

Two-Sided Learning and the Ratchet Principle

Gonzalo Cisternas*

MIT Sloan

May 2017

Abstract

I study a class of continuous-time games of learning and imperfect monitoring. A long-run player and a market share a common prior about the initial value of a Gaussian hidden state, and learn about its subsequent values by observing a noisy public signal. The long-run player can nevertheless control the evolution of this signal, and thus affect the market's belief. The public signal has an additive structure, and noise is Brownian. I derive conditions for a solution to an ordinary differential equation to characterize behavior in which the long-run player's equilibrium actions depends on the history of the game only through the market's correct belief. Using these conditions, I demonstrate the existence of equilibria in pure strategies for settings in which the long-run player's flow utility is nonlinear. The central finding is a learning-driven ratchet principle affecting incentives. I illustrate the economic implications of this principle in applications to monetary policy, earnings management, and career concerns.

Keywords: learning, private beliefs, ratchet effect, Brownian motion.

JEL codes: C73, D82, D83.

1 Introduction

Hidden variables are at the center of many economic interactions: firms' true fundamentals are hidden to both managers and shareholders; workers' innate abilities are unobserved by both employers and workers themselves; and growth and inflation trends are hidden to both

*Email: ggeistern@mit.edu. Earlier versions of this paper were circulated under the title "Two-Sided Learning and Moral Hazard." I would like to thank Yuliy Sannikov for his invaluable advice, and Dilip Abreu, Alessandro Bonatti, Hector Chade, Eduardo Faingold, Bob Gibbons, Leandro Gorno, Tibor Heumann, Andrey Malenko, Iván Marinovic, Stephen Morris, Marcin Peski, Juuso Toikka, Larry Samuelson, Mike Whinston and audiences at Columbia, Harvard-MIT, MIT Sloan, NYU Stern, Stanford GSB, Toulouse School of Economics, UCLA, UCSD and the University of Minnesota for their feedback. Also, I would like to thank three anonymous referees for very valuable suggestions that helped improve the paper.

policymakers and market participants. In those settings, economic agents face common uncertainty regarding payoff-relevant states that underlie the economic environment, and eliminating such uncertainty can be prohibitively costly, or simply impossible; agents thus *learn* about such states simultaneously as decisions are being made, and the incomplete information they face need not ever disappear. This paper is concerned with examining strategic behavior in settings characterized by such forms of fundamental uncertainty.

When agents learn about economic environment, behavior can be influenced the possibility of affecting the beliefs of others. The set of questions that can be asked in such contexts is incredibly rich. In financial markets, is it possible for markets to hold correct beliefs about firm's fundamentals in the presence of earnings management? In labor markets, what are the forces that shape workers' incentives when they want to be perceived as highly skilled? In policy, how is a central bank's behavior shaped by the possibility of affecting markets' beliefs about the future evolution of inflation? The challenge in answering these questions lies on developing a framework that is tractable enough to accommodate both Bayesian updating to capture ongoing learning, and imperfect monitoring to capture strategic behavior.

To make progress towards the understanding of games of learning and imperfectly observable actions, I employ continuous-time methods using Holmström's (1999) *signal-jamming* technology as the key building block. In the setting I study, there is a long-run player and a market (i.e., a population of small individuals) who, starting from a common prior, learn about an unobserved Gaussian *fundamentals* process by observing a public signal. The long-run player can nevertheless influence the market's belief about the fundamentals by taking unobserved actions that affect the evolution of the publicly observed state. As in Holmström (1999), actions and the fundamentals are perfect substitutes in the signal technology, and thus the long-run player cannot affect the informativeness of the public signal (i.e., there is no experimentation). Using Brownian information, I study Markov equilibria in which the long-run player's behavior depends on the history through the belief about the hidden state.

In an equilibrium in pure strategies, the market must anticipate the long-run player's actions at all times; beliefs thus coincide on the equilibrium path. However, allowing for belief divergence is critical to determine the actions that arise along the path of play. Consider, for instance, the earnings management example. To show that an equilibrium in which the market holds a correct belief exists, it must be verified that the payoff that the manager obtains by reporting earnings as conjectured by the market dominates the payoff under any other strategy. But if the manager deviates, the market will misinterpret the report; at those *off-path histories*, both parties' beliefs about the firm's fundamentals differ.

Crucially, when actions are hidden, deviations from the market's conjectured behavior lead the long-run player's belief to become private. Moreover, this private information is

persistent, as it comes from a learning process. As I will explain shortly, the combination of hidden actions and private information off the path of play severely complicates the equilibrium analysis in virtually every setting that allows for learning and imperfect monitoring with frequent arrival of information.¹

To address this difficulty, I follow a first-order approach to studying Markov equilibria in settings where (i) affecting the public signal is costly and (ii) the long-run player’s flow payoff is a general—in particular, nonlinear—function of the market’s belief. Specifically, I construct a necessary condition for equilibria in which on-path behavior is a differentiable function of the common belief, and then provide conditions under which this necessary condition is also sufficient. The advantages of this approach are both conceptual and technical. First, the necessary condition uncovers the forces that shape the long-run player’s behavior in any Markov equilibrium, provided that an equilibrium of this form exists. Second, this approach offers a tractable venue for demonstrating the existence of such equilibria despite the intricacies of off-path private beliefs affecting behavior.

Economic contribution. The main finding of this paper pertains to a *ratchet principle* affecting incentives. Consider a manager who evaluates boosting a firm’s earnings report above analysts’ predictions. The immediate benefit from this action is clear: abnormally high earnings lead the market to believe that the firm’s fundamentals have improved. Crucially, the manager understands that this optimism is incorrect, as the observation of high earnings was a consequence of altering the report. He then anticipates that subsequent manipulation will be required to maintain the impact on the firm’s value, as his private belief about the firm’s fundamentals indicates that the firm would otherwise underperform relative to the market’s expectations. Equally important, if the market expects firms with better prospects to manage their earnings more aggressively, this underperformance can become even more acute. In either case, exhibiting good performance results in a more demanding incentive scheme to be faced tomorrow—i.e., a learning-driven ratchet principle emerges.²

In this paper, ratchet effects—implications on behavior of the ratchet principle just described—do not relate to reduced incentives for information revelation, as in models with ex ante asymmetric information (e.g., Laffont and Tirole, 1988): this is because the long-run player is unable to affect the informativeness of the public signal, which implies that the speed of learning is exogenous. Instead, these effects are captured in the form of distorted levels of costly actions relative to some benchmarks. More generally, their appearance is

¹Holmström’s original setting is unique in this respect, as the linearity in payoffs assumed in his model makes incentives independent of the value that beliefs may take.

²Weitzman (1980) refers to the ratchet principle as the “tendency of planners to use current performance as a criterion in determining future goals” (p.302). In Section 3 I show how a market revising its expectations about future values of a public signal is in fact a target revision from the long-run player’s perspective.

the outcome of a fundamental tension between Bayesian updating and strategic behavior, and hence, they are not exclusive to the case of a Gaussian hidden state. Specifically, since beliefs are revised based on discrepancies between observed and expected signal realizations, actions that lead to abnormally high signals are inherently costly from a dynamic perspective: by creating higher expectations for tomorrow's signals, such actions require stronger future actions to generate a sustained effect on beliefs.

Applications. I first revisit Holmström's (1999) seminal model of career concerns, which is a particular instance of linear payoffs within the class of games analyzed. In this context, I show that the form of ratcheting previously described is embedded in the equilibrium that he finds. Importantly, by precisely quantifying the strength of this force, I show how off-path private information—via a ratcheting mechanism—plays an important role in limiting the power of market-based incentives in the equilibrium found by Holmström.

A key advantage of this paper is its ability to accommodate nonlinear flow payoffs, which can be a defining feature of many economic environments. In an application to monetary policy, I consider a setting in which a price index carries noisy information about both an unobserved inflation trend and the level of money supply, and a central bank can affect employment by creating inflation surprises. The central bank's trade-off between output and inflation is modeled via a loss function that is quadratic in employment (or output) and money growth. In such a context, I show that the ratchet principle can induce a monetary authority to exhibit a stronger commitment to low inflation. Intuitively, while unanticipated inflation can be used to boost employment in the short run, it also leads the market to overestimate future inflation and, hence, to set excessively high nominal wages. This in turn puts downward pressure on future hiring decisions, which makes inflation more costly compared to settings in which the inflation trend is observed or simply absent.

Finally, I study more subtle ratchet effects in an application that analyzes managers' incentives to boost earnings when they have a strong short-term incentive to exceed a zero-earnings threshold, captured in marginal flow payoffs that are single-peaked and symmetric around that point. In such a context, I show that firms that expect to generate positive earnings can inflate reports more actively than firms at, or below, the threshold, despite their managers having weaker myopic incentives and being unable to affect firms' market values. Intuitively, the market anticipates that successful manipulation by firms with poor (good) past performance will lead to stronger (weaker) myopic incentives in the future. Anticipating higher expectations of earnings management by the market, firms with poor profitability find it more costly to inflate earnings relative their successful counterparts. The distortion thus takes the form of a profile of manipulation that is skewed towards firms that have exhibited better performances in the past.

Technical contribution. In the class of games analyzed, learning is conditionally Gaussian and stationary, and hence, beliefs can be identified with posterior means. Moreover, a nonlinear version of the Kalman filter applies. It is then natural to look for Markov perfect equilibria (MPE) using standard dynamic programming tools, with the market and long-run player’s beliefs as states. However, the combination of hidden actions and hidden information off the path of play results in the long-run player’s value function no longer satisfying a traditional Hamilton-Jacobi-Bellman (HJB) equation. In fact, the differential equation at hand does not even have the structure of a usual partial differential equation (PDE); to the best of my knowledge, no existence theory applies.

Implicit in the HJB approach is that, by demanding the determination of the long-run player’s full value function, the method requires exact knowledge of the long-run player’s off-path behavior to determine the actions that arise along the path of play; however, the difficulty at hand is precisely that the long-run player can condition his actions on his private information in complex ways as his own belief changes. Exceptions are settings in which the long-run player’s flow payoff is linear in the market’s belief (e.g., Holmström, 1999), as in those cases the long-run player’s optimal behavior is independent of the past history of play. However, it is exactly in those linear environments that the differential equation delivered by the HJB approach has a trivial solution. If the goal is then to analyze settings that naturally involve nonlinearities, solution methods for linear environments do not apply.

The technical advantage of the first-order approach is that the *ratcheting equation*—the necessary condition for equilibrium behavior—makes bypassing the exact computation of off-path payoffs possible. In fact, this ordinary differential equation (ODE) offers a method to guess for Markov equilibria without knowing how exactly the candidate equilibrium might be supported off the path of play. Importantly, provided that it is verified that a deviation from a solution to the ratcheting equation is not profitable, leaving off-path behavior unspecified in the equilibrium concept is no disadvantage: equilibrium outcomes (i.e., actions and payoffs) are determined exclusively by the actions prescribed by strategies along the path of play.

Therefore, for sufficiency, instead of computing off-path payoffs exactly, I approximate them. Specifically, building on the optimal contracting literature, I bound off-path payoffs in a way that parallels sufficiency steps in relaxed formulations of principal-agent problems (Williams, 2011; Sannikov, 2014) to derive a *verification theorem* for Markov equilibria (Theorem 1). The theorem involves the ratcheting equation and the ODE that characterizes the evolution of the (candidate, on-path) payoff that results from inducing no belief divergence, i.e., two ODEs rather than a PDE. The key requirement is that the *information rent*—a measure of the value of acquiring private information about the continuation game—associated with the solution of the system at hand cannot change too quickly.

The advantage of this verification theorem—relative to both the HJB approach and the contracting literature—is its tractability. Using this result, I determine conditions on primitives that ensure the existence of Markov equilibria in two classes of games exhibiting nonlinearities: linear quadratic games and games with bounded marginal flow payoffs (Theorems 2 and 3), which host the applications I examine. These three results address the belief divergence challenge, and the continuous-time approach is critical for their derivation.

Related literature. Regarding the literature on the ratchet effect, Weitzman (1980) illustrates how revising production targets on the basis of observed performance can dampen incentives in planning economies; both the incentive scheme and the revision rule are exogenous in his analysis. Freixas et al. (1985) and Laffont and Tirole (1988) in turn endogenize ratcheting by allowing a principal to optimally revise an incentive scheme as new information about an agent’s hidden type is revealed upon observing performance; the main result is that there is considerable pooling. As in Weitzman (1980), my analysis focuses on the size of equilibrium actions, rather than on their informativeness. In line with the second group of papers, the strength of the ratcheting that arises in any specific setting is an equilibrium object: by conjecturing the long-run player’s behavior, the market effectively imposes an endogenous moving target against which the long-run player’s performance is evaluated.

Concurrently with this paper, Bhaskar (2014), Prat and Jovanovic (2014), and Bhaskar and Mailath (2016) identify ratchet principles in principal-agent models with symmetric uncertainty: namely, that good performance can negatively affect an agent’s incentives if it leads a principal to overestimate a hidden technological parameter. My analysis differs from these papers along two dimensions. First, I show that market-based incentives can lead to quite rich behavior on behalf of a forward-looking agent; instead, the contracts that these papers analyze implement either minimal or maximal effort. Second, I show that, in games of symmetric uncertainty, the ratchet principle is also determined by a market revising its expectations of future behavior, in addition to revising its beliefs about an unobserved state.³

This paper belongs to a broader class of games of ex ante symmetric uncertainty in which imperfect monitoring leads to the possibility of divergent beliefs. In the reputation literature, Holmström (1999) finds an equilibrium in which a worker’s equilibrium effort is identical on and off the path of play, in part consequence of the assumed linearity in payoffs.⁴ In Board and Meyer-ter-Vehn (2014), private beliefs matter non-trivially for a firm’s

³Also in the context of symmetric uncertainty, Meyer and Vickers (1997) study a model of regulation in which ratcheting is modeled explicitly via an exogenous incentive scheme that reduces payments to more efficient firms. Martinez (2009) instead identifies the potential appearance of endogenous ratchet-like forces in a model of career concerns with piecewise linear wages.

⁴Kovrijnykh (2007), Martinez (2006, 2009) and Bar-Isaac and Deb (2014) study nonlinearities in models of career concerns with finite horizon. Except in the two-period model of Bar-Isaac and Deb (2014), where

investment policy, and the existence of an equilibrium is shown via fixed-point arguments; my approach is instead constructive and focused on pure strategies. Private beliefs also arise in strategic experimentation settings involving a risky arm of two possible types and perfectly informative Poisson signals. Since beliefs are deterministic in this case, the equilibrium analysis is tractable (Bergemann and Hege (2005) derive homogeneity properties of off-path payoffs and Bonatti and Hörner (2011, 2016) apply standard optimal control techniques), and the ratcheting I find is absent, as the observation of a signal terminates the interaction.

To conclude, this paper contributes to a growing literature that analyzes dynamic incentives exploiting the tractability of continuous-time methods. Sannikov (2007), Faingold and Sannikov (2011) and Bohren (2016) study games with imperfect monitoring in which the continuation game is identical on and off the equilibrium path. In contrast, as in the current paper, in the principal-agent models of Williams (2011), Prat and Jovanovic (2014), and Sannikov (2014), deviations lead the agent to obtain private information about future output. All these contracting papers derive measures of information rents and general sufficient conditions that validate the first-order approach they follow. Such sufficient conditions involve endogenous variables, and their verification is usually done both *ex post* (i.e., using the solution to the relaxed problem) and in specific settings. The sufficient conditions that I derive can be instead mapped to primitives for a large class of economic environments.

1.1 Outline

Section 2 presents the model and Section 3 derives necessary conditions for Markov equilibria. Section 4 explores applications. Section 5 states the verification theorem, and Section 6 contains the existence results. Section 7 concludes. All proofs are relegated to the Appendix.

2 Model

A long-run player and a population of small players (the *market*) learn about a hidden state $(\theta_t)_{t \geq 0}$ (the *fundamentals*) by observing a public signal $(\xi_t)_{t \geq 0}$. Their evolution is given by

$$d\theta_t = -\kappa(\theta_t - \eta)dt + \sigma_\theta dZ_t^\theta, \quad t > 0, \quad \theta_0 \in \mathbb{R}, \quad (1)$$

$$d\xi_t = (a_t + \theta_t)dt + \sigma_\xi dZ_t^\xi, \quad t > 0, \quad \xi_0 = 0. \quad (2)$$

In this specification $(Z_t^\theta)_{t \geq 0}$ and $(Z_t^\xi)_{t \geq 0}$ are independent Brownian motions, and σ_θ and σ_ξ are strictly positive volatility parameters. The fundamentals follow a Gaussian diffusion

sufficiency reduces to static second-order conditions, the question of existence of equilibria is not addressed.

(hence, Markov) process where $\kappa \geq 0$ is the rate at which $(\theta_t)_{t \geq 0}$ reverts towards the long-run mean $\eta \in \mathbb{R}$.⁵ The public signal (2) carries information about the fundamentals in its drift, but it is affected by the long-run player's choice of action a_t , $t \geq 0$. These actions take values in an interval $A \subseteq \mathbb{R}$, with $0 \in A$, and they are never directly observed by the market.

The monitoring technology (2) is the continuous-time analog of Holmström's (1999) *signal-jamming* technology, and a key property of it is that it satisfies the full-support assumption with respect to the long-run player's actions.⁶ Thus, the only information that the market has comes from realizations of $(\xi_t)_{t \geq 0}$; let $(\mathcal{F}_t)_{t \geq 0}$ denote the corresponding public filtration, and $\xi^t := (\xi_s : 0 \leq s \leq t)$ any realized public history.

I will examine equilibria in pure strategies in which the long-run player's behavior *along the path of play* is, at all instants of time, a function of the current public history ξ^t , $t \geq 0$. The formal notion of a pure public strategy for the long-run player is defined next; I refer to any such pure public strategy simply as a *strategy* thereafter.

Definition 1. A (pure public) strategy $(a_t)_{t \geq 0}$ is a stochastic process taking values in A that is also progressively measurable with respect to $(\mathcal{F}_t)_{t \geq 0}$, and that satisfies $\mathbb{E} \left[\int_0^t a_s^2 ds \right] < \infty$, $t \geq 0$. A strategy is feasible if, in addition, (2) admits a unique (in probability law) solution.⁷

Everyone shares a prior that θ_0 is normally distributed, with a variance γ^* that ensures that learning is stationary—in this case, the Gaussian structure of both the fundamentals and noise permits posterior beliefs to be identified with posterior means; I defer the details to Section 3.1. Crucially, in order to interpret the public signal correctly, the market needs to *conjecture* the long-run player's equilibrium behavior; in this way, the market can account for how the latter agent's actions affect the evolution of the public signal. Thus, let

$$p_t^* := \mathbb{E}^{a^*} [\theta_t | \mathcal{F}_t]$$

denote the mean of the market's posterior belief about θ_t given the information up to time

⁵When $\kappa = 0$, $(\theta_t)_{t \geq 0}$ corresponds to a Brownian martingale. In the $\kappa \neq 0$ case, this process is usually referred to as an *Ornstein-Uhlenbeck* (or mean-reverting) process.

⁶This is a consequence of Girsanov's theorem, which states that changing the drift in the public signal induces an equivalent distribution over the set of paths of $(\xi_t)_{t \geq 0}$.

⁷Formally, the game takes place in the following filtered probability space $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ (for reference $C(E)$ denotes the set of continuous functions from $E \subseteq \mathbb{R}$ to \mathbb{R}): (i) $\Omega = C(\mathbb{R}_+)$ is the set of sample paths of $(\xi_t)_{t \geq 0}$; (ii) \mathcal{F}_t is the canonical σ -algebra on $C([0, t])$; and (iii) \mathbb{P} is the probability measure on $C(\mathbb{R}_+)$ induced by the long-run player's equilibrium actions via (2). The solution concept for (2) is in a *weak* sense: given the sample space considered, this amounts to the existence of a probability measure on $C(\mathbb{R}_+)$ that is consistent with (2) under $(a_t)_{t \geq 0}$; the uniqueness requirement on such a measure in turn ensures that the outcome of the game is uniquely defined. A strategy is thus a function $a : \mathbb{R}_+ \times C(\mathbb{R}_+) \rightarrow A$, i.e., a mapping connecting (t, ξ) -pairs with actions. Progressive measurability implies that $a_t(\xi)$ depends only on ξ^t , $t \geq 0$, (i.e., $(a_t)_{t \geq 0}$ is *adapted* to $(\mathcal{F}_t)_{t \geq 0}$). Finally, the integrability condition suffices for standard filtering equations to hold.

$t \geq 0$ under the assumption that the feasible strategy $(a_t^*)_{t \geq 0}$ is being followed. In what follows, the market's conjecture $(a_t^*)_{t \geq 0}$ is fixed, and I refer to the corresponding posterior mean process $(p_t^*)_{t \geq 0}$ as the *public belief* process.

The market behaves myopically given its beliefs about the fundamentals and equilibrium play.⁸ Specifically, there is a measurable function $\chi : \mathbb{R} \times A \rightarrow \mathbb{R}$ such that, at each time t , the market takes an action $\chi(p_t^*, a_t^*)$ that affects the long-run player's utility. As a result, the total payoff to the long-run player of following a feasible strategy $(a_t)_{t \geq 0}$ is given by

$$\mathbb{E}^a \left[\int_0^\infty e^{-rt} (u(\chi(p_t^*, a_t^*)) - g(a_t)) dt \mid p_0 = p \right], \quad (3)$$

where $p_0 = p$ denotes the prior mean of θ_0 . In this specification, the notation $\mathbb{E}^a[\cdot]$ emphasizes that a strategy $(a_t)_{t \geq 0}$ induces a distribution over the paths of $(\xi_t)_{t \geq 0}$, thus affecting the likelihood of any realization of $(p_t^*)_{t \geq 0}$. Also, $u : \mathbb{R} \rightarrow \mathbb{R}$ is measurable, and $r > 0$ denotes the discount rate. Finally, affecting the public signal is costly according to a convex function $g : A \rightarrow \mathbb{R}_+$ such that $g(0) = 0$, $g'(a) > 0$ for $a > 0$, $g'(a) < 0$ for $a < 0$ (i.e., increasing the rate of change of the public signal in either direction is costly at increasing rates).

Mild technical conditions on u , χ and g that are used for studying equilibria characterized by ODEs are presented next—these conditions are not needed for examining pure-strategy equilibria at a general level (Definition 2 below), and they are discussed at the end of this section (Remark 1). Let $C^k(E; F)$ be the set of k -times differentiable functions from $E \subset \mathbb{R}^n$ to $F \subset \mathbb{R}$, $n \geq 1$, with a continuous k -th derivative; I omit k if $k = 0$, and F if $F = \mathbb{R}$.

Assumption 1. (i) *Differentiability:* $u \in C^1(\mathbb{R})$, $\chi \in C^1(\mathbb{R} \times A)$ and $g \in C^2(A; \mathbb{R}_+)$ with

$$\rho := (g')^{-1} \in C^2(\mathbb{R}).$$

(ii) *Growth conditions:* the partial derivatives χ_p and χ_{a^*} are bounded in $\mathbb{R} \times A$, and u , u' and g' have polynomial growth.⁹ (iii) *Strong convexity:* $g''(\cdot) \geq \psi$ for some $\psi > 0$.¹⁰

As is standard in stochastic optimal control, a strategy $(a_t)_{t \geq 0}$ is *admissible* for the long-run player if it is feasible and

$$\mathbb{E}^a \left[\int_0^\infty e^{-rt} |u(\chi(p_t^*, a_t^*)) - g(a_t)| dt \mid p_0 = p \right] < \infty,$$

⁸The market can correspond to a sequence of short-run players, or a continuum of identical forward-looking agents who only maximize ex ante flow payoffs over $[t, t + dt)$. The latter can occur if, in the (unmodeled) game played amongst them, each agent is unable to affect any payoff-relevant state.

⁹ $f : \mathbb{R} \rightarrow \mathbb{R}$ is said to have polynomial growth if there is $C > 0$ and $j \in \mathbb{N}$ such that $|f(p)| \leq C(1 + |p|^j)$ for all $p \in \mathbb{R}$. When $j = 2$ ($j = 1$) it is said that f has quadratic (linear) growth.

¹⁰A quadratic cost function satisfies all the conditions on $g(\cdot)$.

(cf., Pham, 2009). In this case, it is said that $(a_t, a_t^*)_{t \geq 0}$ is an admissible pair.

Definition 2. A strategy $(a_t^*)_{t \geq 0}$ is a pure-strategy Nash equilibrium (NE) if $(a_t^*, a_t^*)_{t \geq 0}$ is an admissible pair and

- (i) $(a_t^*)_{t \geq 0}$ maximizes (3) among all strategies $(a_t)_{t \geq 0}$ such that $(a_{t \geq 0}, a_t^*)_{t \geq 0}$ is an admissible pair, and
- (ii) $(p_t^*)_{t \geq 0}$ is constructed via Bayes' rule using $(a_t^*)_{t \geq 0}$.

In a (pure-strategy) NE, the long-run player finds it optimal to follow the market's conjecture of equilibrium play while the market is simultaneously using the same strategy to construct its belief. Thus, along the path of play, (i) the long-run player's behavior is sequentially rational, and (ii) the long-run player and the market hold the same belief at all times. Allowing for belief divergence is, nevertheless, a critical step towards the determination of the actions that arise along the path of play, and at those off-path histories the long-run player can condition his actions on more information than that provided by the public signal; Sections 3 and 5 are devoted to this equilibrium analysis. It is important to stress, however, that for the analysis of equilibrium outcomes (i.e., actions and payoffs), leaving behavior after deviations unspecified in the equilibrium concept is without loss, as the full-support monitoring structure (2) makes this game one of unobserved actions.¹¹

The focus is on equilibria that are *Markov* in the public belief with the property that actions are interior, and the corresponding policy (i.e., the mapping between beliefs and actions) and payoffs exhibiting enough differentiability, as defined next:

Definition 3. An equilibrium is Markov if there is $a^* \in C^2(\mathbb{R}; \text{int}(A))$ Lipschitz such that $(a^*(p_t^*))_{t \geq 0}$ is a NE, and (3) under $a_t = a^*(p_t^*)$, $t \geq 0$, is of class $C^2(\mathbb{R})$ as a function of $p \in \mathbb{R}$.

In a Markov equilibrium, behavior depends on the public history only through the common belief according to a sufficiently differentiable function—such equilibria are natural to analyze due to both the Markovian nature of the fundamentals and the presence of Brownian noise. Importantly, the long-run player's realized actions are, at all time instants, a function of the complete current public history ξ^t via the dependence of p_t^* on ξ^t (i.e., $a_t^* = a^*(p_t^*[\xi^t])$). Moreover, if $a^*(\cdot)$ is nonlinear, such path dependence will also be nonlinear.

¹¹Since the market cannot detect deviations, its information sets are indexed by the partial realizations of the public signal. Thus, along the path of play of any equilibrium in which the market's belief is correct, actions are a function of the current public history. But since all such sets are reached from a time-zero perspective, it follows that the Nash equilibrium concept suffices to characterize the outcome of the game.

The rest of the paper proceeds as follows. Necessary and sufficient conditions for Markov equilibria given a general best response $\chi_t := \chi(p_t^*, a_t^*)$, $t \geq 0$, are stated in Sections 3 and 5, respectively. The applications that employ nonlinear flow payoffs (Sections 4.2 and 4.3) and the existence results (Section 6) in turn specialize on the case $\chi_t = \chi(p_t^*)$; as argued in Section 3, this restriction is the natural one for studying traditional ratchet effects.

Remark 1 (On MPE). Any Markov equilibrium can be extended to MPE (with the market’s and the long-run player’s belief as states) provided an off-path Markov best response exists; the hurdle for showing such existence result is only technical, as the equilibrium analysis I perform does not restrict the long-run player’s behavior off the path of play.¹² Importantly, if a MPE exists and the value function is of class C^2 , the associated policy when beliefs are aligned in fact coincides with the policy of a Markov equilibrium (Remark 6, Section 5).

Remark 2 (On Assumption 1 and the Lipschitz property). The differentiability and growth conditions in Assumption 1 are used to obtain necessary conditions for Markov equilibria in the form of ODEs. On the other hand, the strong convexity assumption on $g(\cdot)$ permits the construction of Lipschitz candidate equilibria using solutions to such ODEs. The Lipschitz property in turn guarantees that the long-run player’s best-response problem (via the market’s conjecture of equilibrium play) is well defined in the sufficiency step. All these conditions can be relaxed, but the extra generality brings no new additional economic insights.¹³

3 Equilibrium Analysis: Necessary Conditions

To perform equilibrium analysis, one has to consider deviations from the market’s conjecture of equilibrium behavior and show that they are all unprofitable. After a deviation occurs, however, there is belief divergence, and long-run player’s belief becomes private. As I show in Section 5, the combination of hidden actions and persistent hidden information off the path of play leads traditional dynamic-programming methods to become particularly complex when the task is to find MPE.

¹²More precisely, the traditional approach to showing the existence of optimal (Markov) policies for stochastic control problems of infinite horizon is via HJB equations. However, for the class of games under study, such HJB approach raises additional complexities relative to standard decision problems (Section 5). Observe, however, that such an off-path best response always exists in settings where the set of actions is finite, and the horizon discrete and finite.

¹³Relatedly, the twice differentiability of (3) along the path of play imposed in Definition 3 is usually guaranteed under much weaker conditions. In fact, if $\kappa = 0$ (and thus beliefs are a martingale in equilibrium), the property follows from the continuity of the flow payoff and the integrability requirement on any admissible strategy (Theorem 4.4.9 and Remark 4.4.10 in Karatzas and Shreve, 1991). Establishing an analogous general result for $\kappa > 0$ is beyond this paper’s interest; thus, it is maintained as part of the definition, and verified to hold in each case via an indirect approach (i.e., via the *Feynman-Kac formula*; Remark 3.5.6 in Pham, 2009).

In order to bypass this complexity, I take a *first-order approach* to performing equilibrium analysis in the Markov case. First, I derive a necessary condition for Markov equilibria: namely, if deviating from the market's conjecture is not profitable, the value of a *small* degree of belief divergence must satisfy a particular ODE (Section 3.2). Second, I establish conditions under which a solution to this ODE makes the creation of any degree of belief asymmetry suboptimal, thus validating the first-order approach (Section 5.2). Importantly, this approach is critical for uncovering the economic forces at play.

3.1 Laws of Motion of Beliefs and Belief Asymmetry Process

Standard results in filtering theory state that, given a conjecture $(a_t^*)_{t \geq 0}$, the market's belief about θ_t given the public information up to t is normally distributed (with a mean denoted by p_t^*).¹⁴ In the case of the long-run player, he can always subtract—regardless of the strategy followed—the effect of his action on the public signal to obtain $dY_t := \theta_t dt + \sigma_\xi dZ_t^\xi = d\xi_t - a_t dt$, $t \geq 0$. Since $(\theta_t, Y_t)_{t \geq 0}$ is Gaussian, it follows that his posterior belief process is also Gaussian; denote by $(p_t)_{t \geq 0}$ the corresponding mean process.

In order for learning to be stationary, I set the common prior to have a variance equal to

$$\gamma^* = \sigma_\xi^2 \left(\sqrt{\kappa^2 + \sigma_\theta^2 / \sigma_\xi^2} - \kappa \right) > 0.$$

In this case, both the market and the long-run player's posterior beliefs about θ_t have variance γ^* at all times $t \geq 0$, and hence, $(p_t^*)_{t \geq 0}$ and $(p_t)_{t \geq 0}$ become their sufficient statistics, respectively. Observe also that γ^* is independent of both conjectured and actual play. In fact, because of the additively separable structure of the public signal, a change in the long-run player's strategy shifts the distribution of the public signal without affecting its informativeness, i.e., there are no experimentation effects.¹⁵

Lemma 1. *If the market conjectures $(a_t^*)_{t \geq 0}$, yet $(a_t)_{t \geq 0}$ is being followed, then*

$$dp_t^* = -\kappa(p_t^* - \eta)dt + \frac{\gamma^*}{\sigma_\xi^2} [d\xi_t - (p_t^* + a_t^*)dt] \quad \text{and} \quad (4)$$

$$dp_t = -\kappa(p_t - \eta)dt + \frac{\gamma^*}{\sigma_\xi} dZ_t, \quad t \geq 0, \quad (5)$$

¹⁴See Theorem 11.1 in Liptser and Shiryaev (1977). Formally, the pair (θ_t, ξ_t) is *conditionally* Gaussian, meaning that $\theta_t | \mathcal{F}_t$ is normally distributed despite $(\xi_t)_{t \geq 0}$ not being necessarily Gaussian. The latter occurs if a_t^* is a nonlinear function of $(\xi_s)_{s < t}$, $t \geq 0$, which can be in turn the result of a nonlinear Markov strategy. A nonlinear version of the Kalman-Bucy filter applies in this case.

¹⁵More generally, under a common normal prior with variance $\gamma^o \geq 0$, Theorem 12.1 in Liptser and Shiryaev (1977) shows that both posterior beliefs have a variance $(\gamma_t)_{t \geq 0}$ that satisfies $\dot{\gamma}_t = -2\kappa\gamma_t + \sigma_\theta^2 - \gamma_t^2 / \sigma_\xi^2$, $t > 0$, $\gamma_0 = \gamma^o$, i.e., the speed of learning is exogenous. It is easy to verify that γ^* is the unique strictly positive stationary solution of this ODE.

where $Z_t := \frac{1}{\sigma_\xi} \left(\xi_t - \int_0^t (p_s + a_s) ds \right) = \frac{1}{\sigma_\xi} \left(Y_t - \int_0^t p_s ds \right)$, $t \geq 0$, is a Brownian motion from the long-run player's perspective. Moreover, $(\xi_t)_{t \geq 0}$ admits the representation $d\xi_t = (a_t + p_t)dt + \sigma_\xi dZ_t$, $t \geq 0$, from his standpoint.

Proof: Refer to Theorems 7.12 and 12.1 in Liptser and Shiryaev (1977). \square

The right-hand side of (4) offers a natural orthogonal decomposition for the local evolution of the public belief: the trend $-\kappa(p_t^* - \eta)dt$, in the market's time t -information set, plus the residual 'surprise' process

$$d\xi_t - \mathbb{E}^{a^*}[d\xi_t | \mathcal{F}_t] = d\xi_t - (a_t^* + p_t^*)dt, \quad (6)$$

which is unpredictable from the market's perspective. Positive (negative) realizations of this surprise process convey information that the fundamentals are higher (lower), and the responsiveness of the public belief to this news is constant and captured by the sensitivity

$$\beta := \gamma^* / \sigma_\xi^2 = \sqrt{\kappa^2 + \sigma_\theta^2 / \sigma_\xi^2} - \kappa.^{16} \quad (7)$$

In the absence of news, the market adjusts its beliefs at rate κ , i.e., in the same way that the fundamentals change absent any shocks to their evolution.

The long-run player's belief $(p_t)_{t \geq 0}$ has an analogous structure, with the Brownian motion $Z_t = \frac{1}{\sigma_\xi} \left(\xi_t - \int_0^t (p_s + a_s) ds \right) = \frac{1}{\sigma_\xi} \left(Y_t - \int_0^t p_s ds \right)$ (or, equivalently, the surprise process $\sigma_\xi Z_t$) now providing news about $(\theta_t)_{t \geq 0}$; the last equality stresses that the realizations of $(Z_t)_{t \geq 0}$ are independent of the strategy followed and, thus, that $(p_t)_{t \geq 0}$ is exogenous.¹⁷ In contrast, the public belief is controlled by the long-run player through his actions affecting the surprise term (6) via the realizations of $(\xi_t)_{t \geq 0}$.

To see how deviations from $(a_t^*)_{t \geq 0}$ affect the public belief, Lemma 1 states that the public signal follows $d\xi_t = (a_t + p_t)dt + \sigma_\xi dZ_t$ from the long-run player's perspective. Plugging this into (4), straightforward algebra yields that $\Delta_t := p_t^* - p_t$ satisfies

$$d\Delta_t = [-(\beta + \kappa)\Delta_t + \beta(a_t - a_t^*)]dt, \quad t > 0, \quad \Delta_0 = 0. \quad (8)$$

It is clear from (8) that deviations from $(a_t^*)_{t \geq 0}$ can lead to belief asymmetry ($\Delta \neq 0$): in this case, the long-run player's belief is private, as the correction $d\xi_t - a_t dt$ used to obtain dY_t is

¹⁶Thus, beliefs are less responsive to such news when κ , σ_ξ and $1/\sigma_\theta$ grow. In particular, higher rates of mean reversion lead to a more concentrated long-run distribution of the fundamentals, and hence, to less responsiveness to news.

¹⁷In particular, observe that $p_t = e^{-\beta(\kappa+\beta)t + \kappa\eta} \int_0^t e^{-(\kappa+\beta)(t-s) + \beta \int_0^t e^{-(\beta+\kappa)(t-s)} dY_s}$, so p_t is a linear function of the realizations $(Y_s)_{s \leq t}$.

incorrectly anticipated by the market. In particular, an upward deviation on the equilibrium path leads the market to hold an excessively optimistic belief about the fundamentals (i.e., $\Delta_t = p_t^* - p_t > 0$), consequence of underestimating the contribution of the long-run player's action to the public signal. I refer to $(\Delta_t)_{t \geq 0}$ as the *belief asymmetry* process.

Starting from a common prior, however, beliefs remain aligned on the equilibrium path (i.e., $\Delta_0 = 0$ and $a_t^* = a_t$, $t \geq 0$, imply $\Delta \equiv 0$). In particular, both parties expect any surprise realization (6) to decay at rate κ on average along the path of play, as the common belief evolves according to $dp_t = -\kappa(p_t - \eta)dt + \beta\sigma_\xi dZ_t$ at any on-path history going forward (see eqn. (5)).

3.2 Necessary Conditions: The Ratcheting Equation

Consider the Markov case. In order to understand the form of ratcheting that arises in this model, it is useful to interpret $(\xi_t)_{t \geq 0}$ as a measure of performance (e.g., output) and the market's best response $\chi(\cdot, \cdot)$ as a payment that rewards high performance. For expositional simplicity, suppose that the long-run player is simply paid based on the market's belief about the fundamentals, $\chi(p^*, a^*) = p^*$; this can occur if, for instance, the fundamentals reflect an unobserved payoff-relevant characteristic of the long-run player (e.g., managerial ability).

In this case, the dynamic of the public belief (4) is effectively an incentive scheme, i.e., a rule that determines how payments are revised in response to current performance:

$$\underbrace{dp_t^*}_{\text{change in payments}} = \underbrace{-\kappa(p_t^* - \eta)dt}_{\text{exogenous trend}} + \underbrace{\beta}_{\text{sensitivity}} \times \left[\underbrace{d\xi_t}_{\text{performance}} - \underbrace{(p_t^* + a^*(p_t^*))dt}_{\text{target}} \right].$$

Central to this scheme is the presence of a *target* in the form of expected performance: the long-run player will positively influence his payment if and only if realized performance, $d\xi_t$, is above the market's expectation, $\mathbb{E}^{a^*}[d\xi_t | \mathcal{F}_t] = (p_t^* + a^*(p_t^*))dt$. But observe that the market's updated belief feeds into the target against which the long-run player's performance is evaluated tomorrow. Moreover, an upward revision of such target leads to a more demanding incentive scheme to be faced in the future—a *ratchet principle* ensues.¹⁸

In continuous time, the distinction between today and tomorrow disappears. It is then natural to define a *ratchet* as the (local) sensitivity of the performance target with respect

¹⁸The way in which the public belief (4) is written (i.e., with $(p_t^* + a^*(p_t^*))dt$ displayed as a target, or with $-\beta(p_t^* + a^*(p_t^*))$ in the drift) is immaterial: the point is that, to defeat or accelerate the natural reversion to the mean, $d\xi_t$ must be greater than $(p_t^* + a^*(p_t^*))dt$, and the same logic follows. Also, specializing to $\chi(p^*, a^*) = p^*$ is without loss. In fact, $dZ_t^* := [d\xi_t - (p_t^* + a^*(p_t^*))dt]/\sigma_\xi$ is a Brownian motion from the market's perspective, so, using Ito's rule, $(\chi(p_t^*, a^*(p_t^*)))_{t \geq 0}$ has innovations also driven by $d\xi_t - (p_t^* + a^*(p_t^*))dt$.

to contemporaneous realized performance $d\xi_t$, namely,

$$\text{Ratchet} := \frac{d(p_t^* + a^*(p_t^*))}{d\xi_t} = \left[1 + \frac{da^*(p^*)}{dp^*} \right] \Big|_{p^*=p_t^*} \times \underbrace{\frac{dp_t^*}{d\xi_t}}_{=\beta} = \beta + \beta \frac{da^*(p_t^*)}{dp^*}.^{19} \quad (9)$$

To understand the implications of this ratchet principle on incentives, consider the following strategy $(a_t)_{t \geq 0}$: the long-run player deviates from $(a_t^*)_{t \geq 0}$ for the first time at time t by choosing $a_t > a_t^*$, and he then matches the market's expectation of performance thereafter. Intuitively, this deviation helps illustrate the strength of the dynamic cost of exhibiting high performance through quantifying the extra effort cost that the long-run player must bear to avoid disappointing the market after strategically surprising the latter.

Matching the market's expectation of performance at all times after a deviation occurs amounts to equating the drift of $(\xi_s)_{s > t}$ from the market's perspective. Thus, the long-run player must take actions according to

$$\underbrace{a_s + p_s}_{\substack{\text{LR player's expectation} \\ \text{of performance at instant } s > t}} = \underbrace{a(p_s^*) + p_s^*}_{\substack{\text{market's expectation} \\ \text{of performance at instant } s > t}} \Rightarrow a_s = a^*(p_s + \Delta_s) + \Delta_s, \quad s > t.$$

The term $a^*(p_s + \Delta_s)$ captures how the long-run player adjusts his actions to match the market's expectation of future behavior. The isolated term Δ_s in turn captures how his actions are modified due to holding a private belief off the path of play. Specifically, since an upward deviation makes the market overly optimistic about the fundamentals, the long-run player anticipates that he will have to exert more effort than expected by the market to match all future "targets," as his private belief indicates that the fundamentals are lower.

If the long-run player does not deviate from $a^*(\cdot)$, $p_t = p_t^*$ holds at all times, and effort is costly according $(g(a^*(p_t)))_{t \geq 0}$ in this case. To compute the corresponding cost under $(a_t)_{t \geq 0}$, let $\epsilon := a_t - a^*(p_t^*) > 0$ denote the size of the initial deviation. From the dynamic of belief asymmetry (8) it follows that $\Delta_{t+dt} = \beta \epsilon dt$, and hence, using that $a_s = a^*(p_s + \Delta_s) + \Delta_s$,

$$\Delta_s = e^{-\kappa(s-t)} \beta \epsilon dt > 0, \quad \forall s > t. \quad (10)$$

That is, the initial stock of belief asymmetry created, $\beta \epsilon dt$, decays at rate κ under this deviation. Thus, the extra cost that the long-run player must bear to match the market

¹⁹This notion of sensitivity is with respect to *realizations* of $(\xi_t)_{t \geq 0}$, and such realizations are driven by $(\theta_t)_{t \geq 0}$ (not by $(p_t^*)_{t \geq 0}$). See remark 3 for more details on this sensitivity.

expectation of performance at time $s > t$ corresponds, for $\epsilon > 0$ small, to

$$g(a^*(p_s + \Delta_s) + \Delta_s) - g(a^*(p_s)) = \underbrace{g'(a^*(p_s)) \times \left[1 + \frac{da^*(p_s)}{dp^*}\right]}_{\text{ratchet}} \beta \epsilon e^{-\kappa(s-t)} dt + O(\epsilon^2), \quad (11)$$

and the ratchet (9) naturally appears. From (11), sustaining performance becomes more costly as the strength of the ratchet grows if positive effort is being exerted (i.e., $g'(a) > 0$), as this requires more subsequent effort to match the market's perceived distribution of $(\xi_t)_{t \geq 0}$.

If $a^*(\cdot)$ is a Markov equilibrium, this type of deviation cannot be profitable. Thus, the extra cost of effort at time t (i.e., $g'(a^*(p_t))\epsilon$) must equate the change in the long-run player's continuation payoff. The latter value consists of extra effort costs in (11), plus the extra stream of payments $(\Delta_t)_{t \geq 0}$ consequence of the public belief increasing from $(p_s)_{s > t}$ to $(p_s + \Delta_s)_{s > t}$. The next proposition formalizes this discussion for a general $\chi(\cdot, \cdot)$ as in the baseline model; recall that $\rho := (g')^{-1}(\cdot)$ and let $\sigma := \beta\sigma_\xi$ denote the volatility of the common belief along the path of play.

Proposition 1 (Necessary Conditions). *Consider a Markov equilibrium $a^*(\cdot)$. Then, $g'(a^*(p)) = \beta q(p)$, where*

$$q(p) := \mathbb{E} \left[\int_0^\infty e^{-(r+\kappa)t} \left[\frac{d}{dp^*} [u(\chi(p^*, a^*(p^*)))] \Big|_{p^*=p_t} - g'(a^*(p_t)) \left(1 + \frac{da^*(p_t)}{dp^*} \right) \right] dt \Big| p_0 = p \right] \quad (12)$$

and $dp_t = -\kappa(p_t - \eta)dt + \sigma dZ_t$, $p_0 = p$. The corresponding equilibrium payoff is given by

$$U(p) := \mathbb{E} \left[\int_0^\infty e^{-rt} [u(\chi(p_t, \rho(\beta q(p_t)))) - g(\rho(\beta q(p_t)))] dt \Big| p_0 = p \right]. \quad (13)$$

Proof: See the Appendix. □

The previous result states a constraint on the structure of any Markov equilibrium. Specifically, if $a^*(\cdot)$ is a Markov equilibrium, the resulting dynamic gain from the deviation under study, $q(p)$, must satisfy $g'(a^*(p)) = \beta q(p)$, through which current and future behavior are linked; β in turn represents the sensitivity of the public belief to current performance. In (12), (i) the ratchet negatively contributes to the value of the deviation whenever $g'(a^*(p))(1 + da^*/dp^*) > 0$, and (ii) κ in the discount rate reflects that the additional payments $(\Delta_t)_{t \geq 0}$ from the deviation decay at that rate. Finally, the equilibrium payoff (13) follows from plugging $a^*(\cdot) = \rho(\beta q(\cdot))$ in (3).

Observe that $q(p)$ is, by definition, the extra value to the long-run player of inducing a small degree of initial belief asymmetry that vanishes at rate $\kappa > 0$, when the current common

belief is p ; thus, $q(\cdot)$ is a measure of marginal utility in which, starting from a common belief, future beliefs do not coincide.²⁰ Proposition 1 opens the possibility of finding Markov equilibria via solving for this measure of marginal utility, and the next result is central to the subsequent analysis in this respect.

Proposition 2 (ODE Characterization: Actions and Payoffs). *Consider a Markov equilibrium $a^*(\cdot)$. Then, $a^*(\cdot) = \rho(\beta q(\cdot))$, where $q(p)$ defined in (12) satisfies the ODE*

$$\left[r + \kappa + \beta + \beta^2 \rho'(\beta q(p)) q'(p) \right] q(p) = \frac{d}{dp} [u(\chi(p, \rho(\beta q(p))))] - \kappa(p - \eta) q'(p) + \frac{1}{2} \sigma^2 q''(p). \quad (14)$$

The long-run player's payoff (13) in turn satisfies the linear ODE

$$rU(p) = u(\chi(p, \rho(\beta q(p)))) - g(\rho(\beta q(p))) - \kappa(p - \eta)U'(p) + \frac{1}{2} \sigma^2 U''(p), \quad p \in \mathbb{R}. \quad (15)$$

Proof: See the Appendix. □

The previous result offers expressions for the pair (q, U) defined by (12)–(13) in the form of a system of ODEs. The U -ODE (15) is a standard linear equation that captures the local evolution of a net present value.²¹ Instead, the q -ODE (14) is a nonlinear equation that captures local evolution that the value of a small degree of belief asymmetry must satisfy in equilibrium. I refer to (14) as the *ratcheting equation*; this equation is novel.

To understand this equation, notice first that the long-run player faces a dynamic decision problem given any $a^*(\cdot)$. Thus, (14) behaves as an Euler equation in the sense that it optimally balances the forces that determine his intertemporal behavior. The right-hand side of (14) consists of forces that strengthen his incentives: myopic benefits (the first term) and cost-smoothing motives (the second and third terms); the larger either term, the larger $q(p)$, everything else equal.²² The left-hand side instead consists of forces that weaken his incentives: the rate of mean reversion κ (the higher this value, the more transitory any change in beliefs is) and the ratchet $\beta + \beta da^*/dp^* = \beta + \beta^2 \rho'(\beta q(\cdot)) q'(\cdot)$.

The novelty of (14) lies on the ratcheting embedded in it altering its structure relative to traditional Euler equations in dynamic decision problems, and this has economic implications. In fact, (14) is an equation for marginal utility in which the anticipation of

²⁰As I show in Remark 6 in Section 5, if the long-run player's value function $V(p, p^*)$ is sufficiently differentiable, then $q(p) = V_{p^*}(p, p)$.

²¹This equation is usually referred to as an *arbitrage equation*: the interest earned on the present value (left-hand side) must equate the current flow (first term on the right) plus the expected capital gains (the expected change in the present value; the second term on the right). See, Dixit and Pindyck (1994).

²²Using Ito's rule, $\mathbb{E}[dq(p_t)/dt|p_t = p] = -\kappa(p - \eta)q'(p) + \frac{1}{2}\sigma^2 q''(p)$. Thus, if the value of affecting the public belief is expected to increase, then, because $g(\cdot)$ is convex, it is optimal to frontload effort.

stronger (weaker) incentives tomorrow dampens (strengthens) today's incentives. This is seen in the interaction term $\beta^2 \rho'(\beta q(\cdot)) q'(\cdot) q(\cdot)$ on left-hand side of (14), where larger values of $da^*/dp^* = \rho'(\beta q(\cdot)) q'(\cdot)$ put more downward pressure on $q(p)$ (and vice-versa), everything else equal; in traditional Euler equations, the opposite affect arises (see also Remark 4).

To conclude this section, two observations. First, notice that since the market perfectly anticipates the long-run player's actions in equilibrium, no belief asymmetry is created along the path of play. As a result, the long-run player bears the ratcheting cost of matching the market's revisions of $a^*(p_t)$ as the common belief changes, but not the ratcheting cost of explicitly accounting for belief divergence. The potential appearance of the latter cost nevertheless affects on-path payoffs through the long-run player's equilibrium behavior.²³

Second, notice that the strength of the ratcheting that arises in any economic environment is endogenous via da^*/dp^* , and the latter can strengthen or weaken incentives depending on its sign. Importantly, if the market's best response depends on a^* , the term $\beta da^*/dp^*$ also accompanies $(u \circ \chi)'$ on the right-hand side of (14), thus distorting the strength of the traditional ratchet principle (understood as a target revision). For this reason, the applications in sections 4.2 and 4.3, and the existence results in Section 6, eliminate such dependence. Conditions for global incentive compatibility (Section 5) are instead derived for a general χ , so as to complement the analysis of this section. In what follows, I sometimes refer to $\beta da^*/dp^* = \beta^2 \rho'(\beta q(\cdot)) q'(\cdot)$ and β as the endogenous and exogenous ratchets, respectively, to emphasize the type of force under analysis.

The next three remarks are technical, and not needed for the subsequent analysis.

Remark 3 (On ratchets and learning). The identification of a ratchet follows from the public belief (4) admitting a representation in terms of the surprise process $d\xi_t - (a_t^* + p_t^*) dt$ —such *innovation processes* play a central role in representation results for beliefs in optimal filtering theory beyond the Gaussian case (cf., Theorem 8.1 in Liptser and Shiryaev, 1977). The ratchet (9) as a sensitivity measure follows from a notion of derivative of p_t^* with respect to the realization ξ^t that determines it, with ξ^t an element of $C([0, t])$ (i.e., a stochastic, or *Malliavin* derivative). Under that type of derivative (denote it by $(D_s \cdot)_{s \leq t}$ for fixed t , with $D_s p_t^*[\xi^t]$ the change in p_t^* resulting from a marginal increase in the time- s realization), $D_t p_t^*[\xi^t] = \beta$, and the chain rule applies (Appendix A in Di Nunno et al., 2009).

Remark 4 (On ratcheting and Euler equations). By the envelope theorem, the change in the optimizer that results from a small change in the current state does not contribute

²³Formally, differentiate (15) to obtain the following ODE for the long-run player's on-path marginal utility $U'(\cdot)$: $[r + \kappa]U'(p) = \frac{d}{dp}[u(\chi(p, \rho(\beta q(p))))] - \beta q(p) \frac{da^*(p)}{dp^*} - \kappa(p - \eta)U'(p) + \frac{1}{2}\sigma^2 U''(p)$. The ratcheting equation can also be written as $[r + \kappa + \beta]q(p) = \frac{d}{dp}[u(\chi(p, \rho(\beta q(p))))] - \beta q(p) \frac{da^*(p)}{dp^*} - \kappa(p - \eta)q'(p) + \frac{1}{2}\sigma^2 q''(p)$. Comparing the left-hand sides of these ODEs confirms that the ratcheting cost $\beta q(p)$ is absent in $U'(\cdot)$.

to marginal utility along the optimal trajectory in a dynamic decision problem. In the class of games analyzed, this holds too, but there is the also effect of a small change in p^* (or, equivalently, Δ) affecting the market's conjecture, which is correct in equilibrium. The resulting equation for marginal utility with respect to p^* (when beliefs are aligned) then exhibits the ratcheting term $-q(p)\beta da^*/dp^* = -\beta^2 q(p)\rho'(\beta q(p))q'(p)$ which effectively acts as a change in the long-run player's action that has a (negative) first-order impact on marginal utility, an effect that is absent in decision problems. While Euler equations do exhibit interaction terms of similar structure, these arise from a change in marginal utility while keeping the decision maker's action fixed; but if actions positively affect the controlled state, the sign is the opposite. An interaction term of that nature is absent in (14) due to the long-run player's action being offset by the market's conjecture along the path of play.

Remark 5 (On deviations that yield the ratcheting equation). The ratcheting ODE (14) can be derived using two other deviations. After a first upward deviation, the long-run player:

1. Chooses $a_t = a^*(p_t^*)$ forever after. In this case, the long-run player does not bear the extra cost of explicitly correcting for Δ in his effort decision, but $(\Delta_t)_{t \geq 0}$ decays at rate $\beta + \kappa$; in (12), κ and $g'(a^*(p_s))(1 + da^*(p_s)/dp^*)$ change to $\beta + \kappa$ and $g'(a^*(p_s))da^*(p_s)/dp^*$, respectively. Intuitively, since the long-run player underperforms in this case, he expects the market to be disappointed more often, and hence to correct its belief faster than the rate at which shocks dissipate, explaining the extra β . Ratcheting is then costly because changes in payments are more transitory.
2. Chooses $a_t = a^*(p_t)$ forever after. In this case, the long-run player does not account for the market's incorrect belief about a^* or for Δ , but belief asymmetry decays, to a first-order approximation, at rate $\beta + \kappa + da^*(p_s)/dp^*$; in (12), κ and $g'(a^*(p_s))(1 + da^*(p_s)/dp^*)$ change to $\beta + \kappa + da^*(p_s)/dp^*$ and 0, respectively. In particular, if $da^*(p_s)/dp^* > 0$, the long-run player does not incur any extra after the deviation, but the additional payment now vanishes even faster, and vice-versa.²⁴

In either case, the extra costs that arise due to changes in payments being more transitory coincide with the extra effort costs needed to match the market's expectation of performance under the original deviation.

4 Applications

In this section, I study *ratchet effects*, i.e., equilibrium consequences of the ratchet principle. The first two applications (career concerns, Section 4.1; and monetary policy; Section 4.2)

²⁴I am grateful to an anonymous referee for suggesting this deviation.

focus on the exogenous ratchet β , whereas the last one (earnings management, Section 4.3) focuses on da^*/dp^* . Nonlinearities naturally appear in the last two settings, and all the examples rely on the ratcheting equation (14) to flesh out properties of equilibrium behavior.

4.1 Career Concerns

I revisit Holmström’s (1999) model of career concerns to illustrate how the ratcheting identified in the previous section is embedded in the equilibrium that he finds. Thus, when employers learn about workers’ abilities, the possibility of employers ratcheting their expectations of future performance can undermine workers’ reputational incentives.

A large number of firms (the market) compete for a worker’s labor (the long-run player). Interpret $(\xi_t)_{t \geq 0}$ as output, $(a_t)_{t \geq 0}$ as effort, and $(\theta_t)_{t \geq 0}$ as the worker’s skills. The worker is risk neutral ($u(\chi) = \chi$) and the market *spot*: at the beginning of “period” $[t, t + dt)$, the worker is paid the market’s expectation of production over the same period, namely,

$$\text{wage at } t := \lim_{h \rightarrow 0} \frac{\mathbb{E}^{a^*}[\xi_{t+h} | \mathcal{F}_t] - \xi_t}{h} = a_t^* + p_t^* =: \chi(p_t^*, a_t^*).$$

Note that surplus over $[t, t + dt)$, $d\xi_t - g(a_t)dt$, is maximized at $a^e > 0$ satisfying $g'(a^e) = 1$.

The ratcheting equation offers a simple method to solve for the equilibrium found by Holmström. In fact, it is easy to verify that (14) admits a constant solution q defined by $[r + \kappa + \beta]q = 1$ in this case. Thus, there is a constant equilibrium a^* satisfying

$$g'(a^*) = \beta q(p) = \frac{\beta}{r + \kappa + \beta}, \text{ where } \beta = \frac{\gamma^*}{\sigma_\xi^2} = \sqrt{\kappa^2 + \sigma_\theta^2 / \sigma_\xi^2} - \kappa.$$

In this equilibrium, β in the numerator captures the sensitivity of the market’s belief to output surprises. The rate of mean reversion explicitly appears in the denominator dampening incentives: as κ increases, changes in beliefs—and hence, changes in wages—have less persistence. Finally, β in the denominator corresponds to the ratchet (9): in a deterministic equilibrium $da^*/dp^* = 0$, i.e., the market never revises its conjecture of equilibrium behavior. Observe that this term negatively contributes to generating inefficiently low effort.

To see why there is a ratchet effect embedded in this equilibrium, notice that, along the path of play, a one-time output surprise of unit size makes both the market and the long-run player expect an additional output (hence, wage stream) of value $\beta/(r + \kappa)$ relative to the counterfactual history in which the surprise did not take place: the common belief reacts with sensitivity β , and this effect vanishes at rate κ on average. However, if the same surprise is the outcome of extra hidden effort, the worker only internalizes $\beta/(r + \beta + \kappa)$ in terms of extra utility. The intuition for this result follows from the logic of the previous section. In

fact, since the market incorrectly ratchets up its expectations of output after the deviation, producing an extra one-time output surprise of unit size—and thus, a wage stream of size $\beta/(r + \kappa)$ —requires more effort than an isolated one-time unit effort increase, as the worker’s off-path private information indicates that he would underperform otherwise.²⁵

4.2 Ratcheting and Commitment in Monetary Policy

This section shows that in economies where agents learn about hidden components of inflation, the possibility of a market ratcheting up its expectations about future prices can induce a monetary authority to exhibit more *commitment*. In particular, if employment responds to unanticipated changes in the price level, monetary policy as an instrument to boost employment can be less aggressive relative to settings where inflation trends are observed or absent. In contrast to the previous application, the potential appearance of a ratcheting cost now can have a positive impact on an equilibrium outcome (namely, on inflation).

The (log) price index $(\xi_t)_{t \geq 0}$ of an economy is given by $d\xi_t = (a_t + \theta_t)dt + \sigma_\xi dZ_t^\xi$, where $(a_t)_{t \geq 0}$ denotes the economy’s money growth rate process and $(\theta_t)_{t \geq 0}$ corresponds to a hidden inflation trend that evolves according to

$$d\theta_t = -\kappa\theta_t dt + \sigma_\theta dZ_t^\theta. \quad (16)$$

Intuitively, $(Z_t^\theta)_{t \geq 0}$ represents shocks beyond the central bank’s control that move the economy’s inflation trend $(\theta_t)_{t \geq 0}$ away from a publicly known long-run inflation target that has been normalized to zero (i.e., $\eta = 0$ in (1) in the baseline model). Such unobserved shocks vanish, on average, at a rate $\kappa \geq 0$.²⁶

Crucially, the central bank has a commitment problem with respect to its long-term inflation goal: in an attempt to boost short-run employment, the monetary authority cannot refrain from injecting money into the economy, which results in an effective trend of size $a_t + \theta_t$, $t \geq 0$. In line with a sizable literature on transparency in monetary policy (see, for instance, Cukierman and Meltzer, 1986; and Atkeson et al. 2007), I assume that the public does not observe the money growth rate process $(a_t)_{t \geq 0}$ directly.

Employment responds to unexpected inflation as in traditional Phillips curves. Specifi-

²⁵It is easy to verify that in Holmström’s model a^* is also optimal off the path of play, thus implying that ratcheting is equally costly at all different levels of private beliefs. Intuitively, if the worker is relatively more pessimistic and the market updates it beliefs upwards, the worker expects to underperform more frequently; but if the worker is instead relatively more optimistic, he then expects the market to be positively surprised less often. In each case, these ratcheting costs are independent of the level of the worker’s own private belief due to the model being fully linear and additive in both beliefs.

²⁶Models of inflation that allow for unobserved trends have been used to explain statistical properties of U.S. postwar inflation data. See, for instance, Stock and Watson (2007) and Cogley et al. (2010).

cally, (log) employment n_t evolves according to

$$dn_t = -\kappa_n n_t dt + \nu(d\xi_t - (a_t^* + p_t^*)dt). \quad (17)$$

where $\kappa_n \geq 0$ and $\nu > 0$. Intuitively, workers and firms set nominal wages at the beginning of $[t, t + dt)$ (i.e., before the price level is realized) taking into account their expectations of inflation $(a_t^* + p_t^*)dt$; high realizations of the price index (i.e., $d\xi_t - (a_t^* + p_t^*)dt > 0$) then reduce real wages, thereby inducing hiring. Finally, the impact of such unanticipated shocks on employment vanishes at rate κ_n : since employment locally reverts to zero in this case, I interpret the latter value as the (normalized) natural level of (log) employment.

To obtain a version of this model that can be directly analyzed with the results presented in this paper, I assume that (i) $\kappa_n = \kappa \geq 0$, (ii) $\nu = \beta$, and $n_0 = p_0^*$. In this case, $n_t = p_t^*$ at all times, and thus the setting fits in the baseline model of Section 2.²⁷ It is important to stress, however, that (i)–(iii) are by no means critical for the subsequent analysis. In fact, the commitment result presented under this choice of parameters also holds for the general specification (16)–(17), and the corresponding equilibria can be computed using analogous methods; the parametric restriction is thus purely driven by expositional reasons.²⁸

The monetary authority trades off the benefits of affecting employment with the effects that money growth has on the price level. These preferences are captured by

$$\mathbb{E} \left[\int_0^\infty e^{-rt} \left(-\frac{n_t^2}{2} - \psi \frac{a_t^2}{2} \right) dt \right], \quad (18)$$

with $\psi > 0$ the relative weight that the central bank attaches to the impact of money on inflation, and where the central bank's target of (log) employment coincides with the natural level. Observe that these preferences are nonlinear in $n_t = p_t^*$, and that the monetary authority has a myopic incentive to boost employment when $n < 0$.²⁹

²⁷Notice that when $\kappa_n = \kappa$, $p_t^* = e^{-\kappa t} p_0 + \beta \int_0^t e^{-\kappa(t-s)} [d\xi_s - (a_s^* + p_s^*)ds]$ and $n_t = e^{-\kappa t} n_0 + \nu \int_0^t e^{-\kappa(t-s)} [d\xi_s - (a_s^* + p_s^*)ds]$ hold at all times. The result then follows from $n_0 = p_0^*$ and $\nu = \beta$.

²⁸The analysis that follows shows that money has a more transitory effect on employment when $(\theta_t)_{t \geq 0}$ is hidden than when it is observed, thus leading to weaker incentives in the first case. For general (16)–(17), employment follows $dn_t = [-\kappa_n n_t - \nu \Delta_t + \nu(a_t - a_t^*)]dt + \nu \sigma_\xi dZ_t$ from the central bank's perspective if $(\theta_t)_{t \geq 0}$ is hidden, with $\Delta_t := p_t^* - p_t$ as in (8). Instead, $dn_t = [-\kappa_n n_t + \nu(a_t - a_t^*)]dt + \nu \sigma_\xi dZ_t$ when the trend is observed. Thus, when beliefs are aligned, increasing the supply of money above the market's expectations in the hidden case leads to the creation of a strictly positive Δ that puts additional downward pressure on employment relative to the observable case, and the same logic follows. Finally, since the environment is linear-quadratic, solving for a model involving (16)–(17) can be done analytically.

²⁹Quadratic loss functions naturally appear in second-order approximations of households' utilities in general equilibrium models, and they are widely used in the “discretion versus commitment” literature; see, for instance, Gali (2008) for an exposition that covers both topics. The wedge in employment is, in many instances, equivalently measured in terms of an output gap.

Before entering the analysis, observe that since in equilibrium the market will anticipate the policy $(a_t^*)_{t \geq 0}$ chosen by the monetary authority, money will have no impact on employment (i.e., (17) evolves as if uncontrolled on the path of play), but if $a_t^* > 0$, inflation is created. The central bank's commitment problem is thus a traditional one (e.g., Kydland and Prescott, 1977): the central bank would like to commit to a zero money growth rule, but, once the market forms expectations accordingly, incentives to deviate from it appear.

4.2.1 Observable Benchmark

Suppose that the inflation trend is observable—the environment then becomes one of imperfectly observable actions only. In fact, the ability to observe θ_t allows the market to remove it from (17) (i.e., $p_t^* = \theta_t$) and, using that $\sigma := \beta\sigma_\xi$, the Phillips curve (17) becomes

$$dn_t = [-\kappa n_t + \beta(a_t - a_t^*)]dt + \sigma dZ_t^\xi. \quad (19)$$

Intuitively, because $(\theta_t)_{t \geq 0}$ is perfectly observed, workers can index their nominal wages to it, which leads real wages to become independent of the current level of the inflation trend.

In equilibrium, the market's conjecture about money growth must be correct. I assume that $r + 2\kappa > 2\beta/\sqrt{\psi}$, which ensures the existence of equilibria in which money growth is linear in the current level of employment.

Proposition 3. *In any linear equilibrium, $a^{*,o}(n) = \frac{\beta}{\psi}\alpha^o n$, where $\alpha^o < 0$.*

Proof: See the Appendix. □

The intuition is simple: since the central bank wants to drive employment towards its ideal target, the money supply must increase (decrease) if n_t is below (above) 0. The functional form comes from (i) behavior being characterized by $a^{*,o}(n) = (g')^{-1}(\beta q(n)) = \beta q(n)/\psi$ when $g(a) = \psi a^2/2$, and (ii) the marginal benefit of boosting employment, $q(n)$, being linear in n this linear-quadratic game.³⁰

4.2.2 Hidden Case

In this case, the market cannot remove $(\theta_t)_{t \geq 0}$ from the Phillips curve, and the latter becomes

$$dn_t = [-\kappa n_t + \beta(a_t - a_t^*) - \beta(p_t^* - p_t)]dt + \sigma dZ_t. \quad (20)$$

³⁰Specifically, there are two linear equilibria as in Proposition 3 in this observable case. However, it is only when the market expects the more moderate policy to arise in equilibrium that the full-commitment policy $a \equiv 0$ is an admissible strategy for the central bank, which is a necessary requirement for discussing the value of commitment (see the proof for details). In any case, the equilibrium policy for the hidden case that is presented next is less steep than both of them, and Figure 2 depicts the less aggressive one.

Dynamic (20) differs from (19) only for the presence of $-\beta(p_t^* - p_t)$, which captures how employment is now affected by the market's incorrect expectation of inflation after a deviation from $(a_t^*)_{t \geq 0}$ occurred. In particular, as $p_t^* - p_t$ grows, employment decays faster.

Because $n_t = p_t^*$ at all times, we can use the ratcheting equation (14) to compute equilibria. The next proposition relies on an existence result for linear equilibria in a class of linear-quadratic games (Section 6.1). As before, assume that $r + 2\kappa > 2\beta/\sqrt{\psi}$.

Proposition 4. *If $(\theta_t)_{t \geq 0}$ is hidden, there exists a linear equilibrium $a^{*,h}(n) = \frac{\beta}{\psi}\alpha^h n$, $\alpha^h < 0$, such that $|\alpha^h| < |\alpha^o|$.*

Proof: See the Appendix. □




Figure 1: Equilibrium policies in the observable and hidden cases. Parameter values: $r = \sigma_\xi = \sigma_\theta = \psi = 1$ and $\kappa = 0.3$.

In the equilibrium found, the monetary policy rule is less aggressive than in the observable benchmark. Thus, the monetary authority exhibits more commitment, as the equilibrium policy is pointwise closer to the full commitment rule. This in turn results in a lower inflationary bias over the region where it is tempting to boost employment (i.e., $n < 0$).

To understand the result, start with the observable case. In this setting, the impact that an unanticipated change in the price level has on employment decays at rate κ , and any off-path history has an on-path counterfactual characterized by the same history of price realizations. A deviation by the central bank is interpreted as a shock to the price level, and hence, changes in the rate of growth of money have the same impact on employment.

In the hidden case, however, there is an identification problem, as an unanticipated changes could also be the outcome of changes in $(\theta_t)_{t \geq 0}$. In particular, if the central bank increases money above the market's expectation, the market will overestimate the value of the trend, and hence the monetary authority will find it more costly to surprise the economy with inflation relative to the observable case. Intuitively, the term $-\beta(p_t^* - p_t)$ present in (20)

captures how, in response to exceedingly high forecasts of inflation, workers ratchet up their future demands for nominal wages, which in turn puts downward pressure on future hiring. The monetary authority then anticipates that, in order to generate an effect on employment that decays at rate κ , more inflation than in the observable case is needed. Inflation thus becomes more costly, thereby inducing more commitment.

To conclude, observe that both equilibrium policies explicitly depend on the level of employment, and hence, the endogenous ratcheting costs are present in both the observable and unobservable cases. Interestingly, this form of ratcheting has now a positive impact on incentives: since the equilibrium policy is negatively sloped (i.e., $da^*/dp^* < 0$), the market actually lowers its expectation of money growth as employment approaches zero from the left, thus incentivizing the creation of more inflation in that region, everything else equal. It is easy then to deduce that the wedge between the two equilibrium policies is entirely driven by the ratcheting cost that appears off the path of play. In fact, using the ratcheting equation,³¹ it can be verified that the coefficients α^o and α^h in each linear equilibrium satisfy

$$\left[r + 2\kappa + \beta^2 \frac{\alpha^o}{\psi} \right] \alpha^o = -1 \quad \text{and} \quad \left[r + 2\kappa + \beta + \beta^2 \frac{\alpha^h}{\psi} \right] \alpha^h = -1,$$

and thus the difference is driven by the ratchet β that affects incentives in the hidden case.

4.3 Ratcheting and Thresholds in Earnings Management

This application examines managers' incentives to boost firms' earnings reports when they face strong myopic incentives to exceed a zero-earnings threshold. The main finding is that firms that are expected to exceed the threshold can actually inflate financial statements more aggressively than those firms expected to underperform, despite their managers having weaker myopic incentives. Central to this result is the endogenous ratchet $\beta da^*/dp^*$.

A firm's (cumulative) earnings report process $(\xi_t)_{t \geq 0}$ is given by $d\xi_t = (a_t + \theta_t)dt + \sigma_\xi dZ_t^\xi$. In this specification, a_t denotes the degree of earnings manipulation exerted by the firm at time $t \geq 0$, and $(\theta_t)_{t \geq 0}$ the firm's unobserved fundamentals. The latter are assumed to evolve according to a Brownian martingale $d\theta_t = \sigma_\theta dZ_t^\theta$.

I assume that the firm pays its dividends far in the future and that its earnings management practices are based on accounting techniques exclusively (e.g., discretionary accruals, typically difficult to observe). In this case, boosting financial statements imposes no real costs to the firm in the short- or medium-run, enabling the analysis to isolate learning-driven ratchet effects. The market then tries to undo the manager's actions when assessing

³¹From (19) and (20) it is easy to see that the marginal value of boosting employment in the observable case satisfies an ODE analogous to the ratcheting equation in which β in the left-hand side does not appear.

short-term performance. Specifically, the market expects the firm’s “natural” earnings over $[t, t + dt)$ to take the value $\mathbb{E}^{a^*} [d\xi_t - a_t^* dt | \mathcal{F}_t] = \mathbb{E}^{a^*} [\theta_t | \mathcal{F}_t] dt = p_t^* dt$.

The manager is risk neutral and affecting earnings entails private costs captured by $\psi a_t^2/2$, $\psi > 0$. In addition, he is rewarded according to a wage process $(\chi(p_t^*))_{t \geq 0}$ with $\chi(\cdot)$ strictly increasing, and thus managers who run firms that are perceived to have better fundamentals receive higher wages.³² Observe that $\chi' > 0$ implies that the manager always has a myopic incentive to inflate earnings.

The model represents a situation in which a manager, in any period $[t, t + dt)$, can influence an accounting division using only the information that he has up to time t ; i.e., before the financial information over $[t, t + dt)$, $dY_t = \theta_t dt + \sigma_\xi dZ_t^\xi$, is processed by such division. This attempts to capture a firm with a strong internal control system that limits the management’s direct involvement in the creation of financial statements, but that is not invulnerable to management pressures. The manager then learns about the firm’s profitability over $[t, t + dt)$ when a report $d\xi_t$ is produced by the firm’s accounting department (moment at which he infers $dY_t = d\xi_t - a_t dt$); but once this occurs, the report cannot be eliminated or modified before releasing it to the public. Finally, $\psi a^2/2$ captures that persuading the accounting division to inflate earnings by a can be costly at increasing rates (e.g., convex opportunity cost of resources allocated to this practice, or reluctance to engage in “creative” accounting).

4.3.1 Linear Benchmark

Suppose that the manager’s flow payoff is linear according to $\chi(p^*) = \alpha p^*$, $\alpha > 0$. In this case, the ratcheting equation (14) admits a constant solution given by $q(p) = \alpha/(r + \beta)$. As a result, $a^* = (g')^{-1}(\beta q(p)) = \beta \alpha / \psi (r + \beta)$, where I used that $g'(a) = \psi a$. Because actions are constant in this equilibrium, the endogenous ratchet $\beta da^*/dp^*$ has no effect on incentives.

Given any nonlinear $\chi(\cdot)$, it is then natural to define its *linear benchmark policy* as

$$p \mapsto \frac{\beta \chi'(p)}{\psi (r + \beta)}.$$

In fact, if the market’s belief is p , $\beta \chi'(p)/\psi (r + \beta)$ captures the incentives that would arise in a linear environment of constant myopic incentives given by $\alpha = \chi'(p)$, $p \in \mathbb{R}$. As I show next, this policy is a useful benchmark for illustrating the non-trivial effect that the endogenous ratchet $\beta da^*/dp^*$ can have on incentives in settings where nonlinearities are present.

To conclude this subsection, observe that, in this linear case: (i) as the strength of the

³²If $(\theta_t)_{t \geq 0}$ represents managerial ability, p_t^* is a measure of the manager value: the market expects future performance to take the form $\mathbb{E}^{a^*} \left[\int_t^\infty e^{r(s-t)} (d\xi_s - a_s^* ds) \middle| \mathcal{F}_t \right] = \mathbb{E}^{a^*} \left[\int_t^\infty e^{r(s-t)} p_s^* ds \middle| \mathcal{F}_t \right] = p_t^*/r$ at $t \geq 0$. The independence of $\chi(\cdot)$ from a^* reflects that the market tolerates some degree of earnings management.

manager's myopic incentives α increases, earnings are inflated more aggressively; and (ii) managers of different firms should exert the same degree of manipulation regardless of the performance of the individual firms they operate.³³

4.3.2 Nonlinear Flow Payoffs: The Importance of Thresholds

There is a large body of evidence contradicting that earnings management is uniform across different levels of performance. In particular, it has been documented that manipulation is particularly strong around some key *thresholds* or *benchmarks*: managers try to avoid (i) reporting losses, (ii) reporting negative earnings growth, and (iii) failing to meet analyst forecasts.³⁴ To capture such incentives, I consider a single-peaked marginal utility function:

Assumption 2. $\chi \in C^3(\mathbb{R})$. χ' is strictly positive, symmetric around zero, and strictly increasing in $(-\infty, 0)$, with $\chi'(p) \rightarrow 0$ as $p \rightarrow -\infty$. Also, $\chi'''(0) < 0$.

As in the linear case, the manager has a myopic incentive to boost reported earnings across all levels of performance ($\chi' > 0$). However, this incentive is now stronger when the market expects the firm to generate zero true earnings over the next period (i.e., when $p_t^* = 0$). I refer to this level of earnings as the *zero-earnings threshold*.³⁵

Observe that, by positively influencing the market's belief, a manager standing at $p^* < 0$ would face stronger myopic incentives compared to a manager standing at the corresponding symmetric point $|p^*| > 0$, thus suggesting that the incentives to inflate earnings should be stronger to the left of zero. However, recall that managers cannot succeed at misleading the market in equilibrium. In addition, when $\kappa = 0$, beliefs evolve as a martingale along the path of play, and hence, they are unpredictable. Since the manager's myopic incentives are strongest at $p^* = 0$, the equilibrium behavior of the public belief suggests that manipulation should be maximized at the zero-earnings threshold. No conclusive answer is obtained by appealing to traditional dynamic programming arguments or looking at primitives.

The next result uses an existence result from Section 6.2 pertaining to *bounded* solutions to the ratcheting equation. Interestingly, the earnings management policy in any such equilibrium has a different structure from the previously ones just discussed.

³³See Stein (1989) for another linear model of earnings management in which equilibrium behavior is independent of performance but where manipulating earnings entail real costs to the firm.

³⁴See Burgstahler and Dichev (1997), Degeorge et al. (1999), and Dichev et al. (2013) for statistical and survey-based approaches to identify this type of practice.

³⁵The symmetry assumption is to illustrate distortions more clearly, whereas the strict concavity of χ' at zero stresses that the myopic incentive at the threshold is acute. This type of marginal incentives can be microfounded through a smooth *S*-shaped contract whose slope is maximized at $p^* = 0$.

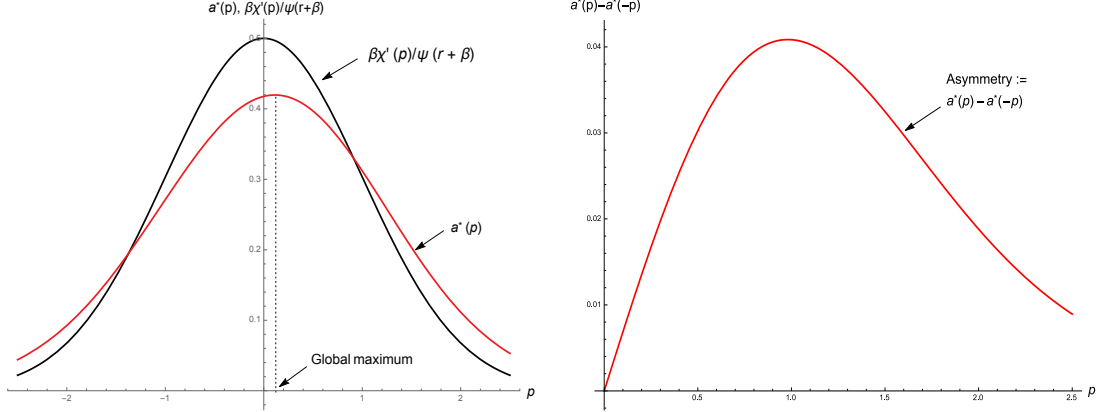


Figure 2: Left panel: the equilibrium policy $a^*(\cdot)$ and the linear benchmark $\beta\chi'(p)/\psi(r+\beta)$ around zero. Right panel: equilibrium policy degree of skewness to the right, as measured by $p \mapsto a^*(p) - a^*(-p)$. $a^*(\cdot)$ is constructed via a solution $q(\cdot)$ to the ratcheting equation on the truncated domain $[-10, 10]$ with parameter values: $r = \sigma_\xi = \sigma_\theta = \psi = 1$ and $\chi'(p^*) = e^{-0.5p^{*2}}$.

Proposition 5. *In any Markov equilibrium $a^*(p) = \beta q(p)/\psi$, with $q(\cdot)$ a bounded solution to the ratcheting equation:³⁶ (i) $q \in (0, \chi'(0)/(r + \beta))$ and $q(p) \rightarrow 0$ as $p \rightarrow \pm\infty$; (ii) $q'(0) > 0$ and $q''(0) < 0$; (iii) q is maximized strictly to the right of zero; and (iv) the manipulation policy is skewed to the right of zero, i.e., $q(p) \geq q(-p)$ for all $p > 0$.*

Proof: See the Appendix. □

Proposition 5 uncovers two interesting distortions around the threshold. First, incentives are depressed at zero relative to the corresponding linear benchmark of slope $\alpha = \chi'(0)$ (i.e., $q(0) < \chi'(0)/(r + \beta)$). Second, the policy is maximized to the right of zero, despite those managers' weaker myopic incentives than those exactly at the threshold, and despite them being unable to truly affect the value of the firms they operate. See figure 2.

These two distortions are the consequence of the endogenous ratcheting costs imposed by the market. To see this, observe that since the manager's myopic incentives become stronger as p approaches zero from the left, the market will conjecture a strictly increasing manipulation profile in this region ($da^*/dp^* > 0$). Incentives thus fall below the linear benchmark in a neighborhood to the left of zero, as successfully influencing the market's belief leads firms that are likely to fail meeting the threshold to face a more demanding incentive scheme. Now, suppose that the market's conjecture is actually maximized at zero. In this case, anticipating that the market would revise its conjecture of manipulation downwards to the right of zero, the manager would be incentivized to boost earnings at the zero-earnings threshold. The market must therefore ratchet up its expectation of behavior

³⁶An equilibrium of this sort will exist if $\chi'(0) < \psi\sqrt{2r\sigma_\xi^2(r\sigma_\xi + \sigma_\theta)^2/4\sigma_\theta^2}$. See Theorem 3.

at zero ($q'(0) > 0$) to assess the firm's fundamentals correctly, resulting in a policy that is maximized to the right of zero.

Furthermore, the consequences of this endogenous ratcheting extend throughout the entire belief space, yielding a manipulation profile that is skewed to the right of the threshold, i.e., $q(p) > q(-p)$, $p > 0$. In fact, observe that in the absence of $\beta da^*/dp^* = \beta^2 q'(p)/\psi$ on the left of (14), the ratcheting equation becomes $(r + \beta)q(p) = \chi'(p) + \sigma^2 q''(p)/2$. By the symmetry of χ' , however, this ODE admits a symmetric solution around zero: cost smoothing (q'') and the ratchet β thus only affect the level of the incentives created, not their skewness. The ratchet effect discovered is thus more subtle than in the previous applications: the incentives to maintain high earnings are, on average, stronger than the incentives to build them up.³⁷

5 Equilibrium Analysis: Sufficient Conditions

This section establishes conditions that validate the use of the ratcheting equation for finding Markov equilibria. More precisely, I present a *verification theorem* that involves the system of ODEs (14)–(15): if there is a pair (q, U) that solves the previous system and satisfies a particular second-order condition, then $\rho(\beta q(\cdot))$ is a Markov equilibrium.

Before stating the theorem, it is instructive to illustrate the importance of the first-order approach relative to traditional dynamic programming methods. To this end, observe first that if the market conjectures a Markov strategy $a^*(\cdot)$, the public and private beliefs follow

$$\begin{aligned} dp_t^* &= [-\kappa(p_t^* - \eta) + \beta(p_t - p_t^*) + \beta(a_t - a^*(p_t^*))]dt + \sigma dZ_t, \quad \text{and} \\ dp_t &= -\kappa(p_t - \eta)dt + \sigma dZ_t, \end{aligned}$$

respectively. Thus, both states are Markov, and $(p_t^*)_{t \geq 0}$ is controlled by the long-run player.³⁸ It follows that given any private history (ξ^t, a^t) , the current value of the pair (p_t, p_t^*) contains all the information that is relevant for future decision making. In other words, the long-run player's best-response problem to $a^*(\cdot)$ becomes a standard stochastic control problem.

A simple modification of this best-response problem makes dynamic programming meth-

³⁷The following additional properties depicted in the left panel of figure 1 are established in the proof of Proposition 5: q is strictly decreasing to the right of the global maximum; at the global maximum, q is below its linear counterpart; there exists a point to the right of the maximum where q changes from being concave to convex; and, finally, if q is convex and decreasing, q is above its linear counterpart.

³⁸From Lemma 1, the dynamic for $(p_t^*)_{t \geq 0}$ follows directly from $dp_t^* = -\kappa(p_t^* - \eta)dt + \beta\sigma_\xi[d\xi_t - (a^*(p_t^*) + p_t^*)dt]$ and $d\xi_t = (a_t + p_t)dt + \sigma_\xi dZ_t$ from the long-run player's perspective.

ods also applicable to the equilibrium problem. More precisely, if $V(p, p^*)$ solves

$$\begin{aligned}
rV(p, p^*) &= \sup_{a \in A} \left\{ u(\chi(p^*, a^*(p^*))) - g(a) - \kappa(p - \eta)V_p(p, p^*) + \frac{\sigma^2}{2}V_{pp}(p, p^*) + \sigma^2V_{pp^*}(p, p^*) \right. \\
&\quad \left. + [-\kappa(p^* - \eta) + \beta(p - p^*) + \beta(a - a^*(p^*))]V_{p^*}(p, p^*) + \frac{\sigma^2}{2}V_{p^*p^*}(p, p^*) \right\} \quad (21) \\
s.t. \quad \arg \max_{a \in A} \{ a\beta V_{p^*}(p, p) - g(a) \} &= a^*(p) \quad (22)
\end{aligned}$$

(subject to transversality conditions), then $V(p, p^*)$ is the long-run player's value function, and $a^*(p) = \rho(\beta V_{p^*}(p, p))$ is a Markov equilibrium (provided that the induced public strategy is feasible). In fact, constraining the HJB equation (21) to (22) ensures that the market anticipates the long-run player's behavior when beliefs are aligned. Notice, however, that solving (21)–(22) leads to perfect knowledge of $(p, p^*) \mapsto \rho(\beta V_{p^*}(p, p^*))$, the long-run player's optimal strategy on and off the equilibrium path. In other words, the HJB approach, through the exact computation of off-equilibrium payoffs, implicitly requires the full specification of an MPE to determine on-path behavior.

The key difficulty with (21)–(22) is that it does not correspond to a standard PDE, as $V(p, p^*)$ depends on $V_{p^*}(p^*, p^*)$: that is, this differential equation has a *non-local* structure, in the sense that it involves the unknown and one of its derivatives each evaluated at different points. This is because the full-support monitoring technology leads the market to believe that the long-run player is always on-path taking actions according to $\rho(\beta V_{p^*}(p^*, p^*))$. If a deviation occurs, however, he can also condition his actions on his private belief, and $\rho(\beta V_{p^*}(p, p^*))$ and $\rho(\beta V_{p^*}(p^*, p^*))$ need not coincide; in those cases, incentives for *double deviations* (i.e., deviations after deviations) appear. To the best of my knowledge, no general existence theory for this type of equation is available.³⁹

The technical importance of the ratcheting equation thus lies on opening a venue to find Markov equilibria bypassing all these difficulties. In fact, as a self-contained object, this equation suggests that computing off-path payoffs exactly is by no means strictly necessary to determine equilibrium behavior. This is confirmed in the next subsection.

Remark 6 (The ratcheting equation as a necessary condition for MPE). Applying the envelope theorem to (21)–(22) yields that $q(p) := V_{p^*}(p, p)$, which characterizes on-path incentives via (22), satisfies the ratcheting equation (14). Thus, the ratcheting equation is a necessary condition for MPE when the value function is sufficiently differentiable. The ODE

³⁹In the linear case, there is solution to (21)–(22) that is additively separable in p and p^* , so the non-localness disappears: $V_{p^*}(p^*, p^*)$ becomes a constant, and is thus independent of (p^*, p^*) . The economic implication is that the long-run player finds it optimal to take the same action on and off the path of play. Observe also that the technical complexities just described are present in any setting involving two-sided learning and imperfect monitoring in which one belief can be actively controlled.

for $U(p) := V(p, p)$ then follows from evaluating (21)–(22) at $p = p^*$.

5.1 A Verification Theorem for Markov Equilibria

As before, let $a^*(p) := \rho(\beta q(p))$ with $q(\cdot)$ a solution to the ratcheting equation. Off the path of play, the long-run player can condition his actions on both $(Z_t)_{t \geq 0}$ (which drives $(p_t)_{t \geq 0}$) and $(\Delta_t)_{t \geq 0}$; denote by $\mathbb{F}^{(Z, \Delta)} := (\mathcal{F}_t^{Z, \Delta})_{t \geq 0}$ the corresponding filtration. In this context, a strategy $(\hat{a}_t)_{t \geq 0}$ is said to be *admissible* if (i) it is $\mathbb{F}^{(Z, \Delta)}$ -progressively measurable; (ii) $\mathbb{E} \left[\int_0^t (\hat{a}_s)^2 ds \right] < \infty$, $t \geq 0$; and (iii) $\mathbb{E} \left[\int_0^\infty e^{-rt} |u(\chi(p_t + \hat{\Delta}_t^*, a^*(p_t + \hat{\Delta}_t^*))) - g(\hat{a}_t)| dt \right] < \infty$, where $(\hat{\Delta}_t^*)_{t \geq 0}$ denotes the solution to (8) under the pair $(a^*(\cdot), (\hat{a}_t)_{t \geq 0})$.⁴⁰

Theorem 1 (Verification Theorem). *Suppose that (q, U) of class $C^2(\mathbb{R})$ solves (14)–(15) and that $a^*(\cdot) := \rho(\beta q(\cdot))$ is interior. Moreover, assume that*

- (i) *There exist C_1, C_2 , and $C_3 > 0$ such that $|U(p)| \leq C_1(1 + |p|^2)$ (quadratic growth), $|U'(p)| \leq C_2(1 + |p|)$ (linear growth), and $|q(p_1) - q(p_2)| \leq C_3|p_1 - p_2|$ (Lipschitz);*
- (ii) $\lim_{t \rightarrow \infty} \mathbb{E}[e^{-rt} U(p_t + \hat{\Delta}_t^*)] = \lim_{t \rightarrow \infty} \mathbb{E}[e^{-rt} q(p_t + \hat{\Delta}_t^*)(\hat{\Delta}_t^*)] = \lim_{t \rightarrow \infty} \mathbb{E}[e^{-rt} U'(p_t + \hat{\Delta}_t^*)(\hat{\Delta}_t^*)] = 0$
for all $(\hat{a}_t)_{t \geq 0}$ admissible, where $(p_t, \hat{\Delta}_t^)_{t \geq 0}$ is the solution to (5)–(8) under $(a^*(\cdot), (\hat{a}_t)_{t \geq 0})$;*
- (iii) *U'' and q' satisfy*

$$|U''(p) - q'(p)| \leq \frac{\psi(r + 4\beta + 2\kappa)}{4\beta^2}, \text{ for all } p \in \mathbb{R}. \quad (23)$$

Then, if $(\rho(\beta q(p_t^)))_{t \geq 0}$ is feasible, $a^*(\cdot) = \rho(\beta q(\cdot))$ is a Markov equilibrium and $U(\cdot)$ its corresponding payoff.*

Proof: See the Appendix. □

Theorem 1 offers a method for finding Markov equilibria that relies on solving a system of equations for (q, U) , in the same way that traditional verification theorems in dynamic programming offer a way to find optimal policies by solving HJB equations. The advantages of this theorem are clear: first, it bypasses all the difficulties encountered in attempting to find

⁴⁰This is the natural notion of admissibility that allows actions to depend on information off the path of play. In general, by definition of $(Z_t)_{t \geq 0}$, its corresponding filtration is only guaranteed to be *contained* in the public filtration. Under a Markov conjecture, however, this notion of admissibility is without loss, as $(Z_t)_{t \geq 0}$ (via $(Y_t)_{t \geq 0}$; Section 3.1), carries all the exogenous information coming from $(\xi_t)_{t \geq 0}$ that enters in the public belief, whereas $(\Delta_t)_{t \geq 0}$ carries the endogenous counterpart pertaining to past play. See Section 7.4.2 in Liptser and Shiryaev (1977, p.276) for a discussion of when such filtrations coincide (in particular, because $(Y_t, \theta_t)_{t \geq 0}$ is Gaussian and $p_t := \mathbb{E}[\theta_t | (Y_s)_{s \leq t}]$, the filtrations of $(Y_t)_{t \geq 0}$ and $(Z_t)_{t \geq 0}$ are identical).

equilibria by computing off-path payoffs exactly; second, it is general, with weak restrictions on payoffs; and third, by involving a system of ODEs, the derivation of qualitative properties of equilibrium behavior becomes considerably simpler.

Regarding the assumptions in the Theorem, the Lipschitz condition on $q(\cdot)$ ensures that a solution to $(p_t^*)_{t \geq 0}$ (and, hence, to $(\Delta)_{t \geq 0}$) exists and is unique under (i) and (ii) of the admissibility concept, and thus, the long-run player's best-response problem is well defined. The conditions on U in turn guarantee that the U -ODE (15) has as a unique solution

$$U(p) := \mathbb{E} \left[\int_0^\infty e^{-rt} [u(\chi(p_t, \rho(\beta q(p_t)))) - g(\rho(\beta q(p_t)))] dt \mid p_0 = p \right],$$

i.e., the payoff from following the market's conjecture $a^*(\cdot) = \rho(\beta q(\cdot))$. The rest of the assumptions are used to construct an upper bound for the long-run player's value function, with the property that it coincides with $U(\cdot)$ on the equilibrium path; but since $U(\cdot)$ can be achieved by following $a^*(\cdot)$ when beliefs are aligned, it follows that inducing no belief asymmetry is optimal. Finally, the feasibility requirement demands that (2) has a unique solution under the public strategy induced by $a^*(\cdot)$ so as to ensure that the outcome of the game is well and uniquely defined. This requirement is verified in the two classes of games studied in Section 6.⁴¹

Condition (23) has an economic interpretation: it is a bound on the rate of change of $U'(p) - q(p)$, and the latter is an *information rent*; i.e., a measure of the value of having private information about the fundamentals. To see why, recall that both U' and q incorporate the costs of matching the changes in the market's conjecture of equilibrium behavior. In addition, $q(p)$ incorporates the cost of adjusting behavior to explicitly account for belief asymmetry; this is absent in $U'(p)$ because beliefs are aligned along the path of play. The difference $q(p) - U'(p)$ is thus a measure of the value of having private information about the fundamentals in the form of a (marginally) more *pessimistic* private belief; $U'(p) - q(p)$ is then the analog measure for the case of a (marginally) more *optimistic* private belief.

Rates of change of information rents appear when approximating the payoff of an alternative strategy using information rents to estimate the continuation value following a partial deviation from $a^*(\cdot)$. This estimation procedure is appropriate because the long-run player has private information off the path of play, and hence, estimating the payoff of an

⁴¹The Lipschitz property of a Markov strategy is strong enough to suggest that verifying the feasibility requirement can be potentially dispensed with in the theorem: namely, by showing that feasibility holds for any Lipschitz $a^*(\cdot)$. This in turn amounts to examining the question of existence of a weak solution to (2) under the corresponding progressively measurable functional $a^*(p_t^*[\cdot]) : C(\mathbb{R}_+) \rightarrow A$ (with $p_t^*[\xi]$ the solution of the public belief process under a $a^*(\cdot)$ as a function of the realized public history) by means of the Girsanov theorem. Refer to Section 5.3.B in Karatzas and Shreve (1991) for more on this topic.

alternative strategy requires accounting for the value of such private information; but the only information available in this procedure is the one conveyed by (q, U) solving (14)–(15). Moreover, as the long-run player’s private belief changes, his continuation payoff varies, and so the payoff of an alternative strategy necessarily depends on $U'' - q'$. When (23) holds, it can be ensured that it is never optimal for the long-run player to induce any degree of belief asymmetry, and thus $U'' - q'$ can be seen as a global measure of the value of such private information. The presence of an absolute value simply reflects the possibility of upward or downward deviations being profitable depending on how the market rewards the long-run player. Finally, (23) can be relaxed, depending on the primitives of the environment at hand; this is the case in the class of linear-quadratic games presented in the next section.

Remark 7 (On the first-order approach in optimal contracting). Similar sufficient conditions have also appeared in the optimal contracting literature, with the corresponding measures of information rents in the form of stochastic processes rather than ODEs, due to the non-Markov nature of such environments. In particular, this paper has built on Williams (2011) and Sannikov (2014) who find sufficient conditions that validate their first-order approaches via using measures of information rents in quadratic bounds for payoffs after deviations.⁴² By recognizing that a full specification of the long-run player’s best response is only sufficient for the analysis of equilibrium outcomes, this paper shows that the methods from the optimal contracting literature are applicable to games of learning and unobserved actions. Importantly, Section 6 goes a step further relative to this literature by showing that (23) can be mapped to primitives for a wide range of settings.

6 Existence of Markov Equilibria

This section uses Theorem 1 to derive two existence results for Markov equilibria that can be computed using the ratcheting equation. This is conducted by (i) proving the existence of solutions to the system (14)–(15) and (ii) verifying that the second-order condition (23) holds. The advantage of the verification theorem is that, by involving ODEs, (23) can be verified on primitives in many economic environments. (It can always be verified ex post.)

The environments under study are (i) linear quadratic games and (ii) games with bounded marginal flow payoffs. In the former, equilibrium behavior is linear in the public belief. In the latter, equilibria are fully nonlinear. As argued earlier, to focus on the question of existence

⁴²Their second-order conditions are one-sided only. This is because Williams (2011) allows for downward deviations only, and in Sannikov (2014) actions do not affect the public signal directly. My construction is closest to Sannikov’s (2014).

of equilibria when the ratchet principle appears in its most traditional form, the next two sections restrict the analysis to the case in which the market's action χ is independent of a^* .

6.1 Linear Quadratic Games

Definition 4. *The environment is said to be linear quadratic if $A = \mathbb{R}$; $g(a) = \frac{\psi}{2}a^2$, $\psi > 0$; and $h(p^*) := u(\chi(p^*, a^*)) = u_0 + u_1p^* - u_2p_t^{*2}$, where $u_0, u_1 \in \mathbb{R}$ and $u_2 \geq 0$.*

The next result shows the existence of a linear Markov equilibrium.

Theorem 2. *Consider a linear quadratic environment. A linear $q(\cdot)$ and a quadratic $U(\cdot)$ solving (14) and (15), respectively, exist if and only if*

$$u_2 \leq \frac{\psi(r + \beta + 2\kappa)^2}{8\beta^2}. \quad (24)$$

In this case, $a^*(p) = \beta[q_1 + q_2p]/\psi$, with

$$q_1 = \frac{\eta\kappa q_2 + u_1}{r + \beta + \kappa + \frac{\beta^2}{\psi}q_2} \quad \text{and} \quad q_2 = \frac{\psi}{2\beta^2} \left[-(r + \beta + 2\kappa) + \sqrt{(r + \beta + 2\kappa)^2 - \frac{8u_2\beta^2}{\psi}} \right] < 0,$$

is a linear Markov equilibrium.

Proof: See the Appendix. □

The linear equilibrium found entails a negative slope q_2 and an intercept q_1 that can take any sign. Suppose that $u_1 = 0$, so flow payoffs are maximized at zero. In this case, for large $|p|$, the long-run player has a myopic incentive to drive the public belief towards the bliss point, and thus, $q_2 < 0$. Interestingly, when $\eta > 0$, q_1 is negative, so the long-run player puts downward pressure on the public belief at the bliss point. The reason is cost smoothing: if $p^* = 0$, the long-run player expects the public belief to revert to $\eta > 0$ with high probability, but when $p^* \in [0, \eta]$, driving p^* back to zero is optimal. Because the disutility of effort is convex, it is optimal to set $a^*(0) < 0$ to distribute such costs optimally over time.

The curvature condition (24) corresponds to a necessary and sufficient condition for a linear solution to the ratcheting equation (14) to exist in the first place; thus, its violation is not an indication that a linear equilibrium ceases to exist due to the value of acquiring private information about the fundamentals becoming too large.⁴³ Instead, the existence

⁴³Recall that verification theorem requires a solution (q, U) to the system (14)–(15) to exist to make the second-order condition (23) applicable. For linear quadratic games, however, the curvature condition (24) is the relevant constraint, as when this one is satisfied, a relaxed version of (23) tailored for this particular class of games can be verified to hold.

problem operates through the endogenous ratcheting channel.

To see this, notice that as u_2 grows, the myopic incentive to drive the public belief towards the bliss point increases in absolute value. Thus, in a linear equilibrium, the market must impose a steeper conjecture $a^*(\cdot)$ to control such incentives. A steeper conjecture, however, also makes the market revise its expectation of performance more drastically given any changes in beliefs. Consequently, an upward (downward) deviation to the left (right) of zero becomes more attractive, as it now leads to a more rapid decrease in the market's expectation of performance tomorrow; in the ratcheting equation, this is captured by the left-hand side $r + \kappa + \beta + \beta da^*/dp^*$ decreasing as da^*/dp^* becomes more negative. Thus, if (24) is violated, a linear conjecture cannot control both (i) myopic incentives and (ii) the incentives to induce the market to ratchet down its expectations of effort.

Finally, if flow payoffs are linear, the curvature condition is trivially satisfied, and thus, a linear equilibrium always exists. In this case, $u_2 = 0$ leads to $q_2 = 0$ in Theorem 2, and hence, to an equilibrium with constant actions defined by

$$q(p) = q_1 = \frac{u_1}{r + \beta + \kappa} \Rightarrow g'(a^*) = \beta \frac{u_1}{r + \beta + \kappa}.$$

Moreover, a linear $U(\cdot)$ solves (15) in this case, which yields $U'' - q' \equiv 0$. Put differently, changes in the long-run player's private information have no value for the long-run player, which is consistent with the equilibrium level of effort found by Holmström (1999) being also optimal off the path of play.

6.2 Bounded Marginal Flow Payoffs

Definition 5. Let $h(p) := u(\chi(p))$. A game is one of bounded marginal flow payoffs if

(i) (Boundedness) $\exists m, M \in \mathbb{R}$ s.t. $-\infty < m := \inf_{p \in \mathbb{R}} h'(p) \leq \sup_{p \in \mathbb{R}} h'(p) := M < \infty$, and

(ii) (Interior actions) A is compact and $\{\beta x \mid x \in [m/(r + \beta + \kappa), M/(r + \beta + \kappa)]\} \subseteq g'(A)$.

I now show that there exists a solution (q, U) to (14)–(15). A solution to (14) satisfying that both q and q' are bounded will be referred to as a *bounded solution*—the focus will be on this type of solution.⁴⁴

⁴⁴Searching for a bounded q is natural given that marginal flow payoffs are bounded. A bounded q' in turn ensures that $a^*(p) := \rho(\beta q(p))$ is Lipschitz and that U' satisfies the transversality and growth conditions of Theorem 1.

Proposition 6 (Existence of Bounded Solutions to the Ratcheting Equation).

There exists $q \in C^2(\mathbb{R})$, a bounded solution to the ratcheting equation, such that

$$q(p) \in \left[\frac{m}{r + \beta + \kappa}, \frac{M}{r + \beta + \kappa} \right], \text{ for all } p \in \mathbb{R}. \quad (25)$$

If, in addition, $h'_+ := \lim_{p \rightarrow \infty} (u \circ \chi)'(p)$ and $h'_- := \lim_{p \rightarrow -\infty} (u \circ \chi)'(p)$ exist, any bounded solution satisfying (25) also verifies that

$$\lim_{p \rightarrow +\infty} q(p) = \frac{h'_+}{r + \beta + \kappa} \text{ and } \lim_{p \rightarrow -\infty} q(p) = \frac{h'_-}{r + \beta + \kappa}. \quad (26)$$

Proof: See the Appendix. □

Bound (25) states that there is a candidate equilibrium that lies in between the equilibria that would arise in environments of linear flow payoffs with slopes m and M . The second part of the proposition in turn asserts that, in settings where marginal flow payoffs become asymptotically constant (as in the earnings management application of Section 4.3), equilibrium behavior converges to the corresponding limit (linear) counterpart (i.e., $u(\chi(p)) = h'_+p$ or h'_-p) as $p \rightarrow \pm\infty$. While this asymptotic property of payoffs is not required for the existence results presented below, it provides useful guidance as to which type of “boundary” conditions to expect when searching for a solution to the second-order ODE (14).

Proposition 7 (Long-Run Player’s Equilibrium Payoff). *Let q denote a bounded solution to (14). The unique solution to the ODE (15) is given by*

$$U(p) = \mathbb{E} \left[\int_0^\infty e^{-rt} [h(p_t) - g(\rho(\beta q(p_t)))] dt \mid p_0 = p \right], \quad (27)$$

where $dp_t = -\kappa(p_t - \eta)dt + \sigma dZ_t$ for $t > 0$ and $p_0 = p$. Furthermore, U has linear growth, and U' is bounded.

Proof: See the Appendix. □

Finally, I establish conditions on the primitives $(r, m, M, \psi, \kappa, \sigma_\theta, \sigma_\xi)$ that ensure that (q, U) as above meets the requirements of Theorem 1. I do this for the case $\kappa = 0$, which simplifies the estimation of the rate of change of information case, as I explain shortly.

Theorem 3 (Existence of Markov Equilibrium). *Suppose that $\kappa = 0$, and let $q : \mathbb{R} \rightarrow \mathbb{R}$*

denote a bounded solution to the ratcheting equation (14). If

$$\frac{M - m}{\psi} \leq \frac{\sqrt{2r\sigma_\xi^2}(r + \beta)^2}{4\beta^2} = \frac{\sqrt{2r\sigma_\xi^2}(r\sigma_\xi + \sigma_\theta)^2}{4\sigma_\theta^2}, \quad (28)$$

$a^*(\cdot) := \rho(\beta q(\cdot))$ is a Markov equilibrium.

Proof: See the Appendix. □

Theorem 3 proves the existence of equilibria in which behavior is a nonlinear function of the common belief for a wide range of economic environments. Condition (28) is relaxed when the public signal is noisy (σ_ξ is large) and when the environment is less uncertain (σ_θ is small), as in this case, beliefs become less responsive to signal surprises. The condition is also relaxed when affecting the public signal is costly (ψ is large), the long-run player is impatient (r is large), and when $M - m$ falls. It is also trivially satisfied when payoffs are linear ($M = m$), as information rents are constant in this case (leading to $U'' - q' \equiv 0$).

The main challenge in the proof of Theorem 3 is the estimation of the rate of change of information rents in terms of primitives. When $\kappa = 0$, the analysis is simplified by the fact that $|U'' - q'|$ can be expressed as an analytic function of q , and the bounds for q in terms of primitives follow from Proposition 6. To the best of my knowledge, no such analytic solution exists when $\kappa > 0$, thus making the estimation more complex.⁴⁵

Importantly, mean reversion can be only expected to reduce the attractiveness of any deviation. In fact, as κ increases, beliefs become less sensitive to new information ($\beta(\kappa) = (\kappa^2 + \sigma_\theta^2/\sigma_\xi^2)^{1/2} - \kappa$ falls with κ), and any shock to beliefs decays faster; a marginal increase in effort has then a smaller and shorter impact on flow payoffs. Moreover, since belief asymmetry decays at rate $\beta + \kappa = (\kappa^2 + \sigma_\theta^2/\sigma_\xi^2)^{1/2}$, the long-run player's informational advantage also disappears faster as κ grows. While higher rates of mean reversion are likely to reduce the strength of the ratcheting performed by the market (as beliefs become more i.i.d.), ratcheting appears only to the extent that beliefs are revised in response to new information, and hence, it is likely to be dominated by the sensitivity effect. This is confirmed in linear environments where incentives are characterized by $\beta q(p) = \beta(\kappa)/(r + \beta(\kappa) + \kappa)$: while the exogenous ratchet term $\beta(\kappa)$ in the denominator decreases with κ , incentives still decay due to beliefs becoming less responsive to signal surprises.

⁴⁵Specifically, $U'' - q'$ satisfies a type of ODE whose solutions are constructed using *confluent hypergeometric functions*, which take the form of power series (cf. Abramowitz and Stegun, 1964).

7 Conclusions

This paper has examined a class of continuous-time games of learning and imperfect monitoring. The contribution is twofold. First, the analysis executed uncovered a learning-driven version of the ratchet principle that naturally appears in settings characterized by common uncertainty and strategic behavior. Second, this paper expanded the class of economic questions that can be studied under the umbrella of signal jamming models beyond linear settings. The applications developed are in fact at the intersection of these areas: they explore ratchet effects in settings that exhibit nonlinearities.

Assuming *ex ante* symmetric uncertainty is a convenient modeling technique to analyze incentives in settings where uncertainty is common to everyone. If the long-run player had *ex ante* superior information about the fundamentals, for instance, his actions could potentially incorporate his private information. Beyond linear quadratic games (linear learning, linear quadratic payoffs) or settings in which the fundamentals take finite values, the long-run player's action would then be a nonlinear function of his private information, which makes handling beliefs technically challenging.⁴⁶ Relatedly, necessary and sufficient conditions for Markov equilibria away from stationary learning can be obtained using identical arguments to the ones employed here. The challenge then becomes to show the existence of solutions to a version of the ratcheting equation that also depends on time as a state.

Finally, I discuss three possible extensions of the model. First, since in any nonlinear Markov equilibrium, actions are a nonlinear function of the complete history of signal realizations, the class of distributions that can be generated for the public signal is quite rich (and not necessarily Gaussian). Thus, the model has the potential to be used to address empirical questions in environments with inherent nonlinearities. Second, pure-strategy equilibria beyond the Markov case could be studied as well, the main difference being that the corresponding necessary and sufficient conditions would involve stochastic processes rather than ODEs. Since Markov equilibria are already a function of the complete public history via the public belief, it is unclear whether this extension produces any new insights. Third, a natural extension involves studying incentives in environments where affecting the informativeness of the public signal is possible, such as when there are complementarities between the fundamentals and actions. While the first-order approach followed here is still applicable in settings beyond the additively separable world, the analysis is complicated by additional experimentation effects. These and other questions are left for future research.

⁴⁶See Bonatti, Cisternas and Toikka (2016) for a linear quadratic model of oligopoly with imperfect public monitoring in which firms have private information regarding their constant marginal costs.

References

- [1] Abramowitz, M. and I. Stegun (1964): *Handbook of Mathematical Functions, with Formulas, Graphs, and Mathematical Tables*. New York: Dover.
- [2] Atkeson, A., V.V. Chari and P. Kehoe (2007): “On the Optimal Choice of a Monetary Policy Instrument,” Federal Reserve Bank of Minneapolis Staff Report 394.
- [3] Bar-Isaac, H. and J. Deb (2014): “What is a Good Reputation? Career Concerns with Heterogeneous Audiences,” *International Journal of Industrial Organization*, 34, 44-50.
- [4] Bergemann, D. and U. Hege (2005): “The Financing of Innovation: Learning and Stopping,” *RAND Journal of Economics*, 36(4), 719-752.
- [5] Bhaskar, V. (2014): “The Ratchet Effect Re-Examined: A Learning Perspective,” Working Paper, UCL.
- [6] Bhaskar, V. and G. Mailath (2016): “The Curse of Long Horizons,” Working Paper, University of Pennsylvania.
- [7] Board, S. and M. Meyer-ter-Vehn (2014): “A Reputational Theory of Firm Dynamics,” Working Paper, UCLA.
- [8] Bohren, A. (2016): “Using Persistence to Generate Incentives in a Dynamic Moral Hazard Problem,” Working Paper, University of Pennsylvania.
- [9] Bonatti, A., G. Cisternas and J. Toikka (2016): “Dynamic Oligopoly with Incomplete Information,” *Review of Economic Studies*, forthcoming.
- [10] Bonatti, A. and J. Hörner (2011): “Collaborating,” *American Economic Review*, 101(2), 632-663.
- [11] Bonatti, A., and J. Hörner (2016): “Career Concerns with Exponential Learning,” *Theoretical Economics*, forthcoming.
- [12] Burgstahler, D. and I. Dichev (1997): “Earnings management to avoid earnings decreases and losses,” *Journal of Accounting and Economics*, 24(1), 99-126.
- [13] Cogley, T., G. Primiceri and T. Sargent (2010): “Inflation-Gap Persistence in the US,” *American Economic Journal: Macroeconomics*, 2(1), 43-69.
- [14] Cohen, S. and R. Elliott (2015): *Stochastic Calculus and Applications, Second Edition*. New York: Birkhäuser.

- [15] Cukierman, A. and A. Meltzer (1986): “A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information,” *Econometrica*, 54(5), 1099-1128.
- [16] Da Prato, G. and J. Zabczyk (1992): *Stochastic Equations in Infinite Dimensions*. Cambridge University Press.
- [17] De Coster, C. and P. Habets (2006): *Two-Point Boundary Value Problems: Lower and Upper Solutions*. Mathematics in Science and Engineering, Vol. 205 (First Ed.). Amsterdam: Elsevier.
- [18] Degeorge, F., J. Patel and R. Zeckhauser (1999): “Earnings Management to Exceed Thresholds,” *The Journal of Business*, 72(1), 1-33.
- [19] Dichev, I., J. Graham, H. Campbell and S. Rajgopal (2013): “Earnings Quality: Evidence From the Field,” *Journal of Accounting and Economics*, 56, 1-56.
- [20] Di Nunno, G., B. Oksendal and F. Proske (2009): “Malliavin Calculus for Lévy Processes with Applications to Finance.” Berlin: Springer-Verlag.
- [21] Dixit, A. and R. Pindyck (1994): *Investment Under Uncertainty*. Princeton: Princeton University Press.
- [22] Faingold, E., and Y. Sannikov (2011): “Reputation in Continuous-Time Games.” *Econometrica*, 79(3), 773-876.
- [23] Freixas, X., R. Guesnerie and J. Tirole (1985): “Planning under Incomplete Information and the Ratchet Effect,” *Review of Economic Studies*, 52(2), 173-191.
- [24] Gali, J. (2008): *Monetary Policy, Inflation, and the Business Cycle*. Princeton: Princeton University Press.
- [25] Holmström, B. (1999): “Managerial Incentive Problems: A Dynamic Perspective,” *The Review of Economic Studies*, 66, 169-182.
- [26] Hörner, J. and L. Samuelson (2014): “Incentives for Experimenting Agents,” *RAND Journal of Economics*, 44, 632-663.
- [27] Karatzas, I. and S. Shreve (1991): *Brownian Motion and Stochastic Calculus*. New York: Springer-Verlag.
- [28] Kovrijnykh, A. (2007): “Career Uncertainty and Dynamic Incentives,” Working paper, University of Chicago.

- [29] Kuo, H-H. (2006): *Introduction to Stochastic Integration*, Universitext, Springer, New York.
- [30] Kydland, F. and E. Prescott (1977): “Rules Rather Than Discretion: The Inconsistency of Optimal Plans,” *Journal of Political Economy*, 85(3), 473-491.
- [31] Laffont, J.J. and J. Tirole (1988): “The Dynamics of Incentive Contracts,” *Econometrica*, 56(5), 1153-1175.
- [32] Liptser, R. and A. Shiryaev (1977): *Statistics of Random Processes I and II*. New York: Springer-Verlag.
- [33] Martinez, L. (2006): “Reputation and Career Concerns,” Mimeo, Federal Reserve Bank of Richmond.
- [34] Martinez, L. (2009): “Reputation, Career Concerns, and Job Assignments,” *The B.E. Journal of Theoretical Economics, Berkeley Electronic Press*, 9(1, Contributions), Article 15.
- [35] Meyer, M. and J. Vickers (1997): “Performance Comparisons and Dynamic Incentives.” *Journal of Political Economy*, 105(3), 547-581.
- [36] Pham, H. (2009): *Continuous-time Stochastic Control and Optimization with Financial Applications*. Berlin: Springer.
- [37] Prat, J. and B. Jovanovic (2014): “Dynamic Contracts when the Agent’s Quality is Unknown,” *Theoretical Economics*, 9(3), 865-914.
- [38] Rogers, L. C. G. and D. Williams (1987): *Diffusions, Markov Processes and Martingales*. Vol. 2. *Ito Calculus*. New York, Wiley.
- [39] Sannikov, Y. (2007): “Games with Imperfectly Observable Actions in Continuous Time,” *Econometrica*, 75(5), 1285-1329.
- [40] Sannikov, Y. (2014): “Moral Hazard and Long-Run Incentives,” Working paper, Princeton University.
- [41] Stein, J. (1989): “Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior,” *Quarterly Journal of Economics*, 104, 655-669.
- [42] Stock, J. and M. Watson (2007): “Why Has U.S. Inflation Become Harder to Forecast?,” *Journal of Money, Credit and Banking*, 39, 3-33.

[43] Weitzman, M. (1980): “The “Ratchet Principle” and Performance Incentives,” *Bell Journal of Economics*, 11(1), 302-308.

[44] Williams, N. (2011): “Persistent Private Information,” *Econometrica*, 79(4), 1233-1275.

Appendix

Throughout this Appendix:

- Instead of looking at the system (p, p^*) , I work with (p, Δ) where $\Delta := p^* - p$ evolves according to $d\Delta_t = [-(\beta + \kappa)\Delta_t + \beta(a_t - a_t^*)]dt$ (eqn. (8)). This avoids carrying the same Brownian motion twice in the off-path analysis.
- Since, $p_t^* = p_t$ along the path of play, and $dp_t = -\kappa(p_t - \eta)dt + \sigma dZ_t$, it follows that $p_t \sim \mathcal{N}(p^o, \sigma^2 t)$ when $\kappa = 0$, and $p_t \sim \mathcal{N}(e^{-\kappa t} p^o + (1 - e^{-\kappa t})\eta, \sigma^2(1 - e^{-2\kappa t})/2\kappa)$ when $\kappa > 0$, from a time-zero perspective. In either case, $\lim_{t \rightarrow \infty} \mathbb{E}[e^{-rt} p_t] = \lim_{t \rightarrow \infty} \mathbb{E}[e^{-rt} p_t^2] = 0$.

Proof of Proposition 1. Consider the strategy $a_t^\epsilon = a_t^* + \Delta_t + \lambda \epsilon_t$, $t \geq 0$, where: (i) a_t^* denotes the market’s current conjecture of equilibrium play at time $t \geq 0$; (ii) Δ_t denotes the current degree of belief asymmetry at time $t \geq 0$; and where (iii) $(\epsilon_t)_{t \geq 0}$ is $(\mathcal{F}_t^{Z, \Delta})_{t \geq 0}$ -progressively measurable and satisfies $\epsilon_t < \bar{\epsilon}$, a.s. for all $t \geq 0$, some $\bar{\epsilon} > 0$.

It is easy to see that the induced process of belief asymmetry, $(\Delta_t^\epsilon)_{t \geq 0}$, is given by

$$\Delta_t^\epsilon = \Delta_t^\epsilon(\lambda) := \lambda \beta \int_0^t e^{-\kappa(t-s)} \epsilon_s ds, \quad t \geq 0,$$

and that the latter grows at most linearly in time (for $\kappa > 0$ it is even bounded). Also, $p_t^* = p_t + \Delta_t^\epsilon(\lambda)$, where $dp_t = -\kappa(p_t - \eta)dt + \sigma dZ_t$, $t \geq 0$.

The payoff of following $(a_t^\epsilon)_{t \geq 0}$ is given by

$$V^\epsilon(\lambda) = \mathbb{E} \left[\int_0^\infty e^{-rt} [u(\chi(p_t + \Delta_t^\epsilon(\lambda), \underbrace{a^*(p_t + \Delta_t^\epsilon(\lambda))}_{a_t^*})) - g(\underbrace{a^*(p_t + \Delta_t^\epsilon(\lambda)) + \Delta_t^\epsilon(\lambda) + \lambda \epsilon_t}_{a_t^\epsilon})] dt \right].$$

Let $\ell(p) = u(\chi(p, a^*(p))) - g(a^*(p))$. Since $a^*(\cdot)$ is Lipschitz, the differentiability and growth conditions in Assumption 1 ensure that $\ell(\cdot)$ is differentiable and that $\ell(\cdot)$ and $\ell'(\cdot)$ have polynomial growth. Thus, letting $L(\omega, t, \lambda) := u(\chi(p_t + \Delta_t^\epsilon(\lambda), a^*(p_t + \Delta_t^\epsilon(\lambda)))) - g(a^*(p_t + \Delta_t^\epsilon(\lambda)) + \Delta_t^\epsilon(\lambda) + \lambda \epsilon_t)$, there are $C > 0$ and $j \in \mathbb{N}$ such that, if $\lambda \in (-\delta, \delta)$, $\delta > 0$,

$$\left| \frac{\partial L}{\partial \lambda} \right| \leq Ct(1 + |p_t + t|^j),$$

as $(\Delta_t^\epsilon)_{t \geq 0}$ grows at most linearly in time. Since $(p_t)_{t \geq 0}$ is Gaussian—with a mean process that is bounded and a variance that grows at most linearly in time—the function $e^{-rt}Ct(1 + |p_t + t|^j)$ is integrable with respect to $d\mathbb{P} \times dt$, where \mathbb{P} is the measure under which $(Z_t)_{t \geq 0}$ (defined in Lemma 1) is a Brownian motion, and dt is the Lebesgue measure; a similar argument shows that $(\omega, t) \mapsto e^{-rt}L(\omega, t, \lambda)$ is also integrable for all $\lambda \in (-\delta, \delta)$. Since $e^{-rt}Ct(1 + |p_t + t|^j)$ does not depend on λ , it follows that $V^\epsilon(\lambda)$ is differentiable over $(-\delta, \delta)$.

Letting $f(p) = u(\chi(p, a^*(p)))$, it is easy to see that

$$\begin{aligned} & \left. \frac{dV^\epsilon}{d\lambda} \right|_{\lambda=0} \\ &= \mathbb{E} \left[\int_0^\infty e^{-rt} \left\{ \left(f'(p_t) - g'(a^*(p_t)) \left[\frac{da^*}{dp^*}(p_t) + 1 \right] \right) \left(\beta \int_0^t e^{-\kappa(t-s)} \epsilon_s ds \right) - g'(a^*(p_t)) \epsilon_t \right\} dt \right] \\ &= \mathbb{E} \left[\int_0^\infty e^{-rt} \epsilon_t \left\{ \beta \int_t^\infty e^{-(r+\kappa)(s-t)} \left(f'(p_s) - g'(a^*(p_s)) \left[\frac{da^*}{dp^*}(p_s) + 1 \right] \right) ds - g'(a^*(p_t)) \right\} dt \right], \end{aligned}$$

where the last equality follows from integration by parts. Using the law of iterated expectations and the Markov property, the object of interest is

$$q(p) := \mathbb{E} \left[\int_t^\infty e^{-(r+\kappa)(s-t)} \left(f'(p_s) - g'(a^*(p_s)) \left[\frac{da^*}{dp^*}(p_s) + 1 \right] \right) dt \middle| p_t = p \right].$$

Also, observe that this expression is finite, consequence both of the growth conditions in Assumption 1 and of $a^*(\cdot)$ being Lipschitz. It follows that in an interior equilibrium $g'(a^*(p_t)) = \beta q(p_t)$ must hold a.s. at all times; otherwise the long-run player can choose ϵ_t such that $\epsilon_t[\beta q(p_t) - g'(a_t^*)] > 0$, $t \geq 0$, thus increasing his payoff. \square

Proof of Proposition 2: Let $f(p) = u(\chi(p, a^*(p)))$ and recall that $dp_t = -\kappa(p_t - \eta)dt + \sigma dZ_t$, $t \geq 0$. From the proof of the Proposition 1, the random variable

$$X := \int_0^\infty e^{-(r+\kappa)t} \left(f'(p_t) - g'(a^*(p_t)) \left[\frac{da^*}{dp^*}(p_t) + 1 \right] \right) dt$$

is integrable. It follows that $Y_t := \mathbb{E}[X | \mathcal{F}_t^Z]$, $t \geq 0$, is a martingale; in particular, a local martingale. By the Martingale Representation Theorem (Theorem 36.5 in Rogers and Williams, 1987), there exists a predictable process $(H_t)_{t \geq 0}$ such that $Y_t = \int_0^t H_s dZ_s$ a.s., $t \geq 0$.

On the other hand, observe that, using the Markov property,

$$Y_t = \int_0^t e^{-(r+\kappa)s} \left(f'(p_s) - g'(a^*(p_s)) \left[\frac{da^*}{dp^*}(p_s) + 1 \right] \right) ds + e^{-(r+\kappa)t} q(p_t).$$

Since $q(\cdot)$ is of class C^2 (consequence of $g'(a^*(p)) = \beta q(p)$ and of a^* and $\rho := (g')^{-1}$ being of class C^2), Ito's rule yields that the drift of the Ito process on the right-hand side of the previous expression must satisfy

$$0 = f'(p_t) - g'(a^*(p_t)) \left[\frac{da^*}{dp^*}(p_t) + 1 \right] - (r + \kappa)q(p_t) - \kappa(p_t - \eta)q'(p_t) + \frac{1}{2}\sigma^2 q''(p_t),$$

as $Y_t = \int_0^t H_s dZ_s$, a.s., $t \geq 0$. Using that $a^*(p) = \rho(\beta q(p))$ it then follows that

$$[r + \kappa + \beta + \beta^2 \rho'(\beta q(p))q'(p)] q(p) = \frac{d}{dp}[u(\chi(p, \rho(\beta q(p))))] \Big|_{p=p_t} - \kappa(p - \eta)q(p) + \frac{1}{2}\sigma^2 q''(p).$$

Regarding $U(p) = \mathbb{E} \left[\int_0^\infty e^{-rt} [u(\chi(p_t, a^*(p_t))) - g(a^*(p_t))] dt \mid p_0 = p \right]$, this function is of class C^2 by definition of a Markov equilibrium. Also, by the growth conditions in Assumption 1 and $a^*(\cdot)$ being Lipschitz, the random variable $\tilde{X} := \int_0^\infty e^{-rt} [u(\chi(p_t, a^*(p_t))) - g(a^*(p_t))] dt$ is integrable. Following the same steps taken to derive the q -ODE (namely, constructing the martingale $\tilde{Y}_t = \mathbb{E}[\tilde{X} | \mathcal{F}_t^Z]$, and then using Ito's rule) yields that

$$rU(p) = u(\chi(p, \rho(\beta q(p)))) - g(\rho(\beta q(p))) - \kappa(p - \eta)U'(p) + \frac{1}{2}\sigma^2 U''(p).$$

□

Proof of Theorem 1: Suppose the market constructs beliefs using $a^*(\cdot) := \rho(\beta q(\cdot))$, with $q(\cdot)$ as in the theorem. Off the path of play, the private and the public belief are given by

$$\begin{aligned} dp_t^* &= [-\kappa(p_t^* - \eta) + \beta(p_t - p_t^*) + \beta(a_t - a^*(p_t^*))]dt + \sigma dZ_t, \text{ where} \\ dp_t &= -\kappa(p_t - \eta)dt + \sigma dZ_t. \end{aligned}$$

Notice that $|(a^*)'(p)| = |\beta q'(p)|/|g''(a^*(p))| \leq \beta C/\psi$, where C is the Lipschitz constant of $q(\cdot)$. Thus, the previous system has a drift and volatility that are globally Lipschitz, which guarantees that it admits a strong solution for any strategy that satisfies (i) and (ii) in the admissibility concept of Section 5.2 (Theorem 1.3.15 in Pham 2009). The long-run player's optimization problem is thus well-defined over the set of admissible strategies.

Take any solution (q, U) as in the theorem. Consider the function

$$U(p + \Delta) + [q(p + \Delta) - U'(p + \Delta)]\Delta + \frac{\Gamma}{2}\Delta^2. \tag{A.1}$$

I will show that, for a suitably chosen Γ , the assumptions in the theorem ensure that this

function is an upper bound to the long-run player's payoff under any admissible strategy.

More concretely, given an admissible strategy $\hat{a} := (\hat{a}_t)_{t \geq 0}$, define the process

$$\hat{V}_t := \int_0^t e^{-rs} [h(p_s + \hat{\Delta}_t) - g(\hat{a}_s)] ds + e^{-rt} \left\{ U(p_t + \hat{\Delta}_t) + [q(p_t + \hat{\Delta}_t) - U'(p_t + \hat{\Delta}_t)] \hat{\Delta}_t + \frac{\Gamma}{2} \hat{\Delta}_t^2 \right\},$$

where $h(p) := (u \circ \chi)(p, a^*(p))$, and $\hat{\Delta}$ denotes the belief asymmetry process under the pair $(a^*(p_t^*), \hat{a})$. Applying Ito's rule to \hat{V} ,

$$\begin{aligned} \frac{d\hat{V}_t}{e^{-rt}} &= [h(\hat{p}_t^*) - g(\hat{a}_t)] dt - r \left\{ U(\hat{p}_t^*) + [q(\hat{p}_t^*) - U'(\hat{p}_t^*)] \hat{\Delta}_t + \frac{\Gamma}{2} \hat{\Delta}_t^2 \right\} dt \\ &\quad + \underbrace{\left\{ U'(\hat{p}_t^*) [-\kappa(\hat{p}_t^* - \eta) - \beta \hat{\Delta}_t + \beta(\hat{a}_t - a^*(\hat{p}_t^*))] + \frac{1}{2} \sigma^2 U''(\hat{p}_t^*) \right\}}_{(A)} dt \\ &\quad + \hat{\Delta}_t \underbrace{\left\{ q'(\hat{p}_t^*) [-\kappa(\hat{p}_t^* - \eta) - \beta \hat{\Delta}_t + \beta(\hat{a}_t - a^*(\hat{p}_t^*))] + \frac{1}{2} \sigma^2 q''(\hat{p}_t^*) \right\}}_{(B)} dt \\ &\quad - \hat{\Delta}_t \underbrace{\left\{ U''(\hat{p}_t^*) [-\kappa(\hat{p}_t^* - \eta) - \beta \hat{\Delta}_t + \beta(\hat{a}_t - a^*(\hat{p}_t^*))] + \frac{1}{2} \sigma^2 U'''(\hat{p}_t^*) \right\}}_{(C)} dt \\ &\quad + [q(\hat{p}_t^*) - U'(\hat{p}_t^*)] [-(\beta + \kappa) \hat{\Delta}_t + \beta(\hat{a}_t - a^*(\hat{p}_t^*))] dt \\ &\quad + \Gamma \hat{\Delta}_t [-(\beta + \kappa) \hat{\Delta}_t + \beta(\hat{a}_t - a^*(\hat{p}_t^*))] dt + \text{Brownian term,} \end{aligned}$$

where I have used that $\hat{p}_t^* := p_t + \hat{\Delta}_t$ evolves according to $d\hat{p}_t^* = (-\kappa(\hat{p}_t^* - \eta) + \beta(\hat{a} - a^*(\hat{p}_t^*)) - \beta \hat{\Delta}_t) dt + \sigma dZ_t$. Now, using (15) and (14) yields

$$\begin{aligned} (A) &= rU(\hat{p}_t^*) - h(\hat{p}_t^*) + g(a^*(\hat{p}_t^*)) + U'(\hat{p}_t^*) [-\beta \hat{\Delta}_t + \beta(\hat{a}_t - a^*(\hat{p}_t^*))] \\ (B) &= \left[r + \beta + \kappa + \beta \frac{da^*(\hat{p}_t^*)}{dp} \right] q(\hat{p}_t^*) - h'(\hat{p}_t^*) + q'(\hat{p}_t^*) [-\beta \hat{\Delta}_t + \beta(\hat{a}_t - a^*(\hat{p}_t^*))] \\ (C) &= (r + \kappa) U'(\hat{p}_t^*) - h'(\hat{p}_t^*) + \underbrace{g'(a^*(\hat{p}_t^*))}_{=\beta q(\hat{p}_t^*)} \frac{da^*(\hat{p}_t^*)}{dp^*} + U''(\hat{p}_t^*) [-\beta \hat{\Delta}_t + \beta(\hat{a}_t - a^*(\hat{p}_t^*))] \end{aligned}$$

with the last equality coming from the fact that U is three times differentiable. Consequently,

$$\begin{aligned} \frac{d\hat{V}_t}{e^{-rt}} &= [g(a^*(\hat{p}_t^*)) - g(\hat{a}_t) + g'(a^*(\hat{p}_t^*))(\hat{a}_t - a^*(\hat{p}_t^*))] dt \\ &\quad + \beta [\Gamma + q'(\hat{p}_t^*) - U''(\hat{p}_t^*)] \hat{\Delta}_t (\hat{a}_t - a^*(\hat{p}_t^*)) dt \\ &\quad - [\beta(q'(\hat{p}_t^*) - U''(\hat{p}_t^*)) + \Gamma \left(\frac{r}{2} + \beta + \kappa \right)] \hat{\Delta}_t^2 dt + \text{Stochastic integral.} \end{aligned}$$

Using that g is strongly convex and that $I := U' - q$, it follows that

$$\begin{aligned} \hat{V}_t - \hat{V}_0 &\leq \int_0^t e^{-rs} \left(-\frac{\psi}{2} (\hat{a}_s - a^*(\hat{p}_s^*))^2 + \beta[\Gamma - I'(\hat{p}_s^*)] \hat{\Delta}_s (\hat{a}_s - a^*(\hat{p}_s^*)) \right. \\ &\quad \left. - \left[\Gamma \left(\frac{r}{2} + \beta + \kappa \right) - \beta I'(\hat{p}_s^*) \right] \hat{\Delta}_s^2 \right) ds + \text{Stochastic integral}. \end{aligned}$$

The integrand of the Lebesgue integral is a quadratic function of $(\hat{\Delta}, \hat{a} - a^*(\hat{p}^*))$. This quadratic will be non-positive whenever Γ is such that

$$\frac{\psi}{2} \left[\Gamma \left(\frac{r}{2} + \beta + \kappa \right) - \beta I'(\hat{p}_t^*) \right] - \frac{\beta^2 [\Gamma - I'(\hat{p}_t^*)]^2}{4} \geq 0 \quad (\text{A.2})$$

over the set $\mathcal{I} := \{I'(p) \mid p \in \mathbb{R}\}$. It is clear that if \mathcal{I} is unbounded, no $\Gamma \in \mathbb{R}$ satisfies (A.2) over the whole set \mathcal{I} . Consequently, $|I'(\cdot)|$ must be bounded for the quadratic bound to hold. Let $\bar{I} := \max\{|\sup(\mathcal{I})|, |\inf(\mathcal{I})|\} < \infty$ and set $\Gamma = \bar{I}$; in this case (A.2) becomes a concave quadratic in the variable $I'(\cdot)$. It is easy to see that (A.2) holds over $[-\bar{I}, \bar{I}]$ (hence, over \mathcal{I}) if $\bar{I} \leq \psi(r + 2\kappa + 4\beta)/4\beta^2$ ((A.2) holds strictly at $I' = \bar{I}$, whereas at $I' = -\bar{I}$ if the previous inequality holds; thus, (A.2) holds in between). A sufficient condition for \hat{V} to be a supermartingale is therefore that $|I'(p)| \leq \psi(r + 2\kappa + 4\beta)/4\beta^2 := \Gamma$.

Remark 8. *The bound on the derivative of information rents, $I'(p)$, can be relaxed if one has more information about the values that $I(\cdot)$ takes. In particular, it is easy to check that*

- (I) if $\mathcal{I} = \{\bar{I}\}$, $\bar{I} > 0$, (A.2) holds when $\Gamma = \bar{I}$;
- (II) if $\mathcal{I} = \{-\bar{I}\}$, $\bar{I} > 0$, (A.2) holds when $\Gamma = \frac{\bar{I}(r+2\kappa-2\beta)}{(r+2\kappa+2\beta)}$ if $\bar{I} \leq \frac{\psi(r+2(\kappa+\beta))^2}{4\beta^2(r+2\kappa)}$;
- (III) if $\mathcal{I} \subseteq [0, \bar{I}]$, $\bar{I} > 0$, (A.2) holds when $\Gamma = \bar{I}$ if $\bar{I} \leq \psi(r + 2\kappa + 2\beta)/\beta^2$. □

With this in hand, a standard localizing argument (which uses (i) in the Theorem) eliminates the stochastic integral through taking expectations, concluding that

$$\mathbb{E} \left[e^{-rt} \left(U(\hat{p}_t^*) + [q(\hat{p}_t^*) - U'(\hat{p}_t^*)] \hat{\Delta}_t + \frac{\Gamma}{2} \hat{\Delta}_t^2 \right) \right] \leq \underbrace{U(p_0)}_{=\hat{V}_0} - \mathbb{E} \left[\int_0^t e^{-rs} [h(\hat{p}_s^*) - g(\hat{a}_s)] ds \right].$$

Using the transversality conditions, the lim sup on the left-hand side in the previous expression is larger or equal than zero. Since $\mathbb{E} \left[\int_0^t e^{-rs} |h(\hat{p}_s^*) - g(\hat{a}_s)| ds \right] < \infty$, applying the dominated convergence theorem on the right-hand side yields that $\mathbb{E} \left[\int_0^t e^{-rs} [h(\hat{p}_s^*) - g(\hat{a}_s)] ds \right]$

converges to $\mathbb{E}[\hat{V}_\infty] := \mathbb{E} \left[\int_0^\infty e^{-rs} [h(\hat{p}_s^*) - g(\hat{a}_s)] ds \right]$. Hence

$$\mathbb{E}[\hat{V}_\infty] = \mathbb{E} \left[\int_0^\infty e^{-rt} [h(\hat{p}_t^*) - g(\hat{a}_t)] ds \right] \leq U(p_0).^{47}$$

Now, take any solution $U \in C^2(\mathbb{R})$ to the ODE (15) satisfying a quadratic growth condition. Then, $|\mathbb{E}[e^{-rt}U(p_t)]| \leq e^{-rt}C(1 + \mathbb{E}[p_t^2]) \rightarrow 0$ as $t \rightarrow \infty$. The Feynman-Kac representation (Remark 3.5.6. in Pham (2009)) yields that U is of the form

$$U(p) = \mathbb{E} \left[\int_0^\infty e^{-rt} (h(p_s) - g(\rho(\beta q(p_s)))) ds \right]$$

with $dp_t = -\kappa(p_t - \eta)dt + \sigma dZ_t$, $t > 0$, $p_0 = p$. Hence, from a time-zero perspective, the payoff from following the market's conjecture is an upper bound to the long-run player's payoff under any admissible strategy. In particular, it is also an upper bound to any public strategy (Section 2): since the market's conjecture is Markov, the process $(Z_t)_{t \geq 0}$ carries all the information relevant for decision-making that is conveyed by the public signal, which makes the long-run player weakly better off over the set of strategies that condition both on $(Z_t)_{t \geq 0}$ and $(\Delta_t)_{t \geq 0}$ than over the set of strategies that condition only on $(\xi_t)_{t \geq 0}$. Finally, since $a^*(\cdot)$ attains this upper bound, and $(a^*(p_t^*[\xi]))_{t \geq 0}$ (i.e., the long-run player's actions seen as a function of the realized public history) is feasible by assumption, it follows that $a^*(\cdot)$ is a Markov equilibrium. This concludes the proof. \square

Proof of Theorem 2: It is straightforward to verify that $U(p) = U_0 + U_1p + U_2p^2$ and $q(p) = q_1 + q_2p$ solve the system of ODEs (14)-(15) if and only if they solve the system:

$$\begin{aligned} (U_0) : 0 &= rU_0 - u_0 - \eta\kappa U_1 + \frac{\beta^2}{2\psi}q_1^2 - \sigma^2U_2 \\ (U_1) : 0 &= (r + \kappa)U_1 - u_1 + \frac{\beta^2}{\psi}q_1q_2 - 2\eta\kappa U_2 \\ (U_2) : 0 &= (r + 2\kappa)U_2 + \frac{\beta^2}{2\psi}q_2^2 + u_2 \\ (q_1) : 0 &= \left(r + \kappa + \beta + \frac{\beta^2}{\psi}q_2 \right) q_1 - \eta\kappa q_2 - u_1 \\ (q_2) : 0 &= (r + \beta + 2\kappa)q_2 + \frac{\beta^2}{\psi}q_2^2 + 2u_2. \end{aligned}$$

⁴⁷Observe that the same argument applied over $[t, T]$ and letting $T \rightarrow \infty$ would lead to $\mathbb{E} \left[\int_t^\infty e^{-r(s-t)} [h(\hat{p}_s^*) - g(\hat{a}_s)] ds \right] \leq U(p_t + \Delta_t) + [q(p_t + \Delta_t) - U'(p_t + \Delta_t)]\Delta_t + \frac{\Gamma}{2}\Delta_t^2$, $t \geq 0$.

The two solutions to the quadratic (q_2) are given by

$$q_2 = \frac{\psi}{2\beta^2} \left[-(r + \beta + 2\kappa) \pm \sqrt{(r + \beta + 2\kappa)^2 - \frac{8u_2\beta^2}{\psi}} \right] < 0.$$

Clearly, both roots are negative. I verify next that the root with the smallest absolute value satisfies the conditions of the Theorem. In this case, $r + \beta + \kappa + \beta^2 q_2 / \psi > 0$, which yields

$$q_1 = \frac{u_1 + \eta\kappa q_2}{r + \beta + \kappa + \beta^2 q_2 / \psi} \quad \text{and} \quad U_2 = -\frac{u_2 + \beta^2 q_2^2 / 2\psi}{r + 2\kappa} = \frac{(r + \beta + 2\kappa)q_2}{2(r + 2\kappa)}.$$

Bound on information rents: The sufficient condition (23) can be improved for linear-quadratic games, as it was derived without imposing any structure on information rents $I = U' - q$. For this class of games, however, information rents are linear, and hence, II in Remark 8 in the Proof of Theorem 1 can be used. More precisely, since

$$I'(p) = 2U_2 - q_2 = \frac{(r + \beta + 2\kappa)q_2}{r + 2\kappa} - q_2 = \frac{\beta q_2}{r + 2\kappa} < 0,$$

we require that

$$|I'| = -\frac{\beta}{r + 2\kappa} q_2 \leq \frac{\psi(r + 2(\kappa + \beta))^2}{4\beta^2(r + 2\kappa)}.$$

But since $-q_2 < \psi(r + \beta + 2\kappa) / 2\beta^2$, the previous condition will be satisfied if $2\beta(r + \beta + 2\kappa) < (r + 2(\beta + \kappa))^2$, which is clearly true.

Transversality conditions: From the proof of Theorem 1 it suffices to show that $\limsup_{t \rightarrow \infty} \mathbb{E}[e^{-rt}[U(p_t + \hat{\Delta}) + [q(p_t + \hat{\Delta}_t) - U'(p_t + \hat{\Delta}_t)]\hat{\Delta}_t + \Gamma\hat{\Delta}_t^2/2] \geq 0$ for any admissible strategy $(\hat{a}_t)_{t \geq 0}$, where, from II in Remark 8 in the Proof of Theorem 1

$$\Gamma = \frac{|I'|(r + 2\kappa - 2\beta)}{r + 2(\beta + \kappa)} = \frac{[q_2 - 2U_2](r + 2\kappa - 2\beta)}{r + 2(\beta + \kappa)}.$$

To this end, observe that $\beta^2 q_2 / \psi + \beta + \kappa + r > 0$, $2(\beta^2 q_2 / \psi + \beta + \kappa) + r > 0$, and that, by admissibility, $C(\hat{a}) := \mathbb{E}[\int_0^\infty e^{-rs} \hat{a}_s^2 ds] < \infty$. Also,

- **Step 1:** $\lim_{t \rightarrow 0} e^{-rt} \mathbb{E}[p_t] = \lim_{t \rightarrow 0} e^{-rt} \mathbb{E}[p_t^2] = 0$ has already been established.
- **Step 2:** $\lim_{t \rightarrow 0} e^{-rt} \mathbb{E}_t[\hat{\Delta}_t] = 0$. Notice that $\hat{\Delta}_t = \int_0^t e^{-(\beta + \kappa + \frac{\beta^2 q_2}{\psi})(t-s)} [\beta \hat{a}_s - \frac{\beta^2}{\psi} (q_1 + q_2 p_s)] ds$.

Let $I_t := \int_0^t e^{-(\beta+\kappa+\frac{\beta^2 q_2}{\psi})(t-s)} \hat{a}_s ds$. By Cauchy-Schwarz's and Jensen's inequalities

$$|e^{-rt} \mathbb{E}[I_t]| \leq (e^{-rt} [1 - e^{-2(\beta+\kappa+\beta^2 q_2/\psi)t}])^{1/2} \underbrace{(e^{-rt} \int_0^t \mathbb{E}[a_s^2] ds)^{1/2}}_{\leq C(\hat{a}) < \infty} \rightarrow 0 \quad (\text{A.3})$$

as $r + 2(\beta + \kappa + \beta^2 q_2/\psi) > 0$. The rest of the terms are straightforward.

- **Step 3:** $\lim_{t \rightarrow 0} e^{-rt} \mathbb{E}[p_t \hat{\Delta}_t] = 0$. Applying Ito's rule to $e^{(\frac{\beta^2 q_2}{\psi} + \beta + 2\kappa)t} p_t \hat{\Delta}_t$ yields

$$\begin{aligned} p_t \hat{\Delta}_t &= \underbrace{\int_0^t e^{-(\frac{\beta^2 q_2}{\psi} + \beta + 2\kappa)(t-s)} \hat{\Delta}_s [\kappa \eta ds + \sigma dZ_s]}_{I_t :=} + \beta \underbrace{\int_0^t e^{-(\frac{\beta^2 q_2}{\psi} + \beta + 2\kappa)(t-s)} p_s a_s ds}_{J_t :=} \\ &\quad - \underbrace{\frac{\beta^2}{\psi} \int_0^t e^{-(\frac{\beta^2 q_2}{\psi} + \beta + 2\kappa)(t-s)} [q_1 p_s + q_2 p_s^2] dp_s}_{K_t :=} \end{aligned}$$

The argument from step 2 shows that $\lim_{t \rightarrow 0} e^{-rt} \mathbb{E}[J_t] = 0$. Showing that $\lim_{t \rightarrow 0} e^{-rt} \mathbb{E}[K_t] = 0$ is straightforward. As for I_t , the stochastic integral has zero mean, so only $L_t := \int_0^t e^{(\frac{\beta^2 q_2}{\psi} + \beta + 2\kappa)(s-t)} \hat{\Delta}_s ds$ is left. The inequality in display (A.3) can be used to show that the integral that depends on $(\hat{a}_t)_{t \geq 0}$ vanishes once discounted by e^{-rt} . It is trivial to show that the remaining terms go to zero as well once discounted.

- **Step 4:** $\lim_{t \rightarrow \infty} e^{-rt} \mathbb{E}[(\hat{\Delta}_t)^2] = 0$. From the previous steps, the analysis is reduced to showing that $\limsup_{t \geq 0} e^{-rt} \mathbb{E}[(q_2 - U_2 + \Gamma) \hat{\Delta}_t^2] \geq 0$. If, $q_2 - U_2 + \Gamma > 0$, this is trivially true. Suppose that this is not the case. Since (i) flow payoffs are bounded by above and (ii) \hat{a} delivers finite utility (by admissibility), it follows that $|\mathbb{E}[\int_0^\infty e^{-rt} u(\chi(p_t + \hat{\Delta}_t)) dt]| < \infty$. Hence, $\limsup_{t \rightarrow \infty} e^{-rt} \mathbb{E}[u(\chi(p_t + \hat{\Delta}_t))] \geq 0$. Using that $\lim_{t \rightarrow \infty} e^{-rt} \mathbb{E}[p_t] = \lim_{t \rightarrow \infty} e^{-rt} \mathbb{E}[\hat{\Delta}_t] = \lim_{t \rightarrow \infty} e^{-rt} \mathbb{E}[p_t^2] = \lim_{t \rightarrow \infty} e^{-rt} \mathbb{E}[p_t \hat{\Delta}_t] = 0$, and that $u_2 < 0$, it can be concluded that

$$\limsup_{t \rightarrow \infty} e^{-rt} \mathbb{E}[u(\chi(p_t + \hat{\Delta}_t))] \geq 0 \Rightarrow 0 \geq \liminf_{t \rightarrow \infty} e^{-rt} \mathbb{E}[(p_t + \hat{\Delta}_t)^2] = \liminf_{t \rightarrow \infty} e^{-rt} \mathbb{E}[(\hat{\Delta}_t)^2].$$

Feasibility: Recall that $a^*(p) = \beta[q_1 + q_2 p]/\psi$, and consider the integral equation in the unknown $P \in C([0, +\infty))$

$$P_t = P_0 + \int_0^t [-\kappa(P_s - \eta) - \beta(P_s + \frac{\beta}{\psi}[q_1 + q_2 P_s])] ds + \beta f_t,$$

where $f \in C([0, +\infty))$ with $f_0 = 0$. Let $\delta = \kappa + \beta + \beta^2 q_2 / \psi$ and $\nu := \kappa \eta - \beta^2 q_1 / \psi$. It is easy to see that the solution to this equation is given by

$$P_t^f = e^{-\delta t} P_0 + \beta f_t + \nu t - \delta \int_0^t e^{-\delta(t-s)} (\nu s + \beta f_s) ds.$$

Importantly, P_t^f determines how the time- t public belief is computed given a realization of the public signal equal to f (as the volatility of the public belief is constant and $\xi_0 = 0$). For a given f , $t \mapsto P_t^f$ is continuous. Moreover, P_t^f depends only on $\{f_s : 0 \leq s \leq t\}$; i.e., it is adapted.

Consider now the canonical space $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P}^0)$ where $\Omega = C(\mathbb{R}_+; \mathbb{R}^2)$, \mathcal{F}_t is the canonical σ -algebra in $C([0, t]; \mathbb{R}^2)$, $t \geq 0$, and \mathbb{P}^0 is the Wiener measure on Ω . Let $\mathbb{E}^0[\cdot]$ denote the corresponding expectation operator, and (B_t^1, B_t^2) a Brownian motion in that space (the coordinate process). Let $(\theta_t)_{t \geq 0}$ satisfying $d\theta_t = -\kappa(\theta_t - \eta)dt + \sigma_\theta dB_t^2$ and $\xi_t := \sigma_\xi B_t^1$, $t \geq 0$. Notice that $p_t[\xi] := P_t^\xi$, $t \geq 0$, is progressively measurable (adapted and continuous; Proposition 1.1.13 in Karatzas and Shreve, 1991), and, moreover, it satisfies the SDE $dp_t = [-\kappa(p_t - \eta)dt - \beta(p_t + a^*(p_t))]dt + \beta \sigma_\xi dB_t^1$, which is linear with constant coefficients. As a result, the pair $(p_t, \theta_t)_{t \geq 0}$ is Gaussian under \mathbb{P}^0 , and can be written as

$$p_t = m_p(t) + \underbrace{\int_0^t e^{-(\kappa + \beta + \beta^2 q_1 / \psi)(t-s)} d\xi_t}_{J_t^1} \quad \text{and} \quad \theta_t = m_\theta(t) + \underbrace{\int_0^t e^{-\kappa(t-s)} \sigma_\theta dB_t^2}_{J_t^2}$$

with $m_p(\cdot)$ and $m_\theta(\cdot)$ deterministic and continuous, and J_t^1 and J_t^2 of the Wiener type.⁴⁸ In particular, the latter are progressively measurable and centered (i.e., zero-mean) Gaussian.

Let $X_t := [a^*(p_t[\xi]) + \theta_t] / \sigma_\xi$, $t \geq 0$, and notice that this process is also progressively measurable (adapted and continuous). By Corollary 3.5.2 in Karatzas and Shreve (2001), if

$$\mathcal{E}_t(X) := \exp \left(\int_0^t X_s dB_t^1 - \frac{1}{2} \int_0^t X_s^2 ds \right), \quad t \geq 0,$$

is a martingale, there exists a unique probability measure \mathbb{P} on $\Omega = C(\mathbb{R}_+; \mathbb{R}^2)$ that is equivalent to \mathbb{P}^0 when restricted to \mathcal{F}_t , $t \geq 0$. Moreover,

$$\begin{pmatrix} Z_t^\xi \\ Z_t^\theta \end{pmatrix} := \begin{pmatrix} B_t^1 \\ B_t^2 \end{pmatrix} - \begin{pmatrix} \frac{1}{\sigma_\xi} \int_0^t [a^*(p_s[\xi]) + \theta_s] ds \\ 0 \end{pmatrix} = \begin{pmatrix} \frac{1}{\sigma_\xi} \xi_t \\ B_t^2 \end{pmatrix} - \begin{pmatrix} \frac{1}{\sigma_\xi} \int_0^t [a^*(p_s[\xi]) + \theta_s] ds \\ 0 \end{pmatrix}$$

⁴⁸When the integrand is continuous and of bounded variation, the Wiener integral coincides almost surely with the path-by-path Riemann-Stieltjes definition of stochastic integral (as in P_t^f). See Theorem 2.3.7 in Kuo (2006).

is a Brownian motion under \mathbb{P} ; i.e., (1)-(2) hold under \mathbb{P} . Thus, it remains to show that $(\mathcal{E}_t(X))_{t \geq 0}$ is a martingale.

To this end, example 15.5.3 in Cohen and Elliott (2015) shows that a sufficient condition for the martingale property to hold is that for any $T > 0$ there exists $\alpha > 0$ such that

$$\sup_{t \in [0, T]} \mathbb{E}^0[\exp(\alpha X_t^2)] < \infty.$$

To show this, let $Y_t := \sqrt{2}[\beta J_t^1/\psi + J_t^2]/\sigma_\xi$. Notice that for $\alpha > 0$ there is $K_{T, \alpha} > 0$ such that $\exp(\alpha X_t^2) \leq K_{T, \alpha} \exp(\alpha Y_t^2)$. Define the random variable $M_T = \sup\{|Y_s| : s \in [0, T]\}$, which is finite. Moreover, since $(Y_s)_{s \in [0, T]}$ is centered Gaussian, it defines a centered Gaussian measure on the Banach space $C[0, T]$ with norm $\sup\{|x_s| : s \in [0, T]\}$. By Fernique's Theorem (Theorem 2.6 in Da Prato and Zabczyk, 1992), there is $\alpha > 0$ such that $\mathbb{E}^0[\exp(\alpha M_T^2)] < \infty$, and the result follows from $\sup_{t \in [0, T]} \mathbb{E}^0[\exp(\alpha Y_t^2)] < \mathbb{E}^0[\exp(\alpha M_T^2)]$. \square

In order to prove the existence results in Propositions 6 and 7, I rely on the following result from De Coster and Habets (2006):

Theorem 4. (De Coster and Habets (2006), Theorem II.5.6) *Consider the second order differential equation $u'' = f(t, u, u')$ with $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ a continuous function. Let α, β of class $C^2(\mathbb{R})$ such that $\alpha \leq \beta$, and consider the set $E = \{(t, u, v) \in \mathbb{R}^3 | \alpha(t) \leq u \leq \beta(t)\}$. Assume that for all $t \in \mathbb{R}$ $\alpha'' \geq f(t, \alpha, \alpha')$ and $\beta'' \leq f(t, \beta, \beta')$. Assume also that for any bounded interval \mathcal{I} , there exists a positive continuous function $I_{\mathcal{I}} : \mathbb{R}^+ \rightarrow \mathbb{R}$ that satisfies*

$$\int_0^\infty \frac{s ds}{\varphi_{\mathcal{I}}(s)} = \infty, \tag{A.4}$$

and for all $t \in \mathcal{I}$, $(u, v) \in \mathbb{R}^2$ with $\alpha(t) \leq u \leq \beta(t)$, $|f(t, u, v)| \leq \varphi_{\mathcal{I}}(|v|)$. Then, the previous ODE has at least one solution $u \in C^2(\mathbb{R})$ such that $\alpha \leq u \leq \beta$.

Remark 9. The proof of this theorem delivers a stronger result when α and β are bounded and $\varphi_{\mathcal{I}}$ is independent of \mathcal{I} . In this case, there is $u \in C^2$ solution to $u'' = f(t, u, u')$ satisfying $\alpha \leq u \leq \beta$ and satisfying that u' is bounded. Refer to p.123 in De Coster and Habets (2006) for the proof of the Theorem and the discussion that addresses this remark.

Proof of Proposition 6: Let $h(p) := u(\chi(p))$.

1) There exists q of class C^2 solution to (14) satisfying (25). To this end, notice that the

ratcheting equation can be written as

$$q''(p) = \frac{2}{\sigma^2} \left[\underbrace{\left(r + \beta + \kappa + \beta^2 \frac{q'(p)}{g''(\rho(\beta q(p)))} \right)}_{:=f(p,q,q')} q(p) + \kappa(p - \eta)q'(p) - h'(p) \right]. \quad (\text{A.5})$$

Let $m := \inf_{p \in \mathbb{R}} h'(p)$ and $M := \sup_{p \in \mathbb{R}} h'(p)$. Take $A, B \in \mathbb{R}$ and notice that

$$\begin{aligned} f(p, A, 0) \leq 0 &\Leftrightarrow (r + \beta + \kappa)A - h'(p) \leq 0 \Leftrightarrow A \leq \frac{m}{r + \beta + \kappa} \\ f(p, B, 0) \geq 0 &\Leftrightarrow (r + \beta + \kappa)B - h'(p) \geq 0 \Leftrightarrow B \geq \frac{M}{r + \beta + \kappa}. \end{aligned} \quad (\text{A.6})$$

Hence, the goal is to find a solution in $J := [m/(r + \beta + \kappa), M/(r + \beta + \kappa)]$ as in (25).

Since g is twice continuously differentiable and strongly convex, there exists $\psi > 0$ such that $g''(\cdot) \geq \psi$. Hence, for a bounded interval $\mathcal{I} \subset \mathbb{R}$, if $p \in \mathcal{I}$ and $u \in J$, it follows that

$$\left| \beta^2 \frac{q'(p)}{g''(\rho(\beta u))} \right| \leq \frac{\beta^2}{\psi} |q'(p)|.$$

Consequently, for any bounded interval \mathcal{I} it is possible to find constants $\phi_0 > 0$ and $\phi_{1,\mathcal{I}} > 0$ s.t.

$$|f(p, u, v)| \leq \varphi_{\mathcal{I}} := \phi_0 + \phi_{1,\mathcal{I}}|v|,$$

when $p \in \mathcal{I}$ and $u \in J$. Since that the right-hand side satisfies (A.4), Theorem 4 ensures the existence of a solution with the desired property.

2) q' bounded. Consider first the $\kappa = 0$ case. Notice that in the previous argument it is possible to choose $\phi_{1,\mathcal{I}} > 0$ independent of \mathcal{I} , so the existence of a solution that in addition has a bounded derivative is also ensured.

As for $\kappa > 0$, notice that showing $\lim_{p \rightarrow \infty} q'(p) = \lim_{p \rightarrow -\infty} q'(p) = 0$ would guarantee that q' is bounded, as q' is continuous. Consider the first limit (the argument for the other limit is analogous). First, it is clear that if $\lim_{p \rightarrow \infty} q'(p)$ exists, then it must be zero; otherwise $|q(p)|$ grows without bound as $p \rightarrow \infty$.

Suppose, towards a contradiction, that $\lim_{p \rightarrow \infty} q'(p)$ does not exist. Clearly, $\lim_{p \rightarrow \infty} q'(p)$ cannot diverge, as this would imply that $|q(p)|$ grows without bound as $p \rightarrow \infty$. Thus, the remaining possibility is that $(q'(p))_{p \geq 0}$ has at least two cluster points. Let c^1 and c^2 denote two of any such points, and suppose that $c := \max\{c^1, c^2\} > 0$. In this case, there is a sequence of local maxima of $(p_n)_{n \in \mathbb{N}}$ of q' such that $q'(p_n) > c - \epsilon > 0$ for large n . Then, $q''(p_n) = 0$, so the left-hand side of (A.5) is zero, but the right-hand side diverges when $\kappa > 0$, as $p_n q'(p_n) \rightarrow \infty$.

dominates $q'(p_n)$. Hence, $q'(p)$ must converge.

If $c := \max\{c^1, c^2\} > 0$ does not hold, then $\hat{c} := \min\{c^1, c^2\} < 0$, and the analogous argument using a sequence of local minima yields that $q'(p)$ must converge. Thus, the limit exist, and hence, it must converge to zero, from where it follows that q' is bounded.

3) Asymptotic properties. Suppose that $\lim_{p \rightarrow \infty} h'(p)$ and $\lim_{p \rightarrow -\infty} h'(p)$ exist. The first result shows the existence of limits:

Lemma 2. *Suppose that both $h'_+ := \lim_{p \rightarrow \infty} h'(p)$ and $h'_- := \lim_{p \rightarrow -\infty} h'(p)$ exist. Then $q_\infty := \lim_{p \rightarrow \infty} q(p)$ and $q_{-\infty} := \lim_{p \rightarrow -\infty} q(p)$ exist.*

Proof: Suppose that $\lim_{p \rightarrow \infty} q(p)$ does not exist. Then $(q(p))_{p \geq 0}$ has at least two different cluster points c^1 and c^2 , one of them different from $\frac{h'_+}{r+\beta+\kappa}$. Without loss of generality, assume that $c := \max\{c^1, c^2\} > \frac{h'_+}{r+\beta+\kappa}$ and call the respective distance $\delta > 0$. Given $\epsilon < \delta/3$, there exists a sequence $(p_n)_{n \in \mathbb{N}}$ of local maxima of $(q(p))_{p \geq 0}$ such that $q(p_n) > c - \epsilon$ for all $n \geq \bar{N}$, some $\bar{N} \in \mathbb{N}$. But evaluating the ratcheting equation in the sequence p_n , for large n , yields

$$\underbrace{q''(p_n)}_{\leq 0} = \frac{2(r + \beta + \kappa)}{\sigma^2} \left[q(p_n) - \frac{h'(p_n)}{(r + \beta + \kappa)} \right] > \delta/3$$

where the right-most inequality comes from the fact that for large n , $|h'(p_n) - h'_+| < \epsilon(r + \beta + \kappa)$. This is a contradiction. The case in which $c := \min\{c^1, c^2\} < \frac{h'_+}{r+\beta+\kappa}$ is analogous if a sequence of local minima is constructed. Consequently, $\lim_{t \rightarrow \infty} q(p)$ exists, and since the argument for the other limit is analogous, i.e. $\lim_{t \rightarrow -\infty} q(p)$ must exist as well. \square

I show now that the limits in (26) hold:

Case $\kappa = 0$: Recall that $\beta = \beta(\kappa)$, so I write $\beta(0)$ in this case. Suppose that $q(p)$ converges to some $L \neq \frac{h'_+}{r+\beta(0)}$ as $p \rightarrow \infty$. If this convergence is monotone, then $q'(p)$ and $q''(p)$ must converge to zero. Using that $q(p)$ is bounded, it follows that

$$\frac{\sigma^2}{2} q''(p) - \beta^2 \frac{q(p)q'(p)}{g''(\rho(\beta q(p)))} \rightarrow 0.$$

But since $\lim_{p \rightarrow \infty} -h'(p) + (r + \beta(0))q(p) \neq 0$, the ratcheting equation would not hold for p large enough, a contradiction. Suppose now that $q(p)$ oscillates as it converges to L . If $L > \frac{h'_+}{r+\beta(0)}$ (which can occur only when $h'_+ < M$), there exists a sequence of local maxima $(p_n)_{n \in \mathbb{N}}$ such that $q'(p_n) = 0$, $q''(p_n) \leq 0$ and

$$q''(p_n) = \frac{2}{\sigma^2} [-h'(p_n) + (r + \beta(0))q(p_n)].$$

But since $(r + \beta(0))q(p_n)$ converges to $L(r + \beta(0)) > h'_+$, the ratcheting equation is violated for n large enough, a contradiction. Equivalently, if $L < \frac{h'_+}{r + \beta(0)}$ (which can occur only when $h'_+ > m$), there is a sequence of minima such that an analogous contradiction holds. Thus, $q(p)$ must converge to $\frac{h'_+}{r + \beta(0)}$. The case $p \rightarrow -\infty$ is identical.

Case $\kappa > 0$: I show that (26) holds in a sequence of steps.

Step 1: $\lim_{p \rightarrow \infty} q'(p) = \lim_{p \rightarrow -\infty} q'(p) = 0$. This follows from the argument used to show that q' is bounded.

Step 2: $\lim_{p \rightarrow \infty} pq'(p) = \lim_{p \rightarrow -\infty} pq'(p) = 0$. Notice that the existence of two cluster points would imply the existence of sequences of local maxima and minima that are separated by a fixed constant; using that h' , q and q' converge, and that $q''(p) = -q'(p)/p$ at a critical point of $pq'(p)$, it would not be possible for (A.5) to hold under both sequences for p large enough. Thus, $\lim_{p \rightarrow \infty} pq'(p)$ either exists or diverges. However, the latter cannot be true, as the ratcheting equation would imply that $q''(p)$, and hence, that q' , diverge as well, a contradiction. Suppose that $\lim_{p \rightarrow \infty} pq'(p) = L > 0$. Then, given $\epsilon > 0$ small and p_0 large enough, it follows that for $p > p_0$

$$q'(p) > \frac{L - \epsilon}{p} > 0 \Rightarrow q(p) > q(p_0) + (L - \epsilon) \log(p/p_0),$$

which implies that $q(p) = O(\log(p))$, a contradiction. The case $L < 0$ is analogous, from where $\lim_{p \rightarrow \infty} pq'(p) = 0$. Finally, the analysis for limit $\lim_{p \rightarrow -\infty} pq'(p) = 0$ is identical.

Step 3: $\lim_{p \rightarrow \infty} q''(p) = \lim_{p \rightarrow -\infty} q''(p) = 0$. Using Step 1 and Step 2, the ratcheting equation implies that $\lim_{p \rightarrow -\infty} q''(p)$ exists. But if this limit is different from zero, then q' diverges as $p \rightarrow \infty$, as $q'(p) = O(p)$, a contradiction. Hence, $\lim_{p \rightarrow -\infty} q''(p) = 0$. The analysis for the other limit is analogous.

Since $q'(p)$, $pq'(p)$ and $q''(p)$ all converge to zero as $p \pm \infty$, the ratcheting equation yields

$$0 = \lim_{p \rightarrow \pm\infty} q''(p) = \lim_{p \rightarrow \pm\infty} [(r + \beta + \kappa)q(p) - h'(p)],$$

concluding the proof. □

Proof of Proposition 7: I first show that, given q a bounded solution to (14), there exists a solution to (15) satisfying a quadratic growth condition; to this end, I apply Theorem 4. I then apply the Feynman-Kac probabilistic representation theorem to show that the unique solution to (15) satisfying a quadratic growth and a transversality condition is precisely the

long-run player's on-path payoff. Finally, I show via first principles that the long-run player's payoff satisfies a linear growth condition, and that it has a bounded derivative when q' is bounded. As in the previous proof, let $h(p) := u(\chi(p))$.

Let $\alpha(p) = -A - Bp^2$. It is easy to see that given any $A, B > 0$, for every bounded interval I there are constants $\phi_{0,I}, \phi_{1,I} > 0$ such that

$$\underbrace{\frac{2}{\sigma^2} | -h(p) + g(\rho(\beta q(p))) + \kappa v(p - \eta) + ru |}_{:=f(p,u,v)} \leq \phi_{0,I} + \phi_{1,I}|v| := \varphi_I(|v|)$$

where $(u, v) \in \mathbb{R}^2$ is such that $|u| \leq A + Bp^2$, and $p \in I$. Observe that the right hand side satisfies the Nagumo condition (A.4). Now, since $G := \sup_{p \in \mathbb{R}} |g(\rho(\beta q(p)))| < \infty$,

$$\underbrace{-h(p) + g(\rho(\beta q(p_t))) - \kappa \alpha'(p)(p - \eta) - r\alpha(p)}_{:=\frac{\sigma^2}{2} f(p, -\alpha(p), -\alpha'(p))} \leq C(1 + |p|) + G - 2B\kappa(p - \eta) - r(A + Bp^2)$$

where I have also used that $\|h'\|_\infty < \infty$ implies that h satisfies a linear growth condition (i.e. there exists $C > 0$ such that $|h(p)| \leq C(1 + |p|)$ for all $p \in \mathbb{R}$). Consequently,

$$\begin{aligned} C(1 + |p|) + G - 2B\kappa(p - \eta) - r(A + Bp^2) &\leq -\frac{\sigma^2}{2} \alpha''(p) = -B\sigma^2 \\ \Leftrightarrow H(p) &:= \underbrace{(C + G + B\sigma^2 + 2B\kappa\eta - rA)}_{(1)} + \underbrace{(C|p| - 2B\kappa p - rBp^2)}_{(2)} \leq 0, \forall p \in \mathbb{R}. \end{aligned}$$

If $\kappa > 0$, (2) ≤ 0 will be automatically satisfied if B satisfies that $2B\kappa > C \Leftrightarrow B > C/2\kappa$. Now, (1) ≤ 0 is guaranteed to hold when A satisfies $rA \geq C + G + 2B\sigma^2/2 + 2B\kappa\eta$. Hence, $H(\cdot)$ is non-positive if A and B satisfy the conditions just stated. If instead $\kappa = 0$, (2) will be violated for $|p|$ small, but choosing B sufficiently large, and then A satisfying the same condition but with enough slackness ensures that $H(p) \leq 0$ for all $p \in \mathbb{R}$.

For $\nu(p) = -\alpha(p)$, notice that

$$\underbrace{-h(p) + g(\rho(\beta q(p_t))) + \kappa \nu'(p)(p - \eta) + r\nu(p)}_{:=\frac{\sigma^2}{2} f(p, \nu(p), \nu'(p))} \geq -C(1 + |p|) - G + 2B\kappa(p - \eta) + r(A + Bp^2).$$

So imposing, $\frac{\sigma^2}{2} \nu''(p) = B\sigma^2 \leq -C(1 + |p|) - G + 2B\kappa(p - \eta) + r(A + Bp^2)$ yields the exact same condition found for α . Consequently, if A, B satisfy the conditions above, α and ν are lower and upper solutions, respectively. Thus, there exist a $U \in C^2(\mathbb{R})$ solution to (15) such that $|U(p)| \leq \nu(p)$, which means that U satisfies a quadratic growth condition. Finally, the fact that $\kappa \geq 0$ and that U has quadratic growth ensures that $\mathbb{E}[e^{-rt}U(p_t)] \rightarrow 0$

as $t \rightarrow 0$. Thus, the probabilistic representation follows from the Feynman-Kac formula in infinite horizon (Pham (2009) Remark 3.5.6.).

The proof concludes by showing that if q' is bounded, (i) U' is bounded and that (ii) U satisfies a linear growth condition. For $p \in \mathbb{R}$ and $h > 0$ let $p_t^h := e^{-\kappa t}(p + h) + (1 - e^{-\kappa t})\eta + \sigma \int_0^t e^{-\kappa(t-s)} dZ_s$, that is, the common belief process starting from $p_0 = p + h$, $h \geq 0$. Notice that $p_t^h - p_t^0 = e^{-\kappa t}h$ for all $t \geq 0$, so

$$\begin{aligned} |U(p + h) - U(p)| &\leq \mathbb{E} \left[\int_0^\infty e^{-rt} (|h(p_t^h) - h(p_t^0)| + |g(\rho(\beta q(p_t^h))) - g(\rho(\beta q(p_t^0)))|) dt \right] \\ &\leq \frac{(\|h'\|_\infty + R)h}{r}, \text{ for some } R > 0, \end{aligned}$$

where I have used that q' is bounded in \mathbb{R} and that $g(\rho(\cdot))$ is Lipschitz over the set $[\frac{\beta m}{r+\beta+\kappa}, \frac{\beta M}{r+\beta+\kappa}]$. Hence, U' is bounded.

Finally, it is easy to see that if $h'(p)$ is bounded, then h' satisfies a linear growth condition. Also, since $q(\cdot)$ is bounded, $G := \sup_{p \in \mathbb{R}} g(\rho(\beta q(p))) < \infty$. When $\kappa > 0$, $p_t = e^{-\kappa t}p_0 + \kappa\eta \int_0^t e^{-\kappa(t-s)} ds + \sigma \int_0^t e^{-\kappa(t-s)} dZ_s$, so

$$|U(p_0)| \leq \mathbb{E} \left[\int_0^\infty e^{-rt} C \left(1 + \kappa\eta t + |p_0| + \left| \int_0^t e^{-\kappa(t-s)} dZ_s \right| \right) + G \right] dt$$

But since $\int_0^t e^{-\kappa(t-s)} dZ_s \sim \mathcal{N}(0, \frac{1-e^{-2\kappa t}}{2\kappa})$, the random part in the right-hand side of the previous expression has finite value. When $\kappa = 0$ the same is true, as $Z_t = \sqrt{t}Z_1$ in distribution. Consequently, there exists $K > 0$ such that $|U(p_0)| \leq K(1 + |p_0|)$. \square

Proof of Theorem 3: Let $h(p) := u(\chi(p))$. Take any bounded solution q to (14).

Step 1: Conditions (i) and (ii) in Theorem 1 hold. From Proposition 7, $U(\cdot)$ has a linear (hence, quadratic) growth condition, and U' is bounded, so the linear growth condition holds too. Also, since q' bounded, q is automatically Lipschitz. Thus, (i) in Theorem 1 holds.

As for condition (ii), I first show that $\lim_{t \rightarrow \infty} e^{-rt} \mathbb{E}[\hat{\Delta}_t] = 0$. Observe that

$$\hat{\Delta}_t = e^{-(\beta+\kappa)t} \underbrace{\hat{\Delta}_0}_{=0} + \beta \int_0^t e^{-(\beta+\kappa)(t-s)} [\hat{a}_s - a^*(p_s + \hat{\Delta}_s)] ds \quad (\text{A.7})$$

and so $|\hat{\Delta}_t| \leq 2 \sup(A)\beta[1 - e^{-(\beta+\kappa)t}]/(\beta+\kappa)$. The result then follows from A being compact.

With this in hand, it is easy to show that all the limits in (ii) holds. This is because $|U(p_t + \hat{\Delta}_t)| \leq C_1(1 + |p_t| + |\hat{\Delta}_t|)$, $|q(p_t + \hat{\Delta}_t)\hat{\Delta}_t| \leq C_2|\hat{\Delta}_t|$ and $|U'((p_t + \hat{\Delta}_t)\hat{\Delta}_t)| \leq C_3|\hat{\Delta}_t|$, for some constants C_1, C_2 and C_3 all larger than zero.

Step 2: Condition (iii) in Theorem 1 holds. From Proposition 7,

$$\mathbb{E} \left[\int_0^\infty e^{-rt} [h(p_t) - g(\rho(\beta q(p_t)))] dt \middle| p_0 = p \right] =: U(p)$$

with $dp_t = -\kappa(p_t - \eta)dt + \sigma dZ_t$, $t > 0$, and $p_0 = p$, is the unique C^2 solution to the ODE (15) satisfying a quadratic growth condition. Because the right-hand side of that ODE is differentiable, U is three times differentiable. Hence, U' satisfies the following ODE:

$$U'''(p) = \frac{2}{\sigma^2} \left[-h'(p) + \beta^2 \frac{q(p)q'(p)}{g''(\rho(\beta q(p)))} + (r + \kappa)U'(p) + \kappa(p - \eta)U''(p) \right], \quad p \in \mathbb{R}. \quad (\text{A.8})$$

Moreover, from the ratcheting equation (14)

$$-h'(p) + \beta^2 \frac{q(p)q'(p)}{g''(\rho(\beta q(p)))} = -(r + \beta + \kappa)q(p) - \kappa(p - \eta)q'(p) + \frac{1}{2}\sigma^2 q''(p).$$

Replacing this into (A.8) yields that $U' - q$ satisfies the ODE

$$(U''' - q'')(p) = \frac{2}{\sigma^2} [-\beta q(p) + (r + \kappa)(U' - q)(p) + \kappa(p - \eta)(U'' - q')(p)], \quad p \in \mathbb{R}. \quad (\text{A.9})$$

But since $U' - q$ is bounded and $\lim_{t \rightarrow \infty} \mathbb{E}[e^{-(r+\kappa)t}(U'(p_t) - q(p_t))] = 0$, the Feynman-Kac formula (Pham 2009, Remark 3.5.6.) yields that the solution to the previous ODE is unique—hence, given by $(U' - q)(\cdot)$ —and that it has the probabilistic representation

$$U'(p) - q(p) = \mathbb{E} \left[\int_0^\infty e^{-(r+\kappa)t} \beta q(p_t) dt \middle| p_0 = p \right], \quad p \in \mathbb{R}. \quad (\text{A.10})$$

Notice that (A.9) implies that $U'' - q' \in C^2$. Also, using that q' is bounded and that $(p_t)_{t \geq 0}$ is mean-reverting or a martingale, it follows that the argument that shows that U' is bounded (proof of Proposition 7) yields that $U'' - q'$ is bounded as well (in particular, $\lim_{t \rightarrow \infty} \mathbb{E}[e^{-(r+2\kappa)t}(U''(p_t) - q'(p_t))] = 0$). Furthermore, differentiating (A.9),

$$(U'''' - q''')(p) = \frac{2}{\sigma^2} [-\beta q'(p) + (r + 2\kappa)(U'' - q')(p) + \kappa(p - \eta)(U''' - q'')(p)], \quad p \in \mathbb{R}. \quad (\text{A.11})$$

The Feynman-Kac formula then yields that

$$U''(p) - q'(p) = \mathbb{E} \left[\int_0^\infty e^{-(r+2\kappa)t} \beta q'(p_t) dt \middle| p_0 = p \right], \quad p \in \mathbb{R}. \quad (\text{A.12})$$

When $\kappa = 0$, the right-hand side of the previous expression—or, equivalently, the solution

to (A.11)—admits an analytic representation in terms of q . In fact, it is easy to see that

$$\frac{\beta}{\sigma^2 \sqrt{\nu}} \left[\int_{-\infty}^p e^{-\sqrt{\nu}(p-y)} q'(y) dy + \int_p^{\infty} e^{-\sqrt{\nu}(y-p)} q'(y) dy \right]. \quad (\text{A.13})$$

where $\nu := 2r/\sigma^2$, is a solution to (A.11) when $\kappa = 0$. Because it is of class C^2 and bounded (hence, it satisfies both quadratic and transversality conditions), Feynman-Kac tell us that its must coincide with $U'' - q'$. Integrating by parts yields

$$U''(p) - q'(p) = \frac{\beta}{\sigma^2} \left[- \int_{-\infty}^p e^{-\sqrt{\nu}(p-y)} q(y) dy + \int_p^{\infty} e^{-\sqrt{\nu}(y-p)} q(y) dy \right]. \quad (\text{A.14})$$

Recalling that $q(\cdot) \in \left[\frac{m}{r+\beta}, \frac{M}{r+\beta} \right]$, that $\nu = 2r/\sigma^2$ and that $\sigma = \beta\sigma_\xi$, it is easy to see that

$$|U''(p) - q'(p)| \leq \frac{M - m}{(r + \beta) \sqrt{2r\sigma_\xi^2}}, \quad p \in \mathbb{R}.$$

But since $\beta = \sigma_\theta/\sigma_\xi$ when $\kappa = 0$, it follows that (23) in Theorem 1 will hold if

$$\frac{M - m}{(r + \beta) \sqrt{2r\sigma_\xi^2}} \leq \frac{\psi(r + 4\beta + 2\kappa)}{4\beta^2} \Big|_{\kappa=0} \Leftrightarrow \frac{M - m}{\psi} \leq \frac{\sqrt{2r\sigma_\xi^2}(r\sigma_\xi + 4\sigma_\theta)(r\sigma_\xi + \sigma_\theta)}{4\sigma_\theta^2}.$$

Since condition (28) is tighter than the one just derived, condition (iii) in Theorem 1 holds.

Remark 10. When $\kappa > 0$, the ODE (A.11) also has a solution of the form $p \mapsto \int_{\mathbb{R}} g(p, y) q'(y) dy$, but the kernel $g(p, y)$ admits no closed-form solution. This is because all solutions to (A.11) are constructed using the corresponding ones for the homogenous problem (i.e. $q' \equiv 0$), which take the form of *confluent hypergeometric functions*; see Abramowitz and Stegun (1964).

Step 3: Feasibility. Let $f \in C([0, +\infty))$ and consider the deterministic integral equation in the unknown $P \in C([0, +\infty))$

$$P_t = P_0 + \int_0^t [-\kappa(P_s - \eta) - \beta(P_s + a^*(P_s))] ds + \beta f_t$$

where $a^*(\cdot) = \rho(\beta q(\cdot))$ with $q(\cdot)$ a solution to the ratcheting equation. Notice that $|(a^*)'(p)| = |\beta q'(p)|/|g''(a^*(p))| \leq \beta C|p|/\psi$ where C is the Lipschitz constant of $q(\cdot)$. Thus, $a^*(\cdot)$ is globally Lipschitz too, and it follows immediately that exists $K > 0$ such that $|a^*(p)|^2 \leq K(1 + p^2)$. From Karatzas and Shreve (1991) p.294 (following display (2.34)), this equation admits a

unique solution P_t^f such that (i) for given f , $t \mapsto P_t^f$ is continuous, and (ii) the process $(P_t^B)_{t \geq 0}$ is adapted when $B := (B_t)_{t \geq 0}$ is a Brownian motion. Trivially, $(P_t^\xi)_{t \geq 0}$ is adapted to $\xi := \sigma_\xi B$, as this amounts replacing B by ξ in the role of f . Observe that P_t^f determines how the time- t public belief is computed given a realization of the public signal equal to f .

Consider now the canonical space $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P}^0)$ where $\Omega = C(\mathbb{R}_+; \mathbb{R}^2)$, \mathcal{F}_t is the canonical σ -algebra in $C([0, t]; \mathbb{R}^2)$, $t \geq 0$, and \mathbb{P}^0 is the Wiener measure on Ω . Let $\mathbb{E}^0[\cdot]$ denote the corresponding expectation operator, and (B_t^1, B_t^2) a Brownian motion in that space (the coordinate process). Let $(\theta_t)_{t \geq 0}$ satisfying $d\theta_t = -\kappa(\theta_t - \eta)dt + \sigma_\theta dB_t^2$ and $\xi_t := \sigma_\xi B_t^1$, $t \geq 0$. Notice that $p_t[\xi] := P_t^\xi$, $t \geq 0$, is progressively measurable (adapted and continuous; Proposition 1.1.13 in Karatzas and Shreve, 1991), and so is $X_t = [a^*(p_t[\xi]) + \theta_t]/\sigma_\xi$, $t \geq 0$. Moreover, since $a^*(\cdot)$ is bounded, given any $T, \alpha > 0$ there exists $K_{T, \alpha} > 0$ such that $\mathbb{E}^0[\exp(X_t^2)] \leq K_{T, \alpha} \mathbb{E}^0[\exp(\alpha Y_t^2)]$, where $Y_t := \frac{\sqrt{2}}{\sigma_\xi} \int_0^t e^{-\kappa(t-s)} dB_s^2$, $t \geq 0$. Since $(Y_t)_{t \geq 0}$ is centered Gaussian, the argument used in the last part of the feasibility step in the proof of Theorem 2 yields that there exists $\alpha > 0$ such that $\sup_{t \in [0, T]} \mathbb{E}^0[\alpha \exp(X_t^2)] < \infty$, and hence, $\mathcal{E}_t(X) := \exp\left(\int_0^t X_s dB_s^1 - \frac{1}{2} \int_0^t X_s^2 ds\right)$, $t \geq 0$, is a martingale. By Corollary 3.5.2 in Karatzas and Shreve (1991), there exists a unique probability measure \mathbb{P} on $\Omega = C(\mathbb{R}_+; \mathbb{R}^2)$ that is equivalent to \mathbb{P}^0 when restricted to \mathcal{F}_t , $t > 0$. Moreover,

$$\begin{pmatrix} Z_t^\xi \\ Z_t^\theta \end{pmatrix} := \begin{pmatrix} B_t^1 \\ B_t^2 \end{pmatrix} - \begin{pmatrix} \frac{1}{\sigma_\xi} \int_0^t [a^*(p_s[\xi]) + \theta_s] ds \\ 0 \end{pmatrix} = \begin{pmatrix} \frac{1}{\sigma_\xi} \xi_t \\ B_t^2 \end{pmatrix} - \begin{pmatrix} \frac{1}{\sigma_\xi} \int_0^t [a^*(p_s[\xi]) + \theta_s] ds \\ 0 \end{pmatrix}$$

is a Brownian motion under \mathbb{P} ; i.e., (1)-(2) hold under \mathbb{P} . This concludes the proof. \square

Proofs of Propositions 3, 4 and 5:

Proof of Proposition 3: Suppose that $r + 2\kappa > 2\beta/\sqrt{\psi}$. I verify that $V(n) = c + \alpha^\circ n^2/2$, $c, \alpha^\circ < 0$ to be determined, satisfies

$$\begin{aligned} rV(n) &= \sup_{a \in \mathbb{R}} \left\{ -\frac{n^2}{2} - \psi \frac{a^2}{2} + V_n(n)[- \kappa n + \beta(a - a^{*,\circ}(n))] + \frac{\sigma^2}{2} V_{nn}(n) \right\} \\ \text{s.t. } \arg \max_{a \in \mathbb{R}} \left\{ a\beta V_n(n) - \psi \frac{a^2}{2} \right\} &= \frac{\beta}{\psi} V_n(n) = a^{*,\circ}(n) \end{aligned}$$

with $\lim_{t \rightarrow \infty} e^{-rt} \mathbb{E}[V(n_t)] = 0$ on the equilibrium path, and $\limsup_{t \rightarrow \infty} e^{-rt} \mathbb{E}[V(n_t)] \geq 0$ under any admissible strategy (cf., Theorem 3.5.3 in Pham, 2009, for these transversality conditions).

To this end, notice that the envelope theorem yields $rV_n = -n - V_n \beta^2 V_{nn} / \psi - \kappa V_{nn} n - V_n \kappa + \sigma^2 V_{nnn} / 2$. Thus, α° satisfies $\beta^2 (\alpha^\circ)^2 / \psi + (r + 2\kappa) \alpha^\circ + 1 = 0$. The condition $r + 2\kappa > 2\beta/\sqrt{\psi}$ then ensures that the previous quadratic has two real solutions given by

$\alpha_{\pm}^o = \psi[-(r + 2\kappa) \pm \sqrt{(r + 2\kappa)^2 - 4\beta^2/\psi}]/2\beta^2 < 0$. Finally, plugging $V(\cdot)$ in the HJB equation and matching the value of the constants at each side yields $c = \alpha^o \sigma^2/2r < 0$.

Under both α_+^o and α_-^o , $(n_t)_{t \geq 0}$ is either mean reverting or a martingale on the path of play, so $\lim_{t \rightarrow \infty} e^{-rt} \mathbb{E}[V(n_t)] = 0$ holds. Fix now $\alpha^o \in \{\alpha_+^o, \alpha_-^o\}$ and the conjecture $a^{*,o}(n) = \beta \alpha^o n / \psi$. Observe that $\limsup_{t \rightarrow \infty} e^{-rt} \mathbb{E}[-n_t^2/2] = 0$ must hold for any admissible strategy; otherwise, the long-run player's discounted flow payoff is bounded away from zero uniformly in $[\bar{t}, \infty)$ some $\bar{t} > 0$, resulting in a total payoff of $-\infty$. Thus, $\limsup_{t \rightarrow \infty} e^{-rt} \mathbb{E}[V(n_t)] = 0$ in both cases. Finally (as an observation), under α_-^o , $-r - 2\kappa - 2\beta^2 \alpha_-^o / \psi = \sqrt{(r + 2\kappa)^2 - 4\beta^2/\psi} > 0$, and so the commitment rule $a \equiv 0$ is not admissible (as its payoff is $-\infty$ in this case). \square

Proof of Proposition 4: Since $r + 2\kappa > 2\beta/\sqrt{\psi}$ implies $r + \beta + 2\kappa > 2\beta/\sqrt{\psi}$, it follows that the curvature condition (24) holds when $u_2 = 1/2$. Consequently, Theorem 2 applies, and hence, a linear equilibrium exists. Setting $u_0 = u_1 = 0$, and $u_2 = 1/2$, the linear equilibrium delivered by the theorem takes the form $a^{*,h}(n) = \beta \alpha^h n / \psi$ where

$$\begin{aligned} 0 > \alpha^h &= \frac{\psi}{2\beta^2} [-(r + 2\kappa + \beta) + \sqrt{(r + 2\kappa + \beta)^2 - 4\beta^2/\psi}] \\ &= \frac{\psi}{2\beta^2} \frac{-4\beta^2/\psi}{(r + 2\kappa + \beta) + \sqrt{(r + 2\kappa + \beta)^2 - 4\beta^2/\psi}} \\ &> \frac{\psi}{2\beta^2} \frac{-4\beta^2/\psi}{r + 2\kappa + \sqrt{(r + 2\kappa)^2 - 4\beta^2/\psi}} = \alpha_+^o > \alpha_-^o. \end{aligned}$$

This concludes the proof. \square

Proof of Proposition 5: Since χ' is bounded, I consider q a bounded solution to the ratcheting equation as in Proposition 6; in particular, q and q' are bounded, $q \in [0, \chi'(0)/(r + \beta)]$ and $\lim_{p \rightarrow \pm\infty} q(p) = 0$.

The proof is divided into 3 steps: steps 1 and 2 prove (ii) in the proposition, whereas step 3 is devoted to (iii), (iv) and to show that $q \in (0, \chi(0)/(r + \beta))$ (open interval). Throughout the proof, when looking at $q(\cdot)$ over \mathbb{R}_- I instead consider the ODE

$$\left[r + \beta - \frac{\beta^2}{\psi} \tilde{q}'(p) \right] \tilde{q}(p) = \chi'(p) + \frac{1}{2} \sigma^2 \tilde{q}''(p), \quad p > 0.$$

which is the ODE that $\tilde{q}(p) = q(-p)$, satisfies for $p \in \mathbb{R}_+$. In particular, $\tilde{q}'(0) = -q'(0)$ and $\tilde{q}''(0) = q''(0)$. In the following, \tilde{q} denotes a solution to this ODE, whereas q denote a solution to the original ratcheting equation, both defined over \mathbb{R}_+ .

Step 1: $q'(0) > 0$ and $0 < q(0) < \chi'(0)/(r + \beta)$. Suppose that $q'(0) < 0$. Then, $\tilde{q} > q$

locally to the right of zero. Notice that $\tilde{q} - q$ is bounded and that $\tilde{q}(p) - q(p) \rightarrow 0$ as $p \rightarrow \infty$. Hence, there exists a \hat{p} at which $\tilde{q} - q$ is maximized. In particular, $\tilde{q}'(\hat{p}) = q'(\hat{p})$ and $\tilde{q}''(\hat{p}) - q''(\hat{p}) \leq 0$. The latter is equivalent to

$$\left[r + \beta - \frac{\beta^2}{\psi} \tilde{q}'(\hat{p}) \right] \tilde{q}(\hat{p}) \leq \left[r + \beta + \frac{\beta^2}{\psi} q'(\hat{p}) \right] q(\hat{p}).$$

Now, since $\tilde{q}(\hat{p}) > q(\hat{p}) \geq 0$ and $\tilde{q}'(\hat{p}) = q'(\hat{p})$ it must be the case that $\tilde{q}'(\hat{p}) = q'(\hat{p}) > 0$ for the previous inequality to hold. But since $q'(0) < 0$, q is strictly decreasing in a neighborhood of zero, so there must exist a strict minimum $\tilde{p} \in (0, \hat{p})$. Consequently,

$$0 \leq \frac{1}{2} \sigma^2 q''(\tilde{p}) = (r + \beta)q(\tilde{p}) - \chi'(\tilde{p}) + \frac{\beta^2}{\psi} q'(\tilde{p})q(\tilde{p}), \quad (\text{A.15})$$

implies that $(r + \beta)q(\tilde{p}) \geq \chi'(\tilde{p})$, as $q'(\tilde{p}) = 0$. Because $q' > 0$ locally to the right of \tilde{p} and χ' is decreasing in \mathbb{R}_+ , it follows that $(r + \beta)q(p) > \chi'(p)$ in a neighborhood to the right of \tilde{p} . But then $q''(p) > 0$ in the same region, as $q' > 0$ (see (A.15)). Thus, q' and q grow to the right, which leads to $(r + \beta)q(p) > \chi'(p)$ growing, and thus to q'' to grow again (as $q' > 0$ has grown). As a result, the existence of a local minimum leads to q'' being strictly bounded away from zero over $[\tilde{p} + \epsilon, \infty)$, some $\epsilon > 0$. Since $q \in C^2(\mathbb{R})$, q' grows indefinitely, and the same happens with q ; a contradiction. Thus, $q'(0) \geq 0$.

Obs: From the previous analysis it follows that q cannot have a strict local minimum in \mathbb{R}_+ , as this leads to q growing indefinitely over \mathbb{R}_+ . In particular, there cannot be a point $p > 0$ such that, in a neighborhood to the right of p , q is strictly decreasing and $\tilde{q} > q$ simultaneously, as this implies the existence of such a strict local minimum.

Suppose now that $q'(0) = 0$. Then, $q(0) = [\chi'(0) + \sigma^2 q''(0)/2]/[r + \beta]$. As a result, $q''(0) \leq 0$, as $q > \chi'(0)/(r + \beta)$ otherwise, which would in turn contradict (i). If $q''(0) < 0$, then $q' < 0$ close to zero, and thus q is strictly decreasing in a neighborhood of zero. Also,

$$q'''(0) = \frac{2}{\sigma^2} q(0)q''(0) < 0 \quad \text{and} \quad \tilde{q}'''(0) = -\frac{2}{\sigma^2} \tilde{q}(0)\tilde{q}''(0) > 0,$$

where I used that q is C^3 at 0, that $\chi''(0) = q'(0) = 0$, and that $q(0) > 0$ (otherwise, 0 is a minimum, a contradiction with $q'' < 0$). Thus, q is strictly decreasing and $\tilde{q} > q$ in a neighborhood of zero; a contradiction with the previous observation. It follows that $q''(0) = 0$ if $q'(0) = 0$. In particular, from the previous display, $q'''(0) = 0$ if $q'(0) = 0$.

Notice that since $q'(0) = q''(0) = 0$, it must be the case that $q(0) = \chi'(0)/(r + \beta)$, i.e., q achieves its maximum value. Because χ' is twice continuously differentiable at zero, and

$\chi'''(0) < 0$, q must be of class C^4 at zero, and hence

$$\frac{\sigma^2}{2} q''''(0) = \underbrace{(r + \beta)q''(0)}_{=0} + \frac{\beta^2}{\psi} \underbrace{[2q''(0)q'(0) + q'(0)q''(0) + q(0)q'''(0)]}_{=0} - \chi'''(0) = -\chi'''(0) > 0.$$

But this implies that q must grow locally to the right of zero, a contradiction with the definition of local maximum. Thus, $q'(0) \neq 0$, from where $q'(0) > 0$. In particular, $0 < q(0) < \chi'(0)/(r + \beta)$; otherwise $q \in [0, \chi'(0)/(r + \beta)]$ is violated in a neighborhood of zero.

Step 2. $q''(0) < 0$. It is clear that $q''(0) \leq 0$. Otherwise, q' is strictly increasing at zero and, since χ' decays in \mathbb{R}_+ , $q'' > 0$ everywhere, which means that (applying the same logic used in step 1) q grows without bound (see (A.15)). Suppose that $q''(0) = 0$. Then,

$$\frac{\sigma^2}{2} q'''(0) = \underbrace{(r + \beta)q'(0)}_{>0} + \underbrace{\frac{\beta^2}{\psi}(q'(0))^2}_{>0} + \underbrace{\frac{\beta^2}{\psi}q(0)q''(0)}_{=0} - \underbrace{\chi''(0)}_{=0} > 0$$

Then, $q'' > 0$ slightly to the right of zero, which means that q' keeps growing locally. Because χ' decreases over \mathbb{R}_+ , q'' , q' and q grow indefinitely over the same interval (same argument as in step 1), which is a contradiction. Thus, $q''(0) < 0$.

Step 3: Global maximum to the right of zero and skewness ((iii) and (iv)). The existence of a maximum over $(0, +\infty)$ is ensured by $q(p) \rightarrow 0$ as $p \rightarrow +\infty$, q being bounded, and $q(\cdot)$ growing locally to the right of zero. If there is another strict local maximum to the right of zero, there must be a strict local minimum in between; a contradiction. Thus, there is a unique strict maximum in $(0, +\infty)$, which I denote by \hat{p} . Also, observe that q cannot be flat over an interval of strict positive measure, as this violates the ratcheting equation (due to χ' being strictly decreasing); thus, q must be strictly increasing (decreasing) before (after) \hat{p} . It remains to show that there is no $p \in (-\infty, 0)$ such that $q(p) \geq q(\hat{p})$ and the skewness property. I start with the latter.

Towards a contradiction, suppose that there is $p > 0$ such that $\tilde{q}(p) = q(-p) > q(p)$. Since \tilde{q} is below q in a neighborhood to the right of zero, \tilde{q} must have crossed q somewhere in $(0, p)$. Suppose that this crossing point is in the region $(\hat{p}, +\infty)$. Thus, there must exist $\bar{p} \in (\hat{p}, +\infty)$ such that q is strictly decreasing and $\tilde{q} > q$ in a neighborhood to the right of \bar{p} , but this contradicts the observation stated in step 1. Thus, \tilde{q} cannot cross q strictly to the right of the local maximum.

Suppose now that \tilde{q} crosses q for the first time at $\bar{p} \in (0, \hat{p}]$. Since $\tilde{q} = q$ at 0 and \bar{p} , and $q > \tilde{q}$ in between, there is $p^\dagger \in (0, \bar{p})$ such that $q - \tilde{q}$ is maximized at p^\dagger over the closed interval $[0, \bar{p}]$; thus, $q'(p^\dagger) = \tilde{q}'(p^\dagger)$ and $q''(p^\dagger) \leq \tilde{q}''(p^\dagger)$. Using the ratcheting equation and

the symmetry of χ' , the last inequality leads to

$$\left[r + \beta + \frac{\beta^2}{\psi} q'(p^\dagger) \right] q(p^\dagger) \leq \left[r + \beta - \frac{\beta^2}{\psi} \tilde{q}'(p^\dagger) \right] \tilde{q}(p^\dagger).$$

But since $q(p^\dagger) > \tilde{q}(p^\dagger) \geq 0$, it must be that $q'(p^\dagger) < 0$ for the previous inequality to hold. However, this contradicts that $q' \geq 0$ over $[0, \hat{p}]$. It follows that $q \geq \tilde{q}$ over \mathbb{R}_+ ; i.e., (iv) holds. Furthermore, this last argument shows that $\tilde{q} < q$ over $[0, \hat{p}]$, \hat{p} included. Because $q(\hat{p}) > q(p) \geq \tilde{q}(p)$ for $p > \hat{p}$, it follows that \hat{p} is the global maximum over \mathbb{R} .

To finish with step 3, at the global maximum \hat{p} , $q'(\hat{p}) = 0$ and $q''(\hat{p}) \leq 0$. As a result, $(r + \beta)q(\hat{p}) = \chi'(\hat{p}) + \sigma^2 q''(\hat{p})/2 \leq \chi'(\hat{p}) < \chi'(0)$, as χ' is strictly decreasing in $(0, +\infty)$ and $\hat{p} > 0$. Also, if there is p that attains 0, it must be that $q'(p) = 0$ and $q''(p) \geq 0$, as 0 is a minimum. But this implies that $0 = (r + \beta)q(p) = \chi'(p) + \sigma^2 q''(p)/2 > 0$, a contradiction. This shows that $0 < q < \chi'(0)/(r + \beta)$.

To conclude, some further properties of $q(\cdot)$ depicted in Figure 2. First, from the first paragraph in step 3, $q(\cdot)$ is strictly decreasing to the right of the global maximum \hat{p} ; and from the last paragraph of step 3, $q(\hat{p}) \leq \chi(\hat{p})/(r + \beta)$. Second, there must exist $p \in [\hat{p}, \infty)$ such that $q''(p) > 0$. Otherwise, if $q''(\cdot) \leq 0$ in $[\hat{p}, \infty)$, the fact that $q(\cdot)$ is strictly decreasing implies that there exists $p > \hat{p}$ such that $q'(p) < 0$, and hence $q'(\cdot)$ is bounded away from zero in $[p, \infty)$. But this in turn implies that $q(p)$ will cross zero eventually, a contradiction. The ratcheting equation then yields that if $q'(p) < 0$ and $q''(p) \geq 0$ for $p > \hat{p}$, then $q(p) > \chi'(p)/(r + \beta)$. This concludes the proof. \square