Venture Capital Investments and Learning over the Firm Life Cycle*

Fernando Riveiro Formoso†

June 15, 2023

Find the most recent version of the paper here.

Abstract

The life cycle of young, high-risk entrepreneurial projects and the financing of innovation has become increasingly important for economists, companies and policy-makers. The objective of this paper is to understand how high-risk firms learn over time about their unknown quality, and how this affects firm decisions and financing. I develop a model of the firm that imitates realistic features of young, high-risk companies, such as uncertainty about a firm's own quality, staged financing, exit strategies, and the realisation of period cash-flows that yield information about the firm's unobserved quality. The model captures empirical patterns of innovative firms documented in the venture capital literature – namely, delayed exit decisions and investments into companies being contingent on firm-level results over their life. I find that the ability to learn makes investment sensitive to period cash-flows in the model. A high initial quality uncertainty reflects into exit and investment strategies and may motivate firms to perform growth investments. In this context, a higher learning ability increases firm value substantially by motivating experimentation and contingent staged financing.

JEL codes: G24, D21, D25, M13, O16.

Keywords: staged financing, firm life cycle, investment under uncertainty, learning process, venture capital.

---

*I am grateful to my advisor Andrés Erosa for his advice and guidance, as well as Belén Jerez for her comments and help with this project. This paper has also benefited from discussions with Juan José Dolado, Matthias Kredler, Johannes Schneider, Felix Wellschmied and Emircan Yurdagul. I also thank participants in the Macroeconomics Reading Group, the Microeconomics Reading Group, and the PhD Student Workshop at Universidad Carlos III de Madrid for their helpful comments. All errors are mine. Financial support from Ministerio de Ciencia e Innovación (Spain), grant PRE2018-085625, is gratefully acknowledged.

†Universidad Carlos III de Madrid. Department of Economics. Calle Madrid 126, 28903, Getafe (Spain). Email: friveiro@eco.uc3m.es. Personal website: https://sites.google.com/view/friveiroformoso.

1 Introduction

The life cycle of firms and the dynamics of firm growth and financing have recently become of increasing interest to economists (Luttmer, 2011; Pugsley et al., 2021). In particular, a key issue is how financial resources should be allocated to young firms over their first years of activity, so that they grow and make sound contributions to the economy. This issue is crucial when we talk about innovative, high-risk entrepreneurial projects, whose underlying quality and future prospects are unknown at birth. The objective of this paper is to shed light on a feature of the firm life cycle that has not been studied in much detail in the context of high-risk firms and their financing: their ability to learn about the company's uncertain quality, as firm-level results are observed over time. This paper studies how learning affects life-cycle investment and exit decisions, as well as firm value, in contexts of high uncertainty.

I present a model of the firm that imitates realistic features of high-risk, innovative companies – namely, uncertainty about a firm's own quality, staged investments in time, exit decisions and firm-level results at different ages. In this one-agent model, the owner of an entrepreneurial project carries out investments over the firm life cycle until it decides to terminate the project or to sell it to the market. While injecting funds over time, the owner learns about the uncertain quality of her firm in a Bayesian manner. The learning process is possible because the agent receives period cash-flows that convey information about the true quality of the firm. Given the high degree of uncertainty about the project, this information serves to update beliefs about the firm's unobserved quality and thus affects investment and exit decisions of the agent.

The model is capable of capturing some firm-level regularities of innovative companies – namely, it replicates empirical patterns that have been documented by the venture capital literature. By *venture capital*, I refer to a type of financial intermediation consisting of a financier, the *venture capitalist*, that buys shares of a private company. In exchange, the company receives not only funding, but also, and importantly, monitoring, networking and expertise from the venture capitalist. Venture capital is focused on the growth of young firms, having as its final goal the exit of the venture capitalist by means of the sale of the shares previously acquired[1]. The venture capital industry in the United States, often considered the paradigm for a developed venture capital industry[2], has become an important vehicle for the financing of young, innovative firms[3].

[1]Usually via an initial public offering (IPO) or mergers and acquisitions (M&A).

[2]See Gompers and Lerner (2001) for a historical overview.

[3]For the sake of example, successful firms that received venture capital financing at different points of their life cycle, such as Cisco Systems, Apple, Google, Starbucks or Yahoo, are well-known.

Some of the practices and conventions within venture-backed companies have been documented by Kaplan and Strömberg (2003), among others. They report that investments made by venture capitalists are contingent on firm-level results or cash-flows, and thus made in a staged manner over the firm life cycle (staged financing). Additionally, exit strategies are carefully chosen by venture capitalists, and contracts used in deals are often sophisticated convertible securities. Some explanations for these practices have been proposed in the literature using principal-agent models to study moral hazard problems (Bergemann and Hege, 1998; Schmidt, 2003; Repullo and Suarez, 2004), control rights (Marx, 1998) and tax motives (Gilson and Schizer, 2003; Ollivierre, 2010). However, little is known about how the learning process inherent to innovative, high-risk projects affects the life cycle of venture-backed companies. In this paper, I abstract from the contracting problem between a financier and an entrepreneur and I rationalise firm-level patterns of venture capital investments, namely staged financing and exit strategies, by modelling a single agent that learns about her firm's uncertain returns over time.

I study the properties of the model and I arrive at two theoretical results. First, the possibility of learning from period cash-flows provides value to high-risk projects. In particular, I find that the ability to learn, jointly with the capability of terminating the project at every period, gives the agent an option value of waiting and updating her beliefs. This is possible if cash-flows are informative, to some extent, about the true quality of the project. Second, the ability to learn from these signals renders optimal investment decisions contingent on period cash-flow realisations, which is consistent with documented patterns at Kaplan and Strömberg (2003). Should we turn signals into completely uninformative ones (so that there is no learning), optimal investment would be unaffected by realised cash-flows.

I numerically solve the model and simulate it to better understand the implications of its theoretical properties. I find that investment and exit policies change over the life cycle of the firm, as more information is revealed and uncertainty decreases. Importantly, if the noise of the signal is sufficiently low, a higher degree of uncertainty results into a higher value from experimentation and a higher sensitivity of investment to cash-flow realisations. This in turn translates into a high positive contemporaneous correlation between simulated investments and cash-flows, in line with the empirical finding that fund injections and firm-level results go hand in hand. Next, I use the model to perform quantitative experiments. If we consider an environment in which learning is impossible, then cash-flow realisations do not provide information to the agent and there is no valuable belief updating. As a consequence, investment is completely insensitive to cash-flow realisations. Instead, if cash-flows are highly informative about project quality, this motivates the owner of a risky firm to continue running it and to carry out contingent investments. This translates into a greater value for firms, particularly when little is known about firm quality at birth. My findings are supportive of the idea that, when the degree of uncertainty about a firm's own quality is intense, a superior ability to learn about firm-level results over the life cycle renders the capacity to perform contingent growth investments valuable – thus being particularly interesting for innovative entrepreneurial projects.

Literature review. This paper relates to different areas of the economics and finance literature. First and foremost, it relates to the theoretical literature on venture capital. This strand of the economic literature aims to rationalise different practices in the venture capital industry and proposes alternative explanations on why venture capital is capable of increasing the value of firms receiving this type of financing. This literature points to agency problems (Bergemann and Hege, 1998; Schmidt, 2003; Cornelli and Yosha, 2003; Repullo and Suarez, 2004), the split of control rights between entrepreneurs and venture capitalists (e.g. the capability of terminating entrepreneurial projects) (Aghion and Bolton, 1992; Marx, 1998; Hellmann, 1998; Jovanovic and Szentes, 2013), and the expertise, the ability to screen entrepreneurial projects, and the reputation of venture capitalists (Ueda, 2004; Sørensen, 2007; Piacentino, 2019) as key determinants of contractual practices and investment and exit dynamics within venture capital projects. These mechanisms are proposed to explain the use of convertible securities, staged financing, and the prevalence of venture-backed firms in the IPO market.

Bergemann and Hege (1998) propose a model to study termination decisions in a context where entrepreneurs and venture capitalists are subject to moral hazard in the allocation of funds and learn about project returns in the event of exit. They conclude that pure equity contracts cannot maximise the net present value of projects since it generates inefficiently early termination. A convertible security structure that mixes debt and equity (both retained by the venture capitalist) and links debt to the liquidation value if the project is unsuccessful is an efficient contract in their setting. Jovanovic and Szentes (2013) also study the dynamic implementation of venture-backed projects as well as the way entrepreneurs and venture capitalists match and split rents from projects in a context in which the venture capitalist has an incentive to terminate a project if she is not sufficiently optimistic about its future prospects and move on to finance a new entrepreneur. Given competition among different venture capitalists, the optimal contract is a simple equity contract, and venture capitalists add value to projects by facilitating financing and monitoring. Although these two works are the main references of this paper, none of them explicitly models the role of intermediate results as signals that help agents engaging in high-risk projects to learn about future returns over their implementation and the way this may shape financing and exit decisions over time altogether. I contribute to this literature by proposing an explanation of staged financing contingent on intermediate results based on learning ability, which is in turn a characteristic of venture capitalists that is thought of as capable of increasing firm value.

Second, this paper relates to the growing macroeconomic literature on venture capital. There are many recent contributions to this literature aiming to assess the economy-wide impact of the venture capital industry. In order to generate substantial aggregate effects of venture capital, these works consider particular value-adding features of venture capital financing, namely the ability to attract superior entrepreneurial talent (Opp, 2019), the expertise in product development (Ateş, 2018), and the degree of assortative matching between entrepreneurs and capitalists (Akcigit et al., 2022). Differently from these works, this paper develops a detailed model of the firm with learning over the life cycle, as in Jovanovic (1982) and Guvenen (2007), to rationalise a superior form of fund injection that is contingent on the realisation of intermediate results and is value-enhancing at the firm level, and thus may add to previously explored explanations of the aggregate effects of the venture capital industry.

Finally, this paper is motivated by the empirical literature on venture capital finance. The seminal article in this field is Kaplan and Strömberg (2003), which is the first one to exhaustively document the contingent nature of venture capital investments. Following this seminal work, other papers such as Kaplan and Strömberg (2004), Lerner and Schoar (2005), Cumming (2008), and more recent surveys by Gompers et al. (2020) and Ewens et al. (2022) have nurtured this part of the finance literature by providing more empirical evidence on venture capital contracts and investments. Inspired by this evidence, I propose a learning explanation to rationalise some of these empirical patterns.

Layout. This paper is structured as follows. In section 2, I briefly revise trends and practices in the US venture capital market. In section 3, I present an investment model of the firm and I discuss theoretical properties of the model. In section 4, I perform a quantitative exploration of the model and I discuss the main results. In section 5, I conclude.

2 The US Venture Capital Industry: a Brief Overview

In order to motivate this paper, let us briefly discuss trends and recent results of the US venture capital industry, as well as firm-level practices within venture-backed companies.

Industry trends. Over the recent decades, the venture capital industry has become increasingly important in the United States, both in terms of resources allocated to it and economic outcomes from venture-backed firms. Metrick and Yasuda (2010) and NVCA (2022) provide data on capital raised by venture capital funds on a yearly basis. In 1983, total resources allocated to venture capital funds were $2.9 billion and accounted for 0.08% of the US gross domestic product. In 2021, these funds were $131.2 billion and represented 0.57% of US GDP. Between 1983 and 2021, funds allocated to venture capital have increased steadily, even during the world pandemic in 2020, with the only exception of the 2001 burst after the dotcom boom. Regarding summary output measures from the US venture capital industry, venture-backed firms broke several all-time records in 2021, as accounted by NVCA (2022). A record number of venture-backed firms, 14,411 companies, received a record amount of funding, $322 billion. In addition, the number of venture-backed exits (181 initial public offerings and 1,357 mergers and acquisitions) reached a historical peak as well. Importantly, the National Venture Capital Association also reports the relatively high prevalence of venture-backed firms in terms of employment: between 1990 and 2020, employment growth at venture-backed companies was eight times higher than that of non-venture-backed ones, and the seven largest publicly traded companies in 2021 have received venture capital financing at some point of their life cycle. As it has also been studied by Akcigit et al. (2022), these firms are disproportionately larger than their non-venture-backed counterparts in terms of employment.

Firm-level practices. The success of venture capital as a financing device for young, innovative firms in the United States is usually linked to the sophistication of venture capitalists. Practices at the firm level that prevail in the US venture capital industry were exhaustively documented by Kaplan and Strömberg (2003) in a seminal empirical study. Some of these practices are common to other venture capital markets, but others are distinct in Europe, Canada or developing countries (Cumming, 2008; Ollivierre, 2010; Lerner and Schoar, 2005). Here, I briefly describe some of these practices, namely staged financing, exit strategies and security types.

Staged financing consists in the injection of funds into firms being split into different periods of their life cycle. Staged financing has been of special interest to the theoretical literature, e.g Neher (1999) and Cornelli and Yosha (2003). According to the empirical study in Kaplan and Strömberg (2003), injection of funds of venture capitalists into firms is found to often be contingent in the realisation of favourable intermediate results or milestones within the company. The punchline of their study is that, within a venture-backed firm, cash-flow rights and other control rights that are important to the relationship between companies and financiers are separated and made contingent on different states of the world, given the high risk of entrepreneurial projects – namely, contingent on project performance[4]. This translates into venture capitalists deciding to inject money into high-risk venture-backed firms over different stages of their life.

Exit strategies refer to the decision of a venture capitalist to sell its shares of the company to the market, typically via an initial public offering (IPO) or a merger or an acquisition (M&A); or, alternatively, to liquidate the project if it turns out to be a failure. On one side, regarding the time path of successful venture-backed companies, Metrick and Yasuda (2010) report that the average sale period in the US venture capital market ranges between 3 and 7 years since the first venture capital investment takes place. This indicates that venture capitalists spend a large amount of time injecting funds and monitoring companies over their life cycle, and choose carefully when to bring the firm to the market. On the other side, a non-negligible fraction of venture-backed firms are terminated without a successful sale. In a recent survey by Gompers et al. (2020) made to venture capitalists and venture-backed firms in the United States, it has been documented that the average venture capital intermediary reports that 32% of its exits are failures (being 15% of its exits initial public offerings and 53% mergers or acquisitions). All this indicates the importance that an expected sale after some firm life periods has on within-firm decisions, but also the high frequency of failures and terminations among venture-backed companies.

Finally, Kaplan and Strömberg (2003) report the widespread usage of convertible securities in the US industry. This type of security can be roughly seen as a combination of debt and equity. Under convertible securities (also called convertible preferred stock), entrepreneurs and venture capitalists split the returns they receive from projects according to an equity share and the venture capitalist accumulates liquidation rights as she invests into the project. In the event of an exit or a liquidation, the venture capitalist can exercise the right to convert her liquidation rights into simple equity and get paid. This class of contracts are though of in the literature as having desirable efficiency properties, as in Bergemann and Hege (1998), Marx (1998) and Schmidt (2003). Importantly, they are overwhelmingly used in the US industry, to the point of a variety of classes of convertibles accounting for 95.8% of all investment rounds in the sample of Kaplan and Strömberg

[4]This contingency guideline also shows up in other characteristics of venture capital contracts. As other examples of contingency, and thus contract sophistication, venture capitalists may vest entrepreneurs with larger cash-flow rights as performance milestones are satisfied. The convertible security structure of contracts, in turn, allows venture capitalists to increase their cash-flow rights in case of exit or liquidation of the project as they invest funds into the firm. Also, investors have an implicit control right over the firm through their ability to stage investments in time. For more on this, see Kaplan and Strömberg (2003).

(2003). Interestingly, convertible securities have not been the traditional type of security in other countries, where simpler contracts are used very often in their national venture capital markets[5]. Thus, security sophistication has been indeed a special element of the US venture capital industry over the last years.

Some explanations for these practices have been proposed in the literature, e.g. agency problems (Bergemann and Hege, 1998; Schmidt, 2003; Repullo and Suarez, 2004), control rights (Marx, 1998) and tax reasons (Gilson and Schizer, 2003; Ollivierre, 2010). This paper relates widespread industry practices, namely staged financing and exit strategies, to a feature that is inherent to innovative, high-risk projects: the learning process over the life cycle about their true quality, which is uncertain at firm birth. Since these projects have high return uncertainty, they might require a specialist, (who may be, for instance, a venture capitalist), to learn about the unknown quality of the firm over the life cycle while injecting funds to grow the project towards a successful sale. I present a model in which a firm learns about its uncertain quality over the life cycle. My model is informative on how the ability to learn relates to contingent staged financing, exit strategies, and outcomes of high-risk firms.

3 A Model of the Firm

3.1 The Environment

In this section, I describe the problem of an agent that owns an entrepreneurial project, or a firm. The agent may live and invest in her firm during several periods. I refer to the life cycle of a firm as project implementation.

Physical environment. At a given point in time, an exogenous mass 1 of age-0 agents is born, each of them owning one firm. Let us first describe what agents know at the moment of firm birth (prior to $t = 0$). An agent, indexed by i, starts its life with a draw π_{i0} from a normal population distribution, $\pi_{i0} \sim N(\Gamma_0, \Sigma_0)$, where Γ_0 and Σ_0 are the exogenously given cross-sectional mean and variance of π_{i0}, respectively. Thus, firms are ex-ante heterogeneous in π_{i0}. Draw π_{i0} represents the quality of the entrepreneurial project – that is, the capacity of the firm to generate cash-flows. Importantly, I assume

[5] Cumming (2008) considers a sample for the European industry and reports that only 32.3% of all investments were using some class of convertible security – indeed, it is common stock that is more generalised in European venture capital markets. A similar pattern has been observed in Canada (Ollivierre, 2010). In developed countries, less sophisticated securities such as common stock or even plain debt contracts are prevailing due to legal constraints that impede enforcement of more complicated securities (Lerner and Schoar, 2005).

that agents do not observe their initial draw π_{i0} at birth, and it remains unobserved over the life cycle. Nevertheless, agents know the population distribution of initial quality. Thus, the cross-sectional distribution of π_{i0} determines initial beliefs of agents. At age 0, every firm i believes that its unobserved quality is distributed according to a normal distribution with mean $\hat{\pi}_{i0} = \Gamma_0$ and variance $\sigma^2_{i0} = \Sigma_0$. All firms in the economy have the same initial beliefs, although they are heterogeneous in their underlying unobserved quality draw[6]. This is equivalent to assuming that an agent owning a firm i is born with unobserved initial quality π_{i0}, and a duple $(\hat{\pi}_0, \sigma^2_0)$ representing initial prior beliefs that is equal for all firms. I take this bidimensional object as a primitive of the model.

I study the firm life cycle after firm birth. Time is discrete and infinite, $t = 0, 1, ...$, agents are risk neutral and discount future payoffs at an exogenous discount rate $r > 0$. In what follows, I focus on a specific firm owner and thus I omit subscript i for expositional purposes. At an age $t \geq 0$, an agent owning a firm enters the period with a prior belief $(\hat{\pi}_t, \sigma^2_t)$ about her firm's quality. The agent may decide to continue running her project or, alternatively, liquidate it or sell it to the market. If she decides to keep her firm at t, she observes period cash-flows, which are informative about the true quality of the project and thus enable the agent to update her prior beliefs, as we shall see. Once period cash-flows realise, the agent can invest in her firm to improve its unobserved quality.

Within an age $t \geq 0$ of an agent's life, there are four stages, in chronological order:

1. **Exit decision**: the agent decides whether to keep her firm, to terminate it, or to sell it to the market in exchange of a price (*keep/termination/sale*). Upon termination or sale, the agent leaves the economy.

2. **Cash-flow realisation**: upon keeping the firm, the agent receives period cash-flows.

3. **Belief updating**: given the cash-flow realisation, prior beliefs are updated.

4. **Investment decision**: after updating beliefs, the agent chooses how much to invest in her firm. The unobserved quality of the project evolves accordingly.

At the end of age t, the agent ends up with posterior beliefs $(\hat{\pi}_{t+1}, \sigma^2_{t+1})$.

Exit and investment decisions. Agents make exit and investment decisions over the life cycle. At the beginning of age $t \geq 0$, the agent makes a discrete decision, either to

[6] These assumptions are made to represent an innovative, cutting-edge industry composed of firms whose true quality is unknown at the moment of birth and may differ across firms. In this environment, projects are inherently risky and no firm has privileged information about its true quality.

cease the activity of the firm (*termination*), to sell the firm to the market (*sale*), or to continue running the firm (*keep*). If the agent chooses *termination* at age t, the firm stops its activity, the agent leaves the economy and she gets a value of zero. If the agent chooses *sale*, the firm is sold and the agent gets the discounted expected value of future cash-flows generated by the firm in exchange, and leaves the economy. The sale process is costly, and these costs are represented by an amount $C_{IPO} > 0$ that the agent has to pay if she chooses *sale*. Importantly, in the event of a sale or a termination, no further investments are made in the project. The original firm owner is the only individual that is capable of exerting effort to make the firm grow. Upon a *keep* decision at period t, the firm observes the realisation of a random variable that I label period cash-flows, $CF_t \in \mathbb{R}$, which represents intermediate revenues or results taking place during the firm life cycle. After the realisation of CF_t, the agent makes a non-negative investment, k_t. Investment is worthy to the firm in that it allows to increase the unobserved quality of the firm at age t, π_t. More specifically, π_t increases by $B > 0$ per unit of investment. Keeping the project and exerting investments has costs. First, there is a fixed cost of operation c_{op} to be paid whenever the agent chooses *keep*, regardless of the level of investment. Second, the agent pays a price p_K per unit of investment good bought. Third, there is an increasing and convex investment cost $c(k) = \frac{\psi}{2}k^2$, whose intensity is parameterised by $\psi > 0$.

Learning process. At birth, the agent does not observe the true quality of the project. However, over the life cycle, as she chooses to keep her project, she receives period cash-flows that are informative about its unobserved quality. Let us focus on the learning process that characterises project implementation, aiming to represent the dynamic experimentation that is inherent to high-risk, innovative projects. I do this by assuming that the firm updates its prior beliefs about project quality in an optimal, Bayesian manner, expressed as a Kalman filtering problem (Kalman, 1960). At birth, the firm has an initial prior belief about the distribution of π_0 that is assumed to be normal, $N(\hat{\pi}_0, \sigma_0^2)$. Under the Kalman filter setting, if the variable whose distribution we are updating is normal at any period t, the updated distribution at period $t+1$ is also normal. Therefore, the learning process describes the evolution of the duple $(\hat{\pi}_t, \sigma_t^2)$.

Two key ideas underlie belief updating. First, period cash-flows CF_t observed upon a *keep* decision are an imperfect signal about the true quality of the project, π_t. This is reflected in the following "observation equation", which represents the observable variable CF_t as a linear expression in the unobservable quality π_t and a transitory shock ε_t, which can be understood as a measurement error. The observation equation reads:

$$CF_t = \pi_t + \varepsilon_t \tag{1}$$

where $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$. The variance of ε_t, σ_ε^2, is non-negative and represents the intensity of the measurement error. If $\sigma_\varepsilon^2 > 0$, the firm is prevented from knowing with certainty whether an observed CF_t was due to its project's quality π_t, or just luck. Second, I assume that the unknown quality evolves in time according to a linear "law of motion equation". The level π_{t+1} is affected by the investment exerted at period t, so that it helps to improve the project's quality, and by its lagged value π_t. This allows us to generate persistent effects of investment in the model. The law of motion for project quality is:

$$\pi_{t+1} = \pi_t + Bk_t \tag{2}$$

where B is strictly greater than zero. The law of motion of uncertain quality makes it explicit that firm growth is only possible if effort is put at a cost. Equations (1) and (2) determine how agent's beliefs evolve over the firm life cycle. We can use equation (1) to get the distribution of CF_t conditional on our prior quality distribution, at a period t. This conditional distribution has cdf $G_t(CF_t|\hat{\pi}_t)$ and, similarly to the prior distribution, it is normal: $CF_t \sim N(\hat{\pi}_t, \sigma_t^2 + \sigma_\varepsilon^2)$. In order to derive equations for the evolution of expected quality $\hat{\pi}_t$ and quality uncertainty σ_t^2, we apply the Kalman filter. As it is discussed in Kalman (1960) and Perla et al. (2022), the Kalman filter consists in recursively finding a predictor for the unknown variable using observed imperfect measures of it, thus yielding laws of motion for its expected value and variance. This form of representing learning is becoming increasingly used in the macroeconomics literature[7]. In our context, the Kalman filter yields a point estimate for the unknown quality, which is the conditional expectation for π_t, given past observations for cash-flows and investment injections. This estimate is the mean squared error minimiser among all Borel functions on $\mathbb{R}$ with bounded variance[8]. The derivation of the Kalman filter equations is discussed in Appendix A.1. Under imperfect observability of the true quality of the project, the evolution of beliefs $(\hat{\pi}_t, \sigma_t^2)$ is characterised by the following Kalman filter equations:

$$\hat{\pi}_{t+1} = \hat{\pi}_t + Bk_t + \kappa_t(CF_t - \hat{\pi}_t) \tag{3}$$

[7]See Guvenen (2007), Baley and Veldkamp (2021) and Farboodi and Veldkamp (2021) for applications of the Kalman filter in macroeconomics.

[8]In fact, it can be shown that predictions using the Kalman filter when the true state π_t is unobservable perform relatively well if we compare them with those from the optimal predictor (in terms of minimising the squared error) $\mathbb{E}[\pi_{t+1}|\pi_t]$ under perfect observability of π_t, which is the predictor a rational agent would use. See Perla et al. (2022) for a simple application of this idea.

$$\sigma_{t+1}^2 = (1-\kappa_t)\sigma_t^2 \tag{4}$$

and

$$\kappa_t = \frac{\sigma_t^2}{\sigma_t^2 + \sigma_\varepsilon^2} \tag{5}$$

where κ_t is the so-called Kalman gain. The Kalman gain provides an idea on how much we update the distribution mean from realisations of period cash-flows, or the sensitivity of $\hat{\pi}_{t+1}$ to different magnitudes of the observed CF_t, as it is reflected in equation (3). Importantly, it depends positively on quality uncertainty σ_t^2 and negatively on the intensity of the measurement error σ_ε^2: the updating of $\hat{\pi}_t$ is more intense when quality is very uncertain and when the imperfect measure is more accurate. At any period t, the three Kalman filter equations (3)-(5) give us the duple $(\hat{\pi}_{t+1}, \sigma_{t+1}^2)$ that characterises the updated quality distribution. Notice that the dynamics of the Kalman gain κ_t are fully determined by the dynamics of quality uncertainty σ_{t+1}^2. The Kalman gain, in turn, affects the dynamics of both the distribution mean and variance.

Value functions. Let us now discuss the problem faced by an agent that owns a firm, as well as optimal exit and investment policies. The agent chooses the investment amounts and the exit strategies to maximise the discounted expected value from her firm. There are no frictions, e.g. informational asymmetries, conflicts of interest, or hold-up problems.

First, in order to determine the value the agent gets in the event of a sale, it is necessary to make a distinction between a firm that is held by the agent and a firm that is held by the market (i.e. it has been previously sold) at age t. The key difference between these two firms is that the market is unable to perform productive investments to increase the unobserved project quality, while the agent is capable of doing so. Additionally, the agent has to pay sale costs in the event of a sale as well as period operation and investment costs, while the market is free of those. I assume, however, that both the agent and the market have the ability to terminate projects, and that, upon not terminating, projects generate period cash-flows and beliefs are updated accordingly; that is, both have the ability to learn.

Consider a firm with beliefs $(\hat{\pi}_t, \sigma_t^2)$ that has been sold to the market at some previous age. The value of an age-t firm held by the market is:

$$W_t(\hat{\pi}) = \max\{0, M_t(\hat{\pi})\} \tag{6}$$

where $M_t(\hat{\pi})$ is the expected value prior to the cash-flow realisation at t:

$$M_t(\hat{\pi}) = \int \big(CF + \frac{1}{1+r} W_{t+1}(\hat{\pi}')\big) dG_t(CF \mid \hat{\pi}) \tag{7}$$

In this expression, $CF = \pi + \varepsilon$, and beliefs evolve according to equations (3)-(5) – namely, $\hat{\pi}' = \hat{\pi} + \kappa_t(CF - \hat{\pi})$, since no investment can be exerted by the market-held firm. If we know market values W_t and M_t, we are able to determine the value of a firm that is kept by its original owner, who has the ability to make period investments. The value of an age-t firm held by an agent with beliefs $(\hat{\pi}_t, \sigma_t^2)$ is:

$$V_t(\hat{\pi}) = \max\{ \underbrace{0}_{\text{termination}}, \underbrace{M_t(\hat{\pi}) - C_{IPO}}_{\text{sale}}, \underbrace{\int U_t(\hat{\pi}, CF) dG_t(CF|\hat{\pi})\}}_{\text{keep}} \tag{8}$$

where the value of keeping the project and receiving a cash-flow realisation CF is:

$$\begin{aligned} U_t(\hat{\pi}, CF) = & \max_{k \geq 0, \hat{\pi}' \in \mathbb{R}} \left\{ CF - c(k) - p_K k - c_{op} + \frac{1}{1+r} V_{t+1}(\hat{\pi}') \right\} \\ & \text{s.t.} \qquad \hat{\pi}' = \hat{\pi} + Bk + \kappa_t(CF - \hat{\pi}) \end{aligned} \tag{9}$$

and also subject to the exogenous evolution of quality uncertainty and the Kalman gain in equations (4) and (5), given initial beliefs. In words, at age t, the agent can sell the firm to the market, in which case she pays a sale cost C_{IPO} and gets the market value of the firm $M_t(\hat{\pi})$ – that is, the expected discounted value of future cash-flows[9]. By selling the project, the agent will be saving investment and operation costs. Otherwise, she can keep the project, which allows her to exert investments that increase its profitability. In the meanwhile, the agent updates her beliefs as new cash-flows arrive. Regarding belief updating, the evolution of the expected quality $\hat{\pi}$ is endogenous, for it depends on the investment decision of the agent. On the other side, notice that, given σ_0^2, equations (4) and (5) evolve exogenously and in a deterministic way. Time is a state in equations (6)-(9) (and thus the time index t) because both quality uncertainty σ_t^2 and the Kalman gain κ_t are strictly positive and vary over the firm life cycle due to belief updating, provided $0 < \sigma_\varepsilon < \infty$. Nevertheless, if we let t increase, σ_t^2 and κ_t decrease in t and infinitely old firms (either held by their initial owners or by the market) have $\sigma_\infty^2 = \kappa_\infty = 0$, thus facing a stationary problem where there is no belief updating.

We want to find the optimal investment and exit decisions made by the agent over the firm life cycle. These are investment policies $g_t^k(\hat{\pi}, CF) \in [0, \infty)$, as well as the discrete exit strategy $g_t^{exit}(\hat{\pi}) \in \{termination, sale, keep\}$. To find these policies, I solve the model

[9] We have to take into account that the market still has the option to terminate the firm, so the value of the market-held firm is bounded below.

numerically using value function iteration. Further details on the algorithm used to find the policy functions are presented in Appendix A.2. In the numerical solution, whose properties are discussed in the next section, the optimal exit policy for the agent $g_t^{exit}(\hat{\pi})$ is characterised by *expected quality thresholds* for $\hat{\pi}_t$, $\{\underline{\hat{\pi}}_t, \overline{\hat{\pi}}_t\}_{t=0}^{\infty}$ with $\underline{\hat{\pi}}_t < \overline{\hat{\pi}}_t$, such that the agent terminates the project at age t if $\hat{\pi}_t < \underline{\hat{\pi}}_t$ (when she thinks the project does not have enough quality), sells the project to the market at age t if $\hat{\pi}_t > \overline{\hat{\pi}}_t$ (when she thinks the project has enough quality), and keeps the project otherwise, for expected qualities lying in between the upper and lower thresholds[10]. By backward induction, we can eventually find the value that the agent gets from owning the firm at birth, $V_0(\hat{\pi}_0)$, taking as given the initial prior belief.

3.2 Learning, Investment and Exit Decisions

Previous to the quantitative exploration of the model solution, it is necessary to understand the mechanisms operating in the model presented above. In particular, it is paramount to study how *learning* affects exit choices and firm investment over the life cycle. This section illustrates the role of learning. A *learning* scenario is one where the firm receives period cash-flows that are partially informative about the true quality of projects, so that $\sigma_\varepsilon^2 > 0$, and beliefs $(\hat{\pi}_t, \sigma_t^2)$ are updated as cash-flows realisations arrive. Belief updating is possible as far as the signal is not too noisy, i.e. if $\sigma_\varepsilon^2 < \infty$. Otherwise, I will be talking about a *non-learning* scenario, where beliefs are not updated over time.

3.2.1 Learning Affects Exit Decisions

Let us first illustrate how belief updating affects discrete decisions. To do this, I abstract here from investment and I consider a firm that has been sold to the market, which does not have the ability to invest[11]. This firm is capable of updating beliefs over time and make termination decisions.

First of all, timely differences in belief updating are mapped to timely differences in σ_t^2 and κ_t, whose exogenous evolution affects the way cash-flow realisations generate surprises and change expected quality $\hat{\pi}_t$ over time. Figure 1 shows how the standard deviation σ_t and the Kalman gain κ_t evolve in time according to equations (4) and (5), departing from an initial belief of $\sigma_0 = 900$. As we can see, both σ_t and κ_t have a decreasing profile in time, and they converge to zero as time goes to infinity. As a consequence, belief updating

[10]Similarly, market-held projects have an optimal termination policy characterised by an expected quality threshold, such that projects are terminated if $\hat{\pi}_t$ lies below a cutoff.

[11]Considering investment and the problem of the original firm owner does not change the main take-away of this section regarding learning, value and discrete decisions.

Figure 1: Belief updating: standard deviation and Kalman gain

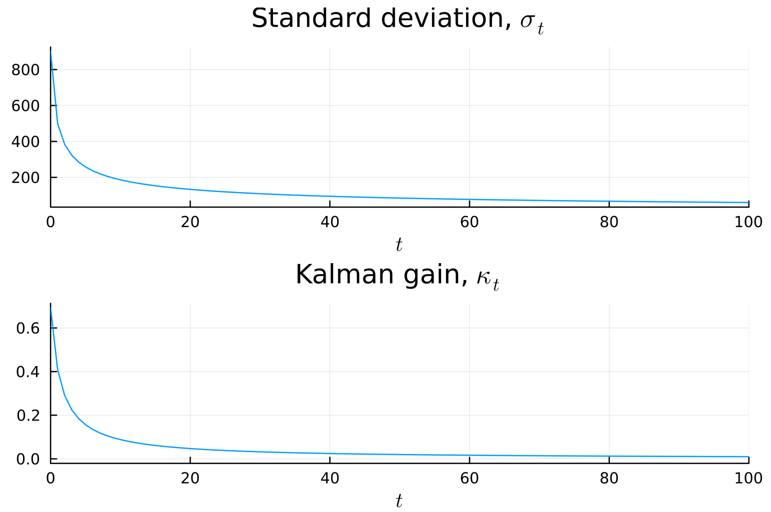

Notes: the two graphs respectively show the life-cycle evolution of the standard deviation characterising beliefs of a firm, σ_t, and the corresponding Kalman gain, κ_t, for 100 periods of life. The standard deviation and Kalman gain profiles have been calculated using equations (4) and (5), respectively. I assume an initial standard deviation of $\sigma_0 = 900$ and a measurement error of $\sigma_\varepsilon = 600$.

will be less notorious for old firms than for young firms, since quality uncertainty fades away with age. Indeed, let us discuss first what happens to infinitely old firms in terms of value. Consider the stationary value of a firm that has been sold to the market. Given $\sigma_\varepsilon^2 > 0$ and equations (4) and (5), if we let time go to infinity we get $\kappa_\infty = \sigma_\infty^2 = 0$. Thus, infinitely old firms know with certainty their true quality (that is, $\hat{\pi}_\infty = \pi_\infty$) and do not update their beliefs. As a consequence, the stationary value of a firm of (expected) quality $\hat{\pi}$ is:

$$W_\infty(\hat{\pi}) = \max\{0, \frac{\hat{\pi}}{1-\beta}\} \tag{10}$$

where $\beta = 1/(1+r)$. Figure 2(a) plots the stationary value $W_\infty(\hat{\pi})$ for different values of $\hat{\pi}$. From the graph, it is clear that the optimal exit decision of the stationary firm consists of terminating if its quality lies below zero, and keep on receiving cash-flows forever otherwise.

Let us now see what happens over the life cycle of a firm that has been sold to the market. At ages $t < \infty$, both the quality uncertainty σ_t^2 and the Kalman gain κ_t are strictly positive, as far as the noise of the signal does not go to infinity. As a consequence, there is a non-trivial belief updating over time. In particular, expected quality by the firm evolves according to $\hat{\pi}_{t+1} = \hat{\pi}_t + \kappa_t(CF_t - \hat{\pi}_t)$. The value of this age-$t$ firm is given by equation (6). Figure 2(b) shows this value function for different firm ages (namely,

Figure 2: Market value: stationary and life cycle

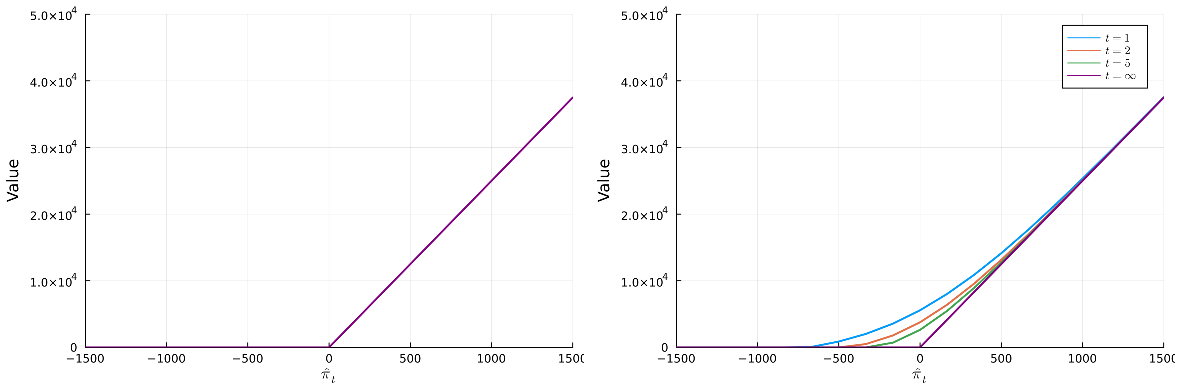

(a) Stationary value

(b) Life-cycle value

Notes: subfigure (a) plots the stationary value of a firm that has been sold to the market in equation (10), as a function of the expected quality of the firm. Subfigure (b) compares this stationary value to different life-cycle values (for ages 1, 2 and 5) of a firm that has been sold to the market in equation (6). Life-cycle values are again plotted as a function of expected quality.

ages 1, 2 and 5). Differently from the stationary function, value has a curvy, strictly convex shape. Importantly, at younger ages, value is weakly higher than the stationary value, and it decreases as the firm ages. The differences in value over the firm life cycle ultimately come from differences in belief updating over time, which translate into timely differences in the Kalman gain.

To illustrate the effects learning has on value, let us consider an age-0 firm whose expected value is slightly below zero, e.g. $\hat{\pi}_0 = -100$, and let us compare it with an age-∞ firm with $\hat{\pi}_\infty = -100$. As we see from Figure 2(b), the age-0 firm is getting a positive value and decides not to terminate the project at age 0. In turn, the infinitely old firm with $\hat{\pi}_\infty = -100$ decides to terminate the project and gets a value of zero. The difference between these two firms is that the young firm has room for belief updating. At age 0, project quality is yet pretty uncertain, since σ_0^2 is high. Consequently, the Kalman gain at age 0 is also large, given equation (5). Therefore, choosing to keep the project alive and to observe a new cash-flow observation CF_0 is going to yield an optimistic belief update (i.e. an increase in $\hat{\pi}_0$) with a very high probability.

Still, there is also a high probability that CF_0 is low, and that beliefs are updated pessimistically. However, the possibility to terminate generates a positive option value. Should the firm receive a very bad cash-flow realisation, it can avoid receiving a time-lasting stream of poor results by just leaving the economy. Put simply, young firms can cut negative cash-flow streams if they expect them, and so do old firms; but, differently from old firms, there is room for positive belief updating for young firms, and thus for attaining higher future cash-flows. As a result, the value of an age-0 firm can be positive even if it starts its life expecting a relative poor quality. This is entirely due to the possibility of learning under uncertainty. Indeed, if we consider a non-learning environment such that period cash-flows are completely uninformative, then $\kappa_t = 0$ for all t, no belief updating is possible, and firm value at ages $t < \infty$ coincides with the stationary value shown in Figure 2(a).

3.2.2 Learning Affects Investment

Let us now come back to the problem of the agent who owns a firm and has the ability to perform growth investments on it. Let us illustrate how belief updating over the life cycle shapes investment decisions. For illustration purposes, I discuss a simplified version of the model. The results from this model extend to the more general setting, which is simulated in section 4. I find that learning makes investment by the agent contingent on intermediate cash-flows, in line with documented empirical patterns discussed in section 2.

Here, I present a three-period version of the model of the firm. There is an agent that owns a project and lives for three periods: $t = 0, 1, 2$. Initial beliefs at age 0 are described by a duple $(\hat{\pi}_0, \sigma_0^2)$, and may be updated over time. At period 2, the agent chooses whether to terminate the project or to sell it to the market, leaving the economy right after in either case. I assume that the discrete decision only takes place at age 2, and not before, but the mechanism discussed here is also present in a model with discrete decisions every period, as shown in section 4. If the agent decides to sell the project to the market, she gets the expected value of period cash-flows at age 2 – that is, $\hat{\pi}_2$. For exposition purposes, I abstract from operation and sale costs. Leaving these considerations aside, the timing of events is similar to that in section 3.1.

Let us first discuss the non-learning case – i.e. cash-flows are completely uninformative about the true quality of projects, and prior beliefs are never updated. An implication of this is that the Kalman gain equals zero at every period $t = 0, 1, 2$ of firm life. At period 2, the firm with expected quality $\hat{\pi}_2$ has value $V_2(\hat{\pi}_2) = \max\{0, \hat{\pi}_2\}$. Thus, in period 1 we have:

$$V_1(\hat{\pi}_1) = E_{\varepsilon_1}[U_1(\hat{\pi}_1, CF_1)]$$

where

$$U_1(\hat{\pi}_1, CF_1) = \max_{k_1 \geq 0} \left\{ CF_1 - \frac{\psi}{2} k_1^2 - p_K k_1 + \frac{1}{1+r} \max\{0, \hat{\pi}_2\} \right\}$$

subject to equation (3), which reads $\hat{\pi}_2 = \hat{\pi}_1 + Bk_1$ in this case, and where $CF_1 = \pi_1 + \varepsilon_1$. Recall that fund injection k_1 is chosen after cash-flow CF_1 is realised. I adopt a backward induction perspective. Conditional on the agent choosing to terminate at period 2, then she gets value $CF_1 - \frac{\psi}{2}k_1^2 - p_K k_1$ and she optimally chooses $k_1^T = 0$. Otherwise, conditional on choosing sale at period 2, the firm gets value $CF_1 - \frac{\psi}{2}k_1^2 - p_k k_1 + \frac{1}{1+r}(\hat{\pi}_1 + Bk_1)$ and chooses $k_1^S = \frac{1}{\psi}(\frac{1}{1+r}B - p_K)$, which is greater than zero if B is large enough. The sale choice takes place whenever $\frac{1}{1+r}(\hat{\pi}_1 + Bk_1^S) \geq 0$ or, equivalently, when $\hat{\pi}_1 + \frac{B}{\psi}(\frac{1}{1+r}B - p_K) \geq 0$. The optimal investment decision in period 1 is:

$$k_1^*(\hat{\pi}_1) = \begin{cases} \frac{1}{\psi}\big(\frac{1}{1+r}B - p_K\big) & \text{if } \hat{\pi}_1 + \frac{B}{\psi}\big(\frac{1}{1+r}B - p_K\big) \geq 0, \\ 0 & \text{otherwise} \end{cases}$$

Notice that optimal investment is weakly increasing in $\hat{\pi}_1$: the fact that the firm enters the period with a high expected quality may alter the discrete decision towards selling instead of terminating, and thus may generate a positive jump in investment. Indeed, positive investment will take place for any expected quality such that $\hat{\pi}_1 \geq -\frac{B}{\psi}\big(\frac{1}{1+r}B - p_K\big)$. Provided that fund injections have a positive marginal productivity B that is sufficiently large relative to the price of the investment good, the firm would be willing to invest and sell the project right after even for some negative values of $\hat{\pi}_1$. The same logic discussed here can be extended to period 0 to obtain $k_0^*(\hat{\pi}_0)$, also weakly increasing in $\hat{\pi}_0$. Nevertheless, staged financing is not contingent on period cash-flows – i.e CF_t does not affect optimal investment k_t^*.

To arrive at this important result, let us now consider a bounded and sufficiently low σ_ε^2, so that there is *learning*, i.e. belief updating over time. In this case, the investment problem at period 1 is similar to that in the non-learning case, but the Kalman gain κ_1 is strictly positive and equation (3) reads $\hat{\pi}_2 = \hat{\pi}_1 + Bk_1 + \kappa_1(CF_1 - \hat{\pi}_1)$, where the extra term depends on CF_1, and on σ_1^2 and σ_ε^2 via the Kalman gain. Following the same logic as before, optimal investment at period 1 is:

$$k_1^*(\hat{\pi}_1, CF_1) = \begin{cases} \frac{1}{\psi}\big(\frac{1}{1+r}B - p_K\big) & \text{if } \hat{\pi}_1 + \frac{B}{\psi}\big(\frac{1}{1+r}B - p_K\big) + \kappa_1(CF_1 - \hat{\pi}_1) \geq 0, \\ 0 & \text{otherwise} \end{cases}$$

Now, the fact that $\kappa_1 > 0$ makes both $\hat{\pi}_1$ and CF_1 capable of generating jumps in investment due to the discrete decision. In particular, the firm will invest and sell the project afterwards if $\hat{\pi}_1 \geq -\frac{1}{1-\kappa_1}\big(\frac{B}{\psi}(\frac{1}{1+r}B - p_K) + \kappa_1 CF_1\big)$. For example, imagine κ_1 is large. This can be caused either by a very uncertain environment (high σ_1^2) or by a

Table 1: Parameterisation of baseline model

Parameter	Definition	Value
Initial beliefs		
$\hat{\pi}_0$	Expected quality at birth	-50
σ_0	Quality uncertainty (standard deviation) at birth	900
Other parameters		
r	Discount rate	0.042
B	Marginal productivity of effort	2
σ_ε	Standard deviation of measurement error	600
ψ	Intensity of convex investment cost	4.5
p_K	Price of investment good	1
c_{op}	Fixed operation cost	400
C_{IPO}	Fixed sale cost	3500

very high learning ability (low σ_ε^2), or both. In this high-κ environment, a negative CF_1 realisation may induce the firm to choose to terminate the project instead of selling it at period 2, and thus may prevent it from investing, even if $\hat{\pi}_1$ is positive and large enough. We can extend this logic to period 0 (although more unpleasantly from an algebraic point of view) and get a similar intuition for $k_0^*(\hat{\pi}_0, CF_0)$. The key insight from this illustrative model is that, jointly with the possibility of terminating or selling the project, *learning is a force that makes investment contingent on period cash-flows.*

As we see next, the mechanism presented here also operates in the infinite-periods model of section 3.1. The ability to learn about firm-level outcomes generates interesting patterns of staged financing, a phenomenon that has been of great interest to the venture capital literature (Neher, 1999; Cornelli and Yosha, 2003). Specifically, learning is shown to be a theoretical mechanism capable of rationalising contingent fund injections, an important feature of real-life US venture-backed firms, as documented by Kaplan and Strömberg (2003).

4 Quantitative Exploration

In this part of the paper, I discuss the numerical solution of the complete, infinite-periods model described in section 3.1, given the baseline parameterisation in Table 1. The theoretical mechanisms described in section 3.2, through which learning affects exit and investment decisions, are operating in the model. In my baseline parameterisation, I consider that agents owning firms start out project implementation with initial beliefs $\hat{\pi}_0 = -50$ and $\sigma_0 = 900$. The assumption on initial expected quality $\hat{\pi}_0$ guarantees that if there were no initial quality uncertainty ($\sigma_0 = 0$), keeping the rest of parameterisation equal, the agent would immediately terminate the project in the first period of firm life. On the other hand, the assumption that $\sigma_0 = 900$ aims at representing a quite intense initial quality uncertainty. I assess whether quality uncertainty is sufficiently high to motivate the agent to keep the firm, to update her beliefs about the project's quality, and to perform growth investments. In particular, I study whether learning is beneficial for the agent in terms of the value of the firm, as well as other simulated moments of interest.

The baseline parameterisation in Table 1 has been chosen for the model to display two reasonable decision-making patterns of agents owning entrepreneurial projects. First, I parameterise the dynamic model such that an infinitely old agent would always decide not to keep the project. That is, the stationary value of an agent of age $t = \infty$ is such that the agent either sells the project to the market, if she knows that the project is profitable, or terminates it otherwise[12]. The parameterisation allows to avoid a situation in which agents keep their projects forever and growth investments are made eternally. This would be at odds with documented venture capital industry facts, where the objective of venture capitalists is precisely to exit the firm after some years, either by cashing out if the firm is profitable or by terminating it if it turns out to be a failure (Metrick and Yasuda, 2010). Second, model parameters are chosen such that, over the first years of the agent's life cycle, the firm is kept by the owner for a non-empty set of initial beliefs – i.e. there is room for the agent to keep the project, invest and learn, and she may not sell or terminate immediately. Again consistently with facts in Metrick and Yasuda (2010), this aims to capture the experimentation process venture capitalists incur when they start providing resources and guidance to high-risk firms.

4.1 Exit and Investment Decisions

Exit policy over the life cycle. Consider the numerical solution for the baseline model in Table 1. Let us first discuss exit policies of agents owning projects in this environment. In Figure 3, we observe the optimal termination threshold $\underline{\hat{\pi}_t}$ and the optimal sale threshold $\overline{\hat{\pi}_t}$ for an agent over the firm life cycle. As we observe, as her firm gets old, the agent optimally keeps it for a smaller set of expected qualities $\hat{\pi}_t$. The upper threshold $\overline{\hat{\pi}_t}$ is decreasing in t, so that the agent is more eager to sell the project to the market as it gets old given a level of $\hat{\pi}_t$. The lower threshold $\underline{\hat{\pi}_t}$ is increasing in t, implying that, for the same level of expected quality, the agent is more willing to terminate the project as age increases.

[12]Regarding implied discrete choices, this stationary value is similar to that shown in Figure 2(a).

Figure 3: Sale and termination thresholds, baseline model

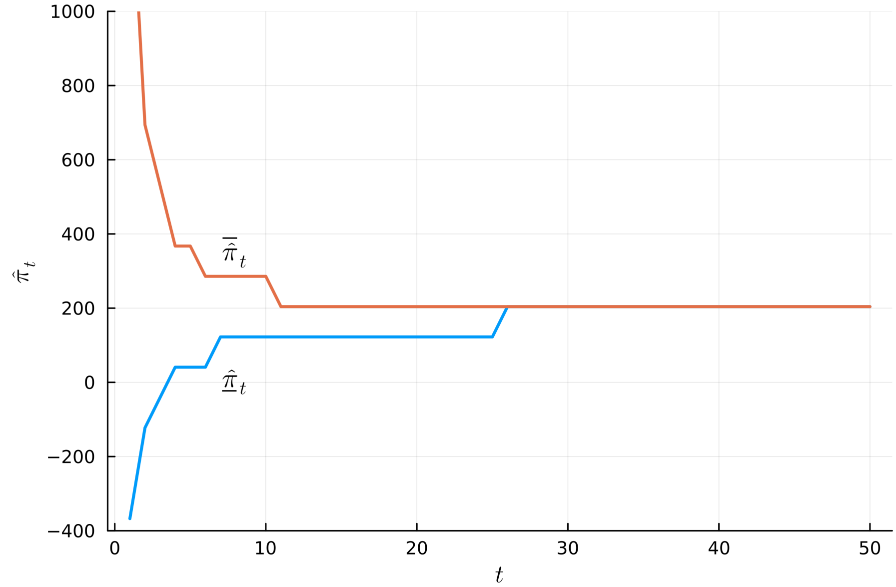

Notes: the graph shows optimal sale $\overline{\hat{\pi}}_t$ and termination $\underline{\hat{\pi}}_t$ thresholds for expected quality $\hat{\pi}_t$ for different periods of an agent's life cycle. The lines are obtained from the numerical solution of the optimisation problem of the agent, given the parameterisation in Table 1. The orange line corresponds to the sale threshold $\overline{\hat{\pi}}_t$ and displays a decreasing pattern in time. The blue line corresponds to the termination threshold $\bar{\hat{\pi}}_t$ and is increasing in time.

Figure 3 shows that agents never keep their firms when they are older than a certain age, under parameter values in Table 1. For a very large t, the sale and the termination threshold coincide. This is due to the fact that σ_t^2 and the Kalman gain are zero for infinite ages, as implied by equations (4) and (5). When firms are very old, they already know their true quality pretty accurately, and cash-flow surprises do not lead to strong updates in their expected quality. For infinitely old agents, similarly to Figure 2(a), the ability to learn does not provide any option value from not terminating the project. This leads them to either terminate or sell, but they never keep and invest.

Figure 4 represents the investment decision of an infinitely old agent ($t = \infty$) that has a sufficiently high expected quality $\hat{\pi}_t$ and corresponding value $V(\hat{\pi}_t)$. This agent does not get any option value from keeping the project, and will only keep it if her investments are sufficiently profitable. As in Figure 2(a), value is linear for high expected qualities, and thus its derivative does not vary with $\hat{\pi}_t$. If the agent has already decided to keep her project, the marginal cost of investment $p_k + c'(k)$ is lower than the marginal value from investing, for small amounts of k, and optimal investment is a fixed positive amount (k^* in Figure 4). Thus, provided that $c(k)$ is convex, the old firm would thus find profitable

Figure 4: Old agents, investment decisions

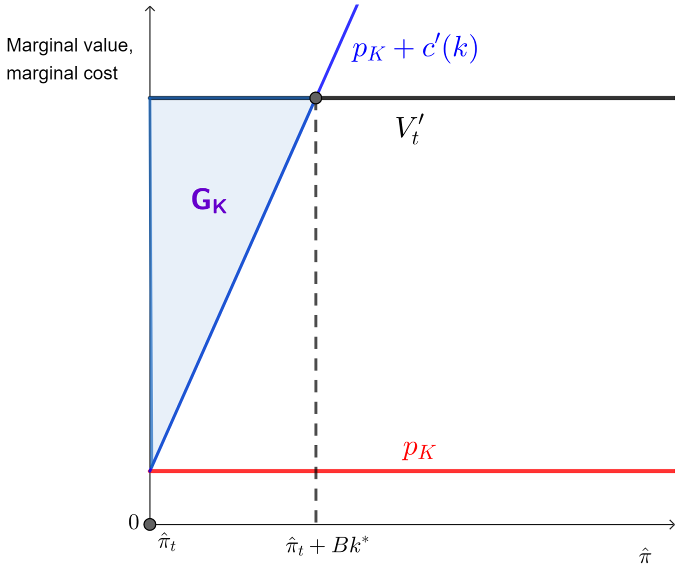

to invest a positive amount[13]. The benefit from investing and increasing the project's quality is represented by the blue area G_K in the graph. Nevertheless, parameters values in Table 1 are chosen such that $c_{op} > G_K$, so that the gains from investing when old are not large enough to cover operation costs. As a consequence, an agent decides not to keep her firm after some age, and she either sells it or terminates it depending on what she knows about her project's quality.

When agents are younger, however, projects are kept for intermediate values of $\hat{\pi}_t$, instead of being terminated or sold. First, the reason why a young agent would prefer to keep her project *instead of terminating* it is that, as far as $\hat{\pi}_t$ is not too low, the high σ_t when young translates into an option value from not terminating the project: given the ability to learn the agent has, a positive updating of $\hat{\pi}_t$ can happen with a large probability, and negative expected cash-flows can be cut down by the possibility of terminating the firm in the future. Importantly, this additional option value from not terminating may compensate the operation cost of keeping, c_{op}. In these cases, once c_{op} is sunk, an agent that keeps her firm finds it profitable to invest as far as her realised period cash-flows CF_t (or, equivalently, her future expected quality $\hat{\pi}_{t+1}$) are high enough.

However, since the market can also learn about the project's quality, ability to learn (a low σ_ε) alone cannot explain why an agent would prefer to keep her project when young *instead of selling* it, if $\hat{\pi}_t$ is sufficiently high. The reason for this is that the sale cost C_{IPO}

[13]Without a convex cost, the problem of the firm would not be well-defined, and the agent would like to keep the project forever and make infinite investments to increase project's quality and value V. The difference between V_t' and p_K in Figure 4 makes that explicit.

Figure 5: Young agents, exit decisions

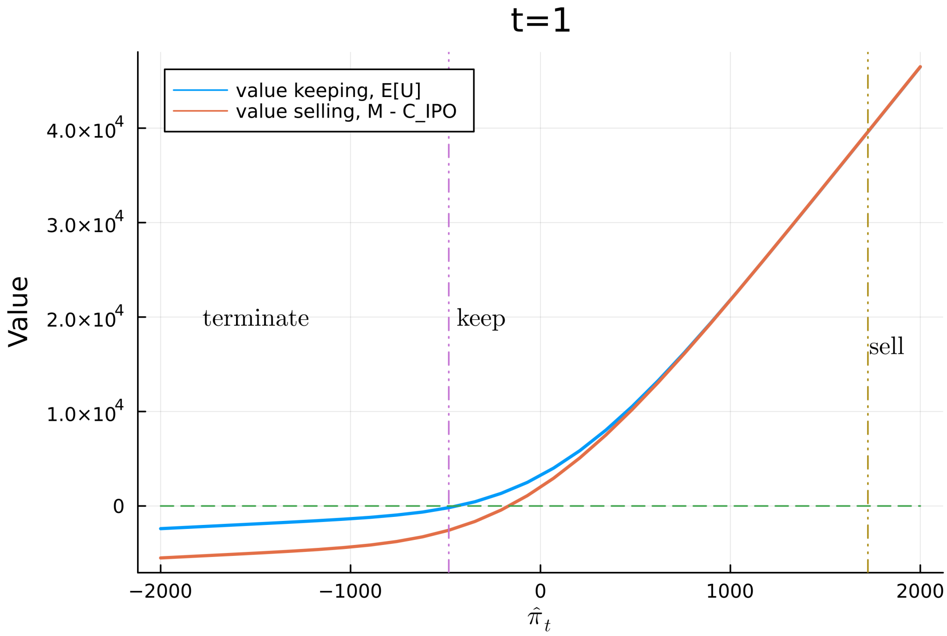

Notes: the graph shows, for an agent of age 1 owning a firm, the value she gets from terminating (dashed, green line), keeping (solid, orange line), and selling her firm to the market (solid, blue line). Sale and termination thresholds $\bar{\hat{\pi}}_t$ and $\underline{\hat{\pi}}_t$ for age 1 and *termination*, *keep* and *sale* regions are made explicit in the graph. Values are plotted against expected quality.

in Table 1 is large – in other words, it is costly to pass the firm to the market and to start market learning. Thus, although once sold to the market a firm is equally capable of updating beliefs, the original owner of the firm can learn from cash-flow realisations in a cheaper manner, for she does not have to pay C_{IPO} while she keeps her firm[14]. As a result, a young agent can delay paying sale costs by keeping her project and still benefit from the option value from learning, thus compensating c_{op} and investing afterwards.

Figure 5 shows the values of keeping and selling for different expected qualities, for an age-1 agent (thus having a high degree of quality uncertainty). For very low values of $\hat{\pi}_t$, expected cash-flows are so low that keeping or selling the project today result in a future termination with a very high probability, so the agent prefers to terminate immediately, given positive c_{op} and C_{IPO}. For very high values of $\hat{\pi}_t$, the probability of the market liquidating the project is almost zero, so the agent prefers not to delay the sale anymore. For expected qualities in between the sale and termination threshold, the agent prefers not to terminate the project and obtains an option value due to high uncertainty and the

[14]Indeed, if I set the sale cost to zero in a quantitative exercise, I find that no young agent would be willing to keep the project. An agent would be willing to sell the firm to the market if $\hat{\pi}_t$ is above some (negative) expected quality.

Figure 6: Investment policy function $g_t^k(\hat{\pi}, CF)$

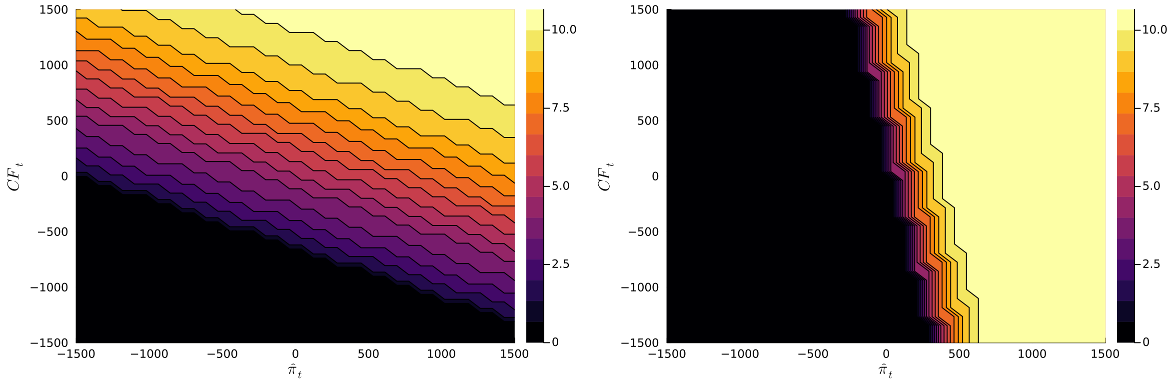

(a) $t = 0$ (b) $t = 5$

Notes: subfigure (a) represents the optimal investment policy at age 0 as a function of expected quality and period cash-flows, $g_0^k(\hat{\pi}, CF)$. Subfigure (b) represents the same object corresponding to age 5. Policy functions are obtained from the numerical solution of the optimisation problem of the agent, given the parameterisation in Table 1. Darker colours represent lower amounts of investment chosen by the agent, and lighter colours represent higher investment amounts.

ability to learn. Still, sale is delayed in that region. On one side, given the high degree of uncertainty, upon a sale today the market may liquidate the project afterwards with a positive probability, which brings down the market price today $M_t(\hat{\pi})$. On the other side, selling the project to the market and starting market learning is expensive, provided that C_{IPO} is large. This sale cost imposes a wedge between the value of keeping and the value of selling. Thus, the agent prefers to keep the project and get extra benefits from investing[15].

Investment policy over the life cycle. In Figure 6, I show the investment policy function $g_t^k(\hat{\pi}, CF)$ of an agent for two periods of time, $t = 0$ and $t = 5$. As argued in section 3.2 for a three-period model, learning makes optimal investment depend positively on CF_t. In the complete, numerically solved model, the investment policy is contingent on period cash-flows for both periods 0 and 5. Indeed, if I simulate the life cycle of an agent that makes decisions according to these policies, I find that the contemporaneous correlation between period cash-flows and investments is equal to 0.75, which denotes that fund injections and intermediate results are comoving in the simulated data. This is in line with documented facts in Kaplan and Strömberg (2003) regarding staged financing. It

[15]If we shut down the ability to invest after keeping ($B = 0$), agents that keep their projects today but sell them right after in the next period with a very high probability (i.e. those with a quite high $\hat{\pi}_t$) would prefer to sell their projects immediately. This indicates that the extra benefits from investing are also a reason for keeping projects instead of selling them.

is worth noting that the frictionless model in this paper generates contingent investment that is induced by the ability to learn and, differently from other papers in the firm dynamics literature, does not rely on financial frictions to yield a high correlation between investments and cash-flows.

As we can see from Figure 6, the degree of sensitivity of the optimal investment policy to period cash-flows changes over the life cycle of the firm. By comparing subfigures 6(a) and 6(b), we observe that the optimal investment policy is more sensitive to CF_t realisations when the firm is very young relative to older periods, given the same value for $\hat{\pi}_t$. For example, imagine a firm that enters period t (either 0 or 5) with $\hat{\pi}_t = 500$. If $t = 0$, a very bad cash-flow realisation (e.g. $CF_t = -1000$) would cause the firm not to invest, thus setting $g_0^k(\hat{\pi}, CF) = 0$, while a moderately good cash-flow realisation (e.g. $CF_t = 10$) would trigger a positive investment, being $g_0^k(\hat{\pi}, CF)$ slightly above 5. If $t = 5$, instead, these two alternative cash-flow realisations would yield optimal investments of around 9 and 10.5 respectively, spanning a smaller range of investment values and thus showing a less exacerbate reaction to intermediate results.

Again, the mechanism causing these different sensitivities is the fact that σ_t^2 is decreasing in time. When the agent is young, there is room for her initial expected beliefs $\hat{\pi}_0$ to vary, given the large initial uncertainty σ_0. This variation in beliefs may induce shifts in exit decisions and trigger strong investment reactions. However, after 5 periods, σ_5 is noticeably smaller than σ_0, as it has been shown in Figure 1. Thus, the older agent has more evidence that the true quality of the project is close to 500, i.e. that the project is sufficiently good. By that period, cash-flow realisations are less likely to cause strong changes in beliefs that induce shifts in discrete decisions, thus not giving room to sizable variations in investment.

It is important to highlight that the shape of the exit and investment policies is possible given our assumption that $\sigma_\varepsilon = 600$. If the agent was not able to learn from cash-flows when keeping her project, equations (4) and (5) would not imply a decreasing pattern in σ_t^2, and the Kalman gain would be equal to zero. As I discuss later in this section, an agent that is completely unable to learn from period cash-flow realisations has an investment policy that is completely insensitive to period cash-flows, regardless of the level of σ_t^2.

4.2 The Role of Uncertainty

Here, I simulate the parameterised model and I discuss how policies and simulated outcomes change when we modify the initial quality uncertainty of the agent, represented by parameter σ_0. The role of initial uncertainty on exit and investment policies can be easily understood along the lines of the discussions of Figures 3 and 6: a higher σ_0 man-

Figure 7: Sale and termination thresholds, different σ_0

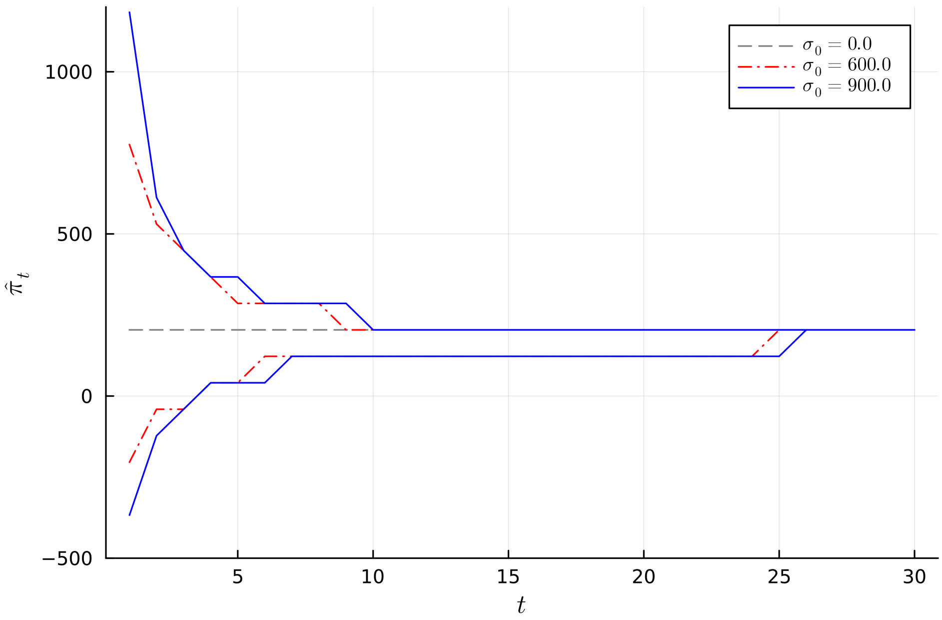

Notes: the graph shows optimal sale $\overline{\hat{\pi}}_t$ and termination $\underline{\hat{\pi}}_t$ thresholds for expected quality $\hat{\pi}_t$ for different periods of an agent's life cycle, considering three distinct levels of initial uncertainty σ_0. The lines are obtained from the numerical solution of the optimisation problem of the agent, given different levels of σ_0. The dashed line line represents the termination-sale threshold when there is no initial uncertainty. The dotted and the solid lines represent sale and termination thresholds for higher levels of initial uncertainty.

ifests through an increase in the Kalman gain, thus making the agent more predisposed to continue running the firm over the life cycle (via an increase in the option value of keeping). Besides, a higher Kalman gain makes investment policy $g_t^k(\hat{\pi}, CF)$ more sensitive to cash-flow realisations. The opposite occurs if the project is less uncertain at birth. Regarding the exit policy, Figure 7 shows how sale and termination thresholds get broader as we consider projects that are initially more uncertain. The higher the initial uncertainty, the higher the likelihood that the agent decides to keep the project over the life cycle.

Simulated moments and value. Let us perform simulations to study how several outcomes are affected by the level of initial uncertainty. For that, I consider different levels of σ_0. For each of these levels, I solve the model (keeping the rest of the parameterisation as in Table 1) and I simulate 500,000 agents that are born with a firm. All of these simulated agents have homogeneous initial beliefs, but heterogeneous unobserved initial qualities, which are drawn from the exogenous population distribution of π_0. I

Figure 8: Exit and investment dynamics and initial quality uncertainty σ_0

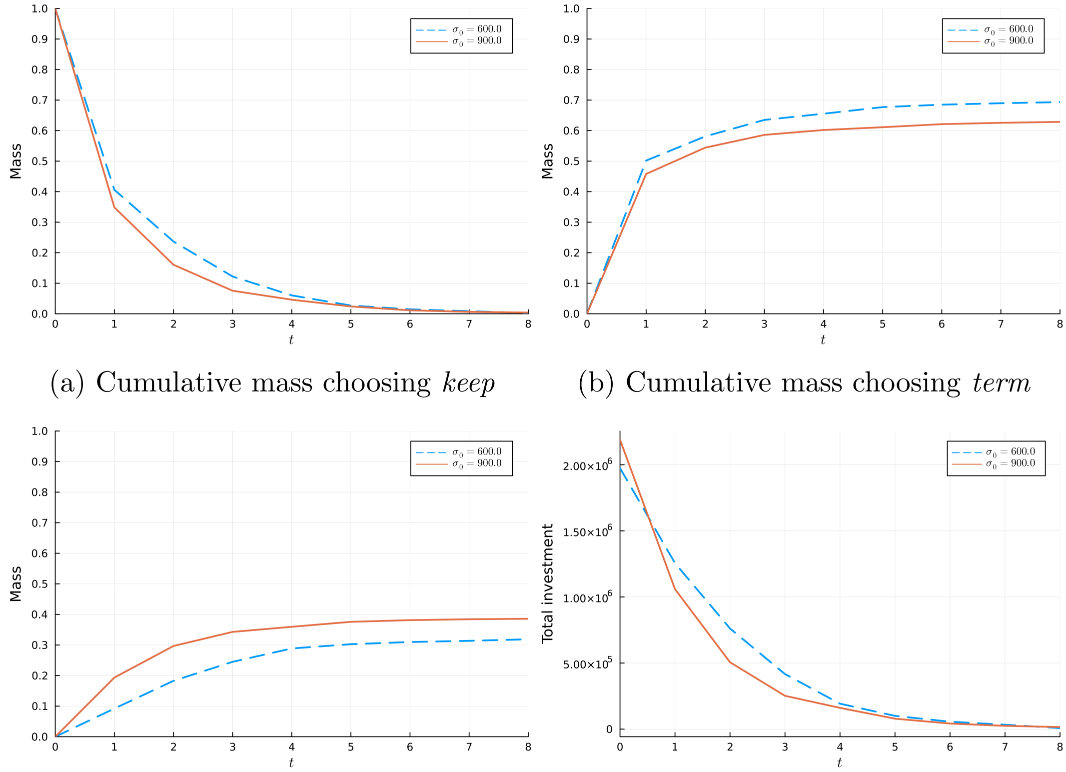

(a) Cumulative mass choosing *keep*

(b) Cumulative mass choosing *term*

(c) Cumulative mass choosing *sale*

(d) Total investment by age

Notes: for every age, and a simulated sample of 500,000 agents, I show the cumulative mass of agents choosing to keep their project (subfigure (a)), to terminate the project (subfigure (b)), to sell the project (subfigure (c)), and the total investment by age in the simulated economy (subfigure (d)). The dashed, blue line corresponds to an initial level of uncertainty $\sigma_0 = 600$. The solid, orange line corresponds to a higher level of uncertainty, $\sigma_0 = 900$.

look at several simulated moments of interest – namely, the mass of projects that are kept/terminated/have been sold by age, the total investment by age, the sale rate (i.e. the percentage of firms that are eventually sold at some period of life), and the value of the agent at firm birth.

First, in Figure 8(a), we see that a slightly larger mass of agents decides to keep their firms at young ages when we consider a level uncertainty $\sigma_0 = 600$ that is relatively low compared to the baseline level of 900. The reason for this is that, although the sale and termination bands are broader when we consider a high σ_0 (see Figure 7), the value of initial uncertainty also affects the population distribution of π_0. The higher the level of uncertainty, the larger the population variance of unobserved qualities across firms, and

Figure 9: Simulated outcomes and initial quality uncertainty σ_0

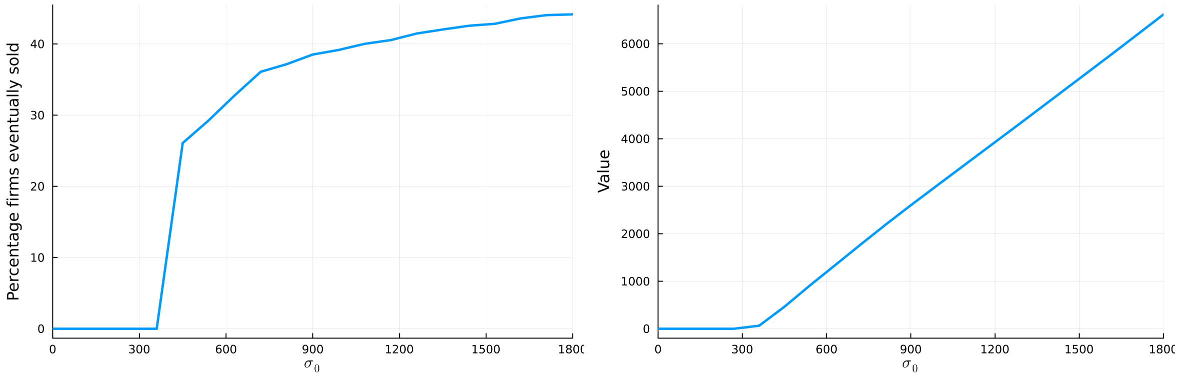

(a) Sale rate

(b) Value at period $t = 0$

Notes: subfigure (a) shows the fraction of simulated agents that eventually sell their firms at some point in their life, for simulated samples of 500,000 agents with different initial uncertainty levels. Subfigure (b) shows the initial value of an agent with initial expected quality $\hat{\pi}_0 = -50$, for different levels of initial uncertainty. I consider a range of levels for σ_0 from 0 to 1800.

thus the larger the mass of firms with extreme values of π_0. This second effect prevails, and thus more firms decide either to terminate or to sell when uncertainty is high.

In Figures 8(b) and 8(c), we see that, when initial quality uncertainty is higher, more projects are sold and less projects are terminated at young ages. As a result of a high σ_0, agents that choose to keep their projects when young receive a high option value from continuing their firms instead of selling or terminating them immediately. These agents carry out contingent investments and eventually manage to sell their projects at high market prices. Finally, Figure 8(d) shows that total investment in the economy is larger for younger firms. As they age, many agents decide to sell them to the market or liquidate them, and this generates a decrease of total investment in age.

In Figure 9, I consider a range of levels of σ_0 from 0 to 1800, and I study how initial uncertainty affects the sale rate of firms, as well as their initial value. Figure 9(a) shows the percentage of firms eventually sold by their owners for different σ_0. If initial uncertainty is very low, the sale rate takes on a value of zero. The low expected quality of projects at birth, $\hat{\pi}_0 = -50$, leads agents to liquidate their firms right away. However, for levels of σ_0 above 350, the probability that a project is eventually sold to the market is monotonically increasing in the initial uncertainty, thus indicating that a sufficiently high risk may induce agents to not terminate them immediately and to carry out growth investments, given the higher likelihood of optimistic belief updatings. Indeed, as shown in Figure 9(b), a higher σ_0 increases the initial value of the agent. A larger σ_0 raises the Kalman gain in the learning process about the project's quality and makes keeping the project and

Figure 10: Investment policy $g_t^k(\hat{\pi}, CF)$ at period $t = 0$, no learning

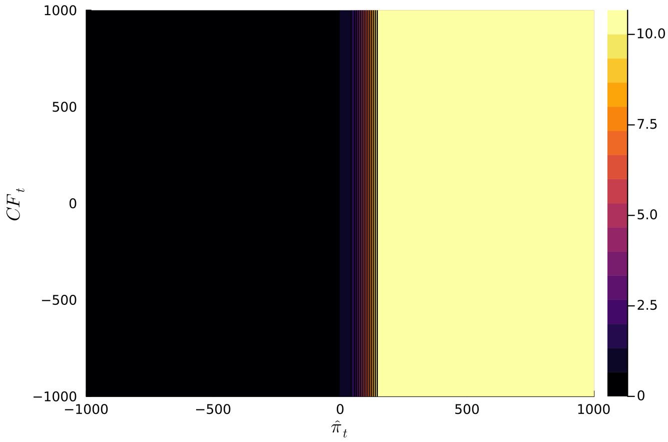

Notes: the graph represents the optimal investment policy at age 0 as a function of expected quality and period cash-flows, $g_0^k(\hat{\pi}, CF)$, in a non-learning scenario. The policy function is obtained from the numerical solution of the optimisation problem of the agent, imposing the prior belief $(\hat{\pi}_0, \sigma_0^2)$ for all t and the rest of the parameterisation in Table 1. Darker colours represent lower amounts of investment chosen by the agent, and lighter colours represent higher investment amounts.

making the appropriate contingent investment more valuable, provided that termination is always possible if belief updating turns out to be pessimistic. Agents owning more uncertain firms thus engage in more valuable experimentation, for they can learn more about its project's prospects and improve upon an immediate sale or termination.

4.3 The Role of Learning

So far, we have assumed that firms can infer information about the true quality of the project from period cash-flows. We represent this ability of firms by a parameter $\sigma_\varepsilon = 600$, which shows up in the "observation equation" (1). This section studies what happens to policies, simulated moments and value when we challenge this parameter assumption.

The non-learning scenario. Let us first depart from the baseline parameterisation and consider instead an extreme non-learning case, where I impose the prior belief $(\hat{\pi}_0, \sigma_0^2)$ for all t (so that there is no belief updating), while keeping the rest of the parameterisation in Table 1. Under the non-learning scenario, an agent does not have access to a useful learning technology that allows her to update the quality distribution. In Figure 10, I show that the investment policy at period $t = 0$ of such an agent is completely insensitive to cash-flow realisations, and is positive only if the agent expects already a high quality for

Figure 11: Exit and investment dynamics and learning σ_ε

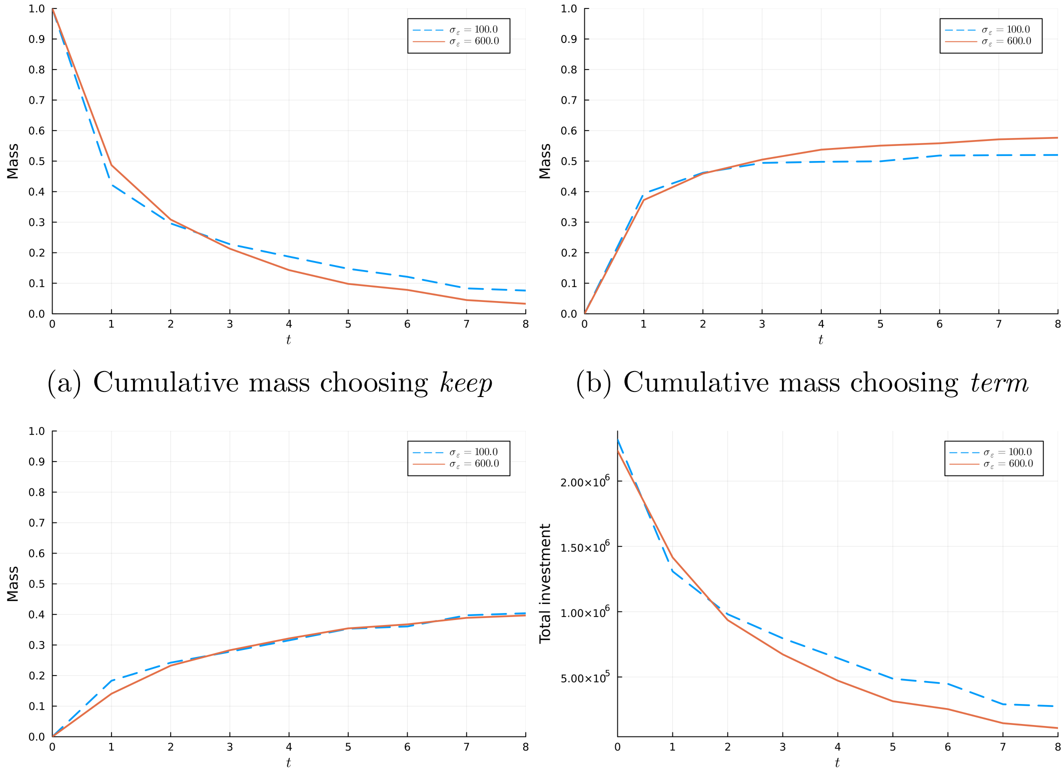

(a) Cumulative mass choosing *keep*

(b) Cumulative mass choosing *term*

(c) Cumulative mass choosing *sale*

(d) Total investment by age

Notes: for every age, and a simulated sample of 500,000 agents, I show the cumulative mass of agents choosing to keep their project (subfigure (a)), to terminate the project (subfigure (b)), to sell the project (subfigure (c)), and the total investment by age in the simulated economy (subfigure (d)). The dashed, blue line corresponds to a high level of learning ability $\sigma_\varepsilon = 100$. The solid, orange line corresponds to a lower level of learning ability, $\sigma_\varepsilon = 600$.

her firm[16]. Underlying this figure is the theoretical mechanism illustrated the three-period model from section 3.2: when the Kalman gain equals zero, no cash-flow realisation can induce a change in discrete decisions, and thus it does not affect investment.

Simulated moments and value. To see what happens if agents have access to better learning technologies relative to a non-learning environment, let us consider different values for σ_ε while keeping the rest of the parameterisation in Table 1, and let us see how simulated outcomes are affected. Figure 11 shows the simulated (for 500,000 firms) mass

[16]This result holds for any positive σ_0. Imposing $(\hat{\pi}_0, \sigma_0^2)$ for all t implies that $\kappa_t = 0$ for all t, thus rendering g_t^k insensitive to cash-flow realisations, regardless of the value of initial uncertainty.

of firms that are kept, terminated and sold over the life cycle, as well as total investment by age, for two different levels of σ_ε. A better learning technology (i.e. a lower σ_ε) induces some of the simulated agents to keep the project for more periods than in the model with more noisy signals. Figure 11(a) shows that almost 10% of agents decide to keep the project after 9 years when $\sigma_\varepsilon = 100$, versus 3% in the baseline $\sigma_\varepsilon = 600$. As a consequence, total investment in the economy at older ages increases (Figure 11(d)). Figure 12 considers a range of values for σ_ε, from 0 to 5000, and how they affect the mass of firms eventually sold and the value at firm birth. As we can see, the lower the noise of period cash-flows, the higher the sale rate in the economy and the higher the value of owning a risky project. For very noisy signals, beliefs are barely updated and waiting to receive cash-flows is useless. Thus, agents decide to terminate the project and get zero value. As we decrease σ_ε, CF_t becomes a more significant signal of π_t, which gives incentives for agents to continue keeping their projects and inject funds in a contingent manner. This increases the chances that projects are eventually sold, and thus increases the value of owning risky firms.

Finally, Figure 13 shows how the initial value of risky companies changes with the noise of period cash-flows, for a high ($\sigma_0 = 900$) and a low ($\sigma_0 = 600$) level of initial uncertainty. As we can see, the increase in firm value due to having access to more informative signals is particularly notorious when the agent owns a more risky project. This gives us a powerful reason to believe that, when projects are highly uncertain, their owners' ability to learn turns out to be a particularly valuable skill. Given a low σ_ε, high-σ_0 firms avoid liquidations and benefit from specialised, contingent funding (which the market is unable to provide), thus increasing their value.

5 Conclusion

In this paper, I present a model of the firm that mimics realistic features of young, innovative entrepreneurial projects – namely, venture-backed companies in the United States. The model explicitly considers uncertain returns, staged financing, exit decisions and intermediate results of firms over the life cycle. In this one-agent model, an agent implements a project by making investment and exit decisions over time. Importantly, the agent receives (as far as the project is not stopped) cash-flows every period that convey information about the true quality of the firm. Therefore, intermediate cash-flows allow for learning to take place over the firm life cycle, and this affects the agent's decisions. By parameterising, solving and simulating the model, I arrive at two main quantitative findings. First, I find that optimal decisions regarding when to continue running the firm,

Figure 12: Simulated outcomes and learning σ_ε

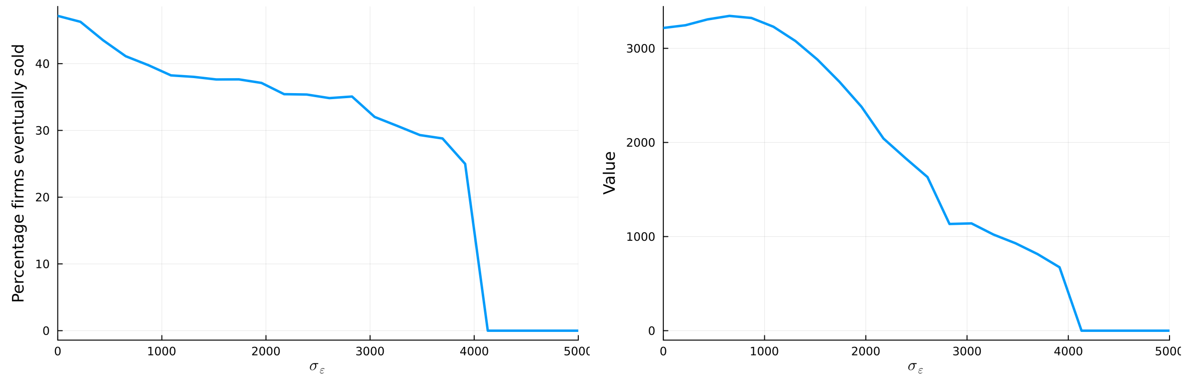

(a) Sale rate

(b) Value at period $t = 0$

Notes: subfigure (a) shows the fraction of simulated agents that eventually sell their firms at some point in their life, for simulated samples of 500,000 agents with different learning abilities, σ_ε. Subfigure (b) shows the initial value of an agent with initial expected quality $\hat{\pi}_0 = -50$, for different levels of learning ability. I consider a range of levels for σ_ε from 0 to 5000.

Figure 13: Value and learning, two levels of initial uncertainty

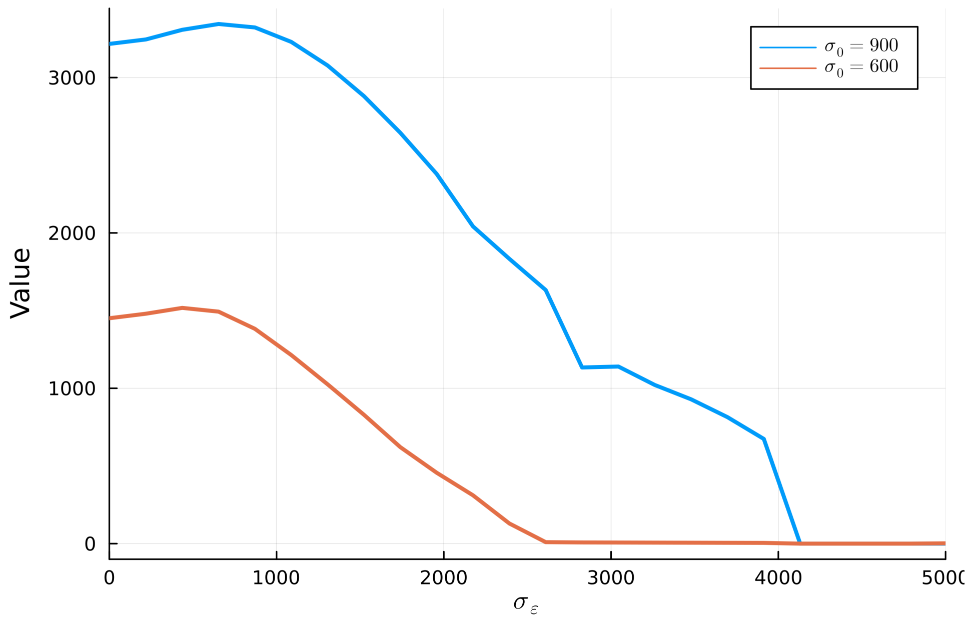

Notes: the graph shows the initial value of an agent with initial expected quality $\hat{\pi}_0 = -50$, for different levels of learning ability. I consider a range of levels for σ_ε from 0 to 5000. The blue line corresponds to a level of initial uncertainty of $\sigma_0 = 900$. The red line corresponds to a lower level of initial uncertainty, $\sigma_0 = 600$.

when to sell it to the market or when to liquidate it depend crucially on the uncertainty regarding the firm's unobserved quality. In this context of uncertain returns, the agent's capability of updating her beliefs about the project (that is, the ability to learn), and of doing it in a cheaper manner than the market (given the large costs implicit in the sale process), is determinant for her to decide when to keep and when to stop or sell the project. Second, I find that, if a sufficiently good learning technology is available to the agent, she uses period cash-flows as informative signals of the true quality of the project and injects funds accordingly. As a result, it is optimal for the agent to make investments that are contingent on realised period cash-flows. Therefore, the single-agent model with learning is capable of rationalising the empirical fact that fund injections into US venture-backed firms are contingent in the realisation of intermediate results (Kaplan and Strömberg, 2003).

From the model simulation, I find that if the noise of the signal is sufficiently low, higher quality uncertainty translates into a higher value from implementing projects, a higher sensitivity of investment to cash-flow realisations, and a high positive contemporaneous correlation between simulated investments and cash-flows, in line with empirically documented patterns. If we shut down learning, keeping high-risk projects is not valuable and investment is completely insensitive to cash-flow realisations. In that situation, improving the learning technology motivates experimentation and contingent investment, thus resulting in value gains for firms. Thus, we may think of venture capital as being a means of financing that is sophisticated enough (at least in the United States) to be close to the kind of the learning agent in my model.

These findings support the idea that a superior ability to learn helps the owners of highly uncertain entrepreneurial projects to increase their value and motivates sophisticated contingent investments. The model in this paper represents a setting where a firm owner chooses investment and exit strategies in order to maximise the net present value of her firm. The agent running the firm is not subject to any sort of contracting friction. However, in reality, we do not have just one agent implementing a high-risk project, but generally we have two parties – e.g. an entrepreneur and a venture capitalist. As it has been well acknowledged by the literature, the fact that two agents implement innovative projects may generate incentive problems (Marx, 1998; Bergemann and Hege, 1998; Cornelli and Yosha, 2003) that are though as motivating the usage of real-life, complicated securities (Kaplan and Strömberg, 2003; Cumming, 2008). Our model in section 3.1 does not inform about contractual choices that different parties may make in order to implement an innovative project. An alternative setting with contracting would consider an entrepreneur that owns the firm, and a venture capitalist that injects funds in it. These two agents implement the project jointly via some financial contract – either simple equity (Cumming, 2008; Ollivierre, 2010) or convertible securities (Kaplan and Strömberg, 2003; Schmidt, 2003). Still, even if this paper studies the life-cycle behaviour of a single agent, I show that a one-agent model is sufficient to study the effects of learning over the life cycle of risky projects, and it gives an idea on what a good contracting environment should yield in terms of investment and exit practices in those firms. As a matter of fact, documented practices in the US venture capital market seem to present both the widespread usage of exit strategies in time and contingent fund injections, two features that the model replicates well. Other important possible extensions of this work, such as the competing role of traditional sources of financing, like banks, or alternative financial contracts, are left for future research.

References

Aghion, P. and Bolton, P. (1992). "An Incomplete Contracts Approach to Financial Contracting". *Review of Economic Studies*, 59(3):473–494.

Akcigit, U., Dinlersoz, E., Greenwood, J., and Penciakova, V. (2022). "Synergizing ventures". *Journal of Economic Dynamics and Control*, page 104427.

Ateş, S. T. (2018). "Beyond Cash: Venture Capital, Firm Dynamics, and Economic Growth". *Unpublished paper. Federal Reserve Board of Governors.*

Baley, I. and Veldkamp, L. (2021). "Bayesian Learning". *NBER Working Paper 29338.*

Bergemann, D. and Hege, U. (1998). "Venture capital financing, moral hazard, and learning". *Journal of Banking & Finance*, 22(6-8):703–735.

Cornelli, F. and Yosha, O. (2003). "Stage Financing and the Role of Convertible Securities". *The Review of Economic Studies*, 70(1):1–32.

Cumming, D. (2008). "Contracts and Exits in Venture Capital Finance". *The Review of Financial Studies*, 21(5):1947–1982.

Ewens, M., Gorbenko, A., and Korteweg, A. (2022). "Venture Capital Contracts". *Journal of Financial Economics*, 143(1):131–158.

Farboodi, M. and Veldkamp, L. (2021). "A Model of the Data Economy". *NBER Working Paper 28427.*

Gilson, R. J. and Schizer, D. M. (2003). “Understanding venture capital structure: A tax explanation for convertible preferred stock”. *Harvard Law Review*, 116(3):874–916.

Gompers, P. and Lerner, J. (2001). “The Venture Capital Revolution”. *Journal of Economic Perspectives*, 15(2):145–168.

Gompers, P. A., Gornall, W., Kaplan, S. N., and Strebulaev, I. A. (2020). “How do venture capitalists make decisions?”. *Journal of Financial Economics*, 135(1):169–190.

Guvenen, F. (2007). “Learning Your Earning: Are Labor Income Shocks Really Very Persistent?”. *American Economic Review*, 97(3):687–712.

Hellmann, T. (1998). “The Allocation of Control Rights in Venture Capital Contracts”. *The RAND Journal of Economics*, 29(1):57.

Jovanovic, B. (1982). “Selection and the Evolution of Industry’. *Econometrica*, 50(3):649.

Jovanovic, B. and Szentes, B. (2013). “On the Market for Venture Capital”. *Journal of Political Economy*, 121(3):493–527.

Kalman, R. E. (1960). “A New Approach to Linear Filtering and Prediction Problems”. *Journal of Basic Engineering*, 82(1):35–45.

Kaplan, S. N. and Strömberg, P. (2003). “Financial Contracting Theory Meets the Real World: An Empirical Analysis of Venture Capital Contracts”. *The Review of Economic Studies*, 70(2):281–315.

Kaplan, S. N. and Strömberg, P. (2004). “Characteristics, Contracts, and Actions: Evidence from Venture Capitalist Analyses”. *The Journal of Finance*, 59(5):2177–2210.

Lerner, J. and Schoar, A. (2005). “Does Legal Enforcement Affect Financial Transactions? The Contractual Channel in Private Equity”. *The Quarterly Journal of Economics*, 120(1):223–246.

Luttmer, E. G. J. (2011). “On the Mechanics of Firm Growth”. *Review of Economic Studies*, 78:1042–1068.

Marx, L. M. (1998). “Efficient venture capital financing combining debt and equity”. *Review of Economic Design 1998 3:4*, 3(4):371–387.

Metrick, A. and Yasuda, A. (2010). *Venture Capital and the Finance of Innovation.* 2nd edition.

Neher, D. V. (1999). "Staged Financing: An Agency Perspective". *The Review of Economic Studies*, 66(2):255–274.

NVCA (2022). "National Venture Capital Association Yearbook". Technical report.

Ollivierre, S. (2010). "The Influence of Taxation on Capital Structure in Venture Capital Investments in Canada and the United States". *University of Toronto Faculty of Law Review*, 68.

Opp, C. C. (2019). "Venture Capital and the Macroeconomy". *The Review of Financial Studies*, 32(11):4387–4446.

Perla, J., Sargent, T. J., and Stachurski, J. (2022). "A First Look at the Kalman Filter".

Piacentino, G. (2019). "Venture Capital and Capital Allocation". *The Journal of Finance*, 74(3):1261–1314.

Pugsley, B., Sedláček, P., and Sterk, V. (2021). "The Nature of Firm Growth". *American Economic Review*, 111(2):547–579.

Repullo, R. and Suarez, J. (2004). "Venture Capital Finance: A Security Design Approach". *Review of Finance*, 8(1):75–108.

Schmidt, K. M. (2003). "Convertible Securities and Venture Capital Finance". *The Journal of Finance*, 58(3):1139–1166.

Sørensen, M. (2007). "How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital". *The Journal of Finance*, 62(6):2725–2762.

Ueda, M. (2004). "Banks versus Venture Capital: Project Evaluation, Screening, and Expropriation". *The Journal of Finance*, 59(2):601–621.

A Appendix

A.1 Derivation of the Kalman Filter Equations

We depart from a generalised form of equations (1) and (2):

$$CF_t = C\pi_t + \varepsilon_t$$

$$\pi_{t+1} = A\pi_t + Bk_t + w_t$$

where $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ and $w_t \sim N(0, \sigma_w^2)$; and from the conditional distribution of CF_t on $\hat{\pi}_t$ that we obtain by taking expectations and variance to expression (1):

$$CF_t \sim N(C\hat{\pi}_t, C^2\sigma_t^2 + \sigma_\varepsilon^2)$$

In the main text, I assume that $A = C = 1$, and that there is no shock in the law of motion equation, so that $\sigma_w^2 = 0$. The first step is to get a **filtering distribution** – that is, the distribution of quality once we have observed the realisation of the imperfect measure CF. This "filtered" quality is denoted by π_t^F, and it is our aim to get an estimate of it, $\hat{\pi}_t^F$. For that purpose, we perform the following regression of the unobserved state $(\pi_t = \pi_t^F - \hat{\pi}_t)$ on the difference between the realisation of cash-flows and their prediction (a "surprise" relative to the expected value of CF_t at period t):

$$\pi_t^F - \hat{\pi}_t = \beta(CF_t - C\hat{\pi}_t) + v_t$$

The estimator that minimises the mean square error (it is in this sense that I talk about optimality here) using the information available is $\hat{\beta} = \dfrac{Cov(\pi_t^F - \hat{\pi}_t, CF_t - C\hat{\pi}_t \mid \hat{\pi}_t)}{Var(CF_t - C\hat{\pi}_t \mid \hat{\pi}_t)}$, from which we find $\hat{\beta} = C\sigma_t^2(C^2\sigma_t^2 + \sigma_\varepsilon^2)^{-1}$. This allows us to write a point estimate for the filtered quality and the corresponding variance:

$$\hat{\pi}_t^F = \hat{\pi}_t + C\sigma_t^2(C^2\sigma_t^2 + \sigma_\varepsilon^2)^{-1}(CF_t - C\hat{\pi}_t)$$

$${\sigma^F}_t^2 = \sigma_t^2 - \sigma_t^2 C^2 \sigma_t^2 (C^2\sigma_t^2 + \sigma_\varepsilon^2)^{-1}$$

which are the parameters that characterise the filtering distribution $\pi_t^F \sim N(\hat{\pi}_t^F, {\sigma^F}_t^2)$. This filtering distribution gives probabilities of different qualities *after* we filter out the prior by the new information provided by the realisation of CF_t, i.e. it is the distribution of π_t conditional on CF_t. In other words, we are applying Bayesian updating to our prior using the realised period cash-flows. The next step is to move from filtering to prediction. I use the filtering distribution and the law of motion (2) to get the **updated**

distribution of quality π_{t+1}. Since new information about CF_t has arrived, we can use the filtered random variable π_t^F instead of π_t for predicting the future value of π_{t+1}. In other words, we turn equation (2) into $\pi_{t+1} = A\pi_t^F + Bk_t + w_t$. By taking expectations and the variance in this expression, we immediately get the two first Kalman filter equations:

$$\hat{\pi}_{t+1} = (A - \kappa_t C)\hat{\pi}_t + Bk_t + \kappa_t CF_t$$

$$= A\hat{\pi}_t + Bk_t + \kappa_t(CF_t - C\hat{\pi}_t)$$

and

$$\sigma_{t+1}^2 = (A^2 - AC\kappa_t)\sigma_t^2 + \sigma_w^2$$

where κ_t is the Kalman gain, whose expression is the third Kalman filter equation:

$$\kappa_t = AC\sigma_t^2(C^2\sigma_t^2 + \sigma_\varepsilon^2)^{-1}$$

A.2 Algorithm to Solve the One-Agent Model

Here, I replicate the value functions I have shown in section 3.1, which correspond to a firm that is held by the market (values W and M) and to a firm held by its original owner (values V and U). In either case, exit and investment decisions are made to maximise the total surplus from the firm.

The value of an age-t firm *held by the market* with period beliefs $(\hat{\pi}_t, \sigma_t^2)$ is:

$$W_t(\hat{\pi}) = \max\{0, M_t(\hat{\pi})\}$$

where $M_t(\hat{\pi})$ is given by:

$$M_t(\hat{\pi}) = \int \big(CF + \frac{1}{1+r} W_{t+1}(\hat{\pi}')\big) dG_t(CF \mid \hat{\pi})$$

where $CF = \pi + \varepsilon$ (observation equation (1)), and beliefs evolve according to equations (3)-(5) – namely, $\hat{\pi}' = \hat{\pi} + \kappa_t(CF - \hat{\pi})$. In turn, the value of an age-t firm *held by its original owner* with period beliefs $(\hat{\pi}_t, \sigma_t^2)$ is:

$$V_t(\hat{\pi}) = \max\{\underbrace{0}_{\text{termination}}, \underbrace{M_t(\hat{\pi}) - C_{IPO}}_{\text{sale}}, \underbrace{\int U_t(\hat{\pi}, CF) dG_t(CF|\hat{\pi})}_{\text{keep}}\}$$

where the value of the *keep* decision is:

$$U_t(\hat{\pi}, CF) = \max_{k \geq 0, \hat{\pi}' \in \mathbb{R}} \left\{ CF - c(k) - p_K k - c_{op} + \frac{1}{1+r} V_{t+1}(\hat{\pi}') \right\}$$
$$\text{s.t.} \qquad \hat{\pi}' = \hat{\pi} + Bk + \kappa_t(CF - \hat{\pi})$$

and, again, subject to equations (4) and (5). Since the evolution of the variance of the project, σ_t^2, and the Kalman gain, κ_t, are fully exogenous, we can use the initial condition on quality uncertainty, σ_0^2, and equations (4) and (5) to get the entire life-cycle path for the variance and the Kalman gain. An important feature of these two variables is that they converge to their respective limit values as t goes to infinity. If the Kalman filter converges at an age T_{conv}, then, for $t > T_{conv}$, the problem of the firm (either market-held or owner-held) is stationary – that is, value functions W, M, V and U do not depend on time from age T_{conv} onwards. That being said, and taking as given sequences $\{\sigma_t^2\}_{t=0}^{\infty}$ and $\{\kappa_t\}_{t=0}^{\infty}$, we can rewrite the optimisation problem faced by the owner-held firm at a generic age $t > 0$ after the period cash-flow has realised as:

$$U_t(\hat{\pi}, CF) = \max_{k \geq 0} \left\{ CF - c(k) - p_K k - c_{op} + \frac{1}{1+r} V_{t+1}(\hat{\pi} + Bk + \kappa_t(CF - \hat{\pi})) \right\}$$

which is a problem where the agent's control k and state $(\hat{\pi}, CF)$ imply a future value V_{t+1}. Importantly, the fact that the Kalman gain varies over time makes this a non-stationary problem for ages $t < T_{conv}$, i.e. until convergence of σ_t^2 and κ_t is achieved. Thus, I index values and policies by time in the exposition.

We want to find the investment policy $g_t^k(\hat{\pi}, CF) \in [0, \infty)$ and the exit policy $g_t^{exit}(\hat{\pi}) \in \{termination, sale, keep\}$. Given the max structure of $V_t(\hat{\pi})$, for convenience, let us express the exit strategy via possibly time-variant thresholds for $\hat{\pi}_t$ that make the agent indifferent between any pair of discrete choices. In the main text, I show that it is indeed the case that policies from the solved model are threshold strategies, given a wide range of alternative parameterisations. I find that these thresholds are $\underline{\hat{\pi}}_t$ and $\overline{\hat{\pi}}_t$, with $\underline{\hat{\pi}}_t < \overline{\hat{\pi}}_t$, such that the firm decides to terminate for $\hat{\pi}_t$ below $\underline{\hat{\pi}}_t$, to sell above $\overline{\hat{\pi}}_t$, and to continue running the firm herself in between these two thresholds. The algorithm I use to solve the model is value function iteration. Iterating on the value function, I get the stationary value for market-held firms and for owner-held firms. For those, I consider that the Kalman gain is equal to its limit value, which I call $\overline{\kappa}$ here. I first get the stationary values W_∞ and M_∞ (as well as the stationary termination-sale policy for a market-held firm) and then, introducing the stationary M_∞ in equation (8) with $\kappa_t = \overline{\kappa}$, I find stationary values V_∞, U_∞ and stationary policies $g_\infty^k(\hat{\pi}, CF)$ and $g_\infty^{exit}(\hat{\pi})$ for the owner-held firm. Having found the stationary functions, I use a backward-induction procedure starting from age $T = T_{conv}$ such that $\kappa_T = \overline{\kappa}$ to period 0, to get the value functions when the economy is not stationary – that is, W_t, M_t, V_t and U_t, and the corresponding non-stationary policies.

The complete value-function-iteration algorithm is:

1. Given σ_0^2, find exogenous sequences $\{\sigma_t^2\}_{t=0}^{\infty}$ and $\{\kappa_t\}_{t=0}^{\infty}$. Find the limit value for the Kalman gain, $\overline{\kappa}$.

2. **Value function iteration (stationary ages)**: for a market-held firm, find the fixed point of:

$$TW_{\infty}(\hat{\pi}) = \max\{0, M_{\infty}(\hat{\pi})\}$$

where

$$M_{\infty}(\hat{\pi}) = \int \left(CF + \frac{1}{1+r} W_{\infty}(\hat{\pi}')\right) dG_{\infty}(CF \mid \hat{\pi})$$

which is an operator that gets TW as a function of W. Notice that I am using the limit value for the Kalman gain, $\overline{\kappa}$, which prevents us from indexing values by time (I index here by ∞ to make explicit the stationary age). Once this is done, for an owner-held firm, and using M_{∞} as an input, find the fixed point of:

$$TV_{\infty}(\hat{\pi}) = \max\{0, M_{\infty}(\hat{\pi}) - C_{IPO}, \int U_{\infty}(\hat{\pi}, CF) dG_{\infty}(CF|\hat{\pi})\}$$

where

$$U_{\infty}(\hat{\pi}, CF) = \max_{k \geq 0} \left\{ CF - c(k) - p_K k - c_{op} + \frac{1}{1+r} V_{\infty}(\hat{\pi} + Bk + \overline{\kappa}(CF - \hat{\pi})) \right\}$$

which is an operator that gets TV as a function of V. Notice the limit value for the Kalman gain, $\overline{\kappa}$. The stationary investment policy is obtained using the Nelder-Mead optimisation routine, which does not rely on derivatives[17]. The stationary policy, characterised by thresholds $\{\underline{\hat{\pi}}_{\infty}, \overline{\hat{\pi}}_{\infty}\}$, is obtained by comparing the stationary values the agent gets if she terminates the firm, if she sells it to the market (in

[17] This routine is more stable than a derivative-based approach, such as an LBFGS routine. This alternative routine, in turn, yields the same results as those in the main text.

exchange of a stationary market price M_∞), and if she keeps it, for different values of $\hat{\pi}$.

3. **Backward induction (non-stationary ages)**: consider age $T = T_{conv}$ such that $\kappa_T = \overline{\kappa}$. At that period, value functions are stationary, i.e. $M_T = M_\infty$, $V_T = V_\infty$, and so on. Then, we iterate backwards on market-held and owner-held values as follows:

   (i) (<u>Period T</u>) Consider a discrete grid for $\hat{\pi}$ and compute $W_T(\hat{\pi})$ and $V_T(\hat{\pi})$ for every value in the discrete grid. Interpolate $W_T(\hat{\pi})$ and $V_T(\hat{\pi})$, to get objects $\tilde{W}_T(\hat{\pi})$ and $\tilde{V}_T(\hat{\pi})$ defined for every $\hat{\pi}_T$ in the real line.

   (ii) (<u>Period $T-1$</u>) In order to solve for U_{T-1}, consider discrete grids for $\hat{\pi}_{T-1}$ and CF_{T-1}. Use the interpolated $\tilde{V}_T(\hat{\pi}_T)$ to solve:

$$U_{T-1}(\hat{\pi}, CF) = \max_{k \geq 0} \left\{ CF - c(k) - p_K k - c_{op} + \frac{1}{1+r} \tilde{V}_T\big(\hat{\pi} + Bk + \kappa_{T-1}(CF - \hat{\pi})\big) \right\}$$

for every pair $(\hat{\pi}_{T-1}, CF_{T-1})$ in the discrete grid. Find the optimal investment policy $g^k_{T-1}(\hat{\pi}, CF)$ for every pair $(\hat{\pi}_{T-1}, CF_{T-1})$ in the discrete grid. Similarly, use the interpolated $\tilde{W}_T(\hat{\pi}_T)$ to find $M_{T-1}(\hat{\pi}, CF)$ for every pair $(\hat{\pi}_{T-1}, CF_{T-1})$ in the discrete grid.

   (iii) For each $\hat{\pi}_{T-1}$ in the discrete grid, interpolate $M_{T-1}(\hat{\pi}, CF)$ with respect to CF. Similarly, interpolate $U_{T-1}(\hat{\pi}, CF)$ with respect to CF. For each $\hat{\pi}_{T-1}$ in the discrete grid, we get $\tilde{M}_{T-1}(\hat{\pi}, CF)$ and $\tilde{U}_{T-1}(\hat{\pi}, CF)$ that is defined for every CF in the real line. Then, take expectations (integrate) these objects with respect to CF to get the value:

$$V_{T-1}(\hat{\pi}) = \max\{0, \tilde{M}_{T-1}(\hat{\pi}) - C_{IPO}, \int \tilde{U}_{T-1}(\hat{\pi}, CF) dG_{T-1}(CF|\hat{\pi})\}$$

   for each $\hat{\pi}_{T-1}$ in the discrete grid (take into account that the distribution of CF_{T-1} is known given state $\hat{\pi}_{T-1}$). By comparison of the three objects within the maximum, we can get quality thresholds $\underline{\hat{\pi}}_{T-1}$ and $\overline{\hat{\pi}}_{T-1}$, which drive the *termination/keep/sale* policy of the firm at period $T-1$. Similarly, we can get value for the market-held firm at period $T-1$, $W_{T-1}(\hat{\pi})$, and the corresponding termination-sale policy for the market.

Again, we can get interpolated objects $\tilde{W}_{T-1}(\hat{\pi})$ and $\tilde{V}_{T-1}(\hat{\pi})$, defined for every π_{T-1} in the real line.

(iv) (Periods $T-2$, ..., 0) Continue backwards until finding $V_0(\hat{\pi}_0)$ and optimal policies $g_t^k(\hat{\pi}, CF)$ and $g_t^{exit}(\hat{\pi})$ for all ages $t \geq 0$.