirical moments in Figure 1, and summary statistics in Table 1. Clearly, all moments exhibit substantial cross-sectional dispersion. The average firm-level depreciation rate, which is arguably the most important determinant of the average investment rate in structural corporate finance models (see Strebulaev and Whited, 2012), varies considerably, ranging from 3.8% to 21.4%. Both average Tobin's Q and the variance of investment exhibit positive skewness. This provides preliminary indication that the cross-sectional distribution of firms may reflect a mixture between a companies characterized by low fixed capital adjustment costs, and a group of firms with substantially higher costs. Skewness in investment ratios, a moment that is often used to identify fixed adjustment costs (e.g., DeAngelo, DeAngelo, and Whited, 2011), is on average positive, but again varies considerably among firms. Absent adjustment costs, the model implies a positive relationship between investment with Q and firm profitability. Interestingly, the sensitivities of investment on Q and profitability exhibit varying signs across firms, again an indication of heterogeneity in adjustment costs.

How does investment policy vary across firms, and how does this variation affect the estimates of the structural parameters in our model? Table 2 shows that firms with faster depreciation exhibit higher Q, average investment, and variance of investment. These effects are consistent with the model's predictions: depreciation affects positively the average investment rate, as firms need to replace capital at a higher rate. In addition, when fixed adjustment costs are present, depreciation makes investment more volatile because it leads to more frequent investment bursts. Average Q and the intercept in the investment regression are positively correlated, which can be explained by the positive relationship of each of these variables with the productivity shock. Interestingly, the skewness and the variance of the investment ratio are cross-sectionally positively correlated. This fact can be caused by fixed adjustment costs, which make investment more skewed due to inactivity and less sensitive to productivity shocks. In turn, this can justify the observed negative correlation between investment skewness and the sensitivity of investment to profitability.

3. Estimation Results

In this section, we present the results of the estimation, discuss the cross-section of structural parameter estimates, and characterize the distribution of capital adjustment costs.

3.1. The Cross-Sectional Distribution of Parameter Estimates

Figure 2 plots the cross-sectional densities of the estimated firm-level parameters. To facilitate a comparison between the two cases of homogeneous and heterogeneous parameters, we also estimate the model assuming that firms are characterized by the same parameter values. To do so, we average the firm specific moments across all firms, estimate their covariance matrix by bootstrap, and repeat the SMM procedure. The resulting set of parameters are shown in Panel A of Table 3. Panel B reports summary statistics of the cross-sectional distribution of parameter estimates, and their cross-correlations are shown in Table 4.

The curvature parameter of the profit function, α, averages 0.83 across all firms, and is close to the 0.89 value obtained by assuming firm homogeneity. In the cross section, however, α ranges from as low as 0.37 to as high as 0.96, exhibiting left skewness. Similarly, we find little bias in the estimation of the productivity shock autocorrelation, ρ_z, assuming homogeneity (0.54), when compared to the average in the cross section (0.58). The same does not hold true for the standard deviation of the innovation of the productivity shock. The cross-sectional average of σ_ε, 0.2, is considerably lower than the estimate obtained by assuming the same parameter values across firms. This bias can be explained by the non-linear relationship between the variance of investment and the variance in the innovation process.[10] For the majority of firms, σ_ε lies between 0.15 and 0.3, but the distribution exhibits fat tails with a kurtosis of 14.48 (Figure 2c).

[10] To illustrate this point, assume away adjustment costs. Then, it can be shown that the optimal capital choice, given by $k' = \left[\frac{z^\rho exp\left(\frac{1}{2}\sigma_\varepsilon^2\right)}{r+\delta}\right]^{\frac{1}{1-\alpha}}$, implies a convex relationship between the variance of the investment ratio $\frac{k'-(1-\delta)k}{k}$ and σ_ε^2. By Jensen's inequality, it follows that the value of σ_ε^2 inferred from the average cross-sectional moment σ_I^2 overestimates the cross-sectional mean of σ_ε^2.

The estimated cross-sectional distribution of the fixed adjustment cost parameter c_0 is bimodal (Figure 2d). While for a 31% of firms c_0 does not exceed 0.005, for another 42% of firms, the parameter lies between 0.015 and 0.03, a considerably higher level than the value obtained by neglecting firm heterogeneity (0.0003). Finally, the quadratic cost parameter c_1 exhibits right skewness, and its average cross-sectional value lies below the estimated c_1 for the representative firm.

In the cross section, the parameter estimates co-vary with firm-level moments in line with the discussion of subsection 2.4. As Table 5 shows, firms with higher curvature parameter in the profit function (i.e., higher α), exhibit lower Q, more volatile investment, and higher sensitivity of investment to Q and to firm profitability. More persistent productivity shocks are associated with a higher autocorrelation in firm profitability, higher average Q, and investment that is more responsive to shocks. The latter effect is reflected in the time-series variance of investment, and its conditional correlation with Tobin's Q and firm profitability. The standard deviation of the productivity innovations is negatively correlated with the persistence of profitability, as well as with the investment-profitability sensitivity. Firms with higher fixed adjustment costs are associated with higher variance and skewness of investment. According to the model, this is explained by the longer periods of inactivity and the more aggressive bursts of investment that higher fixed adjustment costs imply. These bursts are much less pronounced, when adjustment costs are of a convex (quadratic) nature, which justifies their negative cross-sectional correlation with the variance and the skewness of investment.

3.2. The Determinants of the Adjustment Costs

At the firm level, the investment policy depends to a large extent on the magnitude of the capital adjustment costs. In this section, we describe the cross-sectional variation of adjustment costs and we relate their size to firm and industry characteristics. Given the estimated parameters, we simulate the cross-section of firms to find the ratios of adjustment cost to capital and adjustment costs to firm value. We do so under two hypothesis. First, we assume that all firms are characterized by the same parameters, as shown in Panel A of

Table 3. Second, we allow firms to be heterogeneous.

The results in Table 6 reveal a number of interesting patterns. Adjustment costs in the cross section average 1.4% of capital value and 1.0% of firm value, and they are considerably higher than the 0.2% value obtained assuming homogeneity. What explains the bias in the estimates using the homogeneous-parameters assumption? As Figure 1 shows, the average skewness of investment is low. At the same time, the distribution of the variance of investment exhibits a long right tail. The estimate of the fixed adjustment cost parameter c_0 responds non-linearly to these moments. Low values of skewness and variance attract the parameter towards the lower bound at zero, whereas high skewness and variance lead to substantial increases in the c_0 estimate. These two facts, a bounded support of c_0 and skewness in the cross-sectional distribution of the variance of investment, combine to yield an underestimation of the adjustment costs when one assumes a population of firms with homogeneous parameters.

Our estimates show that fixed adjustment costs, as a fraction of capital, are four times as large as convex (quadratic) costs. Their magnitude, however, varies considerably from firm to firm, leading to high dispersion in total adjustment costs (Figure 3). What drives this dispersion? A natural conjecture is that technological heterogeneity across firms is clustered at the industry level. In Table 7, we summarize the average adjustment costs for the 49 Fama-French industries.[11] These costs range from less than 1% for the Beer, Gold and Drugs industries to approximately two percent in the Textiles, Fabricated Products and Wholesale sectors. In an unreported ANOVA analysis, we find that, even though between-sector variation in adjustment costs is statistically significant, 85% of the total variation obtains from firms within the same sector.

To analyze their within-sector variation, in Table 8 we regress the simulated adjustment costs, as a fraction of capital, against a number of firm-level variables. A natural candidate variable, according to the existing literature, is firm size. Cooper and Haltiwanger (2006) analyze a large sample of plant-level data and find evidence of fixed and convex adjustment costs. They use a simulated dataset to argue that the role of fixed adjustment costs in the determination of investment declines at the aggregate level. At the corporate level,

[11] Industry classifications are available on Kenneth French's webpage. In our sample, firms operating regulated industries, financial firms, and quasi-public companies are omitted.

intuition suggests that investment irreversibility and fixed costs associated with individual assets should also be smoothed out. This aggregation effect should be more prevalent in larger firms. Since the investment ratio is decreasing in the capital level, we also expect quadratic costs to be negatively correlated with firm size. Our regression results confirm that fixed, quadratic, and total adjustment costs as a fraction of capital are all negatively related to the time-series average of firm size. These effects are statistically significant at the one percent level and economically large: average fixed adjustment costs are 1.15% of capital value for firms at the bottom size decile and 0.8% for firms at the top decile.

To what extent does aggregation due to firm size help explain the bi-modality of the fixed adjustment costs distribution? To answer this question, we perform a logit regression that relates the probability of low average fixed adjustment costs (less than 0.5% of capital), to firm size and other characteristics. If the negative relationship between fixed adjustment costs and firm size takes place for small rather than large firms, we would expect the corresponding regression coefficient to be insignificant. The results of Panel C in Table 8, however, indicate that the probability of a very low c_0 observation increases in firm size.

A second potential source of adjustment frictions in investment are the transaction costs associated with capital reallocation across firms. To proxy for this effect, we collect information on the M&A expenditure-to-assets ratio for each firm in the sample. Total adjustment costs are significantly positively related to M&A costs at the one percent level, either controlling for industry effects or not. Significance obtains only for quadratic costs however, indicating larger costs of integration associated with mega-mergers. In total, adjustment costs average 1.39% of capital value for firms in the bottom decile of M&A activity and 1.56% for firms in the top decile.

A final characteristic that explains the cross sectional variation of adjustment costs is asset tangibility. In particular, we examine how fixed and quadratic adjustment costs correlate with the average R&D to Sales ratio. We find two opposite effects: While quadratic costs increase for R&D-intensive firms, the contrary holds true for fixed adjustment costs. This finding is consistent with the presence of high labor adjustment costs for R&D intensive firms, which are more likely to invest a large fraction of their assets in human capital. Consistent with this conjecture, Holt, Modigliani, Muth, and Simon (1960) find that a quadratic specification approximates well labor adjustment costs.

4. Conclusion

This paper highlights the importance of accounting for persistent cross-sectional firm heterogeneity when estimating structural models in corporate finance. To do so, we consider a large sample of U.S. companies and estimate separately for each firm the parameters of a dynamic model of corporate investment. This allows us to quantify the cross-sectional distribution of firms' capital adjustment costs, which are the source of investment frictions in our model. We find considerable variation in parameter estimates both for firms within the same industry, and across different industries. Because dynamic investment models with frictions are typically highly non-linear, neglecting heterogeneity creates an estimation bias: the parameter estimates for a representative firm are different from the average parameters across firms. In our case, the average adjustment cost in the cross section (1.4% of capital) exceed the corresponding estimate if one assumes that all firms are ex-ante homogeneous (0.2%).

Being the first paper to estimate a structural investment model at the firm-level, we focus on technological rather than financing frictions – such as costs of external financing or bankruptcy costs. An interesting avenue for future research would be to evaluate empirically models that incorporate both types of frictions (e.g., Gomes, 2001; DeAngelo, DeAngelo, and Whited, 2011), to derive their joint distribution in the cross section of firms. Although such models are characterized by a higher number of state variables, recent developments in numerical integration methods (Judd, Maliar, and Maliar, 2011) and the use of parallel computing (Aldrich, Fernandez-Villaverde, Gallant, and Rubio-Ramirez, 2011) alleviate the computational burden required for estimation.

References

Aldrich, Eric M., Jesus Fernandez-Villaverde, A. Ronald Gallant, and Juan F. Rubio-Ramirez, 2001, Tapping the supercomputer under your desk: Solving dynamic equilibrium models with graphics processors, *Journal of Economic Dynamics and Control* 35, 386 – 393.

Cooper, Russell W., and John C. Haltiwanger, 2006, On the nature of capital adjustment costs, *Review of Economic Studies* 73, 611 – 633.

DeAngelo, Harry, Linda DeAngelo, and Toni M. Whited, 2011, Capital structure dynamics and transitory debt, *Journal of Financial Economics* 99, 235–261.

Eisfeldt, Andrea L., and Tyler Muir, 2012, The joint dynamics of internal and external finance, Working Paper.

Frank, Murray Z., and Vidhan K. Goyal, 2003, Testing the pecking order theory of capital structure, *Journal of Financial Economics* 67, 217 – 248.

Glover, Brent, 2011, The expected cost of default, Working Paper.

Gomes, Joao, 2001, Financing investment, *American Economic Review* 91, 1263 – 1285.

Hamermesh, Daniel S., and Gerard A. Pfann, 1996, Adjustment costs in factor demand, *Journal of Economic Literature* 34, 1264 – 1292.

Hennessy, Christopher A., and Toni M. Whited, 2007, How costly is external financing? Evidence from a structural estimation., *Journal of Finance* 62, 1705 – 1745.

Holt, Charles C., Franco Modigliani, John F. Muth, and Herbert A. Simon, 1960, Planning production, inventories and work force, Prentice-Hall, Englewood Cliffs, New Jersey.

Judd, Kenneth L., Lilia Maliar, and Serguei Maliar, 2011, Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models, *Quantitative Economics* 2, 173 – 210.

Khan, Aubhik, and Julia K. Thomas, 2008, Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics, *Econometrica* 76, 395–436.

McFadden, Daniel, 1989, A method of simulated moments for estimation of discrete response models without numerical integration, *Econometrica* 57, 995–1026.

Morellec, Erwan, Boris Nikolov, and Norman Schurhoff, 2012, Corporate governance and capital structure dynamics, *Journal of Finance* 67, 803 – 848.

Nikolov, Boris, and Toni M. Whited, 2009, Agency conflicts and cash: Estimates from a structural model, Working Paper.

Pakes, Ariel, and David Pollard, 1989, Simulation and the asymptotics of optimization estimators, *Econometrica* 57, 1027–1057.

Riddick, Leigh A., and Toni M. Whited, 2009, The corporate propensity to save, *Journal of Finance* 64, 1729 – 1766.

Stokey, Nancy L., Robert E. Lucas, and Edward C. Prescott, 1989, Recursive methods in economic dynamics, Harvard University Press, Cambridge, Massachussets.

Strebulaev, Ilya A., and Toni M. Whited, 2012, Dynamic models and structural estimation in corporate finance, *Foundations and Trends in Finance* forthcoming.

Tauchen, George, 1986, Finite state Markov-chain approximations to univariate and vector autoregressions, *Economics Letters* 20, 177 – 181.

Welch, Ivo, 2011, A critique of recent quantitative and deep-structure modeling in capital structure research and beyond, Working Paper.

Table 1: Summary statistics of firm-specific empirical moments. This table shows summary statistics for the cross section of firm-specific moments. The sample consists of 1,068 publicly traded firms in Compustat. δ is the firm's average depreciation rate. Q the firm's average Tobin's Q ratio. ρ_{oi} is autocorrelation in firm profitability. σ_I^2 is the time series variance of the firm's investment ratio; $Skew(I)$ is the time series skewness of the investment ratio; β_0, β_Q, and β_{oi} are OLS regression coefficients from a regression of the investment ratio on an intercept, Tobin's Q, and firm profitability. The details of the sample construction and variable definitions can be found in subsection 2.2.

	δ	Q	ρ_{oi}	σ_I^2	Skew(I)	β_0	β_Q	β_{oi}
Mean	0.081	1.498	0.620	0.011	0.845	0.065	0.020	0.209
Median	0.075	1.303	0.662	0.006	0.838	0.062	0.007	0.238
Standard Deviation	0.028	0.627	0.209	0.026	1.180	0.117	0.093	0.576
Skewness	1.403	2.043	-0.941	12.421	-0.909	0.100	2.699	-0.444
Min	0.038	0.764	-0.384	0.000	-5.378	-0.703	-0.354	-3.260
Max	0.214	4.987	0.986	0.542	4.426	0.833	0.864	3.534

Table 2: Spearman rank correlation between firm-level moments. This table shows the Spearman rank correlation matrix between firm-specific empirical moments. The sample consists of 1,068 publicly traded firms in Compustat. δ is the firm's average depreciation rate; Q the firm's average Tobin's Q ratio; ρ_{oi} is autocorrelation in firm profitability; σ_I^2 is the time series variance of the firm's investment ratio; $Skew(I)$ is the time series skewness of the investment ratio; β_0, β_Q, and β_{oi} are OLS regression coefficients from a regression of the investment ratio on an intercept, Tobin's Q, and firm profitability. The details of the sample construction and variable definitions can be found in subsection 2.2.

	δ	Q	ρ_{oi}	σ_I^2	Skew(I)	β_0	β_Q	β_{oi}
δ	1							
Q	0.307	1						
ρ_{oi}	-0.027	0.234	1					
σ_I^2	0.344	0.129	-0.163	1				
Skew(I)	-0.004	-0.028	-0.113	0.270	1			
β_0	0.143	0.175	-0.065	-0.032	0.074	1		
β_Q	0.056	-0.091	-0.062	0.230	0.068	-0.547	1	
β_{oi}	-0.013	-0.024	0.177	-0.040	-0.164	-0.443	-0.274	1

Table 3: Structural parameter estimates. This table shows the structural parameter estimates derived using the Simulated Method of Moments. Panel A shows the estimates obtained by assuming that the same parameter values hold across firms, and by matching the mean moment values in Table 1. Panel B shows summary statistics of the estimates obtained by allowing for heterogeneous parameters across firms, and by matching the firm specific moments summarized in Table 1. The sample consists of 1,068 publicly traded firms in Compustat. The details of the sample construction and variable definitions can be found in subsection 2.2. α is the curvature parameter of the profit function; ρ_z is the autocorrelation of the firm specific log-productivity and σ_ε its variance; c_0 is fixed adjustment cost parameter and c_1 is the convex (quadratic) adjustment cost parameter.

Panel A: Structural parameter estimates assuming firms with homogeneous parameters.

	α	ρ_z	σ_ε	c_0	c_1
Coefficient	0.890	0.544	0.420	0.0003	0.395

Panel B: Summary statistics of firm-level structural parameter estimates.

	α	ρ_z	σ_ε	c_0	c_1
Mean	0.831	0.578	0.207	0.015	0.303
Median	0.865	0.571	0.205	0.020	0.260
Standard Deviation	0.102	0.107	0.059	0.010	0.114
Skewness	-1.796	0.581	1.343	-0.436	3.296
Min	0.372	0.228	0.003	0.000	0.071
Max	0.962	0.990	0.662	0.035	1.657

Table 4: Spearman rank correlation of structural parameters in the cross-section. This table shows the Spearman-rank correlation between the structural parameter estimates. The estimates are obtained using the Simulated Method of Moments for a sample of 1,068 firms in Compustat, allowing for heterogeneous parameters across firms, and matching the firm specific moments summarized in Table 1. α is the curvature parameter of the profit function; ρ_z is the autocorrelation of the firm specific log-productivity and σ_ε its variance; c_0 is fixed adjustment cost parameter and c_1 is the convex (quadratic) adjustment cost parameter.

	α	ρ_z	σ_ε	c_0	c_1
α	1				
ρ_z	-0.128	1			
σ_ε	-0.213	-0.314	1		
c_0	0.302	-0.008	-0.091	1	
c_1	-0.386	-0.105	-0.058	-0.805	1

Table 5: Spearman rank correlation between structural parameters and firm moments. This table shows the Spearman rank correlation matrix between the firm-level structural parameter estimates and the firm's empirical moments. The sample consists of 1,068 publicly traded firms in Compustat. The details of the sample construction and variable definitions can be found in subsection 2.2. Q is the firm's average Tobin's Q ratio; ρ_{oi} is autocorrelation in firm profitability; σ_I^2 is the time series variance of the firm's investment ratio; $Skew(I)$ is the time series skewness of the investment ratio; β_0, β_Q, and β_{oi} are OLS regression coefficients from a regression of the investment ratio on an intercept, Tobin's Q, and firm profitability. α is the curvature parameter of the profit function; ρ_z is the autocorrelation of the firm specific log-productivity and σ_ε its variance; c_0 is fixed adjustment cost parameter and c_1 is the convex (quadratic) adjustment cost parameter.

	Q	ρ_{oi}	σ_I^2	Skew(I)	β_0	β_Q	β_{oi}
α	-0.603	-0.129	0.060	-0.048	-0.097	0.062	0.085
ρ_z	0.083	0.184	0.047	-0.106	-0.301	0.151	0.139
σ_ε	0.110	-0.065	0.200	0.156	0.053	0.045	-0.135
c_0	-0.395	-0.227	0.468	0.215	-0.042	0.062	0.059
c_1	0.501	0.189	-0.493	-0.202	0.156	-0.145	-0.070

Table 6: Summary statistics of simulated adjustment costs. This table shows summary statistics for average simulated adjustment costs scaled by capital and firm value. Total adjustment costs are the sum of fixed and quadratic adjustment costs. The sample consists of 1,068 publicly traded firms in Compustat. In Panel A, all firms' parameters are set at the values reported in Table 3, Panel A. In Panel B, firms have heterogeneous parameters and their values are summarized in Panel B of Table 3.

Panel A: Average simulated adjustment costs using the assumption of homogeneous firms.

	Mean
Fixed AC / Capital	0.0003
Quadratic AC / Capital	0.0019
Total AC / Capital	0.0022
Fixed AC / Firm Value	0.0003
Quadratic AC / Firm Value	0.0017
Total AC / Firm Value	0.0021

Panel B: Summary statistics of cross-section of simulated adjustment costs.

	Mean	Median	St.Dev.	Skew.	Max
Fixed AC / Capital	0.012	0.014	0.008	-0.237	0.027
Quadratic AC / Capital	0.003	0.002	0.002	2.283	0.016
Total AC / Capital	0.014	0.017	0.009	-0.234	0.036
Fixed AC / Firm Value	0.008	0.009	0.007	0.228	0.036
Quadratic AC / Firm Value	0.001	0.001	0.001	1.779	0.006
Total AC / Firm Value	0.010	0.011	0.007	0.221	0.041

Table 7: Summary statistics of adjustment costs by industry. This table shows summary statistics at the sector level for simulated total adjustment costs scaled by capital. Adjustment costs are simulated using the firms-specific parameter values summarized in Panel B of Table 3. The sample consists of 1,068 publicly traded firms in Compustat. Firms are grouped into 49 industries according to the Fama-French classification (FF). Regulated firms, financial firms, and quasi-public firms are excluded from the sample. The details of the sample construction can be found in subsection 2.2

FF	Name	Mean	St. Dev.	FF	Name	Mean	St. Dev.
1	Agriculture	0.0137	0.0111	23	Automobiles and Trucks	0.0154	0.0091
2	Food	0.0123	0.0090	24	Aircraft	0.0131	0.0089
3	Soda	0.0156	0.0125	25	Shipbuilding, Rail. Eq.	0.0115	0.0110
4	Beer	0.0049	0.0080	26	Defense	0.0089	0.0059
5	Tobacco	0.0115	0.0097	27	Gold	0.0086	0.0060
6	Toys	0.0124	0.0107	28	Mines	0.0177	0.0080
7	Entertainment	0.0114	0.0058	30	Oil	0.0131	0.0079
8	Books	0.0108	0.0102	32	Telecommunications	0.0099	0.0105
9	Consumer goods	0.0133	0.0091	33	Personal Services	0.0139	0.0092
10	Clothes & Apparel	0.0153	0.0088	34	Business Services	0.0180	0.0083
11	Healthcare	0.0143	0.0086	35	Computer Hardware	0.0138	0.0103
12	Medical Equipment	0.0093	0.0074	36	Computer Software	0.0159	0.0102
13	Drugs	0.0080	0.0071	37	Electronic Eq.	0.0136	0.0083
14	Chemicals	0.0106	0.0081	38	Measuring Eq.	0.0155	0.0082
15	Rubber & Plastics	0.0157	0.0090	39	Business Supplies	0.0130	0.0081
16	Textiles	0.0203	0.0045	40	Shipping Containers	0.0140	0.0100
17	Construction Mat.	0.0136	0.0083	41	Transportation	0.0168	0.0071
18	Construction	0.0169	0.0076	42	Wholesale	0.0191	0.0070
19	Steel	0.0128	0.0082	43	Retail	0.0146	0.0079
20	Fabricated Prod.	0.0197	0.0031	44	Restaurant and Hotels	0.0127	0.0078
21	Machinery	0.0142	0.0084	49	Other	0.0155	0.0105
22	Electrical Eq.	0.0112	0.0094				

Table 8: Regression analysis of adjustment costs. This table shows the cross-sectional determinants of simulated firm adjustment costs scaled by capital. The sample consists of 1,068 publicly traded firms in Compustat. Firm-specific average adjustment costs are computed by simulation, using the parameters summarized in Panel B of Table 3. Panel A shows summary statistics of the independent variables in the regression, which are defined in subsection 2.2. Panel B shows OLS coefficients obtained by regressing firm-specific average adjustment costs on the time-series averages of firm characteristics. $TotAC$, $FixAC$, $QuadAC$ are, respectively, average firm-specific total, fixed, and quadratic adjustment costs. Panel C shows coefficients of a logit regression of a dummy variable that equals 1 when the average fixed adjustment cost is less than 0.5%, on firm-specific characteristics. *, ** and *** indicate significance at the 10%, 5% and 1% level, respectively. Robust standard errors are in parentheses.

Panel A: Summary statistics of independent variables in the regression.

	Mean	Median	St. Dev.	Skewness
log(Sales)	6.0024	5.9574	1.8140	0.1491
Acquisition Exp.	0.0145	0.0148	0.0101	1.4589
R&D/Sales	0.0198	0.0030	0.0369	2.8870

Panel B: OLS Regression analysis of adjustment costs.

	(1)	(2)	(3)	(4)	(5)
Dependent variable	$TotAC$	$TotAC$	$TotAC$	$FixAC$	$QuadAC$
log(Sales)	-0.0010***	-0.0012***	-0.0012***	-0.0010***	-0.0002***
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(2.8E-05)
Acquisition Exp.		0.0389**	0.0548***	0.0270	0.0278***
		(0.0189)	(0.0193)	(0.0183)	(0.0045)
R&D/Sales		-0.0275***	-0.0075	-0.0255***	0.0180***
		(0.0068)	(0.0089)	(0.0087)	(0.0025)
Constant	0.0201***	0.0210***	0.0205***	0.0170***	0.0035***
	(0.0009)	(0.0010)	(0.0032)	(0.0028)	(0.0007)
Industry effects	No	No	Yes	Yes	Yes
R^2	0.0448	0.0630	0.1483	0.1361	0.4027

Panel C: Logit analysis of fixed adjustment costs.

	Logit coeff.
log(Sales)	0.2407***
	(0.0471)
Acquisition Exp.	-1.2865
	(5.2709)
R&D/Sales	1.0336
	(2.5626)
Constant	-2.1014***
	(0.6784)
Industry effects	Yes
Pseudo-R^2	0.0867

Figure 1: Cross-sectional distribution of empirical moments. This figure shows kernel density estimates of the distributions of firm-specific empirical moments. The sample consists of 1,068 publicly traded firms in Compustat. Sample construction and variable definitions are provided in subsection 2.2.

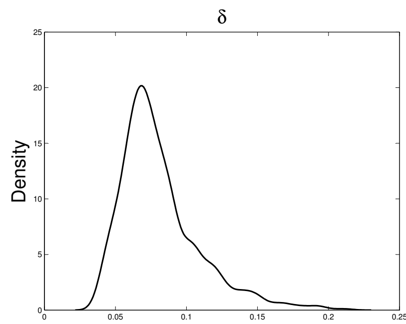

(a) Depreciation rate

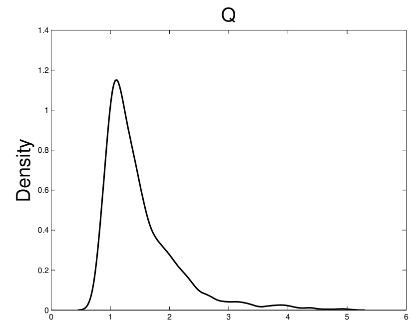

(b) Tobin's Q

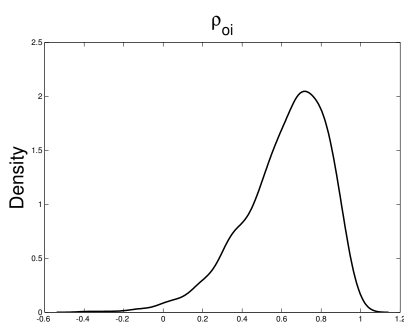

(c) Autocorr. of op. income

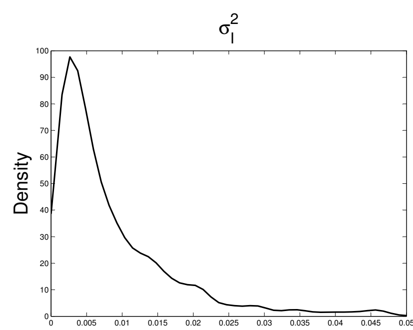

(d) Variance of investment

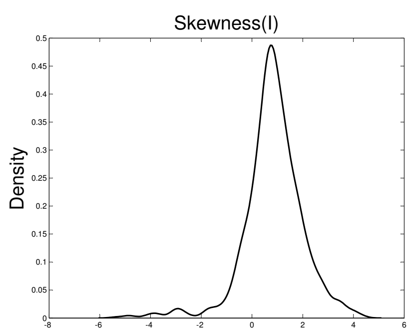

(e) Skewness of investment

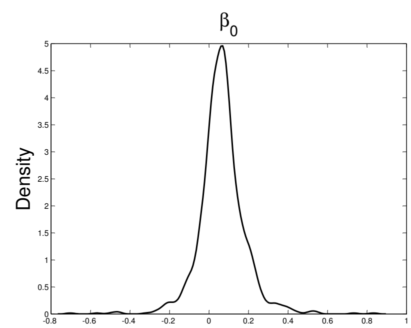

(f) Regression coeff. of constant

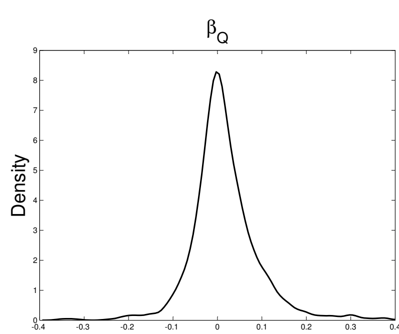

(g) Regression coefficient of Tobin's Q

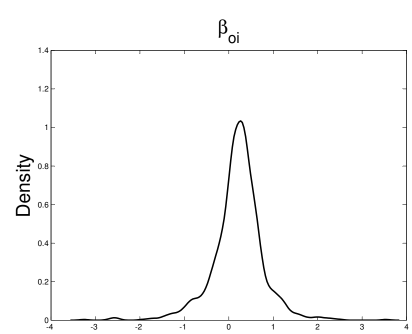

(h) Regression coefficient of operating income

Figure 2: Cross sectional distributions of structural parameters. Estimation is performed using the SMM approach. Summary statistics of the moments to match are in Table 1. The sample consists of 1,068 publicly traded firms in Compustat. Sample construction and variable definitions are provided in subsection 2.2. α is the curvature parameter of the profit function; ρ_z is the autocorrelation of the firm specific log-productivity and σ_ε its variance; c_0 is fixed adjustment cost parameter and c_1 is the convex (quadratic) adjustment cost parameter.

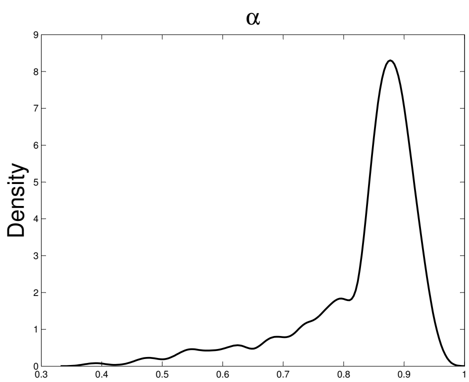

(a) α

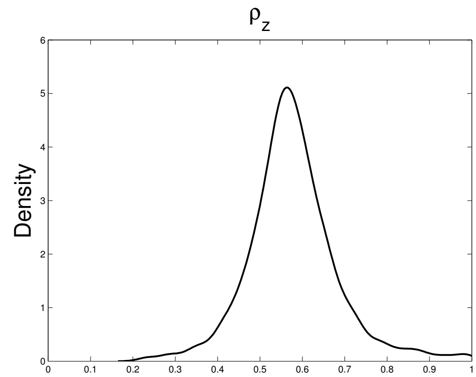

(b) ρ_z

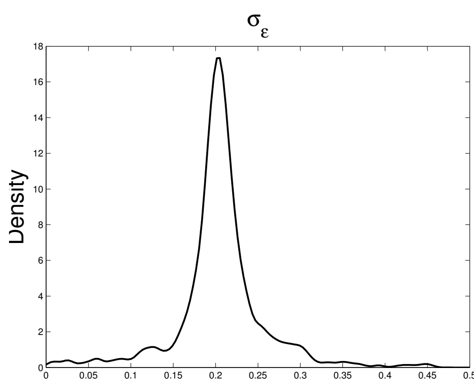

(c) σ_ε

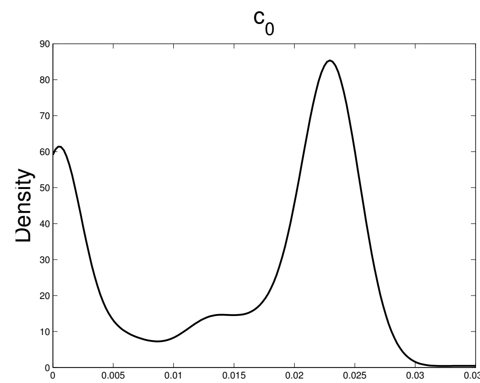

(d) c_0

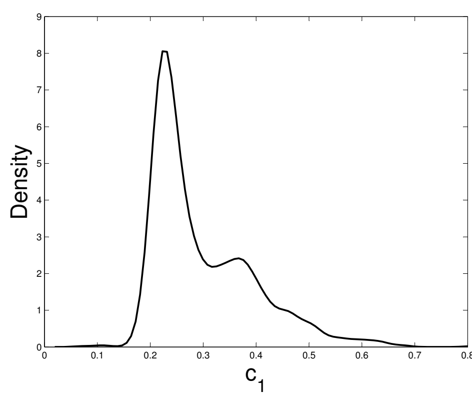

(e) c_1

Figure 3: Distribution of estimated total adjustment costs/capital ratios in the cross-section of firms. The adjustment costs are simulated according to the firm-specific structural parameter estimates obtained using the SMM approach. The sample consists of 1,068 publicly traded firms in Compustat. Sample construction and variable definitions are provided in subsection 2.2.