ESSAYS ON CONTINUOUS-TIME GAMES WITH LEARNING

Gonzalo S. Cisternas Leyton

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
ECONOMICS

Adviser: Yuliy Sannikov

June 2013

Abstract

This dissertation studies the impact of learning about unobserved payoff-relevant variables on economic decisions. In chapter 1, I study a labor market in which employers learn about a worker's unobserved skills by observing output. Skills evolve as a mean-reverting process with a trend that is potentially endogenous due to human capital accumulation. Output is additively separable in the worker's skills and in his hidden effort decision, and is also distorted by Brownian noise. Under general conditions, I show that there is an equilibrium in which effort is a deterministic function of time. This equilibrium is almost always inefficient.

In chapter 2, I study a class of continuous-time games in which one long-run agent and a population of small players learn about a hidden state from a public signal that is subject to Brownian shocks. The long-run agent can influence the small players' beliefs by affecting the signal or by affecting the hidden state itself, in both cases in an additively separable way. The impact of the small players' beliefs on the long-run agent's payoff is nonlinear. At a general level, I derive a necessary condition for Markov Perfect Equilibria in the form of an ordinary differential equation. In a subclass of games with linear-quadratic structure, I obtain closed-form solutions for global incentives through solving a new type of partial differential equation. Applications to procurement and monetary policy in the context of partial information are developed.

In chapter 3, joint with Yuliy Sannikov, a firm's earnings are driven by its stock of capital and by an underlying fundamental process. Earnings are not observable at the moment of investing in capital, thus making fundamentals unobserved. The manager learns about fundamentals by observing a signal which is distorted by Brownian noise. Investment is costly and subject to adjustment costs. We show that the sensitivity of investment to expected earnings increases as uncertainty decays over time if and

only if earnings are a concave function of fundamentals. We also show that the firm's value is always below its corresponding value in the full-information benchmark.

Acknowledgements

I would like to express my deepest gratitude to my advisor Yuliy Sannikov for his guidance, patience, support and invaluable insights through all these years. It has been an honor and a privilege to be his student, and a wonderful experience to learn from him.

I would also like to express my sincere gratitude to Dilip Abreu and Stephen Morris for their tremendous support, thoughtful advice and wisdom. It has been a great pleasure to learn from them, and I am very thankful for all their attention.

Thank you to my friends Nemanja Antic, Ben Brooks, Wei Cui, Josephine Duh, Leandro Gorno, Diogo Guillén, Constantinos Kalfarentzos, Rohit Lamba, Jay Lu, Andrés Maggi and Doron Ravid for allowing me to learn from them on a daily basis, and for making the Fisher Hall basement such a special place. Thank you to Markus Brunnermeier, Hector Chade, Sylvain Chassang, Eduardo Faingold, Henry Farber, Alex Mas and Bruno Strulovici for useful conversations about my work. Thank you to Andrea Nedic, for her infinite patience, constant support and loving care. Thank you to all my non-econ friends, for making Princeton feel like home.

Special thanks to Kathleen DeGennaro and to Laura Hedden for administering the graduate program, and for their contributions to the department.

Finally, thank you to my parents and grandparents, to whom I owe everything.

Contents

Chapter 1

Shock Persistence, Endogenous Skills and Career Concerns

1.1 Introduction

Environmental uncertainty plays an important role in the evolution of workers' perceived skills. While firms can influence an employee's productivity through tailored programs such as compensation schemes, on-the-job training or learning-by-doing, exogenous forces that affect the work environment can also have an important impact on performance. For instance, a worker's productivity may vary because unforeseen events force him to be assigned to a different task at which his productivity changes, or because tasks itself evolve due to technological progress. Moreover, these exogenous changes affect an employer's inference process about a worker's ability, as current performance can become a poor predictor of future one. Thus, in settings where wages are based on perceived skills, the degree of randomness of the environment is thus expected to influence the strategic behavior of a worker whose ultimate goal is to affect his future income stream by building a good reputation.

In this chapter I study how career concerns are shaped by the degree of randomness in the job environment. More specifically, I build on Holmstrom's (1982, 1999) seminal paper of career concerns in order to construct a continuous-time model of reputation that extends his work along two dimensions. First, I allow skills to be any process within the class of Gaussian diffusion (the continuous-time analog of an AR(1) process), with the persistence of shocks to productivity being the measure of environmental uncertainty. Second, I allow the worker to take actions that directly affect productivity. Since these actions have persistent effects on skills, they have the flavor of investments in human capital. The outcome is a very flexible and general framework that provides particularly clean insights on how belief-distortion mechanisms operate, on how wages and effort levels evolve over time, and on the extent to which reputation motives generate socially efficient outcomes.

Traditional career concerns models have focused on issues such as the extent of markets' efficiency (Holmstrom's paper), on how different information structures and multitasking affect incentives (Dewatripont et al. (1999a,b)), on the interplay between implicit incentives and short-term contracts (Gibbons and Murphy (1992)) and even on herding behavior (Scharfstein and Stein (1990)). In this chapter, in turn, I provide a detailed analysis of how different degrees of environmental uncertainty influence the reputational motives faced by individuals in dynamic settings. It is widely understood that forward-looking agents evaluate not only the immediate reputational benefits from their actions, but also their long-run consequences. Such analysis is particularly relevant in settings where market participants have precise estimates of a worker's skills, and hence, convergence to a neighborhood of the stationary learning level occurs relatively fast. I show that the persistence of shocks to productivity (which is what really determines the long-run uncertainty associated to skills) crucially affects the size of the incentives created by career concerns. Moreover, I discover that Holmstrom's classic efficiency result (skills evolving as a random walk and an infinitely patient worker) is truly an exception, with both under- and over-provision of effort as robust equilibrium outcomes for more general skills processes. Inefficiencies within career concerns models with exogenous Gaussian skills are a pervasive phenomena, going beyond discounting and transient-learning considerations.

The choice of a continuous-time framework is largely motivated by the intention to provide clean insights on how the dynamics of learning determine the gains that arise from "signal-jamming". In continuous-time settings of learning with Gaussian processes (Liptser and Shiryaev (1977)), the evolution of posterior beliefs reduce to a stochastic differential equation for the posterior mean and an ordinary differential equation for the posterior second moment. While the latter is completely exogenous, the former depends on the observed path of output and thus is controlled by the worker through his effort decision. In the absence of human capital accumulation, workers evaluate how much effort to exert, taking into account both how responsive beliefs are to new output observations, and how fast these beliefs subsequently decay over time. The strength of these forces are measured by what I call the *sensitivity of beliefs to new information* and the *rate at which beliefs discount past output observations*, respectively. While higher values of the sensitivity process increase the short-term benefits from belief-distortion, higher values of the discounting process make these distortions less persistent, and thus less attractive. In a stationary-learning setting the *sensitivity-discount* ratio corresponds to a measure of the overall responsiveness of beliefs to aggregate information, and this is what determines the benefits from belief-distortion. Most interestingly, this ratio is strictly increasing in the degree of persistence of shocks to skills, allowing us to draw a simple connection between the degree of randomness in the environment and the corresponding incentives created by career concerns: as shocks become more persistent, beliefs are gradually more responsive to overall information, and thus higher effort levels are induced.

I study human capital accumulation in order to understand how the reputational incentives that workers face are influenced by the possibility of the workers becoming endogenously more productive. In such a context, the market's inability to observe skills creates belief-manipulation motives on the workers' side that can be exploited through hidden investment decisions which boost productivity. Most importantly, I argue that investments in human capital are, in general, inefficiently low. This is despite the facts that labor markets are competitive, that there is no limited liability and that workers bear the full cost of training. Distorted incentives to invest in skills arise because, in reputation-driven markets, workers value the option to invest in human capital if and only if it can be used to influence the market's beliefs about skills. Market participants know that due to the persistent effects that human capital has on skills, a temporary additional unit of it today maps into an additional output stream that a more skilled worker is able to produce. Nevertheless, I show that the market is able to anticipate only a fraction of the flow actually realized. Competition then forces the market to pay the expected value of this anticipated stream as an ex-ante premium, and thus it ceases to have any reputational value for the worker. The unanticipated component of this additional output stream is thus attributed to non-observable skills improvement, and it is what determines the worker's marginal private benefit from a temporary additional unit of human capital. In a couple of examples I show how this discrepancy in marginal values actually generates inefficiently low investments in human capital.

The reason why the market is not able to anticipate the entire additional output flow coming from human capital accumulation is purely due to discounting. In fact, the market's belief process discounts past information at rates always higher than the rates at which skills keep track of past productivity shocks. This is because market participants have the history *output not explained by effort* as their only source of information. Such a process is the sum of both signal noise and current skills, which cannot be disentangled. Therefore, given that this process is not truly a martingale, the way in which optimal beliefs filter the information conveyed by this signal is by instantaneously reacting to new output observations, but making these reactions decay relatively quickly. This in turn can be strategically used by the worker to extract additional rents from belief-distortion. Even though it is this *discounting wedge* which makes human capital accumulation valuable for the worker, learning on the market's side is never diffuse enough to induce efficiency.

Following this, I give a detailed study of two types of human capital accumulation technologies that differ in the degree of irreversibility of the investment technology mapping investments into skills. In the weak complementarity case investment occurs through deviating effort to an alternative, but related, activity. Moreover, investments are perfectly reversible such that temporary ones have low persistence effects on output. I show that, under some circumstances, the option to acquire training is delayed: in those environments, the signaling incentives generated by career concerns are so strong that workers initially focus on influencing the market's perception about themselves, and then on investing in skills. Finally, in the strong complementarity case I study the incentives that are created when human capital accumulation arises as a byproduct of final goods production. Furthermore, these investments are more irreversible than in the previous case, so temporary investments have more persistent effects on output. In such a setting I show that effort profiles are always larger than predicted in career concern models with exogenous skills. In addition, as long as there is long-run residual uncertainty concerning beliefs, the effort component associated with human capital accumulation never vanishes.

Wages in the model have a reputational component and, if the worker can directly influence output through the choice of an action, an effort component as well. The latter monotonically decreases, as a consequence of beliefs becoming less responsive to new information over time. Nevertheless, human capital accumulation introduces a positive drift in the beliefs process and, therefore, wages can present increasing and concave profiles on average. That is, the model is able to generate the observed life-cycle pattern of wages through the traditional channel of human capital accumulation (Becker (1964)) and returns to experience (Mincer (1974) and Ben-Porath (1967)). Even more interesting is the result that, because of learning, the posterior mean always locally mean-reverts toward the current true value of skills. In numerical examples I show how reversion towards an stochastically-evolving trend can generate positive autocorrelation of changes in wages under transient learning, and negative autocorrelation in steady state.

This chapter is related to various strands in the economic literature. Regarding career concern models with Gaussian skills, Holmstrom (1982, 1999) provided a formal framework to analyze Fama's (1980) conjecture that competitive markets are sufficient for inducing efficient incentives. He found that stable environments (fixed skills), discounting and transient-learning effects, among other things, can invalidate Fama's claim. Nonetheless, he showed that efficiency is achieved in the random walk case provided the previous conditions are not met. In a static setting Dewatripont et al. (1999b) analyzed the effects of additional tasks on career concerns motives, where they stress the importance that focusing has on incentives. Gibbons and Murphy (1992) in turn studied the effects of short-term linear contracting on incentives in the presence of career concerns. They show (theoretically and empirically) that the sensitivity of optimal wages to performance increases with tenure.

The literature on human capital accumulation is extensive, with Becker (1964) and Mincer (1974) as classic references. Rosen (1972) has emphasized the role of jobs as investment opportunities where workers improve their skills and thus increase their productivity. The idea that workers, through performing tasks, can acquire skills which in turn are valued by the rest of the market can be understood as task-specific human capital (see Gibbons and Waldman (2004, 2005)). Regarding wages' structure, Abowd and Card (1988) found negative first-order autocorrelation in changes of wages using longitudinal date on earnings. They also documented no significant autocorrelation for changes in wages separated for more than two periods. Farber and Gibbons (1996) rejected their pure learning model's prediction that the residuals of wages should evolve as a martingale. In a similar vein, the error measure in my model– the gap between beliefs and true skills– is not a martingale, but instead a mean-reverting process around zero. Closely related to this chapter is the work of Kahn and Lange (2011). They find that combining learning and evolving productivity does a better job at matching the covariance structure of wages in the data used by Baker, Gibbs and Holmstrom (1994a,b) than a pure learning or pure productivity model by themselves would. Finally, this chapter is to some extent related to continuous-time techniques for addressing dynamic incentive problems. In particular, Sannikov (2008) developed a continuous-time framework to analyze a principal-agent interaction from a dynamic programming perspective. In his recursive formulation of the problem, the sensitivity of the agent's continuation-value to new output observations plays a crucial role in shaping the agent's incentives. Similarly, the sensitivity of beliefs to new information determines an important part of the gains from belief-distortion in the model presented here.[1]

In the Section 2 I present the general model. In Section 3 I study how learning in Gaussian settings takes place, I analyze the forces behind belief-distortion and show that the existence of equilibria in deterministic strategies is reduced to a simple optimization problem. In Section 4 I analyze career concerns models with exogenous skills. In Section 5 I add human capital accumulation and discuss the on-equilibrium evolution of wages. I conclude in Section 6. All proofs are relegated to Appendix A.

[1]An important difference between both models is that in Sannikov's problem the principal controls the sensitivity process, while in my model it is completely exogenous.

1.2 The Model

1.2.1 Output Technology, Skills Process and Human Capital

Consider a worker who is able to produce an output $\xi := (\xi_t)_{t\geq 0}$ continuously over time. I assume it obeys the following dynamic

$$d\xi_t = (a_t + \theta_t)dt + \sigma_\xi dZ_t^\xi,\ t \geq 0, \tag{1.1}$$

where $Z^\xi := (Z^\xi)_{t\geq 0}$ is a one-dimensional Brownian motion, $\sigma_\xi > 0$ represents the volatility of the signal's noise component and a_t is the worker's effort choice at time t, which is subject to moral hazard. The term θ_t is a random variable representing some measure of the worker's current skills, $t \geq 0$. Equivalently, it could be interpreted as the value of the worker's ability in a changing environment.[2] The stochastic process $\theta := (\theta_t)_{t\geq 0}$ is not observable by the market participants, and I refer to it as the worker's *skills process*. The output process ξ is a public signal in the economy.[3]

I assume that skills evolve according to the stochastic differential equation (SDE)

$$d\theta_t = (\bar{\theta}_t(a) + \kappa\theta_t)dt + \sigma_\theta dZ_t^\theta,\ t \geq 0. \tag{1.2}$$

Here $Z^\theta := (Z^\theta)_{t\geq 0}$ is a one-dimensional Brownian motion independent from Z^ξ representing shocks that affect true skills (or alternatively, that affect the value of a worker's ability in a changing environment). I refer to these shocks as *productivity shocks* and the parameter $\sigma_\theta > 0$ measures their volatility. The family $(\bar{\theta}_t(\cdot))_{\geq 0}$ is an

[2] In Jovanovic (1979) market participants learn about the quality of a firm-specific match between a worker and a firm. A similar matching interpretation could be applied to the model presented here if the match is understood as task-specific within a changing environment.

[3] The choice of this technology greatly simplifies the analysis and is standard in career concerns papers. The combination of complementarities in the production function, moral hazard and learning, within fully dynamic models is still an open question in the literature.

endogenous trend affecting the way in which skills grow over time (see Assumption 1 below for more details). Finally, the parameter $\kappa \in \mathbb{R}$ will be referred as the *slope* of the skills process. This general specification is able to encompass a wide variety of situations: workers suffering productivity shocks of different persistence levels, environments evolving at different speeds, and workers endogenously adapting to the environment, among others.

To understand the power of this specification, suppose first that skills evolve in a completely exogenous fashion, i.e. $\overline{\theta}_t(\cdot) \equiv \overline{\theta} \in \mathbb{R}$, for all $t \geq 0$. In this case, by varying $\overline{\theta}$, κ and σ_θ the entire class of time-homogeneous Gaussian diffusions is covered.[4] Particular examples are:

(i) Constant skills: $\sigma_\theta = \kappa = \overline{\theta} = 0$;

(ii) Skills evolving as a martingale: $\kappa = \overline{\theta} = 0$, $\sigma_\theta \neq 0$;

(iii) Mean-reverting skills: $\kappa < 0$, $\sigma_\theta \neq 0$;

(iv) Skills growing at a positive rate: $\kappa > 0$, $\sigma_\theta \neq 0$;

(v) I.i.d. skills: $\sigma_\theta = \sigma\sqrt{|\kappa|}$, $\sigma \neq 0$, and $\kappa \to -\infty$. In this case $\theta_t \sim \mathcal{N}(0, \sigma^2/2)$ for all $t \geq 0$.

Situations with non-zero volatility can be interpreted as environments subject to important technological changes, or settings in which workers suffer from non-negligible productivity shocks. The relative strength of these shocks is represented by the drift of the skills process. The martingale specification represents rapidly changing environments and workers who easily adjust to new scenarios. For example, technological changes that put pressure on the worker falling into obsolescence are,

[4] A time-homogeneous Gaussian diffusion $X := (X_t)_{t\geq 0}$ corresponds to an Ito process such that its volatility is constant and its drift is affine with constant coefficients, i.e. it satisfies the SDE $dX_t = (\alpha + \beta X_t)dt + \sigma dZ_t$ for some $\sigma, \alpha, \beta \in \mathbb{R}$, where $Z := (Z_t)_{t\geq 0}$ is a Brownian motion.

on average, immediately canceled out by (unmodeled) skills accumulation. As a consequence, whether the worker will become more or less productive relative to the environment cannot be anticipated, as captured by the zero-drift condition. In the mean-reverting specification (the continuous-time analog of an AR(1) process with root less than 1), the parameter $\overline{\theta}/|\kappa|$ corresponds to a value towards which skills are expected to converge in the long-run. Since skills are driven back to this mean-trend whenever away from it, there is some short-run predictability on the value of the worker's ability. In such a setting, productivity shocks are less persistent than in the martingale formulation and the rate at which they decay is measured by κ. In fact, as this parameter decreases, shocks tend to be less and less persistent, disappearing almost instantaneously in the limit as $\kappa \to -\infty$. A mean-reverting specification of skills may represent a worker subject to daily productivity shocks that temporarily push him away from his human capital level (the mean trend), but which tend to disappear on average as the time horizon expands. Alternatively, it could represent the value of a worker's ability in environments where unmodeled frictions prevent immediate adjustments. In this sense, any advantage or disadvantage relative to the environment is expected to persist in the near future, but is also expected to disappear gradually. Finally, in (iv) (the continuous-time analog of AR(1) processes with root larger than 1), shocks to productivity exhibit higher persistence than in the martingale formulation. This could represent situations in which initial experiences have important consequences on the long-run value of skills. For example, a recently hired worker who is assigned to a negligent mentor may not develop an appropriate understanding of his assigned task, becoming permanently disadvantaged relative to other workers in the same cohort but assigned to more competent instructors. Alternatively, it could represent highly unstable environments characterized by great dynamism and randomness. For instance, in trying to adapt to the rapid changes in the tech industry a worker may either become permanently obsolete or he may develop the exact set of skills which will become essential for a long period of time.

In all these formulations different workers can be identified with different sample paths of Z^θ. That is, unmodeled workers' characteristics that affect skills are summarized in the realization of productivity shocks. This in turn will generate considerable cross-sectional dispersion of skills within each specification (i)-(v). Yet, the likelihood at which highly productive, average or unproductive workers arise will depend on the particular model.

In describing each model, I have emphasized how the persistence of shocks to productivity varies across specifications. The degree of this persistence is captured by the slope κ. In fact, the solution to (1.2) has the form

$$\theta_t = e^{\kappa t}\theta_0 + \int_0^t e^{\kappa(t-s)}\overline{\theta}_s(a)ds + \sigma_\theta \int_0^t e^{\kappa(t-s)}dZ_s^\theta,\ t \geq 0, \tag{1.3}$$

from where it can observed that κ measures the weight given to past productivity shocks. When $\kappa = 0$, all past productivity shocks are given the same weight, so two shocks of the same size at different points in time have the same impact on future skills. When $\kappa < 0$, i.e. skills are mean-reverting, productivity shocks are in fact *discounted* at a rate $|\kappa|$, so their impact on skills tends to disappear as time passes. Finally, when skills grow at a rate $\kappa > 0$, old productivity shocks have more influence on current ability than the most recent ones.

Regarding human capital accumulation, I assume that workers enter the labor market with a human capital stock denoted by $\overline{\theta}^o \in \mathbb{R}$ representing, for instance, different educational levels. Once in a firm, workers can learn from their experience and thus become endogenously more productive. For instance, researchers may become more skilled at their fields of expertise as a consequence of permanent attempts to solve similar problems. Similarly, by constantly monitoring the prospects of the companies they invest in, young traders accumulate experience on how to interpret information coming from markets. This in turn allows them to improve their trading strategies. In environments like the ones just described, acquiring skills is more a result of permanent costly-effort decisions than a result of a choice of investment in its traditional form (upfront payment in exchange for a stream of payoffs). From this perspective, the current effort history of efforts $(a_s : 0 \leq s \leq t)$, $t \geq 0$ can be understood as a measure of the worker's experience at that instant. This experience in turn maps into a value $\overline{\theta}_t(a)$, which I interpret as the worker's *human capital stock* at time $t \geq 0$. This value is, at any point in time, an aggregate measure of both the worker's experience acquired on the job and of past investment in human capital made before entering the labor market. The family of functionals $(\overline{\theta}_t(\cdot))_{t\geq 0}$ captures the technology behind human capital accumulation and satisfies the following conditions:

Assumption 1. (i) *For each $t \geq 0$, $\overline{\theta}_t : \mathcal{M}([0,t], \mathbb{R}_+) \to \mathbb{R}$ where $\mathcal{M}([0,t], \mathbb{R}_+)$ is the set of measurable functions from $[0,t]$ to $\mathbb{R}_+$.*

(ii) *For every $y \in \mathcal{M}([0,t], \mathbb{R}_+)$, the mapping $s \mapsto \overline{\theta}_s(y)$, $0 \leq s \leq t$ is Borel-measurable.*

Part (i) in the previous assumption states that the mapping between experience– as measured by the past history of efforts– and human capital occurs in a deterministic way. Part (ii) simply ensures that integrals are well-defined. This is a very general formulation in which the only restriction I impose is that all the human capital technology is non-stochastic. Later in sections 5.1 and 5.2 I study particular examples. Observe also that because of moral hazard, human capital is private information of the worker at any point in time.

The performance of a worker depends on both his experience and on how he adapts to the environment. Hence, it is skills– the interaction between human capital and productivity shocks– the relevant process for production purposes. My model is particularly interesting in its mean-reverting specification. In such a setting skills evolve around a mean-trend that is permanently changing over time, reflecting the knowledge that workers acquire from their working experience.

The next figure illustrates how a particular realization of productivity shocks (a fixed worker) varies across environments:

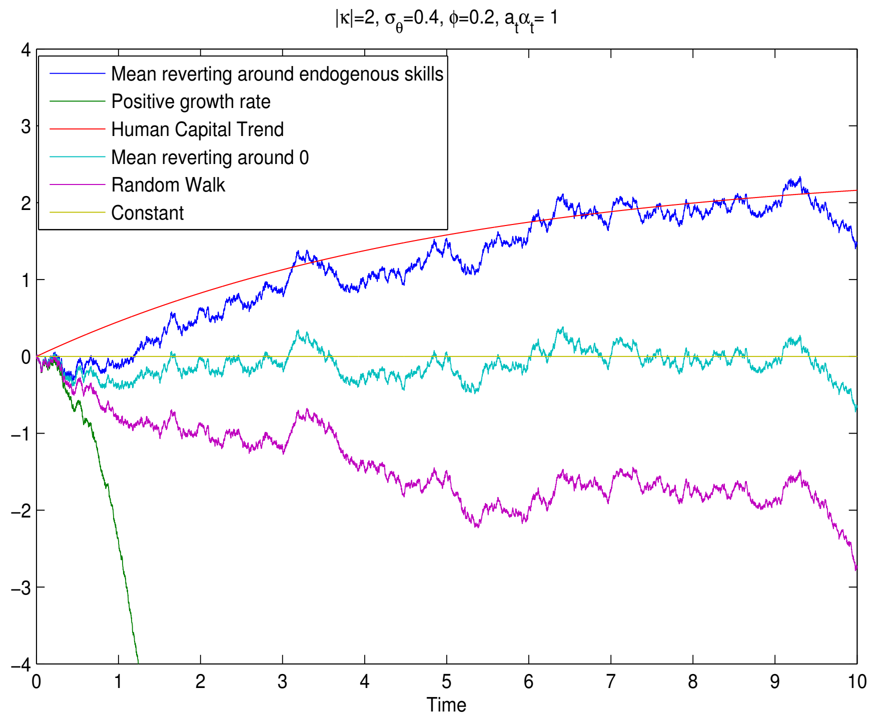

Figure 1.1: Skills models: Random walk, mean-reverting around zero, mean-reverting around endogenous trend, and positive growth rate.

In the figure, the lowest non-divergent path corresponds to a particular realization of a standard Wiener process of volatility $\sigma_\theta = 0.4$. In such a specification, the worker receives, on average, negative productivity shocks throughout the horizon studied. Adding mean-reversion around zero ($\kappa = -2$), instead forces the same sequence of shocks to fluctuate around zero. This may represent unmodeled forces that drive the worker to an average productivity level. By adding human capital accumulation according to an ordinary differential equation governed by $f(t, \overline{\theta}, a) = \alpha_t a_t - \phi\overline{\theta}$, $\phi = 0.2$ and $\alpha_t a_t \equiv 1$, skills now fluctuate around the plotted human capital trend.

Finally, the diverging path corresponds to the case in which skills change at a rate $\kappa = 2$. Observe that in this case a short sequence of negative shocks in the beginning of the worker's career generates a completely different behavior.

1.2.2 Connection with the Literature

Career Concerns: Because of their tractability in discrete-time frameworks, traditional career concerns models involving Gaussian skills (Holmstrom (1999); Dewatripont et al. (1999); Gibbons and Murphy (1982)) have analyzed only two particular specifications: when skills are fixed over time and when they evolve as a random walk. Although interesting in their own, both models are probably too limited to capture a wide variety of real-life features related to the evolution of workers' skills over long horizons. On the one hand, I have argued that workers may suffer non-negligible productivity shocks that influence short-term performance. On the other hand, the random walk specification implicitly assumes both workers and the environment freely adjusting to constant change, which allows for no predictability on the short-run evolution of the value of a worker's skills. The model presented here offers a more robust approach to modeling how skills– or their value in a changing environment– can potentially evolve over time. This allows me to recover both specifications as special cases of my general Gaussian-diffusion model in the absence of human capital accumulation.

Another important feature of these classic specifications is that they constitute landmarks within the class of Gaussian diffusions. First, observe that when skills correspond to a time-zero draw from a normal distribution, the underlying uncertainty in the model is realized at time zero. To the contrary, in any Gaussian specification with non-zero volatility ($\sigma_\theta > 0$) the underlying uncertainty is *gradually revealed* as the stochastic process of skills unfolds over time. As shown in Holmstrom's paper, the way in which uncertainty is resolved has important effects on the long-run reputational motives that arise from career concerns. In fact, he showed that in the fixed "talent" case infinitely long series of observations reveal the true value of skills in the long-run, and thus the career concerns incentives asymptotically disappear. Yet, when skills evolve as a random walk, their inherent unpredictability generates non-negligible long-run residual uncertainty. This maps into career concerns incentives that survive in a steady-state learning level. Second, notice that the martingale specification is at the boundary between the family of Gaussian diffusions that admit a long-run stationary distribution ($\kappa < 0$) and the ones that do not ($\kappa \geq 0$). With respect to learning, the latter family offers a more uncertain environment than the former, even after arbitrarily long sequences of noisy observations of skills. This chapter contributes to the previous literature not only by characterizing incentives for the entire class of Gaussian diffusions, but also by providing a particularly clean characterization on how incentives are shaped by the persistence of shocks to skills. I address this in Section 4.

General Training: Becker (1964) argued that efficient investments in general skills take place when markets are competitive and workers can contribute to their training. Yet, this conclusion relies heavily upon the contractibility– hence, upon the ex-post verifiability– of these investments. This assumption is suitable in the context of formal training methods, since in those programs monitoring problems are typically not an important issue. However, it is less appropriate in settings where training programs are less rigid and where output is noisy. Moreover, if these investments affect output-relevant worker's characteristics that are unobservable to potential employers, it is not clear that such an efficiency result would hold.

With this in mind, I study the reputational incentives that workers face when their investments in acquiring skills are imperfectly monitored. As mentioned earlier, these investments are in the form of costly-effort decisions. Jobs have become much more specific over the last decades and, consequently it is sometimes the task itself (or similar alternative ones) rather than formal training programs which provides the worker with the necessary on-the-job training. This is done in Section 5.

1.2.3 Market Structure and Equilibrium Concept

I maintain the assumptions on preferences and market structure imposed by Holmstrom (1982, 1999). The worker is assumed to be risk neutral and his utility function is separable in consumption and effort.[5] The latter is costly according to a non-negative function $g : \mathbb{R}_+ \to \mathbb{R}_+$, which is strictly increasing, strictly convex, and satisfying $g(0) = 0$ and $g'(0) = 0$. As a consequence, if at time t the manager is paid w_t and exerts effort a_t, he will get a utility flow of $w_t - g(a_t)$, $t \geq 0$.

Regarding market structure, no output-based contracts can be written and the market for workers is perfectly competitive, so firms earn zero expected discounted profits from production. Moreover, no long-term contracting is possible and hence firms earn zero profits at any point in time. Since in this model it is the market that ultimately sets wages, it corresponds to the "principal" in this market-agent interaction.

Because the market cannot observe the worker's skills, it will create estimates of it based on the public signal ξ. The zero-profit condition at every point in time implies that, if the market expects the agent to follow the effort strategy $a^* := (a_t^*)_{t\geq 0}$, it will pay him a flow payoff corresponding to the rate at which current production is expected to change:

$$w_t := \lim_{h\to 0} \frac{\mathbb{E}^{a^*}[\xi_{t+h}|\mathcal{F}_t^\xi] - \xi_t}{h} = \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t^\xi] + a_t^*,\ t \geq 0, \tag{1.4}$$

[5] All the results shown in this chapter are also valid for effort strategies taking multi-dimensional values, but with only one component of effort affecting output ξ. Unless otherwise stated, effort is uni-dimensional.

where $\mathcal{F}_t^\xi$ denotes the public information at time t, and $\mathbb{E}^{a^*}[\cdot|\mathcal{F}_t^\xi]$ is the market's conditional expectation under the assumption that the worker is following the strategy a^*.[6] This shows that wages have a *reputational* component and an effort component, both depending on what the market conjectures is the strategy that the worker will follow.

I allow for the worker to have potentially much more information about himself than the one provided by the public signal ξ. The only restriction I impose is that his posterior beliefs remain Gaussian (in the next section I study the details of the learning process).[7] Let $\mathbb{F} := (\mathcal{F}_t)_{t\geq 0}$ denote the worker's information structure and observe that, in particular, the human capital trend $(\overline{\theta}_t(a))_{t\geq 0}$ belongs to it.

Definition 1.1. *A feasible strategy for the worker corresponds to any* $\mathbb{F}-$*progressively measurable process taking values in* $\mathbb{R}_+$. *Denote this set by* $\mathcal{A}$.[8]

Since $\mathbb{F}^\xi := (\mathcal{F}_t^\xi)_{t\geq 0} \subset \mathbb{F}$, a class of strategies of particular interest is one of feasible strategies adapted to the public signal ξ.

Definition 1.2. *A strategy is public if it corresponds to a feasible strategy that is* $\mathbb{F}^\xi-$*progressively measurable.*

The equilibrium concept corresponds to perfect public equilibria, as defined next:[9]

Definition 1.3. *A perfect public equilibrium corresponds to a public strategy* $a^* := (a_t^*)_{t\geq 0}$ *and a wage process* $w := (w_t)_{t\geq 0}$, *such that:*

[6]Different effort strategies generate different probability measures on the set of paths of ξ. Therefore, if the market conjectures that the worker is following a^* when he has actually chosen the process a, their beliefs will differ.

[7]In fact, as long as screening contracts are not allowed, all the results presented below hold if the skills process is actually observed by the worker.

[8]This definition is up to mild integrability conditions on $\mathcal{A}$ such that the filtering equations exist. Of course, the equilibrium strategy will satisfy them. Hence, I do not constrain $\mathcal{A}$ ex-ante.

[9]More formally, the equilibrium notion corresponds to a sequential equilibrium. That is: (i) given a law of motion of beliefs (derived in the next subsection), the equilibrium strategy must be optimal for the agent; (ii) given the equilibrium action profile, the law of motion of beliefs is obtained via Bayes rule.

(i) Given a^*, the market sets a wage of the form $w_t = \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t^\xi] + a_t^*$ for all $t \geq 0$;

(ii) For any $t \geq 0$ and after any private history $\mathcal{F}_t$, the continuation strategy $(a_s^*)_{s\geq t}$ is optimal for the worker given the wage process in (i):

$$
\begin{aligned}
(a_s^*)_{s\geq t} \in arg \max_{a\in\mathcal{A}_t} \quad & \mathbb{E}^a\left[\int_t^\infty e^{-r(s-t)}(w_s - g(a_s))ds \middle| \mathcal{F}_t\right] \\
s.t. \quad & w_s = \mathbb{E}^{a^*}[\theta_s|\mathcal{F}_s^\xi] + a_s^*, \ \forall s \geq t
\end{aligned} \tag{1.5}
$$

where $\mathcal{A}_t$ is the set of feasible strategies at time t.

Note that this definition emphasizes the fact that, by choosing an effort strategy $a \neq a^*$, the worker induces a distribution over outcome paths that differs from the one anticipated by the market ($\mathbb{E}^a[\cdot]$ operator). This in turn will affect the market's beliefs about how skilled the worker is. For instance, if the latter deviates from a^* at some instant by exerting more effort, this will generate, in expectation, higher output observations and the market will revise its expectations upwards. As a consequence, an increase in effort today will generate, on average, a boost in the reputational component of future wages.

1.3 Preliminary Results

In this section I first explain how learning takes place and I shed light on which forces determine the agent's benefits from distorting the market's beliefs. Next, I show that the existence of *deterministic* equilibria –a particular sub-class of perfect public equilibria based only on the evolution of the posterior second moment– is reduced to the existence of a deterministic solution to a simple optimization problem. This simplification is helpful for understanding the differences between the traditional benefits from "signal-jamming" and the new gains from human capital accumulation.

1.3.1 Learning: Filtering Equations

This section characterizes the market's evolution of beliefs under the conjecture that the worker is following a *public* strategy a^*.[10] Under this assumption, output (1.1) and skills (1.2) are given by

$$\begin{aligned} d\xi_t &= (a_t^* + \theta_t)dt + \sigma_\xi dZ_t^\xi, \\ d\theta_t &= (\overline{\theta}_t(a^*) + \kappa\theta_t)dt + \sigma_\theta dZ_t^\theta. \end{aligned}$$

Since the market does not observe θ, it creates beliefs about the worker's skills based on the observation of ξ. The first relevant question is how are the $\mathbb{F}^\xi-$Brownian motions, i.e. the subjective source of uncertainty under the publicly available information. This question is important because this process will characterize the on-equilibrium evolution of beliefs.

Lemma 1.1. *Suppose the market conjectures that the manager follows a public strategy a^*. Then, the process*

$$Z_t^{a^*} := \frac{1}{\sigma_\xi}\left(\xi_t - \int_0^t (a_s^* + \mathbb{E}^{a^*}[\theta_s|\mathcal{F}_s^\xi])ds\right),\ t \geq 0, \tag{1.6}$$

is an $\mathbb{F}^\xi-$Brownian motion from the market's perspective. In particular, ξ admits a diffusion representation of the form

$$d\xi_t = (a_t^* + \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t^\xi])dt + \sigma_\xi dZ_t^{a^*},\ t \geq 0. \tag{1.7}$$

The process $Z^{a^*} := (Z_t^{a^*})_{t\geq 0}$ is called an innovation process under the public strategy a.

Proof: See Appendix A.

[10] Observe that reducing the analysis of beliefs to public strategies is an *equilibrium restriction*, but not a constraint on the agent's set of feasible strategies.

□

The intuition for this result is straightforward. At any time $t \geq 0$, the expected rate of change in output corresponds to $\mathbb{E}^{a^*}[a_t^* + \theta_t|\mathcal{F}_t^\xi] = a_t^* + \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t^\xi]$, where the last equality comes from the fact that a^* is $\mathbb{F}^\xi-$adapted. As a consequence, the difference $\xi_t - X_t$, $t \geq 0$, must be a martingale under the public information structure. Given ξ's Gaussian form, this translates into the fact that Z^{a^*} has to be a Brownian motion.[11]

Suppose that the market's initial belief is such that $\theta_0|\mathcal{F}_0 \sim \mathcal{N}(m_0, \gamma_0)$ and that it is conjectured that the agent follows a *public* strategy $a^* \in \mathcal{A}$. Given the above assumptions, the conditional distribution of θ_t given the information $\mathcal{F}_t^\xi$ is also a Gaussian for all $t \geq 0$ (Theorem 11.1. in Liptser and Shiryaev (1977)). Denote by $m_t^* := \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t^\xi]$ and $\gamma_t := \mathbb{E}^{a^*}[(\theta_t - m_t^*)^2|\mathcal{F}_t^\xi]$ the market's posterior mean and variance using the public information up to t, respectively, under the assumption that the agent follows a^*. The following result is a standard one in filtering theory (Theorem 12.1. in Liptser and Shiryaev (1977)):

Lemma 1.2. *The market's posterior mean and posterior variance processes,* $m_t^* := \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t^\xi]$ *and* $\gamma_t := \mathbb{E}^{a^*}[(\theta_t - m_t^*)^2|\mathcal{F}_t^\xi]$, $t \geq 0$, *respectively, satisfy the equations*

$$dm_t^* = (\overline{\theta}_t(a^*) + \kappa m_t^*)dt + \frac{\gamma_t}{\sigma_\xi} dZ_t^{a^*}, \tag{1.8}$$

$$\dot{\gamma}_t = 2\kappa\gamma_t + \sigma_\theta^2 - \left(\frac{\gamma_t}{\sigma_\xi}\right)^2, \tag{1.9}$$

where the Z^{a^*} *is defined in Lemma 1.1.*

Proof: See Liptser and Shiryaev (1977).

[11]Standard results in probability theory show that given any $\mathbb{F}^\xi-$progressively measurable strategy a^*, there is measurable map $b : \mathbb{R}_+ \times C(\mathbb{R}_+;\mathbb{R}) \to \mathbb{R}_+$ such that $a_t^* = b(t,\xi)$ and, for each t, $b(t,\xi)$ is $\mathcal{F}_t^\xi-$measurable. As a consequence, $\overline{\theta}_t(a^*) = \tilde{\theta}_t(\xi)$ for some $\tilde{\theta}_t(\xi)$, which is also $\mathcal{F}_t^\xi-$ measurable, $t \geq 0$. This functional representation is needed for applying the filtering techniques from Liptser and Shiryaev (1977).

□

Three interesting features of (1.8) and (1.9) are worth noting. First, the evolution of the posterior mean preserves the stochastic structure of the evolution of skills: since $(\overline{\theta}_t(a^*))_{t\geq 0}$ is adapted to $\mathbb{F}^{\xi}$, the drift of the posterior mean is also affine with the same slope and intercept. Second, the posterior mean's response to unexpected output observations (captured by the innovation process) increases with the size of the mean-square error and decreases with the signal's volatility (σ_{ξ}). This implies that beliefs react more strongly in settings where either less information has been accumulated, or where signals are more accurate. Finally, the mean-square error evolves in a *deterministic* fashion, so, provided that γ_0 is common-knowledge, the entire trajectory of the posterior variance is perfectly anticipated by both parties. This motivates the following equilibrium definition:

Definition 1.4. *A public perfect equilibrium a^* is deterministic if it depends only on the evolution of the market's posterior mean-square error.*

The reason for looking at equilibria with this characteristic is twofold. Firstly, and from a purely technical perspective, the task of finding output-dependent PPE is hard since the worker would also have to take into account the future benefits from distorting the market's conjectured action profile– finding such a fixed-point can become a particularly hard task. Secondly, the main goal of this chapter is to study both how the career concerns motives of workers are affected by the uncertainty in the environment and how these motives evolve throughout the workers' life-cycle. By introducing an additional endogenous output-dependent process, the worker's incentives may be pushed away from the purely reputational ones. Finally, observe that in any deterministic equilibrium the market's anticipated human capital path $(\overline{\theta}_t(a^*))_{t\geq 0}$ is exogenously fixed from the worker's perspective.

The speed of learning is measured by the evolution of $\gamma := (\gamma_t)_{t\geq 0}$. Suppose for the moment that learning is constant, that is, the posterior variance is at the stationary level γ^* given by the solution to the equation $0 = 2\kappa\gamma^* + \sigma_\theta^2 - (\gamma^*/\sigma_\xi)^2$. It is easy to see that, given a fixed slope $\kappa \in \mathbb{R}$, the unique stationary-learning level is given by

$$\gamma^* = \sigma_\xi^2 \left(\sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + \kappa \right), \ \kappa \in \mathbb{R}. \tag{1.10}$$

Thus γ^* corresponds to the market's long-run *residual uncertainty* regarding the worker's skills. Observe that a necessary condition for γ^* to be equal to zero is the absence of unobservable productivity shocks ($\sigma_\theta = 0$). However, as one can see from the case in which skills grow at a strictly positive rate ($\kappa > 0$), this is not sufficient. This may seem counterintuitive since when skills evolve in a deterministic way, a huge part of the model's uncertainty is eliminated. Nonetheless, what really matters in terms of asymptotic learning is the value that skills take in the long-run, and either in the presence of productivity shocks or in a setting without them but in which skills grow at a strictly positive rate, no such a value exists. In these cases, arbitrarily long sequences of noise observations are never accurate enough to fully reveal the true value of skills. The steady-state learning levels for the specifications (i)-(v) are characterized in the following

Proposition 1.1. *In a stochastic environment ($\sigma_\theta > 0$), there is always a strictly positive long-run residual level of uncertainty ($\gamma^* > 0$). Moreover, γ^* is strictly increasing in κ and $\lim_{\kappa \to -\infty} \gamma^* = 0$. In the absence of productivity shocks ($\sigma_\theta = 0$), the long-run residual uncertainty is zero if and only if $\kappa \leq 0$, and equals $2\kappa > 0$ otherwise. Finally, for the i.i.d. case ($\overline{\theta}_t(\cdot) \equiv \overline{\theta} \in \mathbb{R}$, $\sigma_\theta = \sigma\sqrt{|\kappa|}$, $\sigma > 0$ and $\kappa \to -\infty$), $\gamma_{iid}^* := \lim_{\kappa \to -\infty} \gamma^* = \frac{\sigma^2}{2} > 0$.*

Proof: Straightforward.

□

Observe first that for stochastic environments ($\sigma_\theta \neq 0$), the long-run residual uncertainty is strictly increasing in the slope κ (and weakly increasing in deterministic settings). That is, environments in which productivity shocks have a higher persistence generate higher levels of long-run residual uncertainty. This is consistent with the discussion regarding stationary distributions. Second, γ^*_{iid} coincides with the prior variance of the skills process in the i.i.d. case. This is quite intuitive given that in this specification productivity shocks have no persistence and hence, no learning takes place. In the sequel, I assume that the market's initial variance γ_0 is larger or equal than γ^*, so the quality of information improves over time.

1.3.2 Belief Distortion

Because skills are not observable, the market cannot distinguish between output changes coming from the signal noise and output changes that are the consequence of skills' variation across time. Therefore, once the market has conjectured that the worker will follow a particular public strategy, say a^*, the only information that it has available to construct statistics about the worker's skills is the process

$$Y_t := \xi_t - \int_0^t a^*_s ds, \ t \geq 0, \tag{1.11}$$

that is, the *component of output not explained by effort*. Observe that the market's posterior mean (1.8) admits the following representation with respect to $Y := (Y_t)_{t\geq 0}$

$$dm^*_t = (\overline{\theta}_t(a^*) + [\kappa - \beta_t]m^*_t)dt + \beta_t \underbrace{[d\xi_t - a^*_t dt]}_{dY_t :=}, \tag{1.12}$$

where $\beta_t := \gamma_t/\sigma^2_\xi$, $t \geq 0$, is the market's beliefs *sensitivity process*– a measure of the immediate response of beliefs to new information– which is also exogenous. Representation (1.12) is important because it shows how the market's beliefs evolve from the worker's perspective. In fact, the agent is aware that, by deviating from a^*, he can affect the evolution of Y and, as a consequence, neither Z^{a^*} nor m^* are exogenous from his standpoint.[12] Equivalently, (1.12) shows how the market's beliefs are constructed based on the output signal ξ and the conjectured effort strategy a^*, which are both part of the market's information structure.

Market participants know that the signal Y has unknown but, most importantly, non-zero increments. As a consequence, they are aware this process is not truly a martingale and hence, that changes in it are not really a surprise. As a result, the optimal way in which beliefs filter the information conveyed by this signal is by instantaneously reacting to new output observations, but making the reactions decay relatively fast. Signal-jamming indeed carries some informational costs: even though the market reacts to the information provided by Y, past observations of it are discounted more heavily. To see this observe that the solution to (1.12) has the form

$$m_t^* = m_0 e^{\int_0^t (\kappa-\beta_s)ds} + \int_0^t e^{\int_s^t (\kappa-\beta_u)du}\overline{\theta}_s(a^*)ds + \int_0^t e^{\int_s^t (\kappa-\beta_u)du}\beta_s dY_s,\ t \geq 0. \quad (1.13)$$

It is straightforward to see that while skills weigh past productivity shocks according to κ (see eq. (1.3)), skills' estimates weigh the information generated by Y using $\kappa - \beta_t$, $t \geq 0$, which is strictly less that κ. That is, learning creates a wedge between the rate at which skills keep track of past productivity shocks, and the rate at which beliefs keep track of past output observations. I call this wedge the *discounting wedge* that arises from learning. Suppose for example that skills grow at a rate $\kappa \geq 0$. Then, while skills give more weight to old productivity shocks relative to most recent ones using a factor κ, beliefs now *discount* more heavily past information using a discount rate $\delta_t := \beta_t - \kappa > 0$, $t \geq 0$.[13] If instead skills are mean-reverting ($\kappa < 0$), and hence past productivity shocks are discounted according to the fixed rate $|\kappa|$, beliefs will

[12] Of course, Z^{a^*} and, consequently, m^*, are exogenous from the market's perspective. Thus, the market's beliefs do admit a representation with respect to Z^{a^*}.

[13] Recall that $\beta_t := \frac{\gamma_t}{\sigma_\xi^2} > \frac{\gamma^*}{\sigma_\xi^2} = \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + \kappa$.

discount the information generated by Y more heavily using the deterministic discount process $\delta := (\delta_t)_{t\geq 0} > |\kappa|$. The dynamics of the posterior mean are determined by the sensitivity process $\beta := (\beta_t)_{t\geq 0}$ and the discount rate process $\delta := (\delta_t)_{t\geq 0}$. While the first process measures the initial response of beliefs to new information, the second process measures how these initial reactions decay over time. They measure two different aspects of the overall responsiveness of beliefs to aggregate information.

Given that the degree of persistence of the Brownian shocks has important consequences on the long-run behavior of the processes analyzed here, a natural question that arises is how the slope κ affects both β and δ in the long-run. Suppose that learning is stationary, so $\gamma^*(\kappa) = \sigma_\xi^2(\sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + \kappa)$ and, hence, $\delta^*(\kappa) = \sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2}$. From this expression it can be seen that the rate at which beliefs keep track of past output observations is minimized when skills evolve as a martingale. Moreover, although it is intuitive that past output observations are discounted more heavily in the mean reverting specification than in the martingale model (given the less erratic of skills beliefs don't need to keep track of too much information), it is somewhat surprising that instead past observations are given less weight when skills grow at positive rates. The reason is that because in these specifications skills have explosive paths, past output observations become worse predictors for assessing the current value of skills. More interesting is that the stationary discount rate $\delta^*(\kappa)$ depends only on the absolute value of the slope. This occurs because the Kalman-Bucy filter minimizes the distance to the underlying unobservable and thus the weight that estimates attach to past information should not depend on the particular sign of the slope of the skills process. Nonetheless, the way in which the filter differentiates between processes that admit a long-run stationary distribution and the ones that don't is through the sensitivity process. In fact, from Proposition 1.1, $\beta^*(\kappa) := \frac{\gamma^*(\kappa)}{\sigma_\xi^2}$ is increasing in κ. Intuitively speaking, as shocks become more persistent it is more likely that an output surprise is due to skills improvement rather than the consequence of noise. This translates into beliefs reacting more strongly to new output observations. Graphically:

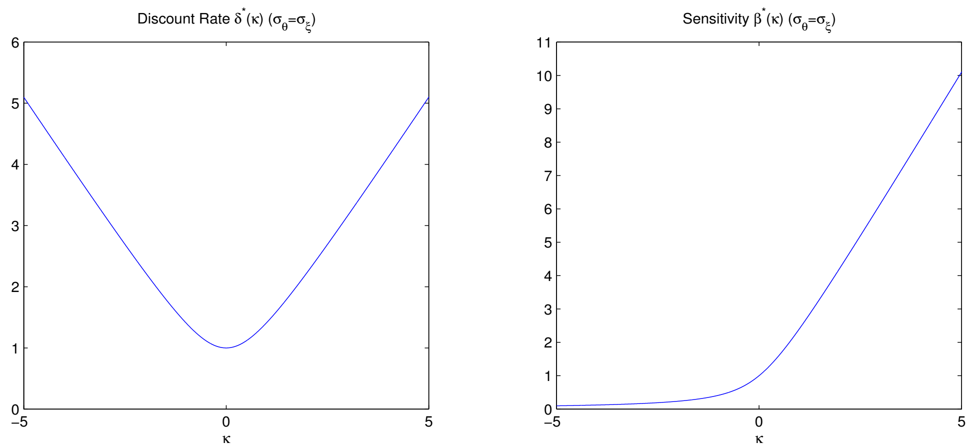

Figure 1.2: Steady-state levels of discount rate and sensitivity of beliefs. While the former does not distinguish between mean-reverting and growth rate processes with the same absolute value of the slope, the optimal filter makes a distinction between them through the long-run sensitivity parameter.

When learning is stationary, a natural measure of the overall responsiveness of beliefs to information is given by the sensitivity-discount ratio

$$w^*(\kappa) := \frac{\beta^*(\kappa)}{\delta^*(\kappa)} = 1 + \frac{\kappa}{\sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2}}. \tag{1.14}$$

The sensitivity-discount ratio is monotonically increasing in κ, as a reflect of beliefs reacting more strongly to information precisely in environments where there is more underlying randomness. Take as a reference the martingale specification $\kappa = 0$. In this case, the intensity at which beliefs react to new information has the same size as the rate at which initial responses decay over time. As a consequence, the sensitivity-discount ratio takes value 1. If skills are instead mean-reverting, productivity shocks will have less persistence than in the martingale model. Given this less erratic behavior, the overall responsiveness of beliefs to information should be lower than in the martingale specification. This results in $w^*(\kappa) < 1$ whenever $\kappa < 0$. Finally, when skills grow at a positive rate $\kappa > 0$ the higher persistence of productivity shocks increase the overall value of information relative to the martingale model. Therefore, the responsiveness of beliefs to information will be higher than in the reference case, which translates into $w^*(\kappa) > 1$ when $\kappa > 0$. It turns out that this measure of belief-responsiveness plays a crucial role in the stationary-learning incentives that arise from career concerns motives.

1.3.3 Problem Reduction: Main Lemma

In this section I first state the general problem the worker solves. Next, I present a lemma that reduces the existence of a PPE in deterministic strategies to the existence of a particular type of solution to a simple deterministic optimization problem. In fact, in order for a solution to this reduced problem to constitute a deterministic PPE, it needs to be independent of the current history of human capital at any point in time. This is because the equilibrium concept used here is quite strong: the market must correctly anticipate the worker's effort strategy after any possible private history observed by the worker.

The worker is allowed to have access to additional sources of information than the one provided by output. This is a reasonable assumption given that skills are an inherent characteristic of the agent himself. I do impose that the stochastic structure of the additional signals must lie within the Gaussian framework. Recall that $\mathbb{F} := (\mathcal{F}_t)_{t\geq 0}$ denotes the agent's information structure (containing, in particular, $\mathbb{F}^\xi$) and that $\mathcal{A}$ denotes the set of feasible strategies for the worker (see Definition 1.1). Given $a \in \mathcal{A}$, the worker's posterior mean evolves as

$$dm_t = (\overline{\theta}_t(a) + \kappa m_t)dt + \sigma_t dZ_t,\ t \geq 0 \qquad (1.15)$$

where $Z := (Z_t)_{t\geq 0}$ is a $\mathbb{F}-$Brownian motion and $\sigma := (\sigma_t)_{t\geq 0}$ is a non-negative process.[14]

Suppose that the worker follows a strategy $a \in \mathcal{A}$. Then, from his own perspective, the process

$$Z_t^a := \frac{1}{\sigma_\xi}\left(\xi_t - \int_0^t (m_t + a_t)dt\right),\ t \geq 0 \tag{1.16}$$

is an $\mathbb{F}-$Brownian motion that is correlated with Z. Therefore, as a direct application of Lemma 1.1 for $\mathbb{F}-$Brownian motions, output can be written as

$$d\xi_t = (m_t + a_t)dt + \sigma_\xi dZ_t^a,\ t \geq 0, \tag{1.17}$$

from the worker's standpoint.

With this in hand, the worker's optimization problem is straightforward: given a deterministic conjecture a^*, after any private history $\mathcal{F}_t$ (which includes, in particular, the current output path $\xi_0^t := (\xi_s)_{0\leq s\leq t}$ and current human capital path $(\overline{\theta}_s)_{0\leq s\leq t}$), the worker solves

$$\begin{aligned}
\max_{a\in\mathcal{A}} \quad & \mathbb{E}^a\left[\int_0^\infty e^{-r(s-t)}(m_s^* - g(a_s))ds\Big|\mathcal{F}_t\right] \\
s.t. \quad & dm_s^* = (\overline{\theta}_s(a^*) - \delta_s m_s^*)ds + \beta_s(d\xi_s - a_s^* ds),\ s > t,\ m_t^* = m^{*,o}, \\
& d\xi_s = (a_s + m_s)ds + \sigma_\xi dZ_s^a,\ s \geq t, \\
& dm_s = (\overline{\theta}_s(a) + \kappa m_s)ds + \sigma_s dZ_s,\ s \geq t,\ m_t = m^o, \\
& \overline{\theta}_s(a),\ s \geq t
\end{aligned} \tag{1.18}$$

[14]A process like this can be generated as follows: suppose the agent also observes d signals of the form $d\xi^i = \theta_t dt + B_t^i dZ_t^i + C_t^i dZ_t^\theta$ where $\{Z^\theta, Z^i : i = 1, ..., d\}$ is a family of independent one dimensional Brownian motions. Then, if the coefficients $\{(B_t^i, C_t^i)_{t\geq 0} : i = 1, ..., n\}$ (which may depend on output) satisfy some measurability and integrability conditions, standard filtering techniques yield that the conditional mean $(m_t)_{t\geq 0}$ evolves as $dm_t = (\overline{\theta}_t(a) + \kappa m_t)dt + \Sigma_t dZ_t^o$ where $\Sigma_t \in \mathbb{R}^d$ and Z_t^o is a d-dimensional innovation process. Letting $dZ_t := \Sigma_t dZ_t^o/||\Sigma_t||$ and $\sigma_t := ||\Sigma_t||$, proves the claim.

where $(a^*_s)_{s\geq t}$, $(\beta_s)_{s\geq t}$, $(\delta_s)_{s\geq t}$, $(\sigma_s)_{s\geq 0}$ and $(\overline{\theta}_s(a^*))_{s\geq t}$ are functions of calendar time only, and $m^{*,o}, m^o$ are given.

In the next lemma I show that looking for deterministic PPE is indeed not a bad guess as long as the family of functionals $(\overline{\theta}_t(\cdot))_{t\geq 0}$ is deterministic, which is part of Assumption 1:

Lemma 1.3. *The existence of a deterministic equilibrium is reduced to the existence of a calendar-dependent solution to the following deterministic optimization problem*

$$\mathcal{P} := \begin{cases} \max_{a\in\mathcal{A}} \int_0^\infty e^{-rt}\left[\beta_t\lambda_t a_t - g(a_t) + \rho_t\overline{\theta}_t(a)\right] dt \\ s.t. \quad \overline{\theta}_t(a),\ t\geq 0, \\ \qquad\quad \gamma_0 \geq \gamma^*, \end{cases}$$

where γ_t, β_t, $\rho_t := \frac{1}{r-\kappa} - \lambda_t$ *and*

$$\lambda_t = \int_t^\infty e^{-\int_t^s (r+\delta_u)du} ds,$$

$t\geq 0$, *are all deterministic functions.*[15]

Proof: See the Appendix.

□

As a consequence, if the starting value of human capital $\overline{\theta}^0$ is common knowledge, the market can perfectly anticipate the worker's optimal actions.

This section connects the model to the standard literature of career concerns without human capital accumulation. I complement the classic inefficiency results

[15]Implicit in this result is the assumption that $r > \kappa$ when there is human capital accumulation (for otherwise, the benefits from experience are unbounded). This assumption is not needed when $(\overline{\theta}_t(\cdot))_{t\geq 0}$ is exogenous, allowing us to perform comparative statics as the discount rate r approaches zero in standard career concerns models. Finally, observe that this constraint has a bite only when $\kappa > 0$, so the random walk and mean-reverting models admit arbitrarily low discount rates even in the presence of partially endogenous skills.

of excessive effort in the early stages of the workers' life with the new finding that there are also long-run inefficiencies which are a robust phenomena within the class of Gaussian diffusions. Human capital is analyzed in Section 5.

1.4 Inefficiencies in Career Concerns Models

Suppose that there is no human capital accumulation, that is, $\overline{\theta}_t(\cdot) \equiv \overline{\theta} \in \mathbb{R}$, for all $t \geq 0$. From Lemma 1.3, the existence of a PPE in deterministic strategies is reduced to finding a solution to $\mathcal{P}$. In this particular case such a solution can be found through pointwise optimization:

Proposition 1.2. *The unique PPE in deterministic strategies $(a_t^*)_{t\geq 0}$ is characterized by the first order condition*

$$g'(a_t^*) = \beta_t \underbrace{\int_t^\infty e^{-\int_t^s (r+\delta_u)du} ds}_{\lambda_t :=}, \ \forall t \geq 0 \tag{1.19}$$

where $\beta_t := \gamma_t/\sigma_\xi^2$ and $\delta_t := \beta_t - \kappa$, for all $t \geq 0$. Moreover,

$$\frac{da_t^*}{dt} \leq 0 \quad \Leftrightarrow \quad \frac{d\gamma_t}{dt} \leq 0.$$

Proof: See the Appendix.

□

To understand this result observe that, from the worker's standpoint, the market's beliefs take the form

$$\begin{aligned} m_t^* \;&=\; m_0 e^{-\int_0^t \delta_s ds} + \int_0^t e^{-\int_s^t \delta_u du} \overline{\theta}_s(a^*) ds + \int_0^t e^{-\int_s^t \delta_u du} \beta_s [m_s ds + \sigma_s dZ_s^a] \\ &\quad + \underbrace{\int_0^t e^{-\int_s^t \delta_u du} \beta_s [a_s - a_s^*] ds}_{(*)}, \end{aligned}$$

where I have plugged into (1.13) the evolution of output from the worker's perspective (1.17). From (*) it is extremely clear how incentives are determined in the model: a marginal increase in effort over $[t, t+1]$ generates the additional wage flow $d_s := \beta_t e^{-\int_t^s \delta_u du}$, $s \geq t$. Standing at t, this *reputational dividend flow* has a net present value of size $\beta_t \lambda_t$, $t \geq 0$, and the worker just equates the marginal cost from exerting effort to the marginal benefit from affecting future wages. As it is clearly seen from this analysis, β is what really determines the short-term gains of belief-distortion while δ determines the long-run benefits from it.

The last part of the proposition states that, as information improves, the incentives to exert effort decay over time. This goes in line with the traditional idea that career concerns motives generate higher returns in environments with more uncertainty. However, there are two opposing forces that make this conclusion non-trivial: when γ_t decreases over time, both the sensitivity of beliefs to new information, β_t, and the rate at which the market's beliefs decay, δ_t, decrease. The first force reduces the short-term benefits from signal-jamming, while the second force makes any reputational gain more persistent over time. Therefore, the result says that the short-term losses from increased precision always outweigh the long-term benefits from more permanent distortions.

The rest of the section is devoted to the analysis of incentives in the long-run. This is relevant for two reasons. First, it is not irrational to think of situations in which falling in a neighborhood of the steady-state level of learning γ^* occurs relatively fast. For instance, some industries suffer from considerable turnover of workers and, therefore, past experiences with former employees may allow firms to have relatively accurate priors about the skills of incoming workers. Second, a stationary-learning environment corresponds to the natural setting in where we can evaluate the degree of efficiency of the incentives created by career concerns motives. To see why this is the case, observe that the first-best effort process $a^e := (a_t^e)_{t\geq 0}$ maximizes the surplus generated by the interaction between the worker and the firm

$$\mathbb{E}\left[\int_0^\infty e^{-rt}(d\xi_t - g(a_t)dt)\right] = \mathbb{E}\left[\int_0^\infty e^{-rt}(\theta_t - g(a_t))dt\right].$$

In the absence of human capital accumulation, efficiency yields constant rule characterized by the condition $g'(a^e) = 1$. As a consequence, given the transient effects that learning has on incentives, a permanent efficient provision of incentives cannot be expected away from steady state.

Theorem 1.1. *Assume that skills evolve according to time-homogeneous Gaussian diffusion of slope $\kappa \in \mathbb{R}$. In a stationary-learning setting optimal effort is constant, say a^*, and characterized by the first order condition*

$$g'(a^*) = \frac{\beta^*(\kappa)}{r + \delta^*(\kappa)} = \frac{\sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + \kappa}{\sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} + r} \tag{1.20}$$

As a consequence:

- *When $\kappa < 0$: $g'(a^*(r)) < 1$, for all $r \geq 0$;*
- *When $\kappa = 0$: $g'(a^*(r)) \nearrow 1$ as $r \to 0$;*
- *When $\kappa > 0$: $g'(a^*(r)) > 1$ for all $r \in [0, \kappa)$ and $g'(a^*(r)) \leq 1$ for all $r \geq \kappa$ (with equality if and only if $r = \kappa$),*

where I made explicit the dependence of effort on the discount rate $r \geq 0$.

Proof: Direct from Proposition 1.2 when learning is stationary.

□

The question of whether or not markets induce efficient incentives captured special attention in the past (see Fama (1980)). In particular, Holmstrom showed that efficiency is achieved when skills evolve as a random walk, learning is stationary and the worker is infinitely patient. The above result proves that the random walk model is truly an exception, and long-run inefficiencies are robust to the entire class of Gaussian diffusions. In fact, observe that the introduction of a non-zero slope in the skills process moves the sensitivity-discount ratio $w^*(\kappa)$ away from 1 (see eq. (1.14)). If for instance skills are mean-reverting, the overall responsiveness of beliefs to new information is too low, and effort levels will be below efficiency. If skills instead grow at a positive rate $\kappa > 0$, the overall responsiveness of beliefs is too high, and effort will be inefficiently high for a patient worker. It is only in the martingale model where the sensitivity-discount ratio has the right size to induce the appropriate incentives. The main message of this analysis is that although the sensitivity-discount ratio is, for any given degree of shock persistence, optimal from a statistical perspective, it is almost never optimal from a social-standpoint.

Given the crucial role that the slope κ plays in the model, I illustrate next how the steady-state effort depends on it, along with its dependence on the ratio of volatilities σ_θ/σ_ξ:

Corollary 1.1. *For $\kappa \in \mathbb{R}$ and $\sigma_\theta, \sigma_\xi > 0$, denote by $a^*(\kappa, \sigma_\theta/\sigma_\xi)$, the stationary-learning effort level that arises in equilibrium. Then, for fixed $r > 0$:*

(i) *Incentives increase with the randomness of the environment: $\frac{da^*(\kappa)}{d\kappa} > 0$, for all $\kappa \in \mathbb{R}$. Also, $\lim_{\kappa \to -\infty} a^*(\kappa, \sigma_\theta/\sigma_\xi) = 0$ and $\lim_{\kappa \to \infty} g'(a^*(\kappa, \sigma_\theta/\sigma_\xi)) = 2$;*

(ii) *Asymptotic efficiency obtains as skills become infinitely volatile relative to signal noise:*

$\lim_{\sigma_\theta/\sigma_\xi \to \infty} g'(a^*(\kappa, \sigma_\theta/\sigma_\xi)) = 1$. For $r > \kappa$ the convergence is from below, ($a^*(\kappa, \sigma_\theta/\sigma_\xi)$ is strictly increasing in σ_θ/σ_ξ) and from above otherwise.

(iii) *Long-run incentives may survive in settings of deterministic skills or infinitely noisy signals: For* $\kappa > 0$ $\lim_{\sigma_\theta/\sigma_\xi \to 0} g'(a^*(\kappa, \sigma_\theta/\sigma_\xi)) = \frac{2\kappa}{r+\kappa}$*. Otherwise, the limit is zero;*

(iv) *Incentives disappear for i.i.d. skills: If* $\sigma_\theta = \sigma\sqrt{|\kappa|}$*,* $\sigma > 0$*, then* $\lim_{\kappa \to -\infty} a^*(\kappa, \sigma_\theta(\kappa)) = 0$*.*

Proof: See the Appendix.

□

Part (i) in the previous corollary is the consequence of a sensitivity-discount ratio that is increasing in κ. Since beliefs are more responsive to overall information when skills are more unstable, there are more benefits from signal-jamming in such environments. Part (ii) states that by expanding the class of stochastic processes for skills, two traditional ideas in career concerns models are now only partially true: that incentives increase both with the uncertainty of the environment (as measured by the volatility of shocks to productivity), and with the precision of the output signal. In fact, in environments where shocks have a high persistence ($\kappa > 0$), a relatively patient agent may find it optimal to decrease effort as the ratio of volatilities σ_θ/σ_ξ increases. This is because the sensitivity-discount ratio is, for very unstable processes, decreasing in the ratio of volatilities: for low values of σ_θ/σ_ξ what really matter is shock persistence, whereas for large values of it the martingale component of skills is the most relevant. Part (iii) comes from Proposition 1.1, which states that the long-run residual variance is non-zero for deterministic skills as long as they grow at strictly positive rates. Finally, part (iv) is just the consequence of the fact that no real learning takes place under i.i.d. skills. I illustrate next how the worker's effort varies with the persistence of the shocks to productivity (part (i) in the previous corollary):

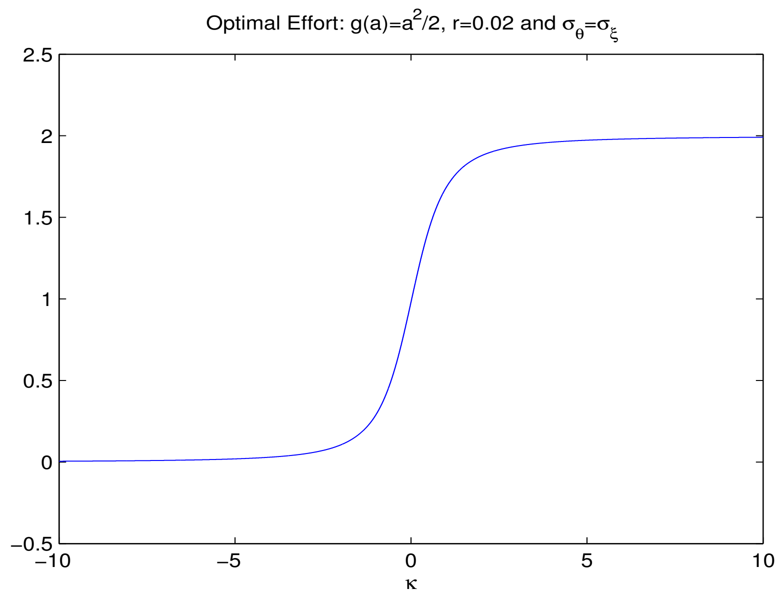

Figure 1.3: Optimal effort as a function of shock persistence.

1.5 Human Capital Accumulation

Recall that, from Lemma 1.3, the existence of a deterministic equilibrium is reduced to the existence of a deterministic solution to $\mathcal{P}$:

$$\mathcal{P} := \begin{cases} \max_{a \in \mathcal{A}} \int_0^\infty e^{-rt} \left[\beta_t \lambda_t a_t - g(a_t) + \rho_t \overline{\theta}_t(a)\right] dt \\ s.t. \quad \overline{\theta}_t(a), \ t \geq 0, \\ \qquad \gamma_0 \geq \gamma^*, \end{cases}$$

where γ_t, β_t, $\rho_t := \frac{1}{r-\kappa} - \lambda_t$ and $\lambda_t = \int_t^\infty e^{-\int_t^s (r+\delta_u)du} ds$, $t \geq 0$, are all deterministic.

The human capital model differs from the standard career concerns setting in that there is a new term in $\mathcal{P}$, $\rho := (\rho_t)_{t \geq 0}$, capturing the reputational benefits from human capital accumulation. This process is deterministic and strictly positive. Furthermore, ρ is independent the particular technology at hand. More specifically, ρ_t corresponds to the marginal private benefit from a *temporary* additional unit of human capital.

To understand how these rents operate, recall that the true value of skills is given by

$$\theta_t = \theta_0 e^{\kappa t} + \int_0^t e^{\kappa(t-s)} \overline{\theta}_s(a) ds + \sigma_\theta \int_0^t e^{\kappa(t-s)} dZ_s^\theta, \ t \geq 0. \tag{1.21}$$

It is easy to see that a temporary marginal increase in human capital at time t boosts skills by an amount $e^{\kappa(s-t)}$ at time $s \geq t$ on average. This in turn translates into an additional output flow of expected net present value $\frac{1}{r-\kappa} > 0$ that a more skilled worker is able to produce. However, since the market's beliefs decay according to $\delta > \kappa$, market participants anticipate an additional output stream of value λ_t only, which is strictly below the one actually generated. The difference in value between these two flows, ρ_t, corresponds to the expected value of a *persistent abnormal output* from the market's perspective. As a result, it is attributed to skills improvement and, hence, it determines the private value from a temporary additional unit of human capital.

The appearance of an unanticipated output flows arising from the worker's hidden investment decision is a phenomenon that takes place in equilibrium. It does not occur in standard signal-jamming models with exogenous skills (Section 4). In those models the component of output explained by effort is perfectly anticipated in equilibrium. Yet, the worker still trapped in exerting effort because the market's high expectations about output force him to do so. When workers can instead secretely invest in acquiring skills, they can make strategic use of the discounting wedge to generate abnormal returns from the market's perspective. These in turn determine the reputational benefits from human capital accumulation.

Even though a fraction of the gains arising from human capital indeed have reputational value, the discounting wedge is never large enough to align marginal private benefits with social ones. As a matter of fact, the efficient effort allocation

$a^e := (a_t^e)_{t\geq 0}$ corresponds to the solution to

$$\mathcal{P}^e := \begin{cases} \max_{a\in\mathcal{A}} \int_0^\infty e^{-rt}\left[a_t - g(a_t) + \frac{1}{r-\kappa}\overline{\theta}_t(a)\right] dt \\ s.t. \quad \overline{\theta}_t(a),\ t\geq 0, \end{cases}$$

Therefore, the social benefit from a temporary additional unit of human capital stock at any time is in fact $\frac{1}{r-\kappa}$, which is always larger than ρ. In a competitive setting the worker internalizes all the benefits and costs from production. This implies that the value of the anticipated output component $\lambda := (\lambda_t)_{t\geq 0}$ is also included in the worker's payments. Yet, competition among firms forces them to incorporate this anticipated value in the form of an ex-ante premium on the worker's wage process. Given this exogenous up-front payment, the associated additional output stream arising from human capital accumulation ceases to have any reputational value for the worker. This type of inefficiency is captured in the following:

Theorem 1.2. *Suppose that skills are unobservable and that they evolve according to (1.2), that is,* $d\theta_t = (\overline{\theta}_t(a) + \kappa\theta_t)dt + \sigma_\theta dZ_t^\theta,\ t\geq 0$, *where the family* $(\overline{\theta}_t(\cdot))_{t\geq 0}$ *satisfies Assumption 1. Suppose instead that output evolves according to* $d\xi_t = \theta_t dt + \sigma_\xi dZ_t^\xi,\ t\geq 0$, *that is, the worker now solves a pure investment problem. Then, the existence of a deterministic equilibrium is reduced to finding a solution to*

$$\mathcal{P}_I := \begin{cases} \max_{a\in\mathcal{A}} \int_0^\infty e^{-rt}\left[\rho_t\overline{\theta}_t(a) - g(a_t)\right] dt \\ s.t. \quad \overline{\theta}_t(a),\ t\geq 0, \end{cases}$$

where $\rho_t = \frac{1}{r-\kappa} - \lambda_t \geq 0$, $t\geq 0$. *However, the efficient investment allocation is given by the solution to*

$$\mathcal{P}_I^e := \begin{cases} \max_{a\in\mathcal{A}} \int_0^\infty e^{-rt}\left[\frac{1}{r-\kappa}\overline{\theta}_t(a) - g(a_t)\right] dt \\ s.t. \quad \overline{\theta}_t(a),\ t\geq 0, \end{cases}$$

That is, it is generally the case that inefficient training is an equilibrium outcome of Gaussian models of reputation.

Proof: See the Appendix.

□

In order to ascertain that inefficiencies are truly a robust phenomenon it is necessary to study different subclasses of non-stochastic human capital accumulation technologies. Nevertheless, given any particular class, it is highly unlikely both $\mathcal{P}_I$ and $\mathcal{P}_I^e$ have the same solution. Consider the two following examples.

Example 1. *Perfectly reversible human capital technologies: Assume that $\overline{\theta}_t(a) = a_t$, $t \geq 0$. Although the human capital trend is potentially discontinuous, these investments have persistent and continuous effects on skills that decay over time, as it can be seen from (1.3)*

$$\theta_t = e^{\kappa t}\theta_0 + \int_0^t e^{\kappa(t-s)} a_s ds + \sigma_\theta \int_0^t e^{\kappa(t-s)} dZ_s^\theta,\ t \geq 0$$

Within this class $g'(a_t^*) = \rho_t < \frac{1}{r-\kappa} = g'(a^e)$. Moreover, $a^* := (a_t^*)_{t\geq 0}$ is decreasing. Finally, the same inequality would hold for more general perfectly-reversible technologies of the form $\overline{\theta}_t(a) = h(t, a_t)$ with h increasing in its second argument.

□

Example 2. *Physical-Capital Technologies: Assume that $(\overline{\theta}_t(a))_{t\geq 0}$, $t \geq 0$, corresponds to the solution to an ordinary differential equation (ODE) of the form*

$$d\overline{\theta}_t = (\alpha_t a_t - \phi\overline{\theta}_t)dt,\ t \geq 0,\ \overline{\theta}_0 = \overline{\theta}^o \geq 0. \tag{1.22}$$

Here $\alpha := (\alpha_t)_{t\geq 0}$ is a positive and uniformly bounded deterministic process representing life-cycle effects from learning-by-doing (typically non-increasing). $\phi \in (0, 1)$ is a *depreciation coefficient. Dynamic-programming arguments used Section 5.3 allow us to conclude that*

$$g'(a_t) = \alpha_t \int_0^\infty e^{-(r+\phi)(s-t)} \rho_s ds < \frac{\alpha_t}{(r+\phi)(r-\kappa)} = g'(a^e).$$

Moreover, $a^* := (a_t^*)_{t \geq 0}$ *is decreasing.*

□

The marginal private benefit from a temporary additional unit of human capital, ρ_t, decays as time goes by. This is because reputation-driven workers value the option to invest in human capital if and only it can be used to distort the market's beliefs. As information becomes more precise, the market's beliefs decay less strongly. This in turn yields abnormal outputs that decrease in size over time, reducing the reputational value that human capital has. Therefore, when the market learns about workers' skills, human capital inefficiencies are expected to worsen as information improves. The previous examples show that this is indeed the case.

Instead, when skills are observable by all market participants a competitive market would set a wage flow process of the form $w_t = \theta_t$, $t \geq 0$. Observe that since it is the worker who actually controls $\theta := (\theta_t)_{t \geq 0}$, the market's conjectured strategy a^* plays no role in the way wages are set. Given this wage process, the worker's problem coincides with the one that maximizes the surplus. That is, the standard efficiency result for general training is recovered.

The reason behind this important discontinuity lies in the fact that, when skills are observable, both parties can implicitly contract on future values of the skills process. Although the worker receives a fixed wage θ_t over the interval $[t, t+dt)$, by investing in human capital at time t the worker affects θ_{t+dt}. The latter random variable in turn determines the worker's flow wage over the next interval of time. A worker standing at time t knows that competition induces the market to implicitly offer a contract that is contingent on all the possible values that θ_{t+dt} can take at time $t + dt$. No ex-ante premia on skills are paid (if they are higher than the ones realized, firms make losses; if they are lower, the worker switches to a different firm). Since skills are observable, this contract is verifiable. Because of competition, this contract is enforceable. As a consequence, efficiency is obtained due to Becker's classic argument.[16] The same efficiency result would hold if skills were not observable, and the investment action monitored but not contractible. In this case, private and public beliefs would be always aligned, and controlled by the worker through his investment decision. Since beliefs would evolve in the same way as skills do (the slope of both processes coincide), incentives are determined by the same optimality conditions as in the skills-observability case, which in turn coincides with the efficient investment allocation.

Before moving on to the next section I would like to summarize the results found here. When skills are observable, hidden investments in human capital do not generate inefficiencies. That is, competitive markets induce workers to take socially efficient actions. However, the market's inability to observe workers' skills creates belief-manipulation motives on the workers' side. These motives can be exploited through hidden investment decisions that boost productivity. Nevertheless, these reputation-driven incentives are never efficient, and their degree of inefficiency is expected to worsen over time.

[16] A similar discontinuity occurs in comparative statics with respect to σ_ξ in career concerns settings with exogenous skills when learning is stationary. As σ_ξ decreases to zero, efficient incentives are induced in the limit (Corollary 1.1, part (ii)). Yet, when $\sigma_\xi = 0$, skills are observable and thus no incentives to exert effort are generated. The same comparative static can be performed in Examples 1 and 2 when learning is stationary. In fact the stationary value of ρ is $\rho^*(\sigma_\xi) = \frac{1}{r-\kappa} - \frac{1}{r-\kappa+\beta^*(\sigma_\xi)}$. But $\beta^*(\sigma_\xi) = \gamma^*(\sigma_\xi)/\sigma_\xi^2 \to 0$ as $\sigma_\xi \to 0$. As a consequence, $\rho^*(\sigma_\xi) \to 0$ as $\sigma_\xi \to 0$.

1.5.1 Weak Complementarity

I assume that human capital accumulation and final-goods production are independent decisions. The worker is allowed to choose an unobservable action profile $a := (a_t^1, a_t^2)_{t\geq 0}$, with the first component affecting output and the second one temporarily boosting skills. Moreover, the impact of the agent's actions on human capital is additively separable: $\overline{\theta}_t(a) := \overline{\theta} + a_t^2,\ \overline{\theta} \in \mathbb{R},\ t \geq 0$. Without loss of generality, $\overline{\theta} = 0$, so skills evolve according to $d\theta_t = (a_t^2 + \kappa\theta_t)dt + \sigma_\theta dZ_t^\theta,\ t \geq 0$. When $\kappa \geq 0$, a_t^2 boosts the rate at which skills grow. If in turn $\kappa < 0$, skills locally mean-revert towards $a_t^2/|\kappa|$. The weak complementarity between human capital accumulation and final-goods production is understood as follows: because of their perfect reversibility, temporary investments in human capital have a low impact on future skills. Therefore, the impact of these investments on output disappears relatively quickly when compared to more irreversible technologies.

Using Lemma 1.3, the worker solves

$$\begin{aligned} \max_{a\in\mathcal{A}} \quad & \int_0^\infty e^{-rt}\left[\beta_t\lambda_t a_t^1 + \rho_t a_t^2 - g(a_t^1, a_t^2)\right] dt \\ s.t. \quad & (a_t^1, a_t^2) \in C \subset \mathbb{R}_+^2,\ t \geq 0, \\ & \gamma_0 \geq \gamma^*, \end{aligned}$$

where $g : C \to \mathbb{R}_+$ represents the disutility of effort and C is the set of feasible values that effort can take.

A case of particular interest is when $g(x, y) = \tilde{g}(x+y)$, some function $\tilde{g} : \mathbb{R}_+ \to \mathbb{R}_+$ strictly increasing and convex, since in such a setting strategic effects coming effort substitutability across tasks are eliminated. A natural questions that arises is whether the agent will actually decide to invest in human capital accumulation, since by affecting output directly, the choice of a^1 probably biases the worker's preferences towards using the traditional signal-jamming channel. As I show next, this may not be the case and, furthermore, delayed training is sometimes optimal.

Proposition 1.3. *Suppose that effort is perfectly substitutable in the cost-of-effort function and that* $C = \{(x, y) \in \mathbb{R}^2_+ | \, x+y \leq R\}$ *for some* $R > 0$*. Then* $\min\{a_t^{1,*}, a_t^{2,*}\} = 0$ *for all* $t \geq 0$ *and:*

(i) *If* $r - \kappa \geq 1$, $a_t^1 > 0$ *for all* $t \geq 0$*. That is, the worker never invests in human capital;*

(ii) *If* $r - \kappa < 1$ *and there is non-zero long-run residual uncertainty* ($\gamma^* > 0$)*, then there exists* $T(\gamma_0) < \infty$ *such that the worker invests in human capital* ($a_t^2 > 0$) *from* $T(\gamma_0)$ *on. Moreover, given any* $\gamma_0 > \gamma^*$*, there exists* $\epsilon > 0$ *such that if* $1 - \epsilon < r - \kappa < 1$*, then* $T(\gamma_0) > 0$*. That is, the worker delays training.*

Proof: See the Appendix.

□

Part (i) eliminates any chance for endogenous accumulation of human capital when skills mean-revert at sufficiently high rates ($\kappa < -1$). In such environments, any investment in human capital vanishes too fast and, hence, the value of the abnormal output generated from it is too low relative to the benefits associated to directly boosting output. In contrast, as part (ii) shows, for relatively low values of the mean reversion coefficient or under a positive grow rate of skills, endogenous experience accumulation may prevail over the standard signal-jamming channel. More interestingly, it is plausible to observe agents that *delay* human capital accumulation. In those cases the worker favors signaling early in his working life since this is actually the fastest channel to quickly build up a reputation. Later on, once information has improved and the market's beliefs are less responsive to new signals, the worker switches to invest in acquiring more skills. Although stylized on its own, this model shows that the decision to delay training is sometimes optimal for relatively impatient agents in environments where productivity shocks have enough persistence. Such a behavior is consistent with the career paths observed, for example, in the banking sector.

1.5.2 Strong Complementarity

Finally, I address the case in which human capital accumulation arises as a byproduct of final goods production. In this setting the worker chooses a unique action profile $a := (a_t)_{t\geq 0}$ affecting both output and human capital. The latter arises as solution to the ODE (1.22) in Example 2, that is, $d\overline{\theta}_t = (\alpha_t a_t - \phi\overline{\theta}_t)dt$. Again, by Lemma (1.3), the agent's problem reduces to an optimal control problem

$$\mathcal{P}^c := \begin{cases} \max\limits_{a\in\mathcal{A}} \int_0^\infty e^{-rt}\left[\beta_t\lambda_t a_t - g(a_t) + \rho_t\overline{\theta}_t(a)\right]dt \\ s.t. \quad d\overline{\theta}_t = (\alpha_t a_t - \phi\overline{\theta}_t)dt,\ t > 0, \\ \qquad\quad \overline{\theta}_0 = \overline{\theta}^o \geq 0. \end{cases}$$

I assume that the worker enters the market with a non-negative human capital stock so as to prevent future levels of it from taking negative values and thus from growing even without the exertion of effort. Finally, in order for the worker's problem to be finite, $\alpha \geq 0$ is assumed to be uniformly bounded and effort can take values in a bounded set $[0, \ell]$, with ℓ large enough. As a consequence, the set of feasible strategies $\mathcal{A}$ corresponds to the set of measurable functions from $\mathbb{R}_+$ to $[0, \ell]$.

Strong complementarity is captured by two features. First, human capital is task-specific: workers become more productive as a consequence of leaning-by-doing. Second, since the mean-trend evolves continuously over time, human capital investments have now a higher degree of irreversibility when compared to the weak complementarity case. This in turn maps into investments having more persistent effects on output relative to perfectly reversible ones.

When human capital evolves as in (1.22) the model become fully separable in effort and the hidden variables (skills and human capital) so a deterministic public equilibrium exists. For this purpose, define the continuation value

$$\mathcal{V}(t,x) \equiv \sup_{a\in\mathcal{A}} \int_t^\infty e^{-rs}\left[a_s\beta_s\lambda_s - g(a_s) + \rho_s\overline{\theta}_s^{t,x}(a)\right] ds,\ (t,x) \in \mathbb{R}_+^2 \tag{1.23}$$

where $(\overline{\theta}_s^{t,x}(a))_{s\geq t}$ is the solution of the ODE (1.22) with initial condition $\overline{\theta}_t = x$. Classic results in dynamic programming state that if $\mathcal{V}$ is smooth enough (of class $C^1(\mathbb{R}_+^2)$) and satisfies a growth condition (see Pham (2009)), then it is the unique solution to the Hamilton-Jacobi equation (HJ)

$$0 = \sup_{u\in A}\left\{e^{-rt}[\beta_t\lambda_t u - g(u) + \rho_t x] + \frac{\partial\mathcal{V}}{\partial t}(t,x) + \frac{\partial\mathcal{V}}{\partial x}(t,x)[\alpha_t u - \phi x]\right\},\ (t,x)\in\mathbb{R}_+^2,$$

satisfying the following transversality condition: $\lim_{t\to\infty}\mathcal{V}(t,\overline{\theta}_t^{0,x}(a)) = 0$ for every $x \in \mathbb{R}_+$ and any control $a \in \mathcal{A}$. Moreover, the optimal effort strategy is given by the maximizer $u^*(\cdot)$ of the right-hand side in (HJ), i.e. $a_t^* = u^*(t,\overline{\theta}_t)$, $t \geq 0$.

The above problem, although linear in its dynamic, it is not straightforward to solve. First, the cost function g corresponds to any strictly increasing and convex function and not necessarily of quadratic form. Second the problem is a non-stationary one. Nevertheless, it has a particularly clean solution:

Proposition 1.4. *In the strong complementarity case a deterministic equilibrium exists. It is characterized by the following first order condition:*

$$g'(a_t^*) = \beta_t\lambda_t + \alpha_t\mu_t \tag{1.24}$$

where $\mu_t := \int_t^\infty e^{-(r+\phi)(s-t)}\rho_s ds$ *is a decreasing function and* $\rho_t = \frac{1}{r+\kappa} - \lambda_t$, $t \geq 0$. *The continuation-value function takes the form*

$$\mathcal{V}(t,x) = e^{-rt}[\eta_t + \mu_t x], \; (t,x) \in \mathbb{R}^2_+, \tag{1.25}$$

where $\eta_t := \int_t^\infty e^{-r(s-t)}[g'(a_s^*)a_s^* - g(a_s^*)]ds$ *for all* $t \geq 0$.

Proof: See the Appendix.

□

As it can be seen from the result, the optimal effort allocation is larger than what predicted by traditional career concerns with exogenous skills. The difference is given by the discounted benefits associated to a marginal increase in human capital μ, adjusted by the rate at which the worker learns from experience, α. The discrepancy between ρ and μ comes from the observation that the former corresponds to the benefit associated to a *temporary* marginal increase in the stock of human capital. But in this model of such a marginal increase disappears only gradually, and thus the shadow value of human capital is given by μ. Since the informational rent that the worker acquires through the human capital channel ρ decreases as information improves, μ is also decreasing. However, as long as there is non-zero residual uncertainty and the marginal productivity of effort in the human capital technology is bounded away from zero, μ will never vanish.

Finally, a comment on the incentives to acquire human capital before entering the job market. The worker's expected discounted benefits from entering it with a human capital level of size $\overline{\theta}$ are given by $\mathcal{V}(0,\overline{\theta}) = \eta_0(\alpha,\phi) + \mu_0(\phi)\overline{\theta}$, where I made explicit the dependence of the initial values η_0 and μ_0 on $\alpha = (\alpha_t)_{t\geq 0}$ and ϕ. Given the long-term effects that it has on the worker's skills, the human capital depreciation rate ϕ indeed affects the worker's schooling choice (the rate at which he learns from work experience is irrelevant at this stage). Suppose the "future worker" makes a static choice on how much education to acquire. The cost of education is given by an increasing and differentiable function $c(e)$ while the benefits from it by an increasing and concave function $\overline{\theta}(e)$. The first-order condition of this problem yields $\mu_0(\phi)\overline{\theta}'(e^*(\phi)) - c'(e^*(\phi)) = 0$, from where we see that $e^*(\phi)$ is be decreasing. That is, individuals that suffer from less human capital depreciation would choose a higher level of education before entering the labor market.

1.5.3 Predicted Path of Wages

Recall that the wage process takes the form $w_t = m_t^* + a_t^1$, $t \geq 0$, where effort either strictly decays over time (away from steady-state) or it remains constant (in a stationary-learning environment). In equilibrium, the posterior mean evolves as

$$dm_t^* = (\overline{\theta}_t(a^*) + \kappa m_t^*)dt + \frac{\gamma_t}{\sigma_\xi} dZ_t^{a^*},\ t \geq 0, \tag{1.26}$$

where Z^{a^*} is a Brownian motion. As a consequence, and from *both parties' perspectives*, wages will, on average, decay in the fixed-skills case and martingale model ($\kappa = \overline{\theta} = 0$).

Some of the human capital accumulation models presented here are able to generate wages with endogenous positive drift. For example, as direct corollary of Proposition 1.3, if effort across tasks were not perfectly substitutable in the cost function, workers would focus both in human capital accumulation and in final-goods production. In such settings, the endogenous accumulation of skills may offset the negative effect that increasingly precise information has on incentives. Also, the strong complementarity case predicts increasing wages close to the stationary-learning level as long as the agent is indeed accumulating human capital. Moving away from steady-state the fast dynamics of the learning structure used here prevent the model from generating concave wages. Nevertheless, the presence of contractual rigidities that prevent the downward adjustment of wages may still give the model validity in explaining the path of wages over the life-cycle through the well-known human capital accumulation channel (Mincer (1974); Ben-Porath (1967)).[17] Moreover, the model is still able to explain the evolution of wages in environments where firms have relatively accurate estimates of workers' skills, in settings where the reputational component of wages is the dominating one (see Dewatripont et al. (1999) for a multi-task analysis motivated by the case of government agencies), or in environments where signal manipulation is not possible.

It is important to emphasize that (1.26) does not correspond to the process actually generating observations from the reputational component of wages. In fact, because of learning, the distribution of wages under the subjective probability measure differs from the distribution of wages under the true data-generating process. More specifically, it is always the case that beliefs locally *mean-revert* towards the contemporaneous *true value of skills* at every point in time, regardless of the specification used for skills. To understand this, observe that under the true probability measure, output is driven by $d\xi = (a_t^* + \theta_t)dt + \sigma_\xi dZ_t^\xi$, so that $\Delta_t := m_t^* - \theta_t$, $t \geq 0$, evolves according to

$$d\Delta_t = -\delta_t \Delta_t + \beta_t dZ_t^\xi - \sigma_\theta dZ_t^\theta,\ t \geq 0, \tag{1.27}$$

implying that whenever the market is optimistic or pessimistic about the worker's ability, the arrival of new information will tend to eliminate this bias.

Traditional pure-learning models (learning about a fixed unobservable) have predicted that increments in wages should be uncorrelated. However, this is only true

[17] In a pure learning model Harris and Holmström (1982) provide an alternative explanation for wages that increase with tenure. When workers are risk averse, firms learn about workers' abilities and long-term contracts can be written, workers pay a risk premium in order to insure themselves from future low-outcome realizations. This premium decreases with the precision of skills' estimates, so senior workers have higher salaries on average.

from the agents' perspective in the economy– that is, given their limited information sets. The econometrician would observe wages coming from the true data generating process, and, because of mean-reversion, some persistence in wages would be observed. This point is in fact not new. In an asset pricing context, Lewellen and Shanken (2002) show how learning about fundamentals can generate predictability in stock returns even when stock prices follow martingales from the market participants' perspectives.

There is a vast literature studying the covariance structure of earnings using data coming from different sources. Relevant to my work are the results from Kahn and Lange (2011), who give empirical support to the idea of combining workers' evolving productivity and employer learning as a way to explain some observed patterns in wages. In independent work, they show how such a model does a better job at matching second moments of wages than a pure-productivity model or a pure-learning model can do on their own.[18] Finally, as any model involving Brownian shocks would, the model predicts that the variance of wages increases over time. More interestingly, by allowing skills to have residual uncertainty, the variance of changes in wages does not decay as fast as it would in a pure learning model with fixed skills (see Farber and Gibbons (1996)). By varying the shock persistence parameter κ, the speed at which the variance of these increments decays over time changes.

1.6 Conclusions

This chapter developed a flexible dynamic model of career concerns involving Gaussian skills and on-the-job experience accumulation. It contributes to the labor markets

[18] From a modeling perspective my setting differs from theirs along two lines. First, the model adds a strategic component to the combined framework of workers' evolving productivity and employer learning. Second, I study AR(1) skills processes with a partially endogenous drift, while they focus on the random walk case with an exogenous growth trend.

literature by studying the effects that learning, evolving skills and human capital accumulation have on incentives. These issues interact significantly in defining the costs and benefits associated to workers' careers. The use of continuous-time techniques is of major importance for elucidating the forces that shape reputation-driven incentives in competitive markets.

I emphasized the primary role that the environment plays for influencing the incentives that arise from career concerns. The persistence of shocks to productivity determines the size of the monetary gains that arise from belief-distortion in stationary-learning settings. This is because shock persistence dictates the overall responsiveness of beliefs to aggregate information. As productivity shocks become more persistent, beliefs become more responsive and, therefore, workers exert more effort. More interestingly, under-provision of effort is not the unique long-run outcome, since inefficiently high effort is an optimal strategy for patient workers in highly unstable environments.

The possibility to secretly invest in human capital creates a new belief-manipulation channel that workers can exploit. Even though workers internalize the full benefits from human capital accumulation, only a fraction of them actually have a reputational value. This is because of the interplay between learning and market competition. Learning about skills allow firms to anticipate part of the additional benefits associated with investments in human capital. Market competition in turn forces firms to incorporate the expected value of these gains as an ex-ante premium in the workers' wage processes. As a result, only the unanticipated component of the monetary benefits associated with human capital accumulation determines the worker's marginal private value for human capital. Inefficiently low investments are expected to be a robust finding and I confirm this in two classes of human capital accumulation technologies.

There are several tangential issues not presently addressed by the present model. With respect to turnover, building a model that incorporates learning, separations, endogenous skills accumulation and moral hazard may seem an attractive challenge, it is unclear whether new substantial insights can be obtained from such a complex structure. Also, I have avoided the analysis of career concerns in the presence of complementarities between skills and effort in the output signal. This has proven itself to be a particularly challenging question, especially because endogenous information asymmetries play a non-trivial role. These and other interesting questions are beyond the scope of this chapter, and are left for future research.

1.7 Appendix A: Proofs

Proof of Proposition 1.1: Since a^* is an $\mathbb{F}^{\xi}-$progressively measurable process, the result is a direct application of Theorem 7.12. in Liptser and Shiryaev (1977).

□

Proof of Proposition 1.3: Suppose that the market conjectures that the manager will follow a deterministic strategy $a^* := (a_t^*)_{t\geq 0}$. Since wages take the form $w_t = m_t^* + a_t^*$, only $(m_t^*)_{t\geq 0}$ matter for incentives. Also, the fact that for each $t \geq 0$, $\theta_t(\cdot)$ is a deterministic functional of paths of the form $(y_s : 0 \leq s \leq t)$, implies that the trajectory of human capital conjectured by the market, $(\overline{\theta}_t(a^*))$, is fixed at time zero and unaffected by the worker's effort choice. The market's beliefs evolve according to

$$dm_t^* = (\overline{\theta}_t(a^*) + \kappa m_t^*)dt + \frac{\gamma_t}{\sigma_\xi} \underbrace{\frac{d\xi_t - (m_t^* + a_t^*)dt}{\sigma_\xi}}_{dZ_t^{a^*}} \tag{1.28}$$

where γ_t follows the dynamic (1.9) and Z^{a^*} is a Brownian motion from the market's perspective. The solution to the above SDE is given by

$$m_t^* = e^{-\int_0^t \delta_s ds} m_0 + \int_0^t e^{-\int_s^t \delta_u du} [\overline{\theta}_s(a^*)ds + \beta_s(d\xi_s - a_s^* ds)] \tag{1.29}$$

where $\beta_t := \frac{\gamma_t}{\sigma_\xi^2}$ and $\delta_t := \beta_t - \kappa$ for all $t \geq 0$. Since from the worker's perspective $(\overline{\theta}_t(a^*))_{t\geq 0}$ and a^* are exogenously given, incentives are determined only by

$$G_t := \int_0^t e^{-\int_s^t \delta_u du} \beta_s d\xi_s. \tag{1.30}$$

Let $(m_t)_{t\geq 0}$ denote the worker's posterior belief process of his own talent when he follows any strategy $a := (a_t)_{t\geq 0}$. Assume that it evolves according to an SDE of the form

$$dm_t = (\overline{\theta}_t(a) + \kappa m_t)dt + \sigma_t dZ_t \tag{1.31}$$

where $Z := (Z_t)_{t\geq 0}$ is a Brownian motion from the worker's standpoint. Moreover, the process $Z_t^a := \frac{1}{\sigma_\xi}\left(\xi_t - \int_0^t (a_s + m_s)ds\right),\ t \geq 0$ is also a Brownian motion from his perspective and is correlated to Z. By Lemma 1.1 we can write output from the worker's perspective as

$$d\xi_t = (m_t + a_t)dt + \sigma_\xi dZ_t^a,\ t \geq 0. \tag{1.32}$$

Inserting this into the expression for G_t gives us how the worker evaluates belief-distortions on the market's side,

$$\begin{aligned} G_t &:= \int_0^t e^{-\int_s^t \delta_u du} \beta_s [(m_s + a_s)ds + \sigma_\xi dZ_s^a] \\ &= \int_0^t e^{-\int_s^t \delta_u du} \beta_s \left[e^{\kappa s} m_0 + \int_0^s e^{\kappa(s-u)} (\overline{\theta}_u(a)du + dZ_u) + a_s ds + \sigma_\xi dZ_s^a \right] \end{aligned}$$

where I used that $m_s = e^{\kappa s}m_0 + \int_0^s e^{\kappa(s-u)}(\overline{\theta}_u(a)du + dZ_u)$, $s \geq 0$. The first term in G_t is unaffected by the effort decision so the worker's optimization problem is reduced to

$$\max_{a\in\mathcal{A}} \mathbb{E}^a \left[\int_0^\infty e^{-rt}\left(\int_0^t e^{-\int_s^t \delta_u du}\beta_s \left\{\int_0^s e^{\kappa(s-u)}(\overline{\theta}_u(a)du + dZ_u) + a_s ds + \sigma_\xi dZ_s^a\right\} - g(a_t)\right) dt\right]$$

For any strategy $a \in \mathcal{A}$ that the worker follows, Z^a and Z are exogenous Brownian motions. Moreover, since of any initial condition γ_0, $(\beta_t)_{t\geq 0}$ and $(\delta_t)_{t\geq 0}$ are uniformly bounded, all the stochastic integrals above will have zero expectation. As a consequence, the problem is reduced to

$$\max_{a\in\mathcal{A}} \mathbb{E}^a \left[\int_0^\infty e^{-rt}\left(\int_0^t e^{-\int_s^t \delta_u du}\beta_s \left\{\int_0^s e^{\kappa(s-u)}\overline{\theta}_u(a)du + a_s ds\right\} - g(a_t)\right) dt\right] \quad (1.33)$$

Integration by parts and the fact that $\delta_t = \beta_t - \kappa$ yield

$$\begin{aligned}&\int_0^t e^{-\int_s^t \delta_u du}\beta_s \int_0^s e^{\kappa(s-u)}\overline{\theta}_u(a)du \\ &= e^{\kappa t}\int_0^t e^{-\kappa s}\overline{\theta}_s(a)ds - e^{-\int_0^t \delta_s ds}\int_0^t e^{\int_0^s \delta_u du}\overline{\theta}_s(a)ds\end{aligned}$$

With this in hand, the manager's objective function has 3 integrals of the form (up to multiplicative constants)

$$I := \int_0^\infty e^{-rt}\left[e^{-\int_0^t \tau_s ds}\int_0^t e^{\int_0^s \tau_u du}\nu_s\right] dt$$

where $\tau = \delta$ or $-\kappa$ and $\nu = a$ or $\overline{\theta}(a)$. Since in any case $r+\tau > 0$, a direct application of Fubini's theorem implies that

$$I = \int_0^\infty e^{\int_0^t \tau_t dt}\nu_t \int_t^\infty \exp^{-\int_0^s (r+\tau_u)du} dsdt = \int_0^\infty e^{-rt}\nu_t \int_t^\infty e^{-\int_t^s (r+\tau_u)du} dsdt$$

Defining $\rho_t = \frac{1}{r-\kappa} - \lambda_t$ and $\lambda_t := \int_t^\infty e^{-\int_t^s (r+\delta_u)du} ds$, the worker will solve

$$\mathcal{P} := \begin{cases} \max_{a \in \mathcal{A}} \int_0^\infty e^{-rt} \left[\beta_t \lambda_t a_t - g(a_t) + \rho_t \overline{\theta}_t(a)\right] dt \\ s.t. \quad \overline{\theta}_t(a), \ t \geq 0, \\ \qquad \gamma_0 \geq \gamma^*, \end{cases}$$

concluding the proof of Lemma 1.3.

□

Proof of Proposition 1.2: We only need to show that $l_t := \beta_t \lambda_t$ is decreases over time, where $\beta_t = \gamma_t / \sigma_\xi^2$, $\lambda_t = \int_t^\infty e^{-\int_t^s (r+\delta_u)du} ds$ and $\delta_t = \beta_t - \kappa$, $t \geq 0$. Observe that

$$\frac{d\log(l(t))}{dt} = \frac{\dot{\gamma}_t}{\gamma_t} + r + \frac{\gamma_t}{\sigma_\xi^2} - \kappa - \frac{1}{\int_t^\infty e^{-\int_t^s (r+\gamma_u/\sigma_\xi^2 - \kappa)du} ds} \tag{1.34}$$

Suppose $\gamma_t > \gamma^*$, which occurs if and only if $\dot{\gamma}_t \leq 0$. Then,

$$\lambda_t = \int_t^\infty e^{-\int_t^s (r+\gamma_u/\sigma_\xi^2 - \kappa)du} ds < \frac{1}{r + \gamma^*/\sigma_\xi^2 - \kappa},$$

implying that

$$\frac{d\log(l(t))}{dt} < \frac{\dot{\gamma}_t}{\gamma_t} + \frac{\gamma_t}{\sigma_\xi^2} - \frac{\gamma^*}{\sigma_\xi^2}$$

Finally, from the ODE that governs γ_t, (see (1.9)), it can be observe that $\dot{\gamma}_t/\gamma_t + \gamma_t/\sigma_\xi^2 = 2\kappa + \sigma_\theta^2/\gamma_t$, so

$$\frac{d\log(l(t))}{dt} < 2\kappa + \frac{\sigma_\theta^2}{\gamma_t} - \frac{\gamma^*}{\sigma_\xi^2} < 2\kappa + \frac{\sigma_\theta^2}{\gamma^*} - \frac{\gamma^*}{\sigma_\xi^2} = 0$$

by definition of γ^*. When $\dot{\gamma}_t \geq 0$ (and so $\gamma_t < \gamma^*$ for all $t \geq 0$) an analogous argument shows that l_t increases over time (the above inequalities just reverse). This concludes the proof.

□

Proof of Corollary 1.1: The only non-trivial assertion in the corollary is the the first part of (i) when $\kappa > 0$. I will show in fact that the derivative is strictly positive for all $\kappa \in \mathbb{R}$. Recall that $\beta^* = \beta^*(\kappa)$, $\delta^*(\kappa) = \beta^*(\kappa) - \kappa = \sqrt{\kappa^2 - \sigma_\theta^2/\sigma_\xi^2} > 0$. It is easy to see that

$$\frac{da^*(\kappa)}{d\kappa} = \frac{(\delta^*(\kappa)+\kappa)(r+\delta^*(\kappa)-\kappa)}{\delta^*(\kappa)(r+\delta^*(\kappa))^2} > 0$$

where the last inequality comes from $\delta^*(\kappa) \pm \kappa > 0$ and $r > 0$.

□

Proof of Proposition 1.2: Comes from the same steps followed in the proof of Lemma 1.3.

□

Proof of Proposition 1.3: Omitting the dependence on time, the first order conditions of the worker's problem correspond to

$$\begin{aligned} \tilde{g}'(a^{1,*}+a^{2,*}) - \beta\lambda - \mu_1 + \mu_3 &= 0 \\ \tilde{g}'(a^{1,*}+a^{2,*}) - \rho - \mu_2 + \mu_3 &= 0 \end{aligned}$$

where $\mu_i \geq 0$ is the lagrange multiplier associated to the constraint $a_i \geq 0$, $i = 1, 2$, and $\mu_3 \geq 0$ the one corresponding to $a_1 + a_2 \leq R$. Therefore, the optimal effort allocation satisfies

$$\begin{aligned} a^{1,*} = (\tilde{g}')^{-1}(\beta\lambda) > 0,\ a^{2,*} = 0 \quad &\Leftrightarrow \quad \beta\lambda > \rho \\ a^{2,*} = (\tilde{g}')^{-1}(\rho) > 0,\ a^{1,*} = 0 \quad &\Leftrightarrow \quad \beta\lambda < \rho \end{aligned}$$

and any (a^1, a^2) s.t. $a^1 + a^2 = (\tilde{g}')^{-1}(\rho)$ when $\rho = \beta\lambda$.

(i) Suppose that $r-\kappa \geq 1$. This yields $\frac{\beta_t+1}{r+\beta_t-\kappa} \geq \frac{1}{r-\kappa}$, for all $t \geq 0$. If $\gamma_0 > \gamma^*$, β_t is strictly decreasing over time, implying that

$$(\beta_t+1)\lambda_t := (\beta_t+1)\int_t^\infty e^{-\int_t^s (r+\beta_u-\kappa)du} ds \underbrace{>}_{(*)} \frac{\beta_t+1}{r+\beta_t-\kappa}, \; t \geq 0.$$

Therefore, $(\beta_t+1)\lambda_t > \frac{1}{r-\kappa}$, and, recalling that $\rho_t := \frac{1}{r-\kappa} - \lambda_t$, this yields that $\beta_t\lambda_t > \rho_t$, $t \geq 0$. As a consequence, it is optimal for the worker to specialize in final-goods production. When $\gamma_0 = \gamma^*$, $\beta_t = \beta^*$ for all t and (*) becomes a weak inequality, so the result still holds.

(ii) Now assume that $r-\kappa < 1$. In steady state $\beta_t = \beta^*$ for all $t \geq 0$ and assume $\beta^* > 0$. Trivially, $(r-\kappa)\beta^* < \beta^*$, so

$$\beta^*\lambda^* = \frac{\beta^*+1}{r+\beta^*-\kappa} < \frac{1}{r-\kappa}.$$

This is equivalent to $\beta^*\lambda^* < \rho^*$, therefore showing that the agent specializes in human capital accumulation in steady-state. By continuity, there exists $T(\gamma_0) \geq 0$ sufficiently large s.t. $\beta_t\lambda_t - \rho_t < 0$ and thus $a_t^{2,*} > 0$ for all $t \geq T(\gamma_0)$.

To conclude, observe that when $r-\kappa = 1$, $\beta_0\lambda_0 - \rho_0 = (\beta_0+1)\lambda_0 - 1$. But, for $\gamma_0 > \gamma^*$, $(\beta_t)_{t\geq 0}$ is strictly decreasing, implying that

$$(\beta_0+1)\lambda_0 = (\beta_0+1)\int_0^\infty e^{-\int_s^t (\beta_u+1)du} ds > 1.$$

As a consequence, given $\gamma_0 > \gamma^*$, there exists $\epsilon > 0$ s.t. for $1-\epsilon < r-\kappa < 1$ we have

$$(\beta_0+1)\lambda_0 - 1 = (\beta_0+1)\lambda_0 - \frac{1}{r-\kappa} > 0,$$

implying that it is optimal to set $a_t^{1,*} > 0$ and $a_t^{2,*} = 0$ in the early stages of the agent's working life. This concludes the proof.

$\square$

Proof of Proposition 1.4: I will find a solution of the form $\mathcal{V}(t,x) = b_t + c_t x$. Plugging this in (HJ) we get

$$\sup_{u\in A}\left\{e^{-rt}[\beta_t\lambda_t u - g(u) + \rho_t x] + \frac{db_t}{dt} + \frac{dc_t}{dt}x + c_t[\alpha_t u - \phi x]\right\} = 0 \tag{1.35}$$

Impose that $(c_t)_{t\geq 0}$ satisfies the ODE $dc_t + (e^{-rt}\rho_t - \phi c_t)dt = 0$ with transversality condition $\lim_{t\to\infty} e^{-\phi t}c_t = 0$. Then,

$$c_t = e^{-rt}\underbrace{\int_t^\infty e^{-(r+\phi)(s-t)}\rho_s ds}_{\mu_t:=},\ t\in\mathbb{R}_+ \tag{1.36}$$

In fact, since $(\rho_t)_{t\geq 0}$ is bounded we get a stronger transversality condition: $c_t \to 0$ as $t\to\infty$. The (HJ) equation then becomes $\sup_{u\in A}\left\{e^{-rt}[\beta_t\lambda_t u - g(u)] + \frac{db_t}{dt} + c_t\alpha_t u\right\} = 0$ which yields the first order condition

$$g'(u_t^*) = \beta_t\lambda_t + \alpha_t\mu_t,\ t\in\mathbb{R}_+$$

and that b_t must satisfy $db_t = e^{-rt}[g(u_t^*) - g'(u_t^*)u_t^*]dt$. Observe that since $(u_t^*)_{t\geq 0}$ is bounded ($(\rho_t)_{t\geq 0}$ and $(\alpha_t)_{t\geq 0}$ are bounded and $(\beta_t\lambda_t)_{t\geq 0}$ is decreasing and non-negative) the last condition has as a solution

$$b_t := e^{-rt}\underbrace{\int_t^\infty e^{-r(s-t)}[g'(u_s^*)u_s^* - g(u_s^*)]ds}_{\eta_t:=},\ t\in\mathbb{R}_+$$

and moreover, $b_t \to 0$ as $t\to\infty$. Therefore, we have found a function $\mathcal{V}(t,x) = b_t + c_t x$ of class $C^1(\mathbb{R}_+\times\mathbb{R})$ such that it satisfies (HJ). We now need to show that is satisfies the transversality condition. To see this, fix an initial condition $x\geq 0$. For any feasible control a the path $t\mapsto \overline{\theta}_t^{t,x}(a)$ takes values in the interval $[0,\max\{\overline{\theta}, \ell K/\phi\}]$, where K

bounds $(\alpha_t)_{t\geq 0}$. As a consequence, the path of human capital remains bounded all the time. Because $b_t, c_t \to 0$ as $t \to \infty$, we trivially conclude that $\lim_{t\to\infty} b_t + c_t\overline{\theta}_t^{t,x}(a) = 0$.

Recall that $\rho_t = \frac{\kappa}{\kappa+r} - \kappa\lambda_t$ where

$$\lambda_t := \int_t^\infty \exp\left(-\int_t^s (r+\delta_u)du\right) ds,\ t \geq 0.$$

Direct calculations show that λ_t satisfies the ODE $d\lambda_t = ([r+\delta_t]\lambda_t - 1)dt$. If γ_t is decreasing over time ($\gamma_0 > \gamma^*$), so will be $\delta_t = \gamma_t^*/\sigma_\xi^2 + \kappa$, which implies that $\lambda_t > \frac{1}{r+\delta_t}$ for all $t \geq 0$. As a consequence λ_t is increasing and, furthermore, bounded above by $\frac{1}{r+\delta}$ (with $\delta := \gamma^*/\sigma_\xi^2 + \kappa$), so it converges. With this in hand, we conclude that the marginal benefit from an extra unit of human capital at time t, ρ_t, decreases over time and will also converge (it is bounded below by zero). Because of this,

$$\mu_t = \int_t^\infty e^{-(r+\phi)(s-t)}\rho_s ds < \frac{\rho_t}{r+\phi}$$

for all $t \geq 0$. Finally, observing that $(e^{rt}\mu_t)_{t\geq 0}$ satisfies the ordinary differential equation $d\mu_t = ([r+\phi]\mu_t - \rho_t)dt$, we conclude.

□

1.8 Appendix B: General Human Capital Technologies

It is not unreasonable to think that there are complementarities in the technology connecting investments in skills and their current stock or level. The problem with studying more general markovian diffusions of the form $d\theta_t = \mu(a_t,\theta_t)dt + \sigma(a_t,\theta_t)dZ_t^\theta,\ t \geq 0$ is that, even though the filtering equations associated to posterior moments of θ given ξ may exist, such a system may not be closed. A tractable way to incorporate such complementarities is by doing so in a deterministic way through an additional state variable. In this section I assume that for any feasible effort strategy $a \in \mathcal{A}$ (a concept to be defined immediately), $(\overline{\theta}_t(a))_{t\geq 0}$ is the solution to the ODE

$$d\overline{\theta}_t = f(t, \overline{\theta}_t, a_t)dt,\ \overline{\theta}_0 = \overline{\theta}^o \geq 0. \tag{1.37}$$

Since in the above specification the agent's optimal action will typically depend on his current human capital stock, we need to relax the equilibrium concept. Even though the worker's effort strategy cannot be correctly guessed once a deviation has taken place, the fact that human capital evolves deterministically allows the market to perfectly anticipate the on-equilibrium effort strategy.

Definition 1.5. *An effort strategy $a := (a_t)_{t\geq 0}$ is of the feedback form is it corresponds to a function $a_t = a(t, \overline{\theta})$, where $\overline{\theta}$ is the agent's stock of human capital at time $t \geq 0$, for some function $a : \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}_+$.*

Given $\overline{\theta}^o$, which I assume common knowledge, and a conjectured feedback control a^*, the market conjectures a human capital trend $(\overline{\theta}_t(a^*))_{t\geq 0}$. This trend is therefore fixed ex-ante and the question is whether a Bayesian nash equilibrium in feedback strategies exists:

Definition 1.6. *An equilibrium of this economy is a feedback effort strategy a^* and a wage process $w := (w_t)_{t\geq 0}$, such that:*

(i) Given a^*, the market sets a wage of the form $w_t = \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t^\xi] + a_t^*$ for all $t \geq 0$;

(ii) a^* is optimal for the manager given the wage process in (i):

$$\begin{aligned} a^* \in arg\max_{a\in\mathcal{A}} \quad & \mathbb{E}^a\left[\int_0^\infty e^{-rt}(w_t - g(a_t))dt\right] \\ s.t. \quad & w_t = \mathbb{E}^{a^*}[\theta_t|\mathcal{F}_t^\xi] + a_t^*,\ \forall t \geq 0. \end{aligned} \tag{1.38}$$

That is, I have eliminated the requirement that the market perfectly anticipates the worker's strategy after all private histories.

Given that the family $(\overline{\theta}(a^*))_{t\geq 0}$ is deterministic, Lemma 1.3 applies and the existence of an equilibrium is reduced to the existence of a feedback control to the optimal control problem

$$\begin{aligned} \max_{a\in\mathcal{A}} \quad & \int_0^\infty e^{-rt}\left[\beta_t\lambda_t a_t - g(a_t) + \rho_t\overline{\theta}_t(a)\right]dt \\ s.t. \quad & d\overline{\theta}_t = f(t,a_t,\overline{\theta}_t)dt,\ t>0,\ \overline{\theta}_0 = \overline{\theta}^o. \end{aligned}$$

I will look for optimal strategies in the following class of controls:

Definition 1.7. *(Feasible Control) A control is said to be feasible if it corresponds to a piecewise continuous*[19] *function of time $a:\mathbb{R}_+\to A:=[0,\ell]$. Denote the set of feasible controls by $\mathcal{A}$.*

Since we want to capture that the agent accumulates more human capital as he becomes more experienced (as measured by how engaged in production the worker has been), I assume that $a\mapsto f(t,\overline{\theta},a)$ is strictly increasing in for all $\overline{\theta}\geq 0$. More generally, this function is required to satisfy this very weak conditions:

Assumption 2. *It is assumed that*

(i) $f:\mathbb{R}_+\times\mathbb{R}\times A\to\mathbb{R}$ is such that $f(t,\cdot,a)\in\mathcal{C}^1(\mathbb{R})$ for all $a\in A$ and $f(t,\overline{\theta},\cdot)$ is differentiable for all $\overline{\theta}\in\mathbb{R}$.

(ii) For all $\overline{\theta}\geq 0$, the function $a\mapsto f(t,\overline{\theta},a)$ is strictly monotone.

[19] A function $\phi:\mathbb{R}\to A\subseteq\mathbb{R}$ is piecewise continuous if for any interval $[a,b]\subset\mathbb{R}$ there exists a finite set of points $a=t_0<t_1<...<t_n=b$ such that ϕ is continuous in $[t_0,t_1]$ and $(t_i,t_{i+1}]$, for $i=1,...,n-1$ and has a finite right hand limit for each t_i, $i=1,...,n$. This definition follows from Halkin 1974.

The following result gives necessary conditions that the any optimal control must satisfy. It is an application of Pontryagin's Maximum Principle for infinite horizon problems (see Halkin (1974)):[20]

Proposition 1.5. *Let $a^* \in \mathcal{A}$ be an optimal control and suppose $a \not\equiv 0, \ell$. Then, there exists a piecewise continuously differentiable function $q : \mathbb{R}_+ \to \mathbb{R}$ such that*

(i) For almost every $t \in \mathbb{R}_+$

$$dq_t = \left\{ q_t \left[r - \frac{\partial f}{\partial \overline{\theta}}(t, \overline{\theta}_t, a_t^*) \right] - \rho_t \right\} dt; \tag{1.39}$$

(ii) For every $t \geq 0$ such that $0 < a_t^* < \ell$, a_t^* satisfies

$$g'(a_t^*) = \beta_t \lambda_t + q_t \frac{\partial f}{\partial a}(t, \overline{\theta}_t(a^*), a_t^*) \tag{1.40}$$

where $\overline{\theta}_t(a^*)$ is the solution to $d\overline{\theta}_t = f(t, \overline{\theta}_t, a_t^*)dt$, $\overline{\theta}_0 = \overline{\theta} \geq 0$.

Proof: By the Pontryagin Maximum Principle for infinite horizon (Halkin (1974)), if $a^* := (a_t^*)_{t \geq 0}$ is an optimal control then there exists $\mu \geq 0$ and a piecewise continuously differentiable function[21] $q : \mathbb{R}_+ \to \mathbb{R}$ s.t.

I. $|(\mu, q_0)| \neq 0$;

II. $\dot{q}_t - rq_t = -\frac{\partial}{\partial x}\mathcal{H}(t, x, a_t^*, \mu, q_t)\big|_{x=\overline{\theta}_t}$, a.s.

III. $\mathcal{H}(t, \overline{\theta}_t, a_t^*, \mu, q_t) \geq \mathcal{H}(t, \overline{\theta}_t, a, \mu, q_t)$, for all $t \geq 0$, $a \in A$.

where the Hamiltonian $\mathcal{H}$ is defined by

$$\mathcal{H}(t, x, a, \mu, y) := \mu \left[a\beta_t \lambda_t - g(a) + \rho_t x \right] + y f(t, x, a).$$

[20] The characterization of Proposition 1.5 is also valid when f is decreasing in effort. That is, what we really require is f to be strictly monotone in effort.

[21] That is, a differentiable function which derivative is piecewise continuous

with ρ_t, β_t and λ_t as in the Proposition.

Now I will prove that, under the hypothesis of the proposition, $\mu \neq 0$. Replacing the expression for the Hamiltonian in II yields the ODE

$$\dot{q}_t = q_t \left[r - \frac{\partial f}{\partial \overline{\theta}}(t, \overline{\theta}_t(a^*), a_t^*)) \right] - \mu \rho_t \tag{1.41}$$

Recall that the set of times where the last ODE does not hold is the set of points at which a^* is discontinuous (at those points q is not differentiable). By definition of piece-wise continuity, for any $T > 0$ there is only a finite number of times less than T at which the optimal control is discontinuous. Therefore, II holds for intervals $[0, t_1]$, $\{(t_i, t_{i+1}] \mid i \in \mathbb{N}\}$ such that their union is the real line. Moreover, since q is continuous, it must that the solution of the above ODE at any subinterval $(t_i, t_{i+1}]$ must satisfy

$$q_{t_i^+} = q_{t_i}$$

where we understand that $q_{t_i^+}$ is the limit as t decreases to t_i of the solution to the ODE (1.41) in $(t_i, t_{i+1}]$ with final condition $q_{t_{i+1}}$, $i \geq 1$. The proof is based on the following

Lemma 1.4. *If $\mu = 0$, then either $a^* \equiv 0$ or $a^* \equiv \overline{a}$.*

Proof of the Lemma :

Suppose $\mu = 0$. Then it must be that the following relationship holds for $t \in [0, t_1^*]$:

$$q_t = q_0 \exp\left(\int_0^t \left[r - \frac{\partial f}{\partial \overline{\theta}}(s, \overline{\theta}_s(a^*), a_s^*) \right] ds \right)$$

where $s \mapsto \overline{\theta}_s(a^*)$ is the trajectory generated by $(a_s^*)_{s \in [0, t_1]}$. If $q_{t_1} = 0$, then $q_0 = 0$, contradicting I. Thus, $q_{t_1} \neq 0$ and therefore q cannot vanish in $[0, t_1]$. Suppose that $q > 0$ over this set. Then, the maximum condition II implies that the optimal control must satisfy

$$a_s^* \in \arg\max_{a \in A} \; q_s f(\overline{\theta}_s(a^*), a), \; \forall s \in [0, t_1]$$

But $q_s f(\overline{\theta}_s(a^*), \cdot)$ is increasing, and thus $a_s^* \equiv \overline{a}$ for all $s \in [0, t_1]$. As a consequence, whenever $\mu = 0$, if $q_{t_1} > 0$ then the optimal control takes the maximum possible value in the first interval. If in turn, $q_{t_1} < 0$ the same reasoning shows that the optimal control will take the minimum value over the same set, this because $q_s f(s, \overline{\theta}_s^*, \cdot)$ would be decreasing for all $s \in [0, t_1]$. In the remainder of the proof, I assume without loss of generality that $q_{t_1} > 0$ (the other case is analogous). If $q_{t_2} \leq 0$, then (1.41) in $(t_1, t_2]$ implies that $q_s \leq 0$ in the same interval. Therefore

$$q_{t_1^+} := \lim_{s \searrow t_1} q_s \leq 0 < q_{t_1}$$

contradicting the fact that q is continuous. Hence, $q_{t_2} > 0$ implying that q is strictly positive in $(t_1, t_2]$ and thus the optimal control must take value $\overline{a}$ over that interval. Proceeding inductively, if $q_{t_1} > 0$ then $q_{t_i} > 0$ for all $i = 0, 2, 3...$ and by the maximum condition $\overline{a}$ is the optimal control. The same reasoning allows us to conclude that when $q_{t_1} < 0$, $a^* \equiv 0$ must be optimal. This concludes the proof.

□

The previous Lemma shows that when an optimal control exists and is neither identically zero nor equal to $\overline{a}$, then $\mu > 0$. When this is the case it is clear that we can assume $\mu = 1$ (equivalently, redefine q as q/μ and note that q/μ satisfies all the conditions of the theorem). This proves part (i) in the proposition. Finally part (ii) is simply the necessary condition that an unconstrained optimum must satisfy. This concludes the proof.

□

The next result states that differentiability of the continuation-value function with respect to the state variable $\overline{\theta}$ ensures that effort is *always* above the career concerns benchmark. For this purpose define

$$\mathcal{V}[t,\overline{\theta}] \equiv \sup_{a\in\mathcal{A}} \int_t^{\infty} e^{-rs} \left[a_s\beta_s\lambda_s - g(a_s) + \rho_s\overline{\theta}_s^{\,t}(a)\right] ds \tag{1.42}$$

where $(\overline{\theta}_s^{\,t}(a))_{s\geq t}$ is the solution of the ODE (1.37). We have the following

Proposition 1.6. *Fix $t \geq 0$ and suppose $\overline{\theta} \in \mathbb{R}$ is such that $\mathcal{V}(t,\overline{\theta}) < +\infty$. Let $(a_s^*)_{s\geq t}$ be the continuation strategy that attains this value. Then, the continuation-value function is increasing in the state variable. Moreover, if $\mathcal{V}(t,\cdot)$ is differentiable in a neighborhood Θ of $\overline{\theta}$, then for any $s \geq t$ such that $\overline{\theta_s}(a) \in \Theta$, $q_s = e^{rs}\frac{\partial\mathcal{V}}{\partial\overline{\theta}}(s,\overline{\theta}_s(a^*)) > 0$.*

Proof of Proposition 1.6: If $\mathcal{V}(t,\overline{\theta^2}) = \infty$ the first part of the Proposition is trivially true. Suppose that is finite. Let $a^* := (a_s^*)_{s\geq t}$ be the optimal control (a function of time) that attains value $\mathcal{V}(t,\overline{\theta^1})$ when starting from the level $\overline{\theta^1} \geq 0$ at time t. In the same vein let $(\overline{\theta}_s(a^*;\overline{\theta}))_{s\geq t}$ denote the path of human capital generated by the feasible control a^* when starting from point $\overline{\theta} \in \mathbb{R}_+$ at time t, that is, the solution to

$$d\overline{\theta}_s = f(s,\overline{\theta}_s,a_s^*),\ s > t,\ \overline{\theta}_t = \overline{\theta}$$

Since the solutions of these two ordinary differential equations cannot cross (they differ only in the initial condition), it must be the that

$$\overline{\theta}_s(a^*;\overline{\theta^2}) > \overline{\theta}_s(a^*;\overline{\theta^1}),\ \forall s \geq t$$

which implies that $V[t,\overline{\theta^2};a^*] > V[t,\overline{\theta^1};a^*] = \mathcal{V}(t,\overline{\theta^1})$, so $\mathcal{V}(t,\cdot)$ is increasing. If, moreover, it turns out to be differentiable in a neighborhood Θ of $\overline{\theta^1}$ and standard perturbation analysis shows that $\frac{\partial\mathcal{V}}{\partial\overline{\theta}}(s,\overline{\theta}_s(a^*)) = \lambda_s$, as long as $\overline{\theta}_s(a^*) \in \Theta$, $s \geq t$, where λ_s is the multiplier associated to the dynamic that governs the ODE of human capital accumulation. Implicit in the formulation of Proposition 1.5 is the fact that $q_t = e^{rt}\lambda_t$, $t \geq 0$. This concludes the proof.

□

Chapter 2

Two-Sided Learning and Moral Hazard

2.1 Introduction

This chapter analyzes a class of continuous-time games in which a long-run agent and a population of small players (a *market*) learn about an unobserved state variable from a public signal that is subject to Brownian shocks. I refer to this unobserved state as the *fundamentals*. In these games, the long-run agent's payoffs are determined by the small players' actions, which in turn depend on their beliefs about the hidden state. Consequently, the long-run agent has the incentive to take costly actions that influence the market's learning process. Using continuous-time techniques, I characterize the long-run agent's behavior in any equilibrium in which the market is able to perfectly anticipate such strategic motives.

Strategic behavior in the context of two-sided learning about an unobserved state variable appears in many economic environments. In labor markets, workers can exert effort so as to manipulate the beliefs of potential employers about their unobserved skills (Holmstrom (1999)). As another example, when customers perceive a firm's product to be of high quality, the firm can reduce its technological investments without suffering considerable reputational losses (Board and Meyer-ter-Vehn (2010a)). Strategic motives also appear in procurement, when a contractor and a government learn about the contractor's efficiency to deliver goods and the contractor desires to sell these goods to the government repeatedly over time. Similarly in monetary policy, when a central bank generates inflationary surprises by affecting an inflation trend that is unobserved by all the agents in the economy.

Dynamic incentives in the context of learning, moral hazard and market interactions have been studied only under perfect competition and risk neutrality. However, these assumptions cease to be appropriate in many interesting settings. For instance, labor market frictions can result in returns to skills that vary nonlinearly across workers' perceived distribution of abilities. In dynamic procurement, the rent that a contractor obtains can be a non-monotone function of perceived efficiency if the government purchasing these goods lacks commitment. In monetary policy, a policymaker may suffer losses from an unobserved inflation trend which moves away from a specific target. Similarly, a firm that invests in product quality can face convex costs of adjustment. Developing methods to understand and quantify the incentives that arise in environments without perfect competition and linearity is thus an important task.

In the class of games I analyze, the market perfectly anticipates the long-run agent's actions when public and private beliefs are aligned. However, the long-run agent's incentives *on the equilibrium path* are determined by the benefits from hypothetical deviations *off the equilibrium path.* Evaluating the long-run agent's off-equilibrium payoffs is challenging because, after a deviation takes place, the long-run agent acquires private information about the evolution of the fundamentals and the market's belief becomes biased. The long-run agent can thus condition his actions on the values that both private and public beliefs take, and the market will construct beliefs using a wrong conjecture about equilibrium play.

I focus on pure-strategy Markov equilibria in which the relevant state variables are the agent's private belief about fundamentals, and the *belief-asymmetry process*, which is a measure of the degree of discrepancy between private and public beliefs. In such a context, I develop a first-order approach for quantifying the value of local deviations off the equilibrium path. In particular, I show that in any equilibrium the value that the long-run agent attaches to inducing a small degree of asymmetry between private and public beliefs is the solution to an ordinary differential equation, which I refer to as the *incentives equation*. The incentives equation captures how the incentives to distort public beliefs vary across different levels of public opinion. Such variations occur when learning, preferences or the market's actions are *nonlinear*.

This necessary condition is obtained under very weak assumptions. The long-run agent's actions can affect the public signal directly (*signal-jamming* games) or instead, affect the evolution of the fundamentals itself (*investment* games with learning). The costs of the long-run agent's actions are required to be convex. The impact of the market's belief on the long-run agent's payoff can be completely general. Two assumptions are nonetheless crucial. First, I restrict to learning processes that admit posterior distributions summarized by a single one-dimensional state variable.[1] Second, the long-run agent's actions must enter additively into the corresponding dynamics. The first modeling assumption is purely for tractability reasons. The second one is economically important, as it prevents the rise of experimentation incentives which might conflict with belief-manipulation motives.

The incentives equation corresponds to an Euler equation subject to an equilibrium condition. Its Euler feature states that the long-run agent must be indifferent to

[1]More specifically, I use both the Kalman-Bucy filter under stationary learning and the Wonham filter for studying signal-jamming games. In analyzing investment with learning, I restrict the analysis to the Kalman-Bucy filter.

the choice between exerting the last unit of belief manipulation now, or postponing it to an instant later. The equilibrium condition in turn requires that this last unit of belief manipulation be perfectly anticipated by the market whenever beliefs are aligned.

The long-run agent has incentives to distort public beliefs because it allows him to obtain an immediate marginal flow gain and to smooth out the costs of his actions due to anticipated changes in the value of belief asymmetry. On the other hand, by postponing belief manipulation the long-run agent avoids the depreciation costs resulting from any stock of belief asymmetry gradually depreciating over time. The rate at which any stock of belief asymmetry depreciates over time in turn depends on the inherent rate of decay of the learning process, and also on the market's conjecture about equilibrium play. This is because the market's expectation of belief manipulation determine the threshold beyond which the long-run agent's actions effectively induce more belief asymmetry. Finally, since the market's expectation of belief manipulation are correct in equilibrium, the rate of return on belief asymmetry turns out to be determined endogenously by the long-run agent's equilibrium strategy.

The equilibrium properties uncovered by this chapter offer novel insights on the structure behind the incentives for belief manipulation. In the linear and additive model of career concerns developed by Holmstrom (1999), workers with the same tenure exert identical effort levels at every point in time, despite their differences in perceived abilities or wages. The incentives equation in turn predicts that any source of nonlinearity will transform perceived ability (hence, wages) into an important determinant of a worker's incentives. In Board and Meyer-ter-Vehn (2010a,b), linear costs of investment generate bang-bang investment policies and thus considerable pooling across different levels of a firm's reputation. Instead, the convex costs of adjustment that I analyze induce investment plans that are gradually executed over time. Finally, in the career concerns model of Bonatti and Hörner (2011), the coarse information structure imposed in that paper prevents the rise of any intra-temporal incentives. On the contrary, when information arrives continuously over time, I show that it is the local change of the market's expectation of belief manipulation with respect to small changes in public beliefs what captures the connection between contemporaneous incentives and conjectures.[2]

The incentives equation, as a necessary condition for on-the-equilibrium path behavior, does not ensure that the long-run agent benefits from large deviation away from the market's beliefs. In order to verify incentive compatibility globally, it is necessary to study the long-run agent's value function off the equilibrium path. However, since the market always constructs beliefs as if the long-run agent had never deviated, the long-run agent's value function satisfies a new type of partial differential equation (PDE) characterized by its particular *nonlocal* structure: the local evolution of the long-run agent's utility *off* the equilibrium path depends on the marginal value of belief asymmetry *along* the equilibrium path. This new class of PDEs correspond to a standard Hamilton-Jacobi-Bellman (HJB) equations that also satisfies the condition that the market must perfectly anticipate the long-run agent's optimal action on the equilibrium path. Since the latter requirement corresponds to a fixed point condition on the value function to be found, verification theorems apply.

I use this crucial insight to show the existence of markovian equilibria for a class of games with linear-quadratic structure: linear learning (Gaussian) and quadratic payoffs. For this class of games, I find closed-form solutions to the PDEs that summarize on- and off-path incentives. Equilibrium actions turn out to be linear in public beliefs, which results in realized actions that are stochastic. Moreover, these equilibria exhibit all the forces that are expected to influence behavior in more nonlinear settings: endogenous rates of return on belief asymmetry and cost-smoothing motives.

[2]Dewatripont et al. (1999)) show that effort and contemporaneous conjectures are strategic complements when skills and effort are complements in the technology that a worker has access to. In the class of games I analyze, such technological complementarities are absent.

Off the equilibrium path, the belief asymmetry process evolves deterministically and vanishes asymptotically.

I exploit the tractability of this linear-quadratic framework in two applications. I first show that in economies where agents learn about an unobserved inflation trend, the size of the inflationary bias created by a central bank that lacks commitment is severely limited by the volatility of both shocks to prices and shocks to trend inflation. I then show how effort smoothing can exacerbate the ratchet effect in contexts where a contractor and a government who lacks commitment learn about a firm's unobserved efficiency of providing goods.

Finally, the methods delivered in this chapter are also applicable to other situations in which the nonlocal PDEs are hard to visualize. This is because any solution to the incentives equation, as a necessary condition for equilibrium incentives, can be used as a valid guess in the dynamics of belief asymmetry in order to solve a standard (local) PDE arising from stochastic control. Numerical comparisons between the initial guess and the policy delivered by the PDE would then validate or eliminate the former as a candidate of Markov perfect equilibrium. While this numerical approach is not followed here, it is a feasible avenue to characterize off-equilibrium incentives in other nonlinear environments.

2.1.1 Literature

In the career concerns literature, Holmstrom (1999) develops a linear and additive model of career concerns in which equilibrium effort is independent of current wages, of future effort and independent of the market's contemporaneous conjectures about effort. In the static model of Dewatripont et al. (1999), introducing complementarities between effort and skills in a worker's output technology generates strategic complementarity between realized effort and conjectured effort. More recently, in the context of coarse information, Bonatti and Hörner (2011) show how current and future effort can become strategic substitutes. The incentives equation I derive here instead predicts that a worker's reputational incentives are affected simultaneously by the current level of wages he receives, by how these wages are expected to evolve over time and by the market's contemporaneous expectations of belief manipulation.

In the context of investment games with belief manipulation, Board and Meyer-ter-Vehn (2010a) characterize a firm's equilibrium investment policy in a context of unobserved product quality, perfectly informative signals and stochastic output technologies. They obtain an expression for the value attached to belief asymmetry in integral form, as valuations in their model can jump. Board and Meyer-ter-Vehn (2010b) instead study a firm's incentives to influence the public belief about a product's quality when the latter is private information to the firm. Because investment costs are linear in both models, there is considerable pooling across different levels of firms' reputations.

In the class of games analyzed in this chapter, the long-run agent's equilibrium incentives incorporate his actions' impact on the "standard" of belief manipulation he will face in the near future. A classic example of such strategic motive is captured by the ratchet effect (Weitzman (1980), Laffont and Tirole (1988) and (1990)). Martinez (2009) finds ratchet effects when studying the career-concerns incentives in the presence of piecewise linear wages. The general class of games I study in fact exhibit "ratchet forces": everything else equal, the higher the local change in the market's conjecture about the level of belief manipulation, the lower the long-run agent's incentives to improve the perception of fundamentals. However, since ratchet forces, cost smoothing and marginal flow benefits interact simultaneously to shape the long-run agent's incentives, the ratchet effect may not always be observed in equilibrium.

To conclude, this chapter also relates to the literature studying dynamic incentives using continuous-time techniques. In Sannikov (2007) and (2008) these techniques are applied to settings of continuous-time games under imperfect monitoring and to agency problems, respectively. Faingold and Sannikov (2011) study reputation dynamics in settings with imperfectly observable actions and one-sided learning about a long-run agent's fixed type. Their characterization of equilibrium incentives is in the form of ODEs subject to a fixed point condition. A similar characterization is found by Bohren (2012) in the context of investment games with imperfectly observable actions and without learning. Finally, Williams (2011) studies agency problems in the presence of persistent private information, deriving necessary and sufficient conditions for incentive compatible contracts. Although his analysis is focused on adverse selection, it shares the feature that whenever unobserved payoff-relevant variables exhibit persistence, an agent's actions can have a long-term impact on his utility.

2.2 Signal-Jamming Games: General Case

2.2.1 Set-up

Consider an economy in which one long-run agent and a population of small players – the *market*– simultaneously learn about a hidden state variable. The unobserved process is denoted by $\theta := (\theta_t)_{t\geq 0}$ and takes values in $\Theta \subseteq \mathbb{R}$. Hereinafter I refer to the hidden state variable as the *fundamental*.

In a signal-jamming game the fundamental is exogenous, and the long-run agent affects the signal from which the market extracts relevant information about the unobserved state.[3] More specifically, there is *public* signal $\xi := (\xi_t)_{t\geq 0}$ that takes the form

$$d\xi_t = (a_t + \theta_t)dt + \sigma_\xi dZ_t,\ t \geq 0, \tag{2.1}$$

[3]The term "signal-jamming" was coined by Fudenberg and Tirole (1986) in the context of firms that take unobserved actions (e.g., price cutting) in order to deter entry in the markets where they operate.

where $Z := (Z_t)_{t \geq 0}$ is a Brownian motion and $\sigma_\xi > 0$ is a volatility parameter. The term a_t represents the degree of signal manipulation exerted by the agent at time $t \geq 0$. The long-run agent's manipulation choices are not observed by the rest of the economy and take values in a set $A \subset \mathbb{R}$.[4]

Since the long-run agent observes his past actions, he also observes the component of the public signal that is not explained by signal manipulation, $Y_t := \xi_t - \int_0^t a_s ds,\ t \geq 0$. By definition

$$dY_t = \theta_t dt + \sigma_\xi dZ_t,\ t \geq 0, \tag{2.2}$$

from where we can see that Y is an exogenous process that is privately observed by the long-run agent. Equations (2.1) and (2.2) yield that signal manipulation has an *additive* structure. In the sequel, $\mathcal{F}_t^Y$ denotes the information generated by the process Y up to time t, and $\mathbb{F}^Y := (\mathcal{F}_t^Y)$ the (completed) filtration associated to Y. The corresponding analogous notation is used to denote the filtration generated by ξ.

The long-run agent uses the information conveyed by Y to construct private beliefs about θ. The *private beliefs process* is denoted by $\rho := (\rho_t)_{t \geq 0}$ where

$$\rho_t(x) := \mathbb{P}(\theta_t \leq x | \mathcal{F}_t^Y),\ x \in \Theta,\ t \geq 0. \tag{2.3}$$

In this definition, $\mathbb{P}(\cdot | \mathcal{F}_t^Y)$ corresponds to the agent's posterior belief about θ_t, given the observations $(Y_s : s \in [0,t])$, $t \geq 0$.

The market instead uses the information conveyed by the public signal ξ for the same purpose. The associated *public belief* process is denoted by $\rho^* := (\rho_t^*)_{t \geq 0}$ and defined analogously by

$$\rho_t^*(x) := \mathbb{P}^*(\theta_t \leq x | \mathcal{F}_t^\xi),\ x \in \Theta,\ t \geq 0. \tag{2.4}$$

[4]Observe that signal structure (2.1) satisfies the full-support assumption with respect the the long-run agent's actions.

The term $\mathbb{P}^*(\cdot|\mathcal{F}_t^\xi)$ denote the market's (subjective) posterior belief about θ_t given the partial observations $(\xi_s : s \in [0,t])$, $t \geq 0$. Private and public beliefs coincide at time t if, for instance, the market has perfectly anticipated the agent's hidden manipulation strategy over the time interval $[0,t]$. Nevertheless, moral hazard can give rise to potential divergence between private and public beliefs.

The small players act myopically because their individual actions are anonymous and they do not affect the population average observed by the long-run agent. Thus, at any instant of time the small players maximize their ex-ante flow payoffs. Ex-ante flow payoffs in turn depend on the small player's beliefs about the current value of the fundamental θ_t and on the action that they conjecture the agent is currently following, a_t^*, $t \geq 0$. I summarize the small players' best-response action through the function

$$\mathtt{b}(\rho_t^*, a_t^*).$$

The long-run agent cares about fundamentals because they influence his payoff through the impact that public beliefs have on the market's actions. Given a market's action profile $b := (b_t)_{t\geq 0}$ and a signal manipulation strategy $a := (a_t)_{t\geq 0}$, the long-run agent's discounted payoffs at time t take the form

$$\int_t^\infty e^{-r(s-t)}(u(b_s) - g(a_s))ds,\ t \geq 0. \tag{2.5}$$

The function $u : \mathbb{R} \to \mathbb{R}$ represents the component of the agent's flow utility that is determined by the market's actions. The costs of signal manipulation are given by a function $g : A \to \mathbb{R}$ which is convex in $a \in A$. For simplicity, I assume that g is differentiable everywhere.

The additive signal structure implies that the only source of asymmetric information in the model comes from the imperfect observability of the long-run agent's actions. As a result, it is natural to define an equilibrium concept entailing a market that perfectly anticipates the agent's strategy:

Definition 2.1. *An equilibrium consists of (i) a manipulation strategy of the long-run agent* $a_t(\xi_s, s \in [0,t], \rho_t)$, $t \geq 0$*; (ii) a public strategy of the market* $b_t(\xi_s, s \in [0,t])$, $t \geq 0$*; (iii) a private belief process of the long-run agent* $\rho_t(\cdot|Y_s, s \in [0,t])$, $t \geq 0$*; and (iv) a public belief process of the market* $\rho_t^*(\cdot|\xi_s : s \in [0,t])$, $t \geq 0$*, such that:*

(a) *Given the market's public strategy* $(b_t)_{t\geq 0}$ *and the long-run agent's belief process* $\rho := (\rho_t)_{t\geq 0}$*, the signal manipulation strategy* $a := (a_t)_{t\geq 0}$ *maximizes the agent's continuation payoff*

$$\mathbb{E}_t\left[\int_t^\infty e^{-r(s-t)}(u(b_s) - g(a_s))dt\right], \tag{2.6}$$

after any history $(\xi_s : s \in [0,t], a_s, s \in [0,t])$, $t \geq 0$*;*

(b) *The market's action at time* t, b_t, *is of the form* $\mathtt{b}(\rho_t^*, a_t^*)$*, where the latter is optimal given its belief* ρ_t^* *and the action* $a_t^*(\xi_s, s \in [0,t]) = a_t(\xi_s, s \in [0,t], \rho_t^*)$, $t \geq 0$*;*

(c) *The agent's beliefs* $\rho_t = \rho_t(\cdot|Y_s, s \in [0,t])$ *are determined by Bayes' rule;*

(d) *The market's beliefs* $\rho_t^* = \rho_t^*(\cdot|\xi_s, s \in [0,t])$ *are determined by Bayes' rule under the assumption that the agent has been playing* $a_t^*(\xi_s, s \in [0,t]) = a_t(\xi_s, s \in [0,t], \rho_t^*)$.

Observe that the agent could in principle follow any strategy $a := (a_t)_{t\geq 0}$ that is progressively measurable with respect to the information generated by ξ and Y. Yet, the histories of the form $(\xi_s : s \in [0,t], \rho_t)$, $t \geq 0$, summarize all the payoff relevant information for the agent, so we can restrict the analysis to strategies as in (i). I refer to them as (feasible) *manipulation* strategies. Part (ii) instead requires that the market's actions depend on the information generated by ξ only.

Concerning the equilibrium conditions, (a) states that the agent's strategy specifies actions both on and off the equilibrium path, and has to maximize future payoffs on and off the equilibrium path. However, the optimality of the market's actions is checked only on the equilibrium path (condition (b)). This is because the signal structure satisfies the full support assumption, which implies that any partial observation $(\xi_s : s \in [0,t])$, $t \geq 0$, is consistent with equilibrium play. Condition (c) and (d) correspond to the consistency requirements that both belief processes ρ and ρ^* must be constructed via Bayes' rule using the strategies specified by equilibrium play. In particular, (d) states that the market's beliefs are always constructed as if the agent were taking the actions prescribed along the equilibrium path.

2.2.2 Learning and Belief Manipulation

The purpose of this section is to present a unified framework for the learning and belief manipulation dynamics in environments where posterior beliefs can be characterized by one-dimensional state variables. Such settings correspond to fundamentals evolving as Gaussian diffusions or fundamentals evolving as two-state Markov-switching processes.[5] This general approach reveals that it is the additivity in the signal-manipulation technology (rather than the particular nature of fundamentals) the key assumption behind the form of the characterization results derived in this chapter.[6]

[5] Going beyond these two classes the analysis becomes intractable, or the set of one-dimensional sufficient statistics cease to have economically meaningful interpretations.

[6] Provided ex-ante payoffs can be written as some function of the private and public posterior mean in each environment, the analysis that follows can be extended to situations in which the long-run agent's flow payoff $u(\cdot)$ also depends on the current state of θ.

Mean-reverting Fundamentals

Suppose first that fundamentals $\theta := (\theta_t)_{t\geq 0}$ evolve as an Ornstein-Uhlenbeck process of the form

$$d\theta_t = -\kappa(\theta_t - \eta)dt + \sigma_\theta dZ_t^\theta,\ t \geq 0. \tag{2.7}$$

The long-run agent creates estimates about θ using the information conveyed by Y (eqn. (2.2)). The following Lemma is a standard result in filtering theory:[7]

Lemma 2.1. Private beliefs' law of motion– Gaussian case. *Consider the system defined by the private signal Y in (2.2) and the fundamentals (2.7). Suppose that the agent's initial prior $\theta_0|\mathcal{F}_0$ is normally distributed $\mathcal{N}(p^o, \gamma^o)$. Then, $\rho_t \overset{cdf}{=} \mathcal{N}(p_t, \gamma_t)$, where the posterior mean process $p_t := \mathbb{E}[\theta_t|\mathcal{F}_t^Y]$ and posterior variance $\gamma_t := \mathbb{E}[(\theta_t - p_t)^2|\mathcal{F}_t^Y]$, $t \geq 0$, satisfy the stochastic differential equation (SDE) and ordinary differential equation (ODE)*

$$dp_t = -\kappa(p_t - \eta)dt + \frac{\gamma_t}{\sigma_\xi}\frac{dY_t - p_t dt}{\sigma_\xi},\ \text{and} \tag{2.8}$$

$$\dot{\gamma}_t = -2\kappa\gamma_t + \sigma_\theta^2 - \left(\frac{\gamma_t}{\sigma_\xi}\right)^2,\ t > 0, \tag{2.9}$$

Finally, the process $Z_t^Y := \frac{1}{\sigma_\xi}\left(Y_s - \int_0^t p_s ds\right)$, $t \geq 0$, is an exogenous $\mathbb{F}^Y$-Brownian motion from the agent's perspective, called the innovation process.

Proof: See Liptser and Shiryaev (1977).

□

Three interesting features of (2.8) and (2.9) are worth noting. First, the evolution of the posterior mean preserves the stochastic structure of the evolution of true

[7] Liptser and Shiryaev (1977) is the main reference for all the filtering results used in this chapter.

fundamentals: the posterior mean p mean reverts toward η at the same rate $\kappa \geq 0$. Second, the posterior mean's response to unexpected signal observations (captured by the innovation process Z^Y) increases with the size of the mean-square error and decreases with the signal's volatility (σ_ξ). This implies that beliefs react more strongly in settings where either less information has been accumulated, or in settings where signals are more accurate. Finally, the posterior variance evolves in a deterministic way, so its entire trajectory is perfectly anticipated at time zero.

The market, on the other hand, does not observe the long-run agent's past actions. Hence, it can only use the information conveyed by ξ to construct estimates about the current value of fundamentals. More specifically, suppose that the market *conjectures* that the long-run agent has been following an $\mathbb{F}^\xi$−progressively measurable process $a^* := (a_t^*)_{t\geq 0}$ as a manipulation strategy. From the market's perspective, the public signal can be written as

$$d\xi_t = (a_t^* + \theta_t)dt + \sigma_\xi dZ_t^*,\ t \geq 0, \tag{2.10}$$

where $Z^* := (Z_t^*)_{t\geq 0}$ is a Brownian motion from its standpoint (potentially different from Z). In equilibrium, this conjecture is correct.

The market's belief about the fundamental account for the potential bias due to signal manipulation on the long-run agent's side. A straightforward application of Lemma 2.1 yields that the public beliefs' posterior mean process $p_t^* := \mathbb{E}^*[\theta_t|\mathcal{F}_t^\xi]$, $t \geq 0$, evolves according to

$$dp_t^* = -\kappa(p_t^* - \eta)dt + \frac{\gamma_t}{\sigma_\xi}\frac{d\xi - (a_t^* + p_t^*)dt}{\sigma_\xi},\ t > 0,$$

where $\frac{1}{\sigma_\xi}\left(\xi_t - \int_0^t (a_s^* + p_s^*)ds\right),\ t \geq 0$, is an innovation process ($\mathbb{F}^\xi$−Brownian motion) from the market's perspective. Intuitively, since at any point in time a fraction of the observed signal is attributable to signal manipulation, only the unexplained increment $d\xi_t - a_t^* dt$ conveys relevant information about fundamentals, $t \geq 0$.[8]

Observe that the agent affects the market's beliefs by controlling the true distribution of the public signal. From his standpoint, the latter evolves according to

$$d\xi_t = (a_t + p_t)dt + \sigma_\xi dZ_t^Y, \ t \geq 0,$$

where Z^Y is the innovation process defined in Lemma (2.1). As a result, public beliefs from the agent's perspective take the form

$$dp_t^* = [-(\kappa + \beta_t)(p_t^* - \eta) + \beta_t(a_t - a_t^*) + \beta_t(p_t - \eta)]dt + \beta_t \sigma_\xi dZ_t^Y, \ t \geq 0, \quad (2.11)$$

where $\beta_t := \gamma_t / \sigma_\xi^2$, $t \geq 0$.

A minimal requirement on both the long-run agent's feasible strategies and on the market's manipulation conjecture is that they must induce a well defined solution to (2.11). I use the following solution concept:

Definition 2.2. Feasible pair. *A strategy $a_t(\xi_s : s \in [0,t], p_t)$, $t \geq 0$, of the long-run agent and a conjecture $a_t^*(\xi_s : s \in [0,t])$, $t \geq 0$, of the market are said to be a feasible pair if the market's belief process (2.11) admits a unique strong solution satisfying*

$$\mathbb{E}\left[\int_0^t |p_s^*|^2 ds\right] < \infty \quad (2.12)$$

[8]Theorem 12.3 in Liptser and Shiryaev (1977) ensures that, up to integrability conditions, given any family measurable of functionals $\{a_t^* : t \geq 0\}$ with $a_t^* : C([0,t)) \to A$, $t \geq 0$, there exists a unique $\mathbb{F}^\xi-$measurable solution $(p_t^*)_{t \geq 0}$, to the SDE in $x := (x_t)_{t \geq 0}$, $dx_t = -\kappa(x_t - \eta)dt + \frac{\gamma_t}{\sigma_\xi} \frac{d\xi - (a_t^*(\xi) + x_t^*)dt}{\sigma_\xi}$, $t > 0$. Moreover, this solution corresponds to $\mathbb{E}^*[\theta_t | \mathcal{F}_t^\xi]$, $t \geq 0$, when ξ evolves as in (2.10).

for all $t \geq 0$.[9][10]

Since effort decisions are subject to moral hazard, the agent can deviate from a^* and, through controlling ξ, distort the market's belief about his skills. Deviations from conjectured strategies create a *discrepancy* between private and public beliefs. Hence, it is convenient to introduce a state variable that measures the wedge between the agent's and the market's perception about fundamentals:

Proposition 2.1. *Suppose that the market conjectures a manipulation strategy* $a^* := (a_t^*)_{t\geq 0}$*, that the long-run agent actually follows* $a := (a_t)_{t\geq 0}$ *and that* (a, a^*) *is a feasible pair. Then, from the long-run agent's perspective, public beliefs can be written as* $p_t^* = p_t + \Delta_t$*, where the process* $\Delta := (\Delta_t)_{t\geq 0}$ *is governed by the ODE*

$$d\Delta_t = [-(\kappa + \beta_t)\Delta_t + \beta_t(a_t - a_t^*)]dt, t > 0, \tag{2.13}$$

with $\beta_t := \gamma_t/\sigma_\xi^2$*,* $t \geq 0$ *and* $\Delta_0 = \Delta^o$*.*

Proof: See the Appendix.

□

The analysis of the dynamics of belief asymmetry (2.13) is crucial for understanding the long-run agent's incentives to deviate off the equilibrium path. Observe first that on the equilibrium path, if the long-run agent manipulates the signal beyond the market's expectations ($a_t > a_t^*$) he then becomes more pessimistic than the market

[9]Solution concepts to SDEs typically include as part of the definition an integrability condition that determines the class of stochastic processes where the solution is to be found. For the particular definition used here, the conditions on a and a^* that ensure existence are: (i) the random processes $(a_t(\cdot, 0))_{t\geq 0}$ and $(a_t^*(\cdot))_{t\geq 0}$ satisfy the integrability condition (2.12) and, (ii) $x \mapsto a_t(\xi, x)$ is Lipschitz, uniformly in ξ and $t \geq 0$.

[10]Equation (2.11) is an SDE with random coefficients, as strategies depend on the paths of ξ. All the results derived in this chapter could be obtained under the weaker requirement that $\int_0^t |p_s^*|^2 ds < \infty$, a.s. for all $t \geq 0$. Such extension adds no economic value to the analysis.

about the current value of fundamentals ($d\Delta_t > 0$). This is because, conditional on both parties observing the same public signal, the long-run agent attributes a higher fraction of the signal to signal-manipulation than the market does.

Second, belief discrepancies have an inherent tendency to disappear as time evolves. More specifically, any stock of belief asymmetry resulting from a one-shot deviation from a^* at time t depreciates at a rate equal to $\kappa + \beta_t$, $t \geq 0$. On the one hand, high rates of mean reversion κ imply that the shocks to fundamentals have low persistence. Consequently, public beliefs give low weight to past signal observations for large values of κ. On the other hand, as information accumulates over time unanticipated changes in the evolution of the public signal are more likely to be the consequence of changes in fundamentals rather than the consequence of noise. Therefore, provided $\gamma_0 = \gamma^o$ large enough, the weight attached to past signals increases over time ($(\beta_t)_{t\geq 0}$ is decreasing for γ^o large).

Finally, observe that the market's conjecture acts as a threshold that the long-run agent has to exceed in order to induce more belief asymmetry. This threshold is endogenous, as a^* must correspond to the long-run agent's strategy on the equilibrium path. Moreover, since a^* is a function of the partial observations of ξ, it is also affected by the long-run agent's manipulation decision. Hence, in evaluating deviations off the equilibrium path, the long-run agent also takes into account the effect that his actions have on the market's future standard of belief manipulation.

In what follows I assume that learning is stationary, that is, $\gamma^o = \gamma^*$, where $\gamma^* = \sigma_\xi^2(\sqrt{\kappa^2 + \sigma_\theta^2/\sigma_\xi^2} - \kappa)$ is the unique stationary solution to (2.9) that is strictly positive. As a result, $\beta_t = \beta := \gamma^*/\sigma_\xi^2$ for all $t \geq 0$.[11]

[11] All the characterization results in this chapter can be easily extended to the case in which learning is away from steady state.

Markov-Switching Fundamentals

If fundamentals $\theta := (\theta_t)_{t\geq 0}$ instead follow a two-state Markov chain, the previous analysis can be easily replicated. Suppose that θ takes values on $\Theta = \{h, \ell\}$, $\ell < h$, and that the associated transition matrix is given by

$$\Lambda := \begin{bmatrix} -\lambda_1 & \lambda_1 \\ \lambda_0 & -\lambda_0 \end{bmatrix}, \tag{2.14}$$

where λ_i corresponds to the transition rate from state i to state j, $i, j \in \{\ell, h\}$.

Denote by π_t the probability that θ_t is in state h at time $t \geq 0$ given the observations $(Y_s : s \in [0,t])$, i.e. $\pi_t := \mathbb{P}(\theta_t = h|\mathcal{F}_t^Y)$, $t \geq 0$. Then, the log-likelihood ratio given the partial observations of Y,

$$p_t := \log\left(\frac{\pi_t}{1-\pi_t}\right),$$

is a sufficient statistic for the posterior distribution of θ_t given $\mathcal{F}_t^Y$, $t \geq 0$. The following result is a Markov-chains counterpart to Lemma 2.1 and Proposition 2.1:

Proposition 2.2. *Suppose that θ is a binary-state markov chain taking values in $\Theta = \{\ell, h\}$, $\ell < h$, and with transition matrix (2.14). Then, the log-likelihood ratio process given the observations of Y, $p := (p_t)_{t\geq 0}$, satisfies the SDE*

$$dp_t = \underbrace{\left[\lambda_1 \frac{e^{p_t}+1}{e^{p_t}} - \lambda_0(1+e^{p_t}) - \frac{\delta^2}{2\sigma_\xi^2}\left(1 - \frac{2e^{p_t}}{1+e^{p_t}}\right)\right]}_{\mu(p_t):=} dt + \frac{\delta}{\sigma_\xi^2}\left(dY_t - \frac{e^{p_t}}{1+e^{p_t}}dt\right) \tag{2.15}$$

where $\frac{1}{\sigma_\xi}\left(Y_t - \int_0^t \frac{e^{p_s}}{1+e^{p_s}}ds\right)$ is an $\mathbb{F}^\xi-$Brownian motion from the agent's standpoint. Moreover, from the agent's perspective, the difference between the public and private *log-likelihood ratio processes*

$$\Delta_t := p_t^* - p_t,\ t \geq 0,$$

evolves according to

$$d\Delta_t = [\phi(p_t, \Delta_t) + \beta(a_t - a_t^*)]dt,\ t > 0,\ \Delta_0 = \Delta^o, \tag{2.16}$$

where $\phi : \mathbb{R}^2 \to \mathbb{R}$ *is defined by*

$$\phi(p, \Delta) = \lambda_1 \left[\frac{e^{p+\Delta} + 1}{e^{p+\Delta}} - \frac{e^p + 1}{e^p} \right] - \lambda_0 [e^{p+\Delta} - e^p]. \tag{2.17}$$

Proof: See the Appendix.

□

Observe that the combination of a finite-state markov chain and Brownian signals generate learning dynamics that are nonlinear. This is in contrast with the linearity of the Kalman-Bucy filter p with respect to Y in the Gaussian case. In particular, this nonlinearity yields a belief-asymmetry process decaying potentially at stochastic rates (ϕ depends on p) even when the long-run agent decides to follow the market's conjectured strategy after having deviated from it in the past. Finally, it is easy to see that whenever private and public beliefs coincide and the agent follows a^*, beliefs remained aligned ($\phi(\cdot, 0) \equiv 0$).

2.2.3 Markov Perfect Equilibrium

Consider an economic environment in which fundamentals evolve as a mean-reverting process (under stationary learning) or as a Markov-switching one. From the previous section, there exist state variables $p := (p_t)_{t \geq 0}$ and $\Delta := (\Delta_t)_{t \geq 0}$ that take values over the entire real line and that fully characterize both the long-run agent's and the market's posterior distribution about fundamentals, $\rho := (\rho_t)_{t\geq 0}$ and $\rho^* := (\rho_t^*)_{t\geq 0}$, respectively (the latter via the relation $p_t^* = p_t + \Delta_t$, $t \geq 0$). Their respective laws of motion take the form

$$dp_t = \mu(p_t)dt + \sigma_p dZ_t^Y,\ p_0 = p^o,\ \text{and} \tag{2.18}$$

$$d\Delta_t = [\phi(p_t, \Delta_t) + \beta(a_t - a_t^*)]dt,\ t > 0,\ \Delta_0 = \Delta^o, \tag{2.19}$$

where μ, σ_p, ϕ and β are given in Lemma 2.1 and Propositions 2.1 and 2.2. I refer to p as the *private beliefs process* and to Δ as the *belief-asymmetry process*, respectively. In the sequel, I assume that the market's best-response action is of the form $\mathtt{b}(p^*, a^*)$ and observe that off the equilibrium path the market's actions can be written as $\mathtt{b}(p_t + \Delta_t, a_t^*)$, $t \geq 0$, from the long-run agent's perspective.

Because I restrict to pure strategies, in any equilibrium the market must perfectly anticipate the long-run agent's action when starting from a common prior. This implies that $\Delta \equiv 0$ on the equilibrium path, leaving $p = p^*$ as the unique payoff-relevant state variable. It is natural to think therefore that, at any point in time, on-path incentives will be a function of the public beliefs process $p^* := (p_t^*)_{t\geq 0}$. This motivates the study of *Markov Perfect Equilibria* in beliefs.

On the equilibrium path, the impact of a first unit of belief asymmetry Δ on the long-run agent's flow payoffs is likely to vary across different levels of public beliefs. This is because his preferences (u), or the market's actions ($\mathtt{b}$) or the learning process (ϕ) could be nonlinear functions. Furthermore, after a deviation takes place, the long-run agent acquires private information about the future evolution of fundamentals and the market's belief about them is biased. The long-run agent can therefore find it profitable to condition his actions on both his private beliefs p and the belief-asymmetry process Δ.

Definition 2.3. *An equilibrium* (a, b, p, p^*) *consisting of a manipulation strategy of the long-run agent* $a := (a_t)_{t\geq 0}$*, a public action profile of the market* $b := (b_t)_{t\geq 0}$*, a private belief process of the long-run agent* $p := (p_t)_{t\geq 0}$*, and a public belief process of the market* $p^* := (p^*_t)_{t\geq 0}$ *is a Markov Perfect Equilibrium (MPE) in beliefs if and only if there exists a measurable function* $\mathtt{a} : \mathbb{R}^2 \to A$ *such that*

(i) *Given any feasible strategy* $\breve{a} := (\breve{a}_s)_{s\geq 0}$ *and any private history* $(\xi_s : s \in [0, t], \breve{a}_s : s \in [0, t])$ *that lead to* $(p_t, p^*_t - p_t) = (\overline{p}, \overline{\Delta}) \in \mathbb{R}^2$*, the agent's action at time* t *is of the form* $a_t = \mathtt{a}(\overline{p}, \overline{\Delta})$*,* $t \geq 0$*;*

(ii) *After all public histories* $(\xi_s : s \in [0, t])$ *that lead to* $p^*_t = \overline{p} \in \mathbb{R}$*,* $b_t = \mathtt{b}(\overline{p}, \mathtt{a}(\overline{p}, 0))$*,* $t \geq 0$*,*

and (a, b, p, p^*) *satisfies (a)-(d) in Definition 2.1.*

The interpretation of the previous Definition should be straightforward. It is worth to emphasize again that since any signal path $(\xi_s : s \in [0, t])$ is consistent with equilibrium actions, the market constructs beliefs assuming that the long-run agent is following $\mathtt{a}(p^*_t, 0)$, $t \geq 0$, that is, as if he had never deviated off the equilibrium path. On the other hand, the strategy $\mathtt{a}(p_t, \Delta_t)$ must be optimal for the long-run agent after all private histories leading to $(p, \Delta) \in \mathbb{R}^2$ when the market construct beliefs using a conjecture of the form $\mathtt{a}(\cdot, 0)$.

2.2.4 Necessary Conditions for Markov Perfect Equilibria

The Agent's Value Function

Suppose that the market conjectures that the long-run agent follows a strategy of the form $(a^*(p^*_t))_{t\geq 0}$ for some measurable function $a^* : \mathbb{R} \to A$. Given this *markovian conjecture* the long-run agent's problem becomes fully recursive in (p, Δ). It consists of choosing a manipulation strategy $a := (a_t(\xi_s : s \in [0,t], p_t))_{t\geq 0}$, such that (i) (a, a^*) is a feasible pair (Definition 2.2) and (ii) at any time t the continuation strategy $(a_s)_{s\geq t}$ maximizes the long-run agent's expected discounted utility

$$\mathbb{E}_t\left[\int_t^\infty e^{-r(s-t)}(u(\mathtt{b}(p_s+\Delta_s, a^*(p_s+\Delta_s))) - g(a_s))ds\right]$$

subject to the dynamics

$$\begin{aligned} d\Delta_s &= [\phi(p_s,\Delta_s) + \beta(a_s - a^*(p_s+\Delta_s))]dt,\ s > t,\ \Delta_t = \overline{\Delta} \\ dp_s &= \mu(p_s)dt + \sigma_p dZ_s^Y,\ s > t,\ p_t = \overline{p}, \end{aligned}$$

after all private histories leading to $(p_t, \Delta_t) = (\overline{p}, \overline{\Delta})$, $t \geq 0$, and for all $(\overline{p}, \overline{\Delta}) \in \mathbb{R}^2$. For notational simplicity I redefine $u(\mathtt{b}(\cdot,\cdot))$ to be $u(\cdot,\cdot)$ unless otherwise stated.

Let $V^{a^*}(\overline{p}, \overline{\Delta})$ denote the long-run agent's *value function* associated to the previous problem, and observe that it explicitly depends on the market's conjecture a^*. As it is standard in the dynamic programming literature, if an optimal control exists and the agent's value function is smooth enough (of class $C^{2,1}(\mathbb{R}^2)$), then V^{a^*} satisfies the Hamilton-Jacobi-Bellman (HJB) equation[12]

$$\begin{aligned} rV^{a^*}(p,\Delta) &= \sup_{a\in A}\Big\{u(p+\Delta, a^*(p+\Delta)) - g(a) + \mu(p)V_p^{a^*}(p,\Delta) + \frac{1}{2}\sigma_p^2 V_{pp}^{a^*}(p,\Delta) \\ &\quad +[\phi(p,\Delta) + \beta(a - a^*(p+\Delta))]V_\Delta^{a^*}(p,\Delta)\Big\},\ (p,\Delta)\in\mathbb{R}^2. \end{aligned} \tag{2.20}$$

Consequently, if an optimal markovian strategy $\hat{a} : \mathbb{R}^2 \to \mathbb{R}$ in response to a^* exists, then it has to satisfy that

$$\hat{a}(p,\Delta) = \arg\max_{a\in A}\{a\beta V_\Delta^{a^*}(p,\Delta) - g(a)\},\ (p,\Delta)\in\mathbb{R}^2. \tag{2.21}$$

[12]Subscripts p and Δ denote partial derivatives with respect to p and Δ, respectively.

This in turn yields the following *equilibrium condition* on the market's conjecture:

$$a^*(p^*) = \arg\max_{a \in A} \{a\beta V_\Delta^{a^*}(p^*, 0) - g(a)\}, \; p^* \in \mathbb{R}^2. \tag{2.22}$$

In other words, when the market constructs its beliefs as if the long-run agent is following the strategy $(a^*(p_t^*))_{t\geq 0}$, the optimal policy delivered by the HJB equation must coincide with the market's conjecture whenever beliefs are aligned.

The following results summarizes the discussion about on and off the equilibrium path incentives:

Theorem 2.1. Global incentives. *Assume that a MPE* $\mathtt{a} : \mathbb{R}^2 \to \mathbb{R}$ *exists and the associated value function* $V^{\mathtt{a}}$ *is of class* $C^{2,1}(\mathbb{R}^2)$*.*[13] *Then, the agent's value function* $V^{\mathtt{a}}(p, \Delta)$ *satisfies the partial differential equation (PDE)*

$$\begin{aligned} rV(p,\Delta) \;\; &= \;\; \sup_{a \in A} \Bigg\{ u(p+\Delta, \mathtt{a}(p+\Delta,0)) - g(a) + \mu(p)V_p(p,\Delta) + \frac{1}{2}\sigma_p^2 V_{pp}(p,\Delta) \\ &\qquad + [\phi(p,\Delta) + \beta(a - \mathtt{a}(p+\Delta,0))]V_\Delta(p,\Delta) \Bigg\}, \;\; (p,\Delta) \in \mathbb{R}^2 \end{aligned} \tag{2.23}$$

$$s.t. \qquad \mathtt{a}(p,0) \in \arg\max_{a \in A} \{\beta V_\Delta(p,0)a - g(a)\}. \tag{2.24}$$

□

When the long-run agent has deviated in the past private and public beliefs do not coincide. This results in value functions characterized by partial differential equations (PDEs). Nevertheless, the type of PDEs that arise in settings involving moral hazard and belief manipulation are non-standard ones. Specifically, the PDE (2.23)-(2.24) is *nonlocal*, as the local behavior of the long-run agent's continuation value around a

[13]More formally, since the market construct beliefs using $\mathtt{a}(\cdot, 0)$ the correct notation should be $V^{\mathtt{a}(\cdot,0)}$. For notational simplicity, I proceed with the abbreviation $V^{\mathtt{a}}$.

point (p, Δ) depends on the value attached to a marginal deviation off the equilibrium at the point $(p + \Delta, 0)$ (eqn. (2.24) in the Theorem). In other words, since after all public histories the market constructs beliefs as if the long-run agent had never deviated, *off-path utility depend on on-path marginal utility*. This is clearly seen at points where on-path incentives are interior:

$$\mathsf{a}(p + \Delta, 0) = (g')^{-1}(\beta V_\Delta(p + \Delta, 0)).$$

This type of non-localness, and consequently, this type of PDE, seems to be new.

Remark 2.1. The importance of Theorem 2.1 lies not only in the insights regarding the determinants of off-path incentives, but also on its applicability for showing the existence of markovian equilibria for the class of games presented. In fact, if V solves the PDE (2.23)-(2.24), then it satisfies the HJB equation (2.20) along with the function

$$a^*(\cdot) := \arg\max_{a \in A}\{a\beta V_\Delta(\cdot, 0) - g(a)\}$$

taken as given in the dynamics of Δ. As a result, verification theorems in dynamic programming apply: provided V satisfies some additional conditions, the policy delivered by the HJB equation is an optimal control. Furthermore, since V satisfies the equilibrium condition (2.24), such a policy is, by construction, a Markov perfect equilibrium. This insight is used in Section 4 to show the existence of markovian equilibria linear in (p, Δ) for a class of games for which the associated PDEs admit analytic solutions.

An ODE for $V_\Delta(p, 0)$

The PDE (2.23)-(2.24) is, in almost every setting, very hard to visualize. However, it is possible to extract properties of equilibrium behavior without the need of fully solving such equation. More specifically, observe that as long as the long-run agent's actions are interior and beliefs are aligned, the long-run agent's incentives are, in any Markov perfect equilibrium, characterized by $V_\Delta(p, 0)$. The next result shows that the value attached to inducing a small discrepancy between private and public beliefs, in fact solves an ODE:

Proposition 2.3. (**On path behavior**) *Assume that a Markov perfect equilibrium* $\mathsf{a} : \mathbb{R}^2 \to \mathbb{R}$ *exists and the associated value function* V^{a} *is of class* $C^{2,1}(\mathbb{R}^2)$. *Furthermore, suppose that at a level of public beliefs* p

(i) *on path incentives are interior and*

(ii) *on path incentives are locally twice continuously differentiable with respect to public beliefs, i.e., there exists a neighborhood* $\mathcal{O}$ *of* p *such that* $\mathsf{a}(\cdot, 0) \in C^2(\mathcal{O})$.

Then, $g'(\mathsf{a}(\cdot, 0)) = \beta V^{\mathsf{a}}_\Delta(\cdot, 0)$, *where* $V^{\mathsf{a}}_\Delta(\cdot, 0)$ *satisfies the ODE in* $p \mapsto V_\Delta(p, 0)$

$$\begin{aligned}\tilde{r}(p, V_\Delta(p,0))V_\Delta(p,0) \quad &= \quad u_p(p, g'^{-1}(\beta V_\Delta(p,0))) + u_a(p, g'^{-1}(\beta V_\Delta(p,0)))\frac{d}{dp}g'^{-1}(\beta V_\Delta(p,0)) \\ &\quad + V_{\Delta p}(p,0)\mu(p) + \frac{1}{2}\sigma_p^2 V_{\Delta pp}(p,0), \; p \in \mathcal{O}, \end{aligned} \tag{2.25}$$

with $\tilde{r}(p, V_\Delta(p,0)) := r - \phi_\Delta(p,0) + \beta \frac{d}{dp} g'^{-1}(\beta V_\Delta(p,0))$.

Proof: Differentiate the PDE (2.23)-(2.24) with respect to Δ and evaluate at $\Delta = 0$.

□

Equation (2.25) corresponds to a necessary condition that on-path incentives must, in any (smooth) Markov perfect equilibrium, satisfy. It takes the form of an ODE for $p \mapsto V_\Delta(p, 0)$, the value attached to inducing a small degree of belief asymmetry across different levels of public beliefs. As a result, it measures the strength of the long-run agent's *local* incentives to deviate off the equilibrium path. I refer to (2.25) as the *incentives equation*.

The incentives equation corresponds to an Euler equation under an equilibrium condition. The Euler feature of this equation states that the long-run agent must be indifferent between exerting signal manipulation "today" (the right-hand side in (2.25)) and delaying signal manipulation to "tomorrow". Exerting signal manipulation today is beneficial for two reasons. First an additional unit of asymmetry between private and public beliefs increases the long-run agent's flow payoffs by

$$\begin{aligned} & \frac{d}{d\Delta} u(p+\Delta, \mathtt{a}(p+\Delta, 0))\Big|_{\Delta=0} \\ = \; & u_p(p, g'^{-1}(\beta V_\Delta(p,0))) + u_a(p, g'^{-1}(\beta V_\Delta(p,0)))\frac{d}{dp} g'^{-1}(\beta V_\Delta(p,0)) \end{aligned}$$

so long as no deviations have taken place. The first term in the previous expression corresponds to the *direct* effect that an additional unit of belief asymmetry has on the market's actions, and consequently, on the long-run agent's flow utility. The second term corresponds to the effect that an additional unit of belief asymmetry has on the market's action through the *indirect* channel of affecting the market's conjecture about the level of signal manipulation. I elaborate more on this point shortly.

The second benefit from exerting signal manipulation today relates to cost smoothing. More specifically, observe that, on the equilibrium path, the expected rate of change of the value associated with an additional unit of belief asymmetry takes the form

$$\begin{aligned} \lim_{h\to 0} \frac{\mathbb{E}_t[V_\Delta(p_{t+h}, \Delta_{t+h})] - V_\Delta(p_t,0)}{h} &= \underbrace{[\phi(p_t,0) + \beta(\mathtt{a}(p_t,0) - \mathtt{a}(p_t,0))]}_{\equiv 0} V_{\Delta\Delta}(p_t,0) \\ &\quad + \mu(p) V_{p\Delta}(p_t,0) + \frac{1}{2}\sigma_p^2 V_{pp\Delta}(p_t,0) \\ &= \mu(p) V_{p\Delta}(p_t,0) + \frac{1}{2}\sigma_p^2 V_{pp\Delta}(p_t,0). \end{aligned}$$

Consequently, if the benefits from signal manipulation are expected to change at high rates in the near future, then because manipulation costs are convex, it is optimal to start investing in belief distortion today. The value attached to belief asymmetry can change because beliefs are locally predictable ($\mu(p) \neq 0$) or because learning about fundamentals is valued differently as public beliefs change (third-order term $V_{\Delta pp}$). In any case, future economic conditions will influence the agent's current decisions if and only if private beliefs p interact in a non-trivial way with the belief-asymmetry process Δ. This is likely to occur in settings where the impact of Δ on the long-run agent's flow utility is nonlinear.

When the agent instead decides to postpone signal manipulation to tomorrow, he saves the costs associated with his investments in belief asymmetry depreciating over time. These depreciation costs are captured in the *required rate of return* on belief asymmetry

$$\tilde{r}(p, V_\Delta(p,0)) = r - \phi_\Delta(p,0) + \beta \frac{d}{dp} g'^{-1}\left(\beta V_\Delta(p,0)\right),\ p \in \mathbb{R},$$

in the left-hand side of (2.25). For any fixed discount rate r, the lower $\tilde{r}(p, V_\Delta(p,0))$ is, the more persistent belief distortions are, and thus the higher the incentives to manipulate public beliefs today.

The required rate of return $\tilde{r}$ is determined endogenously in equilibrium. Recall first that, given any conjecture a^*, the belief-asymmetry process (2.19) evolves according to

$$d\Delta_t = [\phi(p_t, \Delta_t) + \beta(a_t - a_t^*)]dt,\ t \geq 0.$$

Thus, if the market's conjecture a^* were an exogenous function of time, a local deviation at $p_t = p$ would entail a dividend flow that depreciates at a rate equal to $-\phi_\Delta(p,0)$. This in turn would map into a monetary loss discounted at a rate equal to $r - \phi_\Delta(p,0)$.

However, when equilibrium strategies depend on current beliefs, deviations also affect the market's conjecture about equilibrium play. In fact, by inducing a marginal unit of belief asymmetry, the long-run agent locally distorts the market's conjecture by $\frac{d}{d\Delta}a^*(p+\Delta)$ off the equilibrium path. As a result, when a Markov perfect equilibrium $\mathsf{a}(p,\Delta)$ exists and incentives are interior, the equilibrium condition (2.24) yields that

$$\frac{\partial}{\partial\Delta}\mathsf{a}(p+\Delta,0)\Big|_{\Delta=0} := \frac{\partial}{\partial\Delta}g'^{-1}(\beta V_\Delta(p+\Delta,0))\Big|_{\Delta=0} = \frac{\partial}{\partial p}g'^{-1}(\beta V_\Delta(p,0)),$$

on the equilibrium path. The returns from belief asymmetry are thus captured by $\tilde{r}(p, V_\Delta(p,0))$ and not by $r - \phi_\Delta(p,0)$ alone.[14]

Affecting the market's conjecture can strengthen or weaken the long-run agent's incentives to induce belief asymmetry. For instance, if the market expects belief manipulation to be locally increasing in public beliefs, the returns from inducing belief asymmetry could be low. In such a case, the long-run agent would have to fulfill tougher standards of belief manipulation in the future in order to maintain perceived fundamentals at a fixed level, which in turn reduces the benefits from inducing belief asymmetry.[15] In either case, the equilibrium feature of the incentives equation states that, when a Markov perfect equilibrium exists, the additional benefits from affecting the market's conjecture (summarized in the endogenous rate of return $\tilde{r}$) generate belief manipulation incentives that coincides with what the market conjectures will be played on the equilibrium path. As a result, (2.25) does not incorporate any gains associated with having private information about fundamentals.

[14]Since beliefs react to new information according to the sensitivity parameter β, the change on the market's conjecture $\frac{\partial}{\partial\Delta}\mathsf{a}(p+\Delta,0)$ is amplified by β in the equilibrium rate of return $\tilde{r}(p)$.

[15]A similar force appears in Kyle's (1985) model of insider trading. In his model, a trader who has private information about the value of an asset takes into account that trading more aggressively today moves the equilibrium price against him in subsequent trading rounds.

The next Theorem summarizes the previous discussion about incentives on the equilibrium path:

Theorem 2.2. *Assume that a Markov perfect equilibrium exists. Furthermore, suppose that the associated value function* $V(p,\Delta)$ *is of class* $C^{3,1}(\mathbb{R}^2)$*. Then, whenever (on-path) incentives are interior, the agent's (on-path) value function* $V(\cdot,0)$ *satisfies the ODE*

$$rU(p) = u(p,q(p)) - g(g'^{-1}(\beta q(p))) + U'(p)\mu(p) + \frac{1}{2}\sigma_p^2 U''(p),\ p \in \mathbb{R} \qquad (2.26)$$

whereas the agent's (on-path) marginal utility function $V_\Delta(\cdot,0)$ *satisfies the ODE*

$$\begin{aligned}\tilde{r}(p,q(p))q(p) &= u_p(p,g'^{-1}(\beta q(p))) + u_a(p,g'^{-1}(\beta q(p)))\frac{d}{dp}g'^{-1}(\beta q(p)) \\ &\quad + q'(p)\mu(p) + \frac{1}{2}\sigma_p^2 q''(p),\ p \in \mathbb{R}, \qquad (2.27)\end{aligned}$$

where $\tilde{r}(p,q(p)) = r - \phi_\Delta(p,0) + \beta\frac{d}{dp}g'^{-1}(\beta q(p))$*.*

Remark 2.2. The choice of formulating the agent's problem in the (p,Δ)-coordinate system is purely for expositional purposes. It allows us to understand the agent's problem as one of investing in belief asymmetry in the presence of convex costs, with private beliefs playing the role of an exogenous price. Instead, in the (p,p^*)-space on-path incentives are determined by the value attached to a marginal increment in public beliefs along the diagonal $\{(p,p) \mid p \in \mathbb{R}\}$. It is easy to see that the same envelope argument used to derive (2.25) on the corresponding HJB equation yields an ordinary differential equation for $p \mapsto V_{p^*}(p,p)$.

Remark 2.3. Following the analogy between this problem and the literature of investment in the presence of adjustments costs, the resemblance between the incentives equation (2.25) and the ones for the traditional "q" is clear.[16] Nevertheless,

[16]See, for example, Dixit and Pindyck (1994).

what makes the incentives equation distinctive is that the required rate of return on belief asymmetry is determined endogenously and also depends on q. In fact, from the second ODE in the previous theorem, $V_\Delta(p,0)$ satisfies a highly nonlinear version of the traditional equations for "q".

2.2.5 The Incentives Equation: Examples

Linear Environments: Holmstrom's Career Concerns Model

In Holmstrom's model of reputation a risk-neutral worker can produce an output (ξ) using his skills (θ) and effort (a). In continuous-time, the worker's skills are modeled as a martingale $d\theta_t = \sigma_\theta dZ_t^\theta$, $t \geq 0$, whereas output ξ is given by (2.1). If the pool of potential employers is competitive, the worker's wage at time t corresponds to the market's expected output flow at that instant. Hence, the market's action takes the form $\mathtt{b}(a_t^*, p_t^*) = a_t^* + p_t^*$, where a_t^* is the agent's equilibrium effort decision at time t and $p_t^* = \mathbb{E}^*[\theta_t | \mathcal{F}_t^\xi]$, $t \geq 0$. It can be easily verified that the constant function $V_\Delta(p,0) \equiv \frac{1}{r+\beta}$ (with $\beta = \sigma_\theta/\sigma_\xi$) is a solution to the incentives equation (2.25). As a result, both the additivity and linearity imposed by Holmstrom generate value functions V that are fully separable in p and Δ. Thus, current incentives are independent of future effort decisions and independent of contemporaneous equilibrium conjectures.

Nonlinearities 1: Quadratic Payoffs

As in the procurement example in Section 2, suppose that there exists a market friction that makes the long-run agent's payoffs $u(p^*)$ a quadratic loss function of public beliefs p^*. It can be easily checked that a function of the form $a^*(p^*) = \alpha_1 + \alpha_2 p^*$, $p^* \in \mathbb{R}$, solves the incentives equation for suitably chosen parameters α_1 and α_2. In the next Section I show that, if $u(\cdot)$ satisfies a mild curvature condition, such a rule is indeed an equilibrium.

Nonlinearities 2: Discrete-type Space

One of the most relevant features of the incentives equation is that it can shed lights on the shape of equilibrium behavior without the need to fully solve the complex partial differential equations that characterize the long-run agent's value function. This is particularly important in environments that exhibit high nonlinearities, as I illustrate below.

Suppose that an agent's ability θ is a time zero draw from a discrete random variable taking values in $\{0,1\}$. Effort is costly according to the function $g(a)=\frac{a^2}{2}$. Moreover, assume that the agent's wage is given by $p_t^* := \mathbb{P}(\theta_t = 1|\mathcal{F}_t^\xi)$, the posterior probability that his ability is high given the information up to time $t\geq 0$. The agent's flow utility is given by a differentiable function $u:\mathbb{R}\to\mathbb{R}$.

In this case it can be easily checked that $p^*=\frac{p\Delta}{1+p(\Delta-1)}$ with (p,Δ) evolving as

$$dp_t=\frac{p_t(1-p_t)}{\sigma_\xi}dZ_t^Y \quad \text{and} \quad d\Delta_t=\frac{\Delta_t(a_t-a_t^*)}{\sigma_\xi^2}dt,\ t\geq 0,$$

respectively. Moreover, the incentives equation takes the form

$$\left[r+\frac{V_{\Delta p}(p,0)}{\sigma_\xi^4}\right]V_\Delta(p,0)=p(1-p)\left[u'(p)+\frac{V_{\Delta p}(p,0)}{\sigma_\xi^2}+\frac{p(1-p)}{2\sigma_\xi^2}V_{\Delta pp}(p,0)\right], \tag{2.28}$$

for $p\in(0,1)$. Some solutions to this ODE are plotted in the following figure:

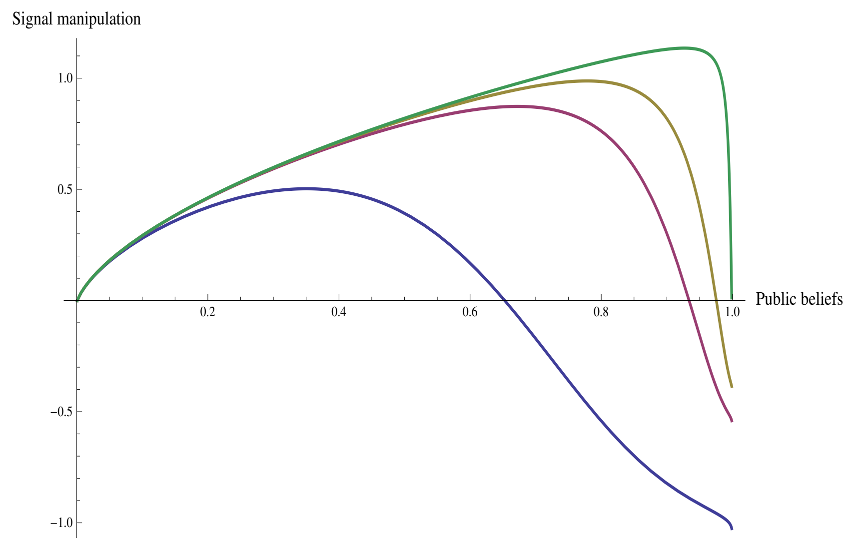

Figure 2.1: Some solutions of the ODE (2.28) when $u(p^*) = p^*$ and $V_\Delta(0, 0) = 0$. Parameter values: $\sigma_\xi = 0.2$, $r = 0.02$.

In particular, it can be shown that there exists a non-negative solution of class C^2 to the boundary value problem defined by the above ODE and the boundary conditions $V_\Delta(0, 0) = V_\Delta(1, 0) = 0$.One would expect equilibrium effort to vanish as public beliefs tend to 0 or 1. This is because public beliefs become unresponsive to new information asymptotically in those limits. Provided effort is constrained to be non-negative and a pure-strategy equilibrium vanishing at the extremes exist, then such equilibrium should look like in the figures below:

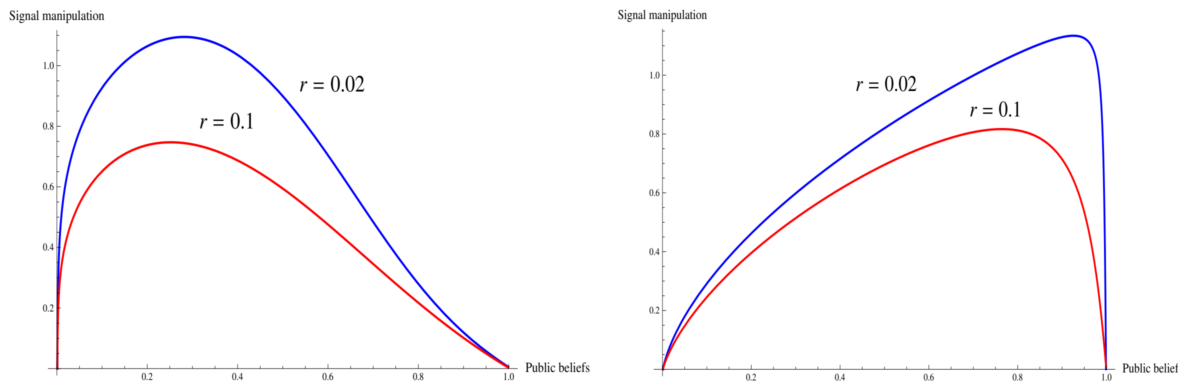

Figure 2.2: Left panel: $u(p^*) = a^* + p^*$. Right panel: $u(p^*) = p^*$. Parameter values: $\sigma_\xi = 1$.

In the right panel, the agent is payed according to the perceived value of his skills only. As a result, the myopic gain from belief distortion is given by

$$\frac{d}{d\Delta}u(p+\Delta)\bigg|_{\Delta=0} = \frac{d}{d\Delta}\left[\frac{p\Delta}{1+p(\Delta-1)}\right]\bigg|_{\Delta=0} = p(1-p),\ p \in (0,1).$$

Moreover, the learning dynamics are also symmetric around $p = 1/2$. Yet, incentives may not exhibit that property. This is because the agent's actions also affect the market's contemporaneous expectations about equilibrium play, and such effect can play against the agent's incentives.[17] For low reputations, inducing a marginal unit of belief asymmetry increases the threshold of belief manipulation that the agent will face in the near future. Moreover, such additional effort is not compensated. Reputational incentives are thus stronger once the agent has already built a reputation, as it is too costly for the agent to risk losing his high reputation and start re-building it again.

In the left panel the agent is also rewarded by effort. Now, the (on-path) myopic gain from belief distortion ceases to be symmetric around $p = 1/2$ and it is given by

$$p(1-p)[1+V_{\Delta p}(p,0)],\ p \in (0,1).$$

[17] In fact, from the incentives equation (2.28) it can be seen that the term $V_{\Delta p}(p,0)V_\Delta(p,0)$ makes the ODE non symmetric around $p = 1/2$.

Whenever equilibrium effort vanishes at the extremes, incentives will be stronger for low reputations, as in this region effort must be increasing. Even though the agent, by exerting more effort, also faces a tougher standard in the future, the market compensates those additional units. Hence, rewarding effort generates incentives for *building* a reputation. Since the costs associated with building a reputation are not as high as in the right panel, an agent with a high reputation can afford to shirk.

2.3 Signal-Jamming Games: Linear-Quadratic Environments

The necessary conditions for utility and marginal utility in Theorem 2.2 are necessary conditions for the incentives that arise in any pure-strategy Markov perfect equilibrium, provided such an equilibrium exists. However, since the incentives equation (2.25) is a local incentive constraint only, it does not ensure that the long-run agent does not benefit from large deviations off the equilibrium path. In order to verify incentive compatibility globally, it is necessary to study the solutions to the PDE (2.23)-(2.24). At this general level, this is an extremely challenging task.

The purpose of this section is to introduce a subclass of *linear-quadratic* games for which (i) Markov perfect equilibria explicitly dependent on beliefs exist, (ii) the associated value functions are smooth and (iii) on- and off-path incentives can be fully characterized. In particular, on the equilibrium path the agent's actions satisfy the incentives equation (2.25). These games have a linear-quadratic structure because learning is Gaussian (linear) and the flow utility that the agent derives from the market's actions is a quadratic loss function of public beliefs.

Quadratic preferences in the context of belief manipulation arise in many economic environments. In procurement for instance, a contractor and a government who interact repeatedly over time can learn about the contractor's efficiency (costs) to deliver goods. Moreover, if the government lacks commitment, the contractor may want to target a particular level of reputation: high enough to be awarded projects, but low enough to avoid triggering rent-extracting actions on the government's side. Consequently, imperfect competition on the demand side can in fact create negative returns from being perceived as more able. Another setting is monetary policy in the context of unobserved components of inflation: the government benefits from inflation being close to a particular target and, at the same time, from inducing inflationary surprises that map into more employment (Kydland and Prescott (1977)). This application is developed in more detail within the class of investment games with learning analyzed in the next section.

The importance of studying games within this subclass is three-fold. First, these games generate PDEs summarizing global behavior (Theorem 2.1) that admit analytic solutions. As a result, on- and off-path dynamics can be fully characterized and comparative statics with respect to key parameters can be easily performed. Second, this type of games exhibit all the forces that are expected to influence on-path behavior in more nonlinear environments (e.g., cost smoothing and endogenous rate of returns on belief asymmetry). Finally, linear-quadratic settings are the framework to study second-order approximations of more nonlinear environments around steady state.

2.3.1 Existence Result

Definition 2.4. *A signal-jamming game is said to be of linear-quadratic form if*

(i) *Fundamentals θ are a mean reverting process: $d\theta_t = -\kappa(\theta_t - \eta)dt + \sigma_\theta dZ_t^\theta$, $t \geq 0$;*

(ii) *The set of feasible action values A is the real line and the cost of signal manipulation is quadratic: $g(a) = \frac{\psi}{2}a^2$, $\psi > 0$;*

(iii) *The long-run agent's flow utility derived from the market's action is a quadratic loss function of public beliefs* $u(\mathtt{b}(p^*, a^*)) = u_0 + u_1 p^* - u_2 p_t^{*2}$, *where* $u_0, u_1 \in \mathbb{R}$ *and* $u_2 \geq 0$.[18]

The next result shows the existence of a *linear* (in beliefs) equilibrium which exhibits all the forces mentioned in the previous sections. The linear quadratic framework presented here is hence tractable enough to obtain analytic solutions, yet at the same time, able to induce rich interactions between learning and incentives.

Theorem 2.3. *Suppose that a linear-quadratic game of signal manipulation is such that*

$$u_2 \leq \frac{\psi(r+\beta+2\kappa)^2}{8\beta^2}. \tag{2.29}$$

Then, a Markov perfect equilibrium in linear strategies exists. In this equilibrium, the agent's value function is given by $V(p,\Delta) = \alpha_0 + \alpha_1 p + \alpha_2 \Delta + \alpha_3 p\Delta + \alpha_4 p^2 + \alpha_5 \Delta^2$, *where* $\alpha_0, \alpha_1 \in \mathbb{R}$, $\alpha_4, \alpha_5 < 0$,

$$\begin{aligned} \alpha_2 &= \frac{\eta\kappa\alpha_3 + u_1}{r+\beta+\kappa+\frac{\beta^2\alpha_3}{\psi}}, \text{ and} \\ \alpha_3 &= \frac{\psi}{2\beta^2}\left[-(r+\beta+2\kappa) + \sqrt{(r+\beta+2\kappa)^2 - \frac{8u_2\beta^2}{\psi}}\right] < 0. \end{aligned} \tag{2.30}$$

The optimal degree of signal manipulation corresponds to $\mathtt{a}(p,\Delta) := \frac{\beta}{\psi}V_\Delta(p,\Delta) = \frac{\beta}{\psi}[\alpha_2 + \alpha_3 p + 2\alpha_5\Delta]$. *On the equilibrium path, signal manipulation takes value* $\mathtt{a}(p,0) = \frac{\beta}{\psi}[\alpha_2 + \alpha_3 p]$.

Proof: See the Appendix.

□

[18] The analysis can be easily extended to the case in which the market's action is also linear in a^*: $u(\mathtt{b}(p^*, a^*)) = u_0 + u_1(k_1 p^* + k_2 a^*) - u_2(k_1 p_t^* + k_2 a_t^*)^2$, $k_1, k_2 \in \mathbb{R}$.

The long-run agent's on-path utility takes the form $V(p,0) = \alpha_0 + \alpha_1 p + \alpha_4 p^2$, $p \in \mathbb{R}$ (see the Appendix for the coefficients' expressions), and the manipulation strategy that arises in equilibrium corresponds to a decreasing function of public beliefs ($\alpha_3 < 0$). This is intuitive as the agent wants to push public beliefs toward the payoff's bliss point $\frac{u_1}{2u_2}$. Graphically:

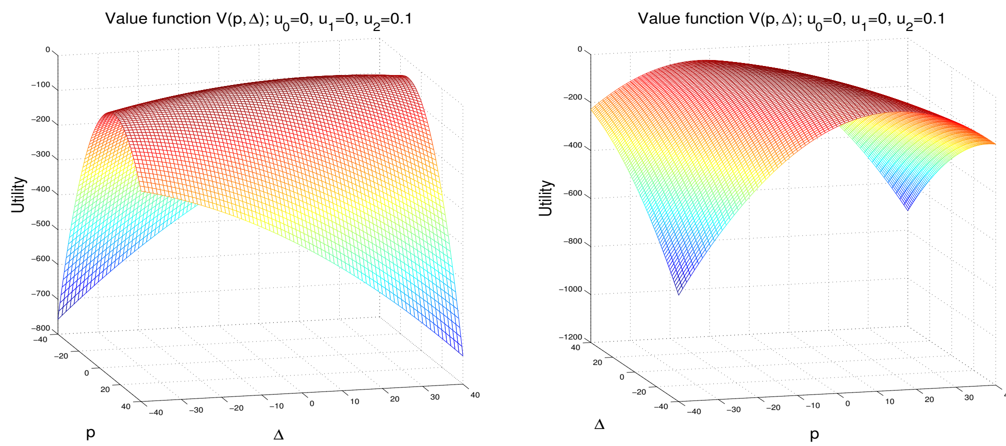

Figure 2.3: The agent's value function from two different angles. Parameter values: $\psi = 1$, $\eta = -1$, $\sigma_\theta = \sigma_\xi = 0.2$, $\kappa = 0.2$.

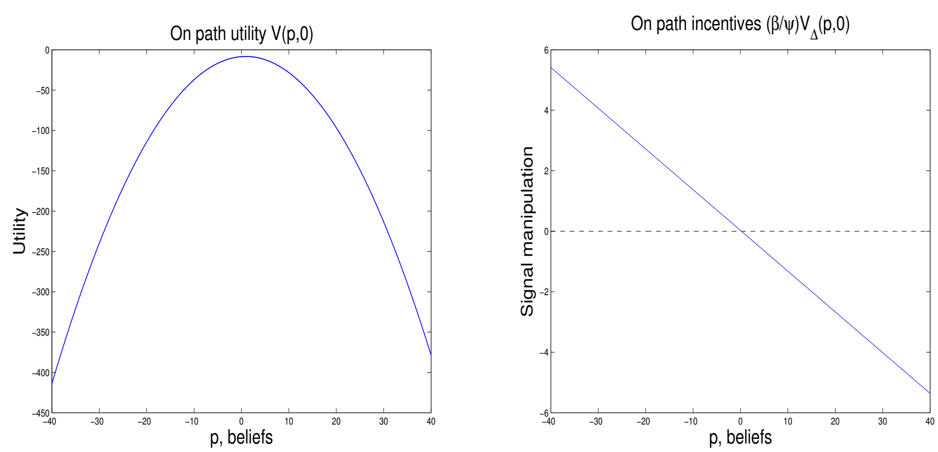

Figure 2.4: The long-run agent's utility and the optimal signal manipulation policy (right panel) on the equilibrium path. Parameter values: $\psi = 1$, $\eta = -1$, $\sigma_\theta = \sigma_\xi = 0.2$, $\kappa = 0.2$, $u_0 = u_1 = 0$, $u_2 = 0.1$.

Since the long-run agent's utility must remain uniformly bounded by above, α_4 and α_5 are strictly negative. This can be seen in Figure 6. The left panel in Figure

7 shows in turn that $\alpha_4 < 0$. Moreover, the choice of parameters satisfy $\eta < 0$ and $u_1 = 0$, we have that $\alpha_2 > 0$. This yields

$$\alpha_1 := \frac{-\frac{\beta^2\alpha_3}{\psi}\alpha_2 + 2\eta\kappa\alpha_4}{r+\kappa} > 0.$$

As a consequence, the long-run agent's (on path) utility is maximized at $-\frac{\alpha_1}{2\alpha_4} > 0$. Finally, the fact that the intercept α_2 is strictly larger than zero means that the agent is exerting strictly positive effort at the payoff's bliss point. I will elaborate more the properties of the linear equilibrium in the next section.

2.3.2 The Structure of the Agent's Incentives

The incentives generated within the class of linear-quadratic games satisfy all the forces identified in the incentives equations. First, the size of marginal flow payoffs drive the size of the long-run agent's incentives: as the myopic gain from belief manipulation decays, equilibrium effort decreases.

Second, the agent's incentives respond to anticipated economic conditions through cost smoothing. Suppose for instance that the linear component of the agent's utility u_1 is zero and that fundamentals mean revert toward $\eta < 0$. Then, since $\alpha_3 < 0$ and $r + \beta + \kappa + \frac{\beta^2\alpha_3}{\psi} > 0$, it can be easily seen that

$$\alpha_2 = \frac{\eta\kappa\alpha_3}{r+\beta+\kappa+\frac{\beta^2\alpha_3}{\psi}} > 0,$$

which means that the agent exerts *strictly positive effort* at the long-run payoffs' bliss point. This is because he anticipates that fundamentals will mean revert to η with high probability, region in which it is optimal to exert signal manipulation. Since the cost of signal manipulation is convex, it is optimal to invest in signal manipulation today.[19]

Finally, the rate at which a marginal unit of belief asymmetry depreciates over time is endogenous. In fact, for the linear-quadratic games studied the required rate of return on belief asymmetry is given by

$$\tilde{r}(p) := r - \phi_\Delta(p,0) + \frac{\beta^2}{\psi} V_{\Delta p}(p,0) = r + \beta + \kappa + \frac{\beta^2}{\psi}\alpha_3$$

from which we conclude that the effect of distorting the market's conjecture on the rate at which belief asymmetry decays over time is *uniform* across all levels of public beliefs. Interestingly, since $\alpha_3 < 0$, the equilibrium rate of return is smaller than $r + \beta + \kappa$ (the rate at which belief asymmetry would decay if equilibrium strategies were deterministic). This is because the agent signal manipulation moves the market conjecture a^* in the direction where it is less costly to keep up with the market's expectations. The ability to reduce the standard of belief manipulation imposed by the market, amplifies the future benefits from belief distortion. All these effects can be seen in the following figure:

[19]If $\kappa = 0$, public beliefs evolve as martingales and thus cost smoothing disappears. Yet for more nonlinear environments third order terms should make cost smoothing considerations important, even when beliefs are unpredictable.

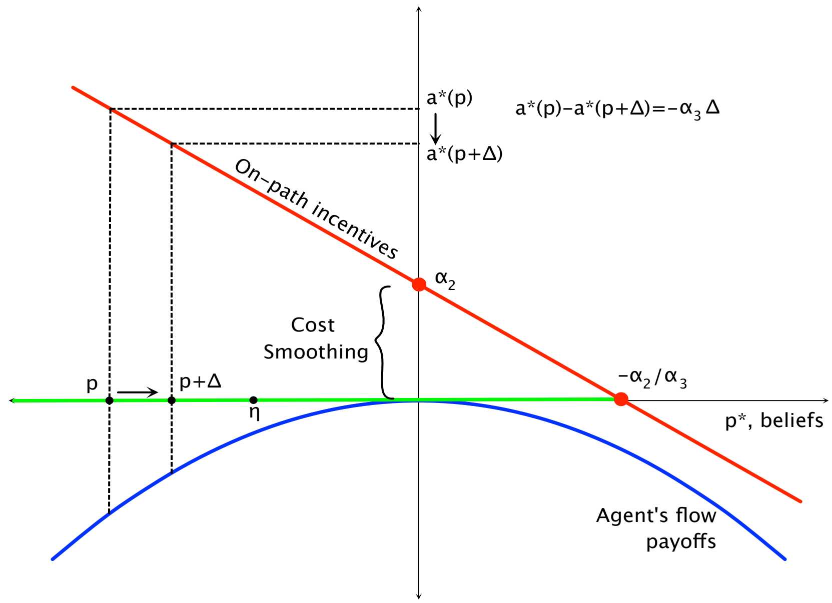

Figure 2.5: Two determinants of the size of incentives: cost smoothing and the effect that distorting the market's conjecture has on rate at which Δ depreciates.

The relevant parameters of the linear-quadratic model correspond to the rate of mean reversion κ, the long-run mean of fundamentals η and the convexity parameter of the effort's disutility function ψ. The sensitivity of equilibrium incentives to these parameters can be seen in the following panel:

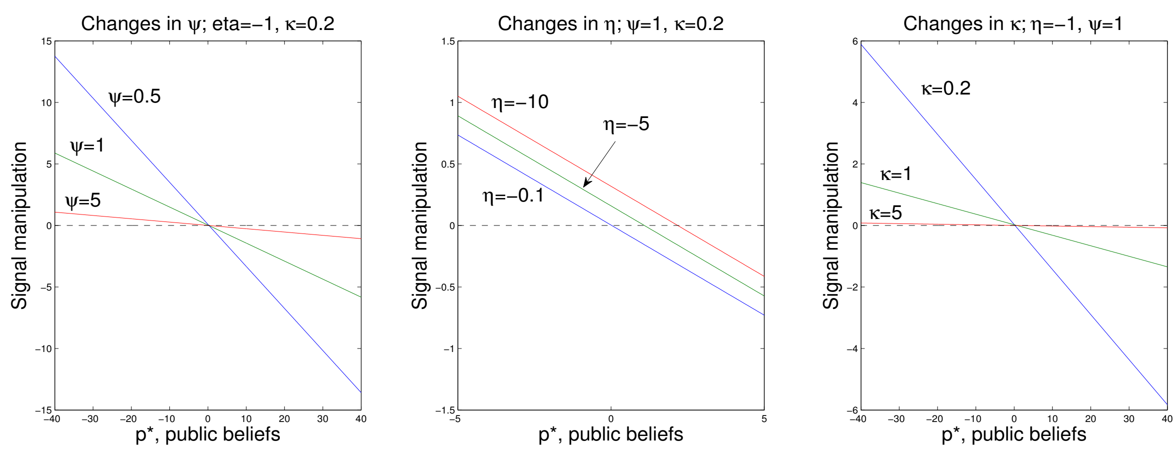

Figure 2.6: Sensitivity of equilibrium incentives to ψ, η and κ, respectively.

In the left panel, more convexity in the costs associated with signal manipulation entail higher costs from actions having high volatility. This results in less steep manipulation policies as ψ increases. The middle panel, instead shows how the distance between the agent's consumption bliss point and the long-run average value of fundamentals affect incentives. As this distance increases, it is more likely that fundamentals will move away from the consumption bliss point, which induces the agent to engage in more signal manipulation. Finally, changes in κ (third panel) can have two effects on incentives. First, as the rate of mean reversion increases there is a pressure towards more effort smoothing (numerator in α_2). However, an increase in κ also makes belief distortion less persistent ($\beta+\kappa = \sqrt{\kappa^2 + (\sigma_\theta/\sigma_\xi)^2}$) and public beliefs become less responsive to new information ($\beta = \sqrt{\kappa^2 + (\sigma_\theta/\sigma_\xi)^2} - \kappa$ decreases in κ).[20] Both less persistence and less sensitivity reduce the incentives for belief manipulation (slope of incentives decay) and the benefits from cost smoothing (denominator in (2.30)). The outcome is less steep incentives and a relatively unchanged intercept (which measure cost smoothing).

I conclude this section with an analysis of the dynamics of signal manipulation, on and off the equilibrium path:

Proposition 2.4. *When the environment is linear-quadratic, the long-run agent's optimal manipulation strategy satisfies the following properties*

(i) *Signal manipulation is positively correlated over time, on and off the equilibrium path. The correlation between signal manipulation at two points in time decreases with time distance.*

(ii) *Off the equilibrium path,*

$$\Delta_t = \Delta^o e^{\rho t},\ t \geq 0,\ \Delta_0 = \Delta^o$$

[20] Recall that in any equilibrium public beliefs evolve according to $dp_t = -\kappa(p_t - \eta)dt + \beta\sigma_\xi dZ_t^Y$, $t \geq 0$.

where $\rho < 0$. That is, discrepancies in beliefs gradually decrease over time, i.e. $\rho < 0$.

Proof: See the Appendix.

□

Part (i) is straightforward, as both parties' beliefs are mean reverting on and off the equilibrium path. The reason behind why (ii) holds has to do with both the fact that the long-run agent is risk averse and the fact that a Markov perfect equilibrium exists. First, risk aversion makes it optimal for the agent to induce dynamics of belief asymmetry that evolve deterministically. Second, since a Markov perfect equilibrium exists, the long-run agent does not benefit from large deviations off the equilibrium path. As a result, any initial stock of belief asymmetry must vanish asymptotically as time goes by, so $\rho < 0$. In particular, starting from a different prior (with the market's prior being common knowledge) both learning processes converge to the same ergodic distribution.

2.3.3 The Curvature Condition

Theorem (2.3) ensures the existence of a linear (in public beliefs) Markov Perfect Equilibrium provided the curvature condition

$$u_2 \leq \frac{\psi(r+\beta+2\kappa)^2}{8\beta^2},$$

holds. In this section I show that the curvature condition is also *necessary* for the existence of such a linear equilibria.

In order to understand the intuition behind this result, suppose that the curvature of the payoff function, u_2, increases. Then, the myopic benefit from inducing belief asymmetry goes up. Consequently, in order to prevent deviations off the equilibrium path, the market has to impose a tougher effort standard, which maps into a steeper (conjectured) effort profile a^*. Steeper conjectures of equilibrium play in turn translate into more persistent benefits from inducing belief-asymmetry. In fact, the required rate of return on belief asymmetry

$$\tilde{r}(p) = r + \beta + \kappa + \frac{\beta^2}{\psi}\frac{da^*}{dp^*}(p^*),\ p \in \mathbb{R},$$

decays as a^* becomes more negatively sloped, which means that the value attached to inducing belief asymmetry increases. This occurs because, by pushing public beliefs toward zero, the agent will face an even lower effort standard tomorrow, which increases the benefits from belief manipulation. Beyond the threshold $\bar{u} = \frac{\psi(r+\beta+2\kappa)^2}{8\beta^2}$, a linear effort schedule cannot control *simultaneously* both the local benefits from a deviation (as measured by marginal flow payoffs) and the long-run benefits from engaging in large deviations off the equilibrium path.

Formally, consider the optimal control problem $\mathcal{P}(\hat{\alpha}_3)$

$$\begin{aligned}
\max_{a \in \mathcal{A}} \quad & \int_0^\infty e^{-rt}\left(-u_2(p_t+\Delta_t)^2 - \frac{\psi}{2}a_t^2\right)dt \\
s.t. \quad & dp_t = -\kappa p_t dt, && (2.31)\\
& d\Delta_t = \left[-\left(\beta+\kappa+\frac{\beta^2}{\psi}\hat{\alpha}_3\right)\Delta_t + \beta a_t - \frac{\beta^2}{\psi}\hat{\alpha}_3 p_t\right]dt, && (2.32)
\end{aligned}$$

where $u_2, \kappa, \beta, \psi > 0$. Observe that the second dynamic corresponds to the belief-asymmetry process when the market conjectures an effort profile of the form $\hat{a}(p^*) = \frac{\beta\hat{\alpha}_3}{\psi}p^* = \frac{\beta\hat{\alpha}_3}{\psi}(p+\Delta)$, with $\hat{\alpha}_3$ a scalar. In other words, this problem corresponds to a deterministic version of the linear-quadratic game previously studied in the case in which u_1, η and the volatility term in the private beliefs process are all zero (β, however, is assumed to depend on $\sigma_\xi > 0$, as in the stochastic game).

Studying this problem is without loss of generality for two reasons. First, u_1 and η do not affect the slope of the effort schedule, which is what at the end of the day matters for the existence of a linear equilibrium. Second, since the original problem has a linear-quadratic structure, any second order term will only affect the level (or constant term) of the agent's value function. Thus, provided a linear best-response to the market's conjecture exists in the original stochastic problem, this one can be found through solving this deterministic version (certainty equivalence principle). In the sequel, $\hat{\alpha}_3 < 0$, which captures the idea that the market wants to prevent belief manipulation towards zero.[21]

The next results show that the previous problem always admits a linear best response. Furthermore, it shows that when the curvature condition is violated, the long-run agent responds *more aggressively* to the market's conjectured effort profile:

Proposition 2.5. *The value function associated with* $\mathcal{P}(\hat{\alpha}_3)$*,* $\hat{\alpha}_3 < 0$*, has the form*

$$V(p,\Delta) = \alpha_3 p\Delta + \alpha_4 p^2 + \alpha_5 \Delta^2,$$

where $\alpha_3 = \alpha_3(\hat{\alpha}_3)$ *and* $\alpha_5 = \alpha_5(\hat{\alpha}_3)$ *are given by*

$$\alpha_3(\hat{\alpha}_3) \;=\; \frac{-2u_2 - \frac{2\beta^2}{\psi}\alpha_5(\hat{\alpha}_3)\hat{\alpha}_3}{r+\beta+2\kappa+\frac{\beta^2}{\psi}\hat{\alpha}_3 - \frac{2\beta^2}{\psi}\alpha_5(\hat{\alpha}_3)} \tag{2.33}$$

$$\alpha_5(\hat{\alpha}_3) \;=\; \frac{r+2\left(\beta+\kappa+\frac{\beta^2}{\psi}\hat{\alpha}_3\right) - \sqrt{\left(r+2\left(\beta+\kappa+\frac{\beta^2}{\psi}\hat{\alpha}_3\right)\right)^2 + \frac{8\beta^2}{\psi}u_2}}{\frac{4\beta^2}{\psi}}. \tag{2.34}$$

Moreover, if $u_2 > \frac{\psi(r+\beta+2\kappa)^2}{8\beta^2}$*, then* $\alpha_3(\hat{\alpha}_3) < \hat{\alpha}_3$*.*

Proof: See the Appendix.

□

When the curvature condition is violated, the sensitivity of the long-run agent's actions to his private information is higher (in absolute value) than the slope of the market's conjectured effort schedule. Consequently, starting from a common prior,

[21]It is straightforward to argue that there is no equilibrium in which on-path effort is given by $\hat{\alpha}_3 p^*$, $p^* \in \mathbb{R}$ with $\hat{\alpha}_3 > 0$.

the long-run agent finds it optimal to generate a positive stock of belief asymmetry from time zero on. This allows him to drive public beliefs quickly towards zero, and the stock of belief asymmetry disappears asymptotically.

2.3.4 Connection with the Literature: Career Concerns

Career concerns models are an example of signal-jamming games in which the actions taken by the long-run agent create value. As argued earlier, in Holmstrom (1999) the long-run agent is risk neutral, learning is Gaussian and wages are linear un public beliefs. All these these linearities result in reputational incentives that do not respond to current wages (e.g., perceived ability) or to the market's contemporaneous conjectures. Since public beliefs evolve as martingales, equilibrium incentives neither respond to future economic conditions, as these are unpredictable.

In Cisternas (2012) I extend Holmstrom's environment to allow for some degree of predictability in future economic conditions. In particular, I show that when skills are modeled as a mean-reverting process (2.7) and learning is stationary, market competition generates incentives that are constant at a value a^H such that

$$g'(a^H) = \frac{\beta(\kappa)}{r + \beta(\kappa) + \kappa}, \tag{2.35}$$

where $\beta(\kappa) = \sqrt{\kappa^2 + (\sigma_\theta/\sigma_\xi)^2} - \kappa$ and $g(\cdot)$ is the effort cost function. An important conclusion of this result is that in contexts of linear preferences and linear learning, introducing predictability in the reputational component of wages does not affect the structure of the incentives created by career concerns; rather, it is only their level that changes.[22] In particular, mean-reversion bounds incentives away from the efficiency (defined as $g'(a^e) = 1$) uniformly across all positive discount rates.

[22] Away from the steady-state level of learning, the dynamics are as in Holmstrom (1999) with incentives driven evolution of the mean-squared error, which is deterministic. Incentives are also deterministic when the agent's actions affect his productivity (an investment game), as in learning-by-doing. See Cisternas (2012) for more details.

The literature on the structure of the incentives created by career concerns is not extensive. In the static model of Dewatripont et. al. (1999b) effort and conjectured effort are *strategic complements* as a result of the complementarity between skills and effort in the production technology. More recently, Bonatti and Hörner (2012) show that current and future effort can become *strategic substitutes*: when an agent is rewarded by effort and builds a reputation through a single output observation, high wages in the future reduce the value that a worker attaches to obtaining a breakthrough today. In Bonatti and Hörner's model beliefs evolve deterministically and therefore, the whole evolution of the game (in the absence of a breakthrough) is fully known at time zero. Finally, in a discrete-time and finite-horizon framework Martinez (2009) finds ratchet effects when studying career-concerns incentives in a context of piecewise linear wages arising from job assignments.

Instead, the linear-quadratic model presented here shows how incentives are affected by contemporaneous conjectures even in the absence of complementarities in the public signal's technology. It also shows how incentives also respond to anticipated future economic conditions through effort smoothing. Finally, it shows that ratchet forces appear through the slope of the market's contemporaneous conjecture of belief manipulation. Most importantly, these features are present in every environment exhibiting some degree of nonlinearity.

2.4 Investment Games

This section outlines the basic ingredients of a class of investment games with learning. Given the analysis performed in the Section 3, setting a framework for the study of markovian equilibria should be conceptually simple.

2.4.1 Model and Learning Dynamics

An investment game consists of a public signal $\xi := (\xi_t)_{t\geq 0}$ of the form

$$d\xi_t = \theta_t dt + \sigma_\xi dZ_t,\ t \geq 0, \tag{2.36}$$

where $\theta := (\theta_t)_{t\geq 0}$ is a hidden state variable affected by the long-run agent's investment decision $a := (a_t)_{t\geq 0}$. For technical reasons I restrict the analysis to fundamentals that evolve according to

$$d\theta_t = (a_t - \kappa\theta_t)dt + \sigma_\theta dZ_t^\theta,\ t \geq 0. \tag{2.37}$$

If $\kappa > 0$ then, at any point in time t, fundamentals locally mean revert to a_t/κ, and this convergence is gradual. If $\kappa = 0$, the long-run agent's action just adds a non-zero drift to the dynamics of θ. Denote by $\theta^a := (\theta_t^a)_{t\geq 0}$ the process of fundamentals that arises when the long-run agent is following an investment strategy a, and let $\xi^a := (\xi_t^a)_{t\geq 0}$ denote the resulting public signal.

In what follows I rapidly move toward a markovian formulation of the long-run agent's problem in which the key points are: 1) deriving filtering equations for the agent's and the market's posterior beliefs, 2) connecting public and private beliefs via a belief-asymmetry process and 3) showing that given any markovian conjecture by the market, a best response in markovian form can be found through solving a well-defined stochastic control problem.[23]

[23] This last point is a non-trivial one and is related to the choice of fundamentals (3.7). In finding optimal markovian policies using a dynamic-programming approach, both the Brownian motions and associated filtrations must be exogenous. This is clearly satisfied in the class of signal manipulation games. However, this is not necessarily the case for investment games, as the information generated by the public signal can depend on the actions taken by the agent though the evolution of θ. Under the specification (3.3)-(3.7) this issue is circumvented through the use of the Separation Principle (Wonham (1968)).

Suppose that the agent follows a strategy $a := (a_t)_{t\geq 0}$ which depends on the current observations of the public signal. As in the signal-jamming case, standard results in filtering theory state that, starting from a Gaussian prior, the posterior distribution of θ_t^a given $\mathcal{F}_t^{\xi^a}$ is normally distributed (Theorem 12.1 in Liptser and Shiryaev (1977)). Moreover, $p_t^a := \mathbb{E}[\theta_t^a|\mathcal{F}_t^{\xi^a}]$ and $\gamma_t^a := \mathbb{E}[(\theta_t^a - p_t^a)^2|\mathcal{F}_t^{\xi^a}]$ evolve according to

$$\begin{aligned} dp_t^a &= (a_t - \kappa p_t^a)dt + \frac{\gamma_t^a}{\sigma_\xi}\frac{d\xi^a - p_t^a dt}{\sigma_\xi} \\ \dot{\gamma}_t^a &= -2\kappa\gamma_t^a + \sigma_\theta^2 - \left(\frac{\gamma_t^a}{\sigma_\xi}\right)^2. \end{aligned} \tag{2.38}$$

where $\frac{1}{\sigma_\xi}\left(\xi_t^a - \int_0^t p_s^a ds\right)$, $t \geq 0$, is an $\mathbb{F}^{\xi^a}$–Brownian motion. Observe that the posterior variance does not depend on a, and hence it is exogenous. As in the previous section, I consider the case in which γ is fixed at the steady-state level of learning γ^*.

However, observe the long-run agent's actions could also depend on his private belief about the unobserved process– and this private belief is in turn affected by the actions he takes. As a result, in order to have well defined posterior beliefs, the following condition is necessary:

Definition 2.5. *A strategy $a := (a_t)_{t\geq 0}$ is said to be feasible if*

(i) *For each t, $a_t : C([0,t)) \times \mathbb{R} \to \mathbb{R}$, $t \geq 0$, and;*

(ii) *There exists a unique $\mathbb{F}^{\xi^a}$– measurable solution $(p_t^a)_{t\geq 0}$ to the SDE in $(x_t)_{t\geq 0}$*

$$dx_t = (a_t(\xi^a, x_t) - \kappa x_t)dt + \frac{\gamma^*}{\sigma_\xi}\frac{d\xi^a - x_t dt}{\sigma_\xi},\ t \geq 0, \tag{2.39}$$

where ξ^a and θ^a are the signal's and fundamental's processes, respectively, under the strategy a.

Part (i) in the previous definition states that at any point in time t the long-run agent can condition his actions on the partial observations of the public signal $(\xi_s : s \in [0,t])$ and also on his beliefs about the current state of fundamentals, $t \geq 0$. Part (ii) is the minimal requirement that the agent strategy must generate posterior beliefs characterized by SDEs.[24]

The market's learning process is similar to the one studied in the class of signal-jamming games. The market conjectures an investment strategy $a^* := (a_t^*)_{t\geq 0}$ which depends on the information generated by ξ, under the assumption that the long-run agent has followed a^*. Denote by $p^{a^*} := (p_t^{a^*})_{t\geq 0}$ the corresponding posterior mean process. The following is an analog result to Proposition 2.1 for the case of investment games, and its proof is straightforward:

Proposition 2.6. *Suppose that the market conjectures a manipulation strategy* $a^* := (a_t^*)_{t\geq 0}$*, while the long-run agent actually follows* $a := (a_t)_{t\geq 0}$*. Then, from the long-run agent's perspective, public beliefs can be written as* $p_t^{a^*} = p_t^a + \Delta_t$*, where the process* $\Delta := (\Delta_t)_{t\geq 0}$ *is governed by the ODE*

$$d\Delta_t = [-(\kappa+\beta)\Delta_t + a_t^* - a_t]dt, t > 0, \tag{2.40}$$

with $\beta := \gamma^*/\sigma_\xi^2$*,* $t \geq 0$ *and* $\Delta_0 = \Delta^o$*.*

Proof: Straightforward.

□

Observe that, as opposed to a signal-jamming game, the gap between the agent's actions and the market's conjecture, $a - a^*$, *reduces* the size of the stock of belief asymmetry. This is because the agent is affecting the evolution of fundamentals,

[24]Such a solution exists when, for instance, a is linear in x: $a_t(\xi;x) = \hat{a}_t(\xi) + x$, for some family of measurable functions $\hat{a}_t : C([0,t)) \to \mathbb{R}$, $t \geq 0$, satisfying mild integrability conditions. Refer to chapters 11 and 12 in Liptser and Shiryaev (1977) for the analysis of this case.

rather than distorting a noisy signal of them. The higher the agent's investment relative to the market's conjecture, the more optimistic the agent will be about current fundamentals compared to the market.

2.4.2 Necessary Conditions for Markov Perfect Equilibria

Off the equilibrium path, the market's beliefs take the form $a^*(p_t^a + \Delta_t)$ where $a := (a_t)_{t\geq 0}$ denotes the strategy followed by the agent, $t \geq 0$, and $a^* : \mathbb{R} \to \mathbb{R}$, is a markovian conjecture. Private beliefs and the belief asymmetry process evolve according to

$$dp_t^a = (a_t - \kappa p_t^a)dt + \beta\sigma_\xi d\overline{Z}_t^a \tag{2.41}$$

$$d\Delta_t = [-(\kappa + \beta_t)\Delta_t + a^*(p_t^a + \Delta_t) - a_t]dt, \ t \geq 0, \tag{2.42}$$

where $\overline{Z}_t^a := \frac{1}{\sigma_\xi}\left(\xi^a - \int_0^t p_s^a ds\right)$, $t \geq 0$, is an $\mathbb{F}^{\xi^a}-$Brownian motion from the agent's perspective. But because the (θ, ξ) is conditionally Gaussian,

$$\overline{Z}_t^a = \overline{Z}_t^{\tilde{a}}, \ t \geq 0, \text{ a.s.},$$

for any pair of feasible strategies a, $\tilde{a}$.[25] Denoting by $\overline{Z}_t^0$ the Brownian motion associated with the strategy 0, observe that the posterior mean process can be written as

$$dp_t^a = (a_t - \kappa p_t^a)dt + \beta\sigma_\xi d\overline{Z}_t^0, \ t \geq 0,$$

where $\overline{Z}^0$ is an exogenous Brownian motion from the agent's perspective.

Under the above conditions the Separation Principle applies. That is, if a best response to a^* in markovian form exists, then it can be found by solving the modified problem

[25] This can be deduced from (3.7) and (2.38).

$$\max_{a \in \mathcal{A}} \quad \mathbb{E}\left[\int_t^{\infty} e^{-r(s-t)}(u(p_s, p_s + \Delta_s) - g(a_s))ds\right]$$
$$s.t. \quad dp_s = (a_s - \kappa p_s)dt + \beta\sigma_\xi d\overline{Z}_s^0, \; s \geq t,$$
$$d\Delta_s = [-(\kappa + \beta_s)\Delta_s + a^*(p_s + \Delta_s) - a_s]ds, \; s \geq t,$$

where I have assumed that the agent's *expected* flow payoff at time t takes the form $u(p_t, p_t^*)$, $t \geq 0$.[26]

The remaining steps toward obtaining necessary conditions for on-path incentives are straightforward. If a best response in markovian form $\hat{a} : \mathbb{R}^2 \to \mathbb{R}$ to a^* exists, and the associated value function $V^{a^*}(p, \Delta)$ is smooth enough then

$$\hat{a}(p, \Delta) = \max_{a \in A} \left\{a[V_p^{a^*}(p, \Delta) - V_\Delta^{a^*}(p, \Delta)] - g(a)\right\}, \; (p, \Delta) \in \mathbb{R}^2.$$

This yields the equilibrium condition

$$a^*(p) = \max_{a \in A} \left\{a[V_p^{a^*}(p, 0) - V_\Delta^{a^*}(p, 0)] - g(a)\right\}, \; p \in \mathbb{R}.$$

The next result states that in games of investment and learning, the benefits from investing in the fundamental are, in any equilibrium, characterized by a system of ODEs: one for the value attached to boosting fundamentals, and another one for the value attached to inducing a small discrepancy between private and public beliefs. This is because the long-run agent affects both fundamentals directly through his

[26] The separation principle states that if we replace the dynamics (2.41)-(2.42) by the corresponding ones under an *exogenous Brownian motion and filtration*, then, provided the modified problem admits an optimal strategy in Markovian form, such a policy is optimal among the class of controls which are measurable with respect to the public signal. See Chapter 16 in Liptser and Shiryaev for an application of the Separation Principle to a linear-quadratic example, or Wonham (1968) for a more general approach.

investment decisions, and the belief asymmetry process indirectly through changes in perceived fundamentals:

Proposition 2.7. *Assume that a Markov perfect equilibrium* $\mathsf{a} : \mathbb{R}^2 \to \mathbb{R}$ *exists and the associated value function* V^{a} *is of class* $C^{2,1}(\mathbb{R}^2)$. *Furthermore, suppose that at a level of public beliefs* $\hat{p}$

(i) *on path incentives are interior and*

(ii) *on path incentives are locally twice continuously differentiable with respect to public beliefs, i.e., there exists a neighborhood* $\mathcal{O}$ *of* $\hat{p}$ *such that* $\mathsf{a}(\cdot, 0) \in C^2(\mathcal{O})$.

Then, $g'(\mathsf{a}(\cdot, 0)) = W^{\mathsf{a}}(\cdot, 0) \equiv V_p^{\mathsf{a}}(\cdot, 0) - V_\Delta^{\mathsf{a}}(\cdot, 0)$ *in* $\mathcal{O}$. *Moreover,* $W^{\mathsf{a}}(\cdot, 0)$ *and* $V_\Delta^{\mathsf{a}}(\cdot, 0)$ *satisfy the system of ODEs in* $p \mapsto (W(p, 0), V_\Delta(p, 0))$

$$\begin{aligned}
(r+\beta)W(p,0) &= u_p(p,p) - \beta V_\Delta(p,0) + [g'^{-1}(W(p,0)) - \kappa p]W_p(p,0) \\
&\quad + \frac{1}{2}(\beta\sigma_\xi)^2 W_{pp}(p,0), \\
\tilde{r}(W(p,0))V_\Delta(p,0) &= u_{p^*}(p,p) + V_{p\Delta}(p,0)[g'^{-1}(W(p,0)) - \kappa p] \\
&\quad + \frac{1}{2}(\beta\sigma_\xi)^2 V_{pp\Delta}(p,0), \;\; p \in \mathcal{O}
\end{aligned}$$

with $\tilde{r}(W(p,0)) := r + \beta + \kappa - \frac{d}{dp} g'^{-1}(W(p,0))$.

Proof: See the Appendix.

□

The interpretation of these equations is as in the class of signal manipulation games. The main difference from that case is that now the agent's incentives are driven by the difference between the benefits from boosting fundamentals ($V_p(p, 0)$) and the benefits from belief manipulation ($V_\Delta(p, 0)$). This is captured in the first ODE for $W(p, 0) = V_p(p, 0) - V_\Delta(p, 0)$. However, since the dynamics of private beliefs and belief asymmetry are inherently different, it is not possible to summarize the long-run agent's incentives in a single ODE.

2.4.3 Application: Monetary Policy and Unobserved Inflation

In this section I revisit the monetary policy application presented in Section 2. This example gauges the effectiveness of inflationary surprises in affecting the employment level of an economy in which agents try to learn about an unobserved component of inflation.

Recall from Section 2.1 that the fundamental (or trend) inflation rate in an economy is given by

$$d\theta_t = (a_t - \kappa\theta_t)dt + \sigma_\theta dZ_t^\theta,\ t \geq 0.$$

The term a_t represents the action of the central bank at time t (e.g. money growth), which is unobserved by the rest of the economy, $t \geq 0$. The Brownian motion $Z^\theta := (Z_t^\theta)_{t\geq 0}$ in turn captures *unobserved* shocks to the state of the economy, thus making fundamental inflation a hidden state variable. Consequently, all the agents in the economy (including the policymaker) learn about θ through observing a public signal $d\xi = \theta_t dt + \sigma_\xi dZ_t,\ t \geq 0$. Finally, the Brownian motion $Z := (Z_t)_{t\geq 0}$ represents domestic shocks to prices beyond the central bank's control.

The central bank's payoffs are given by

$$\int_0^\infty e^{-rt}[k_1(d\xi_t - p_t^* dt) - k_2(\theta_t - \overline{\theta})^2 dt - \frac{\psi}{2}a_t^2 dt],$$

where $p^* := (p_t^*)_{t\geq 0}$ represents the public beliefs about unobserved inflation, and $d\xi_t - p_t^* dt$ the change in employment due to unanticipated realized inflation, $t \geq 0$.[27]

[27] Cukierman and Meltzer (1986) study a central bank's incentives to generate inflationary surprises in a context of imperfectly observable actions and adverse selection. In their model, a central bank's preferences for economic stimulation are subject to private shocks. The market in turn filters the central bank's actions though observing a noisy signal of them. In my setting, the central bank's preferences are common knowledge and the market instead learns about the *mapping* between the central bank's actions and inflation.

Observe that when $\psi \equiv 0$ the bank chooses $\hat{\theta}$ such that it solves

$$\max_{z \in \mathbb{R}} \; k_1(z - p^*) - k_2(z - \overline{\theta})^2.$$

This yields an optimal inflation level of the form $\hat{\theta} = \overline{\theta} + \frac{k_1}{2k_2}$. The market's expectations are thus set at $p^* = \hat{\theta}$ and it is optimal for the bank to choose such a target. The outcome is that no economic stimulus is possible and the economy suffers from high inflation rates.

Optimal Policy: Full-commitment Case

Suppose for the moment that θ is observable. Then, unexpected changes in the signal ξ are driven only by the domestic shocks $Z := (Z_t)_{t \geq 0}$. Consequently, the monetary authority's problem consists of

$$\begin{aligned} \max_{a \in \mathcal{A}} \quad & \mathbb{E}\left[\int_0^\infty e^{-rt}\left(-k_2(\theta_t - \overline{\theta})^2 - \frac{\psi}{2}a_t^2\right) dt\right] \\ s.t. \quad & d\theta_t = (a_t - \kappa\theta_t)dt + \sigma_\theta dZ_t^\theta. \end{aligned}$$

This results in the following

Proposition 2.8. *When fundamentals are observed, the central bank's optimal policy is of the form*

$$a^o(\theta) = \frac{\alpha_1^o + 2\alpha_2^o\theta}{\psi},$$

where

$$\begin{aligned} \alpha_1^o &= \frac{-2k_2\overline{\theta}}{\frac{2\alpha_2^o}{\psi} - (r + \kappa)} \\ \alpha_2^o &= \frac{\psi}{2}\left[(r + 2\kappa) - \sqrt{(r + 2\kappa)^2 + \frac{8k_2}{\psi}}\right] < 0. \end{aligned} \tag{2.43}$$

Also, $a^o(\overline{\theta}) = \frac{\alpha_2^o \kappa}{2\alpha_2^o - r\psi} > 0$ *if* $\kappa > 0$.

Proof: See the Appendix.

$\square$

In other words, since the central bank cannot generate inflationary surprises, it cares about controlling inflation only. This results in a monetary policy rule that is countercyclical.

Observe that if instead θ is hidden, but the central bank has commitment power, then $p \equiv p^*$. Since no economic stimulus is possible, the central bank's problem corresponds to

$$\begin{aligned} &\max_{a \in \mathcal{A}} \quad \mathbb{E}\left[\int_0^\infty e^{-rt}\left(-k_2(p_t - \overline{\theta})^2 - \frac{\psi}{2}a_t^2\right)dt\right] \\ &s.t. \quad dp_t = (a_t - \kappa p_t)dt + \beta\sigma_\xi d\overline{Z}_t^0 \end{aligned}$$

which has the same structure as the one just solved in Proposition 2.8, except for the volatility term. Since the optimal rule in Proposition 2.8 does not depend on the size of the volatility of fundamentals, it is also optimal in this case.[28] Graphically:

[28]Because the payoff function is quadratic, the volatility term only affects the *level* of the central bank's value function.

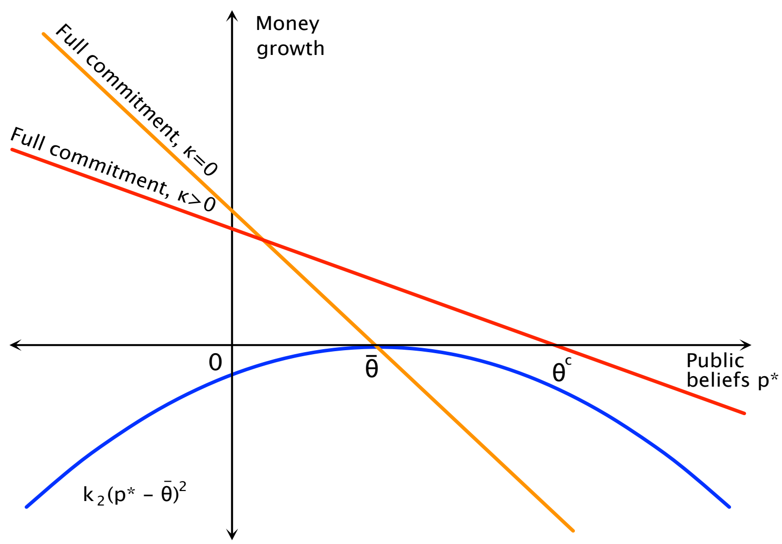

Figure 2.7: Full-commitment rule for $\kappa = 0$ and $\kappa > 0$.

Introducing mean reversion toward zero generates an optimal monetary rule that entails positive money growth at the inflation target $\overline{\theta}$. Furthermore, because the costs of monetary growth are quadratic and it takes time for inflation to move toward the moving-average a_t/κ, it is costly for the central bank to sustain a monetary rule that entails high volatility. As a result, monetary policy is less aggressive than in the random walk case ($\kappa = 0$), and the central bank's monetary rule is less steep. Therefore, whenever the inflation target is large relative to the level which fundamental inflation reverts to, the central bank chooses to exert inflationary pressures over the region $[\overline{\theta}, \theta^c]$ in exchange of more long-run stability.

Optimal Policy: No Commitment

When the central bank lacks commitment, the relevant state variables correspond to the private beliefs process p and the belief-asymmetry process $\Delta := p^* - p$. If learning is stationary, their dynamics are given by

$$\begin{aligned} dp_t &= (a_t - \kappa p_t)dt + \beta\sigma_\xi d\overline{Z}^0_t \\ d\Delta_t &= [-(\beta+\kappa)\Delta_t + (a^*_t - a_t)]dt \end{aligned}$$

with $\beta := \gamma^*/\sigma^2_\xi = \sqrt{\kappa^2 + (\sigma_\theta/\sigma_\xi)^2} - \kappa$.

The central bank's ex-ante flow payoffs take the form $-k_1\Delta_t - k_2(p_t - \overline{\theta})^2 - \frac{\psi}{2}a^2$, $t \geq 0$, (up to an additive constant). Observe that, from the central bank's perspective, the effect on unemployment is determined by Δ. The following result establishes the existence of a monetary policy rule as an equilibrium of the learning game between the monetary authority and the market:

Theorem 2.4. *When θ is unobserved, the central bank's optimal policy takes the form*

$$\alpha^*(p) = \frac{\alpha^o_1 + 2\alpha^o_2 p}{\psi} + \frac{\alpha^*_3}{\psi} \tag{2.44}$$

with α^o_1 and α^o_2 as in the observable case, and

$$\alpha^*_3 = -\frac{k_1}{\frac{2\alpha^o_2}{\psi} - (r+\beta+\kappa)} = \frac{k_1}{r + \sqrt{\kappa^2 + (\sigma_\theta/\sigma_\xi)^2} - \frac{2\alpha^o_2}{\psi}} > 0$$

When $\kappa = 0$, $a^*(\hat{\theta}) < 0$.

Proof: See the Appendix.

□

The previous result shows that moral hazard generates belief-manipulation incentives (inflationary surprises) which are *uniform* across all levels of public beliefs. The size of the shift is given by the term

$$-V_\Delta(p,0) = \alpha_3^* = -\frac{k_1}{r+\beta+\kappa-\frac{2\alpha_2^o}{\psi}} = \frac{k_1}{r+\sqrt{\kappa^2+(\sigma_\theta/\sigma_\xi)^2}-\frac{2\alpha_2^o}{\psi}} > 0,$$

which measures the size of the inflationary bias created by a monetary authority who lacks commitment.

The size of the benefits from inflation surprises affecting employment are driven by the of degree of persistence of Δ. This degree of persistence is determined in equilibrium and can be observed in the denominator for α_3^*. First, belief distortions naturally decay at a rate $\beta+\kappa$. The higher this rate, the lower the effect on employment. This natural rate of depreciation increases with environmental uncertainty σ_θ and decreases with the volatility of domestic shocks σ_ξ. In both cases, beliefs discount past information more heavily, as past prices become less accurate predictor of current fundamentals. Second, by controlling inflation the central bank also affects the market's anticipated level of monetary policy in the near future. As the full commitment rule becomes more aggressive ($|\alpha_2^o|/\psi$ increases) conjectures about monetary growth react more strongly to changes in prices. This is costly for the central bank, as inflation is away from the target $\overline{\theta}$ more frequently. The monetary authority takes into account these reaction by moderating the shift in its optimal rule by $-\frac{2\alpha_2^o}{\psi} > 0$. Graphically:

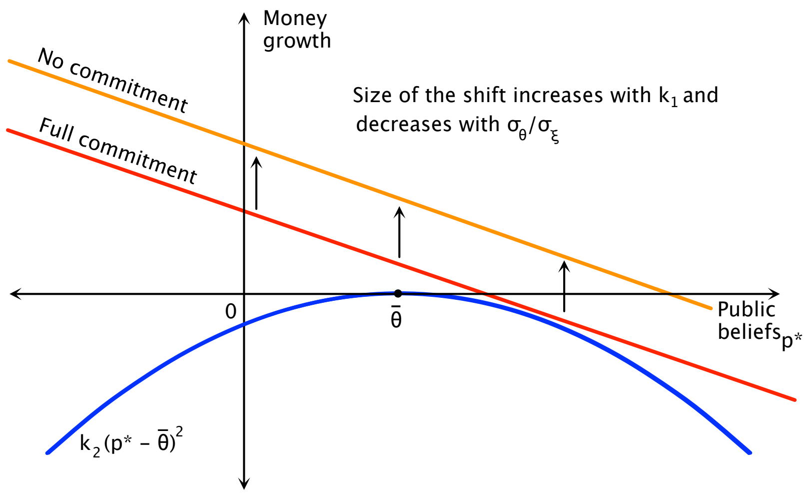

Figure 2.8: Non-commitment rule and the inflationary bias.

The literature on rules versus discretion is large. In the context of perfectly observable actions, Barro and Gordon (1983) show how a central bank's reputational forces can induce more commitment. Instead, when a central bank's actions are imperfectly monitored, Cukierman and Meltzer (1986) uncover the trade-off between transparency and flexibility that a policymaker faces in choosing which monetary policy to follow. Both papers assume adverse selection regarding the central bank's preferences. The model analyzed here instead assumes that the central bank's preferences are common knowledge. Yet, the plausible combination of imperfectly observable actions and learning about an unobserved component of inflation generate similar incentives.

2.5 Discussion

2.5.1 Smoothness and Robustness

More generally, given any well-behaved conjecture $a^*(\cdot)$, the agent's value function V^{a^*} is a *viscosity* solution to the HJB equation (2.23).[29] As a result, even when an equilibrium exists, the resulting value function may not be smooth enough to satisfy the corresponding PDE in a classical sense, thereby invalidating a partial characterization of incentives via ordinary differential equations. Nevertheless, regardless of the differentiability properties of the value function, the *forces that determine incentives remain unchanged*. In fact, in the case of signal-jamming games, it is optimal for the long-run agent to choose a strategy $\mathtt{a}(p, \Delta)$ that solves

$$
\begin{aligned}
\max_{a \in \mathcal{A}} \quad & \mathbb{E}\left[\int_0^\infty e^{-rt}(u(\mathtt{b}(p_t + \Delta_t)) - g(a_t))dt\right] \\
s.t. \quad & dp_t = \mu(p_t)dt + \sigma_p dZ_t^Y,\ t > 0,\ p_0 = p^o, \\
& d\Delta_t = [\phi(p_t, \Delta_t) + \beta(a_t - \mathtt{a}(p_t + \Delta_t, 0))]dt,\ t > 0,\ \Delta_0 = \Delta^o. \quad (2.45)
\end{aligned}
$$

From the previous problem, we see that the forces that shape the agent's incentives are exactly (i) varying marginal flow payoffs, (ii) convex costs of signal manipulation and (iii) the impact of the agent's actions on the market's conjecture about equilibrium play. Whether incentives are characterized by differential equations or by differential inclusions instead, it is just a matter of differentiability.[30] The structure behind the incentives that drive belief manipulation motives remains unchanged.

[29] See Crandall, Ishii and Lions (1990) for a survey of the theory of viscosity solutions to partial differential equations.

[30] For a characterization of optimal controls in terms of value functions that are not smooth enough, see chapter 6 in Yong and Zhou (1999).

The *uniform ellipticity* condition is largely the most well known condition that guarantees the existence of smooth solutions to linear PDEs.[31] Strulovici and Szydlowski (2012) use it to show the existence of smooth solutions to HJB equations in one-dimensional problems. Uniform ellipticity is not applicable to the class of games presented here, as the diffusion matrix is degenerate (only one innovation process, and two dynamics of posterior beliefs). Nonetheless, this condition is only sufficient and not necessary, as shown in the linear-quadratic class of games studied.

2.5.2 Computation: Off-Equilibrium Analysis

The PDE that captures global behavior is, in almost all situations, hard to visualize. However, this chapter offers a *localizing* method to the numerical approximation of markovian equilibria in pure strategies. More specifically, the necessary conditions for incentives derived, for example, in Proposition 2.3, reduce the class of functions over which markovian equilibria are to be found. A good guess on the boundary conditions of the specific problem at hand yields a fully specified conjecture a^* solution of the incentives equation (2.25). Since off the equilibrium path public beliefs take the form $p + \Delta$, it is possible to construct a function $f(p, \Delta) := a^*(p + \Delta)$, which in turn can be used as an input in the corresponding HJB equation. The resulting PDE is therefore *local* and numerical methods to solve these class of equations thus apply. If it turns out that the solution found satisfies the equilibrium condition (2.24) along the equilibrium path, then a^* would correspond to a numerical approximation to a Markov perfect equilibrium. This numerical approach is a feasible avenue to study in more detail off-equilibrium dynamics in more nonlinear settings.

[31] The volatility matrix $\Sigma \in \mathcal{M}_{n\times n}$ of a diffusion $dX_t = \mu(X_t)dt + \Sigma dZ_t$ satisfies the uniform ellipticity condition if there is a constant $\mu > 0$ such that $\sum_{i,j=1}^{n} \Sigma_{ij} d_i d_j \geq \mu \|d\|$ for all $d \in \mathbb{R}^n$, where Σ_{ij} is the $(i, j)-$th component of Σ.

2.6 Conclusions

In this chapter I developed a general class of games that incorporate both learning and belief manipulation as their main features. I provided necessary conditions for Markov Perfect Equilibria at a very general level, and showed the existence of markovian equilibria for a subclass of games with linear-quadratic structure. Most importantly, the methods and results presented here can be used to understand strategic behavior in a wide set of environments, ranging from the determinants of workers' incentives in labor markets, to central banks' behavior in response to unobserved states of the economy.

Rather than re-iterating the novel properties of equilibria that I find and their implications in different economic settings, in these concluding remarks I discuss three important topics to be addressed in the future: existence of equilibria, the symmetry in the model's information structure and the signal's technology assumption.

The question of existence of markovian equilibria in pure strategies is a difficult one. In particular, the combination of imperfectly observable actions and the full support assumption in the signal structure generates off-equilibrium behavior that is summarized in PDEs that are more complex than usual. While a curvature condition on the agent's payoffs ensured the existence of equilibria in the class of linear-quadratic games, it is still unknown whether an analogous condition also ensures existence in more nonlinear environments. Finding explicit sufficient conditions for the existence of pure-strategy Markov equilibria is undoubtedly an important question for future research.[32]

[32]Faingold and Sannikov (2011) find conditions that guarantee the existence of markovian equilibria in a general class of reputation games with one-sided and non-linear learning. However, the methods used in that paper do not carry over to settings where value functions are characterized by partial differential equations. In the environments they analyze the issue of existence is reduced to the study of solutions to second order differential equations.

The results obtained in this chapter rely on the economy's ex-ante symmetric (yet incomplete) information structure about the fundamentals. If the long-run agent for instance were to privately observe θ, he could condition his actions on his private information about the shocks to fundamentals, and the symmetry of the model would break. An equilibrium concept in which a market perfectly anticipates the agent's actions would not be appropriate. Identifying environments that allow for a tractable analysis of belief-manipulation incentives in the presence of ex-ante asymmetric information is another area of research that has wide economic applications. In the class of investment games, the recent work of Board and Meyer-ter-Vehn (2010b) is particularly interesting.

Finally, the results in this chapter also rely on the manipulation technology having an additively-separable structure. Allowing for complementarities between actions and fundamentals (either in the signal or in the fundamentals' process itself) creates another channel for incentives: experimentation. By studying models with an additively-separable structure, I am able to eliminate the experimentation effect and concentrate only on belief manipulation motives. Nonetheless, the model's formulation and the envelope methods used to characterize incentives have a direct analog in such non-separable settings.

2.7 Appendix

Proof of Proposition 2.2: From Lemma 2.1 the agent's and market's posterior variance evolve under the same dynamic (2.9). The public posterior mean is in turn given by

$$dp_t^* = -\kappa(p_t^* - \eta)dt + \frac{\gamma_t}{\sigma_\xi}\frac{d\xi_t - (a_t^* + p_t^*)dt}{\sigma_\xi}, t > 0,\ p_0^* = p^{o*},$$

where a^* the market's conjectured strategy (an $\mathbb{F}^\xi-$ progressively measurable process). Now, when the agent follows a strategy $a := (a_t)_{t\geq 0}$ instead, output evolves, from his perspective, according to

$$d\xi_t = a_t dt + dY_t = (a_t + p_t)dt + \sigma dZ_t^Y, t \geq 0,$$

where we have used that the process Y admits, from the agent's standpoint, the following representation

$$dY_t = p_t dt + \sigma_\xi dZ_t^Y,\ t \geq 0$$

(Theorem 7.12 in Liptser and Shiryaev). Therefore, from the agent's perspective, public beliefs can be written as

$$dp_t^* = [-\kappa(p_t^* - \eta) + \beta_t(a_t - a_t^*) + \beta_t(p_t - p_t^*)]dt + \beta_t \sigma_\xi dZ_t^Y,\ t \geq 0,$$

where $\beta_t := \gamma_t/\sigma_\xi^2$, $t \geq 0$. Since the agent's private beliefs are governed by

$$dp_t = -\kappa(p_t - \eta)dt + \beta_t \sigma_\xi dZ_t^Y,\ t > 0,\ p_0 = p^o,$$

we can conclude that $\Delta := p^* - p$ satisfies, from the agent's perspective,

$$d\Delta_t = [-(\kappa + \beta_t)\Delta_t + \beta_t(a_t - a_t^*)]dt,\ t > 0,\ \Delta_0 = \Delta^o := p^{o*} - p^o.$$

□

Proof of Proposition 2.2: By Theorem 9.1 in Liptser and Shiryaev (1977) the processes $\pi_t := \mathbb{P}(\theta_t = h|\mathcal{F}_t^Y)$ and $\pi_t^* = \mathbb{P}(\theta_t = h|\mathcal{F}_t^\xi)$, $t \geq 0$ evolve according to

$$\begin{aligned} d\pi_t &= (\lambda_1(1-\pi_t) - \lambda_0\pi_t)dt + \frac{\delta\pi_t(1-\pi_t)}{\sigma_\xi}\left(\frac{dY_t - \pi_t dt}{\sigma_\xi}\right), \\ d\pi_t^* &= (\lambda_1(1-\pi_t^*) - \lambda_0\pi_t^*)dt + \frac{\delta\pi_t^*(1-\pi_t^*)}{\sigma_\xi}\left(\frac{d\xi_t - (a_t^* + \pi_t^*)dt}{\sigma_\xi}\right), \end{aligned}$$

where $\delta := h - \ell$ and

$$Z_t^Y := \frac{1}{\sigma_\xi}\left(Y_t - \int_0^t \pi_s ds\right) \quad \text{and} \quad Z_t^{\xi,*} := \frac{1}{\sigma_\xi}\left(\xi_t - \int_0^t (a_s^* + \pi_s^*)ds\right),\ t \geq 0,$$

are exogenous $\mathbb{F}^Y-$ and $\mathbb{F}^\xi-$Brownian motions from the agent's and market's standpoint, respectively.

It is well known starting from any point in $(0,1)$, both processes π and π^* never hit zero or one (Karlin and Taylor (1981)). We can then write $p_t := \log\left(\frac{\pi_t}{1-\pi_t}\right)$ and $p_t^* := \log\left(\frac{\pi_t^*}{1-\pi_t^*}\right)$. A direct application of Ito's rule yields that

$$\begin{aligned} dp_t &= \left(\frac{\lambda_1}{\pi_t} - \frac{\lambda_0}{1-\pi_t} - \frac{\delta^2(1-2\pi_t)}{2\sigma_\xi^2}\right)dt + \frac{\delta}{\sigma_\xi}dZ_t^Y \\ dp_t^* &= \left(\frac{\lambda_1}{\pi_t^*} - \frac{\lambda_0}{1-\pi_t^*} - \frac{\delta^2(1-2\pi_t^*)}{2\sigma_\xi^2}\right)dt + \frac{\delta}{\sigma_\xi}dZ_t^{\xi,*}. \end{aligned} \tag{2.46}$$

But from the agent's perspective,

$$dZ^{\xi,*} = dZ_t^Y + \frac{[a_t - a_t^* + \delta(\pi_t - \pi_t^*)]dt}{\sigma_\xi},\ t \geq 0,$$

which implies that

$$dp_t^* = dp_t + \left[\lambda_1\left(\frac{1}{\pi_t^*} - \frac{1}{\pi_t}\right) - \lambda_0\left(\frac{1}{1-\pi_t^*} - \frac{1}{1-\pi_t} + \frac{\delta}{\sigma_\xi^2}(a_t - a_t^*)\right)\right]dt.$$

Defining $\Delta := p_t^* - p_t$, and observing that $\pi = \frac{e^p}{e^p+1}$ and $\pi^* = \frac{e^{p^*}}{e^{p^*}+1}$, it is easy to see that

$$d\Delta_t = [\phi(p_t, \Delta_t) + \beta(a_t - a_t^*)]dt$$

where $\beta := \delta/\sigma_\xi^2$ and

$$\phi(p, \Delta) = \lambda_1\left[\frac{e^{p+\Delta}+1}{e^{p+\Delta}} - \frac{e^p+1}{e^p}\right] + \lambda_0[e^p - e^{p+\Delta}].$$

Finally, the dynamics of (2.46) as a function of p only are given by

$$dp_t = \left[\lambda_1 \frac{e^{p_t}+1}{e^{p_t}} - \lambda_0(1+e^{p_t}) - \frac{\delta^2}{2\sigma_\xi^2}(1-2\frac{e^{p_t}}{1+e^{p_t}})\right] dt + \frac{\delta}{\sigma_\xi^2}\left(dY_t - \frac{e^{p_t}}{1+e^{p_t}}dt\right), \ t \geq 0.$$

This concludes the proof.

□

Proof of Theorem 2.3: Consider the function

$$V(p,\Delta) = \alpha_0 + \alpha_1 p + \alpha_2 \Delta + \alpha_3 p\Delta + \alpha_4 p^2 + \alpha_5 \Delta^2 \tag{2.47}$$

with α_2 and α_3 as in the theorem (equations (2.30) and (2.56), respectively):

$$\begin{aligned}\alpha_2 &= \frac{\eta\kappa\alpha_3 + u_1}{r+\beta+\kappa+\frac{\beta^2\alpha_3}{\psi}} = \frac{2(\eta\kappa\alpha_3+u_1)}{r+\beta+\sqrt{(r+\beta+2\kappa)^2-\frac{8u_2\beta^2}{\psi}}}, \\ \alpha_3 &= \frac{\psi}{2\beta^2}\left[-(r+\beta+2\kappa)+\sqrt{(r+\beta+2\kappa)^2-\frac{8u_2\beta^2}{\psi}}\right] < 0.\end{aligned}$$

The idea of the proof is as follows: given a market's conjecture of the form $a^*(p^*) = \frac{\beta}{\psi}[\alpha_2 + \alpha_3 p^*]$, suitably chosen scalars $\alpha_0, \alpha_1, \alpha_4$ and α_5 yield that

1. V as above is an upper bound to the agent's problem;
2. The markovian strategy $\mathtt{a}(p,\Delta) := \beta V_\Delta(p,\Delta)$ induces an effort process a^L such that (a^L, a^*) is a feasible pair;
3. V is attained under the markovian control $\mathtt{a} : \mathbb{R}^2 \to \mathbb{R}$.

Observe that, by construction, $a^*(p+\Delta) = \frac{\beta}{\psi}V_\Delta(p+\Delta, 0)$, and thus the equilibrium condition (EC) (see (2.22)) is satisfied.

Before starting to prove the theorem, I make some preliminary observations that will be used repeatedly below:

(i) The private beliefs process $dp_t = -\kappa(p_t - \eta)dt + \beta\sigma_\xi dZ_t^Y$, $t \geq 0$ is such that

$$p_t \sim \mathcal{N}(e^{-\kappa t}p^o + (1 - e^{-\kappa t})\eta, (\beta\sigma_\xi)^2(1 - e^{-2\kappa t})/2\kappa)$$

from a time-zero perspective. Since $\kappa > 0$, $\mathbb{E}[p_t]$ and $\mathbb{E}[p_t^2]$ are uniformly bounded for all $t \geq 0$.

(ii) Denote by $\mathcal{A}^{a^*}(p^o, \Delta^o)$ the set of strategies $a := (a_t)_{t\geq 0}$ for the long-run agent such that (a, a^*) is a feasible pair when the the state variables p and Δ start at (p^o, Δ^o). Then, the integrability condition for p^* in Definition 2.2 (see (2.12)) is equivalent to

$$\mathbb{E}\left[\int_0^t |\Delta_s^{a,a^*}|^2 ds\right] < \infty,\ t \geq 0, \tag{2.48}$$

where Δ^{a,a^*} denotes the solution to (2.13) under the pair (a, a^*). This is because $\Delta := p^* - p$ and p satisfies the above integrability condition too.

(iii) Also,

$$\mathbb{E}\left[\int_0^t |\Delta_s^{a,a^*}|^2 ds\right] < \infty \Leftrightarrow \mathbb{E}\left[\int_0^t e^{-\alpha s}|\Delta_s^{a,a^*}|^2 ds\right] < \infty$$

for all $\alpha > 0$, which allows us to conclude that stochastic integrals against $e^{-\alpha s}\Delta_s^{a,a^*}$, $s \geq 0$, have all zero mean.

In order to show that V is an upper bound, we need to study first the asymptotic behavior of $e^{-rt}\mathbb{E}[V(p_t, \Delta_t^{a,a^*})]$ over the set $\{a \in \mathcal{A}^{a^*}(p^o, \Delta^o) \mid V^{a^*}(p^o, \Delta^o; a) > -\infty\}$, for all $(p^o, \Delta^o) \in \mathbb{R}^2$. This analysis is done in the next three Lemmas and concludes in Corollary 2.1.

Lemma 2.2. *Suppose that the market conjectures an effort strategy of the form* $a^*(p^*) = \frac{\beta}{\psi}[\alpha_2 + \alpha_3 p^*]$, $\alpha_2, \alpha_3 \in \mathbb{R}$. *Then,* $\mathcal{A}^{a^*}(p^o, \Delta^o) \neq \emptyset$ *and there exists* $a \in$

$\mathcal{A}^{a^*}(p^o, \Delta^o)$ *such that* $V^{a^*}(p^o, \Delta^o; a)$ *is finite for all* $(p^o, \Delta^o) \in \mathbb{R}^2$*. Moreover,*

$$V^{a^*}(p^o, \Delta^o; a) > -\infty \;\; \Leftrightarrow \;\; \mathbb{E}\left[\int_0^\infty e^{-rt}\left|u(p_t + \Delta_t^{a,a^*}) - \frac{\psi}{2}a_t^2\right|dt\right] < \infty \quad (2.49)$$

for all $a \in \mathcal{A}^{a^*}(p^o, \Delta^o)$, $(p^o, \Delta^o) \in \mathbb{R}^2$.

Proof: Suppose that the market conjectures the agent is following a manipulation strategy of the form $a^*(p^*) = \frac{\beta}{\psi}[\alpha_2 + \alpha_3 p^*]$. Observe that the pair $(\tilde{a}, a^*)$ with

$$\tilde{a}_t := \frac{\beta}{\psi}[\alpha_2 + \alpha_3 p_t^*] + \frac{\kappa + \beta}{\beta}\Delta_t, \; t \geq 0$$

is a feasible strategy-conjecture pair. In fact, under this pair the discrepancy of beliefs remains constant at its initial value Δ^o, so condition (2.48) is satisfied. Thus, $\mathcal{A}^{a^*}(p^o, \Delta^o)$ is non-empty.

Recall that given the conjecture a^*, the agent solves

$$\begin{aligned}
\max_{a \in \mathcal{A}^{a^*}(p^o, \Delta^o)} \quad & \mathbb{E}\left[\int_0^\infty e^{-rt}(u(p_t^*) - \frac{\psi}{2}a^2)dt\right] \\
s.t. \quad & p_t^* = p_t + \Delta_t, \; t > 0 \\
& d\Delta_t = [-(\kappa + \beta)\Delta_t + \beta(a_t - \frac{\beta}{\psi}(\alpha_2 + \alpha_3 p_t^*))]dt, \; t > 0, \; \Delta_0 = \Delta^o. \\
& dp_t = -\kappa(p_t - \eta)dt + \beta\sigma_\xi dZ_t^Y, \; t > 0, p_o = p^o.
\end{aligned}$$

and denote by $V^{a^*}(p^o, \Delta^o)$ the value of this problem. Observe that $V^{a^*}(p^o, \Delta^o)$ is finite. In fact, by following $\tilde{a}$, the cost of signal manipulation at any time t is of order $O(p_t^2)$, $t \geq 0$. Since from a time-zero perspective $\mathbb{E}[p_t]$ and $\mathbb{E}[p_t^2]$, $t \geq 0$ are uniformly bounded, the agent's total utility under $\tilde{a}$ is finite.

Finally, because $u(\cdot)$ is uniformly bounded by above, we have that for $M > 0$ large enough

$$\begin{aligned}
\mathbb{E}\left[\int_0^\infty e^{-rt}\Big|u(p_t+\Delta_t^{a,a^*}) - \frac{\psi}{2}a_t^2\Big|dt\right] &\leq \frac{M}{r} + \mathbb{E}\left[\int_0^\infty e^{-rt}\Big|u(p_t+\Delta_t^{a,a^*}) - M - \frac{\psi}{2}a_t^2\Big|\right] \\
&\leq \frac{2M}{r} - \mathbb{E}\left[\int_0^\infty e^{-rt}\left(u(p_t+\Delta_t^{a,a^*}) - \frac{\psi}{2}a_t^2\right)dt\right] \\
&= \frac{2M}{r} - V^{a^*}(p^o,\Delta^o;a),
\end{aligned}$$

where we used that $u(\cdot) - M \leq 0$ for M large. As a result,

$$V^{a^*}(p^o,\Delta^o;a) > -\infty \quad \Leftrightarrow \quad \mathbb{E}\left[\int_0^\infty e^{-rt}|u(p_t+\Delta_t^{a,a^*}) - \frac{\psi}{2}a_t|dt\right] < \infty$$

for all $a \in \mathcal{A}^{a^*}(p^o,\Delta^o)$. This concludes the proof.

□

Lemma 2.3. *Let $a \in \mathcal{A}^{a^*}(p^o,\Delta^o)$ be a strategy under which the agent attains finite utility. Then, there are positive constants C_1 and $C_2(a)$ such that*

$$|\mathbb{E}[\Delta_t^{a,a^*}]| < C_1[1 + e^{-(\frac{\beta^2\alpha_3}{\psi}+\beta+\kappa)t}] + C_2(a)[e^{rt}(1+e^{-2(\frac{\beta^2\alpha_3}{\psi}+\beta+\kappa)t})]^{1/2}. \tag{2.50}$$

As a result, $\lim_{t\to\infty} e^{-rt}\mathbb{E}[\Delta_t^{a,a^*}] = 0.$

Proof: Take any such strategy. Under the market's conjecture a^* we can write

$$\Delta_t^{a,a^*} = \Delta^o e^{-(\beta+\kappa+\frac{\beta^2\alpha_3}{\psi})t} + \int_0^t e^{(\beta+\kappa+\frac{\beta^2\alpha_3}{\psi})(s-t)}[\beta a_s - \beta^2(\alpha_2+\alpha_3 p_s)]ds, t \geq 0.$$

Using this and the fact that $\mathbb{E}[p_t]$ is uniformly bounded, we find C_1 s.t.

$$|\mathbb{E}[\Delta_t^{a,a^*}]| \leq C_1[1 + e^{-(\frac{\beta^2\alpha_3}{\psi}+\beta+\kappa)t}] + \int_0^t e^{(\beta+\kappa+\frac{\beta^2\alpha_3}{\psi})(s-t)}\beta\mathbb{E}[|a_s|]ds.$$

Now,

$$I := \int_0^t e^{(\beta+\kappa+\frac{\beta^2\alpha_3}{\psi})(s-t)}\mathbb{E}[|a_s|]ds < \left(e^{rt}\int_0^t e^{2(\beta+\kappa+\frac{\beta^2\alpha_3}{\psi})(s-t)}ds\right)^{1/2}\left(e^{-rt}\int_0^t \mathbb{E}[a_s^2]ds\right)^{1/2}$$

where in the last inequality we used Cauchy-Schwarz's and Jensen's inequalities. But $e^{-rt}\int_0^t \mathbb{E}[a_s^2]ds < C(a) := \int_0^\infty e^{-rs}\mathbb{E}[a_s^2]ds$, which is finite since flow payoffs are bounded by above and a attains finite utility. Therefore

$$I \leq C_2(a)[e^{rt}(1+e^{-2(\frac{\beta^2\alpha_3}{\psi}+\beta+\kappa)t})]^{1/2},$$

for some positive constant $C_2(a)$. This proves (2.50).

Finally, from α_3's definition it is easy to see that $\beta+\kappa+\frac{\beta^2\alpha_3}{\psi}+\frac{r}{2} > 0$. This inequality and the bound just proven above, yield $\lim_{t\to\infty} e^{-rt}\mathbb{E}[\Delta_t^{a,a^*}] = 0$.

□

Lemma 2.4. *Let $a \in \mathcal{A}^{a^*}(p^o, \Delta^o)$ be a strategy under which the agent attains finite utility. Then, $\lim_{t\to\infty} e^{-rt}\mathbb{E}[p_t\Delta_t^{a,a^*}] = 0$. As a consequence, $\liminf_{t\to\infty} e^{-rt}\mathbb{E}[(\Delta_t^{a,a^*})^2] = 0$.*

Proof: Applying Ito's rule to $e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)t}p_s\Delta_t^{a,a^*}$ we obtain the expression

$$\begin{aligned} p_t\Delta_t^{a,a^*} &= e^{-(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)t}p^o\Delta^o + \underbrace{\int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}\Delta_s^{a,a^*}[\kappa\eta ds + \beta\sigma_\xi dZ_s^Y]}_{I_t:=} \\ &+ \beta\underbrace{\int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}p_s a_s ds}_{J_t:=} - \frac{\beta^2}{\psi}\underbrace{\int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}[\alpha_2 p_s + \alpha_3 p_s^2]dp_s}_{K_t:=} \end{aligned}$$

It is clear that the first term in the right-hand side of the previous expression goes to zero when discounted at rate r. Regarding the last term, observe that since $p_t|\mathcal{F}_t^Y \sim \mathcal{N}(e^{-\kappa t}p^o + (1-e^{-\kappa t})\eta, (\beta\sigma_\xi)^2(1-e^{-2\kappa t})/\kappa)$, $\mathbb{E}[p_t]$ and $\mathbb{E}[p_t^2]$ are uniformly bounded for all $t > 0$. As a result $\lim_{t\to\infty} e^{-rt}\mathbb{E}[K_t] = 0$, where we used again that

$\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa+r>0.$

Now, by (2.48) the stochastic integral in I_t has zero mean, for all $t \geq 0$. Hence,

$$\mathbb{E}[I_t] = \eta\kappa \int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}\mathbb{E}[\Delta_s^{a,a^*}]ds,\ t \geq 0.$$

Using the bound (2.50) derived in the previous lemma, we obtain that for some positive constant $C_3(a)$

$$\begin{aligned}
\frac{|\mathbb{E}[I_t]|}{C_3(a)} &\leq \int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}[1+e^{-(\frac{\beta^2\alpha_3}{\psi}+\beta+\kappa)s}]ds \\
&\quad + \int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}[e^{rs}(1+e^{-2(\frac{\beta^2\alpha_3}{\psi}+\beta+\kappa)s})]^{1/2}ds \\
&\leq \int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}ds + e^{-(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)t}\frac{e^{\kappa t}-1}{\kappa} \\
&\quad + e^{-(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)t}\int_0^t [e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa+r/2)s}+e^{(\kappa+r/2)s}]ds.
\end{aligned}$$

Observing that $\frac{\beta^2\alpha_3}{\psi}+\beta+\nu+r>0$ and $\frac{\beta^2\alpha_3}{\psi}+\beta+\nu+r>\beta^2\alpha_3+\beta+\nu+r/2>0$ for $\nu=\kappa, 2\kappa$, we conclude that $\lim_{t\to\infty} e^{-rt}\mathbb{E}[I_t]=0$.

It remains to show that $\lim_{t\to\infty} e^{-rt}\mathbb{E}[J_t] = 0$. Applying the Cauchy-Schwartz inequality twice

$$e^{-rt}|\mathbb{E}[J_t]| \leq \left(e^{-rt}\int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}\mathbb{E}[a_s^2]ds\right)^{1/2}\left(e^{-rt}\int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}\mathbb{E}[p_s^2]ds\right)^{1/2}$$

Since $(\mathbb{E}[p_t^2])_{t\geq 0}$ is uniformly bounded and $\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa+r>0$ the last term in the right-hand side of the previous expression goes to zero as $t\to\infty$. Consequently, it suffices to show that

$$\lim_{t\to\infty} e^{-rt}\int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}\mathbb{E}[a_s^2]ds < \infty.$$

But observe that

$$e^{-rt}\int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa)(s-t)}\mathbb{E}[a_s^2]ds = \int_0^t e^{(\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa+r)(s-t)}e^{-rs}\mathbb{E}[a_s^2]ds$$
$$< \int_0^\infty e^{-rs}\mathbb{E}[a_s^2]ds < \infty$$

where we used that $\frac{\beta^2\alpha_3}{\psi}+\beta+2\kappa > 0$ and the fact that $J^{a^*}(p^o,\Delta^o;a) > -\infty$ when a is feasible. This argument shows that $\lim_{t\to\infty} e^{-rt}\mathbb{E}[p_t\Delta^{a,a^*}] = 0$ for any feasible strategy that yields finite utility.

Finally, since for any $a \in \mathcal{A}(p^o,\Delta^o)$, $V^{a^*}(p^o,\Delta^o;a) > -\infty$, and flow payoffs are bounded by above, we must have that $\mathbb{E}[\int_0^\infty e^{-rt}u(p_t+\Delta_t^{a,a^*})dt]$ must be finite. As a consequence, $\limsup_{t\to\infty} e^{-rt}\mathbb{E}[u(p_t+\Delta_t^{a,a^*})] \geq 0$, because otherwise the agent's expected discounted would be $-\infty$. Using that $\lim_{t\to\infty} e^{-rt}\mathbb{E}[p_t+\Delta_t^{a,a^*}]$ exists and it is equal to zero, we obtain that

$$\limsup_{t\to\infty} e^{-rt}\mathbb{E}[-u_2(p_t+\Delta_t^{a,a^*})^2] \geq 0 \Rightarrow \liminf_{t\to\infty} e^{-rt}\mathbb{E}[(p_t+\Delta_t^{a,a^*})^2] = 0$$

But

$$0 = \liminf_{t\to\infty} e^{-rt}\mathbb{E}[(p_t+\Delta_t^{a,a^*})^2] \geq \underbrace{\liminf_{t\to\infty} e^{-rt}\mathbb{E}[(p_t)^2]}_{=0} + \underbrace{\liminf_{t\to\infty} e^{-rt}2\mathbb{E}[p_t\Delta_t^{a,a^*}]}_{=0}$$
$$+\liminf_{t\to\infty} e^{-rt}\mathbb{E}[(\Delta_t^{a,a^*})^2]$$

This concludes the proof.

□

With these 3 lemmas we obtain the following

Corollary 2.1. *Suppose that the function V in (2.47) is such that $\alpha_5 < 0$. Then, for any $a \in \{a \in \mathcal{A}^{a^*}(p^o, \Delta^o) \mid J^{a^*}(p^o, \Delta^o; a) > -\infty\}$,*

$$\limsup_{t\to\infty} e^{-rt}\mathbb{E}[V(p_t, \Delta_t^{a,a^*})] = 0. \tag{2.51}$$

Proof: Using Lemmas 2.3 and 2.4, we have that, for all $a \in \mathcal{A}^{a^*}(p^o, \Delta^o)$ such that $J^{a^*}(p^o, \Delta^o; a) > -\infty$,

$$\lim_{t\to\infty} e^{-rt}\mathbb{E}[\chi_t] = 0, \ \chi = p, \Delta^{a,a^*}, p\Delta^{a,a^*}, p^2.$$

Thus,

$$\limsup_{t\to\infty} e^{-rt}\mathbb{E}[V(p_t, \Delta_t^{a,a^*})] = \limsup_{t\to\infty} e^{-rt}\alpha_5\mathbb{E}[(\Delta_t^{a,a^*})^2] = \alpha_5 \liminf_{t\to\infty} e^{-rt}\mathbb{E}[(\Delta_t^{a,a^*})^2] = 0,$$

where in the last two equalities we used that $\alpha_5 < 0$ and Lemma 2.4, respectively. This concludes the proof.

□

Now we prove the Theorem:

1. For suitably chosen constants, V is an upper bound to the agent's utility: Recall that $\alpha_2 > 0$ and $\alpha_3 < 0$ are defined by

$$\alpha_2 = \frac{2(\eta\alpha_3 + u_1)}{r + \beta + \sqrt{(r+\beta+2\kappa)^2 - \frac{8u_2\beta^2}{\psi}}}.$$

$$\alpha_3 = \frac{\psi}{2\beta^2}\left[-(r+\beta+2\kappa) + \sqrt{(r+\beta+2\kappa)^2 - \frac{8u_2\beta^2}{\psi}}\right]$$

Now, it is easy to verify that given $a^*(p^*) = \frac{\beta}{\psi}(\alpha_2 + \alpha_3 p^*)$, V as in (2.47) satisfies the HJB equation

$$\begin{aligned} rV(p,\Delta) &= \max_a \Big\{ u_0 + u_1(p+\Delta) - u_2(p+\Delta)^2 - \frac{\psi}{2}a^2 \\ &\quad + -\kappa[p-\eta]V_p(p,\Delta) + \frac{1}{2}\beta^2\sigma_\xi^2 V_{pp}(p,\Delta) \\ &\quad + \left[-\Delta(\beta+\kappa) + \beta a - \frac{\beta^2}{\psi}(\alpha_2 + \alpha_3(p+\Delta))\right] V_\Delta(p,\Delta) \Big\} \end{aligned} \tag{2.52}$$

when $\alpha_0, \alpha_1, \alpha_4$ and α_5 solve the following system of equations:

$$\begin{aligned} (\alpha_0) \ : \ 0 &= r\alpha_0 - u_0 - \eta\kappa\alpha_1 + \frac{1}{2\psi}\beta^2\alpha_2^2 - \beta^2\sigma_\xi^2\alpha_4 \\ (\alpha_1) \ : \ 0 &= r\alpha_1 - u_1 + \kappa\alpha_1 + \frac{\beta^2\alpha_3}{\psi}\alpha_2 - 2\eta\kappa\alpha_4 \\ (\alpha_4) \ : \ 0 &= r\alpha_4 + \frac{1}{2\psi}\beta^2\alpha_3^2 + 2\kappa\alpha_4 + u_2 \\ (\alpha_5) \ : \ 0 &= \left(r + 2\left[\kappa + \beta + \frac{\beta^2\alpha_3}{\psi}\right]\right)\alpha_5 - \frac{2\beta^2}{\psi}\alpha_5^2 + u_2 \end{aligned}$$

In fact, a quadratic guess for V reduces the HJB equation to the above system of equations plus the conditions

$$\begin{aligned} (\alpha_2) \ : \ 0 &= \left(r + \kappa + \beta + \frac{\beta^2\alpha_3}{\psi}\right)\alpha_2 - \eta\kappa\alpha_3 - u_1 \\ (\alpha_3) \ : \ 0 &= (r + \beta + 2\kappa)\alpha_3 + \frac{\beta^2}{\psi}\alpha_3^2 + 2u_2. \end{aligned}$$

The latter are trivially satisfied by our choice of α_2 and α_3. Moreover, given α_3 the equations (α_0), (α_1),(α_2) and (α_4) have a unique solution. For equation (α_5), we choose the its unique negative solution, that is

$$\alpha_5 = \frac{r + 2(\beta+\kappa) + \frac{2\alpha_3\beta^2}{\psi} - \sqrt{\left(r + 2(\beta+\kappa) + \frac{\alpha_3\beta^2}{\psi}\right)^2 + \frac{8\beta^2 u_2}{\psi}}}{\frac{4\beta^2}{\psi}}, \tag{2.53}$$

as the value function must be bounded by above over the whole plane.

Now we show that this choice of V is in fact an upper bound to the agent's utility. Assume that $a \in \mathcal{A}^{a^*}(p^o, \Delta^o)$ is such that $V^{a^*}(p^o, \Delta^o; a) > -\infty$. Consider the process

$$G_t^{a,a^*} = \int_0^t e^{-rs}\left(u(p_s + \Delta_s^{a,a^*}) - \frac{\psi}{2}a_s^2\right) ds + e^{-rt}V(p_t, \Delta_t^{a,a^*}),\ t \geq 0.$$

Observe that

$$\begin{aligned}\frac{dG_t^{a,a^*}}{e^{-rt}} &= \Bigg\{\left(u(p_t + \Delta_t^{a,a^*}) - \frac{\psi}{2}a_t^2 - rV(p_t, \Delta_t^{a,a^*}) + \mathcal{L}^{a,a^*}V(p_t, \Delta_t^{a,a^*})\right) dt \\ &\quad + \frac{1}{2}\beta^2\sigma_\xi^2 V_{pp}(p_t, \Delta_t^{a,a^*}) dZ_t^Y\Bigg\}\end{aligned}$$

where

$$\begin{aligned}\mathcal{L}^{a,a^*}V(p_t, \Delta_t^{a,a^*}) &= -\kappa[p_t - \eta]V_p(p_t, \Delta_t^{a,a^*}) \\ &\quad + \left[-\Delta_t^{a,a^*}(\beta + \kappa) + \beta a_t - \frac{\beta^2}{\psi}(\alpha_2 + \alpha_3(p_t + \Delta_t^{a,a^*}))\right] V_\Delta(p_t, \Delta_t^{a,a^*})\end{aligned}$$

Because,

$$\mathbb{E}\left[\int_0^t e^{-rs}\beta^2\sigma_\xi^2 V_{pp}(p_t, \Delta_t^{a,a^*}) dZ_t^Y\right] = 2\mathbb{E}\left[\int_0^t e^{-rs}\beta^2\sigma_\xi^2 \alpha_5 dZ_t^Y\right] = 0,\ t \geq 0$$

and V satisfies the HJB equation (2.52), we conclude that the drift of G is non-positive, i.e., it is a supermartingale. In particular, $\mathbb{E}[G_t^{a,a^*}] \leq G_0^{a,a^*}$ for all $t \geq 0$ and therefore

$$e^{-rt}\mathbb{E}[V(p_t, \Delta_t^{a,a^*})] \leq V(p^o, \Delta^o) - \mathbb{E}\left[\int_0^t e^{-rs}\left(u(p_s + \Delta_s^{a,a^*}) - \frac{\psi}{2}a_s^2\right) ds\right],\ t \geq 0.$$

Now,

$$\Big|\int_0^t e^{-rs}(u(p_s+\Delta_s^{a,a^*})-\frac{\psi}{2}a_s^2)ds\Big| \leq \int_0^\infty e^{-rs}\Big|u(p_s+\Delta_s^{a,a^*})-\frac{\psi}{2}a_s^2\Big|ds$$

and the latter is integrable for all $a \in \mathcal{A}^{a^*}$ such that $V^{a^*}(p^o,\Delta^o;a) > -\infty$ (Lemma 2.2). Using Corollary 2.1 and the dominated convergence theorem we conclude that

$$0 \leq V(p^o,\Delta^o) - \mathbb{E}\left[\int_0^\infty e^{-rt}\left(u(p_t+\Delta_t^{a,a^*})-\frac{\psi}{2}a_t^2\right)ds\right] = V(p^o,\Delta^o) - V^{a^*}(p^o,\Delta^o;a),$$

that is, V it is an upper bound to the agent's value function.

□

2. The markovian strategy $\mathtt{a}(p,\Delta) := \frac{\beta}{\psi}V_\Delta(p,\Delta)$ induces well defined dynamics of Δ and the process $a_t^L := \mathtt{a}(p_t,\Delta_t^{\mathtt{a}})$, $t \geq 0$ is a feasible strategy:

It is easy to see that the maximum in HJB (2.52) is achieved under $\mathtt{a}(p,\Delta) := \frac{\beta}{\psi}(\alpha_2+\alpha_3 p+2\alpha_5\Delta)$. Then, for any $\Delta^o \in \mathbb{R}$, the dynamics of belief-asymmetry (eq. (2.40)) have as a solution the deterministic function

$$\Delta_t^{\mathtt{a}} \;=\; \Delta^o e^{\rho t},\ t \geq 0. \tag{2.54}$$

with $\rho := \frac{(2\alpha_5-\alpha_3)\beta^2}{\psi} - \kappa - \beta$ as a solution. As a consequence, $(\mathtt{a},\mathtt{a}(\cdot,0))$ satisfies the integrability condition (2.48), so the pair is feasible. Using α_5's definition (2.57) we see that

$$\rho = \frac{r - \sqrt{\left(r+2(\beta+\kappa)+\frac{2\alpha_3\beta^2}{\psi}\right)^2 + \frac{8\beta^2 u_2}{\psi}}}{2}, \tag{2.55}$$

yielding that both $\rho - r$ and $2\rho - r$ are strictly less than zero. It is easy to conclude from here that the agent attains finite utility under a^L.

□

3. V is attained under the markovian control $\mathtt{a} : \mathbb{R}^2 \to \mathbb{R}$. The fact that both ρ and $2\rho - r$ are strictly negative implies that

$$\lim_{t \to 0} e^{-rt} \mathbb{E}[\Delta_t^{\mathtt{a}}] = \lim_{t \to 0} e^{-rt} \mathbb{E}[(\Delta_t^{\mathtt{a}})^2] = 0.$$

Lemmas 2.3 and 2.4 allow us to conclude that the following transversality condition holds:

$$\lim_{t \to \infty} e^{-rt} \mathbb{E}[V(p_t, \Delta_t^{\mathtt{a}})] = 0 \quad (TVC)$$

Finally, observe first that $\mathtt{a}(p, \Delta)$ attains the supremum in the HJB equation. Hence, we have equality in (2.54) for all $t \geq 0$ under the pair (a^L, a^*). Using the dominated convergence theorem and (TVC) we can take limits in (2.54) and conclude that $V(p^o, \Delta^o) = V^{a^*}(p, \Delta; a^L)$. Therefore, a^L is an optimal control.

□

Proof of Proposition 2.4: Recall that off the equilibrium path

$$\Delta_t = \Delta^o e^{\rho t}, \; t \geq 0,$$

where $\rho := \frac{(2\alpha_5 - \alpha_3)\beta^2}{\psi} - \beta - \kappa$. From α_5's definition we can see that

$$\rho = \frac{r - \sqrt{\left(r + 2(\beta + \kappa) + \frac{2\alpha_3 \beta^2}{\psi}\right)^2 + \frac{8\beta^2 u_2}{\psi}}}{2},$$

so $\rho < 0$ if and only if

$$\begin{aligned} 0 &\leq 4(\beta+\kappa)^2 + \frac{4\beta^4\alpha_3^2}{\psi^2} + 4r(\beta+\kappa) + \frac{4r\beta^2\alpha_3}{\psi} + \frac{8\beta^2(\beta+\kappa)\alpha_3}{\psi} + \frac{8\beta^2 u_2}{\psi} \\ \Leftrightarrow 0 &\underbrace{\leq}_{(*)} (\beta+\kappa)^2 + r(\beta+\kappa) + \frac{\beta^3\alpha_3}{\psi} + \frac{\beta^2}{\psi}\underbrace{\left[\beta^2\alpha_3^2 + (r+\beta+2\kappa)\alpha_3 + 2u_2\right]}_{=0,\text{ by }\alpha_3'\text{s definition}} \end{aligned}$$

But, using α_3's definition

$$\frac{\beta^3\alpha_3}{\psi} = \beta\frac{-(r+\beta+2\kappa) + \sqrt{(r+\beta+2\kappa)^2 - \frac{8\beta^2 u_2}{\psi}}}{2}$$

from where we can see that (*) is true. Moreover $2\alpha_5 - \alpha_3 > 0$. This is equivalent to

$$\begin{aligned} [r+2(\kappa+\beta)]^2 &> \left(r + 2(\beta+\kappa) + \frac{2\alpha_3\beta^2}{\psi}\right)^2 + \frac{8\beta^2 u_2}{\psi} \\ \Leftrightarrow 0 &> 4\beta^4\alpha_3^2 + 4r\beta^2\alpha_3 + 8\beta^2(\beta+\kappa)\alpha_3 + 8\beta^2 u_2 \\ \Leftrightarrow 0 &> \beta^3\alpha_3 + \beta^2\underbrace{\left[\beta^2\alpha_3^2 + (r+\beta+2\kappa)\alpha_3 + 2u_2\right]}_{=0,\text{ by }\alpha_3'\text{s definition}} \end{aligned}$$

which is true.

□

Proof of Proposition 2.5: It is easy to conclude that the system (p, Δ) is (i) stabilizable (the belief-asymmetry process is controllable and private beliefs decay to zero) and that (ii) the system is detectable (in the (p, p^*) coordinate system, the "unobserved" component p (i.e. the state variable that does not contribute to the flow payoff) decays to zero). Consequently, the solution of this linear-quadratic regulator problem exists, it is unique and the value function is quadratic (Theorem 12.3. in Wonham (1985)).

Guess a solution of the form:

$$V(p,\Delta)=\alpha_3 p\Delta+\alpha_4 p^2+\alpha_5\Delta^2.$$

Then, $\alpha_i=\alpha_i(\hat{\alpha}_3)$, $i=3,4,5$ are given by

$$\alpha_3(\hat{\alpha}_3) = \frac{-2u_2-\frac{2\beta^2}{\psi}\alpha_5(\hat{\alpha}_3)\hat{\alpha}_3}{r+\beta+2\kappa+\frac{\beta^2}{\psi}\hat{\alpha}_3-\frac{2\beta^2}{\psi}\alpha_5(\hat{\alpha}_3)} \tag{2.56}$$

$$\alpha_4(\hat{\alpha}_3) = \frac{-u_2+\frac{\beta^2\alpha_3(\hat{\alpha}_3)}{2\psi}(\alpha_3(\hat{\alpha}_3)-2\hat{\alpha}_3)}{r+2\kappa} \tag{2.57}$$

$$\alpha_5(\hat{\alpha}_3) = \frac{r+2\left(\beta+\kappa+\frac{\beta^2}{\psi}\hat{\alpha}_3\right)\pm\sqrt{\left(r+2\left(\beta+\kappa+\frac{\beta^2}{\psi}\hat{\alpha}_3\right)\right)^2+\frac{8\beta^2}{\psi}u_2}}{\frac{4\beta^2}{\psi}}. \tag{2.58}$$

Choose the negative root of $\alpha_5(\hat{\alpha_3})$. Then, $\alpha_3(\hat{\alpha}_3)$ and $\alpha_4(\hat{\alpha}_3)$ are uniquely defined. Moreover, it is easy to show that

$$\rho(\hat{\alpha}_3):=\beta+\kappa+\frac{\beta^2}{\psi}\hat{\alpha}_3-\frac{2\beta^2}{\psi}\alpha_5(\hat{\alpha}_3)=\frac{1}{2}\left(\sqrt{\left(r+2\left(\beta+\kappa+\frac{\beta^2}{\psi}\hat{\alpha}_3\right)\right)^2+\frac{8\beta^2}{\psi}u_2}-r\right),$$

yielding that $r+\rho(\hat{\alpha}_3)>0$ and $r+2\rho(\hat{\alpha}_3)>0$. In particular, the denominator of $\alpha_3(\hat{\alpha}_3)$, $r+\kappa+\rho(\hat{\alpha}_3)>0$, so $\alpha_3(\hat{\alpha}_3)<0$.

Under the control $a(p,\Delta)=\alpha_3(\hat{\alpha}_3)p+\alpha_5(\hat{\alpha}_3)\Delta$, the belief-asymmetry process takes the form

$$\Delta_t=\Delta_0 e^{-\rho(\hat{\alpha}_3)t}+\frac{\beta^2(\alpha_3(\hat{\alpha_3})-\hat{\alpha}_3)}{\psi}\int_0^t e^{-\rho(\hat{\alpha}_3)(t-s)}p_s ds$$

with $p_s=p_0e^{-\kappa s}$, $s\geq 0$. Consequently, $e^{-rt}\Delta_t^2$ converges to zero as $t\to\infty$. Since $e^{-rt}p_t^2$ also converges to zero as $t\to\infty$, the conjectured value function satisfies the transversality conditions. It follows that V as above must be the long-run agent's value function.

Now, suppose that $\alpha_3(\hat{\alpha}_3) < \hat{\alpha}_3$. Then, using that the denominator of $\alpha_3(\hat{\alpha}_3)$ is strictly positive, we conclude that this inequality is true if and only if

$$\frac{\beta^2}{\psi}(\hat{\alpha}_3)^2 + (r + \beta + 2\kappa)\hat{\alpha}_3 + 2u_2 > 0,$$

which holds for all values of $\hat{\alpha}_3$ if the curvature condition is violated. This concludes the proof.

□

Proof of Proposition 2.7: Suppose that an equilibrium $\mathsf{a}(\cdot,\cdot)$ exists and that the associated value function V^{a} is of class C^2. Then, V^{a} satisfies the HJB equation

$$\begin{aligned} rV(p,\Delta) \quad = \quad & \max_{a\in A}\Big\{u(p,p+\Delta) - g(a) + [a-\kappa p]V_p(p,\Delta) + \frac{1}{2}\beta^2\sigma_\xi^2 V_{pp}(p,\Delta) \\ & +[-\Delta(\beta+\kappa) + \mathsf{a}(p+\Delta,0) - a]V_\Delta(p,\Delta)\} \end{aligned}$$

along with the equilibrium condition

$$g'(\mathsf{a}(p,0)) = V_p(p,0) - V_\Delta(p,0).$$

whenever incentives are interior at $(p,0)$. If, moreover $\mathsf{a}(\cdot,0)$ is locally twice differentiable around a point $p \in \mathbb{R}$, we can applying the envelope theorem to the above HJB equation with respect to Δ and p and evaluate at $\Delta = 0$. This yields the system of ordinary differential equations (2.43)-(2.43) in the proposition. This concludes the proof.

□

Proof of Proposition 2.8: It is easy to see that $V(\theta) = \alpha_0^o + \alpha_1^o\theta + \alpha_2^o\theta^2$ solves the HJB equation

$$rV(\theta) = \max_{a\in\mathbb{R}} \left\{ -k_2(\theta - \overline{\theta})^2 - \frac{\psi}{2}a^2 + [a - \kappa\theta]V_\theta(\theta) + \frac{1}{2}\sigma_\theta V_{\theta\theta}(\theta) \right\}$$

when

$$\begin{aligned}
\alpha_0^o &= -k_2\overline{\theta}^2 + \frac{\alpha_1^o}{2\psi} + \alpha_2\sigma_\theta^2 \\
\alpha_1^o &= -\frac{2k_2\overline{\theta}}{\frac{2\alpha_2^c}{\psi} - (r+\kappa)} \\
\alpha_2^o &= \frac{\psi}{2}\left[(r+\kappa) - \sqrt{(r+\kappa)^2 + \frac{8k_2}{\psi}}\right]
\end{aligned}$$

In order to show that V is the agent's value function and that $a^o(\theta) = \frac{1}{\psi}V_\theta(\theta)$ an optimal monetary, two things remain to be checked:

1. Any policy a such that

$$\mathbb{E}\left[\int_0^\infty e^{-rt}\left| -k_2(\theta_t^a - \overline{\theta})^2 - \frac{\psi}{2}a_t^2\right| dt\right] < \infty$$

   must also satisfy $\limsup_{t\to\infty} e^{-rt}\mathbb{E}[V(\theta_t^{a^o})] \geq 0$, where θ^{a^o} is the dynamic of θ under the policy a^o.

2. Second, that along the conjectured optimal strategy a^o, $\liminf_{t\to\infty} e^{-rt}\mathbb{E}[V(\theta_t^{a^o})] \leq 0$.

Part 1. implies that V is an upper bound to the agent's utility. Part 2. yields that V is attainable under the Markov control a^o (refer to the proof of Theorem 2.3 for a proof of these statements). Conditions 1. and 2. are proven to be true in the non-commitment case, which corresponds to a slightly more general environment than the one analyzed here. We refer the reader to the proof of Theorem 2.4 below in which it can be seen that the exact same steps performed to show that 1. and 2. indeed hold in that setting can be replicated in the full commitment case.

□

Proof of Theorem 2.4: Suppose the the market conjectures a manipulation strategy of the form $a^*(p^*) = \frac{1}{\psi}[\alpha_1 - \alpha_2 + 2\alpha_3 p^*]$ where

$$\begin{aligned}
\alpha_1 &= -\frac{2k_2\overline{\theta}}{\frac{2\alpha_3}{\psi} - (r+\kappa)} \\
\alpha_2 &= \frac{k_1}{2\alpha_3\psi - (r+\beta+\kappa)} \\
\alpha_3 &= \frac{\psi}{2}\left[(r+\kappa) - \sqrt{(r+\kappa)^2 + \frac{8k_2}{\psi}}\right],
\end{aligned}$$

and observe that $\alpha_1 = \alpha_1^o$ and $\alpha_3 = \alpha_2^o$ the parameters of the full observability case. It is easy to the that the quadratic form

$$V = \alpha_0 + \alpha_1 p + \alpha_2 \Delta + \alpha_3 p^2$$

with $\alpha_0 = -k_2\overline{\theta}^2 + \frac{1}{2\psi}(\alpha_1^2 - \alpha_2^2) + (\beta\sigma_\xi)^2\alpha_3$, solves the HJB equation

$$\begin{aligned}
rV(p,\Delta) &= \max_{a\in\mathbb{R}}\Big\{-k_1\Delta - k_2(p-\overline{\theta}^2)^2 - \frac{\psi}{2}a^2 + [a-\kappa p]V_p(p,\Delta) + \frac{1}{2}\beta^2\sigma_\xi^2 V_{pp}(p,\Delta) \\
&\quad + [-\Delta(\beta+\kappa) + \frac{1}{\psi}(\alpha_1 - \alpha_2 + 2\alpha_3(p+\Delta)) - a]V_\Delta(p,\Delta)\Big\}
\end{aligned}$$

In fact, the right-hand side yields a first order condition of the form

$$\hat{a}(p,\Delta) = \frac{1}{\psi}[V_p(p,\Delta) - V_\Delta(p,\Delta)] = \frac{1}{\psi}[\alpha_1 - \alpha_2 + 2\alpha_3 p] = a^*(p)$$

which is also sufficient. The market's conjecture off the equilibrium path then takes the form $a^*(p+\Delta) = \frac{1}{\psi}[\alpha_1 - \alpha_2 + 2\alpha_3(p+\Delta)]$, as expressed in the last line of the

HJB equation. Inserting the above first order condition along with the corresponding expressions for V, V_p, V_Δ and V_{pp} in the HJB yields the system of equations

$$\begin{aligned}
(\alpha_0):\ r\alpha_0 &= -k_2\overline{\theta}^2 + \frac{1}{2\psi}(\alpha_1^2 - \alpha_2^2) + (\beta\sigma_\xi)^2\alpha_3 \\
(\alpha_1):\ r\alpha_1 &= 2k_2\overline{\theta} - \alpha_1\kappa + \frac{2\alpha_1\alpha_3}{\psi} \\
(\alpha_2):\ r\alpha_2 &= -k_1 + \alpha_2\left[\frac{2\alpha_3}{\psi} - (\beta+\kappa)\right] \\
(\alpha_3):\ r\alpha_3 &= -k_2 - 2\alpha_4\kappa + \frac{2\alpha_4^2}{\psi}.
\end{aligned}$$

The expressions for $\alpha_i,\ i = 0,1,2,3,$ stated above are the unique solution to this system.

In order to show that V is the agent's value function and a^* an equilibrium, two things remain to be checked:

1. Any strategy a such that

$$\mathbb{E}\left[\int_0^\infty e^{-rt}\left| - k_1\Delta_t^{a,a^*} - k_2(p_t^a - \overline{\theta})^2 - \frac{\psi}{2}a_t^2\right| dt\right] < \infty \tag{2.59}$$

must also satisfy $\limsup_{t\to\infty} e^{-rt}\mathbb{E}[V(p_t^a, \Delta_t^{a,a^*})] \geq 0$.

2. Second, that along the conjectured optimal strategy a^*, $\liminf_{t\to\infty} e^{-rt}\mathbb{E}[V(p_t^{a^*}, \Delta_t^{a^*})] \leq 0$.

Part 1. implies that V is an upper bound to the agent's utility. Part 2. yields that V is attainable under the Markov control a^* (refer to the proof of Theorem 2.3 for a proof of these statements).

1. *The transversality condition* $\limsup_{t\to\infty} e^{-rt}\mathbb{E}[V(p_t^a, \Delta_t^{a,a^*})] \geq 0$ *holds for any strategy* a *satisfying (2.59):*

Observe that under the pair (a, a^*), the belief-asymmetry process Δ^{a,a^*} takes the form

$$\Delta^{a,a^*} = e^{-\vartheta t}\Delta_o + [1 - e^{-\vartheta t}]\frac{\alpha_1 - \alpha_2}{\psi\vartheta} + \int_0^t e^{-\vartheta(t-s)}\left[\frac{2\alpha_3}{\psi}p_s^a + a_s\right] ds \tag{2.60}$$

with $\vartheta := \beta + \kappa - 2\alpha_3 > 0$. Moreover, since

$$p_t^a = e^{-\kappa t}p_0 + \int_0^t e^{-\kappa(t-a)}[a_s ds + \beta\sigma_\xi d\overline{Z}_s^0],\ t \geq 0, \tag{2.61}$$

we can use integration by parts and Fubini's theorem to show that the agent's expected discounted utility under strategy a can be written as

$$\begin{aligned} &\mathbb{E}\left[\int_0^\infty e^{-rt}\left(-k_1\Delta_t^{a,a^*} - k_2(p_t^a - \overline{\theta})^2 - \frac{\psi}{2}a_t^2\right)dt\right] \\ = \ &\mathbb{E}\left[\int_0^\infty e^{-rt}\left(-k_2(p_t^a)^2 + C_1 a_s - \frac{\psi}{2}a_s^2\right)dt\right] + C_2 \end{aligned}$$

for some constants C_1 and C_2. As a result, the agent's flow payoffs are uniformly bounded from above by a large constant. Therefore, (2.59) holds if and only if

$$\mathbb{E}\left[\int_0^\infty e^{-rt}(p_t^a)^2 dt\right] < \infty \quad \text{and} \quad \mathbb{E}\left[\int_0^\infty e^{-rt}a_t^2 dt\right] < \infty. \tag{2.62}$$

from where we conclude that $\liminf_{t\to 0} \mathbb{E}[e^{-rt}(p_t^a)^2] = 0$. Since $\alpha_3 < 0$,

$$\limsup_{t\to 0} \mathbb{E}[e^{-rt}\alpha_3(p_t^a)^2] = 0.$$

Now, plugging the expression for p_t^a into (2.60) and using integration by parts, we can find positive constants C_3, C_4 and C_5 such that

$$|\mathbb{E}[\Delta_t^{a,a^*}]| \leq C_3 e^{-\vartheta t} + C_4 \int_0^t e^{-\kappa(t-s)}\mathbb{E}[|a_s|]ds + C_5 \int_0^t e^{-\vartheta(t-s)}\mathbb{E}[|a_s|]ds.$$

The Cauchy-Schwartz's and Jensen's inequalities then yield that

$$e^{-rt}\int_0^t e^{-\lambda(t-s)}\mathbb{E}[|a_s|]ds \leq \underbrace{\left(e^{-rt}e^{-2\lambda t}\int_0^t e^{2\lambda s}ds\right)^{1/2}}_{L_t^1:=}\underbrace{\left(e^{-rt}\int_0^t \mathbb{E}[a_s^2]ds\right)^{1/2}}_{L_t^2:=}.$$

for $\lambda = \kappa, \vartheta > 0$. It is easy to see that $L_t^1 \to 0$ as $t \to \infty$. For L_t^2, observe that

$$e^{-rt}\int_0^t \mathbb{E}[a_s^2]ds < \int_0^t e^{-rs}\mathbb{E}[a_s^2]ds < \int_0^\infty e^{-rs}\mathbb{E}[a_s^2]ds < \infty$$

and thus L^2 is uniformly bounded. This shows that $\lim_{t\to 0} e^{-rt}\mathbb{E}[\Delta_t^{a,a^*}] = 0$. The same argument (Cauchy-Schwartz's and Jensen's inequalities and below) applied to $\mathbb{E}[p_t^a]$ in (2.61) yields $\lim_{t\to 0} e^{-rt}\mathbb{E}[p_t^a] = 0$. We conclude that $\limsup_{t\to\infty} e^{-rt}\mathbb{E}[V(p_t^a, \Delta_t^{a,a^*})] \geq 0$ for any feasible strategy satisfying (2.59).

2. *V is attained under a^*:* $\liminf_{t\to\infty} e^{-rt}\mathbb{E}[V(p_t^{a^*}, \Delta_t^{a^*})] \leq 0.$

If the agent follows the conjectured markov policy a^*, the belief-asymmetry process evolves according to

$$d\Delta_t = [-(\beta + \kappa)\Delta_t + \alpha_3 \Delta]dt,\ t > 0,\ \Delta_0 = \Delta^o.$$

As a consequence, $\Delta_t^{a^*} = e^{-(r+\beta+\kappa+\alpha_3)t}\Delta^o$. Moreover, since $\alpha_3 < 0$, $e^{-rt}\Delta_t^{a^*} \to 0$ as $t \to \infty$. Now the posterior belief process $p_t^{a^*}$ satisfies the SDE

$$dp_t = (\alpha_1 - \alpha_2 + \alpha_3 p_t - \kappa p_t)dt + \beta\sigma_\xi d\overline{Z}_t^0,\ t > 0,$$

i.e. $(p_t^{a^*})_{t\geq 0}$ is mean reverting around $\frac{\alpha_1-\alpha_2}{\kappa-\alpha_3}$. As a result,

$$p_t^{a^*}|\mathcal{F}_0 \sim \mathcal{N}\left(\frac{\alpha_1-\alpha_2}{\kappa-\alpha_3}[1-e^{-(\kappa-\alpha_3)t}] + p^o e^{(\kappa-\alpha_3)t}, (\beta\sigma_\xi)^2\frac{1-e^{-2(\kappa-\alpha_3)}}{2(\kappa-\alpha_3)}\right),\ t \geq 0.$$

with $\kappa - \alpha_3 > 0$, from where we conclude that the agent attains finite utility under the strategy a^* (i.e. it is feasible) and, moreover,

$$\lim_{t\to\infty}\mathbb{E}[p_t^{a^*}] = \lim_{t\to\infty}\mathbb{E}[(p_t^{a^*})^2] = 0.$$

This concludes the proof of the theorem.

□

Chapter 3

Learning, Investment and Adjustment Costs

3.1 Introduction

This chapter studies the effects of the *time-profile of uncertainty* on a firm's investment behavior. In the model we study, a firm makes investments whose cash flows are realized ahead in the future, and the value of these cash flows is driven by an unobserved stochastic process. However, the firm can *learn* about the current value of this underlying fundamental from observing a noisy signal. Consequently, and unlike the existing literature, the firm faces both *current* and *future* uncertainty when making its investment decisions: both the current value of fundamentals and the future evolution of perceived fundamentals are unknown. In such a context, we show how a firm's investment decisions are affected by the way in which uncertainty resolves over time, and how better information generates more value.

Understanding the determinants of firms' investment decisions has been a central topic in macroeconomic research. While in the neoclassical model of Jorgenson (1963) marginal revenue of capital always equates its user cost, subsequent studies have recognized the impact that frictions can have on investment behavior. A large literature modeling internal frictions in the form of convex costs of adjustment (Eisner and Strotz (1963) the earliest reference) has shown how anticipated future economic conditions can generate investment smoothing, something absent in Jorgenson's frictionless model. Another strand of literature has argued that the forward looking behavior of some investment decisions is closely related to their inherent irreversibility (see Dixit and Pindyck (1994) and references therein). In all these environments, investment is a function of Tobin's marginal q, that is, the value of an installed unit of capital (Tobin (1969), Hayashi (1982)).

When investment decisions are affected by future economic conditions, uncertainty plays a crucial role in determining the dynamics of capital. The literature on sequential investment under uncertainty is large (classic examples are Pindyck (1982) and Abel (1983)), yet mostly focused on long-run analysis:[1] production functions exhibit constant returns to scale and firms have perfect information about all current payoff relevant variables (e.g. technological parameters or the law of motion of prices). A large bulk of this literature has concentrated on the case in which it is only the future evolution of stochastic variables what is unknown to the firm at any point in time. Changes in uncertainty in turn refer to changes in parameters that capture the degree of randomness of the stationary environment at hand, such as output price volatility.

In the short run however, firms face physical constraints when making their investment decisions. For instance, some factors of production can be potentially fixed, preventing firms from freely adjusting all production inputs to their desired level. Moreover, firms face informational constraints when choosing how much to invest. For example, firms make ex-ante investments when payoff-relevant variables are realized ahead in the future, or firms make investments in a context of partial information when these payoff-relevant variables are simply unobserved (e.g. the demand for a

[1] An exception is Bertola and Caballero (1994).

firm's product can fluctuate between unobserved states, or the current state of a firm's technological process can be in fact hidden). While it is known that decreasing return to scale will make Tobin's q to depend on the firm's size, it is unknown what is the effect that the *time profile* of uncertainty has on a firm's investment behavior.

This chapter develops a simple model of investment in the presence of adjustment costs in the context of partial information. More specifically, a firm learns about a hidden technological process that affects its cash-flows generating capacity (or more succinctly, its earnings) from observing a noisy signal about the current state of its technology. At the same time, the firm makes investments in capital that affects its future earnings. The model has the following characteristics: (i) investment is perfectly reversible, yet costly according to quadratic adjustment costs satisfying homogeneity of degree one; (ii) the firm's earnings technology has decreasing returns to scale with respect to capital; (iii) the firm earning's technology depends on an unobserved technological process, which we refer to as the firm's *fundamentals*; (iv) learning is Gaussian. Because fundamentals are not observed, the model allows for uncertainty about current economic conditions. Because the evolution of beliefs about the underlying fundamental is driven by an exogenous signal, there is also uncertainty about the future value of perceived fundamentals, which in turn determines the perceived marginal return of capital at all future dates. This model seems to be the first one to allow for these two different types of uncertainty, a distinction first made by Pindyck (1982).

In the particular specification we study, we show how the degree of convexity/concavity of fundamentals in the earnings function crucially determines how investment responds to the time profile of uncertainty. Because the model allows for convex costs of adjustment, the optimal level of investment depends on all future levels of marginal productivity of capital. Thus, when the marginal productivity of capital is concave (convex) in fundamentals, higher uncertainty reduces (increases) the value of installed capital. But since the firm learns about its own fundamentals as information accumulates, this uncertainty will decrease over time. As a result, when earnings are concave (convex) in fundamentals, the *sensitivity of investment to perceived earnings* will be an increasing (decreasing) function of time, everything else equal.

The implications on firm-level investment are particularly interesting. When earnings are linear in fundamentals, investment is independent of the time profile of uncertainty. Thus, all firms with the same size and perceived fundamentals invest the same amount of capital at any point in time. If earnings are instead concave in fundamentals, higher uncertainty will decrease investment relative to the linear benchmark, everything else equal. Consequently, among two firms with the same size, earnings technology and perceived fundamentals, younger firms invest less. The opposite occurs when earnings are convex in fundamentals. This results are in contrast with the traditional steady-state analysis in which two seemingly identical firms exhibit the same investment policies. In the presence of learning about unobserved payoff relevant variables, time in an industry is expected to have non-negligible effects over investment.

The model also provides a useful framework for studying the effect of more precise information structures on the value of the firm. In particular, we show that the firm's value when fundamentals are unobserved (second-best) is always below its *ex-ante* value in the corresponding full-information benchmark (first-best). Consequently, firms would be willing to pay ex-ante an strictly positive amount of money for perfectly informative signals that reveal the true value of fundamentals at all future dates. The possibility of eliminating contemporaneous uncertainty thus generates strictly positive value from a resource-allocation perspective.

Finally, another important aspect of this model is its tractability: we are able to obtain closed for solutions for the dynamic programming problem faced by a firm in the context of (i) time-dependent dynamics (which arise due the firm's learning process) and (ii) adjustment costs that explicitly depend on the firm's size. To our best knowledge, this is the first model that is able to obtain exact solutions for an optimization problem with such characteristics.[2] Because the earnings production function has decreasing returns to scale, investment depends both on the firm's size and on the marginal value of installed capital.

The literature studying the determinants of investment is large, starting with the seminal work of Jorgenson (1963) and the adjustment costs model of Einer and Strotz (1963). Tobin (1969) argued that investment is an increasing function of the ratio of the firm's market value to the replacement cost of capital, or average q. Hayashi (1982) (and also Mussa (1977)) shows how investment is actually determined by marginal q, and he derives conditions under which average q equals marginal q. All these models are deterministic.

Stochastic models of reversible investment in the context of convex costs of adjustment were developed in the 1980's and 1990's. Pindyck (1982) is the first model of investment under adjustment costs that studies the impact of future uncertainty using dynamics that are distorted by Brownian noise. Abel (1983, 1985) studied models involving production functions exhibiting constant returns to scale and convex adjustment costs independent of firm size. He finds that increased uncertainty (as measured by an increase in price volatility) increases investment. In his model, operating profits as a function of capital and the price of the final good only (i.e. maximizing over all other inputs of production) is a convex function of the latter variable. Consequently, his results are consistent with my findings.

The idea of irreversible economic decisions and its consequences on investment goes back to Arrow (1968). More recent studies involving partially irreversible in-

[2]Existing closed-form solutions for investment problems allowing for payoff-relevant state variables other than capital (for instance, prices) assumed both time-homogenous dynamics and costs of investment that are independent of the firm's size. See for instance Abel (1983) and Abel and Eberly (1997).

vestments are Abel and Eberly (1994, 1997 and 1999) and Bertola and Caballero (1994). In all these models there are "inaction regions" where the firm does not invest in capital. For a general treatment of all these type of irreversibilities, see Dixit and Pindyck (1994).

3.2 The Model

Consider a infinitely lived firm that uses its capital to generate revenue. More specifically, I model a firm as an economic entity that is able to generate a cash flow stream at every point in time. Denoting by k_t the firm's stock of capital at time t, the (true) net present value of the cash flow stream generated at time t is given by

$$y_t = \delta_t^{\alpha} \sqrt{k_t},\ t \geq 0, \tag{3.1}$$

where $\alpha > 0$ is a scalar. In this specification, the process $\delta := (\delta_t)_{t\geq 0}$ corresponds to a state variable reflecting current economic conditions. This *fundamental* variable could be external to the firm (e.g. an underlying state of the economy such as the state of demand) or internal to it (e.g. the current state of a firm's technology). In either case, δ captures the idea that the performance of any investment decision depends on forces that are exogenous to the manager. The specification I use corresponds to a geometric Brownian motion

$$\frac{d\delta_t}{\delta_t} = gdt + \sigma dZ_t,\ t \geq 0, \tag{3.2}$$

with g and σ two positive scalars and $Z := (Z_t)_{t\geq 0}$ a Brownian motion.

I refer to the process $y := (y_t)_{t\geq 0}$ as the firm's *realized earnings.* Moreover, I assume that the earnings associated with an investment decision at time t, y_t, are received far ahead in the future, making y_t not readily observable at time $t \geq 0$.

Thus, the current level of the firm's fundamentals at time t, δ_t, is unobserved to the manager at time $t \geq 0$. This is in contrast with the traditional investment models in which operating profits are observable without any lags.

The manager instead observes a noisy signal of current fundamentals. Letting $\theta_t := \log(\delta_t)$, $t \geq 0$, the manager observes the process

$$d\xi_t = \theta_t dt + \sigma_\xi dZ_t^\xi, \ t \geq 0, \tag{3.3}$$

where $Z^\xi := (Z_t^\xi)_{t\geq 0}$ is a Brownian motion independent of Z. Using the information conveyed by $\xi := (\xi_t)_{t\geq 0}$, the manager learns about the current value of fundamentals.

The firm's stock of capital evolves according to the traditional dynamic

$$dk_t = (i_t - \lambda_t k_t)dt, \ t \geq 0, \tag{3.4}$$

where i_t denotes the firm's investment level at time $t \geq 0$, and λ is the rate at which capital depreciates. Adjusting the firm's level of capital is costly according to the function

$$c(i, k) = i + \psi \frac{i^2}{2k}, \tag{3.5}$$

where $\psi > 0$. Observe that the investment required to keep the firm's size (given by its stock of capital) constant is linear in its size k (scale invariance).

Finally, the manager's problem is to choose an investment process $(i_t)_{t\geq 0}$ that maximizes the firm's expected discounted benefits

$$\mathbb{E}\left[\int_0^\infty e^{-rt}\left(\delta_t^\alpha \sqrt{k_t} - i_t - \psi \frac{i_t^2}{2k_t}\right) dt\right] \tag{3.6}$$

subject to (3.2), (3.3) and (3.4), where $r > 0$ is the manager's discount rate.

Two interpretations regarding the timing of earnings and the learning in the model are possible. The first one is that all earnings generated by the firm are received very far in the future (technically, at $+\infty$). In this case, $r > 0$ is represents a preference-driven discount rate only. The second one is that there is an exponentially-distributed random time τ, at which all earnings $(y_t)_{t\leq\tau}$ are received (and, consequently, the whole path $(\delta)_{t\leq\tau}$ is revealed), and the firm stops operating at τ. In this case, $r > 0$ can be decomposed into a time-preference component, and the rate at which this day of reckoning is expected to arrive.

3.3 Learning

In this Section we show that the firm's learning problem has a Markovian structure which allows us to use dynamic programming.

A direct application of Ito's rule yields that the process $\theta_t := \log(\delta_t)$, $t \geq 0$, evolves according to

$$d\theta_t = (g - \sigma^2/2)dt + \sigma_t dZ_t,\ t \geq 0. \tag{3.7}$$

Consequently, the system (3.7)-(3.3) consisting of the dynamics of fundamentals and the signal ξ, respectively, is a Gaussian system. When starting from a Gaussian prior, the Kalman-Bucy filter applies. Letting $m_t := \mathbb{E}[\theta_t|\mathcal{F}_t^\xi]$ and $\gamma_t := \mathbb{E}[(\theta_t - m_t)^2|\mathcal{F}_t^\xi]$ denote the posterior first and second moment of θ_t given the information $\mathcal{F}_t^\xi$, $t \geq 0$, respectively, we have the following:

Lemma 3.1. *Consider the system defined by the firm's fundamentals (3.7) and the signal ξ (3.3). Suppose that the agent's initial prior $\theta_0|\mathcal{F}_0$ is normally distributed $\mathcal{N}(m^o, \gamma^o)$. Then, $\theta_t|\mathcal{F}_t^\xi \sim \mathcal{N}(m_t, \gamma_t)$, where the posterior mean process $m_t := \mathbb{E}[\theta_t|\mathcal{F}_t^\xi]$ and posterior variance $\gamma_t := \mathbb{E}[(\theta_t - m_t)^2|\mathcal{F}_t^Y]$, $t \geq 0$, satisfy the* *stochastic differential equation (SDE) and ordinary differential equation (ODE)*

$$dm_t = (g - \sigma^2/2)dt + \frac{\gamma_t}{\sigma_\xi}\frac{d\xi_t - m_t dt}{\sigma_\xi}, \text{ and} \tag{3.8}$$

$$\dot{\gamma}_t = \sigma^2 - \left(\frac{\gamma_t}{\sigma_\xi}\right)^2, \; t > 0, \tag{3.9}$$

respectively. Finally, the process $Z'_t := \frac{1}{\sigma_\xi}\left(\xi_s - \int_0^t m_s ds\right)$*,* $t \geq 0$*, is an exogenous* $\mathbb{F}^\xi$*-Brownian motion from the manager's perspective, called the innovation process.*

Proof: See Liptser and Shiryaev (1977).

□

Three interesting features of (3.8) and (3.9) are worth noting. First, the evolution of the posterior mean preserves the stochastic structure of θ: the drift of the posterior mean remains unchanged. Second, the posterior mean's response to signal surprises (captured by the innovation process) increases with the size of the mean-square error and decreases with the signal's volatility (σ_ξ). This implies that beliefs react more strongly in settings where either less information has been accumulated, or where signals are more accurate. Finally, the mean-square error evolves in a *deterministic* fashion.

From the manager's perspective, the exogenous component of the firm's expected earnings is driven by the process $\delta'_t := \mathbb{E}[\delta_t|\mathcal{F}^\xi_t]$, $t \geq 0$. Using θ's definition, it is easy to see that

$$\delta'_t = \mathbb{E}[\exp(\theta_t)|\mathcal{F}^\xi_t] = \exp\left(m_t + \frac{\gamma_t}{2}\right),$$

which yields that

$$\mathbb{E}_t[\delta^\alpha_t] = \exp\left(\alpha m_t + \frac{1}{2}\alpha^2\gamma_t\right) = (\delta'_t)^\alpha \exp\left(\frac{\alpha(\alpha-1)}{2}\gamma_t\right), \; t \geq 0. \tag{3.10}$$

Moreover, applying the product rule to $\delta'_t = \exp(m_t)\exp(\gamma_t/2)$ and using the previous Lemma, it is easy to show that

$$\frac{d\delta'_t}{\delta'_t} = gdt + \sigma'_t dZ'_t,\ t \geq 0, \tag{3.11}$$

where Z' is a Brownian motion from the manager's perspective, and $\sigma'_t := \sigma^2 - \left(\frac{\gamma_t}{\sigma_\xi}\right)^2$.

The manager thus solves

$$\max_{(i_t)_{t\geq 0}} \mathbb{E}\left[\int_0^\infty e^{-rt}\left((\delta'_t)^\alpha \exp\left(\frac{\alpha(1-\alpha)}{2}\gamma_t\right)\sqrt{k_t} - i_t - \psi\frac{i_t^2}{2k_t}\right)dt\right] \tag{3.12}$$

s.t. (3.11) and the dynamics of capital (3.4).

In the next section we show that the previous problem admits an analytic solution in the state variables (t, δ', k).

3.4 The Value of the Firm

Let $V(t, \delta', k)$ denote the firm's *value* at time t when perceived fundamentals are $\delta' \geq 0$ and the firm's size is $k \geq 0$. It is well know that when V is regular enough, it must satisfy the following Hamilton-Jacobi-Bellman (HJB) equation: For all (t, δ', k)

$$\begin{aligned} rV(t,\delta',k) &= \max_{i\in\mathbb{R}}\left\{\exp\left(\frac{\alpha(1-\alpha)\gamma_t}{2}\right)(\delta')^\alpha\sqrt{k} - i - \psi\frac{i^2}{2k} + \frac{\partial V}{\partial t}(t,\delta',k)\right. \\ &\left. +(i-\lambda k)\frac{\partial V}{\partial k}(t,\delta',k) + \delta' g\frac{\partial V}{\partial \delta'}(t,\delta',k) + \frac{1}{2}(\sigma'_t\delta')^2\frac{\partial^2 V}{\partial \delta'^2}(t,\delta',k)\right\}. \end{aligned} \tag{3.13}$$

The optimal investment rule given is

$$i = \frac{k}{\psi}\left[\frac{\partial V}{\partial k}(t,\delta',k) - 1\right],\ \text{for all}\ \ (t,\delta',k), \tag{3.14}$$

and the HJB equation becomes

$$\begin{aligned} rV(t,\delta',k) &= \exp\left(\frac{\alpha(1-\alpha)\gamma_t}{2}\right)(\delta')^{\alpha}\sqrt{k} + \frac{k}{2\psi}\left[1+\left(\frac{\partial V}{\partial k}\right)^2\right] - k\left(\lambda+\frac{1}{\psi}\right)\frac{\partial V}{\partial k} \\ &\quad + \frac{\partial V}{\partial t}(t,\delta',k) + \delta' g\frac{\partial V}{\partial \delta'}(t,\delta',k) + \frac{1}{2}(\sigma_t'\delta')^2\frac{\partial^2 V}{\partial \delta'^2}(t,\delta',k). \quad (3.15) \end{aligned}$$

The next result shows that the previous partial differential equation (PDE) admits a closed-form solution which turns out to be the value of the firm:

Theorem 3.1. *Suppose that $r > 2\alpha g + \alpha(2\alpha-1)\sigma^2$. Then, the firm's value takes the form*

$$V(t,\delta',k) = \beta_1 k + \beta_{2,t}(\delta')^{2\alpha} + \beta_{3,t}(\delta')^{\alpha}\sqrt{k} \quad (3.16)$$

where

$$\beta_1 = \psi\left[r+\lambda-\frac{1}{\psi}-\sqrt{\left(r+\lambda-\frac{1}{\psi}\right)^2-\frac{1}{\psi^2}}\right] \quad (3.17)$$

$$\beta_{3,t} = \frac{e^{\frac{\alpha(\alpha-1)}{2}\gamma_t}}{\rho}, \text{ and} \quad (3.18)$$

$$\beta_{2,t} = e^{\alpha(2\alpha-1)\gamma_t}\int_t^{\infty} e^{-[r-2\alpha g-\alpha(2\alpha-1)\sigma^2](s-t)}e^{-\alpha(2\alpha-1)\gamma_s}\frac{\beta_{3,s}^2}{8\psi}ds, \ t \geq 0, \quad (3.19)$$

with $\rho = r + \frac{\lambda\psi+(1-\beta_1)}{2\psi} - \alpha g - \frac{\alpha(\alpha-1)}{2}\sigma^2 > 0$.

Proof: See the Appendix.

□

The proof of the theorem relies on verification theorems for dynamic programming. We first show that an additively separable function of the form

$$V(t,\delta',k) = \beta_{1,t}k + \beta_{2,t}(\delta')^{2\alpha} + \beta_{3,t}(\delta')^{\alpha}\sqrt{k},$$

solves the PDE whenever $(\beta_{i,t})_{t\geq 0}$, $i = 1, 2, 3$ solve a particular system of differential equations. Then, we show that β_1, $\beta_{2,t}$ and $\beta_{3,t}$ as in the theorem, satisfy suitable transversality conditions (ensuring that V corresponds to the firm's value) as long as $r > 2\alpha g + \alpha(2\alpha - 1)\sigma^2$ holds. The latter condition ensures that the firm's value if finite.

3.5 Capital Dynamics and Optimal Investment Rule: The Time-Profile of Uncertainty

Theorem 3.1 delivers a particularly clean solution of the firm's problem. The optimal investment rule is given by

$$i(t, k, \delta') = \frac{k}{\psi}\left[\frac{\partial V}{\partial k}(t, \delta', k) - 1\right] = \frac{1}{\psi}\left[k(\beta_1 - 1) + \frac{\beta_{3,t}}{2}(\delta')^{\alpha}\sqrt{k}\right], \tag{3.20}$$

where

$$\beta_{3,t} = \frac{\exp\left(\frac{\alpha(\alpha-1)}{2}\gamma_t\right)}{\rho}, \ t \geq 0,$$

captures how the evolution of uncertainty affects the firm's investment behavior. We refer to this deterministic function as the *sensitivity of investment to expected earnings*.

When a young firm is very uncertain about its fundamentals, uncertainty will decrease over time as information accumulates ($(\gamma_t)_{t\geq 0}$ is strictly decreasing when γ_0 is large). Hence, if the earnings technology is concave in fundamentals ($\alpha < 1$), the sensitivity of investment to expected earnings will be an increasing function of time. Hence, higher uncertainty reduces investment, everything else equal. The opposite result holds when earnings are a convex function of fundamentals ($\alpha > 0$). Finally, when earnings are linear in fundamentals, investment is not affected by the future evolution of uncertainty.

The next proposition summarizes this discussion:

Proposition 3.1. *Starting from $k_0 > 0$, the optimal stock of capital at any time t is given by*

$$k_t = \left[e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}t} k_0 + \frac{1}{4\psi} \int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}(t-s)} \beta_{3,s}(\delta'_s)^{\alpha} ds \right]^2, \; t \geq 0. \tag{3.21}$$

Finally, $\alpha \leq 1$ if and only if $i(s, \delta', k) \leq i(t, \delta', k)$, $0 \leq s < t$.

Proof: See the Appendix.

□

The previous result has implications on the cross sectional distribution of investment profiles observed in an industry. When earnings are concave in fundamentals, two seemingly identical firms (same size and perceived investment opportunities) will have different investment profiles, with older firms growing faster. The contrary occurs when earnings are convex in fundamentals. Finally, when earnings are linear in fundamentals, two seemingly identical firms should be expected to grow at similar rates, irrespective of differences in their time in an industry.

The intuition for the results stems the monotone evolution of uncertainty over time. When earnings are concave in fundamentals, the expected marginal productivity of capital at all future increases with time in an industry, everything else equal. Hence, capital becomes more valuable for this type of firms as they learn more about their unobserved fundamentals.

3.6 First-Best: The Value of Information

An important question in the analysis presented in this chapter is whether better information adds value to the firm or not. Intuitively, better information structures allow firms to allocate capital more optimally across different states of the world, and this in turn yields a higher value.

In this section we show that the value of the firm found in the previous section (second-best) is always below the firm's *ex-ante* value in the corresponding full-information benchmark (first-best). The difference between both values represents the firm's willingness to pay for the resolution of contemporaneous uncertainty *at all future dates*, given its current beliefs about fundamentals.

We can assume that $\alpha = 1$ without loss of generality. This is because we will compare the firm's value across two different information structures at any given point in time, and because δ^{α} also follows a geometric Brownian motion.

From Theorem 3.1, the value of the firm takes the form

$$V(t, \delta', k) = \beta_1 k + \beta_{2,t}(\delta')^2 + \beta_{3,t}\delta'\sqrt{k},$$

with β_1, $(\beta_{2,t})_{t\geq 0}$ and $(\beta_{3,t})_{t\geq 0}$ as in the Theorem, and $\delta' = \exp\left(m + \frac{\gamma}{2}\right)$. In particular, $\beta_{3,t} = \beta_3 := \frac{1}{\rho}$ and

$$\beta_{2,t} = \frac{e^{\gamma_t}}{8\psi\rho^2}\int_t^{\infty} e^{-(r-2g-\sigma^2)(s-t)}e^{-\gamma_s}ds$$

when $\alpha = 1$.

It is easy to see that when δ is perfectly observed, the firm's value V^{fb} is given by

$$V^{fb}(\delta, k) = \beta_1^{fb}k + \beta_2^{fb}\delta^2 + \beta_3^{fb}\delta\sqrt{k}$$

with $\beta_1^{fb} = \beta_1$, $\beta_2^{fb} = \frac{1}{8\psi\rho^2}\frac{1}{r-2g-\sigma^2}$ and $\beta_3^{fb} = \beta_3 := \frac{1}{\rho}$ (that is, the corresponding values of β_i, $i = 1, 2, 3$ when $\gamma \equiv 0$). The firm's ex-ante value at time t is thus given by

$$\mathbb{E}_t[V^{fb}(\delta_t, k_t)] = \beta_1 k_t + \beta_2^{fb}\mathbb{E}_t[\delta_t^2] + \beta_3\mathbb{E}_t[\delta_t]\sqrt{k_t},$$

where k_t denotes the stock of installed capital right before observing the full path $(\delta_s)_{s>t}$, $t > 0$. Since $\log(\delta_t) \sim \mathcal{N}(m_t, \gamma_t)$ the firm's ex-ante value at time t in the full-information benchmark takes the form

$$\mathbb{E}_t[V^{fb}(\delta_t, k_t)] = \beta_1 k_t + \beta_2^{fb} \exp\left(2m_t + 2\gamma_t\right) + \beta_3 \exp\left(m_t + \frac{\gamma_t}{2}\right)\sqrt{k_t},\ t > 0.$$

Hence,

$$\mathbb{E}_t[V^{fb}(\delta_t, k_t)] - V(t, \delta_t', k_t) = \beta_2^{fb} \exp\left(2m_t + 2\gamma_t\right) - \beta_{2,t} \exp\left(2m_t + \gamma_t\right)\ t \geq 0.$$

From the expression for $\beta_{2,t}$ we obtain that

$$\beta_{2,t} < \frac{e^{\gamma_t}}{8\psi\rho^2} \int_t^\infty e^{-(r-2g-\sigma^2)(s-t)} ds = e^{\gamma_t} \beta_2^{fb},$$

yielding $\mathbb{E}_t[V^{fb}(\delta_t, k_t)] - V(t, \delta_t', k_t) > 0$.

3.7 Conclusions

This chapter developed a tractable model of investment in the context of partial information. A firm learns about an unobserved process which drives the evolution of its earnings. At the same time, the firms invest in capital, which also affects its cash-flows generating capacity. Since the contemporaneous optimal level of capital is always unknown, the firm faces uncertainty about current economic conditions at the moment of making a decision. But the firm also faces uncertainty about the future, in the sense that the future evolution of perceived fundamentals (hence, the value of the marginal productivity of capital at all future dates) cannot be perfectly anticipated.

In this context, we showed how learning about a firm's technology generates investment behavior that varies across a firm's lifecycle. More specifically, since uncertainty decreases as information accumulates, the sensitivity of investment to perceived current economic conditions vary with time in an industry. The main prediction of this result is that two firms with similar size and expected earnings, but different time in an industry, should exhibit different investment profiles, something absent in traditional dynamic models of investment under uncertainty. We also showed that having perfect information generates higher benefits (in terms of allocation of investment resources) than the second-best case involving learning about fundamentals. Consequently, more imprecise information about fundamentals reduces a firm's value.

Closed-form solutions for non-linear parabolic PDEs are certainly an exception. Consequently, the analytic results derived in this chapter depend heavily on the functional forms assumed. However, the economic properties of the solution found here are expected to hold in more general environments characterized by: (i) earnings technologies that are increasing and concave in capital and (ii) convex costs of adjustment that are homogenous of degree one. Deriving qualitative properties of the firm's value function along with numerical solutions to the associated PDEs is a promising approach to understanding the dynamics generated in these more general environments.

Finally, it would be interesting to connect the predictions generated by this model both to the empirical evidence regarding investment within particular industries, and to the accounting literature studying how the information contained in cash flows, earnings and market valuations impacts investment. These and other questions are left for future research.

3.8 Appendix

Before stating the proof of the Theorem, it is instructive to show the following Lemmas:

Lemma 3.2. *For all* $t \geq 0$*:*

$$\mathbb{E}[(\delta_t')^\alpha] = (\delta_0')^\alpha \exp\left(\frac{\alpha(\alpha-1)}{2}(\gamma_0 - \gamma_t)\right) \exp\left(\alpha g t + \frac{\alpha(\alpha-1)}{2}\sigma^2 t\right)$$

Proof: Observe first that

$$\delta_t' = \delta_0' \exp\left(gt - \frac{1}{2}\int_0^t (\sigma_u')^2 du + \int_0^t (\sigma_u') dZ_u'\right),\ t \geq 0,$$

with $\delta_0' = \mathbb{E}[\delta_0] = \mathbb{E}[\exp(\theta_0)] = \exp(m_0 + \frac{\gamma_0}{2})$, solves the SDE $d\delta_t' = \delta_t' g dt + \delta_t' \sigma_t' dZ_t'$, $t \geq 0$. Now, since

$$\int_0^t (\sigma_u') dZ_u' \sim \mathcal{N}\left(0, \int_0^t (\sigma_u')^2 du\right),\ t \geq 0,$$

we have that

$$\mathbb{E}[(\delta_t')^\alpha] = (\delta_0')^\alpha \exp\left(\alpha g t - \frac{\alpha}{2}\int_0^t (\sigma_u')^2 du\right) \exp\left(\frac{\alpha^2}{2}\int_0^t (\sigma_u')^2 du\right),\ t \geq 0.$$

Using that $\int_0^t (\sigma_u')^2 du = \sigma^2 t + \gamma_0 - \gamma_t$ we conclude.

□

Lemma 3.3. *Suppose that* $r > 2\alpha g + \alpha(2\alpha - 1)\sigma^2$ *and let*

$$\beta_1 := \psi \left[r + \lambda - \frac{1}{\psi} - \sqrt{\left(r + \lambda - \frac{1}{\psi}\right)^2 - \frac{1}{\psi^2}} \right].$$

Then, the ordinary differential equation

$$0 = \dot{\beta}_{3,t} - \beta_{3,t}\left[r + \frac{\lambda\psi + (1-\beta_1)}{2\psi} - \alpha g - \frac{\alpha(\alpha-1)}{2}(\sigma_t')^2\right] + \exp\left(\frac{\alpha(1-\alpha)\gamma_t}{2}\right)$$

with transversality condition

$$\lim_{t\to\infty} \beta_{3,t} \exp\left(-\int_0^t \left[r + \frac{\lambda\psi + (1-\beta_1)}{2\psi} - \alpha g - \frac{\alpha(\alpha-1)}{2}(\sigma_s')^2\right] du\right) = 0$$

admits $\beta_{3,t} = \rho^{-1} e^{\frac{\alpha(\alpha-1)}{2}\gamma_t}$, $t \geq 0$, *with* $\rho := r + \frac{\lambda\psi+(1-\beta_1)}{2\psi} - \alpha g - \frac{\alpha(\alpha-1)}{2}\sigma^2$, *as a solution.*

Proof: Observe first that since α and ψ are strictly positive, $r > 2\alpha g + \alpha(2\alpha-1)\sigma^2$ implies

$$\underbrace{r + \frac{2}{\psi} + \sqrt{\left(r + \lambda + \frac{1}{\psi^2}\right)^2 - \frac{1}{\psi^2}}}_{=2\left(r+\frac{\lambda\psi+1-\beta_1}{2\psi}\right)} > 2\alpha g + \alpha(\alpha-1)\sigma^2.$$

Consequently,

$$\rho = r + \frac{\lambda\psi + 1 - \beta_1}{2\psi} - \alpha g - \frac{\alpha(\alpha-1)}{2}\sigma^2 > 0.$$

Now, by definition of σ', $(\sigma_t')^2 := \left(\frac{\gamma_t}{\sigma_\xi}\right)^2 = \sigma^2 - \dot{\gamma}_t$, $t \geq 0$. Hence, the ODE can be written as

$$0 = \dot{\beta}_{3,t} - \beta_{3,t}\left[\rho + \frac{\alpha(\alpha-1)}{2}\dot{\gamma}_t\right] + \exp\left(\frac{\alpha(1-\alpha)\gamma_t}{2}\right).$$

It is easy to see that $\rho^{-1}e^{\frac{\alpha(\alpha-1)}{2}\gamma_t}$, $t \geq 0$, solves this ODE. To check the transversality condition, notice that

$$\exp\left(-\int_0^t \left[r + \frac{\lambda\psi + (1-\beta_1)}{2\psi} - \alpha g - \frac{\alpha(\alpha-1)}{2}(\sigma_s')^2\right] du\right) = e^{-\rho t} e^{\frac{\alpha(\alpha-1)}{2}(\gamma_0 - \gamma_t)}.$$

Since $\gamma_t \searrow \gamma^* = \sigma\sigma_\xi$ and $\rho > 0$, we conclude.

□

Lemma 3.4. *Under the assumptions of the previous Lemma, the ODE*

$$0 = \dot{\beta}_{2,t} - \beta_{2,t}[r - 2\alpha g - \alpha(2\alpha-1)(\sigma_t')^2] + \frac{\beta_{3,t}^2}{8\psi}$$

with transversality condition

$$\lim_{t\to\infty} \beta_{2,t} \exp\left(-\int_0^t [r - 2\alpha g - \alpha(2\alpha-1)(\sigma_s')^2] ds\right) = 0$$

admits

$$\beta_{2,t} = \frac{e^{\alpha(2\alpha-1)\gamma_t}}{8\psi} \int_t^\infty e^{-(r-2\alpha g-\alpha(2\alpha-1)\sigma^2)(s-t)} e^{-\alpha(2\alpha-1)\gamma_s} \beta_{3,s}^2 ds, \ t \geq 0,$$

as a solution.

Proof: Since $(\sigma_t')^2 = \sigma^2 - \dot{\gamma}_t$, $t \geq 0$, the ODE can be written as

$$0 = \dot{\beta}_{2,t} - \beta_{2,t}[r - 2\alpha g - \alpha(2\alpha-1)\sigma^2 + \alpha(2\alpha-1)\dot{\gamma}_t] + \frac{\beta_{3,t}^2}{8\psi},$$

from where it is clear that $\beta_{2,t}$ as in the Lemma solves this equation.

Finally, the transversality condition takes the form

$$\lim_{t\to\infty} \frac{e^{\alpha(2\alpha-1)\gamma_t}}{8\psi} \int_t^\infty e^{-(r-2\alpha g-\alpha(2\alpha-1)\sigma^2)s} e^{-\alpha(2\alpha-1)\gamma_s} \beta_{3,s}^2 ds,$$

which is equivalent to

$$\lim_{t\to\infty} \int_t^\infty e^{-(r-2\alpha g-\alpha(2\alpha-1)\sigma^2)s} ds,$$

since both $(\gamma_t)_{t\geq 0}$ and $(\beta_{3,t})_{t\geq 0}$ are uniformly bounded. The last limit equal zero because of the condition $r - 2\alpha g - \alpha(2\alpha-1)\sigma^2 > 0$.

□

Now we can prove the Theorem:

Proof of Theorem 3.1: Consider the PDE (3.15) given by

$$
\begin{aligned}
rV(t,\delta',k) &= \exp\left(\frac{\alpha(1-\alpha)\gamma_t}{2}\right)(\delta')^{\alpha}\sqrt{k} + \frac{k}{2\psi}\left[1+\left(\frac{\partial V}{\partial k}\right)^2\right] - k\left(\lambda+\frac{1}{\psi}\right)\frac{\partial V}{\partial k} \\
&\quad + \frac{\partial V}{\partial t}(t,\delta',k) + \frac{\partial V}{\partial \delta'}(t,\delta',k)\delta' g + \frac{1}{2}(\sigma_t'\delta')^2\frac{\partial^2 V}{\partial \delta'^2}(t,\delta',k),
\end{aligned}
$$

and conjecture a solution of the form

$$
V(t,\delta',k) = \beta_{1,t}k + \beta_{2,t}(\delta')^{2\alpha} + \beta_{3,t}(\delta')^{\alpha}\sqrt{k}. \tag{3.22}
$$

The PDE is then reduced to solving the following system of ordinary differential equations:

$$
0 = \dot{\beta}_{1,t} - \beta_{1,t}\left[r+\lambda+\frac{1}{\psi}\right] + \frac{1}{2\psi}(1+\beta_{1,t}^2) \tag{3.23}
$$

$$
0 = \dot{\beta}_{2,t} - \beta_{2,t}[r - 2\alpha g - (2\alpha-1)\alpha(\sigma_t')^2] + \frac{\beta_{3,t}^2}{8\psi} \tag{3.24}
$$

$$
\begin{aligned}
0 &= \dot{\beta}_{3,t} - \beta_{3,t}\left[r + \frac{\lambda\psi + (1-\beta_{1,t})}{2\psi} - \alpha g - \frac{\alpha(\alpha-1)}{2}(\sigma_t')^2\right] \\
&\quad + \exp\left(\frac{\alpha(1-\alpha)\gamma_t}{2}\right)
\end{aligned} \tag{3.25}
$$

For the first ODE, consider its smallest stationary solution:

$$
\beta_1 = \psi\left[r+\lambda-\frac{1}{\psi} - \sqrt{\left(r+\lambda-\frac{1}{\psi}\right)^2 - \frac{1}{\psi^2}}\right]
$$

For the second and third ODEs, consider the solutions derived in Lemmas (3.4) and (3.3), respectively.

In order to show that V as defined above is the firm's value function, it suffices to show that

(i) The investment rule $i(t,k,\delta') = \frac{k}{\psi}\left[\frac{\partial V}{\partial k}(t,\delta',k) - 1\right]$ induces a well-defined law of motion of capital and;

(ii) Along the capital path induced by the previous investment rule, the following transversality condition holds

$$\lim_{t\to 0} e^{-rt}\mathbb{E}[V(t,\delta_t,k_t)] = 0 \tag{3.26}$$

For (i), it is easy to show that, under the (candidate to) optimal investment rule, capital takes the form

$$k_t = \left[e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}t}k_0 + \frac{1}{4\psi}\int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}(t-s)}\beta_{3,s}(\delta'_s)^\alpha ds\right]^2,\ t \geq 0,$$

so (i) holds (see the proof of Proposition 3.1 for a proof of the previous expression). In order to show (ii), I will prove that each component of $e^{-rt}\mathbb{E}[V(t,\delta',k)]$ vanishes asymptotically.

Part 1: $\lim_{t\to 0} e^{-rt}\mathbb{E}[\beta_{3,t}(\delta'_t)^\alpha\sqrt{k_t}] = 0$.

Since $\beta_{3,t}$ is uniformly bounded, we can reduce the analysis to studying

$$\lim_{t\to 0} e^{-rt}\mathbb{E}[(\delta'_t)^\alpha\sqrt{k_t}].$$

Regarding the first term in this expression (see the expression of capital above), notice that

$$\begin{aligned}&\lim_{t\to\infty} e^{-\left(r+\frac{\lambda\psi+1-\beta_1}{2\psi}\right)t}\mathbb{E}[(\delta'_t)^\alpha]\\ =\ & e^{-\left(r+\frac{\lambda\psi+1-\beta_1}{2\psi}\right)t}\exp\left(\alpha g t + \frac{\alpha(\alpha-1)}{2}\sigma^2 t\right)(\delta'_0)^\alpha\exp\left(\frac{\alpha(\alpha-1)}{2}(\gamma_0-\gamma_t)\right)\end{aligned}$$

where the last equality comes from Lemma 3.2. But from Lemma (3.3)

$$r + \frac{\lambda\psi + 1 - \beta_1}{2\psi} - \alpha g - \frac{\alpha(\alpha-1)}{2}\sigma^2 > 0,$$

so the limit above is actually zero.

Regarding $\lim_{t\to\infty} e^{-rt}\mathbb{E}\left[(\delta'_t)^\alpha \int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}(t-s)}\beta_{3,s}(\delta'_s)^\alpha ds\right]$, it suffices to analyze

$$\lim_{t\to\infty} \underbrace{e^{-rt}\mathbb{E}\left[\int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}(t-s)}\mathbb{E}[(\delta'_t)^\alpha(\delta'_s)^\alpha]ds\right]}_{:=L_t}.$$

From Lemma 3.2

$$\begin{aligned}\mathbb{E}[(\delta'_t)^\alpha(\delta'_s)^\alpha] &= (\delta'_0)^{2\alpha}\exp\left(\alpha g t + \alpha g s - \frac{\alpha}{2}\int_0^t(\sigma'_u)^2du - \frac{\alpha}{2}\int_0^s(\sigma'_u)^2du\right)\\ &\times \underbrace{\mathbb{E}\left[\exp\left(\alpha\int_0^t(\sigma'_u)dZ'_u + \alpha\int_0^s(\sigma'_u)dZ'_u\right)\right]}_{=\mathbb{E}\left[\exp\left(2\alpha\int_0^s(\sigma'_u)dZ'_u\right)\exp\left(\alpha\int_s^t(\sigma'_u)dZ'_u\right)\right]},\ t\geq 0. \end{aligned} \tag{3.27}$$

But since $\int_0^s(\sigma'_u)dZ'_u$ and $\int_s^t(\sigma'_u)dZ'_u$ are independent and normally distributed random variables, we have that

$$\begin{aligned}\mathbb{E}[(\delta'_t)^\alpha(\delta'_s)^\alpha] &= (\delta'_0)^{2\alpha}\exp\left(\alpha g t + \alpha g s - \frac{\alpha}{2}\int_0^t(\sigma'_u)^2du - \frac{\alpha}{2}\int_0^s(\sigma'_u)^2du\right)\\ &\exp\left(2\alpha^2\int_0^s(\sigma'_u)^2du\right)\exp\left(\frac{\alpha^2}{2}\int_s^t(\sigma'_u)^2du\right),\ t\geq 0. \end{aligned} \tag{3.28}$$

This results in

$$\begin{aligned}L_t &= e^{-\left(r+\frac{\lambda\psi+1-\beta_1}{2\psi}-\alpha g t+\frac{\alpha}{2}\sigma^2\right)t}e^{\frac{\alpha}{2}(\gamma_t-\gamma_0)}\times\\ &\int_0^t e^{-\left(\frac{\lambda\psi+1-\beta_1}{2\psi}-\alpha g t+\frac{\alpha}{2}\sigma^2\right)s}e^{\frac{\alpha}{2}(\gamma_s-\gamma_0)+2\alpha^2\sigma^2 s+2\alpha^2(\gamma_0-\gamma_s)+\frac{\alpha^2}{2}\sigma^2(t-s)+\frac{\alpha^2}{2}(\gamma_s-\gamma_t)}ds.\end{aligned} \tag{3.29}$$

Ignoring all the terms involving γ (it is uniformly bounded), we are left with analyzing

$$\begin{aligned}\tilde{L}_t &= e^{-\left(r+\frac{\lambda\psi+1-\beta_1}{2\psi}-\alpha g t+\frac{\alpha}{2}\sigma^2\right)t}\int_0^t e^{\left(\frac{\lambda\psi+1-\beta_1}{2\psi}+\alpha g-\frac{\alpha}{2}\sigma^2\right)s}e^{2\alpha^2\sigma^2 s+\frac{\alpha^2}{2}\sigma^2(t-s)}ds \\ &= e^{-\left(r+\frac{\lambda\psi+1-\beta_1}{2\psi}-\alpha g t-\frac{\alpha(\alpha-1)}{2}\sigma^2\right)t}\int_0^t e^{\left(\frac{\lambda\psi+1-\beta_1}{2\psi}+\alpha g-\frac{\alpha}{2}\sigma^2(1-3\alpha)\right)s}ds \\ &\propto e^{-(r-2\alpha g-\sigma^2\alpha(2\alpha-1))t}-e^{-(r+\frac{\lambda\psi+1-\beta_1}{2\psi}-\alpha g t-\frac{\alpha(\alpha-1)}{2}\sigma^2)t} \end{aligned} \tag{3.30}$$

and both term go to zero as $t \to \infty$ (Lemma 3.3). This concludes part 1.

□

Part 2: $\lim\limits_{t\to\infty} e^{-rt}\mathbb{E}[\beta_{2,t}(\delta'_t)^{2\alpha}] = 0.$

Observe that since $r > 2\alpha g + \alpha(2\alpha-1)\sigma^2 > 0$, and $(\beta_{3,t})_{t\geq 0}$, $(\gamma_t)_{t\geq 0}$ are uniformly bounded, $(\beta_{2,t})_{t\geq 0}$ will be uniformly bounded as well.

Now, from Lemma 3.2,

$$\mathbb{E}[(\delta'_t)^{2\alpha}] = (\delta'_0)^{2\alpha}\exp\left(\alpha(2\alpha-1)(\gamma_0-\gamma_t)\right)\exp\left(2\alpha g t+\alpha(2\alpha-1)\sigma^2 t\right),$$

and since $r > 2\alpha g + \alpha(2\alpha-1)\sigma^2$, we conclude.

□

Part 3: $\lim\limits_{t\to\infty} e^{-rt}\mathbb{E}[k_t] = 0.$

Recall that

$$\begin{aligned} k_t &= e^{-\frac{\lambda\psi+1-\beta_1}{\psi}t}k_0+\frac{e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}t}}{2\psi}\int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}(t-s)}\beta_{3,s}(\delta'_s)^\alpha ds \\ &\quad+\left(\frac{1}{4\psi}\int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}(t-s)}\beta_{3,s}(\delta'_s)^\alpha ds\right)^2,\ t\geq 0. \end{aligned} \tag{3.31}$$

By definition of β_1,

$$r + \frac{\lambda\psi + 1 - \beta_1}{\psi} = \frac{2}{\psi} + \sqrt{\left(r + \lambda - \frac{1}{\psi}\right)^2 - \frac{1}{\psi^2}} > 0$$

so $\lim_{t\to\infty} e^{-rt} e^{-\frac{\lambda\psi+1-\beta_1}{\psi}t} k_0 = 0$. Regarding the second term, using Lemma 3.2 and that $(\beta_{3,t})_{t\geq 0}$ and $(\gamma_t)_{t\geq 0}$ are both uniformly bounded, it suffices to study the limiting behavior of

$$\begin{aligned} & e^{-\left(r+\frac{\lambda\psi+1-\beta_1}{2\psi}\right)t} \int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}(t-s)} e^{\left(\alpha g + \frac{\alpha(\alpha-1)}{2}\sigma^2\right)s} ds \\ \propto \quad & e^{-\left(r+\frac{\lambda\psi+1-\beta_1}{2\psi} - \alpha g - \frac{\alpha(\alpha-1)}{2}\sigma^2\right)} - e^{-\left(r+\frac{\lambda\psi+1-\beta_1}{\psi}\right)t} \end{aligned} \tag{3.32}$$

and both terms go to zero as $t \to \infty$.

Finally, in order to study

$$L_1 := e^{-rt}\mathbb{E}\left[\left(\int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}(t-s)} \beta_{3,s}(\delta_s')^\alpha ds\right)^2\right]$$

we can apply the Cauchy-Schwartz inequality to obtain

$$L_1 \leq e^{-rt}\left(\int_0^t \mathbb{E}[(\delta_s')^{2\alpha}] ds\right)\left(\int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{\psi}(t-s)} \beta_{3,s}^2 ds\right)$$

Observe that since $(\beta_{3,t})_{t\geq 0}$ is uniformly bounded and

$$\frac{\lambda\psi + 1 - \beta_1}{\psi} = \frac{2}{\psi} - r + \sqrt{\left(r + \lambda - \frac{1}{\psi}\right)^2 - \frac{1}{\psi^2}} > 0$$

the limit $\lim_{t\to\infty} \int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{\psi}(t-s)} \beta_{3,s}^2 ds$ is finite. Thus, it suffices to show that

$$\lim_{t\to\infty} e^{-rt} \int_0^t \mathbb{E}[(\delta_s')^{2\alpha}] ds = 0.$$

From Part 2, it is equivalent to study

$$\lim_{t\to\infty} e^{-rt}\int_0^t e^{-\left(2\alpha g+\alpha(2\alpha-1)\sigma^2\right)s}ds = \lim_{t\to\infty} e^{-\left(r-2\alpha g-\alpha(2\alpha-1)\sigma^2\right)t} - e^{-rt} = 0.$$

This concludes the proof of the Theorem.

$\square$

Proof of Proposition 3.1: Given the investment rule

$$i(t,\delta',k) = \frac{1}{\psi}\left[(\beta_1-1)k + \frac{\beta_{3,t}}{2}(\delta')^\alpha\sqrt{k}\right],$$

the dynamics of capital can be written as

$$dk_t = \left[-\frac{\psi\lambda+1-\beta_1}{\psi}k_t + \frac{\beta_{3,t}}{2}(\delta'_t)^\alpha\sqrt{k_t}\right]dt,\ t\geq 0.$$

This implies that $x_t = \sqrt{k_t}$ satisfy the ODE

$$dx_t = \frac{1}{2}\left[-\frac{\psi\lambda+1-\beta_1}{\psi}x_t + \frac{\beta_{3,t}}{2}(\delta'_t)^\alpha\right]dt,$$

which has as a solution

$$\sqrt{k_t} = x_t = e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}t}k_0 + \frac{1}{4\psi}\int_0^t e^{-\frac{\lambda\psi+1-\beta_1}{2\psi}(t-s)}\beta_{3,s}(\delta'_s)^\alpha ds,\ t\geq 0.$$

The rest of the results are straightforward.

$\square$

Bibliography

[1] Abel, A. (1983): "Optimal Investment Under Uncertainty," *American Economic Review, Vol. 73, No. 1, pp. 228-233.*

[2] Abel, A. (1985): "A Stochastic Model of Investment, Marignal q and the Market Value of the Firm," *International Economic Review, Vol. 84, No. 5, pp. 1369-1384.*

[3] Abel, A. and J. Eberly (1994): "Optimal Investment with Costly Reversibility," *The Review of Economic Studies, Vol. 26, No. 2, pp. 302-325.*

[4] Abel, A. and J. Eberly (1997): "An Exact Solution for the Investment and Value of a Firm Facing Uncertainty, Adjustment Costs and Irreversibility," *Journal of Economic Dynamics and Control, Vol. 21, pp. 831-852.*

[5] Abel, A. and J. Eberly (1999): "The Effects of Irreversibility and Uncertainty on Capital Accumulation," *Journal of Monetary Economics, Vol. 44, pp. 339-377.*

[6] Abowd, J. and Card, D. 1989: "On the Covariance Structure of Earnings and Hours Changes," *Econometrica, Vol. 57, No 2, pp. 411-445.*

[7] Arrow, K. (1968): "Optimal Capital Policy with Irreversible Investment," in J.N. Wolfe, ed., Value Capital and Growth, papers in honor of Sir John Hicks, Edinburgh University Press, Edinburgh, pp. 1-19.

[8] Baker, G. Gibbs, M. and Holmstrom, B. 1994: "The Internal Economics of the Firm: Evidence from Personnel Data," *The Quarterly Journal of Economics, Vol. 109, No. 4, pp. 921-55.*

[9] Baker, G. Gibbs, M. and Holmstrom, B. 1994: "The Wage Policy of a Firm," *The Quarterly Journal of Economics, Vol. 109, No. 4, pp. 881-919.*

[10] Barro, R. and D. Gordon (1983): "Rules, discretion and reputation in a model of monetary policy," *Journal of Monetary Economics, Vol. 12, No. 1, pp. 101-121.*

[11] Becker, G. 1964: *Human Capital; a theoretical and empirical analysis, with special reference to education.* New York: Columbia University Press.

[12] Ben-Porath, Y. 1967: "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy, Vol. 75, No. 4, Part 1, pp. 352-365.*

[13] Bergemann, D., and U. Hege (2005): "The Financing of Innovation: Learning and Stopping," *RAND Journal of Economics, Vol. 36, No. 4, pp. 719-752.*

[14] Bertola, G. and R. Caballero (1994): "Irreversibility and Aggregate Investment" *The Review of Economic Studies, Vol. 61, No. 2, pp. 223-246*

[15] Board, S., and M. Meyer-ter-Vehn (2010a): "A Reputation Theory of Firm Dynamics," Working Paper, UCLA.

[16] ——. (2010b): "Reputation for Quality," Working Paper, UCLA.

[17] Bohren, A. (2012): "Stochastic Games in Continuous Time: Persistent Actions in Long-term Relationships," Working Paper, UC San Diego.

[18] Bonatti, A., and J. Hörner (2011): "Career Concerns with Coarse Information," Cowles Foundation Discussion Paper 1831, Cowles Foundation for Research in Economics, Yale University, revised Jan. 2012.

[19] Cisternas, G. (2012): "Shock Persistence, Endogenous Skills and Career Concerns," Working Paper, Princeton.

[20] Cukierman, A., and A. Meltzer (1986): "A Theory of Ambiguity, Credibility and Inflation Under Discretion and Asymmetric Information," *Econometrica, Vol. 54, No. 5, pp.1099-1128.*

[21] Crandall, M., Ishii, H., and P.-L. Lions (1992): "User's Guide to Viscosity Solutions of Second Order Partial Differential Equations," *Bulletin of the American Mathematical Society, Vol. 27, pp. 1-67.*

[22] Dewatripont, M., I. Jewitt, and J. Tirole (1999a): "The Economics of Career Concerns, Part I: Comparing Information Structures," *Review of Economic Studies, Vol. 66, No. 1, pp. 199-217.*

[23] ——. (1999b): "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies," *Review of Economic Studies, Vol. 66, No. 1, pp. 199-217.*

[24] Dixit, A., and R. Pindyck (1994): *Investment Under Uncertainty*. New Jersey: Princeton University Press.

[25] Eisner, R. and R. Strotz (1963): "Determinants of Business Investment," in Commission on Money and Credit, *Impacts of monetary policy*, Englewoods Cliffs, NJ: Prentice Hall, pp. 59-337.

[26] Faingold, E., and Y. Sannikov, (2010): "Reputation in Continuous-Time Games," *Econometrica, Vol. 79, No. 3, pp. 773-876.*

[27] Fama, E., 1980: "Agency Problem and the Theory of the Firm". *Journal of Political Economy, 88, No. 2, pp. 288-307.*

[28] Farber, H. and Gibbons, R. 1996: "Learning and Wage Dynamics," *The Quarterly Journal of Economics, Vol. 111, No. 4, pp. 1007-1047.*

[29] Fleming, W., and R. Rishel (1975): *Deterministic and Stochastic Optimal Control*. New York: Springer-Verlag.

[30] Fleming, W., and M. Soner (2006): *Controlled Markov Processes and Viscosity Solutions*. New York: Springer.

[31] Fudenberg, D., and J. Tirole (1986): "A "Signal-Jamming" Theory of Predation," *RAND Journal of Economics, The RAND Corporation, Vol. 17, No. 3, pp. 366-376.*

[32] Gibbons, R. and Murphy, K., 1992: "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *The Journal of Political Economy, Vol. 100, No. 3, pp. 468-505*

[33] Gibbons, R. and Waldman, M. 2004: "Task-Specific Human Capital," *The American Economic Review, Vol. 94, No. 2, pp. 203-207.*

[34] ——. 2006. "Enriching a Theory of Wage and Promotion Dynamics within Firms," *Journal of Labor Economics, Vol. 24, No. 1, pp. 59-107.*

[35] Harris, M. and B. Holmstrom (1982): "A Theory of Wage Dynamics," *The Review of Economic Studies, Vol. 49, No. 3, pp. 315-333.*

[36] Hayashi, F. (1982): "Tobin's Marginal q and Average q: A Neoclassical Interpretation," *Econometrica, Vol. 50, No. 1, pp. 213-224.*

[37] Holmstrom, B. (1979): "Moral Hazard and Observability," *The Bell Journal of Economics, Vol 10, No. 2, pp. 74-91.*

[38] ——. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *The Review of Economic Studies, 66, pp. 169-182.*

[39] Holmstrom, B. and P. Milgrom (1991): "Multi-Task Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design". *Journal of Law, Economics and Organizations, Vol. 7, pp. 24-52.*

[40] Jorgenson, D. (1963): "Capital Theory and Investment Behavior," *American Economic Review, Vol. 53, No. 2, pp. 247-259.*

[41] Jovanovic, B., 1979: "Job Matching and the Theory of Turnover," *The Journal of Political Economy, Vol. 87, No. 5, Part 1, pp. 972-990.*

[42] Kahn, L. and F. Lange (2011): "Employer Learning, Productivity and the Earnings Distribution: Evidence from Performance Measures". IZA Discussion Papers 5054, Institute for the Study of Labor (IZA).

[43] Karlin, S., and H. Taylor (1981): *A Second Course in Stochastic Processes.* London: Academic Press.

[44] Kushner, H., and P. Dupuis (2001): *Numerical Methods for Stochastic Control Problems in Continuous Time.* New York: Springer-Verlag.

[45] Kydland, F. and E. Prescott (1977): "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy, University of Chicago Press, Vol. 85, No. 3, pp 473-491.*

[46] Kyle, A. (1985): "Continuous Auctions and Insider Trading," *Econometrica, Vol. 53, No. 6, pp. 1315-1335.*

[47] Laffont, J.J. and J. Tirole (1988): "The Dynamics of Incentive Contracts," *Econometrica, Vol. 56, No. 5, pp. 1153-1175.*

[48] Laffont, J.J. and J. Tirole (1990): "Adverse Selection and Renegotiation in Procurement," *Review of Economic Studies, Vol. 57, No. 4, pp. 597-625.*

[49] Lazear, E., (1995): *Personnel Economics*. Cambridge, Massachussets: The MIT Press.

[50] Lewellen, J. and J. Shanken (2002): "Learning, Asset-Pricing and Market Efficiency," *The Journal of Finance, Vol. 42, No. 3, pp. 1113-1145.*

[51] Liptser, R., and A. Shiryaev (1977): *Statistics of Random Processes I and II.* New York: Springer-Verlag.

[52] Lucas, R. (1975): "An Equilibrium Model of the Business Cycle," *Journal of Political Economy, Vol. 83, No. 6, pp. 1113144.*

[53] Martinez, L. (2009): "Reputation, Career Concerns, and Job Assignments," *The B.E. Journal of Theoretical Economics, Berkeley Electronic Press, Vol. 9, No. 1. (Contributions), Article 15.*

[54] Mincer, J., (1974): *Experience, Schooling and Earnings*. New York: National Bureau of Economic Research.

[55] Mussa, M. (1977): "External and Internal Adjustment Costs and the Theory of Aggregate and Firm Investment," *Econometrica, Vol. 44, No. 174, pp. 163-178.*

[56] Pham, H. (2009): *Continuous-time Stochastic Control and Optimization with Financial Applications*. Berlin: Springer.

[57] Prendergast, C. (1998): "What Happens within Firms? A Survey of Empirical Evidence on Compensation Policies." *National Bureau of Economic Research Working Paper Series No. 5802, pp. 329-356.*

[58] Pyndick, R. (1982): "Adjustment Costs, Uncertainty and the Behavior of the Firm," *The American Economic Review, Vol. 72, No. 3, pp. 415-427.*

[59] Rosen, S., (1972): "Learning and Experience in the Labor Market," *The Journal of Human Resources, Vol. 7, No. 3, pp. 326-342.*

[60] Sannikov, Y. (2007): "Games with Imperfectly Observable Actions in Continuous Time," *Econometrica, Vol. 75, No. 5, pp. 1285329.*

[61] Sannikov, Y. (2008): "A Continuous-Time Version of The Principal-Agent Model," *The Review of Economic Studies, 75, pp. 957-984.*

[62] Scharfstein, D. and J. Stein, (1990): "Herd Behavior and Investment," *The American Economic Review, Vol. 80, No. 3, pp. 465-479.*

[63] Strulovici, B., and M. Szydlowski (2012): "On the Smoothness of Value Functions and the Existence of Optimal Strategies," Working Paper, Northwestern University.

[64] Tobin, J. (1969): "A General Equilibrium Approach to Monetary Theory," *Journal of Money, Credit and Banking, Vol. 1, No. 1, pp. 15-29.*

[65] Williams, N. (2011): "Persistent Private Information," *Econometrica, Vol. 79, No. 4, pp. 1233-1275.*

[66] Yong, J. and X. Zhou (1999): *Stochastic Controls*. New York: Springer-Verlag.

[67] Wonham, W. (1968): "On the Separation Theorem of Stochastic Control," *SIAM Journal on Control and Optimization, Vol. 6, No. 2, pp: 312-326.*

[68] Wonham, W. (1985): *Linear Multivariate Control: A Geometric Approach.* Berlin: Springer-Verlag.