Essays in financial economics

Golec, P.M.

Chapter 1

Why do firms become less dynamic after they go public? Evidence from structural estimation

1.1 Introduction

In the subsequent years after going public, firms cut investment. This is not all that surprising however. An important reason firms do an initial public offering (IPO) is because they discover major growth opportunities, and therefore want to raise capital to be able to grow. The fact that their investment declines after going public could simply reflect their gradual growth path towards their new efficient scale (Clementi, 2002). With convex adjustment costs, early on they grow faster when their marginal product of capital is high, and later on more slowly when it is low.

Interestingly, over half of the investment decline is unrelated to profitability or size, so it can not only be about firms growing towards their efficient scale. What then is the explanation for the conditional investment decline? While there is an understanding that firms become more rigid (Asker et al., 2014), are learning Pástor et al. (2008) and become less profitable after they go public (Chemmanur et al., 2009), the relative importance of the different factors is yet unknown.

This paper quantifies the contributions of five sources towards the conditional investment decline. In addition to growing towards their efficient scale (GTES), rigidity, learning and declining profitability, I also consider the role of decreasing shock volatility as firms age and interactions between the different sources.

Measuring the relative importance of these mechanisms represents a challenge. Investment is endogenous, and firms' rigidity and their beliefs are hard to measure. Although one can estimate the reasons' directional effects using reduced-form empirical techniques, evaluating their magnitudes requires estimating or calibrating an economic model.

These challenges lend themselves to a structural estimation approach. I present a dynamic firm investment model which focuses on firms' life after going public. The model is a convex adjustment cost framework that I extend across five dimensions. (i) Firms' capital stock at IPO is lower than what is long-term efficient, i.e., they are growing towards their efficient scale (GTES), (ii) they are learning about their efficient scale, (iii) their volatility of productivity declines over time, (iv) their adjustment costs increase over time, i.e., they become more rigid and (v) their productivity is abnormally high at the time of IPO and then reverts to the mean.[1]

I estimate the model using the simulated method of moments (SMM) and Compustat data on U.S. public firms. Although the model contains different sources for the decline in conditional investment, they all generate different predictions for the joint dynamics of investment, Tobin's Q and profitability. These predictions enable the estimation procedure to separate them from one another. The model fits the dynamics well. With the estimated model in hand, I decompose the fitted conditional investment decline into the different sources.

Estimates indicate that no single source alone is responsible for the decline in conditional investment. There are three important mechanisms at work, each contributing roughly one third. Interestingly, one is an interaction of two sources, GTES and increasing rigidity, which explains 31 percent of the decline in conditional investment. The fact that profitability and size explain almost half of the decline in both investment and also Tobin's Q implies that GTES must be important. I also find that firms' adjustment costs increase by a factor of 3.5 during the first several years after going public. In the model they must increase quite substantially to fit the declining sensitivity of investment to profitability. This allows firms to grow faster towards their efficient scale than they could otherwise. To build intuition, it is useful to compare two otherwise identical firms which are at different points in their post-IPO life-cycle. Both have the same marginal product of capital and profitability, but the firm that recently went public will invest more aggressively, due to lower adjustment costs. Thus there is a discrepancy in their investment rates that is not explained by profitability or size.

Learning about the efficient scale accounts for 37 percent of the conditional investment decline. To match the pronounced decline in Q conditional on size and profitability in the data, the option value from learning must be pretty high in the model. Intuitively, firms earlier in their post-IPO life-cycle are more uncertain about their efficient scale, so there is (still) a chance of learning to in fact be highly productive or unproductive. The option value comes from productivity being skewed, so the upside is more valuable. Skewness in productivity implies that firms which revise their beliefs upwards invest more than firms which revise their beliefs downward divest. The result is a positive net investment on average, the magnitude of which decreases as firms over time gradually learn about their efficient scale. Interestingly, firms learn mostly from sources other than their productivity, as the conditional Q in the data declines too fast compared to how noisy firms' productivity is.

Decreasing shock volatility over time accounts for the remaining 32 percent of the decline in conditional investment. This mechanism also operates through skewness in productivity, similarly to learning. When the volatility of productivity is high, firms react more to positive productivity shocks than negative because there is a higher chance that another productivity shock next period could produce abnormally high profits. Since the volatility of productivity is correlated with the time since IPO but not with productivity itself, this mechanism also generates a declining conditional investment pattern.

This paper contributes to different literatures. In the firm dynamics literature, it

[1] Section 1.3.4 contains a discussion about other potential explanations that a priori can be ruled out in explaining the conditional investment decline.

is well known that conditional on size, young firms grow faster and conditional on age, small firms grow faster (see Decker et al. (2014) for a review). These results come from data on private companies and the focus is in employment growth. My contribution to that literature is twofold. First, I quantify which mechanisms are important in generating the conditional growth relationship in the context of public firms, capital investment and time since IPO rather than age. It turns out that focusing on public firms is what enables me to do this analysis in the first place, since accounting for the dynamics of Q is important to distinguish between some of the mechanisms. Second, although not the focus of this paper, to my knowledge I am the first to theoretically propose GTES in combination with increasing rigidity as an explanation for the conditional investment decline.

There is a literature on the (operational) underperformance of firms after they go public (e.g., Jain and Kini (1994); Chemmanur et al. (2009)). My contribution is to quantify the relative importance of the two main explanations provided in the literature. According to the theory of Pástor et al. (2008), it is optimal for firms to go public when their expected future profitability is sufficiently high, i.e., after a sequence of positive productivity shocks. The post-IPO drop in operating performance then simply reflects mean-reversion in productivity. This is in contrast to the theory of Clementi (2002), which attributes the drop in operating performance to firms growing towards their efficient scale. There, when firms discover a growth option, their capital stock is relatively low compared to their productivity and profitability is high. As they grow, profitability decreases. I find that 55% of the decline in profitability post-IPO is driven by GTES, and 45% by mean-reversion in productivity, validating both theories. It would be difficult to obtain these estimates without a structural model, because productivity is hard to measure.

The paper also contributes to the literature on learning, an important mechanism in my model. My contribution is twofold. First, I revisit two measures previously interpreted as evidence of learning. Pástor and Veronesi (2003) interpret the decline in Q in the years after IPO, conditional on other observables, as evidence of learning. I find that indeed, 66 percent of the decline is due to learning, most of the rest coming from growth options related to declining volatility. Moyen and Platikanov (2013) interpret the decline in the sensitivity of investment to profitability as evidence of learning. I find that only 20 percent of that decline is related to learning, the rest coming mostly from increasing rigidity. Second, I am the first to quantify how much firms learn about their efficient scale of operations after going public and how this affects investment. There are papers quantifying learning in other contexts, for example entrepreneurial choice (Catherine, 2018), learning about CEO quality (Taylor, 2010), learning about disaster risk (Hennessy and Radnaev, 2016), or learning about trading skills (Linnainmaa, 2011). The paper closest to mine on the learning front is David et al. (2016), who study how much firms learn and the relationship with misallocation across firms. While in their model firms learn about future shocks and are in a steady state, I explicitly study how firms gradually learn about their (persistent) quality post-IPO and the dynamics thereof.

The remainder of the paper is organized as follows. The next section describes the data and documents the (conditional) investment decline. In section 1.3 I present the model. In section 1.4 I discuss the estimation of the model while in section 1.5

decompose the conditional investment decline into different mechanisms. I perform robustness exercises in section 1.6. Section 1.7 concludes.

1.2 Stylized facts

This section first describes the data and construction of variables. Then I document the (conditional) investment decline in the subsequent years after firms go public.

1.2.1 Data and summary statistics

The data source is the annual Compustat database. Compustat contains financial data on (most) publicly listed U.S. firms. An observation is defined at the firm-year-since-IPO level. The data source for the IPO date is also Compustat.[2] In the data, years since IPO is correlated with fiscal year for observations with long time since IPO. If there are structural shifts in firm dynamics over time, this could introduce bias. To mitigate this bias, I remove observations with more than 20 years since IPO.

Furthermore, I remove all regulated utilities (SIC 4900-4999), financial firms (SIC 6000- 6999), and quasi-governmental and nonprofit firms (SIC 9000-9999). I also remove firm-year observations with missing capital stock (*ppegt*), assets (*at*), operating income (*oibdp*) and sales (*sale*). Furthermore I also remove observations with negative sales, negative capital stock and negative assets. Compustat sometimes also includes a few observations before firms go public, which I delete too. All ratios are winsorized at the 1% level. In line with the corporate investment literature I remove firms from the sample if their capital stock in any one year is smaller than one million (2010) dollars.[3]

Table 1.1: Summary Statistics

	Obs.	Mean.	St. dev.	25th	Median	75th
Fiscal year	65816	2002	7.536	1997	2002	2008
Year of IPO	65816	1995	6.678	1991	1995	1999
Years since IPO	65816	6.654	5.424	2	5	10
investment rate	57351	0.210	0.272	0.061	0.122	0.241
scaled operating income	57351	0.289	1.508	0.051	0.238	0.595
sales-to-capital ratio	57351	1.135	1.204	0.407	1.163	1.870
capital stock (millions)	65816	1370	10832	12	59	329
log of Q	65816	1.503	1.332	0.489	1.375	2.380

This Table reports the summary statistics used in the analysis. For each variable, I present its mean, median, 25th and 75th percentiles, and its standard deviation, as well as the number of non-missing observations for this variable. All variables are defined in Table A.31. The sample period is from 1960 to 2018.

The final sample is a panel data set with 65'816 firm-year-since-IPO observations

[2] I find that this source practically has the same coverage as the widely used data collected by Jay Ritter, available on his website.

[3] All dollar variables are deflated by the CPI index provided by the St. Louis Fed (2010 dollars).

from 1960 to 2018.[4] Table 1.1 reports the summary statistics. The firms in the sample are skewed in size and being public are obviously large, the average capital stock being 1370 million dollars and the median being 59 million dollars. They are also rather profitable, operating income being 29% of the capital stock on average. Q, the ratio of market valuation to the capital stock, is also pretty high. It is an unbalanced panel. On average 12% of firms exit the sample in a given year, so the sample is skewed towards observations with fewer years since IPO.

1.2.2 Declining investment

Figure 1.1: Investment decline after IPO.

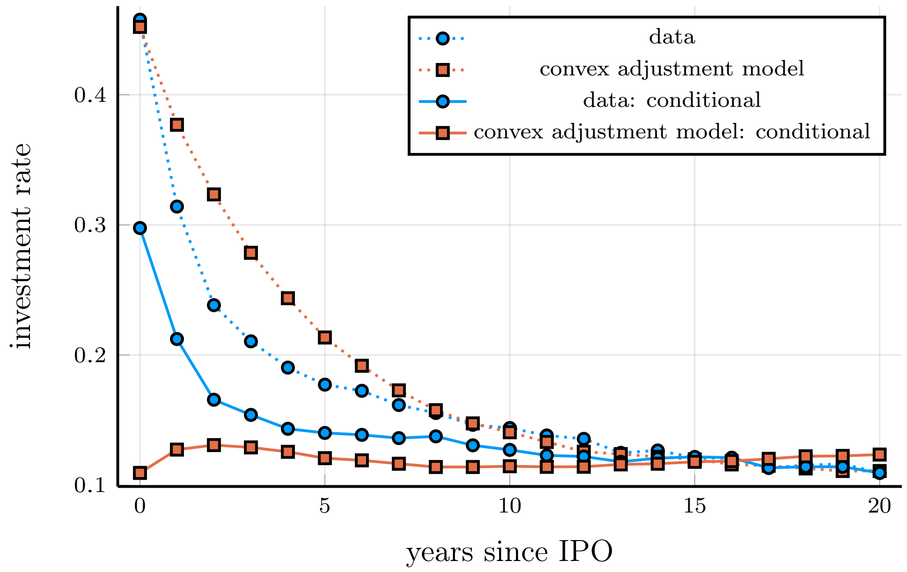

The dashed lines plots the investment rate and the solid lines plots the conditional investment rate, for the data (circles) and a standard neoclassical model with convex adjustment costs (squares) (Clementi, 2002; Cooper and Haltiwanger, 2006). Conditional investment is the age dummies obtained from regressing the investment rate on years-since-IPO dummies, the log capital stock, the log sales-to-capital ratio, scaled operating income and squared terms thereof while using firm fixed effects (column 2 in Table A.32).

The dashed blue line in Figure 1.1 shows that firms' capital investment rate decreases in the subsequent years after they go public. The solid blue line plots firms conditional investment rate, which is obtained by regressing the investment rate on years-since-IPO dummies while also controlling for size and two different measures of profitability. The area under the conditional investment curve is 60 percent of the area under the unconditional investment curve. More than half of investment can thus not be explained by size or proftiability. The red lines are obtained from simulating a standard convex adjustment cost model (see e.g., Clementi (2002); Cooper and Haltiwanger (2006)) where firms start with a lower than steady-state capital stock. This produces a declining investment pattern, due to a decreasing marginal product of capital as they converge to their steady state capital stock. I obtain the solid red line by running the

[4] Figure A.41 shows that most initial public offerings in the sample are concentrated in the period from the mid eighties to 2001 when the dot-com bubble burst.

conditional investment regression on the sample obtained from the simulated model. No matter the parameter values in the convex adjustment model, the line is always flat. Intuitively, the state variables in that model are only productivity and the capital stock, so all the variation in investment is soaked up by profitability (which is highly correlated with productivity) and size.

1.3 The model

In this section I present a dynamic investment model which contains different sources of the conditional decline in investment. The model is a neoclassical convex adjustment cost model (see e.g., Cooper and Haltiwanger (2006)). It starts at a point in time which corresponds to the IPO. I extend it model across five dimensions. First, I assume that firms start with a smaller than efficient capital stock. Even though the model of Clementi (2002) where firms grow towards their efficient scale can not explain the conditional investment decline, it can explain the unconditional investment decline. Omitting this (likely) first-order mechanism could introduce a bias when I estimate the model. Furthermore it could interact with other sources. Second, the firm dynamics literature proposes learning as a source for the conditional investment decline in the context of private firms (see e.g., Jovanovic (1982); Arkolakis et al. (2018)). In the model I therefore assume that firms are learning about their efficient scale. Third, according to the theory of Pástor et al. (2008), firms optimally choose to go public when their productivity is temporarily high. This implies that after the IPO it tends to revert back to the mean. Although this mechanism is not an obvious source for the conditional investment decline, including it in the model allows me to quantify the sources of the post-IPO underperformance documented in the literature (Jain and Kini, 1994). Therefore I assume that firms start on average with an abnormally high productivity. Fourth, Asker et al. (2014) show that public firms invest less and respond less to profitability than (similar) private firms, and interpret this as differences in adjustment costs. It is possible that this phenomenon reflects the process of firms gradually becoming more rigid after they go public. Therefore I assume in the model that firms' adjustment costs increase over time.[5] Fifth, in the data firms' volatility of profitability decreases in the subsequent years after they go public. Thus I assume that the volatility of productivity decreases over time.

After presenting the model, I also provide intuition behind how learning and volatility can produce the conditional investment pattern. Finally, I discuss other potential mechanisms and why I chose to omit them.

1.3.1 Cash flows

I begin by laying out my assumptions about cash flows. Consider an infinitely lived firm making capital investment decisions in discrete time. The firm operates at decreasing returns to scale. Sales at time t are $A_t K_t^{\alpha}$. The variable A_t reflects the time-varying productivity or demand for the firm's products, K is the capital stock and α the returns

[5] I am agnostic about the sources of increasing rigidity, for example short-termism (Holmström, 1999) or organizational rigidity (Ferreira et al., 2012).

to scale parameter.[6] Log productivity is the sum of a transitory component z_t and a permanent component μ, i.e., $\ln A_t \equiv a_t = z_t + \mu$. The transitory component follows an AR(1) process with decreasing volatility:

$$z_{t+1} = \rho_z z_t + \varepsilon_{z,t+1}, \text{ where } \varepsilon_{z,t+1} \sim N(0, f_z(t)\sigma_z^2) \tag{1.1}$$

$$\text{and } f_z(t) = 1 + \frac{b_z}{1+t}, \tag{1.2}$$

where ρ_z is the autocorrelation coefficient and $\varepsilon_{z,t+1}$ is an independently distributed random variable with a normal distribution. It has mean 0 and a variance that depends on t. The function $f_z(t)$ is a scaling factor which is decreasing in time t. It converges to one, which means that the variance of $\varepsilon_{z,t+1}$ converges to σ_z^2. The parameter b_z determines by how much volatility decreases over time.

Firm investment in physical capital is defined as:

$$I_t = K_{t+1} - (1-\delta)K_t, \tag{1.3}$$

where δ is the depreciation rate of capital. When the firm invests, it incurs adjustment costs, which can be thought of as profits lost as a result of the process of investment. These adjustment costs are convex in the rate of investment net of depreciation and increase over time. They are given by:

$$\Phi(I_t, K_t, t) \equiv f_\gamma(t)\frac{\gamma}{2}\left(\frac{I_t}{K_t} - \delta\right)^2 K_t, \quad \text{where } f_\gamma(t) = 1 - \frac{b_\gamma}{1+t}. \tag{1.4}$$

The parameter γ is the curvature of the adjustment cost function. It is scaled by the factor $f_\gamma(t)$, which increases over time and converges to one. The parameter b_γ determines by how much adjustment costs increase over time.

The firm maximizes the sum of cash flows, which are discounted at a constant rate β. Cash flows are operating income minus capital expenditures, where operating income is sales minus adjustment costs. Thus cash flows are given by:

$$C_t = A_t K_t^\alpha - \Phi(I_t, K_t, t) - I_t. \tag{1.5}$$

1.3.2 Initialization

The model starts at $t = 0$, which corresponds in the empirical setting to the time of IPO. At $t = 0$, the firm draws it's initial and permanent productivity, forms beliefs about it's type and chooses it's initial capital stock. It's permanent productivity μ is drawn from a normal distribution with mean 0 and variance σ_μ^2. It's initial transitory productivity z_0 is also drawn from a normal distribution but with mean μ_{z0} and variance $(1 + b_z)m_z^{-1}$, where $m_z = \frac{1-\rho_z^2}{\sigma_z^2}$. Thus if $\mu_{z0} > 0$, then the firm starts with an on average higher productivity than in the long-run. The variance of z_0 is just the long-run variance of z_t adjusted for the fact that earlier it is higher.

[6] I focus on physical capital, not other production factors. One can thus think of α as a combination of the firms market power and the relative importance of physical capital compared to other production factors (under the assumption that they are frictionlessly adjustable). It is easy to show this analytically in a setting with constant returns to scale, capital and labor as production factors and monopolistic competition (see e.g., Catherine et al. (2018)).

Following Alti (2003), the firm chooses it's initial capital stock K_0 to maximize the discounted sum of future cash flows. To capture the fact that small firms grow faster than larger firms, it faces an additional deadweight cost $C_0 \geq 0$ for purchasing capital at $t = 0$ (over and above the rental rate of capital). It's initial capital stock will thus naturally be smaller than it's expected long-run efficient capital stock.[7]

1.3.3 Beliefs

I assume that the firm does not separately observe it's permanent and transitory productivity, but only the sum of the two. At $t = 0$ it forms beliefs about μ (and z_0) upon observing productivity a_0. The firm's initial belief $\widehat{\mu}_0$ about it's permanent productivity is:

$$\widehat{\mu}_0 = a_0(1 + b_z)^{-1} m_z P_0, \text{ where } P_0 = \frac{1}{\frac{1}{\sigma_\mu^2} + \frac{m_z}{1+b_z}}, \tag{1.6}$$

which is a product of its initial productivity a_0, the precision of the initial productivity signal $(1 + b_z)^{-1} m_z$ and a factor P_0. The latter is the remaining uncertainty the firm faces after observing the productivity signal.

The firm gradually learns about μ over time. At the beginning of each period, when its productivity changes, it updates its beliefs about μ. When a is unexpectedly high, then it revises its beliefs upward. When a is unexpectedly low, then it revises it's beliefs downward. How aggressively it it does so depends firstly on the signal precision $(1 + f_z(t))^{-1} m_z$, which depends positively on the volatility and persistence of the transitory shock. A higher b_z means that the signal is noisier early on, which slows down learning. Secondly, it also depends on how much uncertainty P_t the firm still faces. Early on when P high it reacts a lot to news. Over time P declines as the firm learns about μ.

I also assume that the firm observes a private signal l about μ in each period (except at birth). This signal is orthogonal to productivity and captures all other sources the firm is learning from, for example soft information. It is given by:

$$l_t = \mu + \varepsilon_{l,t}, \text{ where } \varepsilon_{l,t} \sim N(0, \sigma_l^2), \tag{1.7}$$

where σ_l reflects the noisiness of the signal.

Summarizing, The value of the firm and it's capital investment choice thus depends on four state variables. First, it's capital stock K, second, it's (log)productivity a, third, the belief about it's permanent productivity $\widehat{\mu}$, and fourth, time t. I solve the firm's problem with value function iteration. Appendix A.1 contains details about the procedure, including the recursive formulation of the firm's maximization problem and it's choice of initial capital stock.

1.3.4 Discussion

In what follows, I first provide intuition on how learning and volatility can generate a declining conditional investment pattern. Then I discuss some of the functional forms I used in the model and which mechanisms I chose not to include.

[7] The parameter C_0 does not have a clear interpretation, it could mean physical installation costs, higher costs of external finance pre-IPO, or how large the initial growth option of the firm is. Loosely speaking, it is just a measure of how far the firm is from its expected long-run efficient capital stock.

How learning and volatility generate net investment

As firms learn, some revise their beliefs about μ upward and some downward. How does this result in net investment? Since productivity is log-normally distributed, a positive shift in beliefs of μ coincides with a bigger increase in (perceived long-term) productivity than the decrease in productivity resulting from a negative shift in beliefs, ceteris paribus.[8] Since productivity is linked to size, a firm with a positive belief shock invests more than a firm with a negative belief shock divests, ceteris paribus. Furthermore, for the same reason a firm that received a positive belief shock in the past also responds stronger to a positive belief shock in the future, being further to the right in the productivity distribution. The more firms are learning, the stronger this mechanism. Since learning decreases over time, this results in a declining investment pattern. A similar logic applies to how volatility produces net investment. Those firms with a positive transitory shock invest more than those with a negative shock disinvest, ceteris paribus. Since volatility also decreases over time, this also results in a declining investment pattern.

The question remains how these two mechanisms can generate a decline in conditional investment. Regarding learning, since private signals are unobservable to the econometrician, controlling for profitability will not capture that. Second, as previously argued, the higher the volatility of productivity, the stronger the investment asymmetry. Essentially the omitted variable in the conditional investment regression is how strongly firms update their beliefs about μ in the case of learning and how volatile productivity is in the case of the volatility mechanism. Both these variables are correlated with time, which is what produces the variation in conditional investment.

Functional forms for rigidity and volatility

It is plausible that firms become more rigid and their volatility declines as they become larger, rather than over time as I assumed in the model. But parameterizing firms' rigidity or volatility as a function of size makes it impossible for either mechanism to match the decline in investment conditional on size and profitability, the main stylized fact.[9]

Other possible mechanisms

In the firm dynamics literature (which studies mostly private firms) a popular way to rationalize the declining growth as firms age (conditional on size) is selection through exit. There, when a firm's productivity is too low for it to operate profitably (e.g., due to fixed costs) it will exit (Jovanovic, 1982). Public firms however mostly exit the Compustat sample not due to liquidations, but rather due to acquisitions, mergers or other types of delistings (Grullon et al., 2017, Figure 5). Also, for private firms the likelihood of exiting is higher for younger firms (Decker et al., 2014; Arkolakis et al.,

[8] This is illustrated in Figure A.42 in the appendix.

[9] Using only one parameter to capture the increase in adjustment costs (b_Y) or decrease in shock volatility (b_z) is also a rather rigid parametrization, putting a lot of structure. Having a more flexible form with more parameters could provide a better model fit to the data. Section 1.4.3 however shows that the model does quite well as is. I also experimented with a different parametrization than the reciprocal of one plus firm age, using the log of one plus firm age. This gave very similar results in the estimation.

2018), which is consistent with models with fixed costs where firms may exit if they learn they are unprofitable (Jovanovic, 1982). For public firms however there is no declining pattern, if anything firms that more recently went public have a lower chance of exit.[10] For these reasons I implicitly assume that firms are large enough compared to their fixed costs to operate profitably. This is intuitive for public firms.

Financing frictions could also produce declining investment as firms age (Cooley and Quadrini, 2001). I omit them for multiple reasons. First, I find that the decline in conditional investment, my primary object of interest, is not explained by measures of financing constraints such as cash flows, leverage or cash holdings. Second, public firms have widespread access to capital markets. The literature on the existence and effects of financing frictions on investment for public firms is inconclusive at best (see e.g., Farre-Mensa and Ljungqvist (2016)). Furthermore, existing structural models only find modest effects of financing frictions on investment (Hennessy and Whited, 2007; Warusawitharana and Whited, 2015).

Frictions in product markets could also play a role. In Gourio and Rudanko (2014) for example, it takes time for firms to build a customer base, which translate into sluggish sales expansion. This friction slows down capital growth and reduces the sensitivity of capital investment to productivity shocks. The amount of customers a firm has could be an omitted variable, which is correlated with the time since IPO. Firms which recently went public however grow faster (conditional on size) and react more to shocks than firms which have been public for longer. The friction therefore produces exactly the opposite of what is required to match the declining sensitivity of investment to profitability I find in my sample. More generally, this argument applies to any friction that dampens the sensitivity of investment to profitability more for firms that recently went public.

Next, firms may continuously become less innovative from the moment on that they are founded, implying they get less new ideas over time. If a new idea warrants capital investment then this mechanism could produce a decline in investment. I show in section 1.6.1 however that the post-IPO dynamics in investment and conditional investment are related to the number years since the IPO, not the number of years since founding.

A more nuanced theory that is in line with this fact is the one of Ferreira et al. (2012). There firms go public precisely to exploit existing ideas, which on the other hand reduces their innovative capacity. One can think of an extension of the theory where this tradeoff manifests itself gradually during the first few years after the IPO. Similarly, firms may still have some ideas left over from the time when they used to be private and may implement them gradually when they go public until they run out. If implementing a new idea leads to capital investment (i.e., an increase in size), then this mechanism could produce the conditional investment decline. However, the implementation of an idea should also increase the productivity of a firm if it warrants an increase in firm size. It turns out however that productivity declines after firms go public (Chemmanur et al., 2009, Figure 1).

Theories of the going-public decision revolving around competition could also affect post-IPO dynamics (Bhattacharya and Ritter, 1983; Maksimovic and Pichler, 2001). Firms' market share is flat in the years after they go public however, making it unlikely (Chemmanur et al., 2009, Figure 6).

[10] See Figure A.43 in the appendix.

1.4 Estimation and identification

In this section I discuss the estimation of the model. This involves calibrating two parameters and estimating the rest using the simulated method of moments (SSM). I will also discuss in detail how my choice of moments pins down the parameters in the model. Lastly, I discuss the parameter estimates and how well the model fits the data.

1.4.1 Estimation

I estimate most of the structural parameters of the model using SMM. This involves simulating the model and calculating interesting moments from the simulated data and actual data. SMM then finds the model parameters that make the actual and simulated moments as close as possible. Appendix A.2 provides technical details on the estimation procedure.

I calibrate two parameters. First, following Warusawitharana and Whited (2015) I estimate the risk-free interest rate, r, as the average real Baa interest rate over the (observation-weighted) sample period.[11] This translates into a discount factor β of 0.95. Second, as previously mentioned firms sometimes exit the sample. Without any adjustments this could introduce bias towards firms with a long time since IPO. I therefore assume that firms face an exogenous probability π of exiting the simulated sample. Since most firms exit Compustat due to acquisitions or delistings, I assume that the chance of exiting does not distort firms' decisions, i.e., it does not affect firms' discount factor. I calibrate $\pi = 0.12$ to match the exit rate in the data.

I then estimate the following 11 parameters using SMM: the capital depreciation rate, δ; the capital initialization cost, C_0; the baseline volatility and autocorrelation of the productivity process, σ and ρ; the returns to scale parameter, α; the quadratic adjustment cost parameter, γ; the dispersion in firm quality, σ_μ; the volatility of the private signal, σ_l; the mean initial transitory productivity, μ_{z0}; the volatility scaling parameter, b_z and the rigidity scaling parameter, b_γ.

1.4.2 Model identification

In this section I discuss my choice of moments and how they help pin down the model parameters. How well the estimation works depends on whether the moments change a lot depending on the parameters and that the model is identified in a statistical sense. That is, all I require for the model to be identified is that the Jacobian matrix (how the moments change with the parameters) is invertible. But for the procedure to be successful and the estimates credible, exactly how individual moments contribute to pinning down the parameters must be as transparent as possible. The entire analysis is about local identification, in the sense that I operate around the main SMM estimate for $(\delta, C_0, \sigma_z, \rho_z, \alpha, \gamma, \sigma_\mu, \sigma_l, \mu_{z0}, b_z, b_\gamma)$ – which I discuss in detail in what follows.

Standard moments

The average investment rate of firms with more than 10 years since IPO is highly informative about the depreciation rate of capital δ. The higher δ, the larger the required investment rate to maintain the capital stock.

[11] I obtained the data from the Federal Reserve Bank of St. Louis FRED database.

The average scaled operating income of firms with more than 10 years since IPO helps pin down the returns to scale parameter α. In a hypothetical environment with constant returns to scale and monopolistic competition, higher competition translates into a higher α and intuitively also into lower scaled operating income. Thus they are inversely related.

The dynamics of the log sales-to-capital ratio are informative about the parameters driving the TFP process. Specifically, the autocorrelation thereof is informative about the autocorrelation of productivity ρ_z. The variance of log sales-to-capital growth of firms with long time since IPO pin down the baseline volatility of productivity σ_z. The same moment but for firms with five or less years since IPO helps pin down the volatility scaling parameter b_z.[12]

Capturing firm dynamics

The moments I discuss next are related to the dynamics of variables such as investment, profitability and Q after the IPO. Here, a delicate issue arises whether to take a parametric or non-parametric approach to describe the data in terms of moments, and how to implement it. Figure 1.1 for example which showed the stylized facts plots 21 years-since-IPO dummies, which could all be used as moments. Such a granular, non-parametric approach is problematic. The SMM procedure will put more weight on the first few years because there are more observations. This results in a great fit in the first few years but a terrible fit later on. I experimented with forming age groups that have roughly the same number of observations, which works better. But a problem that remains with providing detailed information about slopes at different points in time is that the dynamics of some variables such as Q will be matched (too) well, and the dynamics of other variables such as profitability will be matched poorly.

I also experimented with parametric approaches. For example, I attempted to fit the decline in the variables simply using the reciprocal of one plus firm age.[13] This approach worked relatively well, due to its simplicity and because empirically the decline in the variables of interest is close to proportional to the reciprocal of one plus firm age.

Summarizing, there is a delicate tradeoff between providing a lot of information to the SMM procedure and not matching the total decline of some variables and providing little information, matching poorly the speed of decline. I settle for the the parametric approach, using the reciprocal of one plus firm age for all variables. The only deviation of this concerns the decline in conditional Q and the sensitivity of investment to profitability where I use a more granular approach, described later in detail. What provides piece of mind is that when I granularly compare in section 1.4.3 the dynamics of the simulated model to the data, the fit is pretty good.

[12] Note that both the sales-to-capital ratio and scaled operating income could be classified as measures of profitability. The latter also includes costs of goods sold and SG&A. In Compustat, SG&A is a combined measure for fixed overhead costs, the wage bill, marketing costs and training costs. The log sales-to-capital ratio is my preferred measure of profitability to capture the dynamics of productivity, since it is least contaminated by the firms decisions. When mentioning profitability, from now on I am referring to scaled operating income when discussing the level of profitability, and in all other cases I am referring to the log of the sales-to-capital ratio.

[13] Using the natural log of one plus firm age gave similar results.

Novel moments

Next, I regress the log of the sales-to-capital ratio on a constant and the reciprocal of one plus firm age. The age coefficient captures the decline in profitability as firms age and is informative about the average initial transitory productivity μ_{z0}. Due to mean-reversion, intuitively productivity and thus profitability declines over time.

I obtain additional moments from an empirical policy function regression (EPF), which estimates an approximation the firms policy function, i.e., investment as a function of the state variables. Note that not all state variables are observable to the econometrician, so the EPF will be incomplete. This is only a problem to the extent that it may be imprecise, as any bias will be reflected equally in the estimates of the EPF based on the empirical and simulated sample. Specifically, I regress investment growth on profitability growth, log capital stock growth, and interactions with years-since-IPO group dummies.[14] In the EPF regression, the (un-interacted) coefficient on profitability is highly informative about the adjustment cost parameter γ. The higher γ, the less investment reacts to productivity shocks, which implies a lower regression coefficient.

Which moments are informative about C_0, σ_μ, b_γ and σ_l remains to be discussed. I normalize C_0 to isolate GTES from the other sources, which facilitates the identification and counterfactual experiments later on. Details are in Appendix A.2.1.

How the parameters affect the informative moments is shown in Figure 1.2. The decline in the sensitivity of investment to profitability as firms age is highly informative about the increasing adjustment cost parameter b_γ. This is intuitive as firms with lower adjustment costs respond more to shocks. Note that also C_0 affects the decline in the investment-profitability sensitivity because early on firms have a high marginal product of capital (because K is low), which makes them react more to changes in profitability. This is obvious if one takes the partial derivative of AK^α/K with respect to K and productivity. The smaller K, the larger is the resulting term. Apart from C_0, also σ_μ affects that moment positively. When firms are learning, then current profitability also has an informational component. They react more to profitability because it provides a signal about their efficient scale, as opposed to merely reacting to transitory shocks. Importantly for what follows, the other moments in Figure 1.2 are not affected by b_γ.

Next, I regress the log of Q on equally-sized year-since-IPO groups, the log capital stock, the log of the sales-to-capital ratio and squared terms thereof. I also use firm fixed effects to account for persistent heterogeneity in firms' profitability and size. I target different years-since-IPO group dummies in the estimation to capture the speed of decline in conditional Q, which is informative about the precision of the private signal. The total decline in conditional Q is almost flat in C_0, which helps the SMM procedure to separate C_0 from (σ_μ, σ_l). The flatness comes from the fact that C_0, representing GTES, operates through profitability and size which I control for.

Next, I regress the log of Q on a constant and the reciprocal of one plus the number of years since IPO. The coefficient on the latter captures the decline in unconditional Q over time and is informative about C_0. If firms start with a lower than long-run efficient capital stock, then future cash flows are relatively high compared to their current size, so early on Q is high. Over time, as they grow, the denominator of Q increases and

[14] I sort all observations across years since IPO and form five groups that all have (roughly) the same number of observations. The resulting groups are 0 and 1; 2 and 3; 4 to 6; 7 to 11; 12 to 20. Table A.33 in the appendix shows this regression for the data.

Figure 1.2: Comparative statics of critical parameters at the SMM estimates.

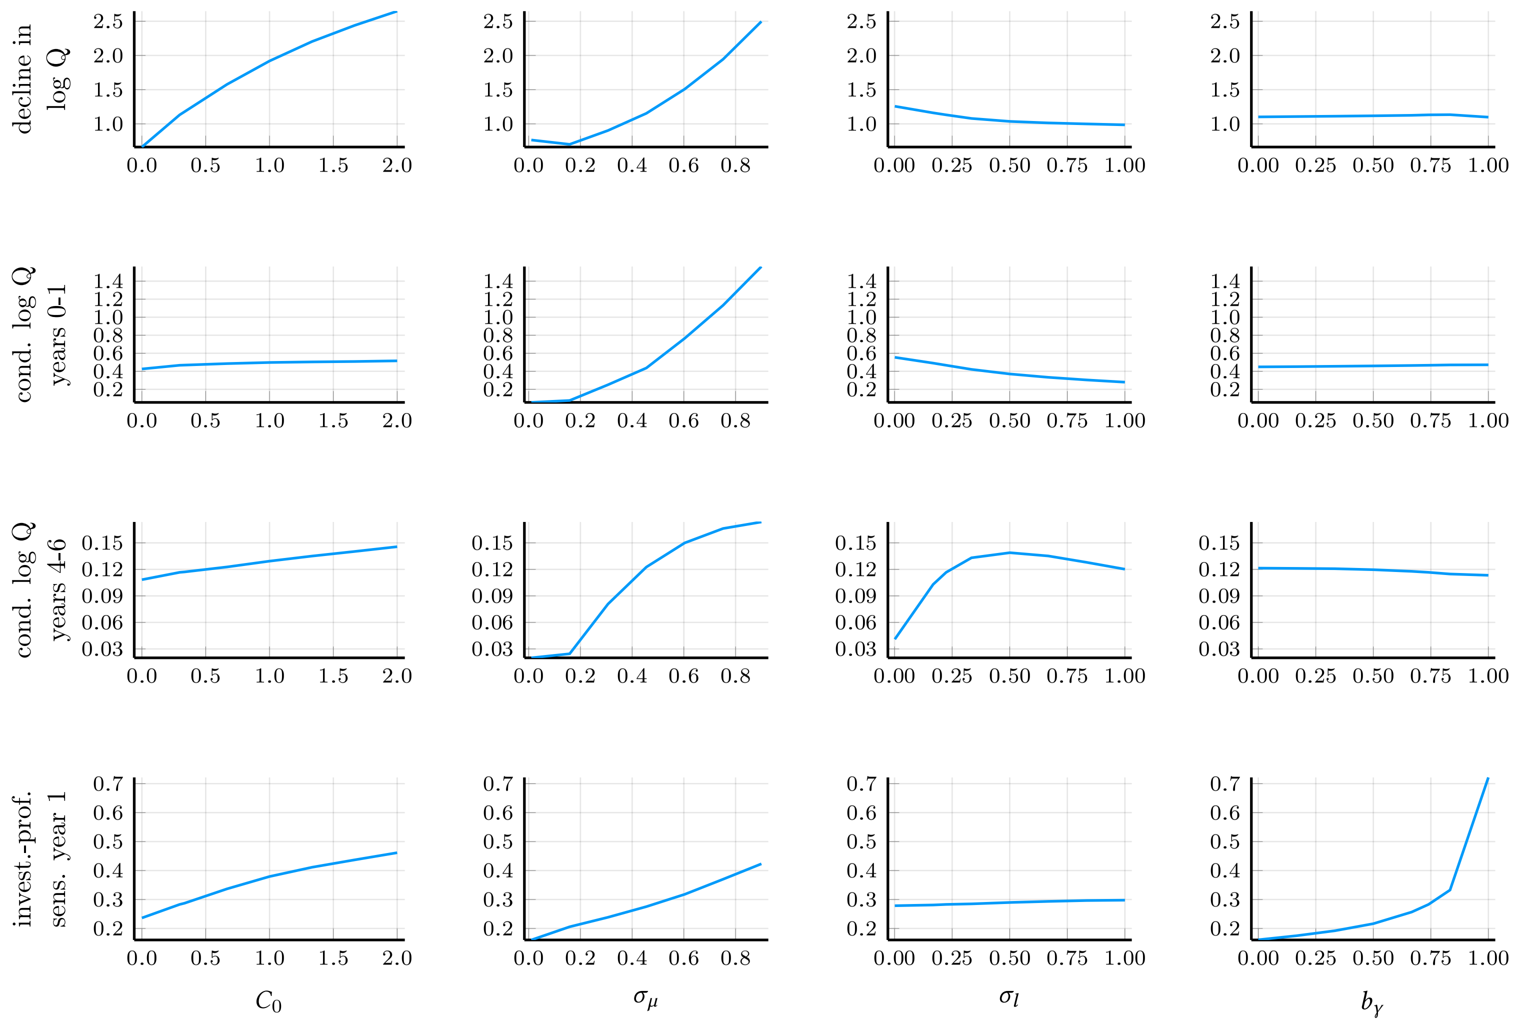

In these figures only one parameter is changed at a time, while the others are held constant at the SMM estimates. *Decline in log Q* is obtained from regressing log Q on the reciprocal of one plus years-since-IPO. *Cond. log Q years 0-1* is obtained from regressing log Q on equally-sized groups of years-since-IPO dummies, the log capital stock, the log sales-to-capital ratio and squared terms thereof while using firm fixed effects. *0-1* and *4-6* refers to 0 and 1 and 4 to 6 years since IPO, respectively. C_0 is the (normalized) initial deadweight cost of capital, σ_μ is the dispersion in quality, σ_l is the noisiness of the private signal and b_γ is the adjustment cost scaling factor.

it shrinks. The decline in unconditional Q is also affected by σ_μ, but as previously mentioned the decline in conditional Q helps separate them.

Finally, comparing the total decline in conditional Q and the decline in conditional Q from years 4 to 6 provides information on the speed of decline in conditional Q. The latter is highly informative about the noisiness of the private signal σ_l. The lower σ_l, the faster firms learn. The total decline in conditional Q is informative about the total amount of learning/uncertainty, σ_μ.

I target two additional moments, the decline in investment and the decline in conditional investment. I use a parametric approach for both.

Limitations of the model

Naturally there are features of the data that the model is not designed to explain. If I force the model to match them nevertheless, it could do poorly in matching the features of the data I am actually interested in. In what follows I discuss this.

First, in the data the autocorrelation of investment is relatively low. It is well known that non-convex adjustment costs are necessary to match this fact (Cooper and Haltiwanger, 2006). Experimenting with the inclusion of non-convex adjustment costs revealed however that they don't play a first-order effect in the mechanisms that I am studying. Intuitively, non-convex adjustment costs have more of a bite in models with financing frictions for example, where firms may have difficulty financing lumpy investment outlays out of profits (see e.g., Riddick and Whited (2009)). I am already estimating quite a lot of parameters. Including more would increase complexity without adding any insights. Therefore I don't target the autocorrelation of investment and the coefficient on capital stock growth in the EPF regression (but I do control for it).

Second, in the data the correlation of investment with other observables such as profitability is relatively low. The model can not match this because there are only productivity shocks. Since the correlation can be decomposed into the variance of investment, variance of profitability and covariance of investment and profitability, I therefore can only match two out of the three. I must target the variance of profitability to pin down the TFP parameters. If I choose to match the entire variation of investment, then in the model it will come mostly from variation in profitability, whereas in reality there must be other unobservable sources (or measurement error). By instead matching the covariance of investment with profitability, I can at least capture that relationship, while admitting that I can not match the entire variation of investment.

Third, I do not attempt to match the level of Q. In the model, the firm is risk-neutral so stock prices and Q don't contain any risk premia. This is clearly at odds with reality.

1.4.3 Goodness of fit

Figure 1.3 plots all variables of interest against years since IPO, comparing the data to the simulated sample obtained from the SMM estimates.[15] The model fits the data well. While it matches the total decline in investment, it declines slightly too fast. The

[15] Table A.34 in the appendix reports the fit of the moments targeted in the estimation. Table A.35 in the appendix contains the Jacobian matrix, i.e., the derivatives of the moments with respect to the parameters.

Figure 1.3: Model fit.

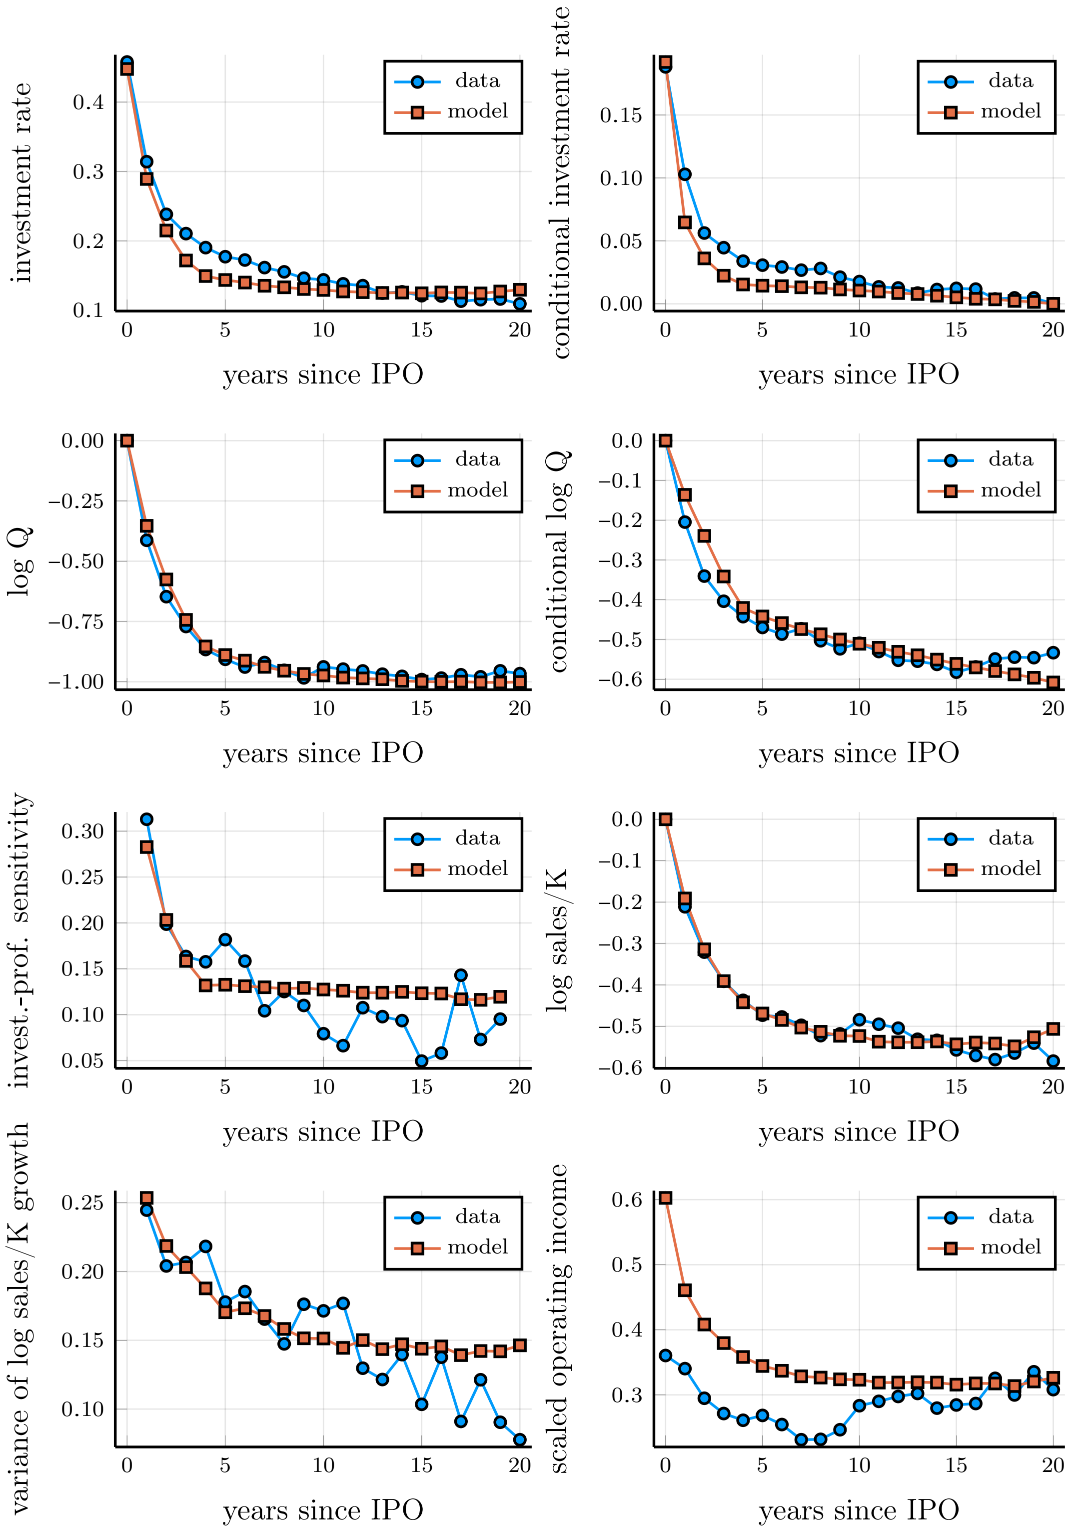

This Table plots the variables of interest for the data and simulated sample at the SMM estimates. The scale on the y-axis of the conditional investment rate starts with zero because that specification uses firm fixed effects, which removes the mean. The decline in log Q, conditional log Q and log sales/K are normalized (start at zero) because I do not attempt to match the means of those variables. I also do not attempt to match the decline in scaled operating income. The investment to profitability sensitivity is obtained by regressing investment growth on the growth of the log sales to capital ratio and log capital growth for each year since IPO.

area between the orange and blue curve represents 31 percent of the total area under the blue curve.

Although my primary interest is explaining the decline in investment, it is surprising how well the model fits the dynamics of other variables too. It has a bit of difficulty matching the decline in the sensitivity of investment profitability from year four on. Also, conditional Q declines too slow compared to the data. The model also has no chance matching the different decline in the sales to capital ratio and scaled operating income. The reason is that they are constructed almost the same in the model, but in the data scaled operating income also contains sales, general and administrative costs (*SG&A*).

1.4.4 Results

Table 1.2: Parameters estimated with SMM.

	Par	Estimate	Std. Err.
depreciation rate	δ	0.138	0.002
(normalized) initial C.o.C.	C_0	0.296	0.055
volatility of productivity	σ_z	0.337	0.007
autocorrelation of productivity	ρ_z	0.800	0.010
returns to scale	α	0.506	0.023
adjustment costs	γ	3.267	0.243
dispersion in quality	σ_μ	0.460	0.011
private signal noisiness	σ_l	0.224	0.074
initial mean log productivity	μ_{z0}	0.400	0.034
volatility scaling parameter	b_z	1.702	0.143
adjustment cost scaling parameter	b_γ	0.740	0.037

Table 1.2 reports the estimated parameters and their respective standard errors. All parameters are significantly different from zero.

The estimated volatility and autocorrelation of TFP, as well as the depreciation rate of capital are values in line with the literature. The returns to scale parameter $\alpha = 0.51$ is lower than is typically calibrated, which is usually in the range of 0.6 to 0.7 (see e.g., Gomes (2001)).[16] In a robustness exercise in section 1.6.3, I show however that my results are not that sensitive to the estimate of α.

The estimated adjustment costs are $\gamma = 3.3$, which is higher than what what people typically find in the firm dynamics literature (Cooper and Haltiwanger, 2006) and structural corporate finance literature (DeAngelo et al., 2011). The reason is that in these studies, people target the variance of investment, which is relatively high, whereas I target the covariance of investment and profitability, which is relatively low. Alti (2003) for example calibrates γ to match regression coefficients as I do and gets a similar value as I do. Interestingly, I estimate that adjustment costs increase quite a

[16] In the structural corporate finance literature, people sometimes obtain estimates of close to 0.8 (see e.g., DeAngelo et al. (2011)). The reason is that in their sample average profitability is lower, which translates into a higher α. In Compustat, depending on the sample selection and winsorization, estimates of average scaled operating income can be sensitive to outliers.

lot over time. The adjustment cost scaling parameter $b_\gamma = 0.74$ together with $\gamma = 3.3$ implies that firms that just went public have roughly 3.5 times lower adjustment costs compared to 10 years later. This is in line with the reduced-form evidence of Asker et al. (2014), who compare private and public firms. The average estimated adjustment cost averaged over all observations in the sample is 2.5.

The parameter $C_0 = 0.3$ capturing GTES does not have a direct interpretation. It is informative however to translate it into how much firms are growing towards their efficient capital stock after their IPO (assuming they know their efficient scale). This turns out to be 0.67 times an order of a magnitude.

The estimate $b_z = 1.7$ implies that the volatility of productivity significantly decreases over time. The standard deviation of the shocks decreases from 0.56 at $t = 0$ to 0.36 10 years later, a 35 percent reduction.

Furthermore, the estimates imply that when firms go public, uncertainty about their efficient scale is quite high. This on the one hand has to do with $\sigma_\mu = 0.46$ being quite high and on the other hand firms having quite volatile productivity at the IPO, which makes it a rather uninformative signal. The private signal $\sigma_l = 0.22$ is estimated to be quite precise, because the fast decline in conditional Q implies that uncertainty about μ must be resolving quickly compared to the noisiness of productivity. The weight put on the private signal is high at $m_L = 18.4$, whereas the weight put on the profitability signal $m_x = 0.2$ is low. Taylor (2010) also finds that non-earnings information is much more informative in the context of boards learning about their CEOs skill.

It is informative to compare the uncertainty related to learning from the long-run variation in transitory productivity, which depends on σ_z and ρ_z. For a firm trying to forecast its transitory productivity z many years from now, the 95% confidence interval lies in $[-1.77, 1.77]$. If the firm at t=0 estimates what its true μ is, then the 95% confidence interval lies in $[-1.32, 1.32]$. The intervals look similar but imply quite different variations in size. The standard deviation in the long-run log-capital stock coming from uncertainty about μ is 0.80, much larger than the variation coming from transitory shocks, which is 0.30. The reason for this is that μ represents permanent productivity compared to z which is transitory, so firms react much less to changes in z compared to changes in beliefs about μ.

The abnormal initial productivity μ_{z0} is estimated at 0.4. To put this into perspective, it is 71% of the long-run variation (standard deviation) of transitory shocks. Thus at IPO, on average firms transitory productivity z is at the 76th percentile of their long-run transitory productivity (as opposed to the 50th percentile for $\mu_{z0} = 0$).

1.5 What explains firm's post-IPO dynamics?

Previously, I established that my model fits the data well. Now I use the estimated model to investigate which mechanisms are quantitatively important in explaining firms' post-IPO dynamics. Of particular interest are the sources of the decline in conditional investment. To do this, I run different counterfactual experiments where each source is either activated (set to the SMM estimate) or deactivated. I then compare the dynamics of these counterfactuals with the dynamics of the baseline estimation. All five mechanisms are directly related to five specific model parameters, described as follows. If $C_0 = 0$, then firms are not growing towards their efficient scale. Because of the normalization of C_0 discussed in Appendix A.2.1, setting $C_0 = 0$ only affects

GTES. If the firm receives a perfect signal about μ when it is born then there is no learning. If $b_\gamma = 0$, then adjustment costs are constant. If $b_z = 0$, then the volatility of productivity is constant. If $\mu_{z0} = 0$, then firms on average don't have an abnormal productivity at IPO. Due to interactions there are many possible combinations possible, so in what follows I only discuss those that illustrate the main mechanisms at work. This is the outcome of an iterative procedure where I started with examining all combinations and removing one at a time to get to the drivers of the dynamics.[17]

First I analyze which mechanisms explain the decline in conditional investment. Figure 1.4 shows that no source alone comes close to matching it, so it must come from combinations of sources and interactions between them. Eyeballing the Figure, it looks like learning and volatility together could account for the slow and steady the decline from year four on. This is indeed the case, as is shown in Figure 1.5. During the first four years, learning explains 31 percent and volatility 21 percent of investment. Of the total conditional investment, learning contributes 37 percent, and volatility 32 percent.[18]

Figure 1.4: Conditional investment - individual sources.

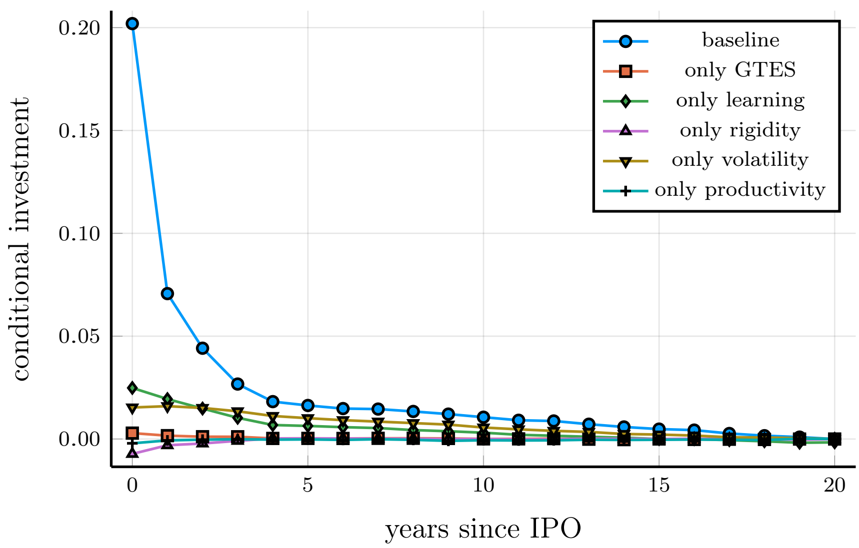

The circle line plots the simulated conditional investment rate at the SMM estimates, i.e., when all sources are activated. The other lines are simulations from the model when only one source at a time is activated.

Figure 1.5 shows that the interaction of GTES and rigidity explains the remaining 32 percent of conditional investment. The intuition is as follows. GTES creates a strong motive for firms to invest, but it is entirely explained by profitability (and size). In combination with asymmetric adjustment costs over time however, early on firms respond more to fundamentals than firms long since their IPO. Therefore two firms with the same profitability and size can have very different investment rates. When they want to grow, the firm that recently went public will invest more aggressively.

[17] Table A.36 in the appendix shows the results of all combinations of counterfactual experiments.

[18] These percentages are calculated as the integral of the respective curves divided by the integral of the curve of the baseline simulation at the SMM estimates.

Figure 1.5: Conditional investment - main mechanisms.

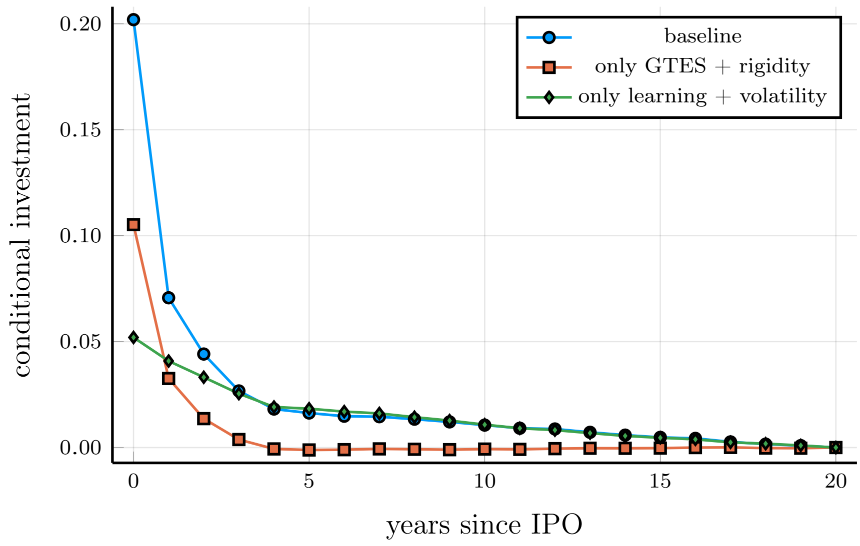

The effect of GTES in combination with rigidity is most prominent in the first four years after IPO, where it explains 48 percent of investment. Apart from the interaction between GTES and rigidity, I find that other interaction effects between the mechanisms are quite small.

Figure 1.6: Unconditional investment.

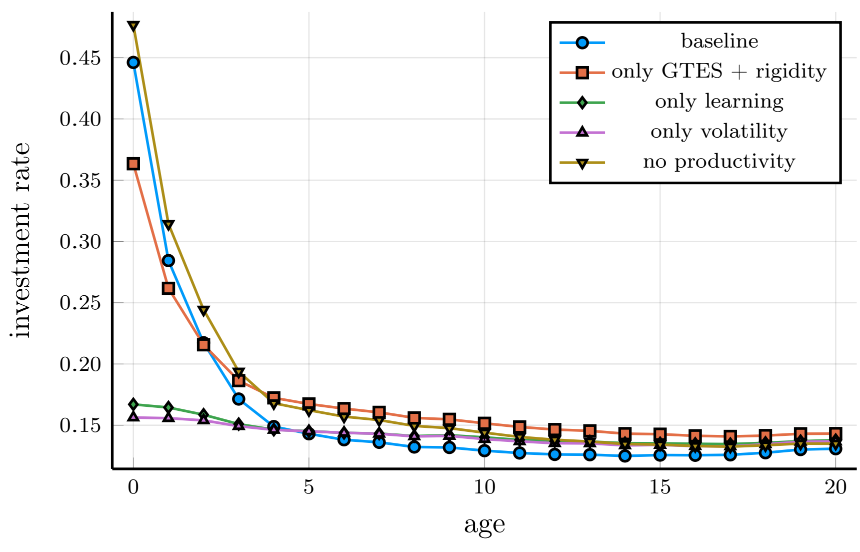

No productivity means that all sources are activated except high abnormal initial productivity, which is set to zero.

Next, Figure 1.6 illustrates that GTES is the dominant mechanism responsible for the decline in unconditional investment. Rigidity interacts with GTES only by changing the timing of investment, i.e., how fast firms grow towards their efficient scale. However, alone it does not generate a declining investment pattern because for firms that recently went public there is no particular investment motive compared to that it could accelerate. The effects of learning and volatility are small, and there is little interaction between them and GTES.

Interestingly, Figure 1.6 also shows that without an abnormally high productivity at IPO, the investment decline would actually be more pronounced. The reason is that if firms (in the model) start with an abnormally high productivity, then they choose a larger capital stock to take advantage of it. This means they have to grow less towards their efficient scale. Put differently, if firms have abnormally high (transitory) productivity at IPO, then optimally they must have grown before the IPO and ceteris paribus they will shrink after the IPO when productivity mean-reverts. In the data of course there are other mechanisms in play that overall incentivize them to grow.

Figure 1.7: Unconditional Q.

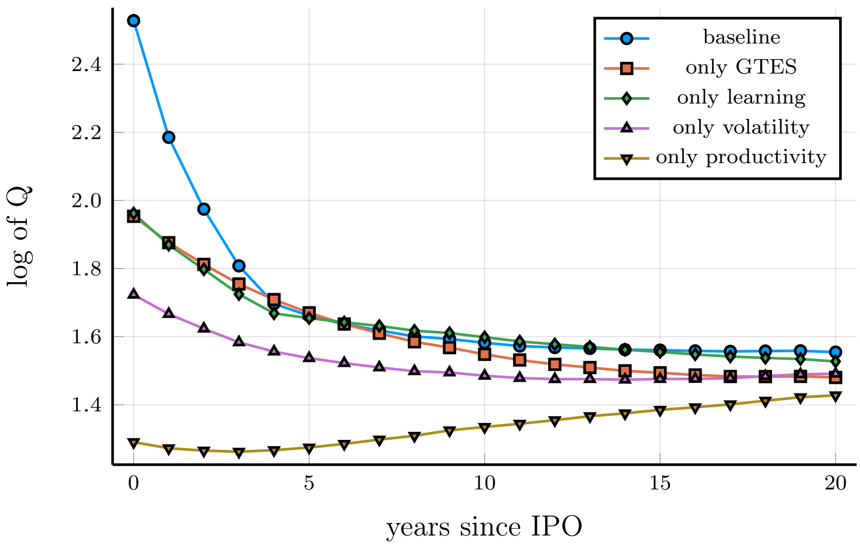

Next, Figure 1.7 shows that GTES, learning and volatility produce a declining pattern in Q, whereas abnormal initial productivity increases it. GTES contributes 40 percent towards the decrease in log Q, the reason being that size in the denominator increases. Learning and volatility explain 39 percent and 21 percent, respectively, by creating an implicit option value which decreases over time (see section 1.3.4). Abnormal initial productivity dampens the decline in log Q by 16%. The intuition is the same as previously discussed. Firms choose a larger capital stock at t=0, which means that their current size in the denominator is high compared to future profitability.

Next I turn towards the decline in conditional Q, shown in Figure 1.8. Only learning and volatility create a meaningful decline in conditional Q, learning (66%) being more important than volatility (26%). The combination of GTES and rigidity plays a very small role (8%). Pástor and Veronesi (2003) interpret the decline in conditional Q as evidence of learning. For the most part, my model confirms this.

The main sources behind the decline in profitability are shown in Figure 1.9. GTES contributes 55 percent and abnormal initial productivity 45 percent, the other mechanisms are unimportant. Evaluating the two competing theories of Clementi (2002) and Pástor et al. (2008), both seem to be roughly equally important in generating the

Figure 1.8: Conditional Q.

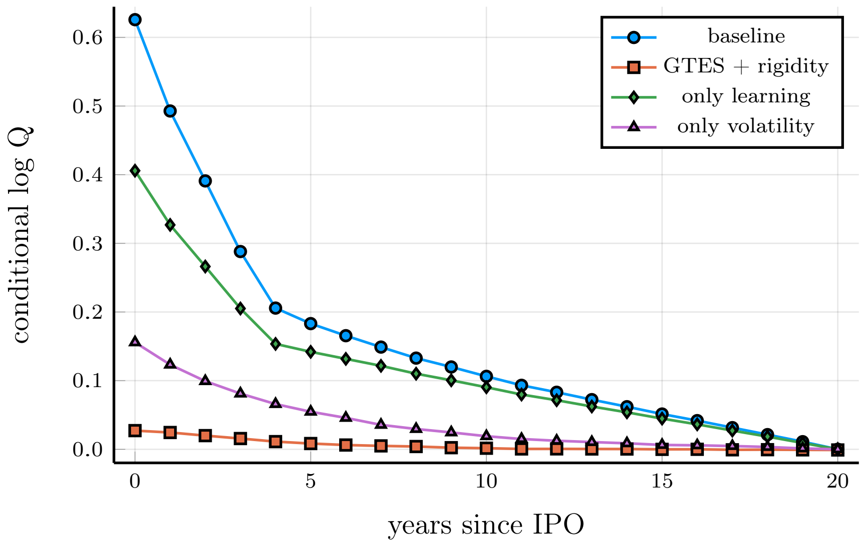

Figure 1.9: Profitability.

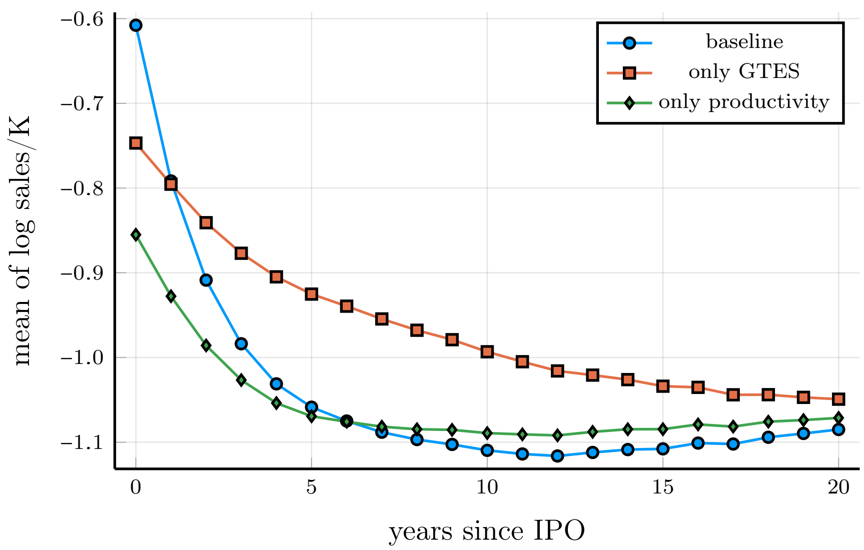

decline in profitability after firms' IPO. Profitability thus declines both because firms grow in size and because productivity mean-reverts.

Finally, I also discuss the causes of the decline in the sensitivity of investment to profitability, shown in Figure 1.10. It almost all comes from increasing rigidity (63%), a bit from learning (20%) and GTES (17%). In the learning literature, it is common to interpret a declining investment-to-profitability sensitivity as evidence of learning from cash flows (e.g., Moyen and Platikanov (2013)). My estimates paint a different picture. First, I find that firms learn mostly from other signals than profitability, so learning can not be responsible for the decline. The bulk of the decline rather comes from from increasing rigidity. This is in line with Asker et al. (2014), who show that public firms respond less to fundamentals than private firms. My analysis underlines that this is a change which happens within firm, and at least part of it occurs gradually during the

Figure 1.10: Sensitivity of investment to profitability.

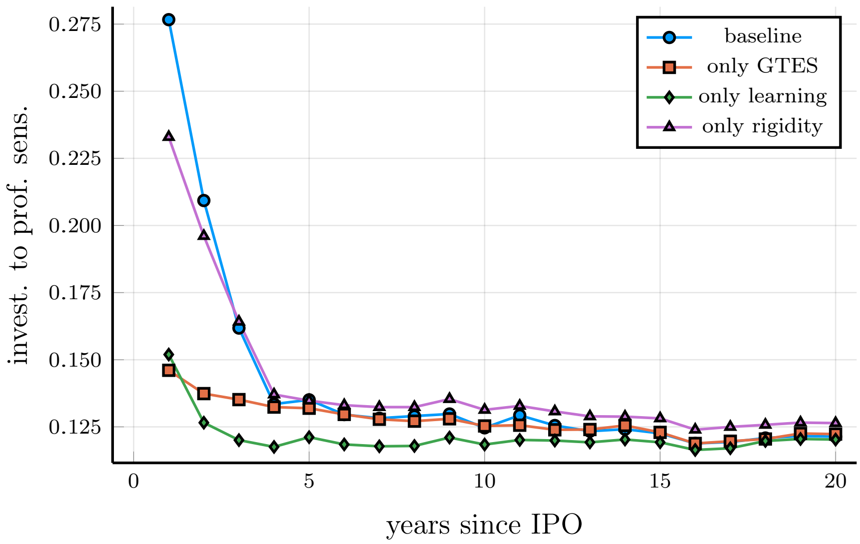

The investment to profitability sensitivity is obtained by regressing investment growth on the growth of the log sales to capital ratio and log capital growth for each firm age.

first few years after the IPO.

1.6 Robustness

In this section I perform three robustness exercises. First, I show that the investment decline is not just an artifact of firms' life-cycle since founding but is directly related to the IPO. The second is aimed at better understanding where the identification is coming from and the third gauges how sensitive my results are to the returns to scale parameter α.

1.6.1 Founding age vs years since IPO

My model explains firms' dynamics after their IPO. A natural question to ask is whether these dynamics simply are a manifestation of firms' life-cycle dynamics from when they are founded to when they dissolve, sometime in between being the IPO. It is well known from the firm dynamics literature that private firms receive large productivity shocks (Catherine, 2018) and grow in spikes (Cooper and Haltiwanger, 2006). Permanent productivity shocks basically are changes in the efficient scale of operations (GTES). Chemmanur et al. (2009) show that firms' sales growth, the investment rate, and profitability and productivity increase before and then decrease after the IPO, i.e., they spike around the IPO.

In the backlight of these previous findings, the interpretation of my approach is as follows. I study the process of firms converging to their efficient scale after having received a particularly large permanent shock. For the firms in my sample, that shock was large enough to incentivize them to go public, for example to access a broader capital market.[19] I decompose the subsequent dynamics of investment and other variables

[19] Another reason for firms to go public is for diversification purposes.

into different mechanisms.

Unfortunately, Compustat does not have reliable pre-IPO data (see the discussion in Pástor et al. (2008)) to confirm the aforementioned intuition in my sample. However, I merge it with Jay Ritter's dataset on firm founding dates. With the merged dataset, which has roughly a third of the number of observations, I run a horserace between years since founding and years since IPO. Table A.37 in the Appendix confirms that indeed years since IPO explains the decline in investment and conditional investment after the IPO much better than years since founding.

1.6.2 Identification

To better understand the model and the identification of key parameters, I estimate it under the assumptions that (i) there is no growing towards the efficient scale ($C_0 = 0$), (ii) there is no learning (firms get a perfect signal at $t = 0$), (iii) adjustment costs are constant ($b_\gamma = 0$), (iv) the volatility of productivity is constant ($b_z = 0$), and (v) have no abnormal initial productivity ($\mu_{z0} = 0$). Table A.38 in the Appendix shows the parameter estimates of this exercise and Table A.39 the fitted moments.

When there is no GTES, α is estimated much higher than in the baseline. The model can not match average profitability and the decline in conditional Q is too strong. This is surprising, as a priori the model is expected to do badly in matching the decline of unconditional Q since this moment is important in pinning down C_0, but this is not the case. If however I fix α at the baseline value and rerun the estimation, then the dispersion in firm quality σ_μ and adjustment costs are estimated to be quite high. Learning is doing all the heavy lifting in matching the the dynamics of Q. The model also can not match the decline in the sensitivity of investment to profitability. This shows that not including GTES into the estimation would strongly bias the results towards the other mechanisms, in particular learning.

If there is no learning, then α and C_0 are estimated much higher than in the baseline. Due to the high α, profitability is too low in the simulated sample, while the other moments fit rather well. However, one would expect that without learning the model would have difficulty matching the decline in conditional Q rather than average profitability. The reason is that the SMM procedure by increasing α is choosing to match the decline in conditional Q rather than profitability. To double-check I rerun the estimation without learning and fix α at the baseline SMM estimate (and do not target average profitability). The estimate for C_0 then is closer to the baseline estimate and as suspected the model can not match the decline in conditional Q. This exercise confirms that indeed the decline in conditional Q plays an important role in pinning down how much firms learn.

When adjustment costs are constant, then as expected the SMM procedure has difficulty matching the declining sensitivity of investment to profitability. The estimated parameters are close to the benchmark, except that adjustment costs γ are lower to make firms that recently went public more flexible. But then firms far from IPO react too much to profitability. This exercise confirms that indeed the decline in the investment-to-profitability sensitivity plays an important role in pinning down by how much adjustment costs increase over time.

If volatility is kept constant, then unsurprisingly the model can not match the decreasing shock volatility as firms age. As a consequence, the SMM procedure estimates a higher shock volatility σ_z than in the baseline, doing the splits between matching the volatility of profitability of firms that recently versus long since went public. This shows that b_z is important in matching the decline the volatility of profitability as firms age.

When there is no high abnormal initial productivity, then the estimated returns to scale parameter α is lower than in the baseline and C_0 higher, because they need to do the heavy lifting to match the decline in profitability. Nevertheless the model still has difficulties matching the pronounced decline in profitability. This shows that μ_{z0} is important to be able to match that moment, and omitting it can bias the results towards GTES.

1.6.3 Returns to scale parameter

The previous robustness section raises the question whether my results are sensitive to the estimate of the returns to scale parameter α. As a robustness check, I estimate the model for different fixed levels of α. The parameter estimates for this exercise are reported in Table A.310 in the Appendix. They are quite stable across specifications. Only the dispersion in quality σ_μ, which determines how much firms have to learn, is different. Table A.311 in the Appendix however shows that the conclusion - that each mechanism contributes roughly one third to the conditional investment decline (section 1.5) - does not change much, apart from when α is as high as 0.8.

1.7 Conclusion

This paper studies why firms gradually cut investment in the years after they go public, conditional on their profitability and size. There are three different causes, each explaining roughly one third of the decline. First, the combination of firms growing towards their efficient scale and increasing rigidity. Second, firms are learning about their efficient scale. Third, their shock volatility is decreasing over time. As a byproduct of the structural estimation, I also examine the causes of the underperformance of firms after they go public. I find that both mean-reversion in productivity and increase in size are important, validating the theories of Pástor et al. (2008) and Clementi (2002), respectively. I find that learning is an important determinant of Q conditional on size and profitability, confirming the interpretation of Pástor and Veronesi (2003). Previous researchers also interpret a declining sensitivity of investment to profitability as evidence of learning. My analysis cautions against this, as I find it is mostly driven by increasing rigidity.

In line with Asker et al. (2014), I find that firms become more rigid over time after they go public. There can be different reasons for this, for example short-termism (Holmström, 1999) or a restructuring into a more hierarchical and formal organization to exploit existing ideas (Ferreira et al., 2012). Future research could answer this question by taking parts of the current model and expanding it across those dimensions.