Games of Private Information and Learning

by

Aaron M. Kolb

Department of Economics
Duke University

Date: ______________________
Approved:

____
Attila Ambrus, Supervisor

____
Brendan Daley

____
Curtis Taylor

____
David McAdams

____
Vincent Conitzer

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2016

Abstract

Games of Private Information and Learning

by

Aaron M. Kolb

Department of Economics
Duke University

Date: ____________________________
Approved:

____
Attila Ambrus, Supervisor

____
Brendan Daley

____
Curtis Taylor

____
David McAdams

____
Vincent Conitzer

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics in the Graduate School of Duke University
2016

Abstract

This dissertation studies private information and learning in games. Chapter 1 considers a dynamic game of strategic adoption, in which a patient buyer faces a seller of privately known quality. The buyer learns about the seller gradually through an exogenous news process, and the seller is able to exit the market or privately upgrade its quality at any time. I discover two novel kinds of reputational dynamics in equilibrium: a resetting barrier, where low types upgrade at low states inducing discrete upward jumps in reputation, and skew Brownian motion, where low types exit continuously at intermediate states, creating a permeable barrier for reputation. Contrary to the classic lemons result, the VC prefers this private information environment to one with symmetric information. The rich strategic interaction between the startup and VC implies, somewhat surprisingly, that players may benefit from increases in their own cost parameters.

Chapter 2 applies a similar information structure to an adversarial setting of optimal market entry timing. A player of privately known strength chooses when to enter a market, and an incumbent chooses whether to compete or concede. Information about the potential entrant's type is revealed publicly according to an exogenous news process and the timing of entry. I analyze stationary equilibria using the public belief as a state variable. No equilibria in pure strategies exist, and smooth-pasting conditions need not hold. Under both D1 and a novel refinement, the informed player has nondecreasing value functions and her strategy has the following structure: for

high states, both types enter with certainty; for a possibly empty interval of intermediate states, no type enters; and for low states, the high type enters while the low type mixes. I obtain closed form solutions and analyze comparative statics for such equilibria. The welfare effects of the presence of news, relative to no news, depend on the starting belief; however, for a fixed equilibrium, a marginal increase in news quality always helps the informed player regardless of her type and always hurts total welfare.

Chapter 3 explores private information in a delegation setting. We investigate competition in a delegation framework, with a coarsely informed principal. Two imperfectly informed and biased experts simultaneously propose action choices. A principal with a diffuse prior, and only able to ordinally compare the two proposals, has to choose one of them. In equilibrium, experts may exaggerate their biases, and moreover, such an equilibrium may maximize the principal's welfare. We show that having a second expert can benefit the principal, even if the two experts have the same biases or if one expert is known to be unbiased. In contrast with other models of expertise, in our setting the principal prefers experts with equal rather than opposite biases. The principal may also benefit from commitment to an "element of surprise," making an ex post suboptimal choice with positive probability. A methodological contribution of our paper is characterizing restrictions on the set of strategies which allows a formal generalization of ex ante expected payoffs to games with diffuse prior.

Contents

Abstract **iv**

List of Tables **x**

List of Figures **xi**

Acknowledgements **xiii**

1 Reputation and Strategic Adoption **1**

1.1 Introduction 1

1.2 Baseline Model 9

1.2.1 Equilibrium Definition 11

1.2.2 Equilibrium Construction 14

1.3 Endogenous Quality 21

1.3.1 Static Benchmark 21

1.3.2 Dynamic Upgrading 23

1.3.3 Resetting Equilibrium 27

1.3.4 Skew-Resetting Equilibrium 32

1.3.5 Section Summary 39

1.4 Variations on Baseline Model 40

1.4.1 Symmetric Incomplete Information 40

1.4.2 Commitment 42

1.4.3 Commitment with Symmetric Incomplete Information 43

1.5 Comparative Statics 44
1.6 Discussion 48
1.6.1 Applications 48
1.6.2 Alternative Models 49
1.6.3 Conclusion 50
1.7 Appendix A: Baseline Model 50
1.7.1 Stationary Weakly Monotone Equilibium 50
1.7.2 Uniqueness of Parameters 51
1.7.3 Proof of Theorem 1.7.1 56
1.8 Appendix B: Endogenous Quality 66
1.8.1 Belief Updating 67
1.8.2 Uniqueness of Belief Processes 68
1.8.3 Proof of Theorem 1.3.1 70
1.8.4 Skew-Resetting Equilibrium 83
1.9 Appendix C: Variations and Comparative Statics 87
1.9.1 Symmetric Incomplete Information 87
1.9.2 Commitment 93
1.9.3 Commitment with Symmetric Incomplete Information 96
1.9.4 Comparative Statics 97
2 Optimal Entry Timing **102**
2.1 Introduction 102
2.1.1 Related Literature 107
2.2 Model 109
2.3 Equilibria 112
2.3.1 Static Benchmark and Preliminary Results 115

2.3.2 Interval Equilibria 117
2.3.3 Nondecreasing Value Functions 123
2.3.4 Weakly Belief Order Preserving Equilibria 124
2.3.5 D1 Equilibria 126
2.4 Comparative Statics 128
2.5 Discussion 133
2.5.1 Conclusions 133
2.5.2 Variations of the Model 134
2.6 Appendix A: Proofs of Main Results 135
2.6.1 Proof of Proposition 2.3.2 135
2.6.2 Proof of Theorem 2.3.1 136
2.6.3 Proof of Theorem 2.3.2 138
2.6.4 Proof of Theorem 2.3.3 138
2.7 Appendix B: Supporting Results 139
2.7.1 Lemmas for All Equilibria 139
2.7.2 Interval Equilibria 143
2.7.3 NDVF Equilibria 144
2.7.4 WBOP Equilibria 146

3 A Delegation-Based Theory of Expertise 148

3.1 Introduction 148
3.2 Model 159
3.3 Equilibria in Pure Strategies 161
3.3.1 Best Response to Constant Markup Strategies 161
3.3.2 Simple delegation 163
3.3.3 Upward equilibrium 164

3.3.4 Downward equilibrium 165

3.3.5 Principal-Optimal Equilibrium 166

3.4 Extensions 168

3.4.1 Bonus Payments 168

3.4.2 Ex Ante Commitment by the Principal 172

3.4.3 Unselected Expert Indifferent over Actions 173

3.5 Comparative Statics 177

3.5.1 Equally Biased Experts 178

3.6 Conclusion 182

3.7 Appendix A: Microfoundations for Ex Ante Expected Payoffs with a Diffuse Prior 183

3.7.1 Overview 184

3.7.2 Formal Analysis of Weak Admissibility 186

3.7.3 Formal Analysis of Strong Admissibility 191

3.7.4 Proofs for Section 3.7.2 194

3.7.5 Proofs for Section 3.7.3 198

3.8 Appendix B: Proofs for Sections 3.3-3.5 201

3.8.1 Proofs for Sections 3.3 and 3.4.1 201

3.8.2 Proofs for Sections 3.4.2 and 3.4.3 211

3.8.3 Proofs for Section 3.5 218

Bibliography **220**

Biography **227**

List of Tables

2.1 Lump Sum Payoffs 110

2.2 Strategy Categories 141

List of Figures

1.1 Heuristic timeline for period t 10

1.2 Reflecting Equilibrium. 15

1.3 Gap between value functions in reflecting equilibrium. 23

1.4 Resetting Equilibrium 28

1.5 Sample path in skew-resetting equilibrium. 35

1.6 Value functions in skew-resetting equilibrium. 38

1.7 Boundaries between reflecting, resetting, and skew-resetting equilibrium regions. Reflecting occurs northeast of the upper curve. Resetting equilibrium occurs southwest of the upper curve, and skew-resetting occurs in between. 40

1.8 Resetting equilibrium: buyer value and equilibrium thresholds for fixed p and varying k. 45

1.9 Reflecting equilibrium: An example of V_H nonmonotonic in c. 46

2.1 Timeline for period t 105

2.2 Value Functions in Maximal IE 122

2.3 Value functions for equilibria with a fixed upper boundary $\alpha > z^*$ and delay regions of varying sizes, where $\alpha - \beta^* = \delta^*$. On the left, type H has increasing value in δ for $\delta < \delta^*$ and decreasing value in δ for $\delta > \delta^*$. On the right, type L always benefits as δ increases. 133

3.1 $k_1 > 0, k_2 < 0, k_1 + k_2 > 0$. The principal chooses the lower offer a_1, which lies closer to her expectation $\frac{s_1+s_2}{2}$. 162

3.2 Depending on the biases (b_1, b_2), the principal's strategy under commitment is either simple delegation to expert 2 (shaded region) or choosing the lower offer (unshaded). 168

3.3 The Expected Bias and the Principal's Payoff (exogenous bonus) for $b = 0.5, \sigma = 2$ 181

3.4 The Principal's Payoff (bonus paid by Principal) for $b = 1$, $\sigma = 1$ on the left; for $b = 2$, $\sigma = 0.5$ on the right 182

Acknowledgements

I would like to thank my main advisor, Attila Ambrus, and Brendan Daley for their endless patience and commitment to my research success. I also thank the rest of my committee – Curtis Taylor, David McAdams, and Vincent Conitzer – for their long-running support and eagerness to discuss research ideas even in the early stages. I have benefited greatly from frequent interactions with others in the theory community at Duke Economics and Fuqua, of which there are too many to name. I also thank my classmates, who have made the graduate school experience such an enjoyable one.

I am indebted beyond words to my wife, Jennifer, for undertaking this journey with me, for making innumerable sacrifices, and for constant cheerfulness and trust in me.

1

Reputation and Strategic Adoption

1.1 Introduction

There is a large existing literature on exit-like behavior in reputation games. A unifying result in this literature is that at low reputations, weak types are indifferent between mimicking strong types and revealing themselves by giving up, whether that be by selling (Daley and Green, 2012), not selling (Bar-Isaac, 2003), quitting (Gul and Pesendorfer, 2012), exiting a market (Board and Meyer-ter Vehn, 2014) or neglecting to take an aggressive action (Kolb, 2015). In most of this literature, types are fixed from the beginning – forever strong or forever weak. In reality, a weak type may have an incentive to become strong, even if doing so is costly and not directly observed. In this paper, I propose a new model of reputation with an informed player or "seller" who can exit or privately upgrade his type at any time. Although there are existing models of hidden investment in quality, such as Board and Meyer-ter Vehn (2013, 2014) and Dilmé (2014), my model produces completely new reputational dynamics. In equilibrium, a weak seller may upgrade with such intensity as to induce discrete jumps in his reputation, and he may be most likely to exit when his reputation is intermediate, not low.

I extend the existing reputation literature in a second dimension by modeling a long-lived and strategic uninformed player or "buyer." Existing models of reputation, including all of those mentioned above, have typically assumed for simplicity that buyers are short-lived or myopic and earn zero payoffs. However, these assumptions are inappropriate for many applications in which buyers can wait for the right opportunity to buy, and they preclude the analysis of buyer behavior and welfare. My innovation here is to model a buyer as a forward-looking agent who can buy at any time for a fixed price. Contrary to the classic lemons result, I find that the buyer actually prefers an environment in which the seller knows his type to one in which players begin symmetrically uninformed. Moreover, I am able to study commitment power for the buyer, and I show that these information environments – symmetric and asymmetric – have opposite implications for the role of commitment power and thus for the optimal design of institutions.

The model is general and has many applications, but for illustrative purposes, suppose the seller is a technology startup company of high or low quality, and the buyer is a venture capitalist ("VC"). The startup's goal is to obtain funding from the VC, and the VC's goal is to fund or "adopt" a high quality startup. The startup knows its own quality from the outset, as its founders understand the viability of their business plan, their user base, and their personal ambitions and limitations. The VC does not initially know the startup's quality and must become sufficiently impressed to be willing to adopt. The startup faces various ongoing costs, including operating costs such as server fees and office rental and opportunity costs from foregoing more stable income. At any time, the startup may shut down and exit the market if its prospects are dim, or it may privately improve its quality by upgrading its product or by hiring better talent.[1] The VC learns about the startup's quality

[1]That upgrading is unobserved is a natural assumption when buyers cannot monitor all of a

through journalism, informative advertising, and observing performance of trial versions; collectively, we refer to these channels as "news." At any time, as long as the startup has not exited, the VC can adopt the startup at a fixed price for a net return that is positive or negative depending on the startup's true quality. To evaluate the startup, the VC must account for three factors: news, the fact that the startup has not yet exited, and the possibility that the startup may have upgraded.

I begin the analysis with a baseline version of the model where upgrading is not possible, and I show that, subject to mild criteria, there exists a unique equilibrium. The buyer adopts above a certain threshold, and at a lower threshold, a weak seller mixes between exit and continuing. Intuitively, a weak seller anticipates bad news on average and becomes discouraged sooner than a strong seller. Continuing is thus a signal of strength, and the reputational benefit of continuing exactly offsets the cost. Equilibrium thus exhibits a reflecting barrier, which is consistent with the existing literature (Bar-Isaac, 2003; Daley and Green, 2012; Gul and Pesendorfer, 2012). This equilibrium persists when upgrading is possible but sufficiently expensive.

When upgrading is sufficiently cheap, a weak seller need not exit, and may instead randomize over upgrading. A novel equilibrium dynamic arises in which the seller's reputation faces a *resetting barrier*. Whenever his reputation reaches a certain lower threshold, a weak seller upgrades with a large probability. Upgrading is unobserved, but in equilibrium, the buyer correctly accounts for the seller's upgrading frequency. The seller's reputation instantly resets at a higher level, and from there it resumes continuous evolution. This behavior is new to the reputation literature, and it helps to explain several empirical patterns, one being the idea of business "comebacks." The case of General Motors and its 2014 ignition switch recalls serves as one example of this phenomenon. As of April 2014, GM had recalled 6 million vehicles and reported 13 deaths as a result of ignition switch failures. In response to public scrutiny, the

seller's actions, or when those actions are difficult to interpret.

company's new CEO, Mary Barra, announced on April 15, 2014 at the New York International Auto Show the creation of a new "global product integrity" division dedicated to preventing such failures in the future. By the end of that trading week, GM's share price had risen 6.4% versus 2.7% for the S&P 500 index, despite the fact that the effects of the new division had yet to be seen.[2]

A distinguishing feature of these resetting equilibria is that value functions are U-shaped in reputation, which implies that the seller sometimes benefits from a falling reputation. This feature fits the broad observation that "bad news can be good news." Berger et al. (2010) find empirically that, while bad publicity indeed hurts sales for well-known products, it helps sales for unknown products by increasing awareness. My model offers one explanation for why bad publicity does in fact increase awareness: a product with a bad reputation is likely to undergo improvement, and thus it becomes more relevant to potential buyers.[3]

Put another way, a U-shaped value function means that the seller is worst off at an intermediate reputation. As upgrade costs increase, the weak seller's U-shaped value function in the equilibrium above shifts down and eventually reaches zero. Thus for intermediate costs of upgrading, equilibrium involves resetting as described above and further novel behavior: exit at an intermediate state. This result also has empirical support. Fontana and Nesta (2009) study determinants of firm survival in a high-tech industry and find that medium-quality firms are least likely to survive. They reason that in markets with two quality tiers, low-quality firms are shielded from competition with high-quality firms, whereas medium-quality firms are not. My model offers an alternative explanation for lower survival rates of medium-quality

[2]This result is consistent with the modeling assumption of hidden investment. Since a seller in such a situation has an incentive to create the *appearance* of investment in quality regardless of whether it is actually occurring, any signal of this investment is very noisy.

[3]There are other, complementary explanations for why bad publicity helps sales; for example, buyers may have limited memory and need to be reminded of their options, or they may have heterogeneous preferences.

firms: low-quality firms are expected to make large improvements and "leapfrog" over medium-quality firms. More generally, the result that medium-quality types behave differently from high- and low-quality types also occurs in the three-type model of Feltovich et al. (2002), although there the situation is reversed in that medium types exert the highest effort.

A third contribution of this paper lies in the techniques used to characterize the two equilibrium dynamics mentioned above. The exit of low types with intermediate reputation creates a stochastic process known as *skew Brownian motion.* Skew Brownian motion was first developed rigorously by Harrison and Shepp (1981) and has since been used in diverse fields such as particle physics (Zhang, 2000), population biology (Cantrell and Cosner, 1999), fluid dynamics (Appuhamillage et al., 2010) and finance (Corns and Satchell, 2007). To my best knowledge, skew Brownian motion is new to the economics literature, and my results may further serve as microfoundations for its use in financial asset pricing models as well. In my model, the conditioning of beliefs upon not observing exit at this intermediate state exerts an upward force on the belief process, but the rate of exit is not high enough to form a reflecting barrier as in Daley and Green (2012), Gul and Pesendorfer (2012) and the baseline model of this paper. Instead, this intermediate state serves as a permeable barrier for the belief process, and thus conceptually, skew Brownian motion is a generalization of reflected Brownian motion. A further novelty is that the asymmetry of the belief process at the skew point implies that value functions of the informed player are only piecewise convex, and form a nondifferentiable, concave kink at that point. Intuitively, when the belief starts from this point it is more likely to move up than down conditional on no exit, and to offset this the gain from moving up must be smaller than the loss from moving down.

In addition to skew Brownian motion, resetting barriers are also new to economics, at least in the form of reputational dynamics. Dixit (1993) discusses resetting barriers applied exogenously to a stochastic process and gives necessary boundary conditions for value functions. Dumas (1991) shows that the optimal regulation of a stochastic process involves resetting if regulation has a fixed cost component. However, this result does not extend to my analysis since the seller can influence the belief process only through his type, not directly. Resetting barriers have since been considered as exogenous policy instruments in both biostatistics (Grigg and Farewell, 2004) and finance (Sircar and Xiong, 2007).

After analyzing equilibria with endogenous quality as above, I return to the baseline model in order to analyze the role of private information and its interaction with commitment power. I consider three variations of the baseline model. First, I consider the case of symmetric incomplete information. In this variation, both players learn together, and hence there is nothing to signal by remaining in the market. Thus when the seller's reputation gets sufficiently low, he exits with certainty and the game ends. The buyer is always worse off as a result, because low types exit too late and high types exit with positive probability.[4] This result contrasts with the classic lemons result of Akerlof (1970), where private information by the seller hurts the buyer. For the seller in my model, private information reduces errors but delays adoption by raising standards. For high starting beliefs, the standard-raising effect dominates so the seller ex ante prefers a symmetric environment.

In the second variation, I allow the buyer to commit to an adoption threshold. Under commitment, the buyer internalizes the impact of her threshold on the lower exit barrier. Since she benefits from inducing low types to exit sooner, the optimal commitment threshold, when it exists, is weakly higher than that of the baseline model. Furthermore, since the above benefit is proportional to the fraction of low types, the commitment threshold is a decreasing function of the starting belief – that

[4]Bar-Isaac (2003) also compares symmetric incomplete information to private information, focusing on asymptotic learning.

is, adoption standards are higher when more low types are present. For sufficiently high starting beliefs, the benefit of commitment vanishes. For sufficiently low starting beliefs, the buyer faces an open set problem and would like to commit to a threshold "as high as possible." As with the nonexistence result of Mirrlees (1974), the optimal payoff can be approximated arbitrarily closely.

In the third variation, I combine commitment power and symmetric information. As the buyer would like to encourage an uninformed seller to remain longer, the results are opposite those of the second variation; the optimal commitment threshold is lower than the competitive one, and is an increasing function of the starting belief. This result has a similar flavor to Georgiadis et al. (2014), where a manager with limited commitment power keeps extending a project as it nears completion, and where optimal project size decreases as managerial commitment power increases. In summary, the comparisons between the three variations of the baseline model have implications for the optimal design of institutions, and may fit anecdotal evidence of notoriously difficult group selection methods in the form of hazing rituals or boot camps in some environments and relatively soft methods in others.

A number of surprising comparative statics arise. A high quality seller sometimes benefits from an increase in operating costs since this magnifies the positive signal of remaining in the market, which shortens the time to adoption, despite the fact that the equilibrium adoption standard also increases. With endogenous quality, the buyer may benefit from an increase in the adoption cost, since this mimics commitment to a higher equilibrium adoption threshold, which leads low types to upgrade sooner. Similarly, in the baseline model with symmetric information, both players can sometimes benefit from increases in their own discount rates. In general, the rich strategic interaction and dynamic setting allow for positive indirect effects of changes in parameter values to outweigh negative direct effects.

Several recent papers analyze endogenous seller quality and buyer learning in a continuous-time setting, but with qualitatively different results. Board and Meyer-ter Vehn (2013) develop a model of hidden investment using stochastic technology shocks and extend this model to market exit in Board and Meyer-ter Vehn (2014), and Dilmé (2014) allows for instantaneous upgrades and downgrades in quality. An important difference from the current paper is that in these papers there are no reputational jumps and value functions are increasing, so there is no exit at intermediate states.[5] Furthermore, these papers do not model strategic buyer behavior, so there is no basis for comparison to the current paper in terms of buyer behavior and welfare. A technical difference is that my paper features Brownian rather than Poisson learning. Despite this, I am able to characterize equilibria in closed form, whereas Board and Meyer-ter Vehn (2014) establish equilibrium existence using a fixed point theorem.

This paper has similarities in terms of both applications and model features to a number of other papers. Bergemann and Hege (1998, 2005) consider venture capital financing with learning about project quality. Frick and Ishii (2015) study an optimal adoption problem with a population of forward-looking buyers and information externalities. Related to the seller's side, Bohren (2012) studies more generally actions with persistent effects in dynamic games. Cisternas (2012) allows for manipulation of both signals and underlying fundamentals in a game of two-sided learning. Tirole (2014) proposes a class of games in which players make hidden choices over information structures to be accessed when choosing actions in a later stage. Daley and Green (2015) look at bargaining between two long-lived players using a similar information structure as the one in this paper, but with exogenous types.

Our analysis applies to a wide variety of settings. In a political setting, a special interest group engages in costly lobbying to induce a decision making authority to adopt a certain policy; the special interest group can improve its proposal by, say,

[5]Of these papers, only Board and Meyer-ter Vehn (2014) allows for market exit, and the Poisson arrival of shocks implies that beliefs evolve continuously, which rules out jumps and the other novel results of my model. Dilmé (2014) considers only continuous dynamics by choice.

consulting with relevant experts. In a job market setting, a job applicant attends career fairs or accumulates minor credentials until being hired by an employer; or, an intern works an entry-level job with low pay in the hope of earning a full-time offer. In higher education, a Ph.D. student foregoes job opportunities in order to work towards a degree, which is granted only upon the recommendation of her advising committee. In the either setting, the employee or student can take hidden actions to improve his knowledge and skills. In general, the model especially fits situations where prices tend to be rigid and little bargaining or rejection of offers takes place, so that trade is closely approximated as adoption. The model can be readily extended to capture specific features of such applications.

The paper is organized as follows. Section 1.2 presents the baseline model, equilibrium concept, and an existence and uniqueness result. Section 1.3 introduces endogenous quality. Section 1.4 discusses three variations to the baseline model and Section 3.5 contains comparative statics. Section 1.6 concludes with a discussion of some other applications and alternative models.

1.2 Baseline Model

In this section I lay out the baseline model and provide an existence and uniqueness result. I extend the baseline model to endogenous quality in Section 1.3. The game is played between two players in continuous time over an infinite horizon. The seller has quality $\theta \in \Theta := \{H, L\}$, which is private information. The buyer has a commonly known prior $p_0 = \Pr(\theta = H)$. The game ends when either the seller exits, giving both players 0, or when the buyer adopts.[6] To continue, the seller must pay a flow cost $c > 0$. If both players act at the same time, we assume that the seller's

[6] We use "adopt" and "buy" as synonyms.

exit decision prevails.[7] If the buyer adopts, the seller receives a lump sum payoff of 1 (regardless of his type) and the buyer receives $\mathbb{1}\{\theta = H\} - k$, where $k \in (0, 1)$ represents the cost of adoption. While the game continues, information arrives according to a Brownian motion with type-dependent drift:

$$dX_t = \mu_\theta dt + \sigma dW_t, \tag{1.2.1}$$

where without loss of generality we assume $\mu_H > \mu_L$. Figure 1.1 is a heuristic extensive-form depiction of the model.

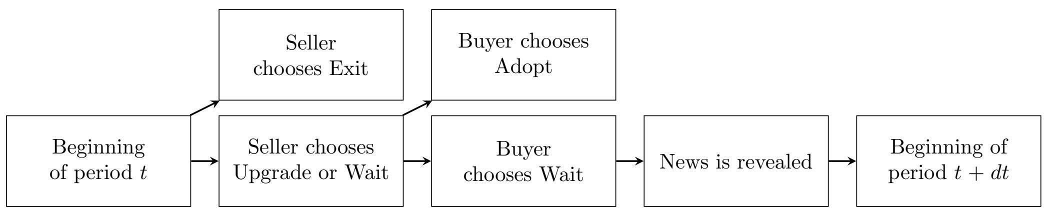

FIGURE 1.1: Heuristic timeline for period t

The seller's type is a random draw from the probability space $(\Theta, 2^\theta, \kappa)$, where κ is the probability measure on Θ that assigns probability p_0 to $\theta = H$. The standard Brownian motion W_t is defined on a canonical probability space, $(\Omega, \mathcal{F}, \mathbb{Q})$. Let $(\mathcal{F}_t)_{t\geq 0}$ be the filtration generated by the news process, where $\mathcal{F}_t = \sigma\left(\{X_s : 0 \leq s \leq t\}\right)$. Since the game ends at adoption or exit, the public history that informs either player's decision at time t is fully captured by $\mathcal{F}_t$. Define the product space $(\Omega', \mathbb{H}, \mathbb{P}) := (\Omega \times \Theta, \mathcal{F} \times 2^\Theta, \mathbb{Q} \times \kappa)$.

Formally, a pure strategy is a stopping time, and a mixed strategy is a stopping time on a state space enlarged to include a randomization device. Any such mixed strategy induces a cumulative distribution function along each path. For the analysis, it is more convenient work with these CDFs, so in some abuse of language we define

[7] This assumption captures the spirit of the extensive form interpretation of the game and simplifies payoff expressions. In equilibrium under this assumption, no such ties occur, and the equilibrium we characterize would be an equilibrium under any alternative tie-breaking rule.

stategies as CDFs, following the precedent of Daley and Green (2012) and Kolb (2015). That is, a strategy for the seller is a type-dependent family of stochastic processes, $\{(S_t^{\theta,t_0})_{t\geq t_0}\}_{t_0\in\mathbb{R}_+}$, where each $(S_t^{\theta,t_0})_{t\geq 0}$ is a stochastic process that forms a CDF over exit times along each path.

Definition 1.2.1. *An entry strategy for type θ of the seller is a family of $\mathcal{F}_t$-adapted stochastic processes $\{(S_t^{\theta,t_0})_{t\geq t_0}\}_{t_0\in\mathbb{R}_+}$ such that*

- *For all $t_0 \geq 0$, the process $(S_t^{\theta,t_0})_{t\geq t_0}$ is a right-continuous, nondecreasing, $\mathcal{F}_t$-adapted process taking values in* $[0, 1]$.
- *For all $t_1 \geq t_0 \geq 0$, let $S_{t_1-}^{\theta,t_0} := \lim_{s\uparrow t_1} S_s^{\theta,t_0}$ and specify that $S_s^{\theta,t_0} = 0$ for all $s < t_0$. If $S_{t_1-}^{\theta,t_0} < 1$, then for all $t \geq t_1$,*

$$S_t^{\theta,t_1} = \frac{S_t^{\theta,t_0} - S_{t_1-}^{\theta,t_0}}{1 - S_{t_1-}^{\theta,t_0}}.$$

Similarly, a strategy for the buyer is $(B_t^{t_0})_{t\geq 0}$, a collection of CDFs over adoption times. Pure strategies by this interpretation are strategies whose CDFs take values in $\{0, 1\}$.

We define the support of the seller's strategy $S^{\theta,s}$, denoted $\mathcal{S}^{\theta,s} := \text{supp}(S^{\theta,s})$, as the set of $\mathcal{F}_t$-adapted stopping times τ such that for all ω, $\tau(\omega) < \infty$ implies $S_{\tau(\omega)+\epsilon}^{\theta,s}(\omega) > S_{\tau(\omega)-\epsilon}^{\theta,s}(\omega)$ for all $\epsilon > 0$ and $\tau(\omega) = \infty$ implies $S_\infty^{s,\theta} > S_t^{s,\theta}$ for all $t < \infty$. We define $\mathcal{B}^s := \text{supp}(B^s)$ analogously, using ρ to denote an arbitrary stopping time for the buyer.

1.2.1 Equilibrium Definition

In this section we formally define the optimal stopping problems for the players and the equilibrium concept.

From each time s, given a strategy B^s for the buyer, the seller faces an optimal stopping problem that depends on his type, which we denote $\text{MAX}^{\theta,s}$:

$$\max_{\tau^\theta \geq s} \mathbb{E}^\theta \left[\underbrace{\int_0^{\tau^\theta} -c(1-e^{-r(\rho-s)}) + e^{-r(\rho-s)} dB^s_{t-}}_{\text{Buyer adopts first}} \underbrace{-c(1-e^{-r\tau^\theta})(1-B^s_{\tau^\theta -})}_{\theta \text{ exits first}} |\mathcal{F}_s \right]. \tag{$\text{MAX}^{\theta,s}$}$$

The expectation above is with respect to news paths, conditional on θ and the public history up to time s.

Given a strategy $S^{\theta,s}$ for the seller, the buyer's optimal stopping problem from s is

$$\max_{\rho \geq s} \mathbb{E}\left[e^{-r(\rho-s)}(\mathbb{1}\{\theta = H\} - k)(1 - S^{\theta,s}_\rho)|\mathcal{F}_s \right]. \tag{$\text{MAX}^{B,s}$}$$

For belief updating, it is most convenient to work with log-likelihood transformations. Given a belief $p_t \in [0,1]$, we write $Z_t = z(p_t) := \ln \frac{p_t}{1-p_t}$ for the transformed belief process and $p(z) := \frac{e^z}{1+e^z}$ for the inverse transformation. When Bayes' rule applies, the belief evolves according to

$$Z_{t_1} = Z_{t_0} + \phi \left[\frac{2\mu_\theta - \mu_H - \mu_L}{2\sigma}(t_1 - t_0) + W_{t_1} - W_{t_0} \right] + \ln \frac{1 - S^{H,t_0}_{t_1 -}}{1 - S^{L,t_0}_{t_1 -}}. \tag{1.2.2}$$

We focus on stationary equilibria, which allows value functions to depend only on the state variable and is standard in the dynamic games literature. To rule out bad equilibria, I further narrow the focus to stationary *monotone* equilibria, which have two additional properties: H never exits and the buyer plays a threshold strategy. To motivate the first of these, suppose the starting belief is just under z_m, and consider the following strategy profile. Both types of the seller exit immediately after any history. If at any time the seller has not yet exited, the off-path belief is set to $-\infty$.

The buyer never buys after any history. This constitutes an equilibrium,[8] but it relies on the threat of low off-path beliefs to keep both L and H from continuing. However, since H expects better news in the future than L, remaining in the game off-path should be a sign of strength.

The main result of this section, Theorem 1.2.1, is the existence and uniqueness of stationary monotone equilibrium and its full specification. In the appendix, we state and prove a stronger version of this result. There we define a larger class of *stationary weakly monotone equilibria* (SWME) which allows for all stationary buyer strategies, not just threshold strategies. The unique equilibrium of Theorem 1.2.1 is also unique within this larger class.

Given these preliminaries, we define a stationary monotone equilibrium.

Definition 1.2.2. *A stationary monotone equilibrium is an $\mathcal{F}_t$-adapted public belief process $\{Z_t\}_{t\geq 0}$, a type-dependent strategy for the seller and a threshold strategy for the buyer such that:*

1. *Seller Optimality: for all $s \geq 0$, all strategies $\tau^\theta \in \mathcal{S}^{\theta,s}$ solve* $(\mathrm{MAX}^{\theta,s})$.
2. *Buyer Optimality: for all $s \geq 0$, all strategies $\rho \in \mathcal{B}^s$ solve* $(\mathrm{MAX}^{B,s})$.
3. *Bayesian Consistency: for all real $t_1 \geq t_0 \geq 0$ and $\omega \in \Omega$ the beginning-of-period belief Z_{t_1} satisfies* (1.2.2).
4. *Stationarity: Z is a time-homogeneous, $\mathcal{F}_t$-Markov process.*
5. *Monotonicity:*
   (a) *For all real $s \geq 0$, $\mathcal{S}^{H,s} = \{\tau_H\}$, where $\tau_H \equiv +\infty$.*
   (b) *There exists $\alpha \in \mathbb{R}$ such that for all real $s \geq 0$, $\mathcal{B}^s = \{\rho^s_\alpha\}$, where $\rho^s_\alpha := \inf\{t \geq s : Z_t \geq \alpha\}$.*

[8] We have not yet formally defined equilibrium, so we are appealing to a naive notion of equilibrium here.

Condition 5 has several advantages. First, it implies that all histories in which the seller has not exited are on-path. Second, it ensures that the last term of (1.2.2) is nondecreasing, so that over time, the seller's reputation Z_t weakly increases relative to his reputation based only on news, $\widehat{Z}_t$. Third, conditions 4 and 5b together imply stationarity of the seller's problem ($\text{MAX}^{\theta,s}$); likewise, 4 and 5a imply stationarity of the buyer's problem ($\text{MAX}^{B,s}$). It follows that the current state z is a summary statistic for each player's problem, and all value functions depend only on z.

Remark 1. *Condition 5 is what allows us to call condition 4 stationarity. Without conditon 5, condition 4 is not enough to guarantee that value functions depend only on the state variable, and something more explicit like condition 1.7.1 of Definition 1.7.1 would be necessary. The buyer has no impact on the belief process, and she is patient, so her strategy could be nonstationary, affecting the seller's problem without violating condition 4. Moreover, multiple distinct seller strategies can produce the same law of motion for the belief path (for example, if both seller types mix with equal rates), and nonstationarity could arise through this multiplicity, which affects the buyer's problem. By contrast, in Daley and Green (2012), there is a continuum of identical, competitive buyers offering exactly expected values for prices. Thus condition 4 in their model implies that the price path is a time-homogeneous, Markov process. In turn, this implies that both the buyers' and the seller's problems are stationary.*

1.2.2 Equilibrium Construction

We begin with some informal arguments in order to motivate a formal equilibrium specification. Suppose that the buyer adopts immediately as soon as the state reaches some threshold α. At states below α, the seller is either exiting or waiting for the state to increase to α. We claim that starting from states below α, there must be a positive probability of exit before adoption for some seller type. If not, then beliefs below α are driven solely by news, and for sufficiently low beliefs, both types would strictly prefer to exit due to the positive flow costs. Now type H has an advantage over L: H knows that on average he earns better news and that his reputation will reach α sooner. Thus at some state β below α, L is willing to exit while H is not. It cannot be that L exits with certainty at β, since by remaining in the market an instant longer, he could mimic H and become adopted immediately. It follows that L must mix over exit at β, and in general must always weakly prefer to continue. Since the state fluctuates according to a Brownian diffusion, mixing must occur at β precisely so as to create a *reflecting barrier*, as in Daley and Green (2012), Gul and Pesendorfer (2012) and in discrete time, Bar-Isaac (2003).

We now formally define a *reflecting equilibrium*, which consists of two parameters $\beta < \alpha$ and has the following key properties:

- The buyer adopts immediately for $z \geq \alpha$.
- Seller L mixes over exit for $z \leq \beta$.
- The belief process spends no time below β.

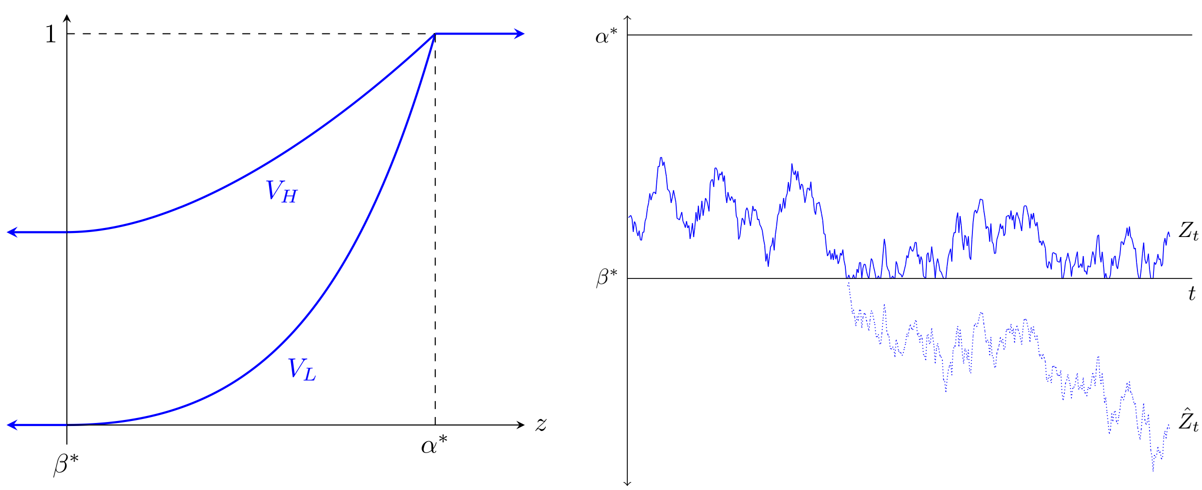

(a) Seller Value Functions (b) Sample Path with Reflecting

FIGURE 1.2: Reflecting Equilibrium.

The main result of this section, Theorem 1.2.1, is that a unique stationary monotone equilibrium exists, and it is indeed a reflecting equilibrium. Figure 1.2 shows the seller's value functions and a sample path of reputation, with and without conditioning on no exit, in a reflecting equilibrium.

Definition 1.2.3 (Reflecting Equilibrium). *For any pair $(\beta, \alpha) \in \mathbb{R}^2$, a reflecting equilibrium, denoted $\Xi(\beta, \alpha)$, is an equilibrium with strategy profile and belief process as follows:*[9]

$$\begin{aligned} Z_t &= \widehat{Z}_t + L_t \\ L_t &= \sup_{0 \le s \le t} (\beta - \widehat{Z}_s)^+ \\ S_{t_1}^{H,t_0} &= 0 \\ S_{t_1}^{L,t_0} &= 1 - e^{-(L_{t_1} - L_{t_0})} \\ B_{t_1}^{t_0} &= \mathbb{1}\{\exists s \in [t_0, t_1] : Z_s \ge \alpha\}. \end{aligned}$$

In what follows, we construct a reflecting equilibrium and informally show how its parameters are uniquely determined. Later, we argue that any equilibrium must be a reflecting equilibrium. The fully detailed arguments are contained in the appendix. For all t such that $Z_t \in (\beta, \alpha)$, both types continue, the final term of (1.2.2) vanishes, yielding

$$dZ_t^H = d\widehat{Z}_t^H := \frac{\phi^2}{2} dt + \phi dW_t \tag{1.2.3}$$

$$dZ_t^L = d\widehat{Z}_t^L := -\frac{\phi^2}{2} dt + \phi dW_t. \tag{1.2.4}$$

For $Z_t \in (\beta, \alpha)$, the value to both seller types is approximated as the expected

[9] We use $Y^+ := \max\{0, Y\}$.

discounted value at time $t + dt$ less the flow cost paid over the interval $[t, t + dt)$:

$$V_H(Z_t) \approx -crdt + (1 - rdt)\mathbb{E}^H[V_H(Z_{t+dt})|Z_t].$$

Using Ito's Lemma, we derive the Hamilton-Jacobi-Bellman (HJB) equation for type H:

$$V_H(Z_t) = -crdt + (1 - rdt)(V_H(Z_t) + \frac{\phi^2}{2}dtV_H'(Z_t) + \frac{\phi^2}{2}V_H''dt)$$

$$\implies V_H(z) = -c + \frac{\phi^2}{2r}V_H'(z) + \frac{\phi^2}{2r}V_H''(z), \tag{1.2.5}$$

and likewise for type L:

$$V_L(z) = -c - \frac{\phi^2}{2r}V_L'(z) + \frac{\phi^2}{2r}V_L''(z). \tag{1.2.6}$$

The right sides in (1.2.5) and (1.2.6) each consist of the flow cost, a first order term representing the effect of reputational drift, and a second order term representing the volatility driven by Brownian motion.

These differential equations have solutions of the form[10]

$$v_H(z) = C_1^H e^{(m-1)z} + C_2^H e^{-mz} - c \tag{1.2.7}$$

$$v_L(z) = C_1^L e^{mz} + C_2^L e^{(1-m)z} - c, \tag{1.2.8}$$

where $m = \frac{1}{2}\left(1 + \sqrt{1 + \frac{8r}{\phi^2}}\right) > 1$.

Since adoption is immediate at α, both seller types earn 1 at α. Since L mixes at β, he must be indifferent which requires his continuation value be 0. These give

[10]We use lowercase v_H and v_L to distinguish the solutions to the ODEs, and V_H and V_L to denote the true value functions, which overlap with v_H and v_L only in $[\beta, \alpha]$. Since v_H and v_L have nice properties on all of $\mathbb{R}$, the distinction is often useful.

the following three boundary conditions:

$$V_H(\alpha) = 1 \tag{1.2.9}$$

$$V_L(\alpha) = 1 \tag{1.2.10}$$

$$V_L(\beta) = 0, \tag{1.2.11}$$

By construction, the belief process must move immediately upward upon reaching the reflecting barrier β. Since this effect must be incorporated into each seller type's continuation value at β, it is necessary that both value functions be flat at that point, giving two more boundary conditions:

$$V'_H(\beta) = 0 \tag{1.2.12}$$

$$V'_L(\beta) = 0. \tag{1.2.13}$$

For a given $\alpha \in \mathbb{R}$, the above equations can be used to solve for the four constants in (1.2.7) and (1.2.8) and β all as functions of α. We use $\beta^*(\alpha)$ to denote the seller's best response function, and we show that $\beta^*(\alpha)$ is linear with slope 1.

For the buyer, it is often convenient to use untransformed beliefs $p \in [0,1]$, which evolve inside (β, α) as

$$dp_t = \phi\Phi_t(\mathbb{1}\{\theta = H\} - p_t)dt + \Phi_t dW_t, \tag{1.2.14}$$

where $\Phi_t := \phi p_t(1 - p_t)$. To a first order approximation, the HJB equation is

$$\begin{aligned} V_B(p) &\approx (1 - rdt)\mathbb{E}[V_B(p_{t+dt})|p_t = p] \\ &= (1 - rdt)[V_B(p) + \frac{\phi^2}{2}p^2(1-p)^2 V''_B(p)dt] \\ \implies rV_B(p) &= \frac{\phi^2}{2}(p(1-p))^2 V''_B(p). \end{aligned} \tag{1.2.15}$$

The stopped process $p_{t\wedge\tau}$, where $\tau := \inf\{t \geq 0 : p_t \notin (\beta, \alpha)\}$, evolves solely based on news and is thus a martingale, so there is no first order drift term in (1.2.15).

The solution is of the form

$$V_B(p) = C_1^B(1-p)\left(\frac{p}{1-p}\right)^m + C_2(1-p)\left(\frac{p}{1-p}\right)^{1-m} \tag{1.2.16}$$

$$\iff V_B(z) = \frac{C_1^B e^{mz} + C_2^B e^{(1-m)z}}{1+e^z}. \tag{1.2.17}$$

At $z = \alpha$, the buyer adopts, so we have the value matching condition

$$V_B(\alpha) = p(\alpha) - k = \frac{e^\alpha}{1+e^\alpha} - k. \tag{1.2.18}$$

At $z = \beta$, L exits stochastically, and conditional on no exit, the belief process is instantly regulated upward from β. The value lost, $V_B(\beta)$, weighted by the fraction of low types $(1 - p(\beta)) = (1 + e^\beta)^{-1}$, must exactly offset the gain from the upward regulation, $V_B'(\beta)$, which gives the *Robin* boundary condition[11]

$$V_B'(\beta) = \frac{V_B(\beta)}{1+e^\beta}. \tag{1.2.19}$$

For a third and final necessary condition, we argue that the optimality condition for the buyer is smooth-pasting at α:

$$V_B'(\alpha) = p'(\alpha) = \frac{e^\alpha}{(1+e^\alpha)^2}. \tag{1.2.20}$$

Note that $V_B'(p(\alpha))$ cannot be greater than 1 since V is bounded below by $p - k$. In addition, if $V_B'(p(\alpha)) < 1$, then her value function has a convex kink at α and she can exploit this by an arbitrarily small delay.

The buyer's payoff function is single-peaked in α so it has a unique global maximizer denoted $\alpha^*(\beta)$. In the appendix we show that $\alpha^{*\prime}(\beta) \in (0, 1)$ for all $\beta \in \mathbb{R}$.

[11] For an alternative derivation, observe that at $z < \beta$, L exits with an atom of probability $1 - e^{z-\beta}$. If there is no exit, the belief instantly jumps to β and the buyer earns $V_B(\beta)$. Hence for $z \leq \beta$, $V_B(z) = [1 - (1 - p(z))(1 - e^{z-\beta})]V_B(\beta) = \frac{e^{z-\beta}+e^z}{1+e^z} V_B(\beta)$. Imposing differentiability w.r.t. z at β yields (1.2.19).

Using this and the fact that $\beta^{*\prime}(\alpha) = 1$, the best response functions $\beta^*(\alpha)$ and $\alpha^*(\beta)$ have a unique intersection point, denoted (β^*, α^*).

In the appendix, we show that any stationary monotone equilibrium must be of the reflecting form $\Xi(\beta, \alpha)$, and as sketched above, its parameters are unique. The argument for uniqueness of form goes as follows. There must exist some state, α at which the buyer weakly prefers to adopt; otherwise, she weakly prefers to wait forever, which yields zero equilibrium payoff, and at sufficiently high states, she can adopt immediately for a positive payoff. Now the set of states at which adoption can occur is bounded below by the static threshold, z_m, so we define α_0 as the infimum of such states. Similarly, there must a state $\beta < \alpha_0$ at which L weakly prefers to exit. If not, then by Bayes' rule, the belief process only updates based on news, and therefore reaches arbitrarily low states. At sufficiently low states, L would strictly prefer to exit, rather than wait for the belief to return to α_0. We let $\beta_0 < \alpha_0$ denote the supremum of such β. We also argue that once the belief reaches β_0, it never goes below β_0. A priori, there may exist states β' *above* α_0 at which $V_L(\beta') = 0$. We argue that the set of such β' is bounded above, otherwise the buyer would adopt arbitrarily soon at sufficiently high β', and then L would strictly prefer to continue there. Hence, we let β_1 be the supremum of these β'. Lastly, we argue that for any state z, there exists $\alpha' > z$ at which the buyer *strictly* prefers to adopt, and we define α_1 as the infimum of such $\alpha' > \beta_1$. Now α_0 represents a "best case scenario" for L, and hence determines a lower bound for β_0. Likewise, β_1 represents a best case scenario for the buyer, and thus determines an upper bound for α_1. The rest of the argument establishes that these bounds collapse, that is $\alpha_0 = \alpha_1 = \alpha^*$ and $\beta_0 = \beta_1 = \beta^*$, and that equilibrium strategies in fact match $\Xi(\beta^*, \alpha^*)$.

Theorem 1.2.1. *There is a unique stationary monotone equilibrium, namely* $\Xi(\beta^*, \alpha^*)$.

1.3 Endogenous Quality

In this section I endogenize quality by allowing L to privately upgrade at a cost $K \in (0, 1)$ and instantaneously become H. If he does so, both players benefit: the news process drift becomes μ_H thereafter and the buyer earns $1 - k$ by adopting. I provide an equilibrium characterization and show that seller strategies depend qualitatively on K.

When K is high, the above equilibrium persists. Numerically, it appears that the same is true for poor news quality. Intuitively, poor news quality means that the extra good news earned by a high type takes longer to arrive, and this reduces the benefit of upgrading. On the other hand, when K is low or news quality is high, upgrading must occur. Upgrading at low reputations can even eliminate exit from equilibrium and produce a resetting (as opposed to reflecting) barrier for the firm's reputation.

For intermediate K, the resetting barrier remains, but some exit must occur at an intermediate reputation to keep the low type's value function nonnegative. This mixing produces a *skew Brownian motion*, where the high type's value function has a single concave kink.

I first consider a benchmark case in which the private upgrade option is available as a one-time static decision prior to time 0 and the start of the news process. In Section 1.3.2, extend this to allow private upgrading at any time on-path.

1.3.1 Static Benchmark

Augment the original model as follows. First, nature chooses the seller's initial type $\theta_- \in \{H, L\}$ from a common prior $p_- = \Pr(\theta_- = H)$. The seller observes θ_- and chooses a type $\theta_0 \in \{H, L\}$. If $\theta_- = L$ and $\theta_0 = H$, the seller pays a lump-sum cost K. Formally, a mixed strategy for the seller consists of an exit strategy together with a probability $q \in [0,1]$ of investment by L. The buyer then forms an ex interim belief p_0, and the game continues as in the baseline model. An equilibrium of the augmented game is thus a strategy profile that satisfies the conditions of Definition 1.2.2 along with

$$p_0 = p_- + (1 - p_-)q \tag{1.3.1}$$

$$q \in \arg\max_y y(V_H(p_0) - K) + (1 - y)V_L(p_0), \tag{1.3.2}$$

where V_H and V_L are the value functions from $\Xi(\beta^*, \alpha^*)$. For a fixed (log-likelihood) belief z, the seller's indifference condition is

$$V_H(z) - V_L(z) = K. \tag{1.3.3}$$

In the reflecting equilibrium, the difference $V_H(z) - V_L(z)$ is convex at $z = \beta^*$, is single-peaked in z and attains its maximum in (β^*, α^*), as illustrated by Figure 1.3. Intuitively, the incentive to become H is nonmonotonic because at very high states, adoption is impending regardless of the seller's type, and for very low states, adoption is heavily discounted, regardless of his type. It is for intermediate states that becoming H has the largest marginal benefit. This result is a special case of Lemma 1.3.1, which we postpone for the sake of exposition.

Define $K^* := \max_z V_H(z) - V_L(z)$ and $K_{\beta^*} := V_H(\beta^*) - V_L(\beta^*)$, and note that $0 < K^{\beta^*} < K^* < 1$.

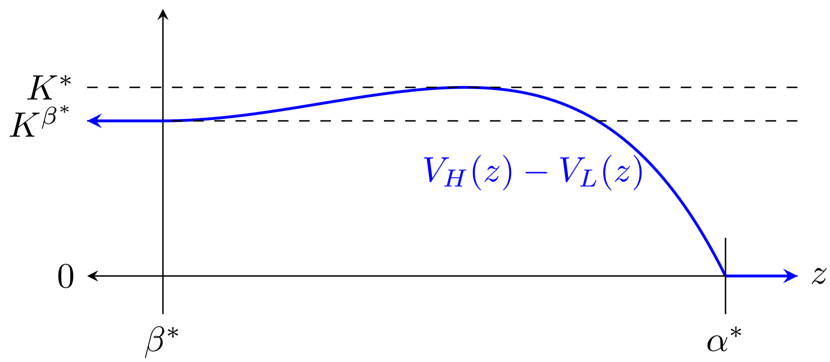

FIGURE 1.3: Gap between value functions in reflecting equilibrium.

Proposition 1.3.1. *Equilibrium of the game with static upgrading consists of the unique baseline equilibrium* $\Xi(\beta^*, \alpha^*)$ *together with the following:*

- *If* $K > K^*$*, then equilibrium is unique for all* z_-*, with* $q = 0$ *and* $z_0 = z_-$*.*
- *If* $K < K^{\beta^*}$*, then there is a unique solution* $z' \in (\beta^*, \alpha^*)$ *to* (1.3.3)*. If* $z_- \geq z'$*, equilibrium is unique with* $q = 0$ *and* $z_0 = z_-$*. If* $z_- < z'$*, equilibrium is unique with* $q = \frac{p'-p_-}{1-p_-}$ *and* $z_0 = z'$*.*
- *If* $K \in (K^{\beta^*}, K^*)$*, there are two distinct solutions* $z'', z''' \in (\beta^*, \alpha^*)$ *to* (1.3.3)*, where* $z'' < z'''$*. If* $z_- > z'''$*, then equilibrium is unique with no upgrading. If* $z_- \in (z'', z''')$*, then equilibrium is unique with* $z_0 = z'''$*. If* $z_- < z'''$*, then there are three equilibria, with* $z_0 \in \{z_-, z'', z'''\}$*.*

1.3.2 Dynamic Upgrading

We now allow the seller to choose upgrade dynamically. Heuristically, we interpret the upgrade decision as occurring at the beginning of the period, prior to the exit decision.[12]

Formally, let $\theta_t \in \{H, L\}$ denote the seller's true type at the beginning of period t. Let θ_- be the seller's exogenously assigned type at the beginning of the game.

[12] Again, no ties occur in the equilibria we construct, and these equilibria would persist under any tie-breaking rule.

A pure strategy for either type is a pair of stopping times (τ, ν), where τ is the time of exit and ν the time of upgrading, each adapted to the filtration $\mathcal{F}_t^W := \sigma\left(\{W_s : 0 \leq s \leq t\}\right)$. Every such ν induces a unique type process defined by[13]

$$(\theta_t^\nu)_{t\geq 0} = \begin{cases} H & \text{if } \theta_- = H \text{ or } \nu \leq t \\ L & \text{otherwise.} \end{cases}$$

Identifying H with 1 and L with 0, this process is binary, nondecreasing and right-continuous and $\theta_0 \geq \theta_-$.

Given any investment time ν, the news process is thus

$$dX_t = \mu(\theta_t^\nu)dt + \sigma dW_t, \tag{1.3.4}$$

where we use $\mu(\theta) := \mu_\theta$ to reduce subscripts.

We let $\mathcal{F}_t^X := \sigma\left(\{X_s : 0 \leq s \leq t\}\right)$, the filtration generated by the news process.[14]

Definition 1.3.1. *A (mixed) strategy for the seller is a family of pathwise joint CDFs,* $\left\{\left(I^{\theta,t_0}(t_1,t_2)\right)_{t_1,t_2\geq t_0}\right\}_{t_0\in\mathbb{R}_+}$ *such that*

- *For all* $t_1, t_2 \geq t_0 \geq 0$, $I^{\theta,t_0}(t_1,t_2)$ *is* $\mathcal{F}^W_{t_1\vee t_2}$*-measurable, and the mapping* $(t_1,t_2) \mapsto I^{t_0}(t_1,t_2)$ *is nondecreasing and right-continuous in both* t_1 *and* t_2*, and takes values in* $[0,1]$*.*

- *For all* $t \geq t_0 \geq 0$*, let* $I^{\theta,t_0}((t,t)-) := \lim_{s\uparrow t} I(s,s)^{\theta,t_0}$ *and specify that* $I^{\theta,t_0}(s,s) = 0$ *whenever* $s < t_0$*. If* $t_0' \geq t_0$ *and* $I^{\theta_{t_0},t_0}((t_0',t_0')-) < 1$*, then for all* $t_1,t_2 \geq t_0'$*,*

$$I^{\theta_{t_0'},t_0'}(t_1,t_2) = \frac{I^{\theta_{t_0},t_0}(t_1,t_2) - I^{\theta_{t_0},t_0}((t_0',t_0')-)}{1 - I^{\theta_{t_0},t_0}((t_0',t_0')-)}.$$

[13]It is notationally and conceptually simplest define the type process and news process for all times, even if the game has ended (that is, even if exit or adoption has occurred before ν).

[14]Note that $\mathcal{F}_t^X$ depends on the realization of the seller's randomization over pure strategies, through the latter's effect on the type process.

- *For all* $t_1, t_2 \geq t_0 \geq 0$, $t_2 < \infty \implies I^{H,t_0}(t_1, t_2) = 0$.

The first property is a straightforward extension of the one from Definition 1.2.1. The second property extends the second property of Definition 1.2.1 to hold across types in case L has upgraded by time t_0'. The third property merely rules out the possibility of H upgrading (to H). It follows that H's maximization problem is just a restatement of (MAX$^{\theta,s}$) for $\theta = H$. We use $\mathcal{I}^{\theta,s}$ to denote the support of type θ's strategy at time s. The third property above is thus equivalent to the statement that for all $s \geq 0$, if $(\tau^H, \nu^H) \in \mathcal{I}^{H,s}$, then $\nu^H = +\infty$. For any time $s \geq 0$, stopping time $\tau \geq s$ and the buyer strategy B, let $G(\tau; B, s)$ denote the maximand in (MAX$^{\theta,s}$) for $\theta = H$. For L at time s, the optimal stopping problem is[15]

$$\max_{\tau,\nu \geq s} \mathbb{E}^L \Bigg[\underbrace{\int_0^{\tau \vee \nu} -c(1 - e^{-r(t-s)}) + e^{-r(t-s)} dB^s_{t-}}_{\text{Buyer adopts first}} + \underbrace{\mathbb{1}\{\tau < \nu\} \left[-c(1 - e^{-r\nu})(1 - B^s_{\tau-}) \right]}_{L \text{ exits first}} + \underbrace{\mathbb{1}\{\nu \leq \tau\} \left[-c(1 - e^{-r\nu}) + e^{-r\nu}(-K + G(\tau; B, \nu)) \right] (1 - B^s_{\nu-})}_{L \text{ invests first}} | \mathcal{F}^W_s \Bigg]. \tag{QMAX^{L,s}}$$

Observe that the above problem nests the problem of H at the investment time ν, in the event that investment occurs before exit or adoption.

[15]Integration in QMAX^{L,s} is done with respect to time, t.

The buyer's optimal stopping problem is

$$\max_{\rho \geq s} \mathbb{E}\Bigg[\underbrace{e^{-r(\rho-s)}(1-k)\left(p_s+(1-p_s)I^{L,s}((\rho,\infty],\rho)\right)}_{\theta_s=H \text{ or } \theta_s=L \text{ and } L \text{ upgrades before } \rho}$$

$$\underbrace{-e^{-r(\rho-s)}kI^{L,s}((\rho,\infty]^2)}_{\theta_s=L \text{ and } L \text{ does not upgrade before } \rho} \quad |\mathcal{F}_s^X\Bigg]. \qquad (\text{QMAX}^{B,s})$$

Definition 1.3.2. *A stationary monotone equilibrium is a strategy profile* (I,B) *and a pair of belief processes* Z^-, Z *such that*

1. *Seller Optimality: for all* $s \geq 0$*, all* $(\tau,\nu) \in \mathcal{I}^{L,s}$ *solve* $(\text{QMAX}^{L,s})$.
2. *Buyer Optimality: for all* $s \geq 0$*, all strategies* $\rho \in \mathcal{B}^s$ *solve* $(\text{QMAX}^{B,s})$.
3. *Bayesian Consistency: for all* $t_1 \geq t_0 \geq 0$ *and* $\omega \in \Omega$ *such that* $I^{H,t_0}([t_0,t_1),\infty) \cdot I^{L,t_0}([t_0,t_1),\infty) < 1$*, the beginning-of-period belief* $Z^-_{t_1}$ *satisfies* (1.8.1). *If* $I^{H,t_1}(t_1,\infty) \cdot I^{L,t_1}(t_1,\infty) < 1$*, the middle-of-period belief* Z_{t_1} *solves* (1.8.2).
4. *Stationarity:* Z^- *and* Z *are time-homogeneous,* $\mathcal{F}_t^X$*-Markov process.*
5. *Monotonicity:*
   (a) *For all real* $s \geq 0$*, if* $(\tau_H,\nu_H) \in \mathcal{I}^{H,s}$*, then* $\tau_H \equiv +\infty$.
   (b) *There exists* $\alpha \in \mathbb{R}$ *such that for all real* $s \geq 0$*,* $\mathcal{B}^s = \{\rho^s_\alpha\}$*, where* $\rho^s_\alpha := \inf\{t \geq s : Z_t \geq \alpha\}$.

Since we have already shown that there is a maximum gap of K^* between value functions in the reflecting equilibrium $\Xi(\beta^*,\alpha^*)$, if $K \geq K^*$, never upgrading is dynamically optimal for L. The following proposition is immediate and requires no proof.

Proposition 1.3.2. *For $K \geq K^*$, the reflecting equilibrium $\Xi(\beta^*, \alpha^*)$, under the interpretation that L never upgrades, remains an equilibrium under Definition 1.3.2.*

1.3.3 Resetting Equilibrium

In the spirit of the baseline equilibrium, a natural conjecture is that an equilibrium exists for the game with endogenous quality, in which a reflecting barrier at β is created through investment instead of exit. However, this conjecture is false. For such an equilibrium to exist, it would be necessary that $V_H'(\beta) = V_L'(\beta) = 0$ and $V_H(\beta) = V_L(\beta) + K$, in addition to the same conditions at the adoption threshold as before. For any solution to such conditions, it must be that $V_H(z) - V_L(z)$ is convex at $z = \beta$, and thus the gap between value functions would exceed K for z sufficiently close to β from above. Type L would then strictly prefer to invest, which contradicts the optimality of waiting inside (β, α), and so the purported equilibrium unravels. This result is stated formally in the following lemma.

Lemma 1.3.1. *Let v_L, v_H and β satisfy* (1.2.5) *and* (1.2.6) *and the following conditions:*

$$v_H(\beta) > v_L(\beta)$$
$$v_L'(\beta) = v_H'(\beta) = 0.$$

Then $v_H''(\beta) - v_L''(\beta) > 0$. In particular, if $v_L(\beta) = v_H(\beta) - K$ for some $K > 0$, then for any $\alpha > \beta$ there exists $z \in (\beta, \alpha)$ such that $v_H(z) - v_L(z) > K$.

As shown in the static benchmark section, for sufficiently low investment cost K, the gap between value functions must exceed K. This means that the baseline equilibrium ceases to exist for low K – investment must occur in equilibrium. In what follows, we show that the above problems can be reconciled in equilibrium if there is a *resetting barrier* instead of a reflecting barrier. That is, instead of investing in a continuous fashion at β, L invests with an atom of probability such that upon reaching β, the belief immediately jumps or "resets" to some $z^* \in (\beta, \alpha)$, where it then resumes a continuous path. This is the first economics paper in which a player's reputation follows a resetting path.

We now introduce a *resetting equilibrium*, which is defined by a triple (β, z^*, α) such that $\beta < z^* < \alpha$ and has the following structure:

- For $z \geq \alpha$, adoption is immediate.
- For $z \leq \beta$, L upgrades with an atom of probability.
- Once $z \leq \beta$, the belief immediately resets to z^*.

The main result of this section is that a resetting equilibrium exists if and only if the upgrade cost is below a certain cutoff, and when it exists, it is uniquely identified. Definition 1.3.3 formalizes the concept above. Figure 1.4 illustrates the seller's value functions and a sample reputation path in a resetting equilibrium.

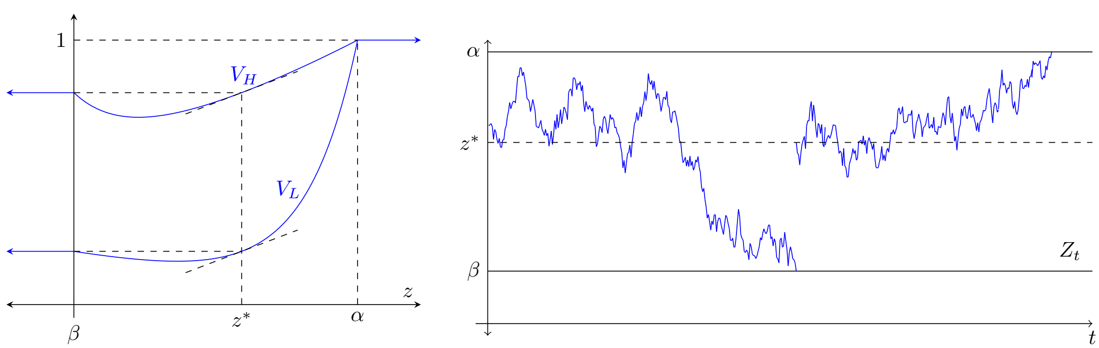

(a) Seller Value Functions (b) Sample Path with Resetting

FIGURE 1.4: Resetting Equilibrium

Definition 1.3.3. *A resetting equilibrium, denoted $\mathcal{R}(\beta, z^*, \alpha)$, is an equilibrium* *with strategy profile and belief process defined as follows:*

$$Z_t^- = \widehat{Z}_t + Q_{t-}$$

$$Z_t = \widehat{Z}_t + Q_t$$

$$Q_t = \sum_{s \in [0,t]: Z_s \leq \beta} (z^* - Z_s)$$

$$I^{H,t_0}(s,t) = 0 \qquad \textit{for all } s, t \in [t_0, \infty)$$

$$I^{L,t_0}(s,t) = \begin{cases} 1 - e^{-(Q_t - Q_{t_0})} & \textit{if } t \in [t_0, \infty) \textit{ and } s = \infty \\ 0 & \textit{otherwise.} \end{cases}$$

$$B_{t_1}^{t_0} = \mathbb{1}\{\exists s \in [t_0, t_1] : Z_s \geq \alpha\}.$$

In such an equilibrium, a number of conditions must hold for both players. Boundary conditions (1.3.5), (1.3.6), (1.3.11) and (1.3.12) are the same as before. Since resetting is instantaneous, all player's value functions must be the same immediately before and after the reset, as expressed in (1.3.7), (1.3.8) and (1.3.13). For L to be indifferent to investment, it must be that difference in value between continuing as H or as L must be exactly K, which is (1.3.9).[16] Finally, since the belief process is continuous at z^*, it is necessary that the difference $V_H(z) - V_L(z)$ reaches a local maximum at z^*. Due to the form of the solutions, this condition guarantees that

[16]Given (1.3.7)and (1.3.8), condition (1.3.9) could be stated equivalently as $V_H(\beta) = V_L(\beta) + K$.

$V_H(z) - V_L(z) \leq K$ for all $z \in [\beta, \alpha]$.

$$V_H(\alpha) = 1 \tag{1.3.5}$$

$$V_L(\alpha) = 1 \tag{1.3.6}$$

$$V_H(\beta) = V_H(z^*) \tag{1.3.7}$$

$$V_L(\beta) = V_L(z^*) \tag{1.3.8}$$

$$V_H(z^*) = V_L(z^*) + K \tag{1.3.9}$$

$$V_H'(z^*) = V_L'(z^*) \tag{1.3.10}$$

$$V_B(\alpha) = p(\alpha) - k \tag{1.3.11}$$

$$V_B'(\alpha) = p'(\alpha) \tag{1.3.12}$$

$$V_B(\beta) = V_B(z^*). \tag{1.3.13}$$

In total, there is a nonlinear system of 9 equations and 9 unknowns: there are three relevant belief states β, z^*, α and given these there are two constants each for H, L, and the buyer that determine their value functions. Figure 1.4 shows value functions for the seller in a resetting equilibrium. Note that at β, L is indifferent between continuing at z^* as type L or paying the upgrade cost to "jump tracks" and continue at z^* as H. In addition to the conditions above, a necessary condition for equilibrium is that $V_L \geq 0$, otherwise L would strictly prefer to exit. We show (Theorem 1.3.1) that a solution to this system exists and is an equilibrium, provided that K is sufficiently small.

The proof of existence and uniqueness of resetting equilibrium is an order magnitude more difficult than that of $\Xi(\beta^*, \alpha^*)$, and is contained in the appendix. Nonetheless, the steps in the proof provide an illustration of how the various parameters work together, and we outline these steps here. Most of the proof consists of establishing uniqueness the seller's parameters for a fixed α; once those are in place, it is relatively easy to show that the buyer's payoff is single-peaked in the adoption threshold, and that there is a single fixed point for players' best responses.

Substituting (1.2.7) and (1.2.8) into (1.3.5)-(1.3.10), we obtain

$$C_1^H e^{(m-1)\beta} + C_2^H e^{-m\beta} = C_1^H e^{(m-1)z} + C_2^H e^{-mz} \tag{1.3.14}$$

$$C_1^L e^{m\beta} + C_2^L e^{(1-m)\beta} = C_1^L e^{mz} + C_2^L e^{(1-m)z} \tag{1.3.15}$$

$$C_1^H e^{(m-1)z} + C_2^H e^{-mz} = C_1^L e^{mz} + C_2^L e^{(1-m)z} + K \tag{1.3.16}$$

$$(m-1)C_1^H e^{(m-1)z} - mC_2^H e^{-mz} = mC_1^L e^{mz} + (1-m)C_2^L e^{(1-m)z} \tag{1.3.17}$$

$$C_1^H e^{(m-1)\alpha} + C_2^H e^{-m\alpha} = 1 + c \tag{1.3.18}$$

$$C_1^L e^{m\alpha} + C_2^L e^{(1-m)\alpha} = 1 + c. \tag{1.3.19}$$

1. For each α, there is at most one pair (β, z^*) with $\beta < z^* < \alpha$ that solve (1.3.14)-(1.3.19).

    a. By translation invariance, normalize $z^* = 0$. Parameterize the system (1.3.14)-(1.3.19) using $V_L(0) = C_1^L + C_2^L - c = x$.

    b. There is an increasing function $\beta(x)$, defined on $((m-1)K - c, \infty)$, such that for each x, $\beta(x)$ is the unique β that solves (1.3.14)-(1.3.17).

    c. Relax (1.3.18) and (1.3.19) by replacing $1 + c$ with $y + c$.

    d. Given x, there is a unique pair $(\alpha(x), y(x))$ of values for (α, y) that solve the relaxed system.

    e. The function $\alpha(x)$ is decreasing and $y(x)$ is increasing.

    f. We have $\lim_{x\uparrow 1-K} y(x) > 1$. The limit $\lim_{x\downarrow(m-1)K-c} y(x) < 1$ if and only if $\frac{m^m K}{(m-1)^{m-1}} < 1 + c$, which is thus a necessary and sufficient condition for existence and uniqueness of (x, β, α) solving the system with $y = 1$.

2. Establish a cutoff rule for the upgrade cost that determines whether V_L is globally nonnegative.

a. Using K as an independent variable, there is a differentiable function $v_L(K)$ that gives the global minimum of V_L pinned down by Step 1.

b. $v_L(K)$ is strictly decreasing, positive for sufficiently small K, and negative as $\lim_{K \uparrow K_0}$ where K_0 is defined implicitly as the unique solution to $\frac{m^m K}{(m-1)^{m-1}} = 1 + c$. Hence, there exists a unique cutoff K^{**} such that $v_L(K^{**}) = 0$, and $v_L(K) \geq 0$ if and only if $K \leq K^{**}$.

3. Derive the buyer's response to a given equilibrium resetting strategy of the seller, uniquely determined by β, as a strictly increasing function $\alpha(\beta)$ with derivative strictly less than 1. It follows that there is a unique fixed point (β^*, z^*, α^*).

Theorem 1.3.1. *There exists a cutoff K^{**} such that a resetting equilibrium exists if and only if $K \leq K^{**}$, and when it exists, it is unique.*

1.3.4 Skew-Resetting Equilibrium

For intermediate upgrade costs, neither the baseline equilibrium nor the resetting equilibrium of the previous section exists. In this situation, K is small enough that investment must occur in equilibrium, but too large for resetting equilibrium as V_L would strictly prefer exit at intermediate beliefs. I show that a hybrid of the previous two types of equilibria exists for intermediate values of K. Type L must exit with some probability at some intermediate state between the resetting boundary and the adoption boundary, but not so frequently as to create a reflecting barrier. In particular, L can mix so that the equilibrium belief process follows a *skew Brownian motion*, which is a generalization of reflected Brownian motion.

For a large class of stochastic processes, one can define a continuous, nondecreasing stochastic process, known in the probability literature as *local time*. The local time of a real-valued process X at a point $x \in \mathbb{R}$ is defined as

$$\ell_t^X(x) := \lim_{\epsilon \downarrow 0} \frac{1}{2\epsilon} \text{meas}\{s \in [0,t] : |X_s - x| \leq \epsilon\}. \tag{1.3.20}$$

In essence, the local time of X at x is a scaled measure of the time the process has spent at x. In many applications, a process undergoes some regulation at a single point, and therefore must recursively satisfy an equation involving its own local time. For concreteness, consider the stochastic differential equation

$$X_t = W_t + \delta \ell_t^X(0), \tag{1.3.21}$$

where W_t is standard Brownian motion, and $|\delta| \leq 1$. Harrison and Shepp (1981) show that a solution to (1.3.21) exists if and only if $|\delta| \leq 1$ and is unique when it exists. Such a solution is known as *skew Brownian motion.*[17] Le Gall (1984) generalizes this result to stochastic differential equations involving local times on larger sets of points.

For $\delta = 0$, (1.3.21) collapses and the solution is simply $X_t = W_t$. As noted by Harrison and Shepp (1981), for $\delta = 1$ (resp. $\delta = -1$), the solution is upward-reflected (resp. downward-reflected) Brownian motion. Hence skew Brownian motion generalizes both standard Brownian motion and reflected Brownian notion.

Remark 2. *Harrison and Shepp (1981) further show that X_t solving* (1.3.21) *can be constructed as the limit of discrete-time modified random walks. These random walks have standard, symmetric behavior everywhere except at the origin, where the probability of an upward step is adjusted according to δ. This interpretation is not precisely appropriate for this model, since the upward regulation at the point of skew is the result of Bayes' rule, which must be applied deterministically whenever the*

[17] Originally, the concept of skew Brownian motion was introduced by Ito and McKean (1965) less rigorously as a random walk with the sign of excursions flipped at the origin. Harrison and Shepp (1981) show formally that the solutions to (1.3.21) in fact coincide with the concept of Ito and McKean (1965).

belief hits that point. However, the random walk of Harrison and Shepp (1981) can be adapted to the present model as follows. Instead of increasing the probability of an upward step at the origin, insert a deterministic, fractional upward step, so that the expected value of the adjustment is the same. By the law of large numbers, these random walks must have the same limit.

We now define and construct a skew-resetting equilibrium, which is characterized by five parameters $\beta, \hat{z}, z^*, \alpha$ and δ as follows:

- For $z \geq \alpha$, adoption is immediate.
- For $z \leq \beta$, L upgrades with an atom of probability
- Once $z \leq \beta$, the belief immediately resets to z^*.
- At $z = \hat{z}$, L exits at instantaneous rate $\delta \in (0, 1)$.

Definition 1.3.4 formalizes this idea.

Definition 1.3.4. *A skew-resetting equilibrium, denoted $\Psi(\beta, \hat{z}, z^*, \alpha, \delta)$, is the strategy profile (I, B) and pair of belief processes (Z^-, Z) such that:*

$$Z_t^- = \widehat{Z}_t + Q_{t-} + L_{t-}$$

$$Z_t = \widehat{Z}_t + Q_t + L_t$$

$$Q_t = \sum_{s \in [0,t]: Z_s \leq \beta} (z^* - Z_s)$$

$$L_t = \delta \phi^{-2} \ell_t^Z(\hat{z})$$

$$I^{H,t_0}(s,t) = 0 \qquad \textit{for all } s, t \in [t_0, \infty)$$

$$I^{L,t_0}(s,t) = 1 - e^{-(Q_t - Q_{t_0} + L_s - L_{t_0})} \qquad \textit{for all } s, t \in [t_0, \infty)$$

$$B_{t_1}^{t_0} = \mathbb{1}\{\exists s \in [t_0, t_1] : Z_s \geq \alpha\}.$$

The parameters β, z^*, and α have the same meaning as in the resetting equilibrium. At $\hat{z}$, type L becomes indifferent to exit and mixes. The belief process is regulated upward at $\hat{z}$, but unlike reflected Brownian motion, it immediately reaches states strictly above and states strictly below $\hat{z}$. Figure 1.5 shows the belief process with (Z) and without $(\hat{Z})$ regulation. Note that the process $L = Z - \hat{Z}$ is only increasing when Z is at the point $\hat{z}$.

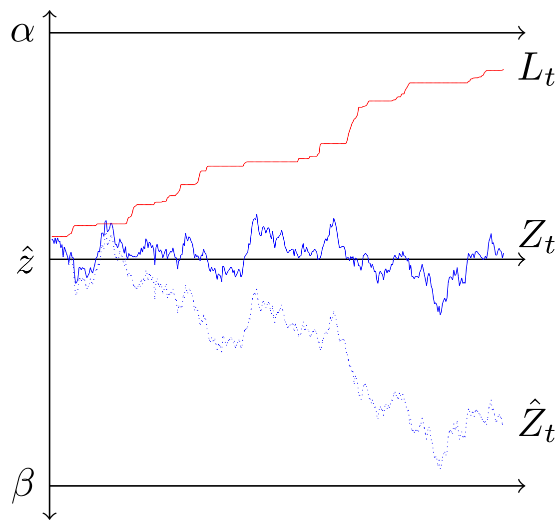

FIGURE 1.5: Sample path in skew-resetting equilibrium.

A priori, due to the inherent asymmetry of skew Brownian motion, value functions need not be differentiable at $\hat{z}$. In fact, in a skew-resetting equilibrium, V_H is not differentiable at $\hat{z}$, and moreover, V_H is only piecewise convex. As Figure 1.6 illustrates, V_H has a concave kink at $\hat{z}$.

Although differentiability need not hold in general at points of skew, the left- and right-derivatives of each value function must jointly satisfy certain conditions related to the intensity of regulation. In the current model, the intensity of exit determines a parameter $\gamma \in (1/2, 1)$ that measures the degree of skew at $\hat{z}$, and value functions must jointly satisfy the conditions

$$\gamma V_H'(\hat{z}+) = (1-\gamma)V_H'(\hat{z}-) \tag{1.3.22}$$

$$\gamma V_L'(\hat{z}+) = (1-\gamma)V_L'(\hat{z}-). \tag{1.3.23}$$

A heuristic interpretation of these conditions is as follows. Using the language of a random walk, starting from $\hat{z}$, the belief process, after incorporating the exogenous news process and conditional on no exit, moves up a step with probability γ and down with probability $1-\gamma$. The value function at $\hat{z}$ must incorporate any expected gain from these movements, and due to the unbounded variation of Brownian motion, these movements are immediate, so the net expected gain must be zero. For type H, the high probability (γ) of an upward step attaches to a small increase in value ($V_H'(\hat{z}+)$) to the right of $\hat{z}$ in order to offset the low probability ($1-\gamma$) of a downward step which yields a large decrease in value ($V_H'(\hat{z}-)$). For type L, the same logic applies, but since $V_L(\hat{z}) = 0$ and V_L is bounded below by 0, his value function must be flat on both sides: $V_L'(\hat{z}+) = V_L'(\hat{z}-) = 0$. It follows that V_L is smooth inside (β, α) and its closed form (1.2.8) involves just two constants. On the other hand, V_H is composed piecewise of two distinct convex functions, which we denote $V_{H,-}$ and $V_{H,+}$, and involves four constants.

In total, the seller's boundary conditions in skew-resetting equilibrium are as follows.

$$V_{H,+}(\alpha) = 1 \tag{1.3.24}$$

$$V_L(\alpha) = 1 \tag{1.3.25}$$

$$V_{H,-}(\beta) = V_{H,+}(z^*) \tag{1.3.26}$$

$$V_L(\beta) = V_L(z^*) \tag{1.3.27}$$

$$V_{H,+}(z^*) = V_L(z^*) + K \tag{1.3.28}$$

$$V'_{H,+}(z^*) = V'_L(z^*) \tag{1.3.29}$$

$$V_L(\hat{z}) = 0 \tag{1.3.30}$$

$$V'_L(\hat{z}) = 0 \tag{1.3.31}$$

$$V_{H,+}(\hat{z}+) = V_{H-}(\hat{z}-). \tag{1.3.32}$$

On the buyer's side, we use $V_{B,+}$ and $V_{B,-}$ to denote the value function restricted to $[\hat{z}, \infty)$ and $(-\infty, \hat{z}]$, respectively. By construction, this requires (1.3.36). The usual value matching and smooth pasting conditions, (1.3.33) and (1.3.34), apply at the adoption threshold. Condition (1.3.37) adapts (1.3.13) to the piecewise construction in light of the fact that z^* lies to the right of $\hat{z}$ and β lies to the left. Equation (1.3.35) is conceptually a hybrid of (1.2.19) and (1.3.22). The term on the left side is the value lost in the event of exit, which occurs with probability $1 - p(\beta) = \frac{1}{1+e^{\beta}}$ times instantaneous exit rate $(2\gamma - 1)$. On the right side is the expected instantaneous change in value upon no exit: the gain $V'_{B,+}(\hat{z}+) > 0$ is weighted by $\gamma > \frac{1}{2}$

and the loss of magnitude $V'_{B,-}(\hat{z}-) > V'_{B,+}(\hat{z}+)$ is weighted by $1-\gamma < \frac{1}{2}$.

$$V_{B,+}(\alpha) = p(\alpha) - k \tag{1.3.33}$$

$$V'_{B,+}(\alpha) = p'(\alpha) \tag{1.3.34}$$

$$\frac{2\gamma-1}{1+e^{\beta}} V_{B,+}(\hat{z}) = \gamma V'_{B,+}(\hat{z}+) - (1-\gamma) V'_{B,-}(\hat{z}-) \tag{1.3.35}$$

$$V_{B,-}(\hat{z}) = V_{B,+}(\hat{z}) \tag{1.3.36}$$

$$V_{B,-}(\beta) = V_{B,+}(z^*). \tag{1.3.37}$$

By translation invariance, one of the parameters can be set to 0 without loss of generality. The values $V'_{H,-}(\hat{z}-)$ and $V'_{H,+}(\hat{z}+)$ that emerge can be used to back out the value of γ that solves (1.3.22). In turn, this pins down the skew parameter, $\delta = 2\gamma - 1$, which describes L's exit strategy.

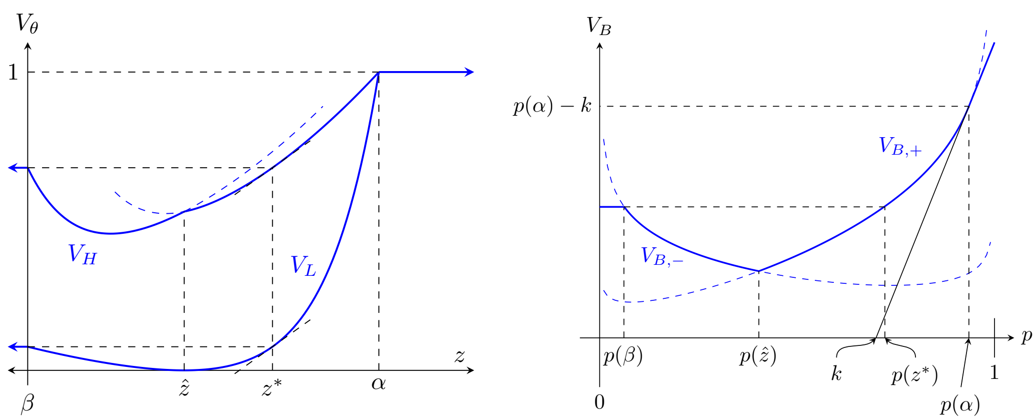

(a) Seller Value Functions (b) Buyer Value Function

FIGURE 1.6: Value functions in skew-resetting equilibrium.

Theorem 1.3.2. *The cutoffs satisfy $K^{**} < K^*$, and a skew-resetting equilibrium exists if and only if $K \in (K^{**}, K^*)$. When it exists, it is unique.*

1.3.5 Section Summary

We have defined three particular forms of equilibrium – reflecting, resetting, and skew-resetting. The main results of this section, namely Proposition 1.3.2 and Theorems 1.3.1 and 1.3.2, show that among these three forms, an equilibrium exists and is uniquely pinned down. The particular form of this equilibrium depends on the upgrade cost K as follows:

- For small K, the equilibrium is resetting.
- For intermediate K, the equilibrium is skew-resetting.
- For large K, the equilibrium is reflecting.

As K increases, the equilibrium morphs from resetting to reflecting. When K reaches K^{**} from below, L begins exiting stochastically at the intermediate state $\hat{z}$. This exit then forms a permeable barrier for the belief process, and the rate of exit increases as K goes from K^{**} to K^{*}. Once K reaches K^{*}, exit is with full intensity, causing the belief process to reflect upward: the skew point $\hat{z}$ thus becomes the original reflecting barrier, β^{*}.

Although we state the above results in terms of cutoffs for upgrade costs, numerical simulations indicate that similar cutoffs exist for news quality, with higher news quality (i.e., lower m) and lower upgrade costs acting as substitutes. In the

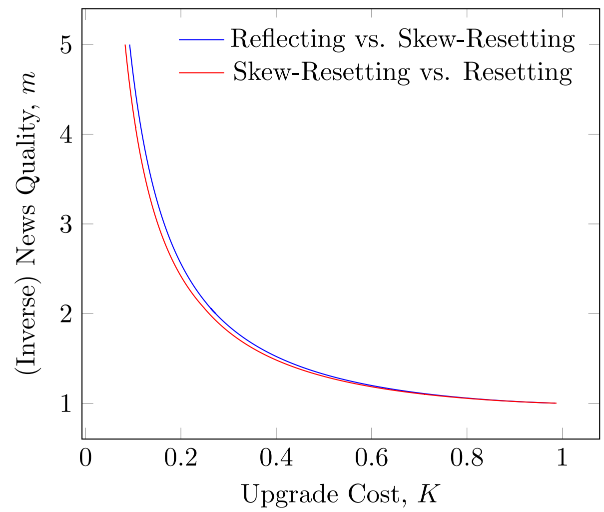

FIGURE 1.7: Boundaries between reflecting, resetting, and skew-resetting equilibrium regions. Reflecting occurs northeast of the upper curve. Resetting equilibrium occurs southwest of the upper curve, and skew-resetting occurs in between.

1.4 Variations on Baseline Model

In this section, I return to the baseline model and present three variations in order to analyze the roles of private information and commitment power, and the interaction between the two.

1.4.1 Symmetric Incomplete Information

We now model the seller's type as incomplete information to both players, who have a common prior z_0 and learn together from the news process. Value functions for both players, V_S and V_B, satisfy the ODE (1.2.15) but with different boundary conditions. At some threshold α, the buyer adopts, and at some threshold β, the seller exits. As Proposition 1.4.1 states, both thresholds are lower than the respective thresholds in the baseline model. To see why, first consider a fixed lower threshold, β. Since the seller does not know his type, both types exit at β, which ends the game. This reduces the buyer's value to 0 at β, reducing the value of delay. Thus the buyer's best response function shifts down. Next, consider a fixed adoption threshold, α. For any given belief, a the seller who does not know his type is more optimistic than type L in the baseline model. This implies that the seller's best response function also shifts down.

The buyer is unambiguously worse off when incomplete information is symmetric, compared to the baseline model of private information. This result stands in stark contrast to the classic lemons result of Akerlof (1970), in which a used car buyer is worse off when the seller has private information about the car's quality. The reason for this result is that the buyer prefers for L to exit, but under symmetric incomplete information, L remains in the game longer, that is, at lower beliefs. Moreover, the buyer prefers for H to remain, but H exits with positive probability, unlike in the baseline model.[18] An implication of this result is that, in a richer model, the buyer would benefit from the opportunity to pay some small cost to finance the seller's effort to determine his type, even without any requirement that the seller report his findings.

For the seller and hence social welfare, the welfare comparison to the baseline model depends on the starting belief. There are two forces acting in opposite directions. With private information, the seller avoids unnecessary exit should he be H and avoids excessive continuation costs should he be L. On the other hand, as noted above, the adoption threshold is higher under private information, which hurts the seller. The seller prefers to maintain incomplete information for high prior beliefs, but prefers the private information environment for low priors. For high priors, the cost of delaying adoption under private information is borne in the near future, while the benefit of correctly exiting or not exiting is discounted heavily as the lower boundary is farther away.

Given α, V_S must satisfy the value matching condition $V_S(\beta) = 0$ and the smooth

[18] To the latter point, we acknowledge that in the baseline model we specifically focus on equilibria without exit by H, but we conjecture that a priori weaker restrictions of scope would produce the same equilibria in that model.

pasting condition $V_S'(\beta) = 0$. Similarly, given β, V_B must satisfy $V_B(\alpha) = p(\alpha) - k$ and $V_B'(\alpha) = p'(\alpha)$. We use $\bar{V}(z) := p(z)V_H(z) + (1 - p(z))V_L(z)$ to denote the ex ante expected utility for the seller in the private information setting.

Proposition 1.4.1. *There exists a unique stationary equilibrium of the game with symmetric incomplete information, with thresholds* $\beta^{*,sym} < \beta^*$ *and* $\alpha^{*,sym} < \alpha^*$.

Corollary 1.4.1. *For all beliefs, the buyer is worse off than under private information by the seller. For the seller and social welfare, the comparison depends on the starting belief. There exists* $\hat{z} \in (\beta^{*,sym}, \alpha^{*,sym})$ *such that for* $z_0 < \hat{z}, V_S(z_0) < \bar{V}(z_0)$ *and for* $z_0 \in (\hat{z}, \alpha)$, $V_S(z_0) > \bar{V}(z_0)$.

1.4.2 Commitment

In this section we allow the buyer to commit to choosing a threshold α and adopting at the first time $Z_t \geq \alpha$. Given α, the seller faces the same problem as in the baseline model, and thus his response is the same; L exits stochastically at $\beta^*(\alpha)$ to create a reflecting barrier for the belief process conditional on no exit. Now recall that $b^*(\alpha)$ is increasing. Hence by marginally raising the adoption theshold α, the buyer induces L to exit sooner. The buyer weakly benefits from this effect since she does not want to adopt L. Put differently, committing to a higher threshold raises the reflecting barrier, which raises the future belief process in the sense of first-order stochastic dominance, improving the buyer's prospects. It follows that the optimal threshold under commitment, which we denote α^{**} when it exists, is at least α^*. For some parameter values, specifically those for which $x^* = \alpha^* - \beta^*$ is below a certain cutoff, the buyer's value is strictly increasing in the commitment threshold, which results in an open set problem. To see why, suppose the starting belief is z_0 and we restrict the buyer's choice of thresholds to the set $[z_0, z_0 + x^*]$, that is the set of thresholds α for which $z_0 \in [\beta(\alpha), \alpha]$. Further suppose that the constrained optimal threshold is the corner $z_0 + x^*$. Then further increases in α come without cost to the buyer: conditional on either type, the expected time until adoption is the same (i.e., there is no additional delay), but now L exits at time 0 with a larger atom of probability. It follows that the buyer prefers to set a threshold "as high as possible", but in these cases we can characterize her supremum payoff, which is finite.[19]

Surprisingly, for parameter values such that α^{**} is finite and above z_0, the commitment threshold is decreasing in z_0. By the above discussion, the benefit of commitment lies in inducing low types to exit sooner. Put simply, this benefit diminishes when there are fewer low types. More specifically, a higher starting belief has two implications. First, exit by L becomes less likely, as the belief process must fall farther to reach the lower boundary. Second, it is less likely that $\theta = L$ in the first place. In fact, as starting beliefs approach the competitive threshold, α^*, the benefit of delay under commitment vanishes and we have $\alpha^{**} \searrow \alpha^*$.

Proposition 1.4.2. *For starting beliefs $z_0 \geq \alpha^*$, adoption is immediate in the unique equilibrium under commitment. For large x^* and intermediate z_0, there is a unique equilibrium with adoption threshold $\alpha^{**}(z_0)$ and reflecting barrier at $\beta^{**}(z_0) = \alpha^{**}(z_0) - x^*$; moreover, $\alpha^{**}(z_0) > \alpha^*$ and α^{**} is decreasing in z_0. For all remaining cases, the buyer's supremum payoff is approximated by committing to arbitrarily high thresholds.*

1.4.3 Commitment with Symmetric Incomplete Information

By the logic of the previous two sections, the buyer benefits from a higher exit threshold when under private information but benefits from a lower exit threshold under symmetric information. Under private information, commitment power helps the buyer by allowing her to raise her adoption threshold and thereby increase the exit

[19] Bolton (1987) discusses such an open set problem under commitment in the context of crime deterrence, where the optimal punishment may be as harsh as possible.

threshold. Under symmetric information, this is of no advantage. Under symmetric information and commitment *together*, the buyer optimally commits to a threshold $\alpha^{**,sym}(z_0)$ that is lower than the competitive one, $\alpha^{*,sym}$. The benefit of committing to a lower threshold lies in encouraging the seller to remain in the game longer, while the cost is in adopting more low types. For low starting beliefs, the benefit is weighted highly, since exit is very likely, and any lowering of the exit threshold greatly increases the chances of reaching the adoption threshold. It follows that $\alpha^{**,sym}$ is increasing in the starting belief, but bounded above by $\alpha^{*,sym}$.

Proposition 1.4.3. *Under commitment and symmetric incomplete information together, if $z_0 \geq \alpha^{*,sym}$, then any threshold $\alpha \leq z_0$ (immediate adoption) is optimal, and if $z_0 \leq \beta^{sym}(z_m)$, then any threshold $\alpha \geq z_m$ (immediate exit) is optimal. For $z_0 \in [\beta^{sym}(z_m), \alpha^{*,sym}]$, the buyer's optimal threshold policy is given by a strictly increasing, compact-valued, upper-hemicontinuous and almost everywhere single-valued correspondence, $\boldsymbol{\alpha} : [\beta^{sym}(z_m), \alpha^{*,sym}] \rightrightarrows [z_m, \alpha^{*,sym}]$, with $\max \boldsymbol{\alpha}(z_0) < \alpha^{*,sym}$ for all $z_0 < \alpha^{*,sym}$.*[20]

1.5 Comparative Statics

In this section I provide show how equilibrium thresholds and utilities change with respect to underlying parameters.

Cost of Adoption, k: An increase in k causes the buyer to raise her adoption threshold, and has a negative direct effect on her utility. Indirectly, this raises the seller's exit threshold, which benefits the buyer. In the baseline model, the direct effect always dominates, and the buyer is made strictly worse off for all starting beliefs. The seller is made weakly worse off in equilibrium, and strictly so for intermediate starting beliefs. With endogenous quality, however, increase in k may help either

[20] By strictly increasing, we mean that for all $z_0' > z_0$, $\min \boldsymbol{\alpha}(z_0') \geq \max \boldsymbol{\alpha}(z_0)$.

player. Intuitively, in a resetting equilibrium, the indirect benefit to the buyer of raising the seller's threshold is larger for two reasons: investment by L is better than exit for the buyer, and investment occurs with an atom of probability, whereas exit occurs continuously in the baseline model. The net effect on sellers of increasing k is a translation of value functions to the right; since these value functions are U-shaped in a resetting equilibrium, the change helps for some starting beliefs and hurts for others.

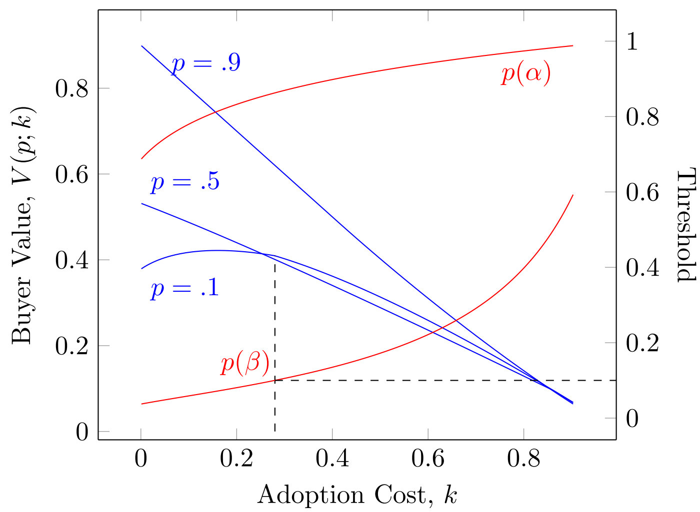

FIGURE 1.8: Resetting equilibrium: buyer value and equilibrium thresholds for fixed p and varying k.

Under symmetric information, an increase in k results in an upward shift of both thresholds, leaving both players weakly worse off.

Proposition 1.5.1. *In the baseline model, under both private and symmetric information, all value functions are weakly decreasing in k. Moreover, all thresholds are increasing in k.*

Operating Cost, c: In the baseline model, and increase in c has a direct negative effect on L and causes him to exit sooner, raising the reflecting barrier. The buyer is thus unambiguously made weakly better off in equilibrium, and raises her adoption threshold, which has two unsurprising implications: (i) L is made weakly worse off at all beliefs (and strictly at some), and (ii) there always exist belief at which H is made strictly worse off. However, the fact that the reflecting barrier moves closer to the adoption barrier exerts a positive externality on H: starting from β^*, the adoption occurs stochastically sooner. We have numerical evidence that for certain values of parameters and starting belief z, this positive indirect effect outweighs the direct effect, and $V_H(z)$ is nonmonotonic in c. Figure 1.9 shows one such example.[21]

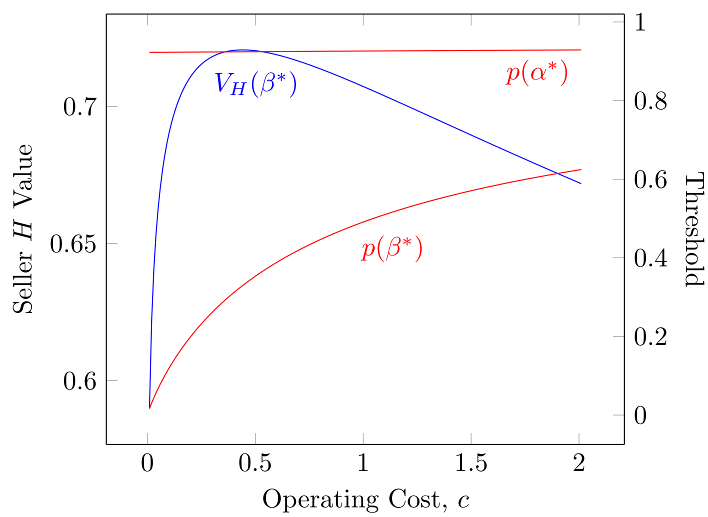

FIGURE 1.9: Reflecting equilibrium: An example of V_H nonmonotonic in c.

Under symmetric information, the indirect effect on the buyer of an increase in c reverses; the buyer is made worse off and lowers her adoption threshold, which in turn helps the seller. The balance of this benefit and the direct disbenefit depends on the belief. For beliefs just below the adoption threshold, the marginal benefit to the buyer of the lower adoption threshold is large, while the marginal disbenefit of higher operating costs is small.

Proposition 1.5.2. *In the baseline model under private information, V_B is weakly*

[21] Although β^* itself is not fixed, V_H is constant in z for $z \leq \beta^*$, so it suffices to observe any nonmonotonicity in $V_H(\beta^*)$.

increasing in c while V_L is weakly decreasing and V_H is nonmonotonic. Under symmetric information, V_B is again weakly increasing in c and V_S is nonmonotonic.

Discount rates: For clarity it is helpful to decouple the discount rates of the two players; we start with $r_S = r_B = r > 0$ and consider increasing r_S or r_B in isolation. Increasing r_S has a similar direct effect to increasing c, raising the exit threshold and thus under private information, raising the adoption threshold and helping the buyer. However, in the previous analysis, raising c hurts type H relatively less than it hurts type L, since H attains adoption sooner; but raising r_S hurts H more than L because the discounting is applied to a larger value. It follows that the benefit to H of shrinking the delay region does not overcome the disbenefit of extra discounting, and both types are made worse off. We now turn to r_B, starting private information. In the baseline model, increasing r_B is similar to increasing k in that it has a negative direct effect on the buyer's utility. On the other hand, it shifts down the buyer's best response function, which has an indirect effect of shifting down the seller's exit threshold. Under private information, this indirect is harmful and reinforces the direct effect, so the buyer suffers while the seller benefits.

Under symmetric information, raising r_S causes S to exit sooner, which hurts B and lowers the adoption threshold. Consequently, S may be worse off or better off depending on the starting belief. Raising r_B helps S as under private information, but the indirect downward shift of the seller's exit threshold positively affects the buyer by increasing the likelihood of reaching the adoption barrier first. For beliefs near the adoption threshold, this benefit is negligible and the negative effect above dominates. For beliefs near the exit threshold, the positive effect dominates. In the starkest example, starting at exactly the exit threshold, a small increase in the discount rate makes adoption possible at all.

Proposition 1.5.3. *In the baseline model under private information, the following* *relatioships hold (weakly): V_B is decreasing in r_B while both V_H and V_L are increasing; V_B is increasing in r_S, while both V_H and V_L are decreasing. Under symmetric information, V_S is increasing in r_B and V_B is nonmonotonic; V_B is decreasing in r_S and V_S is nonmonotonic.*

1.6 Discussion

In this section I discuss several applications, alternative models and give a concluding summary with directions for future research.

1.6.1 Applications

Although the model has been described in the context of a startup and venture capitalist, and more generally buyers and sellers of goods or services, there are many other realistic applications. The model fits the situation of a lobbyist for some special interest who exerts effort to influence a politician to adopt a certain position. The lobbyist knows whether the politician would ultimately agree with the special interest and benefit from the position, and the politician is likely to adopt the position once she is sufficiently convinced of its value. Moreover, the lobbyist can exert effort to improve the quality of the position, making it more aligned with the politician's preferences. For example, the lobbyist can consult health economists to better formulate a health care proposal.

Another application is a job market. The product is the job candidate himself, who knows the value of his skills, and pays a cost to engage in networking activities or acquire credentials that signal his value to a potential employer. The job candidate can also upgrade his skills through, say, independent research or informal project involvement, which may not be immediately observable by future employers.

A third application is a dating or marriage market. Presumably there are fundamentals that determine mate quality, such as health and financial security, and individuals engage in many kinds of costly activities to signal these qualities. They can also take less conspicuous actions that improve the fundamentals.

1.6.2 Alternative Models

A natural extension of the model is to include two or more buyers. If all news is public and the buyers face a preemption problem, with adoption only for the first buyer to act, then the buyers compete away their profits, and the adoption threshold drops to the static threshold. The seller's best response would not change. On the other hand, if buyers privately observe different (but correlated) news processes, there would be a "winner's curse" effect that offsets the preemption incentive, and it is unclear whether the resulting threshold would be higher or lower than with one buyer. A difficulty of this model would be that a buyer's inference, based on having not observed adoption by the other buyer, depends both on the observed news path and the calendar time; as time elapses, the variance of the rival buyer's possible belief increases, so the negative inference given no adoption becomes more severe. The answer to this question would have important implications for applications in which a seller must choose how many buyers to approach.

Although we discuss commitment to threshold rules, it is possible that a buyer might commit *not* to buy a good if the reputation ever falls below a certain threshold. This would provide the buyer a more direct channel through which to induce earlier exit by L, without having to raise the adoption threshold. Presumably, the seller would benefit from this ability, and would commit to a policy involving both an adoption rule and a "deal breaker" rule.

In the spirit of endogenous quality, one might consider endogenous private information; that is, the seller might begin uninformed and have the option to pay some cost to privately learn his type along the path of play. This adds a dimension to the state variable, as the buyer must assess whether the seller is H or L and informed or not.

1.6.3 Conclusion

We have analyzed a model of both strategic adoption and endogenous quality, extending the current literature in two dimensions. In the baseline model, despite additional technical difficulties, we have shown the existence and uniqueness of equilibrium under mild refinement. We have compared the baseline model to three variations and found that a rich interaction exists between commitment power and the information structure. When the seller has private information, the buyer leverages commitment power by raising her adoption threshold, but under symmetric incomplete information, this relationship reverses. After extending the model to endogenous quality, we have constructed a type of equilibrium new to reputation games, in which the belief process follows a discontinuous resetting pattern. For certain parameter values, we have shown that further novel behavior arises, namely skew Brownian motion.

1.7 Appendix A: Baseline Model

1.7.1 Stationary Weakly Monotone Equilibium

Definition 1.7.1. *A stationary weakly monotone equilibrium (SWME) is an $\mathcal{F}_t$-adapted public belief process $\{Z_t\}_{t\geq 0}$ and strategy profile satisfying conditions 1-3 of Definition 1.2.2 and the following:*

4. *Seller Stationarity: condition 4 of Definition 1.2.2.*

5. *Buyer Stationarity: there exist sets $\mathcal{C}$ and $\mathcal{D}$ and functions $f^C, f^D : \mathbb{R} \to \mathbb{R}_+$ such that:*

   (a) $\mathcal{C} \cup \mathcal{D} = \mathbb{R}$

   (b) $B_{t_1}^{t_0} = 1 - \exp\left(\int_{t_0}^{t_1} -f^C(Z_t) + \ln(1 - f^D(Z_t))\mathbb{1}\{Z_t \in \mathcal{D}\}dt\right).$

6. Weak Monotonicity: condition 5a of Definition 1.2.2.

Essentially, condition 5 says that strategies can be decomposed into continuous intensities and discontinuous jumps that depend only on the current state, Z_t.[22]

We now state and prove a stronger version of Theorem 1.2.1, from which Theorem 1.2.1 immediately follows as a corollary.[23]

Theorem 1.7.1. *There is a unique SWME, and it is the reflecting equilibrium* $\Xi(\beta^*, \alpha^*)$.

The proof is divided into two parts. In Section 1.7.2, we show that there is a unique solution (β^*, α^*) to the system of boundary conditions outlined in Section 1.2. In Section 1.7.3, we show that any SWME equilibrium must be a reflecting equilibrium, that in particular the unique candidate is $\Xi(\beta^*, \alpha^*)$, and that this candidate is in fact an SWME.

1.7.2 Uniqueness of Parameters

The next two lemmas refer to conditions stated in Section 1.2 and for now stand as purely algebraic claims; that these conditions are necessary in equilibrium is shown as part of the proof of Lemma (1.7.9).

Lemma 1.7.1. *There is an increasing, linear function* $\beta^* : \mathbb{R} \to \mathbb{R}$ *such that* $\beta^*(\alpha)$ *is the unique solution* $\beta \leq \alpha$ *to* (1.2.8) *and* (1.2.10)-(1.2.11).

---

[22]The technically inclined reader may note that there exist certain pathological distributions, such as the Cantor distribution, that are not mixtures of discrete and continuous distributions. We rule out those possibilities in our analysis.

[23]To see that Buyer Stationarity is weaker than condition 5b, note that any cutoff rule α can be represented by $\mathcal{C} = (-\infty, \alpha)$, $\mathcal{D} = [\alpha, \infty)$, $f^C \equiv 0$ and $f^D \equiv 1$.

Proof. By substitution of (1.2.8) into (1.2.10)-(1.2.11), we have

$$C_1^L e^{m\alpha} + C_2^L e^{(1-m)\alpha} - c = v_L^\alpha \tag{1.7.1}$$

$$C_1^L e^{m\beta} + C_2^L e^{(1-m)\beta} - c = v_L^\beta \tag{1.7.2}$$

$$mC_1^L e^{m\beta} + (1-m)C_2^L e^{(1-m)\beta} = 0, \tag{1.7.3}$$

where $v_L^\alpha, v_L^\beta \in [0,1]$ are arbitrary but fixed constants.[24] Solving for the constants C_1^L, C_2^L in (1.7.1) and (1.7.2) yields

$$C_1^L = M_L \left[(v_L^\alpha + c)e^{\alpha(m-1)} - (v_L^\beta + c)e^{\beta(m-1)} \right], \tag{1.7.4}$$

$$C_2^L = M_L \left[-(v_L^\alpha + c)e^{\alpha(m-1)+(2m-1)\beta} + (v_L^\beta + c)e^{\beta(m-1)+(2m-1)\alpha} \right], \tag{1.7.5}$$

where $M_L = (e^{(2m-1)\alpha} - e^{(2m-1)\beta})^{-1}$. Substituting these into (1.7.3) and simplifying yields

$$J(x; v_L^\alpha, v_L^\beta) := m[(v_L^\alpha + c) - (v_L^\beta + c)e^{(1-m)x}] + (m-1)[(v_L^\alpha + c) - (v_L^\beta + c)e^{mx}] = 0, \tag{1.7.6}$$

where $x = \alpha - \beta$. Since J is strictly decreasing in x (and hence strictly increasing in β) it has a unique and positive solution, which we denote x^*. This yields $\beta^*(\alpha; v_L^\alpha, v_L^\beta) = \alpha - x^*$.[25] Note also that J is increasing in v_L^α and decreasing in v_L^β, and thus by implicit differentiation, $\beta^*(\alpha; v_L^\alpha, v_L^\beta)$ is decreasing in v_L^α and increasing in v_L^β for fixed α. □

Since most of our analysis is concerned with $v_L^\alpha = 1$ and $v_L^\beta = 0$, we record for future convenience the specialized form of (1.7.6):

$$J(x) := m[(1+c) - ce^{(1-m)x}] + (m-1)[(1+c) - ce^{mx}] = 0. \tag{1.7.7}$$

Lemma 1.7.2. *There is an increasing function* $\alpha^* : \mathbb{R} \to \mathbb{R}$ *with* $\frac{\partial a^*}{\partial \beta} \in (0,1)$ *such*

[24] Later it will be shown that in the unique baseline equilibrium, $v_L^\alpha = V_L(\alpha) = 1$ and $v_L^\beta = V_L(\beta) = 0$, but the generality will be useful for later analysis.

[25] When convenient, we suppress dependence on v_L^α and v_L^β, and then it is assumed that $v_L^\alpha = 1$ and $v_L^\beta = 0$ unless otherwise specified.

that $\alpha^*(\beta)$ *is the unique solution* $\alpha \geq \beta$ *to* (1.2.17), (1.2.18), (1.2.19) *and* (1.2.20). *Moreover, the induced function* $V_B(z) := \frac{C_1 e^{mz} + C_2^B e^{(1-m)z}}{1+e^z}$ *is strictly increasing for* $z \geq \beta$.

Proof. By substitution of (1.2.17) into (1.2.18), (1.2.19) and (1.2.20), we get

$$\frac{C_1^B e^{m\alpha} + C_2^B e^{(1-m)\alpha}}{1+e^\alpha} = \frac{e^\alpha}{1+e^\alpha} - k \quad (1.7.8)$$

$$\frac{(m(1+e^\beta) - e^\beta)C_1^B e^{m\beta} + ((1-m)(1+e^\beta) - e^\beta)C_2^B e^{(1-m)\beta}}{(1+e^\beta)^2} = \frac{C_1^B e^{m\beta} + C_2^B e^{(1-m)\beta}}{(1+e^\beta)^2} \quad (1.7.9)$$

$$\frac{(m(1+e^\alpha) - e^\alpha)C_1^B e^{m\alpha} + ((1-m)(1+e^\alpha) - e^\alpha)C_2^B e^{(1-m)\alpha}}{(1+e^\alpha)^2} = \frac{e^\alpha}{(1+e^\alpha)^2}. \quad (1.7.10)$$

Combining (1.7.8) and (1.7.9) yields

$$C_1 = Mme^{(1-m)\beta} \quad (1.7.11)$$

$$C_2 = M(m-1)e^{m\beta}, \quad (1.7.12)$$

where $M = \frac{e^\alpha - k(1+e^\alpha)}{(m-1)e^{m\beta+(1-m)\alpha} + me^{m\alpha+(1-m)\beta}}$. By combining (1.7.8)-(1.7.10), or equivalently (1.7.10)-(1.7.12), and simplifying, we obtain the condition

$$0 = K_1 e^{m\alpha} + K_2 e^{(m-1)\alpha} + K_3 e^{(1-m)\alpha} + K_4 e^{-m\alpha} =: K(\alpha;\beta). \quad (1.7.13)$$

where $K_1 = -m(m-1)(1-k)e^{(1-m)\beta} < 0$, $K_2 = km^2 e^{(1-m)\beta} > 0$, $K_3 = m(m-1)(1-k)e^{m\beta} > 0$ and $K_4 = -k(1-m)^2 e^{m\beta} < 0$. The function $K(\alpha;\beta)$ is a Dirichlet polynomial with two sign changes in the coefficients, and thus has either 0 or 2 real roots. Moreover, it is easy to check that $K(\beta;\beta) > 0$ and that $K(\alpha;\beta) \to -\infty$ as $\alpha \to \pm\infty$. It follows that there is a unique root denoted $\alpha^*(\beta) > \beta$. Note also that by straightforward calculation, $K(z_m;\beta) > 0$, and thus $\alpha^*(\beta) > z_m$, where $z_m = \ln \frac{k}{1-k}$ is the static threshold.

For the second claim, we use implicit differentiation. From the above arguments, it must be that $\frac{\partial K(\alpha;\beta)}{\partial \alpha}|_{\alpha=\alpha^*(\beta)} < 0$. Towards the other claims, differentiate $K(\alpha;\beta)$ with respect to α and then subtract m times (1.7.13) to obtain

$$\frac{\partial K(\alpha;\beta)}{\partial \alpha}|_{\alpha=\alpha^*(\beta)} = -D_1 - D_2 + D_3,$$

where we denote $D_1 := m(m-1)(2m-1)(1-k)e^{(1-m)\alpha+m\beta}$, $D_2 := km^2e^{(m-1)(\alpha-\beta)}$ and $D_3 := 2mk(m-1)^2e^{-m(\alpha-\beta)}$.

Next, differentiate with respect to β and add $(m-1)$ times (1.7.13) to obtain

$$\frac{\partial K(\alpha;\beta)}{\partial \beta}|_{\alpha=\alpha^*(\beta)} = (2m-1)(m-1)e^{-m(\alpha-\beta)}[m(1-k)e^{\alpha}-(m-1)k]|_{\alpha=\alpha^*(\beta)}. \quad (1.7.14)$$

Since $\alpha^*(\beta) > \frac{k}{1-k}$ and $m > 1$, the expression above is strictly positive. It follows that

$$\frac{\partial a^*(\beta)}{\partial \beta} = -\left.\frac{\frac{\partial K(\alpha;\beta)}{\partial \beta}}{\frac{\partial K(\alpha;\beta)}{\partial \alpha}}\right|_{\alpha=\alpha^*(\beta)} > 0.$$

For the other bound, we have

$$\begin{aligned}\frac{\partial a^*(\beta)}{\partial \beta} &= -\left.\frac{\frac{\partial K(\alpha;\beta)}{\partial \beta}}{\frac{\partial K(\alpha;\beta)}{\partial \alpha}}\right|_{\alpha=\alpha^*(\beta)} \\ &= \frac{D_1 - (2m-1)(m-1)^2ke^{-m(\alpha-\beta)}}{D_1 + D_2 - D_3}|_{\alpha=\alpha^*(\beta)}.\end{aligned}$$

Note that the term D_1 appears in both the numerator and denominator, and in the denominator we have $D_2 - D_3 > 0$; this establishes that the fraction above is less than 1.

For the last claim, differentiate and rearrange terms to see that $V_B'(z)$ is proportional to $m(m-1)(e^{(1-m)\beta+(m+1)z} - e^{m\beta+(2-m)z}) + m^2e^{(1-m)\beta+mz} - (m-1)^2e^{m\beta+(1-m)z}$. Since $m > 1$, it is easy to check that this is positive for $z \geq \beta$. □

Corollary 1.7.1. *There exists a unique pair (β^*, α^*) such that $\alpha^* = \alpha^*(\beta^*)$ and*

$\beta^* = \beta^*(\alpha^*)$.

Proof. The inverse function $[\beta^*]^{-1}(\beta) = \beta + x^*$ is linear with slope 1, while $\alpha^*(\beta)$ has slope strictly less than 1, so these functions can intersect at most once. Since $\alpha^*(\beta) > z_m$, for sufficiently low β, we have $\alpha^*(\beta) > [\beta^*]^{-1}(\beta)$. On the other hand, by direct computation,

$$\begin{aligned} 2K([\beta^*]^{-1}(\beta); \beta) = K(\beta + x^*; \beta) & \\ &= -m(m-1)(1-k)e^{(m-1)x^*} + km^2 e^{(m-1)x^* - (\beta + x^*)} \\ &\quad + m(m-1)(1-k)e^{-mx^*} - k(1-m)^2 e^{-mx^* - (\beta + x^*)} \\ &\to m(m-1)(1-k)[e^{-mx^*} - e^{(m-1)x^*}] < 0, \end{aligned}$$

where the limit is taken as $\beta \to \infty$. Hence for large β, $K(\beta; \beta) > 0 = K(\alpha^*(\beta); \beta) > K([\beta^*]^{-1}(\beta); \beta)$ and thus $\alpha^*(\beta) < [\beta^*]^{-1}(\beta)$. By the Intermediate Value Theorem, there exists a unique β^* where the graphs intersect: $\alpha^*(\beta^*) = [\beta^*]^{-1}(\beta^*)$. Defining $\alpha^* := \alpha^*(\beta^*)$, we have the unique pair (β^*, α^*) as desired. □

Lemma 1.7.3. *For each α, there exist unique C_1^H, C_2^H such that V_H solves* (1.2.7)*,* (1.2.9) *and* (1.2.12) *for $\beta = \beta^*(\alpha)$.*

Proof. Combining (1.2.7) with (1.2.9) and (1.2.12) yields

$$C_1^H e^{(m-1)\alpha} + C_2^H e^{-m\alpha} = 1 + c$$

$$(m-1)C_1^H e^{(m-1)\beta} - mC_2^H e^{-m\beta} = 0.$$

The solutions are $C_1^H = m(1+c)[me^{(m-1)\alpha} + (m-1)e^{(m-1)\alpha - (2m-1)x^*}]^{-1}$, $C_2^H = \frac{m-1}{m} e^{(2m-1)(\alpha - x^*)} C_1^H$. □

Corollary 1.7.2. *For all $\alpha > 0$, $V_H(\beta^*(\alpha)) > 0$.*

Proof. Evaluating (1.2.7) using the solutions above and $\beta^*(\alpha) = \alpha - x^*$, we have

$$V_H(\beta^*(\alpha)) \propto -mce^{(m-1)x^*} + (2m-1)(1+c) - (m-1)ce^{-mx^*}.$$

By subtracting $0 = (2m-1)(1+c) - mce^{(1-m)x^*} - (m-1)ce^{mx^*}$ from (1.7.7), we get $V_H(\beta^*(\alpha)) \propto P(x^*) := \frac{e^{mx^*}-e^{-mx^*}}{m} - \frac{e^{(m-1)x^*}-e^{-(m-1)x^*}}{m-1}$. It is easy to verify that $P(x)$ has a triple root at $x = 0$ and is strictly increasing, so in particular $P(x^*) > 0$ and the claim follows. □

1.7.3 Proof of Theorem 1.7.1

Consider any equilibrium, Ξ.

For any starting belief $Z_\emptyset$, define $\mathcal{Z}_\emptyset := \{z : \mathbb{P}^H_{Z_\emptyset}(\inf\{t \geq 0 : Z_t = z\} < \infty) > 0\}$, that is $\mathcal{Z}_\emptyset$ is the set of states that the process Z reaches with positive probability under $\mathbb{P}^H_{Z_\emptyset}$ (and hence also under $\mathbb{P}^L_{Z_\emptyset}$. We say that $\mathcal{Z}_\emptyset$ is the set of *reachable* states from $Z_\emptyset$. For any reachable state z, we write $\mathbb{P}^L_z$ for the law of Z continuing from $Z_{t_0} = z$. Since Z is a time-homogeneous Markov process, without loss of generality we assume $t_0 = 0$.[26]

Lemma 1.7.4. *There exists $\alpha \in \mathcal{Z}_\emptyset$ such that $V_B(\alpha) = p(\alpha) - k$.*[27]

Proof. Suppose by way of contradiction that $V_B(z) > p(z) - k$ for all $z \in \mathcal{Z}_\emptyset$. We derive a contradiction from this by showing that the buyer's strategy must be $\rho \equiv \infty$, which yields $V_B(Z_0) = 0$. Suppose $\rho \in \mathcal{B}^0$ and $\mathbb{P}_{Z_0}(\rho < \infty) > 0$. We have

$$V_B(Z_0) = \mathbb{E}_{Z_0}\left[e^{-r\rho} \underbrace{(p(Z_\rho) - k)}_{\text{expected payoff at } \rho} \underbrace{(p(Z_0) + (1-p(Z_0))(1-S^{L,0}_\rho)}_{\text{probability of no exit by } \rho}\right]. \quad (1.7.15)$$

Now let $\widehat{\rho} \in \mathcal{B}^\rho$ be arbitrary, and note that $\widehat{\rho} \geq \rho$ almost surely and $\widehat{\rho} > \rho$ with positive probability, since otherwise immediate exit is supported at time ρ, which

[26] Strictly speaking, this is abuse of notation because for a fixed starting belief, $Z_0 = Z_\emptyset$ deterministically. However, this convention reduces subscripts and simplifies the exposition.

[27] That is, the buyer's value at α is the value of immediate adoption.

yields payoff 0. Since $\rho \wedge \widehat{\rho}$ is feasible, we have

$$V_B(Z_0) = \mathbb{E}_{Z_0}\left[e^{-r\rho}\underbrace{\mathbb{1}\{\rho < \widehat{\rho}\}V_B(Z_\rho) + \mathbb{1}\{\rho = \widehat{\rho}\}(p(Z_\rho) - k)}_{\text{expected payoff at }\rho}\underbrace{(p(Z_0) + (1 - p(Z_0))(1 - S^{L,0}_\rho)}_{\text{probability of no exit by }\rho}\right]. \tag{1.7.16}$$

By the assumption, the expression in (1.7.16) is strictly greater than that in (1.7.15), a contradiction; thus we have shown the existence of α such that $V_B(\alpha) = p(\alpha) - k$. □

Observe that the set of α satisfying Lemma 1.7.4 is bounded below by the static threshold $z_m = \ln \frac{k}{1-k}$, and define $\alpha_0 := \inf\{\alpha \in \mathcal{Z}_\emptyset : V_B(\alpha) = p(\alpha) - k\}$.

Lemma 1.7.5. *There exists a reachable state $\beta < \alpha_0$ such that $V_L(\beta) = 0$, and starting from $Z_0 = \beta$, $Z_t \geq Z_0$ for all $t \geq 0$, $\mathbb{P}^L_{Z_0}$-almost surely.*

Proof. We divide the proof into two cases. First, suppose there exists $z \in \mathcal{Z}_\emptyset$ such that by setting $Z_0 = z$, Z is discontinuous in t at $t = 0$ with positive probability under $P^L_{Z_0}$. In this case, by almost-sure continuity of $\widehat{Z}$, it must be that Q is discontinuous in t at $t = 0$ with positive probability. Since $Q_0 = 0$ and $Q_{0+} = \lim_{t\downarrow 0} Q_t = \lim_{t\downarrow 0} \ln\left(\frac{1}{1-S^{L,0}_{t-}}\right) = \ln\left(\frac{1}{1-S^{L,0}_0}\right)$, it must be that $S^{L,0}_0 > 0$ with positive probability. Now by Blumenthal's 0-1 Law (Durrett, 2010, Theorem 8.2.3), $\mathbb{P}^L_{Z_0}(S^{L,0}_0 > 0) = 1$ and thus $S^{L,0}_0$ is a deterministic, strictly positive quantity, corresponding to an atom of immediate exit by L. By definition, immediate exit is in the support of L's strategy: $\tau_0 \in \mathcal{S}^{L,0}$, where $\tau_0 \equiv 0$. It follows that $V_L(Z_0) = 0$. Moreover, $V_L(Z_{0+}) > 0$, and since Z can only have upward jumps, the lower bound property follows.

In the remaining case, we have that for all $z \in \mathcal{Z}_\emptyset$, whenever $Z_{t_0} = z$, Z is $\mathbb{P}^L_{Z_{t_0}}$-almost surely continuous at t_0. We consider two subcases: (a) $\mathcal{Z}_\emptyset$ is unbounded below, and (b) $\mathcal{Z}_\emptyset$ is bounded below.

Case (a): We first establish that there exists $\beta < \alpha_0$ with $V_L(\beta) = 0$. Suppose by way of contradiction that $V_L(z) > 0$ for all reachable $z < \alpha_0$, and consider any reachable $Z_0 < \alpha_0$. Define $\tau_{\alpha_0} := \inf\{t \geq 0 : Z_t \geq \alpha_0\}$. Note that $B^0_{\tau_{\alpha_0}-} = 0$, otherwise there would exist $\alpha' \in (\beta, \alpha_0)$ such that $V_B(\alpha') = p(\alpha') - k$, contradicting the definition of α_0 as the infimum. By continuity of Z, $\tau_{\alpha_0} > 0$ almost surely. We show that $S^{L,0}_{\tau_{\alpha_0}-} = 0$ almost surely. Suppose otherwise, and let $\tau_L \in \mathcal{S}^{L,0}$ be arbitrary. Type L's value is then[28]

$$\begin{aligned} V_L(Z_0) =& \mathbb{P}^L_{Z_0}(\tau_L \leq \tau_{\alpha_0})\mathbb{E}^L_{Z_0}\left[\int_0^{\tau_L} -cre^{-rt}dt | \tau_L \leq \tau_{\alpha_0}\right] \\ &+ \mathbb{P}^L_{Z_0}(\tau_{\alpha_0} < \tau_L)\mathbb{E}^L_{Z_0}\left[\int_0^{\tau_{\alpha_0}} -cre^{-rt}dt + e^{-r\tau_{\alpha_0}}V_L(Z_{\tau_{\alpha_0}}) | \tau_{\alpha_0} < \tau_L\right]. \quad (1.7.17) \end{aligned}$$

Next, take any $\widehat{\tau}_L \in \mathcal{S}^{L,\tau_L}$, and note that $\widehat{\tau}_L \geq \tau_L$ by definition. Moreover, $\widehat{\tau}_L > \tau_L$ with probability $P^L_{Z_0} > 0$, otherwise $V_L(z) = 0$ whenever τ_L is finite and $Z_{\tau_L} = z$, contradicting the counterfactual. Since $\widehat{\tau}_L$ is feasible from $t = 0$, we must have

$$\begin{aligned} V_L(Z_0) \geq& \mathbb{P}^L_{Z_0}(\tau_L < \widehat{\tau}_L \leq \tau_{\alpha_0})\mathbb{E}^L_{Z_0}\left[\int_0^{\tau_L} -cre^{-rt}dt + e^{-r\tau_L}V_L(Z_{\tau_L}) | \tau_L < \widehat{\tau}_L \leq \tau_{\alpha_0}\right] \\ &+ \mathbb{P}^L_{Z_0}(\tau_L = \widehat{\tau}_L \leq \tau_{\alpha_0})\mathbb{E}^L_{Z_0}\left[\int_0^{\tau_L} -cre^{-rt}dt | \tau_L = \widehat{\tau}_L \leq \tau_{\alpha_0}\right] \\ &+ \mathbb{P}^L_{Z_0}(\tau_{\alpha_0} < \tau_L)\mathbb{E}^L_{Z_0}\left[\int_0^{\tau_{\alpha_0}} -cre^{-rt}dt + e^{-r\tau_{\alpha_0}}V_L(Z_{\tau_{\alpha_0}}) | \tau_{\alpha_0} < \tau_L\right]. \quad (1.7.18) \end{aligned}$$

The expansion in (1.7.18) exceeds that of (1.7.17) by

$$\mathbb{P}^L_{Z_0}(\tau_L < \widehat{\tau}_L \leq \tau_{\alpha_0})\mathbb{E}^L_{Z_0}[e^{-r\tau_L}V_L(Z_{\tau_L}) | \tau_L < \widehat{\tau}_L \leq \tau_{\alpha_0}],$$

which is strictly positive under the counterfactual. This is a contradiction, and

[28] We adopt the convention that $Z_\infty = 0$, though this is without loss of generality as the discount factor vanishes.

therefore for all $\tau_L \in \mathcal{S}^{L,0}$, $\tau_L > \tau_{\alpha_0}$, $\mathbb{P}^L_{Z_0}$-almost surely. It follows that P_{Z_0}-almost surely, $S^{L,0}_{\tau_{\alpha_0}} = 0$ and hence for all $t \in [0, \tau_{\alpha_0}]$, $Z_t = \widehat{Z}_t$. Now since $V_L(\cdot)$ is bounded above by 1, we have for all $Z_0 \in Z^m$, $V_L(Z_0) \le \mathbb{E}^L_{Z_0}[-c(1-e^{-r\tau_{\alpha_0}}) + e^{-r\tau_{\alpha_0}}]$. By standard properties of Brownian motion, as $Z_0 \to -\infty$, this upper bound tends to $-c < 0$, contradicting optimality since L can guarantee 0 by exit at time 0, and thus there exists $\beta < \alpha_0$ such that $V_L(\beta) = 0$. If $\mathcal{Z}_\emptyset$ is unbounded below, it contains points below β. For any $Z_0 < \beta$, by continuity of Z, it must be that $V_L(Z_0) < 0$, but this is also a contradiction. Thus we have ruled out the case that Z is continuous and $\mathcal{Z}_\emptyset$ is unbounded below.

The remaining case is (b), that $\mathcal{Z}_\emptyset$ is bounded below. Let $\beta := \inf\{z \in \mathcal{Z}_\emptyset\}$. We first show that $\beta < \alpha_0$. If not, then $\inf_{z\ge\beta} V_B(z) = p(\beta) - k$. We show that there exists $\epsilon > 0$ and starting belief $Z_0 \in [\beta, \beta+\epsilon)$ such that the buyer can profitably deviate by playing $\rho_\epsilon := \inf\{t \ge 0 : Z_t \ge \beta + \epsilon\}$. The buyer's value from playing ρ_ϵ is

$$v(Z_0;\rho_\epsilon) = p(Z_0)\mathbb{E}^H_{Z_0}\left[e^{-r\rho_\epsilon}(1-k)\right] + (1-p(Z_0))\mathbb{E}^L_{Z_0}\left[-e^{-r\rho_\epsilon}k(1-S^{L,0}_{\rho_\epsilon})\right].$$

Take any $Z_0(\epsilon) \in (\beta, \beta+\epsilon^2)$,[29] and let $\epsilon \to 0$. Then $\mathbb{P}^L_{Z_0(\epsilon)}(S^{L,0}_{\rho_\epsilon} > 0) \to 1$, so that there exists $\delta < 1$ such that

$$\begin{aligned}\lim_{\epsilon\to 0} v(Z_0(\epsilon);\rho_\epsilon) &= p(\beta)(1-k) - (1-p(\beta))k\delta \\ &> p(\beta) - k,\end{aligned}$$

a contradiction. Thus we conclude that $\beta < \alpha_0$.

By way of contradiction, suppose now that $V_L(z) > 0$ for all $z < \alpha_0$. Next, consider any $Z_0 \in (\beta, \alpha_0)$, and By the same argument as in case (a), under the counterfactual assumption that $V_L(z) > 0$ for all reachable z, we have $S^{L,0}_t = 0$,

[29] This is just to ensure that $Z_0(\epsilon) \to \beta$ sufficiently fast relative to ϵ.

$\mathbb{P}^L_{Z_0}$-a.s. for all $t < \tau(\beta, \alpha_0) := \inf\{s \geq 0 : \widehat{Z}_s \notin (\beta, \alpha_0)\}$. Thus $Z_t = \widehat{Z}_t$ for all $t < \tau_{(\beta,\alpha_0)} \inf\{s \geq 0 : \widehat{Z}_s \notin (\beta, \alpha_0)\}$. Since $\widehat{Z}_t$ reaches β with positive probability, $\beta \in \mathcal{Z}_\emptyset$, and it must be that β is a reflecting barrier for Z. This implies that Q_t is almost surely strictly increasing at $t = 0$, and immediate exit at $Z_0 = \beta$ is supported (i.e., $\tau_0 \in \mathcal{S}^{L,0}$, where $\tau_0 \equiv 0$), and $V_L(\beta) = 0$, a contradiction, so there must exist some β such that $V_L(\beta) = 0$.

Given that there exists β with $V_L(\beta) = 0$, the same arguments as in the conclusion of case (a) can be repeated. Suppose there exists a reachable state $z < \beta$. By continuity and the fact that there is no adoption below α_0 and hence below β, we have $V_L(Z) < 0$, a contradiction. Thus $\beta = \min\{z \in \mathcal{Z}_\emptyset\}$ and therefore satisfies the second property of the lemma. □

Let β_0 denote the supremum of the (nonempty) set of β satisfying the properties of Lemma 1.7.5.

Lemma 1.7.6. *In any stationary equilibrium, $\beta_0 < \alpha_0$ and β_0 itself satisfies the properties of Lemma 1.7.5.*

Proof. For the first claim, suppose instead that $\beta_0 = \alpha_0$. Starting from $Z_0 = \alpha_0$, the belief process never goes below α_0. A contradiction then follows by the same reasoning as in case (b) of Lemma 1.7.5 that showed that $\alpha_0 > \inf\{z \in Z_\emptyset\}$.

For the second claim, by the same reasoning as in case (b) above, no exit occurs in (β_0, α_0), which implies β_0 is both reachable and a reflecting barrier, so immediate exit is supported at β_0, and hence $V_L(\beta_0) = 0$. □

Lemma 1.7.7. *The set $\{\beta \in \mathcal{Z}_\emptyset : V_L(\beta) = 0\}$ is bounded above.*

Proof. Suppose not, and take any sequence of reachable states $\beta_n \to \infty$. We claim that for any associated sequence $\rho_n \in \mathcal{B}^0$ for each $Z_0 = \beta_n$, $\mathbb{E}_{\beta_n} e^{-r\rho_n} \to 1$. If not, then there exists $\epsilon > 0$ and a subsequence (which we relabel for convenience) $\beta_n \to \infty$

with $\mathbb{E}_{\beta_n} e^{-r\rho_n} < 1 - \epsilon$ for all $n \in \mathbb{N}$. We have $V(\beta_n, \rho_n) \leq (1-\epsilon)p(\beta_n)(1-k)$,[30] but buying immediately yields $p(\beta_n) - k$. For sufficiently small ϵ, buying immediately is a strict improvement. Given that $\mathbb{E}_{\beta_n} e^{-r\rho_n} \to 1$ for arbitrary ρ_n, it follows that for sufficiently large n, L can earn $V_L(\beta_n) > 0$ by never exiting. □

Given Lemma 1.7.7, let $\beta_1 := \sup\{\beta \in \mathcal{Z}_\emptyset : V_L(\beta) = 0\}$. Since $V_L(\beta_1) = 0$, β_1 is the maximum of this set.[31]

Lemma 1.7.8. *For any finite $z \in \mathbb{R}$, there exists a reachable state $\alpha \geq z$ such that for $Z_0 = \alpha$, $B_0^0 = 1$, $\mathbb{P}_{Z_0}$-almost surely.*

Proof. First, suppose there exists $\alpha \geq z$ such that $B_0^0 > 0$ with positive probability (and hence almost surely). We argue that this implies $B_0^0 = 1$. Note that if $B_0^0 > 0$, immediate exit cannot occur; that is, $\tau := \inf\{t \geq 0 : S_t^{L,0} > 0\} > 0$, $P_{Z_0}^L$ almost surely. It follows that $Z_t = \widehat{Z}_t$ for $t \in [0, \tau)$, and by standard properties, $\widehat{Z}_t = 0$ for uncountably many $t \in [0, \tau)$, almost surely. Since the buyer's strategy is stationary, this implies that for any $\delta < 1$, $\inf\{t \geq 0 : B_t^0 > \delta\} = 0$ almost surely, and by right-continuity, $B_0^0 > \delta$ almost surely. It follows that $B_0^0 = 1$, as desired.

If there does not exist $\alpha \geq z$ such that for $Z_0 = \alpha$, $B_0^0 > 0$ with positive probability, we show that a contradiction arises. In this case, $B_t^0 = 1 - \exp\left(\int_0^t -f_2^C(Z_t)dt\right) < 1$ for all $t < \rho_z := \inf\{t \geq 0 : Z_t \leq z\}$ and hence for any $\rho \in \mathcal{B}^0$, $\widehat{\rho} := \rho_z \vee \rho$ is also in $\mathcal{B}^0$. As $Z_0 \to \infty$, the value $V_B(Z_0, \widehat{\rho}) \to 0$, while for sufficiently high Z_0, it is strictly profitable to buy immediately, a contradiction. □

Let α_1 denote the infimum of the set of $\alpha > \beta_1$ satisfying Lemma 1.7.8. Since starting from $Z_0 = \alpha_1$, $\inf\{t \geq 0 : Z_t > \alpha_1\} = 0$ almost surely, α_1 itself satisfies Lemma 1.7.8, so it is also the minimum of this set.

[30]This notation is for the value starting at β_n and playing ρ_n.

[31]To see this, note that either β_1 is a reflecting barrier, or Z goes below β_1 immediately; either case requires that $V_L(\beta_1) = 0$.

Lemma 1.7.9. *In any stationary equilibrium, $\alpha_1 = \alpha_0 = \alpha^*$ and $\beta_1 = \beta_0 = \beta^*$.*

Proof. We prove the following inequalities:

(i) $\beta_0 \geq \beta^*(\alpha_0)$,

(ii) $\alpha_0 = \alpha^*(\beta_0)$,

(iii) $\beta_1 = \beta^*(\alpha_1)$,

(iv) $\alpha_1 \leq \alpha^*(\beta_1)$.

(i) $\beta_0 \geq \beta^*(\alpha_0)$: Fix any feasible $V_L(\alpha_0) \in [0,1]$. For $z \in (\beta_0, \alpha_0)$, V_L satisfies (1.2.7) subject to boundary conditions $V_L(\beta_0) = 0$ and $V_L(\alpha_0)$. Clearly, $V_L(\beta_0+) \geq 0$, otherwise $V_L(z) < 0$ for $z \in (\beta_0, \alpha_0)$ sufficiently close to β_0. By (Harrison, 2013, Chapter 6), a necessary boundary condition at a reflecting barrier with regulation cost $c = 0$ is $V_L'(\beta_0) = c = 0$. From Lemma 1.7.1, the solution to these conditions is $\beta_0 = \beta^*(\alpha_0; v_L) \geq \beta^*(\alpha_0) := \beta^*(\alpha_0; 1)$.

(ii) $\alpha_0 = \alpha^*(\beta_0)$: As (β_0, α_0) is a delay region, V must satisfy (1.2.17) for $z \in (\beta_0, \alpha_0)$. By definition, we have the condition $V_B(\alpha_0) = p(\alpha_0) - k$, which is (1.7.8). Since β_0 is a killing boundary with exponential rate $\kappa = (1 - p(\beta_0)) = \frac{1}{1+e^{\beta}}$ (see Harrison (2013), pp. 159-161), V must satisfy a *Robin* boundary condition that involves both the value and slope at β_0:

$$V_B'(\beta_0) = \kappa V(\beta_0) = \frac{V(\beta_0)}{1 + e^{\beta}}. \tag{1.7.19}$$

This equation is (1.2.19), and its closed form is (1.7.9).

Finally, to see that (1.2.20), and hence (1.7.10), is necessary, observe that $V_B(z) \geq p(z) - k$ for all $z \in (\beta_0, \alpha_0)$, and hence $V_B'(\alpha_0-) \leq \frac{d}{d\alpha}(p(\alpha) - k)|_{\alpha=\alpha_0} = \frac{e_0^{\alpha}}{(1+e_0^{\alpha})^2}$. If $V_B(\alpha_0-) < p(\alpha_0) - k$, then for $Z_0 = \alpha_0$ and sufficiently small $\epsilon > 0$, the buyer can strictly improve by deviating to $\rho := \inf\{t \geq 0 : Z_t \geq \alpha_0 + \epsilon\}$. We have now established that V satisfies (1.7.8)-(1.7.10), and thus by Lemma 1.7.2, $\alpha_0 = \alpha^*(\beta_0)$.

(iii) $\beta_1 = \beta^*(\alpha_1)$: By definition, $V_L(\beta_1) = 0$ giving (1.7.2) for $\beta = \beta_1$ and clearly $V_L(\alpha_1) = 1$, giving (1.7.1) for $v_L = 1$ and $\alpha = \alpha_1$. By familiar arguments, $V_L'(\beta_1+) \geq 0$. Now suppose $V_L'(\beta_1+) > 0$, and set $Z_0 = \beta_1$. Let $\tau = \inf\{t \geq 0 : Z_t \notin (\beta_1 - \epsilon, \beta_1 + \epsilon)$. We have $V_L(\beta_1) \geq \mathbb{E}^L_{\beta_1}[e^{-r\tau}V_L(Z_\tau)]$. Since $V_L \geq 0$, this lower bound is strictly positive for sufficiently small ϵ, a contradiction. Thus $V_L'(\beta_1+) = 0$, giving us (1.7.3). By Lemma 1.7.1, $\beta_1 = \beta^*(\alpha_1)$.

(iv) $\alpha_1 \leq \alpha^*(\beta_1)$: By definition, $V_B(\alpha_1) = p(\alpha_1) - k$, so α_1 satisfies (1.7.8), and as above, $V_B'(\alpha_1-) = \frac{d}{d\alpha}(p(\alpha) - k)|_{\alpha=\alpha_1}$ is necessary, giving (1.7.10). For any $z \in [\beta_1, \alpha_1]$, the buyer can be no better off than if β_1 were a reflecting barrier, and thus $p(\alpha_1) - k \leq p(\alpha^*(\beta_1)) - k$, which implies $\alpha_1 \leq \alpha^*(\beta_1)$, as desired.

We now combine (i)-(iv) to get the desired result. From (i) and (ii), we have $\beta_0 \geq \beta^*(\alpha^*(\beta_0)$ which implies $\beta^0 \geq \beta^*$. In addition, we have $\alpha_0 \geq \alpha^*(\beta^*(\alpha_0)$, so that $\alpha_0 \geq \alpha^*$. Similarly, (iii) and (iv) imply that $\beta_1 \leq \beta^*$ and $\alpha_1 \leq \alpha^*$ Since $\beta_1 \geq \beta_0$ and $\alpha_1 \geq \alpha_0$ by definition, both sets of equalities must hold. □

Lemma 1.7.10. *For all $z \geq \alpha^*$, $V_B(z) = p(z) - k$. The buyer's strategy is the one defined by $\Xi(\beta^*, \alpha^*)$.*

Proof. Consider any $Z_0 > \alpha^*$ and $\alpha' > \alpha^*$ satisfying Lemma 1.7.8, and let $\rho \in \mathcal{B}^0$ be arbitrary. Clearly, $\rho < \infty$ almost surely since adoption is immediate at α^* and α'. Suppose that $\rho > 0$ with positive probability (and hence almost surely). Since L does not exit above α^* (because $\beta_1 = \beta^* < \alpha^*$), $\mathbb{P}_{Z_0}$-a.s. we have $Z_t = \widehat{Z}_t$ for all $t \leq \rho$. Since the process $[0,1]$-space process $p_t = p(\widehat{Z}_t)$ is a martingale for $t \leq \tau_{(\alpha_2,\alpha_1)}$ and $0 < \tau_{(\alpha_2,\alpha_1)} < \infty$ almost surely under $\mathbb{P}_{Z_0}$, the Optional Sampling Theorem

(Harrison, 2013, Corollary A.4) implies that

$$\begin{aligned} V_B(Z_0) &= \mathbb{E}_{Z_0}[e^{-r\rho}(p(Z_\rho) - k)] \\ &< \mathbb{E}_{Z_0}[(p(Z_\rho) - k)] \\ &= p(Z_0) - k, \end{aligned}$$

contradicting optimality of ρ. We therefore conclude that $\rho = 0$ almost surely for all $Z_0 \geq \alpha^*$. Since α^* is the lowest state at which adoption occurs, the buyer's strategy is that of $\Xi(\beta^*, \alpha^*)$. □

Lemma 1.7.11. *For all reachable $Z_0 \leq \beta^*$, $V_L(Z_0) = 0$ and $Z_{0+} = \beta^*$. The seller's strategy is the one defined by $\Xi(\beta^*, \alpha^*)$.*

Proof. Now from any $Z_0 \leq \beta^*$, consider $\tau_{\beta^*} := \inf\{t \geq 0 : Z_t \geq \beta^*\}$. We claim that $\tau_{\beta^*} = \inf\{t \geq 0 : Z_t = \beta^*\}$ holds $\mathbb{P}^L_{Z_0}$-a.s.; if not, then there would exist some $Z_0' \leq \beta^*$ such that $Z_{0+}' > \beta^*$, which would require that $\tau = 0$ be in the support for L from Z_0'; in turn, this would imply that $V_L(Z_0') = V_L(Z_{0+}') = 0$, but since $Z_{0+}' > \beta^*$, $V_L(Z_{0+}') > 0$, a contradiction. In addition, it must be that $\tau_{\beta^*} = 0$ almost surely. Otherwise by the familiar 0-1 argument we have $\tau_{\beta^*} > 0$ almost surely, and since $c > 0$ and $V_L(\beta^*) = 0$, $V_L(Z_0) = \mathbb{E}^L_{Z_0}\left[-c(1 - e^{-r\tau_{\beta^*}}) + e^{-r\tau_L} V_L(\beta^*)\right] < 0$, a contradiction. Thus $Z_{0+} = \beta^*$, which through (1.2.2) implies there is immediate exit with probability $e^{Z_0 - \beta^*}$, in accordance with $\Xi(\beta^*, \alpha^*)$.

For reachable $Z_0 > \beta^*$, the belief process evolves continuously at all times almost surely, and conditional on no exit must reflect off β^*; this requires that $Z_t = \widehat{Z}_t + L_t$, where L_t is of the form given in Definition 1.2.3; together with (1.2.2) this implies that L's strategy is that of Definition 1.2.3 for all $Z_0 > \beta^*$. □

Lemma 1.7.12 (Verification). $\Xi(\beta^*, \alpha^*)$ *is a stationary, monotone equilibrium.*

Proof. Condition 3 follows by construction; given the the seller strategies of $\Xi(\beta^*, \alpha^*)$, the final term of (1.2.2) is $\ln\left(\frac{1-0}{1-(1-\exp(-(L_{t_1}-L_{t_0}))}\right) = L_{t_1} - L_{t_0}$. Condition 4 is immediate from the definition of L_t and the fact that $\widehat{Z}$ is a time-homogeneous, Markov process. The buyer's strategy can be expressed by $\mathcal{C} = (-\infty, \alpha)$, $\mathcal{D} = [\alpha, \infty)$, $f_C(z) = 0$ for all $z \in \mathcal{C}$, and $f_D(z) = 1$ for all $z \in \mathcal{D}$, which satisfies property 5.

Next we verify the optimality of strategies. For $\theta = H, L$, it is immediate that strategies are optimal for $Z_0 \geq \alpha^*$, since they earn their maximum feasible payoffs 1. Now consider type H and $Z_0 < \alpha^*$. Let $\tau_\infty \equiv \infty$ denote H's strategy under $\Xi(\beta^*, \alpha^*)$ as a stopping time, and let $\tau_{\alpha^*} := \inf\{t \geq 0 : Z_t \geq \alpha^*\}$. From any $Z_0 = z < \alpha^*$, we have $V_H(z) = V_H(\tau_\infty; z)$, and it is easy to check that $V_H(z) = V_H(\tau; z)$ for all τ such that $\mathbb{P}^H_z(\tau \geq \tau_{\alpha^*}) = 1$. Now suppose there exists $\tau' \leq \tau_{\alpha^*}$ with probability $\mathbb{P}^H_z > 0$, such that $V_H(\tau'; z) > V_H(z)$. We have

$$\begin{aligned} V_H(\tau'; z) =& \mathbb{P}^H_z(\tau' \leq \tau_{\alpha^*})\mathbb{E}^H_z\left[\int_0^{\tau'} -cre^{-rt}dt + 0|\tau' \leq \tau_{\alpha^*}\right] \\ &+\mathbb{P}^H_z(\tau' > \tau_{\alpha^*})\mathbb{E}^H_z\left[\int_0^{\tau_{\alpha^*}} -cre^{-rt}dt + e^{-r\tau'}1|\tau' > \tau_{\alpha^*}\right]. \end{aligned}$$

By deviating to τ_∞, H replaces the 0 from exit with the continuation value $V_H(Z_{\tau'}) > 0$, and since $\mathbb{P}^H_z(\tau' \leq \tau_{\alpha^*}) > 0$, this deviation is strictly profitable.

By a similar argument, we show that there is no profitable deviation for L for $Z_0 < \alpha^*$. Again, we have $V_L(z) = V_L(\tau_\infty; z) = V_L(\tau_{(\beta^*,\alpha^*)}; z)$ where $\tau_{(\beta^*,\alpha^*)} := \inf\{t \geq 0 : Z_t \notin (\beta^*, \alpha^*)\}$. Let τ' be any strategy such that $V_L(\tau'; z) \geq V_L(z)$. If $\mathbb{P}^L_z(\tau' \leq \tau_{\alpha^*}, Z_{\tau'} > \beta^*) > 0$, then along such paths $V_L(Z_{\tau'}) > 0$ and L can strictly improve by deviating to τ_∞. Since β^* is a reflecting barrier, it follows that $\mathbb{P}^L_z(\tau' \leq \tau_{\alpha^*}, Z_{\tau'} > \beta^*) = 0$, and since $V_L(\beta^*) = V_L(\tau_\infty, \beta^*) = 0$, $V_L(\tau'; z) = V_L(\tau_\infty, z)$. Thus there is no profitable deviation.

For the buyer, we have that $V_B(p) \geq p - k$ and $V_B(p) > p - k$ for all $p < p(\alpha^*)$.

Thus

$$
\begin{aligned}
e^{-rt}(p_t - k) \leq e^{-rt}V_B(p_t) & \\
&= V_B(p_0) + \int_0^t e^{-rs}(\phi(\mathbb{1}\{\theta = H\} - p_s) + 1)\Phi_s V_B'(p_s)dp_s \\
&+ \int_0^t e^{-rs}\Phi_s(V_B'(p_s) - (1-p_s)V_B(p_s))\mathbb{1}\{p_s = p(\beta^*)\}dL_s \\
&+ \int_0^t e^{-rs}(-rV_B(p_s) + \frac{1}{2}\Phi_s^2 V_B''(p_s))ds.
\end{aligned}
$$

Note that (i) $\int_0^\tau e^{-rs}(\phi(\mathbb{1}\{\theta = H\} - p_s) + 1)\Phi_s V_B'(p_s)dp_s$ is a martingale, (ii) L_s is only increasing when $p_s = p(\beta^*)$ and thus when $(V_B'(p_s) - (1-p_s)V_B(p_s)$ vanishes by (1.2.19), and (iii) $-rV_B(p_s) + \frac{1}{2}\Phi_s^2 V_B''(p_s) \leq 0$ for all p_s, with strict inequality iff $p(s) > p(\alpha^*)$. By taking expectations at an arbitrary stopping time τ, it follows that $\mathbb{E}_{p_0}[e^{-r\tau}(p_\tau - k)] \leq V_B(p_0)$. Since $\tau = \tau_{\alpha^*}$ attains equality, τ_{α^*} is optimal. This completes the verification. □

1.8 Appendix B: Endogenous Quality

Proof of Lemma 1.3.1. By substituting $v_L'(\beta) = v_H'(\beta) = 0$ into (1.2.5) and (1.2.6) evaluated at β and then subtracting, we obtain $v_H(\beta) - v_L(\beta) = \frac{\phi^2}{2r}(v_H''(\beta) - v_L''(\beta)) > 0$. The assumption $v_H(\beta) > v_L(\beta)$ then implies the first result. It follows that for sufficiently small $\epsilon > 0$, for all $z \in (\beta, \beta + \epsilon)$, $v_H'(z) > v_L'(z)$ and thus $v_H(z) > v_L(z)$. □

Proof of Proposition 1.3.1. By Lemma 1.3.1, the definitions of K^*, K^{β^*}, z', z'' and z''' as given in Section 1.3.1 are valid. Note that L cannot strictly prefer to upgrade, otherwise $p_0 = 1$ and adoption would be immediate; since upgrades are not observed, L could profitably deviate by not upgrading. Thus $V_H(p_0) - K \leq V_L(p_0)$. If $V_H(p_0) - K < V_L(p_0)$, then there is no investment, and in this case $p_0 = p_-$. It follows there exists an equilibrium without upgrade if and only if one of the following holds (i) $K > K^*$, (ii) $K < K^{\beta^*}$ and $z_- \geq z'$, or (iii) $K \in (K^{\beta^*}, K^*)$ and either $z_- > z'''$ or $z_- < z''$.

If $p_0 > p_-$, L must be indifferent and thus it must be that z_0 solves (1.3.3). If $K > K^*$, there is no solution, and we have already characterized the unique equilibrium. If $K < K^{\beta^*}$, then we must have $z_0 = z'$ and $z_- < z'$. By Bayes' rule, $q = \frac{p'-p_-}{1-p_-}$. If $K \in (K^{\beta^*}, K^*)$, then $z_0 = z''$ and $z_0 = z'''$ each correspond to equilibria with q uniquely determined by Bayes' rule. □

1.8.1 Belief Updating

First, we provide a construction of a belief process under a given strategy profile. For $\theta = H, L$, let g_t^θ denote the density of $N(\mu_\theta t, \sigma^2 t)$. Given strategy I, for any intervals $S_1, S_2 \subseteq [t_0, \infty]$, define $I^{\theta,t_0}(S_1, S_2) := \int_{(s_1,s_2)\in S_1\times S_2} I^{\theta,t_0}(ds_1, ds_2)$. In consistency with the original notation, if $S_1 = [t_0, t]$ for some $t \in [t_0, \infty]$, we write t in place of S_1, and likewise for S_2.

Beliefs at the beginning of t_1, conditional on no exit prior to t_1, must be computed from beliefs at t_0 according to Bayes' rule.

$$p_{t_1}^- = \frac{A+B}{A+B+C}, \tag{1.8.1}$$

where we define locally

$$A := p_{t_0}^- g_{t_1-t_0}^H (X_{t_1} - X_{t_0}) I^{H,t_0}([t_1,\infty],\infty),$$

$$B := (1-p_{t_0}^-)\int_{t_0}^{t_1} g_{t_1-s}^H (X_{t_1} - X_s) g_{s-t_0}^L (X_s - X_{t_0}) I^{L,t_0}([t_1,\infty], ds-),$$

$$C := (1-p_{t_0}^-) g_{t_1-t_0}^L (X_{t_1} - X_{t_0}) I^{L,t_0}([t_1,\infty]^2).$$

Given a beginning-of-period belief $p_{t_1}^-$ at time t_1, we have mid-period belief p_{t_1}

satisfying

$$p_{t_1} = \frac{p_{t_1}^- I^{H,t_1}((t_1,\infty],\infty) + (1-p_{t_1}^-)I^{L,t_1}((t_1,\infty],t_1)}{p_{t_1}^- I^{H,t_1}((t_1,\infty],\infty) + (1-p_{t_1}^-)I^{L,t_1}((t_1,\infty],\infty)} \tag{1.8.2}$$

$$= \frac{A'+B'}{A'+B'+C'}, \qquad \text{where}$$

$$A' := p_{t_0}^- g_{t_1-t_0}^H (X_{t_1}-X_{t_0}) I^{H,t_0}((t_1,\infty],\infty),$$

$$B' := (1-p_{t_0}^-)\int_{t_0}^{t_1} g_{t_1-s}^H(X_{t_1}-X_s) g_{s-t_0}^L(X_s-X_{t_0}) I^{L,t_0}((t_1,\infty],ds),$$

$$C' := (1-p_{t_0}^-) g_{t_1-t_0}^L (X_{t_1}-X_{t_0}) I^{L,t_0}((t_1,\infty]^2).$$

1.8.2 Uniqueness of Belief Processes

Here we prove by construction that there exist unique solutions to the equations that recursively define belief processes in Definitions 1.3.3 and ...

Lemma 1.8.1. *There exists a unique solution (Z^-, Z, Q) to the system of equations defining beliefs in Definition 1.3.3.*

Proof. We provide a procedure for constructing a well-defined solution candidate and argue that (i) this candidate is indeed a solution, and (ii) it is the only solution. where $T:\mathbb{N}_0 \to [0,\infty]$, $J:\mathbb{N}_0\to[0,\infty]$ and $N:[0,\infty)\to\mathbb{N}_0$ are defined by the following procedure:

Initialization step: $T(0)=0,\ J(0)=0,\ N(0)=0.$

Steps $i=1,2,\ldots$:

$$T(i) = \inf\{t \geq T(i-1) : \widehat{Z}_t + \sum_{j=0}^{i-1} J(j) \leq \beta\}$$

$$J(i) = \mathbb{1}\{T(i)<\infty\}(z^* - (\widehat{Z}_{T(i)} + \sum_{j=0}^{i-1} J(j)))$$

$$N(t) = \sup\{n \geq 0 : T(n) \leq t\}.$$

Given these functions, we set

$$
\begin{aligned}
Q_t &= \sum_{i=0}^{N(t)} J(i) \\
Z_t &= \widehat{Z}_t + Q_t \\
Z_t^- &= Z_{t-} \\
&= \widehat{Z}_t + Q_{t-}.
\end{aligned}
$$

□

Lemma 1.8.2. *There exists a unique solution (Z^-, Z, Q, L) to the system of equations defining beliefs in Definition 1.3.4.*

Proof. As before, we define a procedure.

Initialization step:

$$
\begin{aligned}
T(0) &= 0 \\
J(0) &= 0 \\
N(0) &= 0 \\
\widehat{Z}_t^0 &= \widehat{Z}_t \\
\widetilde{Z}_t^0 &= \widehat{Z}_t^0 + \delta \ell_t^{\widetilde{Z}^0}(\hat{z})
\end{aligned}
$$

Steps $i = 1, 2, \ldots$:

$$
\begin{aligned}
T(i) &= \inf\{t \geq T(i-1) : \widetilde{Z}_t^{i-1} \leq \beta\} \\
J(i) &= \mathbb{1}\{T(i) < \infty\}(z^* - \widetilde{Z}_{T(i)}^{i-1}) \\
\widehat{Z}_t^i &= \widehat{Z}_t^0 + (z^* - \widehat{Z}_{T(i)}^0) \\
\widehat{Z}_t^{i,\Delta} &= \widehat{Z}_{t+T(i)}^i \\
\widetilde{Z}_t^{i,\Delta} &= \widehat{Z}_t^{i,\Delta} + \delta \ell_t^{\widetilde{Z}^{i,\Delta}}(\hat{z}) \\
\widetilde{Z}_t^i &= \widetilde{Z}_{t-T(i)}^{i,\Delta} \qquad \text{for all } t \geq T(i) \\
N(t) &= \sup\{n \geq 0 : T(n) \leq t\}.
\end{aligned}
$$

Given the above, we set

$$
\begin{aligned}
Q_t &= \sum_{i=0}^{N(t)} J(i) \\
Z_t &= \widetilde{Z}_t^{N(t)} \\
L_t &= \begin{cases} \delta \ell_t^{\widetilde{Z}^0}(\hat{z}) & \text{if } N(t) = 0 \\ \delta \left(\sum_{i=1}^{N(t)} \ell_{T(i)-T(i-1)}^{\widetilde{Z}^{i-1,\Delta}}(\hat{z}) + \ell_{t-T(N(t))}^{\widetilde{Z}^{N(t)},\Delta}(\hat{z}) \right) & \text{otherwise} \end{cases} \\
Z_t^- &= Z_{t-}
\end{aligned}
$$

□

1.8.3 Proof of Theorem 1.3.1

The proof is divided into steps as outlined in Section 1.3.3.

1. We prove that for any α, there exists at most one pair (β, z^*) that solves (1.3.14)-(1.3.19).

a. By translation invariance, (β, z^*, α), where $\beta < z^* < \alpha$, is a solution to

(1.3.14)-(1.3.19) if and only if $(\beta - z^*, 0, \alpha - z^*)$ is solution. Without loss of generality, we set $z^* = 0$ and solve for a pair (β, α) with $\beta < 0 < \alpha$. Introducing a parameter $x = V_L(0)$, we replace (1.3.14)-(1.3.17) with five equations:

$$C_1^H e^{(m-1)\beta} + C_2^H e^{-m\beta} = x + c + K \tag{1.8.3}$$

$$C_1^L e^{m\beta} + C_2^L e^{(1-m)\beta} = x + c \tag{1.8.4}$$

$$C_1^H + C_2^H = x + c + K \tag{1.8.5}$$

$$C_1^L + C_2^L = x + K \tag{1.8.6}$$

$$(m-1)C_1^H - mC_2^H = mC_1^L + (1-m)C_2^L. \tag{1.8.7}$$

b. Define $\mathcal{X} := ((m-1)K - c, \infty)$. We state and prove the following lemma.

Lemma 1.8.3. *There is a function $\beta : \mathcal{X} \to (-\infty, 0)$ such that for all x in its domain, $\beta(x)$ is the unique value of x solving* (1.8.3)-(1.8.7)*. Moreover, β is differentiable and strictly increasing.*

Proof. The system of equations (1.8.3)-(1.8.6) is linear in the four constants C_1^H, C_2^H, C_1^L and C_2^L. The unique solution is

$$C_1^H = \frac{1 - e^{m\beta}}{1 - e^{(2m-1)\beta}}(x + c + K) \tag{1.8.8}$$

$$C_2^H = \frac{e^{m\beta} - e^{(2m-1)\beta}}{1 - e^{(2m-1)\beta}}(x + c + K) \tag{1.8.9}$$

$$C_1^L = \frac{1 - e^{(m-1)\beta}}{1 - e^{(2m-1)\beta}}(x + c) \tag{1.8.10}$$

$$C_2^L = \frac{e^{(m-1)\beta} - e^{(2m-1)\beta}}{1 - e^{(2m-1)\beta}}(x + c). \tag{1.8.11}$$

Substituting these into (1.8.7) and simplifying yields

$$\begin{aligned} F(\beta,x) :=& -e^{m\beta}(x+c+mK)+e^{\beta}(2m-1)(x+c+K) \\ &-(2m-1)(x+c)+e^{(1-m)\beta}(x+c-(m-1)K) \\ &=0. \end{aligned} \tag{1.8.12}$$

The function F is a Dirichlet polynomial in β of length 4, so it has at most 3 roots, counting multiplicity. It is easy to verify that for all x, (i) $F(0,x)=0$, (ii) $\frac{\partial}{\partial\beta}F(0,x)=0$ and (iii) $\lim_{\beta\to+\infty}F(\beta,x)=-\infty$. Moreover, for all x,

$$\begin{aligned} \frac{\partial^2}{(\partial\beta)^2}F(0,x) &= -2m^3+3m^2-m \\ &= -m(m-1)(2m-1)<0, \end{aligned}$$

where the inequality follows from the fact $m>1$. Hence F is concave in β at 0, which is a double root, and for $\beta<0$ sufficiently close to 0, $F(\beta,x)<0$. The last coefficient, $x+c-(m-1)K$ is positive if and only if $x>(m-1)K-c$. If it is positive, then $\lim_{\beta\to-\infty}=\infty$, and thus there is exactly one remaining root, strictly negative, which we denote $\beta(x)$. Moreover, from these facts it follows that $F(\beta,x)$ crosses 0 from above at $\beta(x)$, that is, $\frac{\partial}{\partial\beta}F(\beta(x),x)<0$.

Differentiating F with respect to x for arbitrary β, we have

$$\frac{\partial}{\partial x}F(\beta,x)=-e^{m\beta}+e^{\beta}(2m-1)-(2m-1)+e^{(1-m)\beta}=:F_x(\beta).$$

We show that $F_x(\beta)>0$ for all $\beta<0$; by implicit differentiation, it then follows that $\beta(x)$ is strictly increasing. Observe that $\mathrm{sign}(F_x)=\mathrm{sign}(\widehat{F}_x)$, where

$$\begin{aligned} \widehat{F}_x(\beta) &:= e^{-\beta/2}F_x(\beta) \\ &= -e^{(m-1/2)\beta}+e^{\beta/2}(2m-1)-e^{-\beta/2}(2m-1)+e^{-(m-1/2)\beta}. \end{aligned}$$

It is easy to verify that $\widehat{F}_x(0)=0$; we now show that $\widehat{F}'_x(\beta)<0$ for all $\beta<0$, which implies that $\widehat{F}_x$ and thus F_x are positive for $\beta < 0$. By direct computation,

$$\widehat{F}_x'(\beta) \propto -e^{(m-1/2)\beta} + e^{\beta/2} + e^{-\beta/2} - e^{-(m-1/2)\beta}. \tag{1.8.13}$$

The right side above can be written as $G(\beta, 1/2) - G(\beta, m - 1/2)$, where $G(\beta, \xi) := e^{\xi\beta} + e^{-\xi\beta}$. For any fixed β, we have $\frac{\partial}{\partial \xi} G(\beta, \xi) = \beta(e^{\xi\beta} - e^{-\xi\beta})$, which is strictly positive when $\beta < 0$ and $\xi > 0$; hence $G(\beta, m-1/2) > G(\beta, 1/2)$ and thus $\widehat{F}_x'(\beta) < 0$, as desired.

We have shown that $\frac{\partial}{\partial \beta} F(\beta(x), x) < 0$ and $\frac{\partial}{\partial x} F(\beta(x), x) > 0$, so by implicit differentiation, $\beta(x)$ is strictly increasing. □

For each $x \in \mathcal{X}$, we have so far shown that there exists a unique vector $(C_1^H, C_2^H, C_1^L, C_2^L, \beta)$ of differentiable functions of x solving (1.8.3)-(1.8.7). We then define the candidate value functions $F_H(z) := C_1^H e^{(m-1)z} + C_2^H e^{-mz} - c$ and $F_L(z) := C_1^L e^{mz} + C_2^L e^{(1-m)z} - c$. Define $\Delta(z) := F_H(z) - F_L(z)$.

c. Relax equations (1.3.18) and (1.3.19) to

$$C_1^H e^{(m-1)\alpha} + C_2^H e^{-m\alpha} = y + c \tag{1.8.14}$$

$$C_1^L e^{m\alpha} + C_2^L e^{(1-m)\alpha} = y + c. \tag{1.8.15}$$

It remains to find a solution (x, α, y) solving the three remaining equations (1.8.14), (1.8.15) and $y = 1$.

d. We state and prove the following lemma.

Lemma 1.8.4. *For all $x \in \mathcal{X}$, a unique solution $(\alpha(x), y(x)$ to (1.8.3)-(1.8.7) and (1.8.14)-(1.8.15) exists.*

Proof. By subtraction, α solves (1.8.14) and (1.8.15) if and only if $z = \alpha$ is a solution to $\Delta(z) = 0$. By expansion and arranging terms in decreasing order of exponents, we have $\Delta(z) = -C_1^L e^{mz} + C_1^H e^{(m-1)z} - C_2^L e^{(1-m)z} + C_2^H e^{-mz}$. It is also useful to define the normalized difference function,

$$\Delta_K(z) := \Delta(z) - K = -C_1^L e^{mz} + C_1^H e^{(m-1)z} - K - C_2^L e^{(1-m)z} + C_2^H e^{-mz}.$$

Since the five constants K and the C_i^θ above are strictly positive, so $\Delta_K(z)$ is Dirichlet polynomial with sign pattern $-, +, -, -, +$ so $\Delta_K(z)$ has exactly 1 or 3 roots. By construction, there is a root at $z = \beta$ and a double root at $z = 0$. Moreover, $\Delta'_K(z) = -mC_1^L e^{mz} + (m-1)C_1^H e^{(m-1)z} + (m-1)C_2^L e^{(1-m)z} - mC_2^H e^{-mz}$ is a Dirichlet polynomial with sign pattern $-, +, +, -$ and hence has exactly 0 or 2 roots. By (1.8.7), $\Delta'_K(0) = -mC_1^L + (m-1)C_1^H + (1-m)C_2^H - mC_2^L = 0$, so $z = 0$ is one root. Since the coefficients of the end terms e^{mz} and e^{-mz} of both negative, $\lim_{z\to\pm\infty} \Delta'_K(z) = -\infty$. It follows that Δ'_K has a root $z' \in (\beta, 0)$ and that Δ_K is decreasing in $(-\infty, z')$, increasing in $(z', 0)$, and decreasing in $(0, \infty)$. Furthermore, $\lim_{z\to+\infty} \Delta_K(z) = -\infty$ so there exists a unique point α in $(0, \infty)$ such that $\Delta_K(z) = -K$ and thus $\Delta(z) = 0$. As noted above, α thus solves (1.8.14) and (1.8.15). □

e. For $i = 1, 2$ and $\theta = H, L$, we introduce the notation $C_{i,x}^\theta := \frac{dC_i^\theta(x)}{dx}$.[32] The following lemma states that the common slope of F_H and F_L at $z = 0$ is strictly increasing in x.

Lemma 1.8.5. *The function $(m-1)C_1^H - mC_2^H$ is strictly increasing in x.*

Proof. By differentiating (1.8.12) and using the chain rule, we obtain

$$(x + c + K)\beta'(x)P_1(\beta) + \beta'(x)KP_2(\beta) = P_3(\beta), \tag{1.8.16}$$

where $P_1(\beta) := me^{m\beta} - e^\beta(2m-1) + (m-1)e^{(1-m)\beta}$, $P_2(\beta) := m(m-1)(e^{m\beta} - e^{(1-m)\beta})$, and $P_3(\beta) := -e^{m\beta} + (2m-1)e^\beta - (2m-1) + e^{(1-m)\beta}$.

[32] A priori, the relationship between x and each constant is ambiguous; x influences the constants directly as well as through β, and these forces can act in opposite directions.

By differentiating (1.8.8) and (1.8.9), we have

$$
\begin{aligned}
(m-1)C^H_{1,x} - mC^H_{2,x} &= (m-1)\frac{\partial}{\partial x}C^H_1 - m\frac{\partial}{\partial x}C^H_2 \\
&\quad + (x+c+K)\beta'(x)[(m-1)\frac{\partial}{\partial \beta}C^H_1 - m\frac{\partial}{\partial \beta}C^H_2] \\
&= \frac{me^{(2m-1)\beta}) - (2m-1)e^{m\beta} + (m-1)}{1-e^{(2m-1)\beta}} \\
&\quad + (x+c+K)\beta'(x)\frac{(1-m)e^{(3m-1)\beta} + (2m-1)e^{(2m-1)\beta} - me^{m\beta}}{(1-e^{(2m-1)\beta})^2} \\
&= \frac{P_1(\beta)e^{(m-1)\beta}}{1-e^{(2m-1)\beta}} - (x+c+K)\beta'(x)\frac{(2m-1)e^{(2m-1)\beta}}{(1-e^{(2m-1)\beta})^2}P_4(\beta),
\end{aligned}
\tag{1.8.17}
$$

where $P_4(\beta) := (m-1)e^{m\beta} - (2m-1) + me^{(1-m)\beta}$. Now (1.8.12) can be arranged to obtain:

$$
x + c + K = \frac{KP_4(\beta)}{P_3(\beta)}. \tag{1.8.18}
$$

Substitute(1.8.18) into (1.8.16) and solve for β' to obtain

$$
\beta'(x) = \frac{(P_3(\beta))^2}{KP_1(\beta)P_4(\beta) + KP_2(\beta)P_3(\beta)}. \tag{1.8.19}
$$

Multiplying (1.8.18) by (1.8.19) yields

$$
(x+c+K)\beta'(x) = \frac{P_3(\beta)P_4(\beta)}{P_1(\beta)P_4(\beta) + P_2(\beta)P_3(\beta)}. \tag{1.8.20}
$$

Finally, substitute (1.8.20) into (1.8.17) to obtain

$$
(m-1)C^H_{1,x} - mC^H_{2,x} = \frac{P_1(\beta)e^{(m-1)\beta}}{1-e^{(2m-1)\beta}} - \frac{(2m-1)e^{(2m-1)\beta}}{(1-e^{(2m-1)\beta})^2}\frac{P_3(\beta)(P_4(\beta))^2}{P_1(\beta)P_4(\beta) + P_2(\beta)P_3(\beta)} \tag{1.8.21}
$$

We claim that $P_1(\beta)P_4(\beta) + P_2(\beta)P_3(\beta)$ is strictly positive for $\beta < 0$. To see this, expand to obtain

$$e^{(1-m)\beta}(2m-1)[(m-1)^2e^{2m\beta}-m^2e^{(2m-1)\beta}+(4m-2)e^{m\beta}-m^2e^{\beta}+(m-1)^2]. \quad (1.8.22)$$

Let $P_5(\beta)$ denote the term in square brackets. Then $P_5(\beta)$ is a Dirichlet polynomial with four sign changes, so it has at most four real roots. Moreover, one can verify that $P_5(0) = P_5^{(1)}(0) = P_5^{(2)}(0) = P_5^{(3)}(0)$, where we use superscript (n) to denote the n^{th} derivative, and thus $\beta = 0$ is a root of multiplicity 4, and there are no other roots. Since the final term is $(m-1)^2 > 0$, $\lim_{\beta\to\infty} P_5(\beta) = \infty$, and thus $P_5(\beta) > 0$ for all $\beta < 0$. The other two factors in the expansion are positive, so $P_1(\beta)P_4(\beta) + P_2(\beta)P_3(\beta)$ is positive for $\beta < 0$.

By multiplying through by the negative factor $-(2m-1)^{-1}e^{-2m\beta}(1-e^{(2m-1)\beta})^2(P_1(\beta)P_4(\beta)+P_2(\beta)P_3(\beta))$, $(m-1)C^H_{1,x} - mC^H_{2,x}$ has the opposite sign as the polynomial

$$\begin{aligned} P_6(\beta) := {} & e^{(3m-1)\beta}(m-1)^3 - e^{(3m-2)\beta}m^3 + e^{(2m-1)\beta}3(2m-1)^2 - e^{m\beta}3m^2(3-2m) \\ & + (e^{(m-1)\beta}3(m-1)^2(3m-1) + e^{\beta}(2m-1)^3) \\ & - (e^{-\beta}(2m-1)^3 + e^{(1-m)\beta}3(m-1)^2(3m-1)) \\ & + e^{-m\beta}3m^2(3m-2) - e^{(1-2m)\beta}3(2m-1)^2 + e^{\beta(2-3m)}m^3 - e^{(1-3m)\beta}(m-1)^3. \end{aligned}$$

The polynomial P_6 has 9 sign changes, and by direct computation, we have $P_6(0) = P_6^{(1)}(0) = \cdots = P_6^{(8)}(0)$, so that $\beta = 0$ is a root of multiplicity 9. As the final coefficient is $-(m-1)^3 < 0$, it must be that $\lim_{\beta\to-\infty} P_6(\beta) = -\infty$, and thus $P_6(\beta) < 0$ for all $\beta < 0$. Since $(m-1)C^H_{1,x} - mC^H_{2,x}$ has the opposite sign, we have shown that $(m-1)C^H_1 - mC^H_2$ is increasing in x. □

Lemma 1.8.6. *The function $\alpha(x)$ is differentiable, decreasing and strictly positive, and $y(x)$ is differentiable and strictly increasing.*

Proof. By differentiation with respect to x,

$$
\begin{aligned}
y &= C_1^H e^{(m-1)\alpha} + C_2^H e^{-m\alpha} = C_1^L e^{m\alpha} + C_2^L e^{(1-m)\alpha} \\
\implies y' &= C_{1,x}^H e^{(m-1)\alpha} + C_{2,x}^H e^{-m\alpha} \\
&\quad + \alpha' \left[(m-1)C_1^H e^{(m-1)\alpha} - mC_2^H e^{-m\alpha}\right] \qquad \text{and} \qquad (1.8.23) \\
y' &= C_{1,x}^L e^{m\alpha} + C_{2,x}^L e^{(1-m)\alpha} \\
&\quad + \alpha' \left[mC_1^L e^{m\alpha} + (1-m)C_2^L e^{(1-m)\alpha}\right],
\end{aligned}
$$

where again we have suppressed dependence of y and α on x. Solving for α' yields

$$\alpha'(x) = \frac{C_{1,x}^H e^{(m-1)\alpha} + C_{2,x}^H e^{-m\alpha} - C_{1,x}^L e^{m\alpha} - C_{2,x}^L e^{(1-m)\alpha}}{mC_1^L e^{m\alpha} + (1-m)C_2^L e^{(1-m)\alpha} - (m-1)C_1^H e^{(m-1)\alpha} + mC_2^H e^{-m\alpha}}. \quad (1.8.24)$$

Since $F_L(z)$ intersects $F_H(z)$ from below at $z = \alpha$, we have $(m-1)C_1^H e^{(m-1)\alpha} - mC_2^H e^{-m\alpha} < mC_1^L e^{m\alpha} + (1-m)C_2^L e^{(1-m)\alpha}$, and so the denominator in (1.8.24) is strictly positive. Next, let $N(\alpha)$ denote the numerator of we show that the numerator in (1.8.24); we show that $N(\alpha)$ is negative, which implies $\alpha'(x) < 0$. Note that differentiating (1.8.5)-(1.8.7) yields

$$C_{1,x}^H + C_{2,x}^H = 1 \quad (1.8.25)$$

$$C_{1,x}^L + C_{2,x}^L = 1 \quad (1.8.26)$$

$$(m-1)C_{1,x}^H - mC_{2,x}^H = mC_{1,x}^L + (1-m)C_{2,x}^L. \quad (1.8.27)$$

Combining (1.8.25)-(1.8.27) yields $C_{1,x}^L = C_{1,x}^H - \frac{1}{2m-1}$. We can rewrite $N(\alpha)$ in terms of $C_{1,x}^H$ as follows:

$$N(\alpha) = C_{1,x}^H(e^{(m-1)\alpha} - e^{-m\alpha}) + e^{-m\alpha} - (C_{1,x}^H - \frac{1}{2m-1})(e^{m\alpha} - e^{(1-m)\alpha}) - e^{(1-m)\alpha}. \quad (1.8.28)$$

Now by substituting (1.8.25) into the inequality $(m-1)C_{1,x}^H > m(C_{2,x}^H)$, we get

$C^H_{1,x} > \frac{m}{2m-1}$. Note that the right side of (1.8.28) is decreasing in $C^H_{1,x}$, so we have

$$N(\alpha) < \frac{m}{2m-1}(e^{(m-1)\alpha} - e^{-m\alpha}) + e^{-m\alpha} - \frac{m-1}{2m-1}(e^{m\alpha} - e^{(1-m)\alpha}) - e^{(1-m)\alpha}$$
$$\propto -(m-1)e^{m\alpha} + me^{(m-1)\alpha} - me^{(1-m)\alpha} + (m-1)e^{-m\alpha} =: N_2(\alpha).$$

It is easily checked that $N_2(0) = N_2'(0) = N_2''(0) = 0$, so that N_2 has a triple root at $\alpha = 0$. By familiar arguments, it follows that $N_2(\alpha) < 0$ for all $\alpha > 0$. Hence we have shown that $N(\alpha) < 0$, so $\alpha'(x) < 0$.

Substituting the expression (1.8.24) into (1.8.23), we get

$$y'(x) = N^H(\alpha) + Q^H(\alpha)\frac{N^H(\alpha) - N^L(\alpha)}{Q^L(\alpha) - Q^H(\alpha)}$$
$$= \frac{N^H(\alpha)Q^L(\alpha) - N^L(\alpha)Q^H(\alpha)}{Q^L(\alpha) - Q^H(\alpha)},$$

where $N^H(\alpha) := C^H_{1,x}e^{(m-1)\alpha} + C^H_{2,x}e^{-m\alpha}$, $N^L(\alpha) := C^L_{1,x}e^{m\alpha} + C^L_{2,x}e^{(1-m)\alpha}$, $Q^H(\alpha) = (m-1)C^H_1e^{(m-1)\alpha} + mC^H_2e^{-m\alpha}$, and $Q^L(\alpha) = mC^L_1e^{m\alpha} + (1-m)C^L_2e^{(1-m)\alpha}$. Note that $N^H(\alpha) - N^L(\alpha) = N(\alpha)$, the numerator of (1.8.24) which is positive as shown above. Since $Q^L(\alpha) > Q_H(\alpha) > 0$, to show that $y'(x) > 0$, it suffices to show that $N^H(\alpha) > 0$. Now $N^H(0) = C^H_{1,x} + C^H_{2,x} = 0$, and $\frac{d}{d\alpha}N^H(\alpha) = (m-1)C^H_{1,x}e^{(m-1)\alpha} - mC^H_{2,x}e^{-m\alpha} > (m-1)C^H_{1,x}(e^{(m-1)\alpha} - e^{-m\alpha}) > 0$, where we have used the facts that $(m-1)C^H_{1,x} > mC^H_{2,x}$ and $C^H_{1,x} > \frac{m}{2m-1} > 0$. □

Lemma 1.8.7. *For sufficiently large x, $y(x) > 1$, while $\lim_{x\downarrow(m-1)K-c} y(x) < 1$ if and only if $K < K_0 : \frac{(1+c)(m-1)^{m-1}}{m^m}$.*

Proof. Since $y(x) > x + K$ by construction, we have $y(x) > 1$ for any $x > 1 - K$.

As $x \downarrow (m-1)K - c$, we get from (1.8.12) that $\beta(x) \to -\infty$. Accordingly, the following limits of (1.8.8)-(1.8.11) hold as $x \downarrow (m-1)K - c$:

$$C_1^H \to mK$$

$$C_2^H \to 0$$

$$C_1^L \to (m-1)K$$

$$C_2^L \to 0.$$

Equating (1.8.14) and (1.8.15) and using these values, we have

$$mke^{(m-1)\alpha^*} = (m-1)ke^{m\alpha^*}$$

$$\iff \alpha^* = \ln\left(\frac{m}{m-1}\right).$$

By continuity, $\lim_{x\downarrow(m-1)K-c} F_H(\alpha(x)) = F_H(\alpha^*) = \frac{m^m K}{(m-1)^{m-1}} - c$. Since $y(x)$ is increasing in x, it follows that $y(x) < 1$ for some $x \in \mathcal{X}$ if and only if $\frac{m^m K}{(m-1)^{m-1}} - c < 1$. $\square$

2. We show that there is a cutoff K^{**} such that V_L is globally nonnegative if and only if $K \leq K^{**}$.

Lemma 1.8.8. *For each $K \in (0, K_0)$, let $v_L(K) := \min_{z\in\mathbb{R}} F_L(z; K)$. Then $v_L(K)$ is strictly decreasing in K.*

Proof. First, consider candidates for equilibria under K and $K' > K$. For any z, $F_L(z; K') - F_L(z; K)$ is a polynomial of length two with at most one root. Translate $F_L(z; K)$ to $\widehat{F}_L(z)$ so that $\widehat{F}_L(\alpha(K); K') = F_L(\alpha(K); K) = 1$; thus $z = \alpha(K)$ is the unique point of intersection of $\widehat{F}_L(z; K')$ and $F_L(z; K)$, and it follows that either $\widehat{F}_L(z; K') > F_L(z; K)$ for all $z < \alpha(K)$ or $\widehat{F}_L(z : K') < F_L(z; K)$ for all $z < \alpha(K)$. We will show that the latter of these is true, and since $\widehat{F}_L(z; K')$ and $F_L(z; K')$ have the same minimum value, this will prove the lemma.

Note that for any distinct K and K', $\widehat{F}_L(z;K')$ and $F_L(z;K)$ cannot have the same minimum, otherwise there would exist two points of intersection. This implies that $v_L(K)$ is either strictly increasing or strictly decreasing. We show that it must be strictly decreasing.

By differentiating (1.8.12) with respect to K, we have

$$\frac{\partial}{\partial K}F(\beta,x;K) = -me^{m\beta} + e^{\beta}(2m-1) - (m-1)e^{(m-1)\beta}.$$

The polynomial above has at most two roots, and it is easily verified that $\beta = 0$ is a double root. Since $\lim_{\beta\to\pm\infty} \frac{\partial}{\partial K}F(\beta,x;K) = -\infty$, we have $\frac{\partial}{\partial K}F(\beta,x;K) < 0$ for all $\beta \neq 0$. Since $F(\beta,x;K)$ is decreasing in β at $\beta = \beta(x;K)$, we have that $\beta(x;K)$ is decreasing in K. Moreover, for sufficiently large K, $\beta(K)$ can be made arbitrarily small, and since $x(K) \leq 1$, this implies that $v_L < 0$ for sufficiently large K. For sufficiently small K, we have $x(K) \to 1-K$ and $\beta(K) \to 0$, which implies that $v_L \to 1-K > 0$. By continuity, there exists some threshold K^{**}, which is a function of other parameters, such that a resetting equilibrium exists if and only if $K \leq K^{**}$. □

3. We now solve the buyer's problem. The buyer's value function is of the form $V_B(z) = \frac{C_1^B e^{mz} + C_2^B e^{(1-m)z}}{1+e^z}$, as in the original model.

Lemma 1.8.9. *There is an increasing, differentiable function* $\alpha : \mathbb{R} \to \mathbb{R}$ *with* $\alpha' \in (0,1)$ *such that, given* β *and* $z^* = \beta + \delta$, $\alpha(\beta)$ *is the unique* $\alpha > \beta$ *that solves* (1.3.11)-(1.3.13).

Proof. Using the closed form above, these become

$$\frac{C_1^B e^{m\alpha} + C_2^B e^{(1-m)\alpha}}{1+e^{\alpha}} = \frac{e^{\alpha}}{1+e^{\alpha}} - k \tag{1.8.29}$$

$$\frac{C_1^B e^{m\beta} + C_2^B e^{(1-m)\beta}}{1+e^{\beta}} = \frac{C_1^B e^{mz^*} + C_2^B e^{(1-m)z^*}}{1+e^{z^*}} \tag{1.8.30}$$

$$\frac{e^{(m+1)\alpha}C_1(m-1) + e^{m\alpha}mC_1 - e^{(2-m)\alpha}mC_2 - e^{(1-m)\alpha}(m-1)C_2}{(1+e^{\alpha})^2} = \frac{e^{\alpha}}{(1+e^{\alpha})^2}. \tag{1.8.31}$$

Combining (1.8.29) and (1.8.30) gives the following values of the constants:

$$C_1 = \frac{e^{\alpha}(1-k)-k}{e^{m\alpha}+Re^{(1-m)\alpha}}$$

$$C_2 = RC_1,$$

where $R := \frac{(1+e^{\beta})e^{mz^*}-(1+e^{z^*})e^{m\beta}}{(1+e^{z^*})e^{(1-m)\beta}-(1+e^{\beta})e^{(1-m)z^*}} > 0$. Equation (1.8.31) is thus equivalent to

$$0 = P(\alpha, R) := A_1(\alpha) + A_2(\alpha) + A_3(\alpha) + R[A_4(\alpha) + A_5(\alpha) + A_6(\alpha)], \tag{1.8.32}$$

where $A_1(\alpha) := e^{(m+1)\alpha}(1-k)(m-1), A_2(\alpha) := e^{m\alpha}[-k(m-1)+(1-k)m-1], A_3(\alpha) := -e^{(m-1)\alpha}km, A_4(\alpha) := -e^{(2-m)\alpha}(1-k)m, A_5(\alpha) := e^{(1-m)\alpha}(2mk-m-k)$, and $A_6(\alpha) := e^{-m\alpha}k(m-1)$. Independent of the sign of A_2 and A_5, the function above is a polynomial with two sign changes and hence either two or zero roots. As $\alpha \to \pm\infty$, $P(\alpha, R) \to +\infty$, and it can be verified that at the static threshold, z_m, we have $P(z_m) < 0$. It follows that there are two distinct roots, one below and one above z_m, and since $P(\cdot, R)$ must cross 0 from below to satisfy second order conditions, we have a well-defined solution α to the system above.

Next, we show that $\alpha'(\beta) < 1$. By writing (1.8.32) as $0 = P(\alpha(\beta), R(\beta))$ and differentiating with respect to β, we have $\alpha'(\beta) = -R'(\beta)\frac{P_R}{P_\alpha}$. We show that (i) $\frac{R'(\beta)}{R(\beta)} \in (0, 2m-1)$ and (ii) $\frac{R \cdot P_R}{P_\alpha} > (2m-1)^{-1}$). For (i), we have that $R(\beta) > 0$ so it suffices to show that (a) $R'(\beta) > 0$ and that (b) $R'(\beta) < (2m-1)R(\beta)$. For (a), write

$R = \frac{X}{Y}$ so that $R' = \frac{X'Y - Y'X}{Y^2}$. By expansion, $X'Y - Y'X = (1+e^z)[e^{mz+(1-m)\beta}(m - 1 + me^\beta) - (2m-1)e^\beta(1+e^z) + e^{(1-m)z+m\beta}(m+(m-1)e^\beta)$. The expression in square brackets is a Dirichlet polynomial in z with sign pattern $+, -, -, +$ and has two or zero real roots. It is easy to verify that it vanishes at $z = \beta$, as does its first derivative w.r.t. z, and thus it is strictly positive for all $z > \beta$. This establishes (a). Next, note that $R' - (2m-1)R = Y^{-2}(X'Y - Y'X - (2m-1)XY)$. By expansion, $(X'Y - Y'X - (2m-1)XY = -e^{mz+(1-m)\beta}(1+e^z)(m+(m-1)e^\beta) + e^z(1+e^\beta)^2(2m-1) - e^{(1-m)z+m\beta}(1+e^z)(m-1+me^\beta)$, which has pattern (in z) $-, -, +, -, -$. Again, it is easily verified that this expression and its first derivative vanishes at $z = \beta$, which gives us (b); together (a) and (b) prove (i). For (ii), we have

$$P_R = A_4(\alpha) + A_5(\alpha) + A_6(\alpha)$$

$$P_\alpha = (m+1)A_1(\alpha) + mA_2(\alpha) + (m-1)A_3(\alpha) + R[(2-m)A_4(\alpha) + (1-m)A_5(\alpha) - mA_6(\alpha)].$$

By using (1.8.32) to eliminate R, we obtain

$$\begin{aligned}(2m-1)R \cdot P_R + P_\alpha &= A_1 - A_3 - (A_6 - A_4)\frac{A_1 + A_2 + A_3}{A_4 + A_5 + A_6} \\ &= -P_R^{-1}(1+e^\alpha)^2 k(1-k)(2m-1).\end{aligned}$$

Recall that $P_\alpha > 0$. If $P_R > 0$, then the first line is positive and the second line negative, a contradiction, so it must be that $P_R < 0$, and thus $(2m-1)R \cdot P_R + P_\alpha > 0$, which implies $\frac{R \cdot P_R}{P_\alpha} > -(2m-1)^{-1}$, as desired. □

Since we have already shown that β and z^* are linear in α with slope 1, the following corollary is immediate.

Corollary 1.8.1. *There exists a unique triple* (β^R, z^*, α^R) *such that* $z^* = \beta^R + \delta$, $\alpha^R = \alpha^*(\beta^*)$ *and* $\beta^* = \beta^*(\alpha^*)$.

Lemma 1.8.10. *The conditions* (1.3.5)-(1.3.13) *are necessary in resetting equilibrium.*

Proof. As noted previously, conditions (1.3.5), (1.3.6), (1.3.11) and (1.3.12) are a repeat of (1.2.9), (1.2.10), (1.2.18) and (1.2.20). The conditions (1.3.7), (1.3.8), and (1.3.13) are analogous to one another, so it suffices to consider just (1.3.7). Starting from $Z_0 = \beta$, by dominated convergence, the fact that V_H is continuous, and that Z is continuous at z^*, we have $V_H(\beta) = \mathbb{E}^H[e^{-rt}V_H(Z_{t\wedge\tau_\alpha}) \to V_H(Z_0+) = V_H(z^*)$.[33] Condition (1.3.9), is immediate since L must be indifferent or not at β and at z^*; investing yields $V_H(z^*) - K$ while not investing yields $V_L(z^*)$. Finally, for condition (1.3.10), suppose instead that $V_H'(z^*) > V_L'(z^*)$. Then for sufficiently small $\epsilon > 0$, $V_H(z) - V_L(z) > K$ for all $z \in (z^*, z^* + \epsilon)$, and L would strictly prefer to invest at such z, contradicting optimality. Likewise, if $V_H'(z^*) < V_L'(z^*)$, L would strictly prefer to invest in some interval $(z^* - \epsilon, z^*)$. □

Proof of Theorem 1.3.1. All that remains is to establish that $\mathcal{R}(\beta^R, z^*, \alpha)$ is indeed an equilibrium. Properties 3 and 4 are satisfied by construction. Clearly, H has no profitable deviation since $V_H(z) > V_L(z) \geq 0$ for all z. As shown in the proof of Lemma 1.8.4, $V_H(z) - V_L(z) \leq K$ for all z, and equality only holds below β^R and at z^*, so L's strategy is indeed optimal. Optimality for the buyer is established by a similar argument to the one given in the proof of Theorem 1.2.1. □

1.8.4 Skew-Resetting Equilibrium

Lemma 1.8.11. *Let α and $\hat{z} < \alpha$ be arbitrary, and for any $w \in \mathbb{R}$, let $V_H(z; w) = C_1^H(w)e^{(m-1)z} + C_2^H(w)e^{-mz}$, where $C_1^H(w), C_2^H(w)$ solve*

$$C_1^H(w)e^{(m-1)\alpha} + C_2^H(w)e^{-m\alpha} = 1 + c$$

$$C_1^H(w)e^{(m-1)\hat{z}} + C_2^H(w)e^{-m\hat{z}} = w + c.$$

Then for any $z \in [\hat{z}, \alpha)$,

- $\frac{\partial}{\partial w}V_H(z; w) \in (0, 1)$

[33] For more on value matching with respect to resetting barriers, see Dixit (1993), p. 26.

- $\lim_{w\to\pm\infty} V_H(z;w) = \pm\infty$
- $\frac{\partial^2}{\partial w \partial z} V_H(z;w) < 0$
- $\lim_{w\to\pm\infty} \frac{\partial}{\partial z} V_H(z;w) = \mp\infty$.

Proof. Rewrite the conditions as

$$C_1^H(w)e^{(m-1)\alpha} + C_2^H(w)e^{-m\alpha} = 1 + c$$
$$C_1^H(w)e^{(m-1)\hat{z}} + C_2^H(w)e^{-m\hat{z}} = w + c.$$

The constants are

$$C_1^H(w) = \frac{(1+c)e^{m\alpha} - (w+c)e^{m\hat{z}}}{e^{(2m-1)\alpha} - e^{(2m-1)\hat{z}}}$$
$$C_2^H(w) = \frac{(w+c)e^{m\hat{z}+(2m-1)\alpha} - (1+c)e^{m\alpha+(2m-1)\hat{z}}}{e^{(2m-1)\alpha} - e^{(2m-1)\hat{z}}}.$$

Now differentiate with respect to w the value at any fixed z, obtaining

$$\frac{\partial}{\partial w} V_H(z;w) = \frac{e^{m\hat{z}}\left[e^{(2m-1)\alpha - mz} - e^{(m-1)z}\right]}{e^{(2m-1)\alpha} - e^{(2m-1)\hat{z}}},$$

which is contained in $(0,1)$ for all $z \in (\hat{z}, \alpha)$. Since the above expression is independent of w, the derivative is bounded away from zero and thus for fixed z, $\lim_{w\to\pm\infty} V_H(z;w) = \pm\infty$.

By inspection, the expression above for $\frac{\partial}{\partial w} V_H(z;w)$ is strictly decreasing in z, so $\frac{\partial^2}{\partial w \partial z} V_H(z;w) < 0$. Since the expression is also independent of w, the claim about limits follows. □

Lemma 1.8.12. *The thresholds K^* and K^{**} satisfy $K^{**} < K^*$.*

Proof. For fixed z, $\frac{d}{dK}\frac{\partial}{\partial z} V_H(z;w(K)) = w'(K)\frac{\partial^2}{\partial w \partial z} V_H(z;w(K)) < 0$ since $w'(K) > 0$ and $\frac{\partial^2}{\partial w \partial z} V_H(z;w) < 0$. □

Now fix α and let $\hat{z}, C_1^L$ and C_2^L solve (1.3.25), (1.3.30) and (1.3.31), which are a restatement of (1.2.10)-(1.2.11) after replacing β with $\hat{z}$. As shown in Section 1.7, given any α, the solution is $C_1^L = M_L(m-1)e^{-m\hat{z}}$, $C_2^L = M_L m e^{(m-1)\hat{z}}$ and $M_L = (1+c)[(m-1)e^{mx^*} + me^{(1-m)x^*}]^{-1}$, and where x^* is the unique solution x to

$$0 = m[(1+c) - ce^{(1-m)x}] + (m-1)[1+c-ce^{mx}] = 0.$$

Let $V_L(z) = C_1^L e^{mz} + C_2^L e^{(1-m)z}$. For any w, let $V_H(z;w)$ be defined as above, and let $\Delta(z;w) := V_H(z;w) - V_L(z)$. Let $\Delta^*(w) := \max_{z\in[\hat{z},\alpha]} \Delta(z;w)$, and let $z^*(w) = \max_{z\in[\hat{z},\alpha]} \Delta(z;w)$. Now for each $w \in \mathbb{R}$, define $G_w : [\hat{z},\alpha] \to \mathbb{R}_+$ by $G_w(z) := \max\{0, \Delta(z;w)\}$. For each $z \in [\hat{z},\alpha]$, as $w \downarrow -\infty$, $G_w(z) \downarrow 0$. By Dini's Theorem, $G_w \to 0$ uniformly as $w \downarrow -\infty$, and thus for any $K > 0$, there exists w such that $G_w(z) < K$ for all $z \in [\hat{z},\alpha]$, and thus $\Delta^*(w) < K$. In addition, $\Delta^*(w) \geq \Delta(\hat{z};w) \to \infty$ as $w \to \infty$. By the envelope theorem, for all w, $\frac{\partial}{\partial w}\Delta^*(w) = \frac{\partial}{\partial w}\Delta(z;w)|_{z=z^*(w)} < 0$. It follows that there exists a unique $w(K)$ such that $\Delta^*(w(K)) = K$, and $w(K)$ is increasing.

If $K < K^*$, then $w(K) < w(K^*)$. By the definition of K^*, $w(K^*)$ is such that $\frac{\partial}{\partial z}V_H(z;w(K^*)) = 0$, and since $\frac{\partial^2}{\partial w \partial z}V_H(z;w) < 0$ for all w, $\frac{\partial}{\partial z}V_H(z;w(K)) > \frac{\partial}{\partial z}V_H(z;w(K^*)) = 0$.

Now given K, set $C_1^{H,+} = C_1^H(w(K))$ and $C_2^{H,+} = C_2^H(w(K))$.

Next, since $V_L(z)$ is strictly convex and $\lim_{z\to\pm\infty} V_L(z) = \infty$, there is a unique $\beta < \hat{z}$ such that $V_L(\beta) = V_L(z^*(w(K)))$. Continuing, there is a unique pair $(C_1^{H,-}, C_2^{H,-})$ that solves

$$C_1^{H,-}e^{(m-1)\beta} + C_2^{H,-}e^{-m\beta} = V_L(\beta) + K$$

$$C_1^{H,-}e^{(m-1)\hat{z}} + C_2^{H,-}e^{-m\hat{z}} = C_1^{H,+}e^{(m-1)\hat{z}} + C_2^{H,+}e^{-m\hat{z}}.$$

Lemma 1.8.13. *Assume that $K < K^*$. Then $V_{H,-}(\beta) < V_{H,+}(\beta)$ if and only if $K > K^{**}$.*

Proof. First, note that if $K \leq K^{**}$, then a resetting equilibrium exists, so $V_{H,-} \equiv V_{H,+}$, so the claim holds. Since resetting equilibrium does not exist for $K > K^{**}$, and in particular for $K \in (K^{**}, K^*)$, by continuity, it must be that either $V_{H,-}(\beta) < V_{H,+}(\beta)$ holds for all $K \in (K^{**}, K^*)$, or $V_{H,-}(\beta) > V_{H,+}(\beta)$ holds for all $K \in (K^{**}, K^*)$. Towards a contradiction, suppose that the latter of these is true. Note that at $K = K^*$, the baseline equilibrium $\Xi(\beta^*, \alpha^*)$ exists, and at $z = \beta^*$, we have $V_H'(\beta^*) = V_L'(\beta^*) = 0$. Let $\bar{z}$ be defined implicitly as the unique $z \neq \beta^*$ such that $V_H'(z) = V_L'(z)$ under the equilibrium $\Xi(\beta^*, \alpha^*)$ for $K = K^*$, and let $\bar{\beta}$ be the unique $z < \beta^*$ such that $V_L(z) = V_L(\bar{z})$. Due to the asymmetry in the drift terms of their ODEs, if $V_H(\bar{\beta}) \leq V_L(\bar{\beta}) + K^*$, then V_H would attain its minimum at a state to the left of that of V_L. It follows that $V_H(\bar{\beta}) > V_L(\bar{\beta}) + K^*$. By supposition and continuity, we have that $V_{H,-}(\beta) > V_{H,+}(\beta) \to V_H(\bar{\beta}) > V_L(\bar{\beta}) + K^*$, and thus $V_{H,-}(\beta) > V_L(\beta) + K$ for K sufficiently close to K^*, contradicting the definition of $V_{H,-}$. □

Lemma 1.8.14. *The conditions* (1.3.24)-(1.3.37) *are necessary in skew-resetting equilibrium.*

Proof. Since Z is continuous at $\hat{z}$ and immediately reaches points below and points above $\hat{z}$, value functions must be continuous at $\hat{z}$. Conditions (1.3.32) and (1.3.36) then follow by construction. All other conditions are a repeat of conditions from resetting equilibrium, adapted to the piecewise definitions where appropriate, with the following exceptions: (1.3.22), (1.3.23), (1.3.30), (1.3.31) and (1.3.35). Condition (1.3.30) is necessary for L to be indifferent to exit. Now Z^L is $(2\gamma-1)$-elastic, γ-skew Brownian motion[34] as in Appuhamillage et al. (2011). By Theorem 2.1 in that paper, we have $(2\gamma-1)V_L(\hat{z}) = \gamma V_L'(\hat{z}+) - (1-\gamma)V_L'(\hat{z}-) = 0$, which by (1.3.30) reduces to (1.3.23). Since $V_L'(\hat{z}+) \geq 0$, $V_L'(\hat{z}-) \leq 0$ and $\gamma < 1$, we have $V_L'(\hat{z}+) = V_L'(\hat{z}-) = 0$,

[34] For the current purpose, we can set $\phi = 1$ WLOG by rescaling time. Drift is second order to the local time regulation and thus has no bearing on the boundary condition.

which is (1.3.31). For H, Z^H is 0-elastic, γ-skew Brownian motion and thus we have (1.3.22) directly from Appuhamillage et al. (2011, Theorem 2.1). Finally, from the seller's perspective, Z is $\frac{(2\gamma-1)}{1+e^{\beta}}$-elastic, γ-skew Brownian motion and the same theorem gives (1.3.35). □

Proof of Theorem 1.3.2. We have already shown that any skew-resetting equilibrium must solve the equations (1.3.24)-(1.3.37), that a solution exists if and only if $K \in (K^{**}, K^{*})$, and that the solution is unique when it exists. That the candidate is indeed an equilibrium follows from verification arguments similar to those given earlier, which we omit here.

□

1.9 Appendix C: Variations and Comparative Statics

1.9.1 Symmetric Incomplete Information

Define an *interval equilibrium*, denoted $I(\beta, \alpha)$, as the belief process $Z = \widehat{Z}$ and (pure) strategy profile such that the seller exits at $\tau_\beta := \inf\{t \geq 0 : Z_t \leq \beta\}$ and the buyer $\tau_\alpha := \inf\{t \geq 0 : Z_t \geq \alpha\}$. In the interval (β, α) both players' value functions are of the form

$$V_i(z) = \frac{C_i^1 e^{mz} + C_i^2 e^{(1-m)z}}{1 + e^z} \qquad i = 1, 2. \tag{1.9.1}$$

We state six boundary conditions to be used later:

$$V_S(\alpha) = 1 \tag{1.9.2}$$

$$V_S(\beta) = 0 \tag{1.9.3}$$

$$V_S'(\beta) = 0 \tag{1.9.4}$$

$$V_B(\alpha) = p(\alpha) - k \tag{1.9.5}$$

$$V_B(\beta) = 0 \tag{1.9.6}$$

$$V_B'(\alpha) = p'(\alpha). \tag{1.9.7}$$

Lemma 1.9.1. *Given α, there is an increasing function $\beta^{sym} : \mathbb{R} \to \mathbb{R}$ such that $\beta^{sym}(\alpha)$ is the unique solution $\beta \leq \alpha$ to* (1.9.1) *and* (1.9.2)-(1.9.4).

Proof. By direct substitution, (1.9.2) − (1.9.4) become

$$\frac{C_1^S e^{m\alpha} + C_2^S e^{(1-m)\alpha}}{1+e^{\alpha}} = 1 + c \tag{1.9.8}$$

$$\frac{C_1^S e^{m\beta} + C_2^S e^{(1-m)\beta}}{1+e^{\beta}} = c \tag{1.9.9}$$

$$\frac{(1+e^{\beta})[mC_1^S e^{m\beta} + (1-m)C_2^S e^{(1-m)\beta}]}{(1+e^{\beta})^2} = \frac{e^{\beta}[C_1^S e^{m\beta} + C_2^S e^{(1-m)\beta}]}{(1+e^{\beta})^2}. \tag{1.9.10}$$

Combining (1.9.8) and (1.9.9) yields

$$\begin{pmatrix} C_1^S \\ C_2^S \end{pmatrix} = (e^{m\alpha+(1-m)\beta} - e^{m\beta+(1-m)\alpha})^{-1} \begin{pmatrix} e^{(1-m)\beta} & -e^{(1-m)\alpha} \\ -e^{m\beta} & e^{m\alpha} \end{pmatrix} \cdot \begin{pmatrix} (1+c)(1+e^{\alpha}) \\ c(1+e^{\beta}) \end{pmatrix}. \tag{1.9.11}$$

Substituting these into (1.9.10) and simplifying, we obtain

$$\begin{aligned} 0 =& e^{m\beta}[c(m-1)e^{(1-m)\alpha}] + e^{(m-1)\beta}[cme^{(1-m)\alpha}] - (1+c)(2m-1)(1+e^{\alpha}) \\ &+ e^{(1-m)\beta}[cme^{m\alpha}] + e^{-m\beta}[c(m-1)e^{m\alpha}] \end{aligned} \tag{1.9.12}$$

$$\begin{aligned} \iff 0 =& e^{m\alpha}[cme^{(1-m)\beta} + c(m-1)e^{-m\beta}] - (1+c)(2m-1)e^{\alpha} - (1+c)(2m-1) \\ &+ e^{(1-m)\alpha}[c(m-1)e^{m\beta} + cme^{(m-1)\beta}] \end{aligned} \tag{1.9.13}$$

For every α, (1.9.12) has two distinct roots, exactly one of which satisfies $\beta < \alpha$. To see this, note that the RHS of (1.9.12) is a Dirichlet polynomial in β with sign pattern $+, +, -, +, +$, and its second derivative has sign pattern $+, +, +, +$, which implies convexity in β. It is easy to verify that the RHS tends to $+\infty$ as $\beta \to \pm\infty$, attains a global minimum of negative value at $\beta = \alpha$, and is decreasing in β for $\beta < \alpha$. Hence (1.9.12) determines a best response function $\beta^{sym}(\alpha)$. Now the first derivative w.r.t. α of (1.9.13) has sign pattern $+, -, -$ and thus the RHS of these equations is decreasing then increasing in α. Since the RHS is negative for $\alpha = \beta$ and positive for large α, it must be increasing in α when equality holds. Together these facts imply that $\beta^{sym}(\alpha)$ is an increasing function. $\square$

Lemma 1.9.2. *There is a decreasing function $\alpha^{sym} : \mathbb{R} \to \mathbb{R}$ such that $\alpha^{sym}(\beta)$ is the unique $\alpha \geq \beta$ solving* (1.9.1) *for $i = 2$ and* (1.9.5)-(1.9.7).

Proof. Again by substitution,

$$\frac{C_1^B e^{m\alpha} + C_2^B e^{(1-m)\alpha}}{1+e^\alpha} = \frac{e^\alpha}{1+e^\alpha} - k \tag{1.9.14}$$

$$\frac{C_1^B e^{m\beta} + C_2^B e^{(1-m)\beta}}{1+e^\beta} = 0 \tag{1.9.15}$$

$$\frac{(1+e^\alpha)[mC_1^B e^{m\alpha} + (1-m)C_2^B e^{(1-m)\alpha}]}{(1+e^\alpha)^2} = \frac{e^\alpha + e^\alpha[C_1^B e^{m\alpha} + C_2^B e^{(1-m)\alpha}]}{(1+e^\alpha)^2}. \tag{1.9.16}$$

Next, we show that $\alpha^{sym}(\beta)$ is a decreasing function. Solving (1.9.14) and (1.9.15) yields Combining (1.9.8) and (1.9.9) yields

$$\begin{pmatrix} C_1^B \\ C_2^B \end{pmatrix} = \frac{e^\alpha - k(1+e^\alpha)}{e^{m\alpha+(1-m)\beta} - e^{m\beta+(1-m)\alpha}} \begin{pmatrix} e^{(1-m)\beta} \\ -e^{m\beta} \end{pmatrix}. \tag{1.9.17}$$

Substituting these solutions into (1.9.16) and simplifying gives

$$\begin{aligned} 0 = G(\alpha) := &\underbrace{e^{m\alpha}[(m-1)(1-k)e^{(1-m)\beta}]}_{A} - \underbrace{e^{(m-1)\alpha}[mke^{(1-m)\beta}]}_{B} \\ &+ \underbrace{e^{(1-m)\alpha}[m(1-k)e^{m\beta}]}_{C} - \underbrace{e^{-m\alpha}[(m-1)ke^{m\beta}]}_{D}. \end{aligned} \tag{1.9.18}$$

Note that $G'(\alpha)$ has exactly two roots, $\alpha = z_m$ and $\alpha = \beta$, and $G'(\alpha) > 0$ for all $\alpha > \max\{\beta, z_m\}$. We argue first that if (β, α) with $\alpha > \beta$ solves (1.9.12) and (1.9.18), then $\beta \leq z_m$. Suppose otherwise. Then $G'(\alpha) > 0$ for all $\alpha > \beta$, and it is easy to verify that $G(\beta) > 0$, so there is no $\alpha > \beta$ that solves (1.9.18). Next, we show that for $\beta \leq z_m$, there exists a unique $\alpha \geq \beta$ such that $G(\alpha) = 0$, which we denote $\alpha^{sym}(\beta)$, and moreover, $\alpha^{sym}(\beta) \geq z_m$. If $\beta = z_m$, then G is increasing on $[z_m, \infty)$ and $G(z_m) = 0$, so $\alpha^{sym}(\beta) = z_m$ is the solution. Next we argue that for $\beta < z_m$, there is a unique $\alpha > z_m$ such that $G(\alpha) = 0$, which we denote $\alpha^{sym}(\beta)$. If $\beta < z_m$, then G has a local maximum at β, is decreasing on (β, z_m) and is increasing on $[z_m, \infty)$, so there is a unique $\alpha > z_m$ and the claim holds.

To see that $\alpha^{sym}(\beta)$ is decreasing for $\beta \leq z_m$, note that $e^{\alpha}(1-k) > k$ and thus $C - D = B - A > 0$. The derivative of the RHS of (1.9.18) with respect to α simplifies to $m(m-1)(e^{\alpha}(1-k)-k)(e^{(m-1)(\alpha-\beta)} - e^{-m(\alpha-\beta)})$, which is strictly positive for $\alpha = \alpha^{sym}(\beta)$ and $\beta < z_m$. The derivative the RHS of (1.9.18) with respect to β is $(m-1)(B-A) + m(C-D) > 0$, so we conclude that $\alpha^{sym}(\beta)$ is decreasing. ☐

Together Lemmas 1.9.1 and 1.9.2 yield the following corollary.

Corollary 1.9.1. *There is a unique pair* $(\beta^{sym,*}, \alpha^{sym,*})$ *solving* (1.9.1) *and* (1.9.2)-(1.9.7).

Proof. Since $\beta^{sym}(\alpha)$ is increasing, $[\beta^{sym}]^{-1}(\beta)$ is also increasing. Now since $\beta^{sym}(x) < x$ for all x, $[\beta^{sym}]^{-1}(z_m) > z_m$, while as shown above, $\alpha^{sym}(\beta) = z_m$. Moreover, it is easy to verify that $[\beta^{sym}]^{-1}(x) \to -\infty$ as $x \to -\infty$, and thus for small β, $\alpha^{sym}(\beta) > z_m > [\beta^{sym}]^{-1}(\beta)$. These together with the facts that $\alpha^{sym}(\beta)$ is decreasing and $\beta^{sym}(\alpha)$ is increasing imply that there exists a unique solution to (1.9.8) – (1.9.16), denoted $(\beta^{*,sym}, \alpha^{*,sym})$. ☐

Lemma 1.9.3. *Any equilibrium under symmetric incomplete information is an interval equilibrium.*

Proof. By a repeat of arguments in the proof of Lemma 1.7.4 and Lemma 1.7.8, there must exist states $\alpha \neq \alpha'$ such that for $z = \alpha, \alpha'$, $V_B(z) = p(z) - k$ and adoption is immediate. We now claim that for any state $z_0 \in (\alpha, \alpha')$, adoption must also be immediate. Suppose not; then there exists some interval $(\alpha_1, \alpha_2) \subseteq (\alpha, \alpha')$ that contains z_0 such that the buyer weakly prefers to wait at all states in (α_1, α_2). Let $\tau = \inf\{t \geq 0 : Z_t \notin (\alpha_1, \alpha_2)\}$. Then τ is almost surely finite and $p_{t\wedge\tau}$ is a bounded martingale, so the optional sampling theorem applies, and we have $\mathbb{E}p_{t\wedge\tau} = p_0$. The value of immediate adoption is $p_0 - k$, while the value of waiting is at most $\mathbb{E}[e^{-r(t\wedge\tau)}(p_{t\wedge\tau} - k)] < p_0 - k$, due to the possibility of exit before time τ. It follows that the buyer would strictly prefer to deviate to immediate adoption. Hence the buyer's strategy is τ_α for some α.

Given the buyer's strategy, we show that the seller must also play some τ_β. By a repeat of arguments in the proof of Lemma 1.7.5, there must exist $\beta < \alpha$ such that $V_L(\beta) = 0$. Let $\beta_0 := \sup\{\beta < \alpha : V_L(\beta) = 0\}$. Clearly, $\beta_0 < \alpha$. Moreover, since Z is continuous, $V_L(z) = 0$ for all $z < \beta_0$. It must be that $V_L(\beta_0) = 0$, otherwise $V_L(z') > 0$ for $z' \in (\beta_0 - \epsilon, \beta_0)$ for sufficiently small $\epsilon > 0$. Finally, exit must be immediate at all $z \leq \beta_0$, otherwise the seller receives a negative expected payoff; and there is no exit at $z > \beta_0$ since $V_L(z) > 0$. □

Proof of Proposition 1.4.1. Given Corollary 1.9.1 and Lemma 1.9.3, it only remains to show that the equations (1.9.2)-(1.9.7) are necessary, and to verify that $I(\beta^{sym,*}, \alpha^{sym,*})$ is an equilibrium. The boundary value conditions (1.9.2)-(1.9.3) and (1.9.5)-(1.9.6) are trivial. To establish (1.9.4), first suppose $V_S'(\beta+) < 0$. Since $V_S(\beta) = 0$, it follows that for sufficiently small $\epsilon > 0$, $V_S(z) < 0$ for $z \in (\beta, \beta + \epsilon)$; but the seller would strictly prefer to exit immediately, contradicting optimality. Next suppose $V_S(\beta+) > 0$. Then for sufficiently small $\epsilon > 0$, the seller can attain strictly positive payoff starting at $Z_0 = \beta$ by exiting at $\tau_\epsilon := \inf\{t \geq 0 : Z_t \leq \beta - \epsilon\}$, which also contradicts optimality. It follows that $V_S(\beta+) = 0$. The argument for the buyer that $V_B(\alpha-) = p'(\alpha)$ is the same as the one given in the proof of Lemma 1.7.9.

Now $I(\beta^{sym,*}, \alpha^{sym,*})$ trivially satisfies properties 3-5 of Definition 1.2.2. The steps to verify properties 1 and 2 are a repeat of those given in the proof of Theorem

1.2.1 and are therefore omitted. □

Proof of Corollary 1.4.1. To show that $\beta^{*,sym}(\alpha) < \beta^*(\alpha)$ for any real α, we evaluate the RHS of (1.9.12) at $\beta^*(\alpha) = \alpha - x^*$. The RHS of (1.9.12) is

$$c(m-1)e^{mx^*} + cme^{\alpha}e^{mx^*} - (1+c)(2m-1)(1+e^{\alpha}) + cme^{(1-m)x^*} + c(m-1)e^{\alpha - mx^*}$$

$$= e^{\alpha}(cme^{(m-1)x^*} - (1+c)(2m-1) + c(m-1)e^{-mx^*}),$$

where we have used (1.7.7). Using (1.7.7) again, this expression equals

$$P(x^*) := cm(e^{(m-1)x^*} - e^{(1-m)x^*}) - c(m-1)(e^{mx^*} - e^{-mx^*}).$$

By straightforward arguments, $P(0) = 0$ and P is strictly decreasing, so $x^* > 0$ implies $P(x^*) < 0$. As shown before, the RHS of (1.9.12) is decreasing in β for $\beta < \alpha$, and thus $\beta^{*,sym}(\alpha) < \beta^*(\alpha)$.

Next, we show that $\alpha^{*,sym}(\beta) < \alpha^*(\beta)$. Suppose otherwise. The difference in value functions is of the form $\frac{Ae^{mz}+Be^{(1-m)z}}{1+e^z}$, which has exactly one root, and hence the value functions intersect exactly once. Observe that $V_B^{sym}(\beta) = 0 < V_B(\beta)$. Then if $\alpha^{*,sym} > \alpha^*$, the value functions would intersect at least twice (in (β, α^*) and in $(\alpha^*, \alpha^{*,sym})$), while if $\alpha^{*,sym} = \alpha^*$, the functions must coincide. In either case, we have a contradiction.

Proof of claim for the buyer: Note that the value functions coincide for all $z \geq \alpha^*$, and that for $z \leq \beta^{*,sym}$, $V_B^{sym}(z) = 0 < V_B(z)$. As we have already shown, $\alpha^{*,sym} < \alpha^*$, and the "value functions" cannot intersect in $(\beta^{*,sym}, \alpha^*)$, otherwise there would be at least two points of intersection on the reals; thus the value under symmetric incomplete information must be strictly lower for all $z < \alpha*$.

Proof of claim for the seller: For $z \geq \alpha^*$, the seller has value 1 in either case. For $z \in (\alpha^{*,sym}, \alpha)$, we have $V_S(z) = 1 > \bar{V}(z)$. For $z \leq \beta^{*,sym}$, $V_S(z) = 0 < \bar{V}(z)$. It remains to consider $z \in (\beta^{*,sym}, \alpha^{*,sym})$. Note that for $z \leq \beta^*$, $\bar{V}(z) = \frac{e^z}{1+e^z}V_H(\beta^*)$.

By straightforward computation, $V_S(z)$ [notation] is of the form $\frac{Ae^{mz}+Be^{(1-m)z}}{1+e^z}$, so it must intersect the function $\frac{e^z}{1+e^z}V_H(\beta^*)$ at zero or exactly two points. By examining limit behavior, it is evident that there are two points of intersection, z_1 and z_2 with $z_1 < \beta^{*,sym} < z_2$. Moreover, $V_S(z)$ must intersect $\bar{W}(z)$ exactly once, at some point $z_3 < \alpha^{*,sym}$. If $z_3 > \beta^*$, then the claim follows with $\hat{z} = z_3$. If $z_3 < \beta^*$, then the claim follows with $\hat{z} = z_2$.

Proof of claim for social welfare: The weighted social welfare functions $\eta\bar{V}(z) + (1-\eta)V_B(z)$ and $\eta V_S^{sym}(z) + (1-\eta)V_B^{sym}(z)$ are both of the form $\frac{Ae^{mz}+Be^{(1-m)z}}{1+e^z}$, therefore intersect exactly once. □

1.9.2 Commitment

Proof of Proposition 1.4.2. For any starting belief z_0 and threshold $\alpha \in [z_0, z_0 + x^*]$, we define the buyer's interim payoff function

$$V^{int}(\alpha; z_0) = \frac{C_1^{int}(\alpha)e^{mz_0} + C_2^{int}(\alpha)e^{(1-m)z_0}}{1+e^{z_0}}, \tag{1.9.19}$$

where $C_1^{int}(\alpha), C_2^{int}(\alpha)$ solve (1.7.11) and (1.7.12) for $\beta = \beta^*(\alpha)$. By expanding and simplifying,

$$V^{int}(\alpha; z_0) = M_0(z_0)(e^{\alpha}(1-k) - k)(me^{m(z_0-a+x^*)} + (m-1)e^{(1-m)(z_0-a+x^*)})$$

$$\implies \frac{d}{d\alpha}V^{int}(\alpha; z_0) = M_0(z_0)[M_1(z_0)e^{m\alpha} + M_2(z_0)e^{(m-1)\alpha} + M_3(z_0)e^{(1-m)\alpha} + M_4(z_0)e^{-m\alpha}],$$

where we define

$$M_0(z_0) = [(1+e^{z_0})((m-1)e^{(1-m)x^*} + me^{mx^*})]^{-1} > 0$$

$$M_1(z_0) = (1-k)m(m-1)e^{(1-m)(z_0+x^*)} > 0$$

$$M_2(z_0) = -k(m-1)^2 e^{(1-m)(z_0+x^*)} < 0$$

$$M_3(z_0) = -(1-k)m(m-1)e^{m(z_0+x^*)} < 0$$

$$M_4(z_0) = km^2 e^{m(z_0+x^*)} > 0.$$

Since the sign pattern is $+,-,-,+$, there are exactly two or zero roots.

Note that $V^{int}(z_0; z_0) = p(z_0) - k = \frac{e^{z_0}}{1+e^{z_0}} - k$ for all z_0. By algebra, $\frac{d}{d\alpha}V^{int}(\alpha; z_0)|_{\alpha=z_0} \leq 0$ if and only if

$$z_0 \geq \ln\left(\frac{k(m^2e^{mx^*} - (m-1)^2e^{(1-m)x^*})}{(1-k)m(m-1)(e^{mx^*} - e^{(1-m)x^*})}\right) =: z_1.$$

Note that z_1 is precisely α^*, the competitive threshold.

For $\alpha > z_0 + x^*$, that is $z_0 < \alpha - x^*$, there is an atom of exit at time 0 and the belief conditional on no exit jumps to $\alpha - x^*$. We have

$$\begin{aligned}\underline{V}^{int}(\alpha; z_0) &= \frac{e^{z_0-(\alpha-x^*)} + e^{z_0}}{1+e^{z_0}} V^{int}(\alpha; \alpha - x^*) \\ &= (2m-1)M_0(z_0)e^{z_0+x^*}(1-k-ke^{-\alpha}).\end{aligned}$$

Note that $\underline{V}^{int}(\alpha; z_0)$ is strictly increasing in α and $\lim_{\alpha\to\infty} \underline{V}^{int}(\alpha; z_0) = (2m-1)M_0(z_0)e^{z_0+x^*}(1-k) =: \underline{V}(z_0)$.

By algebra, $p(z_0) - k \geq \underline{V}(z_0)$ if and only if

$$z_0 \geq \ln\left(\frac{k(me^{mx^*} + (m-1)e^{(1-m)x^*})}{(1-k)((m-1)e^{(1-m)x^*} + me^{mx^*} - (2m-1)e^{x^*})}\right) =: z_2.$$

We now show that $z_1 \geq z_2$ if and only if $x^* \geq \widehat{x}$, where $\widehat{x}$ is a positive constant.

By algebra, the inequality $z_1 > z_2$ is equivalent to

$$m^2 e^{2mx^*} - m^2(2m-1)e^{(m+1)x^*} + 2m(m-1)e^{x^*}$$
$$+(m-1)^2(2m-1)e^{(2-m)x^*} + (m-1)^2 e^{2(1-m)x^*} \geq 0.$$

The left side above is a polynomial $P(x^*)$ with 2 sign changes, and hence it has exactly zero or two roots. The limits are $\lim_{x^* \to \pm\infty} = \infty$, and it is easy to verify that $P(0) = 0$. By algebra, $P'(0)$ factors as $-(2m-1)(m-1)m < 0$. It follows that P has one other root, strictly positive, denoted $\widehat{x}$.

Claim: For all z_0, α such that $z_m < \alpha \leq z_0 + x^*$, $\frac{\partial}{\partial z_0}\left(\frac{V^{int}(\alpha;z_0)}{\underline{V}(z_0)}\right) > 0$. Define $P_3(\alpha) := \frac{e^{\alpha}(1-k)-k)}{(2m-1)(1-k)e^{\alpha}}$. By algebra,

$$\frac{V^{int}(\alpha;z_0)}{\underline{V}(z_0)} = (me^{(m-1)(z_0-\alpha+x^*)} + (m-1)e^{-m(z_0-\alpha+x^*)})P_3(\alpha) \qquad \text{and}$$

$$\frac{\partial}{\partial z_0}\left(\frac{V^{int}(\alpha;z_0)}{\underline{V}(z_0)}\right) = m(m-1)(e^{(m-1)(z_0-\alpha+x^*)} - e^{-m(z_0-\alpha+x^*)})P_3(\alpha),$$

which is positive whenever $z_m < \alpha \leq z_0 + x^*$. Moreover, it is easy to check that $\lim_{z_0\to\infty} \frac{V^{int}(\alpha;z_0)}{\underline{V}(z_0)} = \infty$ and that $\frac{V^{int}(\alpha;z_0)}{\underline{V}(z_0)}|_{z_0=\alpha-x^*} = 1 - \frac{k}{1-k}e^{-\alpha} < 1$. Thus there is a unique cutoff $z_3 > z_m - x^*$ such that $\max_{\alpha \leq z_0+x^*} V^{int}(\alpha; z_0) \geq \underline{V}(z_0)$ if and only if $z_0 \geq z_3$.

We treat the cases $x^* \leq \widehat{x}$ and $x^* > \widehat{x}$ separately.

Case (i): $x^* \leq \widehat{x}$. Here, we have $z_1 \leq z_2$. For $z_0 \geq z_2 \geq z_1$, immediate adoption is optimal. For $z_0 \in (z_1, z_2)$, $\underline{V}(z_0)$ exceeds the value of immediate adoption, which in turn exceeds $\max_{\alpha\in[z_0,z_0+x^*]} V^{int}(\alpha; z_0)$. It follows that $z_3 > z_2$, and thus if $z_0 \leq z_1$, then $z_0 < z_3$ and $\underline{V}(z_0)$ is the supremum.

Case (ii): $x^* > \widehat{x}$. In this case, $z_1 > z_2$. For $z_0 \geq z_1$, immediate adoption is optimal. For $z_0 \in (z_2, z_1)$ there is an interior maximum. This implies that $z_3 < z_2$, so an interior maximum remains optimal for $z_0 \in (z_3, z_1)$. For $z_0 \leq z_3$, the supremum is $\underline{V}(z_0)$. $\square$

1.9.3 Commitment with Symmetric Incomplete Information

Proof of Proposition 1.4.3. Recall the definition of $\beta^{sym}(\alpha)$ as the solution β to (1.9.12). Define

$$V_B^{sym}(\alpha; z_0) := \frac{C_1^B(\alpha)e^{mz_0} + C_2^B(\alpha)e^{(1-m)z_0}}{1+e^{z_0}}, \tag{1.9.20}$$

where $C_1^B(\alpha)$ and $C_2^B(\alpha)$ are given by (1.9.17) using $\beta = \beta^{sym}(\alpha)$.

For $z_0 \leq \beta^{sym}(z_m)$, if $\alpha < [\beta^{sym}]^{-1}(z_0)$ then $\alpha < z_m$ and thus $V_B^{sym}(\alpha; z_0) < 0$. For $\alpha \geq [\beta^{sym}]^{-1}(z_0)$, we have $z_0 \leq \beta^{sym}(\alpha)$, so exit is immediate and the buyer obtains a payoff of zero. Hence, any $\alpha \geq [\beta^{sym}]^{-1}(z_0)$ is optimal.

For $z_0 \geq \alpha^{*,sym}$, we argue that immediate adoption is optimal. To see this, note that the discussion following (1.9.18) and the fact that $\beta^{sym}(\alpha)$ is increasing imply that $\frac{\partial}{\partial z_0} V_B^{sym}(\alpha; z_0)|_{z_0=\alpha}$ is increasing in α. Since $\alpha^{*,sym}$ is the value of α that solves $\frac{\partial}{\partial z_0} V_B^{sym}(\alpha; z_0)|_{z_0=\alpha} = p'(\alpha)$, we have $\frac{\partial}{\partial z_0} V_B^{sym}(\alpha; z_0)|_{z_0=\alpha} > p'(\alpha)$ for all $\alpha > \alpha^{*,sym}$. Now recall that $V_B^{sym}(\alpha; \alpha) = p(\alpha) - k$. Thus if $z_0 \geq \alpha^{*,sym}$ and $\alpha > z_0$, we have $V_B^{sym}(\alpha; z_0) < p(z_0) - k$, so immediate adoption is optimal at z_0.

For the remaining interval we define a correspondence $\boldsymbol{\alpha} : [\beta^{sym}(z_m), \alpha^{*,sym}] \rightrightarrows [z_m, \alpha^{*,sym}]$ by $\boldsymbol{\alpha}(z_0) := \arg\max_{\alpha \in [z_m, \alpha^{*,sym}]} V_B^{sym}(\alpha; z_0)$. We argue that $\boldsymbol{\alpha}$ is singleton-valued. Note that $V_B^{sym}(\alpha; z_0)$ is continuous in each argument, and the set $[z_m, \alpha^{*,sym}]$ is compact. By the Maximum Theorem, $\boldsymbol{\alpha}$ is a compact-valued and upper hemi-continuous correspondence. We argue that $\boldsymbol{\alpha}$ has the following property: for all $z_0', z_0 \in [z_m, \alpha^{*,sym}]$, $z_0' > z_0 \implies \min \boldsymbol{\alpha}(z_0') \geq \max \boldsymbol{\alpha}(z_0)$.

Observe that for any fixed $\alpha, \alpha' \in [z_m, \alpha^{*,sym}]$, $V_B^{sym}(\alpha; z_0)$ and $V_B^{sym}(\alpha'; z_0)$ as functions of z_0 intersect exactly once. To see this, note that $V_B^{sym}(\alpha; z_0) - V_B^{sym}(\alpha; z_0)$ has the same sign as $[C_1^B(\alpha) - C_1^B(\alpha')]e^{mz_0} + [C_2^B(\alpha) - C_2^B(\alpha')]e^{(1-m)z_0}$. The latter is a polynomial in z_0 of length 2 and has at most one root. Now suppose WLOG that $\alpha' > \alpha$. Since $V_B^{sym}(\alpha; z_0)$ is increasing in z_0 for $z_0 \in [\beta^{sym}, \alpha]$, we have $V_B^{sym}(\alpha'; \beta^{sym}(\alpha')) = 0 < V_B^{sym}(\alpha; \beta^{sym}(\alpha'))$. Next, since $\alpha' \leq \alpha^{*,sym}$ we $V_B^{sym}(\alpha'; z_0)$ crosses $p(z_0) - k$ from above at α'. By algebra,

$$G(z_0) := V_B^{sym}(\alpha'; z_0) - (p(z_0) - k) \propto C_1^B(\alpha')e^{mz_0} - (1-k)e^{z_0} + k + C_2^B(\alpha')e^{(1-m)z_0}. \tag{1.9.21}$$

Since $C_1^B(\alpha') > 0 > C_2^B(\alpha')$, this implies that $G(z_0)$ has three roots: α', one root strictly above α', and one root strictly below $\beta^{sym}(\alpha')$. This implies that $G(\alpha) > 0$, so $V_B^{sym}(\alpha'; \alpha) > p(\alpha) - k = V_B^{sym}(\alpha; \alpha)$. By the Intermediate Value Theorem, $V_B^{sym}(\alpha; z_0)$ and $V_B^{sym}(\alpha'; z_0)$ must intersect at some $z_0^* \in (\beta^{sym}(\alpha'), \alpha)$. We have already argued that there is at most one intersection point, so z_0^* is unique.

Now suppose $z_0 < z_0'$, and there exist $\alpha \in \boldsymbol{\alpha}(z_0)$, $\alpha' \in \boldsymbol{\alpha}(z_0')$ such that $\alpha > \alpha'$. As above, there exists z_0^* such that $V_B^{sym}(z; \alpha) \geq V_B^{sym}(z; \alpha')$ if and only if $z > z_0^*$. Hence we have $z_0 \geq z_0^*$ and $z_0' \leq z_0^*$, so $z_0 \geq z_0'$, a contradiction.

To complete the proof, consider any selection from $\boldsymbol{\alpha}$, that is, a function α such that $\alpha(z_0) \in \boldsymbol{\alpha}(z_0)$ for all z_0 in its domain. By the above property, $\alpha(z_0)$ is an increasing function, and hence it is continuous almost everywhere. If $\boldsymbol{\alpha}(z_0)$ is not a singleton, then z_0 is a point of discontinuity of α. Thus the set of z_0 such that $\boldsymbol{\alpha}(z_0)$ is not a singleton is a subset of a measure zero set and therefore has zero measure itself. This implies that $\boldsymbol{\alpha}$ is almost everywhere a singleton. □

1.9.4 Comparative Statics

Proof of Proposition 1.5.1. Consider first the case of private information. Recall that α^* can be written explicitly as $\alpha(k) = \ln \frac{kd}{1-k}$, where $d := \frac{m^2 e^{(m-1)x^*} - (m-1)^2 e^{-mx^*}}{m(m-1)(e^{(m-1)x^*} - e^{-mx^*})} > 1$. Note that $\alpha'(k) = \frac{1}{k(1-k)} > 0$, and thus β^* satisfies $\beta'(k) = \frac{d}{dk}[\alpha(k) - x^*] = \alpha'(k) > 0$, proving the second claim of the proposition. Hence V_H and V_L simply translate right from an increase in k, and since they are increase in z, these function weakly decrease in k. For the buyer, using earlier expressions for C_1, C_2:

$$C_1(k) := \frac{m(e^{\alpha(k)}(1-k)-k)e^{-m(\alpha(k)-x^*)}}{(m-1)e^{(1-m)x^*}+me^{mx^*}},$$

$$C_2(k) := \frac{(m-1)(e^{\alpha(k)}(1-k)-k)e^{(m-1)(\alpha(k)-x^*)}}{(m-1)e^{(1-m)x^*}+me^{mx^*}}.$$

By algebra, the partial derivative of the buyer's value at $\beta = \alpha(k) - x^*$

$$\frac{\partial}{\partial k} V_B(\beta; k) = -\frac{k(2m-1)(d-1)}{(1-k)(1+e^{\beta})[(m-1)e^{(1-m)x^*}+me^{mx^*}]} < 0.$$

Now, it is easy to check that $C_1'(k) < 0$ and $C_2'(k) > 0$, so given $k' > k$, the functions $V_B(z; k)$ and $V_B(z; k')$ intersect exactly once, at some $z' < \alpha(k) - k$, and $V_B(z; k) > V_B(z; k')$ for all $z \geq z'$. It follows that $V_B(\cdot; k')$ lies entirely below $V_B(\cdot; k)$.

For the case of symmetric information, note that by writing the RHS of (1.9.18) as $G(\alpha, k)$ and evaluating at $\alpha = \alpha^{sym}(\beta)$ which solves (1.9.18), we have $G_k(\alpha, k) = G_k(\alpha, k) - k^{-1} \cdot G(\alpha, k) = -k^{-1} \cdot ((m-1)e^{m\alpha+(1-m)\beta} + me^{(1-m)\alpha+m\beta}) < 0$. Since $G_\alpha(\alpha^{sym}(\beta), k) > 0$, $\alpha^{sym}(\beta)$ is increasing in k. Next, write (1.9.13) as

$$\begin{aligned} 0 = \widehat{G}(\alpha, \delta) :=& c(m-1)e^{m\delta} + cme^{\alpha+(m-1)\delta} - (1+c)(2m-1)(e^{\alpha}+1) \\ &+ c(m-1)e^{\alpha-m\delta} + cme^{(1-m)\delta}. \end{aligned} \tag{1.9.22}$$

Now $\widehat{G}_\delta = cm(m-1)[e^{m\delta} - e^{(1-m)\delta} + e^{\alpha+(m-1)\delta} - e^{\alpha-m\delta}] > 0$, and $\widehat{G}_\alpha = e^{\alpha}[cme^{(m-1)\delta} - (2m-1)(1+c) + c(m-1)e^{-m\delta}]$, which is negative when (1.9.22) holds. It follows that the implicit function $\delta(\alpha)$ is increasing, and recalling that $\beta'(\alpha) > 0$, we have now $\beta'(\alpha) = 1 - \delta'(\alpha) \in (0, 1)$. By translation invariance of V_S conditional on θ, the impact on the seller conditional on θ can be decomposed into two effect: a rightward shift in both thresholds of distance $\beta'(\alpha) > 0$, and a further shift in α of distance $\delta'(\alpha) > 0$. Both effects weakly hurt the seller conditional on any $\theta \in \{H, L\}$, so his unconditional value V_S is weakly decreasing in k. $\square$

Proof of Proposition 1.5.2. For private information, note that $J(x;c)$ from (1.7.7) is decreasing in both c and x, and thus $x^*(c)$ is decreasing, or equivalently, $\beta^*(\alpha;c)$ is increasing in c. Since the buyer's best response function $\alpha^*(\beta)$ is increasing, it follows that both equilibrium thresholds β^* and α^* are increasing. Now let $c' > c$ be two distinct cost levels. Then $V_L(z;c') - V_L(z;c) = (C_1^L(c') - C_1^L(c))e^{mz} - (c'-c) + (C_2^L(c') - C_2^L(c))e^{(1-m)z}$ is a Dirichlet polynomial of length three and has at most two roots. Since $\beta^*(c') > \beta^*(c)$, and both $V_L(\cdot;c')$ and $V_L(\cdot;c)$ are strictly convex with double roots at $\beta^*(c')$ and $\beta^*(c)$, respectively, it follows that they intersect at some $z \in (\beta^*(c), \beta^*(c'))$. Moreover, since $\alpha^*(c') > \alpha^*(c)$ and $V_L(\alpha^*(c');c) > V_L(\alpha^*(c);c) = 1 > V_L(\alpha^*(c');c')$, there can be no roots in $(\beta^*(c'), \alpha^*(c'))$ (otherwise this interval would have at two roots, for a total of more than three), and thus $V_L(\cdot;c')$ lies weakly below $V_L(\cdot;c)$. Now the buyer's polynomials $v_B(z;c)$ and $v_B(z;c')$ can intersect at most once. Since these functions are strictly convex and their points of smooth pasting to the myopic value curve satisfy $\alpha^*(c') > \alpha^*(c)$, it follows that they must intersect at a point in $(\alpha^*(c)\alpha^*(c'))$, and it follows that $V_B(\cdot;c')$ lies weakly above $V_B(\cdot;c)$.

For type H, note that as $c \to 0$, $x^* \to \infty$, but α^* is bounded below by z_m; since α^* is monotonic in c, denote $\underline{\alpha} := \lim_{c\to 0} \alpha^*$. Since $x^* \to \infty$, $\beta^* \to -\infty$. From expressions in the proof of Lemma 1.7.3, we have that $\lim_{c\to 0} C_1^H \to \widehat{C}_1^H$, for some positive and finite value $\widehat{C}_1^H$. It follows that $C_2^H e^{(1-m)\beta^*} = \frac{m-1}{m} e^{m\beta^*} C_1^H \to 0$. Now $V_H(\beta^*;c) = C_1^H e^{m\beta^*} + C_2^H e^{(1-m)\beta^*} - c \to 0$ as $c \to 0$, as desired. This also implies that for some fixed z, $V_H(z;c)$ is increasing in c. Moreover, since $V_H(\alpha^*;c) = 1$, $V_H(z;c)$ is increasing in z, and α^* is increasing in c, we have that $\frac{\partial}{\partial c} V_H(z;c)|_{z=\alpha^*(c)} < 0$, so together we conclude that $V_H(\cdot;c)$ is nonmonotonic in c.

For the case of symmetric information, standard arguments show that the seller's best response function $\beta^{sym}(\alpha)$ shifts up, and since $\alpha^{sym}(\beta)$ is decreasing, this implies that β^{sym} is increasing in c and α^{sym} is decreasing. By similar arguments to the ones above, $V_B(\cdot;c')$ lies weakly below $V_B(\cdot;c)$ if $c' > c$. On the other hand, for sufficiently small $\epsilon > 0$ and fixed $z \in (\alpha^{sym}(c) - \epsilon, \alpha^{sym}(c))$, $V_S(z;c) < 1$ while $V_S(z;c') = 1$ for sufficiently large c'. For small $\epsilon > 0$ and fixed $z \in (\beta^{sym}(c), \beta^{sym}(c) + \epsilon)$, $V_S(z;c) > 0$ while $V_S(z;c') = 0$ for sufficiently large c'. Since $v_S(z;c)$ and $v_S(z;c)$ can intersect at most once, they intersect exactly once at some cutoff in $(\beta^{sym}, \alpha^{sym})$. □

Proof of Proposition 1.5.3. We divide the proposition into four parts.

1. Private information, r_B: To see that the best response function $\alpha^*(\beta)$ shifts up as r_B increases, note that for a fixed adoption threshold $\alpha > \beta$, the buyer's value decreases for all $z \in (\beta, \alpha)$. Since $\alpha^*(\beta; r_B)$ satisfies smooth pasting of $V_B(\cdot; r_B)$ onto $p(z) - k$ for r_B, if $r'_B > r_B$, then $V_B(z; r'_B) < p(z) - k$ for $z \in (\alpha^*(\beta) - \epsilon, \alpha^*(\beta))$, and thus $\alpha^*(\beta; r'_B) < \alpha^*(\beta; r_B)$. It follows that both thresholds are decreasing, at the same rate, in r_B and that B is made weakly worse off. Since V_H and V_L are nondecreasing in z, H and L are made weakly better off by the translation left.

2. Private information, r_S: By similar arguments to those used above, increasing r_S decreases payoffs pathwise at all points inside $(\beta^*(\alpha), \alpha)$ and thus $V_L(z; r'_S)$ becomes negative at $\beta^*(\alpha; r_S)$ for $r'_S > r_S$, so β^* must increase with r_S. Since $\alpha^*(\beta)$ is increasing in β, both thresholds shift right. Thus $V_S(\cdot; r_S)$ weakly increases with r_S and V_θ decreases for $\theta = H, L$.

3. Symmetric information, r_B: Clearly $\alpha^{sym}(\beta)$ shifts down as r_B increases. It follows that both equilibrium thresholds are decreasing in r_B, and the seller is weakly better off. By increasing to any $r'_B > r_B$, for small $\epsilon > 0$, the buyer is made worse off at any $z \in (\alpha^{sym} - \epsilon, \alpha^{sym})$ as adoption becomes immediate. For $z = \beta^{sym}$ (and by continuity, an open interval of beliefs), the buyer is made better off by the change as her value becomes strictly positive.

4. Symmetric information, r_S: By similar arguments to part 2, $\beta^{sym}(\alpha)$ shifts up, and thus in equilibrium β^{sym} increases while α^{sym} decreases. Now consider increasing r_S to r'_S. The buyer is made worse off pathwise and thus in expectation. For $z \in (\beta^{sym}, \beta^{sym} + \epsilon)$, the seller is made worse off as his value falls to 0; for $z \in (\alpha^{sym} - \epsilon, \alpha^{sym})$, he is made better off as his value increases to 1. □

2

Optimal Entry Timing

Originally published as http://dx.doi.org/10.1016/j.jet.2015.03.002.

2.1 Introduction

In a wide variety of settings, an agent must choose the optimal time to take an aggressive action against an opponent. The aggressor naturally has private information about her strength, which may be high or low, and her opponent must choose how to respond to the action given his beliefs about the aggressor. If he is confident that the aggressor is weak, he may decide to fight back, knowing that he is likely to win. If instead he believes the aggressor is strong, he may prefer to concede, cutting his losses. Public information about the potential aggressor arrives through a continuous stream of news with noise; that is, both good news and bad news arrive in infinitesimal increments, but the likelihood of good news is correlated with the aggressor's strength. The aggressor's public reputation is influenced both exogenously by this news process and endogenously by the information inferred from the timing of her action. As the optimal response under uncertainty depends on the aggressor's reputation, the aggressor should optimize the timing of her action anticipating the path her reputation might take in the future.

Such games of aggression or "attack" occur in many specific settings. While for expositional purposes I will focus on an economic setting of market entry, it is worth highlighting some other applications of the model.

- A military power or rebel organization decides when to invade a territory in the hope of gaining control or having some demands met. Prior to this decision, exogenous information arrives through central intelligence reports and social media. Its opponent may simply meet these demands, or it can engage in risky physical conflict.

- A politician running for office decides when to launch a harsh negative ad campaign, and its opponent then either quits the race or responds with heightened campaign spending. Or, an activist group may choose when to publicly rally to influence a decision making authority, and the authority responds by changing its stance or not. Information arrives in either case through journalism, social media and polls.

- An injured party may privately threaten a defendant with a lawsuit, and the defendant can try to negotiate out of court or accept a costly legal battle. Information arrives in the form of evidence, legal research, and traditional news outlets insofar as they influence a potential jury pool.

In each of these settings, the aggressor has private information about her strength and can exploit a strong reputation by intimidating an opponent. In this paper, I model such scenarios as a dynamic game of private information in continuous time with an exogenous news process.

Consider a technology firm, firm one, deciding when to enter a market by introducing a highly innovative product. Specifically, this could be a firm that develops an innovative smart phone or a start-up that creates a new way for consumers to rent movies and video games. Its competitor, firm two, is a traditional firm in the same industry. Naturally, firm one has an informational advantage over firm two about the specific features of the innovation and relevant consumer preferences. However, this information is gradually revealed publicly in two ways: an exogenous news process and the fact that firm one has not yet decided to enter. The exogenous news process represents the aggregate information revealed through channels such as financial statements, customer reviews, expert speculation and leaks. Since firm one's entry strategy depends on its private information, firm two also revises its beliefs based on the timing of entry.

Prior to the market entry, firm two receives flow profits from its continuing business while firm one receives nothing. If firm one introduces the product, firm two decides whether to compete with firm one and final payoffs are realized. Conceding and exiting the market is firm two's safe option, and firm one then receives a large payoff in the form of monopoly profits. If firm two chooses to compete, both firms must pay fighting costs in the form of unavoidable sunk costs. The outcome of competition depends on the true type of firm one. Firm one wins and receives a high payoff after competition if and only if it has a high quality product. For simplicity I assume that a single winner emerges with a natural monopoly. Figure 2.1 shows the interpretation of the instantaneous period beginning at time t. In the final section of the paper, I consider variations of the model in which only firm two can enter or either firm can enter.

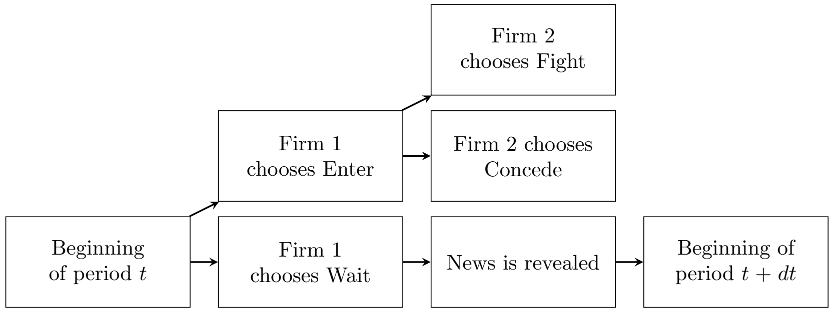

FIGURE 2.1: Timeline for period t

Stationary subgame perfect Bayesian equilibrium is a natural solution concept for this model, using the public belief about firm one's type as the state variable. As in other games of private information and learning, delay arises in equilibrium. The high (i.e., strong) type of firm one anticipates good news to arrive exogenously and improve its reputation, and thus has an incentive to wait. The low type wants to mimic the high type so it can enter and induce firm two to concede.

I first define a class of "interval equilibria" in which strategies have an appealing structure. For high beliefs, both types of firm one enter with certainty. For low beliefs, the high type enters with certainty while the low type mixes to make firm two indifferent to fighting back. For an interval of intermediate beliefs, both types wait for news to arrive, and the belief is driven solely by the exogenous news process. I present a novel equilibrium refinement that belief updating must be *weakly belief order preserving (WBOP)*. The essence of this refinement is that, for any fixed action of firm one, the posterior belief upon observing that action should be nondecreasing in the prior belief. In the main result of the paper, I show that in a particular sense, the class of equilibria surviving this refinement is precisely the class of interval equilibria with nondecreasing value functions. The class of equilibria surviving D1 is found to be a nonempty, proper subclass of the class of WBOP equilibria.

One feature of the model is that a high type of firm one receives positive payoff from entry, even when firm two fights back with certainty. This feature limits the size of no-entry regions that can occur in equilibrium, because in the interior of such a region, a high type may prefer to deviate by entering rather than wait for beliefs to exit this region. It turns out that higher fighting costs, higher news quality and greater patience allow for larger delay regions by increasing the incentive of the high type to wait. For a fixed equilibrium, the value function of each firm one type is weakly decreasing in both the discount rate and the fighting cost. Since higher news quality hastens exit out of the delay region without affecting the distribution over exit boundaries, both firm one value functions are increasing in news quality, while the opposite is true for firm two and total welfare. The total welfare effect of the *presence* of news depends on the starting belief; news can introduce the possibility of fighting on the equilibrium path for some beliefs, but it can introduce the possibility of full concession for other beliefs.

The interpretation of the main findings depends on the externalities imposed by the game in the specific application. First, one can compare social welfare across equilibria for a fixed set of model inputs. In real market entry applications, the introduction of a new product to the market likely exerts a positive net externality on consumers. If this externality is sufficiently large, then equilibria with longer delay are worse for social welfare. In a military application, however, invasion and conflict presumably exert negative net externalities, so delay is socially desirable. For applications in politics, activism, and legal disputes, net externalities likely fall somewhere between these extremes and depend on further details. Second, one can compare social welfare by fixing an equilibrium and varying news quality; then quality of news determines the expected duration of delay, so in this sense higher news quality may help or hurt social welfare depending on the application.

2.1.1 Related Literature

This paper belongs to the literature on dynamic games and stopping problems. With thematic differences, the information structure of the model is closely related to Daley and Green (2012) ("DG"). In their model, a seller has private information about the quality of an object and faces a competitive market of buyers who make continuous offers on the basis of exogenous news and the fact that sale has not yet occurred. A market for lemons arises even if there is common knowledge of gains from trade, since at certain reputation levels a high-value seller may prefer to wait, expecting upward drift from news; in DG, this results in inefficient "no-trade" regions. Similarly, equilibria arise in the current model which contain "no-entry" regions. However, the current model differs from DG in several respects. First, in DG flow payoffs to player 1 depend on her type but payoffs at termination do not; in this model the reverse is true. Second, a feature of the DG model is that two forces encourage the high type seller to delay trade: she not only anticipates good news but also earns higher utility from retaining the object. In the current model, this monotonicity is absent: a high type expects good news but earns a positive payoff even when player 2 fights back. Finally, in my model, player 1 chooses when to end the game, but player 2 makes the final move; in DG, the seller is responsible for both the timing of the end and the final move.

Aside from DG, a variety of papers have looked at inefficiency from delay. In Taylor (1999), a home owner may want to set a high price to weaken the negative inference made by a prospective buyer about its quality upon observing no sale. In Bonatti and Hörner (2011), team members procrastinate on a project that has uncertain potential payoff. In Ortner (2013), a monopolist with stochastic time-varying costs may find it optimal to delay sale in equilibrium. When asymmetric information is added to the baseline model in Ambrus et al. (2013), bidders in an eBay-like auction may delay high bids in order to manipulate their rivals' beliefs.

There is a large existing literature on dynamic games of one-sided incomplete information that is revealed through a player's actions. In the 1960s, Aumann and Maschler developed the first of such models, where player one observes the (persistent) state that determines the true payoff matrix, and player two only learns by observing player one's actions; see Aumann and Maschler (1995). Hörner et al. (2010) generalize this work by allowing the state to follow a Markov process. Dynkin (1969) introduced a general class of games in which players choose when to stop a stochastic process that determines their termination payoffs. Touzi and Vieille (2002) have expanded this literature to include mixed strategies in continuous time.

A large number of recent economics papers have employed continuous time techniques in timing and/or reputation games. Faingold and Sannikov (2011) use a general framework to study the reputation of a large player with imperfectly observed actions facing a continuum of small players. Gryglewicz (2009) models a similar market entry interaction to that of the current paper, where instead of news, a stochastic process drives the value of the market. In Heinsalu (2014), a news process arises endogenously from noisy signalling, but in that model there are no actions other than signalling. Gul and Pesendorfer (2012) model a game between opposing political parties who choose how long to incur costs to spread information. Murto and Välimäki (2011) characterize exit waves in a stopping game with informational externalities. In Moscarini and Squintani (2010), privately informed firms in an R&D race decide when to abandon a project, and a survivor's curse arises in equilibrium.

The current paper also contributes to the literature on market entry. In the model of Loury (1979), firms make investment decisions under rivalry and uncertainty; market failure arises in the form of excessive R&D expenditure and an excess of firms entering the market. In Dixit (1989), firms facing volatile prices enter the market at higher thresholds than they exit, due to sunk costs. In Bergemann and

Välimäki (2002), an entrant adjusts the supply of its innovation as it learns about its quality. In Jovanovic (1981), entrants have private information but they choose entry locations. On the empirical side of the literature, there is much work relating firm and industry characteristics to entry timing. Schoenecker and Cooper (1998) use data from the minicomputer industry in the 1990s to find a positive correlation between early entry and higher R&D intensity, possession of a direct sales force, larger firm size and stronger involvement in a threatened market; the entrant's type in the current model can be thought to aggregate factors such as these. Mitchell (1989) looks at the decision to enter an emerging field from the perspective of an incumbent, which is related to the variations of the current model. Other well-studied phenomena in this literature include preemption, deterrence, first-mover advantage or disadvantage, brand extensions and patent protection. For an extensive review of the market entry literature, see Helfat and Lieberman (2002) and Hauser et al. (2006).

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 defines and characterizes equilibria and discusses several refinements. Section 4 provides comparative statics. Section 2.5 concludes and briefly introduces several variations, which are treated in detail in a separate, supplementary appendix. All proofs not provided in the main text are contained in the appendices of this paper or in the supplement.

2.2 Model

I model the entry game as a dynamic game between two players played in continuous time with infinite horizon. Player 1 ("she") is of privately known type, $\theta \in \Theta := \{H, L\}$, and may be interpreted as a firm of high or low strength; player 2 ("he") is an incumbent that places prior probability p_0 on $\theta = H$, which is common knowledge. Except for in the variations of the model discussed in Section 2.5, only player 1 can enter, and while the game continues, players 1 and 2 earn flow payoffs of 0 and 1, respectively.[1] If player 1 enters, player 2 immediately chooses to compete ("fight") or concede. In either case, the game ends with lump sum payoffs as follows:

Table 2.1: Lump Sum Payoffs

Player 1 Type	Player 2	Payoffs
H or L	Concede	$(1, 0)$
H	Fight	$(1-k, -k)$
L	Fight	$(-k, 1-k)$

The player that retains the market earns a monopoly profit of 1. The constant $k \in (0, 1)$ is interpreted as a fighting cost paid by both players.[2] Players are risk-neutral expected utility maximizers who discount future utility at rate $r > 0$. Before introducing the public news process, I construct the underlying probability space. To allow for randomization by Nature over player 1's type, I introduce a probability space $(\Theta, 2^\Theta, \nu)$, where ν is a probability measure over Θ such that $\nu(H) = p_0$. Let B_t be a standard Brownian motion on $(\Omega, \mathcal{F}, \mathbb{Q})$, the canonical probability space. Assume that $\mathcal{F}$ and 2^Θ are independent, and define the product space $(\Omega', \mathcal{H}, \mathbb{P}) := (\Omega \times \Theta, \mathcal{F} \times 2^\Theta, \mathbb{Q} \times \nu)$. States in Ω' will be referred to as ω' or explicitly as (ω, θ). Public news is defined on this space and arrives according to a Brownian diffusion process given by

$$dX_t = \mu_\theta dt + \sigma dB_t. \tag{2.2.1}$$

Assume that $\mu_H \geq \mu_L$ and $\sigma > 0$, and without loss of generality set $X_0 = 0$. The informativeness of news is given by the signal-to-noise ratio, $\phi := \frac{\mu_H - \mu_L}{\sigma}$. Let $(\mathcal{F}_t)_{t \geq 0}$ be the filtration generated by X: for all $t \geq 0$, $\mathcal{F}_t = \sigma\left(\{X_s : 0 \leq s \leq t\}\right)$. The

[1] Player 2's flow payoff is strategically irrelevant except in the variations of the model, but a unit flow payoff is consistent with the value of the market.

[2] That each type receives payoff 1 in monopoly can be assumed WLOG by normalization. In a more general model, H and L could have different normalized fighting costs but qualitatively the results would be the same; the key feature is the way the types rank the outcomes.

information available publicly at time t is captured by $\mathcal{F}_t$.

Player 2 updates beliefs given the public news, whether entry is occurring in the current period, and the fact that entry has not occurred prior to the current period.

A pure strategy for player 1 specifies when to enter as a function of her type and available information, if she has not entered previously. Fundamentally, a pure strategy for θ is a $\mathcal{F}_t$-adapted stopping time that can take the value $+\infty$. I follow the approach in recent economics literature in interpreting mixed strategies as continuous distribution functions.[3] In the current model, such stopping strategies need to be defined at every time t_0 in order to describe the continuation behavior in case the game is off-path at t_0. The definition below directly corresponds to that in Daley and Green (2012), and I refer the reader to that paper for further discussion.

Definition 2.2.1. *An entry strategy for type θ of player 1 is a family of $\mathcal{F}_t$-adapted stochastic processes $\{(A_t^{\theta,t_0})_{t\geq t_0}\}_{t_0\in\mathbb{R}_+}$ such that*

- *For all $t_0 \geq 0$, the process $(A_t^{\theta,t_0})_{t\geq t_0}$ is a right-continuous, nondecreasing, $\mathcal{F}_t$-adapted process taking values in $[0,1]$.*

- *For all $t_1 \geq t_0 \geq 0$, let $A_{t_1-}^{\theta,t_0} := \lim_{s\uparrow t_1} A_s^{\theta,t_0}$ and specify that $A_s^{\theta,t_0} = 0$ for all $s < t_0$. If $A_{t_1-}^{\theta,t_0} < 1$, then for all $t \geq t_1$,*

$$A_t^{\theta,t_1} = \frac{A_t^{\theta,t_0} - A_{t_1-}^{\theta,t_0}}{1 - A_{t_1-}^{\theta,t_0}}.$$

A pure strategy for θ is one that does not depend on randomization, so that $A_t^{\theta,t_0}(\omega) \in \{0,1\}$ for all $t \geq t_0 \geq 0$ and $\omega \in \Omega$.

[3] In addition to DG, see Gul and Pesendorfer (2012), although the latter is slightly different in that a mixed strategy is a CDF of a distribution over state variable thresholds, which in turn determine pure stopping strategies. For a discussion of the equivalence of various representations of randomized stopping times, see Touzi and Vieille (2002).

Definition 2.2.2. *A response strategy for player 2 is a family of $\mathcal{F}_t$-measurable functions $f_t : \Omega' \to [0,1]$, $t \in \mathbb{R}_+$.*

A strategy for player 2 specifies, for each time t, a probability of fighting back against entry at time t given the public history up to time t. A pure strategy is one that takes only values in $\{0,1\}$.

2.3 Equilibria

The most general form of equilibrium I will consider is stationary subgame perfect Bayesian equilibrium, which I define in this section. Given a fighting strategy f, types H and L at time s face the optimal stopping problems

$$\max_{\tau \geq s} \mathbb{E}^H \left[e^{-r(\tau - s)} (1 - k f_\tau) | \mathcal{F}_s \right] \tag{$AP^{H,s}$}$$

$$\max_{\tau \geq s} \mathbb{E}^L \left[e^{-r(\tau - s)} (1 - (1+k) f_\tau) | \mathcal{F}_s \right]. \tag{$AP^{L,s}$}$$

For all $s \geq 0$, define $\mathcal{A}^{\theta,s}$ to be the support of θ's continuation strategy from time s; that is, the set of stopping times τ such that for all $\epsilon > 0$ and for all $\omega \in \Omega$, $A^{\theta,s}_{\tau(\theta,\omega)+\epsilon}(\omega) - A^{\theta,s}_{\tau(\theta,\omega)-\epsilon}(\omega) > 0$. A strategy $A^{H,s}$ (resp. $A^{L,s}$) is said to solve ($AP^{H,s}$) (resp. ($AP^{L,s}$)) if for all $\tau \in \mathcal{A}^{\theta,s}$, τ solves the corresponding maximization problem.

Player 2 must respond to entry by first updating his belief and then acting optimally given that belief. Since he can guarantee payoff 0 by conceding, player 2's strategy is optimal at s if f_s solves

$$\max_{q \in [0,1]} q \cdot \mathbb{E}[(-k) \mathbb{1}\{\theta = H\} + (1-k) \mathbb{1}\{\theta = L\} | \mathcal{F}_s, \tau = s], \tag{FP_s}$$

where the conditioning on $\tau = s$ incorporates the additional information that player 1 has chosen to enter at time s.

Next I discuss belief updating. It is convenient to use the log-likelihood transformation of beliefs, $Z_t = \ln \frac{p_t}{1-p_t}$.[4] Suppose the game continues at the beginning of some time t. I impose the strong condition, common to most definitions of perfect Bayesian equilibrium, that beliefs must be formed using Bayes' rule whenever possible, even off the equilibrium path. Hence, for all $t_1 > t_0 \geq 0$, whenever $A^{H,t_0}_{t_1-} \cdot A^{L,t_0}_{t_1-} < 1$, the log-likelihood belief satisfies[5]

$$Z_{t_1} = Z_{t_0} + \phi \left[\frac{2\mu_\theta - \mu_H - \mu_L}{2\sigma}(t_1 - t_0) + B_{t_1} - B_{t_0} \right] + \ln \frac{1 - A^{H,t_0}_{t_1-}}{1 - A^{L,t_0}_{t_1-}}. \qquad (2.3.1)$$

The current belief is the sum of three terms: the starting belief, the information derived from the exogenous news process, and the information that entry has not yet occurred. If the game continues at the beginning of some period $t_1 > 0$ such that for all $t_0 \in [0, t_1)$, $A^{H,t_0}_{t_1-} \cdot A^{L,t_0}_{t_1-} = 1$, then (2.3.1) does not apply, but the equilibrium concept and later refinements impose structure on the belief process.

If entry occurs at some time t, then player 2 updates beliefs one final time before choosing a response. These beliefs must be derived from Bayes' rule whenever possible. The equilibrium concept and later refinements impose further structure on belief updating.

Definition 2.3.1. *A stationary subgame perfect Bayesian equilibrium of the entry game is an $\mathcal{F}_t$-adapted public belief process $\{Z_t\}_{t\geq 0}$, a type-dependent entry strategy for player 1, a response strategy for player 2, and a measurable updating function $j : \bar{\mathbb{R}} \to \bar{\mathbb{R}}$ such that:*

1. *Player 1's entry strategy is optimal: for all $s \geq 0$, strategies $A^{\theta,s}$ solve the corresponding problem $(AP^{H,s})$ or $(AP^{L,s})$.*

[4] I use $p(z) := \frac{e^z}{1+e^z}$ to denote the inverse transformation.

[5] For a detailed derivation, I refer the reader to Daley and Green (2012), since the belief process prior to entry here is analogous to the belief process prior to sale there.

2. *Player 2's response strategy is optimal: for all* $s \geq 0$, f_s *solves* (FP_s).

3. *For all* $t_1 > t_0 \geq 0$ *and* $\omega \in \Omega$ *such that* $A^{H,t_0}_{t_1-} \cdot A^{L,t_0}_{t_1-} < 1$, *the beginning-of-period belief* Z_{t_1} *satisfies* (2.3.1).

4. *For all* $t \geq 0$, *if entry occurs at time* t *and* $Z_t = z$, *then the post-entry belief is* $j(z)$, *which is formed by Bayes' rule whenever possible.*[6]

5. *For all* $\theta \in \{H, L\}$ *and* $u \in [0, \infty]$, *there exists a function* $g^{\theta,u} : \bar{\mathbb{R}} \to [0,1]$ *such that for all* $t \geq 0$ *and all* $\omega \in \Omega$, $\mathbb{E}^\theta[A^{\theta,t}_{t+u}(\omega)|\mathcal{F}_t] = g^{\theta,u}(Z_t(\omega,\theta))$.

6. *There exists a function* $f : \bar{\mathbb{R}} \to [0,1]$ *such that for all* $t \geq 0$ *and all* $\omega' \in \Omega'$, $f_t(\omega') = f(j(Z_t(\omega')))$.

7. Z_t *is a time-homogeneous,* $\mathcal{F}_t$*-Markov process.*

Condition 5 above states that player 1's strategy is stationary: the probability of θ entering in the next u units of time should only depend on θ, u, and the current state Z_t. I will say that entry is in the support of θ's strategy at z if $g^{\theta,u}(z) > 0$ for all $u > 0$. Condition 6 imposes stationarity on player 2's strategy and the further requirement that her fighting frequency depend only on the post-entry belief.[7] Condition 7 puts further structure on the belief process off the equilibrium path.[8] For brevity I refer to stationary subgame perfect Bayesian equilibria as simply "equilibria."

I use Ξ to denote an arbitrary equilibrium and $\mathcal{X}$ to denote the class of all equilibria. I use $V_\theta(z;\Xi)$ and $V_2(z;\Xi)$ to denote value functions under equilibrium Ξ, dropping the dependence on Ξ when there is no ambiguity.

[6]At $z = \pm\infty$, $j(z) = z$ in any equilibrium, so the rest of the paper focuses on pre-entry beliefs $z \in \mathbb{R}$.

[7]In Section 5, I discuss an alternative, weaker stationarity requirement that would allow fighting frequencies to depend on the pre-entry belief.

[8]This condition also appears in DG. Condition 5 is not sufficient for condition 7 because the belief process is undefined after off-path waiting.

2.3.1 Static Benchmark and Preliminary Results

In this section I briefly discuss the static version of the entry game, in which there is no news process and player 1 must decide to enter now or never. Define $p^* := 1-k$ so that $p^*(-k) + (1-p^*)(1-k) = 0$, and let $z^* = \ln \frac{p^*}{1-p^*} = \ln \frac{1-k}{k}$. Then z^* is the threshold posterior belief for player 2's indifference; he fights with probability one (zero) for posterior beliefs strictly below (above) z^*. Note that entry is a strictly dominant strategy for H in the static game. If $z_0 > z^*$, then $j(z_0) > z^*$ regardless of L's strategy, so player 2 concedes and L enters with probability 1. If $z_0 \leq z^*$, L must randomize so that $j(z_0) = z^*$, and player 2 must fight with probability $\frac{1}{1+k}$ to make L indifferent. These strategies comprise the unique equilibrium of the static game and are analogous to the strategies of the particular equilibrium of the dynamic game given in Proposition 2.3.2.

There are several features of the static equilibrium discussed above which carry over to the dynamic setting, summarized below. These preliminary results for the dynamic game provide a foundation for the rest of the paper:

- Neither type can be revealed by entry (Lemma 2.3.1).
- No equilibrium exists in pure strategies (Proposition 2.3.1).
- Player 2 fights with probability $\frac{1}{1+k}$ at z^* (Lemma 2.3.2).

Lemma 2.3.1 (No Separation on Entry)**.** *In any equilibrium, at any $z \in \mathbb{R}$, it cannot be that exactly one player 1 type enters with an atom of mass at z or that entry is in the support of exactly one type's strategy.*

Proof. First suppose on the contrary that either $g^{H,0}(z) > 0 = g^{L,0}(z)$ or that entry at z is in the support of H's strategy alone for some $z \in \mathbb{R}$. Then the posterior after entry is $j(z) = +\infty$ and player 2 concedes. L would then strictly prefer to enter at belief z, a contradiction. If instead $g^{L,0}(z) > 0 = g^{H,0}(z)$ or entry is in the support of L's strategy alone, then the belief jumps to $-\infty$ and player 2 fights back; but then L would strictly prefer to wait forever. □

Proposition 2.3.1. *No equilibrium exists in pure strategies.*

Proof. Suppose there is an equilibrium in pure strategies. Then for any belief z, either both types enter with probability 1, both enter with probability 0, or one enters with probability 1 and the other with probability 0. By No Separation on Entry, only the first two cases can occur in equilibrium. For any real $z < z^*$, it must be that both types enter w.p. 0 at z; if both entered w.p. 1, then the posterior after entry would be $j(z) = z < z^*$, player 2 would fight back, and L would strictly prefer to wait. Thus both types enter with probability 0 for beliefs in $(-\infty, z^*)$, so in this region beliefs evolve solely according to news. Recall that type H can obtain at least $1-k > 0$ by entering immediately. Since player 1 can receive positive payoff only after entry, an upper bound on the continuation value to type H from waiting is the expected discounted value of entry, met by pure concession, at the first time the belief is at least z^*. This continuation value converges to zero as the starting belief approaches $-\infty$, so for sufficiently low z, H would strictly prefer to enter immediately, contradicting that players enter with probability 0. □

An extension of the above logic yields Lemma 2.3.2, proven in the appendix.

Lemma 2.3.2. *In any equilibrium, player 2 must play the strategy summarized by the function*

$$f(z) := \begin{cases} 0 & \text{if } z > z^* \\ \frac{1}{1+k} & \text{if } z = z^* \\ 1 & \text{if } z < z^*. \end{cases}$$

The following proposition characterizes a particular equilibrium. While it will be subsumed by a later characterization theorem, I present it here for expositional purposes.

Proposition 2.3.2. *An equilibrium exists in which all entry occurs at $t = 0$. Type H enters with probability one for any starting belief. Type L enters with probability one for any starting belief $z \geq z^*$ and enters with probability $g^{L,0}(z) = exp(z - z^*)$ for all $z < z^*$ so that $j(z) = z^*$. Following any delay, the belief becomes $-\infty$, and type L never enters. Player 2 plays the strategy given in Lemma 2.3.2.*

The equilibrium outcome described in Proposition 2.3.2 is the same as the outcome of the unique Bayesian Nash Equilibrium of the static game. It is also the outcome of any equilibrium that satisfies the interval structure discussed next for the dynamic game with uninformative news, $\mu_H = \mu_L$. The value functions in this equilibrium are

$$V_H(z) = \begin{cases} 1 & \text{if } z > z^* \\ \frac{1}{1+k} & \text{if } z \leq z^* \end{cases}$$

$$V_L(z) = \begin{cases} 1 & \text{if } z > z^* \\ 0 & \text{if } z \leq z^*. \end{cases}$$

While the equilibrium described above remains an equilibrium when news is informative, news also introduces other equilibria, with a period of delay before entry occurs. A high type anticipates good news and thus has an incentive to wait for its arrival, and a low type has an incentive to wait in order to mimic the high type. The rest of this paper analyzes equilibria with such delay.

2.3.2 Interval Equilibria

As mentioned above, one way in which news might play a nontrivial role in equilibrium is through "no-entry" regions of beliefs where both types of player 1 wait for news. Due to the flexibility of off-path beliefs, the entry game has a vast multiplicity of equilibria. In fact, for any discrete set of real numbers, an equilibrium can be constructed that contains disjoint no-entry intervals that cover this set. However, there is a class of equilibria that are more appealing both in their simplicity and their optimality for each player 1 type.[9] Consider the following strategies:

- For $z \geq \alpha$, both types enter immediately with probability one.
- For $z \in (\alpha - \delta, \alpha)$, both types wait.
- For $z \leq \alpha - \delta$, H enters with probability one and L enters with probability such that Bayes' rule yields $j(z) = z^*$.

If an equilibrium satisfies the above properties for some $\alpha \in \mathbb{R}$ and $\delta \in \mathbb{R}_+$, it is said to be an *interval equilibrium.*[10] Later, I provide two distinct refinement concepts that select only interval equilibria.

In the supplementary appendix, I define the process Y_t as the belief process updated only through news, so that in a no entry region the evolution of beliefs reduces to

$$dZ_t^H = dY_t^H = \frac{\phi^2}{2}dt + \phi dB_t \tag{2.3.2}$$

$$dZ_t^L = dY_t^L = -\frac{\phi^2}{2}dt + \phi dB_t. \tag{2.3.3}$$

Note that when $\theta = H$, the belief drifts upward, and when $\theta = L$ it drifts downward.

The following definition fully specifies a candidate for interval equilibrium.[11]

[9] This claim is formalized in Proposition 2.4.5.

[10] Formally, player 1's strategies in an interval equilibrium are of the form given by Definition 2.3.2.

[11] It is important to note that Definition 2.3.2 only specifies one such candidate; even for a fixed (δ, α) there is a large multiplicity of off-path beliefs that can support the given strategy profile in equilibrium.

Definition 2.3.2. *For any $\alpha \in \mathbb{R}$ and $\delta \in \mathbb{R}_+$, define $\Xi(\delta, \alpha)$ to be the belief process, updating rule, and strategy profile as follows:*

$$Z_t(\omega') = \begin{cases} Y_t(\omega') & \text{if } \inf_{s \geq 0}\{s : Y_s(\omega') \notin (\alpha - \delta, \alpha)\} \geq t \\ -\infty & \text{otherwise} \end{cases}$$

$$j(z) = \begin{cases} z & \text{if } z \geq \alpha \\ -\infty & \text{if } z \in (\alpha - \delta, \alpha) \\ z^* & \text{if } z \leq \alpha - \delta \end{cases}$$

$$f_t(\omega') = f(j(Z_t(\omega')))$$

$$A_t^{H,t_0}(\omega) = \begin{cases} 1 & \text{if } \tau^H \leq t \\ 0 & \text{otherwise} \end{cases}$$

$$A_t^{L,t_0}(\omega) = \begin{cases} 1 & \text{if } \tau^L \leq t \text{ and } Z_{\tau^L} \geq \alpha \\ e^{Z_{\tau^L} - z^*} & \text{if } \tau^L \leq t \text{ and } Z_{\tau^L} \leq \alpha - \delta \\ 0 & \text{otherwise,} \end{cases}$$

where $\tau^\theta := \inf\{s \geq t_0 : Z_s(\omega, \theta) \notin (\alpha - \delta, \alpha)\}$ and f is defined as in Lemma 2.3.2.

Note that by construction, the belief process and the updating rule are uniquely determined by Bayes' rule except after entry occurs from within $(\alpha - \delta, \alpha)$ and after waiting occurs at α or above. I will say that two equilibria are *equivalent* if they have the same strategy profile, so that any interval equilibrium is equivalent to some $\Xi(\delta, \alpha)$. Two classes, E and F, of equilibria are said to be equivalent if every equilibrium in E is equivalent to some equilibrium in F and vice versa.

Characterization of Interval Equilibria

To begin, note that any interval equilibrium must satisfy the position constraint

$$z^* \in [\alpha - \delta, \alpha]. \tag{2.3.4}$$

If $z^* < \alpha - \delta$ in some interval equilibrium, a contradiction arises because by construction L mixes such that for any $z \in (z^*, \alpha - \delta)$, $j(z) = z^* < z$, but the belief cannot fall since H enters with probability one. If $z^* > \alpha$, $j(\alpha) = \alpha < z^*$, and L would strictly prefer to wait, contradicting that both types enter with probability one. An implication of (2.3.4) is that if $\delta = 0$ in some interval equilibrium, then $\alpha = z^*$. The equilibrium in Proposition 2.3.2 is equivalent to $\Xi(0, z^*)$ and its value functions have already been characterized. The rest of this section analyzes the value functions for $\Xi(\delta, \alpha)$ with $\delta > 0$ and $\alpha > z^*$.[12]

Above the no-entry region, entry is immediate and is met by concession, so for $z \geq \alpha$, $V_H(z) = V_L(z) = 1$. Below the no-entry region, player 1's strategy is such that the post-entry belief is z^* and player 2 fights with probability $\frac{1}{1+k}$, so for $z \leq \alpha - \delta$, $V_H(z) = 1 - \frac{k}{1+k} = \frac{1}{1+k}$ and $V_L(z) = 0$. Inside the no-entry region $(\alpha - \delta, \alpha)$, both H and L wait. V_H satisfies the Hamilton-Jacobi-Bellman equation

$$
\begin{aligned}
V_H(z) &= \mathbb{E}^H[(1 - r dt) V_H(Z_{t+dt}) | Z_t = z] \\
&= (1 - r dt)\mathbb{E}^H[V_H(z) + dZ_t V_H'(z) + \frac{1}{2}(dZ_t)^2 V_H''(z)] \\
&= (1 - r dt)[V_H(z) + \frac{\phi^2}{2} dt V_H'(z) + \frac{\phi^2}{2} dt V_H''(z)] \\
\implies r V_H(z) &= \frac{\phi^2}{2} V_H'(z) + \frac{\phi^2}{2} V_H''(z), \qquad (2.3.5)
\end{aligned}
$$

by using (2.3.2), Ito's Lemma and the rule $dB_t \cdot dB_t = dt$. Since the belief path for L is only different in that it has the opposite drift term, V_L satisfies

$$
r V_L(z) = -\frac{\phi^2}{2} V_L'(z) + \frac{\phi^2}{2} V_L''(z). \qquad (2.3.6)
$$

[12] The appendix analyzes the knife-edge case $\delta > 0$, $\alpha = z^*$, which requires a different boundary value condition at α.

These differential equations have the closed form solutions:[13]

$$W_H(z) := C_1^H \exp\left(\frac{\lambda-1}{2}z\right) + C_2^H \exp\left(\frac{-\lambda-1}{2}z\right) \tag{2.3.7}$$

$$W_L(z) := C_1^L \exp\left(\frac{\lambda+1}{2}z\right) + C_2^L \exp\left(\frac{-\lambda+1}{2}z\right), \tag{2.3.8}$$

for $z \in (\alpha-\delta, \alpha)$ where $\lambda = \sqrt{1+\frac{8r}{\phi^2}}$ and the C_i^θ are constants. Value Matching at the upper and lower boundaries yields

$$W_H(\alpha-) = 1 \tag{2.3.9}$$

$$W_L(\alpha-) = 1 \tag{2.3.10}$$

$$W_H((\alpha-\delta)+) = \frac{1}{1+k} \tag{2.3.11}$$

$$W_L((\alpha-\delta)+) = 0. \tag{2.3.12}$$

The equations (2.3.7)–(2.3.12) completely determine W_H and W_L given α and δ, which I often denote explicitly by $W_\theta(z;\delta,\alpha)$. In addition to the position constraint (2.3.4), a necessary condition for $\Xi(\delta,\alpha)$ to be an interval equilibrium is

$$W_H(z;\delta,\alpha) \geq 1-k \qquad \text{for all } z \in (\alpha-\delta,\alpha), \tag{2.3.13}$$

since H always has the option to enter and be fought with probability at most one.

Remark 3. *Due to discontinuities in the belief process and updating rule, the value functions here need not satisfy smooth-pasting at the boundaries. For an illustration, consider value functions at the lower boundary, where there may be a convex kink for both types. After any arbitrarily small delay by player 1, the belief immediately becomes* $-\infty$*, and thus entry is fought with probability one. Such deviations entail a discrete loss in utility for player 1. By contrast, in models where beliefs evolve*

[13] In some instances it will be helpful to use properties of these solutions on the entire real line, so I distinguish between W_θ and V_θ.

continuously, a player could exploit a convex kink in the value function by waiting, and thus optimality requires smooth-pasting.

For large values of δ, W_H is strictly convex, and for small values, W_H is strictly increasing. In the appendix, I implicitly characterize δ^{IE}, the maximum value of δ that satisfies (2.3.13). In the case $\alpha = z^*$, H receives only $\frac{1}{1+k}$ at α, so the corresponding maximum $\delta^{IE}_{z^*}$ is smaller.

The value functions V_θ, given below, coincide with the W_θ inside the delay region. I give an expression for V_2, the value function for player 2, in the supplementary appendix. Finally, I define total welfare as $V(z) := p(z)V_H(z) + (1 - p(z))V_L(z) + V_2(z)$. For comparative statics, it will be useful to define these functions for all $(\delta, \alpha) \in \mathbb{R}_{++} \times \mathbb{R}$, even when $\Xi(\delta, \alpha)$ is not an equilibrium:

$$V_H(z; \delta, \alpha) := \begin{cases} 1 & \text{if } z \geq \alpha \\ W_H(z; \delta, \alpha) & \text{if } z \in (\alpha - \delta, \alpha) \\ \frac{1}{1+k} & \text{if } z \leq \alpha - \delta \end{cases}$$

$$V_L(z; \delta, \alpha) := \begin{cases} 1 & \text{if } z \geq \alpha \\ W_L(z; \delta, \alpha) & \text{if } z \in (\alpha - \delta, \alpha) \\ 0 & \text{if } z \leq \alpha - \delta. \end{cases}$$

The value functions for $\Xi(\delta^{IE}, \alpha)$ are shown in Figure 2.2 below.

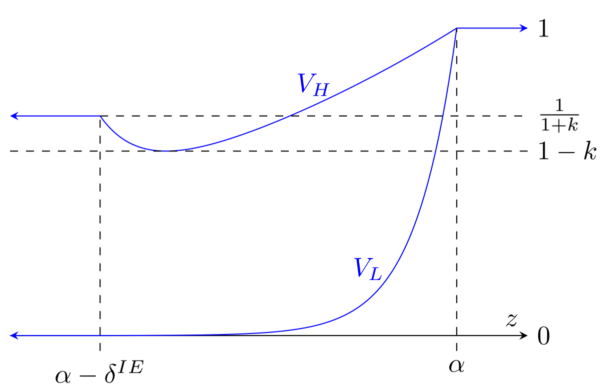

FIGURE 2.2: Value Functions in Maximal IE

Define $S^{IE} := \{(\delta, \alpha) \in \mathbb{R}^2 : \delta \in [0, \delta^{IE}], z^* \in [\alpha - \delta, \alpha)\} \cup \{(\delta, z^*) \in \mathbb{R}^2 : \delta \in [0, \delta^{IE}_{z^*}]\}$. The next theorem is proved in the appendix.[14]

Theorem 2.3.1. *The class of interval equilibria, denoted $\mathcal{X}^{IE}$, is equivalent to $\{\Xi(\delta, \alpha) : (\delta, \alpha) \in S^{IE}\}$, and every element of this set is an interval equilibrium.*

2.3.3 Nondecreasing Value Functions

This subsection considers equilibria with the property of nondecreasing value functions (NDVF) for player 1, which is desirable on two fronts. First, as will be formalized below, all NDVF equilibria have the simple interval structure defined earlier. Second, nondecreasing value functions fit the natural interpretation of upward movement in the news process as good news; as the state variable increases, player 1 weakly benefits. The equilibrium in Proposition 2.3.2 exhibits NDVF, but some interval equilibria do not, as shown in Figure 2.2.

The main intuition behind Theorem 2.3.2 is that, since the belief process is continuous, there are two possible points of exit from the delay region, and under NDVF, the value at the upper exit point must be strictly higher than the value at the lower exit point. This implies that, if the delay region is nonempty, one exit point must lie strictly above z^* and the other weakly below z^*, and thus the delay region is an interval. All interval equilibria have a nondecreasing V_L (Lemma 2.7.12), but the NDVF property for H places an upper bound on the size of the delay region which eliminates some interval equilibria. Lemma 2.7.7 gives the existence of a cutoff $\delta^* > 0$ such that $V_H(z; \delta, \alpha)$ is nondecreasing in z on $\mathbb{R}$ if and only if $\delta \in [0, \delta^*]$. Define $S^{NDVF} := \{(\delta, \alpha) : \delta \in [0, \delta^*], z^* \in [\alpha - \delta, \alpha)\} \cup \{(0, z^*)\}$.

Theorem 2.3.2. *The class of equilibria with nondecreasing value functions, denoted*

[14]Theorems 2.3.2, 2.3.3 and 2.3.4 are stated in an analogous way in order to provide examples of fully specified equilibria under various refinements. I use "class" to emphasize the multiplicity of equilibria due to modifications in off-path beliefs, and "set" to emphasize the bijection between (δ, α) and $\Xi(\delta, \alpha)$.

$\mathcal{X}^{NDVF}$, is equivalent to $\{\Xi(\delta, \alpha) : (\delta, \alpha) \in S^{NDVF}\}$, and every element of this set is an interval equilibrium with nondecreasing value functions.

In the next section, I propose an equilibrium refinement stated only in terms of the belief process and updating function that selects for equilibria with nondecreasing value functions.

Remark 4. *Since player 2 receives a flow payoff of 1, his value function inside the delay region is nonmonotonic in z. To see this, note that his values at the boundaries are $V_2(\alpha - \delta) = (1 - p(z))(1 - \exp(\alpha - \delta - z^*))$ and $V_2(\alpha) = 0$. For beliefs sufficiently close to $\alpha - \delta$, the expected flow payoffs earned before entry increase in z enough to offset the increased probability of exit through α. However, as stated in Proposition 2.4.3, this nonmonotonicity vanishes in the limit as $\phi \to \infty$ and entry becomes immediate.*

2.3.4 Weakly Belief Order Preserving Equilibria

As noted earlier, the concept of stationary equilibrium imposes little structure on off-path beliefs, creating a large multiplicity of equilibria. In particular, beliefs in equilibrium may "cross" in the following way: there may be beliefs z and z' such that $z > z'$ but the posterior belief after entry (resp. delay) from z' is higher than the posterior after entry (resp. delay) from z. In fact, if such a crossing does occur, then z and z' can be found for arbitrarily small $|z - z'|$.[15] Thus for a starting belief of z and some fixed action, a small amount of additional *bad* news arriving immediately prior to the action has the same net effect as *good* news arriving immediately after the action. Belief updating in equilibrium should not be so sensitive to the timing of news and actions as to allow this form of inconsistency, and this idea is formalized in

[15] To see this, fix an action A and let $a(z)$ denote the posterior belief. If $z > z'$ and $a(z') > a(z)$, then for any $z'' \in (z', z)$, $a(z'') > a(z)$ or $a(z'') < a(z')$.

the following definition. For an informal argument in the form of a speech by player 1, see the supplementary appendix.

Definition 2.3.3. *An equilibrium is said to be weakly belief order preserving (WBOP) if the following conditions hold:*

1. $\forall z, z' \in \bar{\mathbb{R}}$, $z > z'$ *implies* $j(z) \geq j(z')$.
2. *For almost all* $\omega' \in \Omega'$ *and all* $t \geq 0$, $Z_{t+}(\omega') := \lim_{\epsilon \downarrow 0} Z_{t+\epsilon}(\omega')$ *is well-defined.*[16]
3. *For almost all* $(\omega_1', \omega_2') \in \Omega' \times \Omega'$ *and all* $t_1, t_2 \geq 0$, *if* $Z_{t_1}(\omega_1') \geq Z_{t_2}(\omega_2')$, *then* $Z_{t_1+}(\omega_1') \geq Z_{t_2+}(\omega_2')$.

The first condition is that, for any distinct beliefs, the resulting beliefs after entry cannot reverse order. The second condition is purely technical. Essentially, the third condition requires that beliefs do not reverse order immediately following decisions to wait.

Surprisingly, it turns out that all WBOP equilibria are interval equilibria, and moreover, the class of WBOP equilibria is equivalent to the class of equilibria with nondecreasing value functions, as stated in Theorem 2.3.3. To illustrate why WBOP rules out equilibria with disconnected delay regions, consider first an interval equilibrium $\Xi(\delta, \alpha)$ with a small delay interval centered at z^*. One can modify this equilibrium to obtain a new, non-interval equilibrium by inserting a small interval (z_1, z_2) of no-entry strictly above α, supported by the off-path belief of $-\infty$. This new equilibrium fails WBOP because at any belief $z \in (\alpha, z_1)$ both types enter, so $j(z) = z$. Likewise, one can modify $\Xi(\delta, \alpha)$ by inserting a small interval (z_3, z_4) below $\alpha - \delta$. The new equilibrium fails WBOP because the belief evolves continuously

[16]This includes $\pm\infty$. Z_t need not be right-continuous, but this condition rules out belief processes for which $\limsup_{\epsilon \downarrow 0} Z_{t+\epsilon}(\omega') > \liminf_{\epsilon \downarrow 0} Z_{t+\epsilon}(\omega')$, where there is not a clear way to define the "next" belief.

following delay in (z_3, z_4), but it drops to $-\infty$ following delay in $(z_4, \alpha - \delta)$. Furthermore, player 1 value functions under WBOP are nondecreasing in equilibrium because entry from below α, off-path or not, must result in a posterior belief of exactly z^*, giving H a minimum value of $\frac{1}{1+k}$.

Since the belief process and updating rule in $\Xi(\delta, \alpha)$ do not satisfy WBOP due to the drop in beliefs that occurs after off-path events, they must be slightly modified to give a full WBOP equilibrium specification. Define $\Xi'(\delta, \alpha)$ to be $\Xi(\delta, \alpha)$ with the modification that no adjustment occurs after waiting is observed (off-path) from beliefs weakly above α, and the belief is set to z^* after entry is observed (off-path) from $z \in (\alpha - \delta, \alpha)$; formally,

$$Z_t(\omega') = \begin{cases} Y_t(\omega') & \text{if } \inf_{s \geq 0}\{s : Y_s(\omega') \leq \alpha - \delta\} \geq t \\ -\infty & \text{otherwise} \end{cases}$$

$$j(z) = \begin{cases} z & \text{if } z \geq \alpha \\ z^* & \text{if } z < \alpha. \end{cases}$$

Theorem 2.3.3. *The class of WBOP equilibria, $\mathcal{X}^{WBOP}$, is equivalent to $\mathcal{X}^{NDVF}$, and $\Xi'(\delta, \alpha)$ is a WBOP equilibrium for all $(\delta, \alpha) \in S^{NDVF}$.*

2.3.5 D1 Equilibria

The D1 equilibrium refinement, introduced in Banks and Sobel (1987) and Cho and Kreps (1987), is not formally defined for this game. However, any formal adaptation must at least include the following: if type θ would benefit from a deviation to enter off-path for a strictly larger set of player 2 fighting frequencies than type θ', then following off-path entry, player 2 must place full probability on type θ.[17]

Definition 2.3.4. *An equilibrium is said to satisfy D1 if whenever entry is off-path from $z \in \mathbb{R}$, the following conditions hold:*

[17] A stronger definition would restrict beliefs after an off-path decision not to enter and would involve much greater technical difficulty given the dynamic, continuous time setting.

1. $\{d \in [0,1] : V_H(z) \leq 1 - kd\} \subsetneq \{d \in [0,1] : V_L(z) \leq 1 - (1+k)d\} \implies j(z) = -\infty$.

2. $\{d \in [0,1] : V_H(z) \leq 1 - kd\} \supsetneq \{d \in [0,1] : V_L(z) \leq 1 - (1+k)d\} \implies j(z) = \infty$.

In any equilibrium, from any state in a no-entry region, it cannot be that full weight is placed on H after entry, because then L would strictly prefer to enter, and thus the second containment above must not hold in any D1 equilibrium. Theorem 2.3.4 states that delay region lengths for D1 equilibria are strictly smaller than and bounded away from δ^*. The intuition behind the strict inequality is that for an equilibrium with a delay region of size $\delta = \delta^*$, $V_H'((\alpha-\delta)+) = 0$ but $V_L'((\alpha-\delta)+) > 0$. This implies that at states marginally above $\alpha - \delta$, H has a stronger incentive to deviate and the D1 requirement fails. Theorem 2.3.4 also states that $\Xi(0, z^*)$ and all nondegenerate interval equilibria with $\alpha > z^*$ and small enough delay region are equivalent to some D1 equilibrium. In fact, there appears to be a threshold rule, so that $\underline{\delta} = \bar{\delta}$ can satisfy the theorem, but I have only numerical support for this conjecture.[18]

Theorem 2.3.4. *There exists a nonempty, proper subset* $S^{D1} \subset S^{NDVF}$ *such that the class of equilibria satisfying D1 is equivalent to* $\{\Xi(\delta, \alpha) : (\delta, \alpha) \in S^{D1}\}$, *and every element of this set is a D1 equilibrium. Furthermore, there exist* $\underline{\delta}, \bar{\delta}$ *with* $0 < \underline{\delta} \leq \bar{\delta} < \delta^*$ *such that* $S^{IE} \cap (\{(\delta, \alpha) : \delta < \underline{\delta}, \alpha > z^*\} \cup \{(0, z^*)\}) \subset S^{D1}$ *and in any D1 equilibrium,* $\delta \leq \bar{\delta}$.

Proof. See the supplementary appendix. □

Let $\mathcal{X}^{D1}$ denote the class of D1 equilibria. The following shorthand summarizes

[18]Calculations done by computer are available from the author upon request.

the relationships between the various equilibrium refinements discussed in this paper:

$$\mathcal{X}^{D1} \subset \mathcal{X}^{WBOP} = \mathcal{X}^{NDVF} \subset \mathcal{X}^{IE}.$$

2.4 Comparative Statics

In this section I provide comparative statics, both within and across equilibria.[19] Proposition 2.4.1 analyzes the delay, defined as the time until the belief process exits the no-entry region. This proposition makes use of the observation that an increase in the signal-to-noise ratio reduces the delay pathwise. The underlying intuition is that the only significance of noise is that it takes time to filter, slowing down the learning process; an increase in the signal-to-noise ratio by a factor $\gamma > 1$ has the same effect as an increase in the amount of news per unit of time by a factor of γ^2, and hence as a rescaling of time.[20] In particular, for both types and independent of α and δ, the quality of news has no impact on the probability, or relative timing, of exit through the barriers of the delay region. As news quality increases to infinity in the limit, delay converges to zero almost surely.

Proposition 2.4.1 also compares expected delay across equilibria. Unsurprisingly, delay increases with the size of the delay region. However, expected delay is non-monotonic and concave in the position of the delay region, as one barrier becomes closer and the other farther away. By translation, an upward shift in the delay region is akin to a downward shift in the starting belief, and thus delay is a concave function of the starting belief. Since the belief for H drifts upward, her delay is maximized at a point lower than that of L.

Proposition 2.4.1 (Delay). *For a given equilibrium $\Xi(\delta, \alpha)$ and starting belief $z \in (\alpha - \delta, \alpha)$, the delay conditional on θ, $\tau_\alpha^\theta \wedge \tau_{\alpha-\delta}^\theta$, is decreasing in ϕ and tends to 0*

[19] All proofs for this section are contained in the supplementary appendix.

[20] This follows from the scaling property of Brownian motion: γB_t has the same distribution as $B_{\gamma^2 t}$. Faingold and Sannikov (2011) report a similar result.

almost surely as $\phi \to \infty$. *The expected delay* $\mathbb{E}^\theta[\tau^\theta_\alpha \wedge \tau^\theta_{\alpha-\delta}|Z_0 = z]$ *is increasing in* δ*, concave in* α*, and converges to 0 uniformly as* $\phi \to \infty$.[21] *Delay is longest for* H *from* $z_H := \alpha - \ln \frac{e^\delta - 1}{\delta}$ *and for* L *from* $z_L := \alpha - \ln \frac{\delta}{1-e^{-\delta}} > z_H$.

Unless otherwise specified, the rest of this section considers interval equilibria with $\alpha > z^*$, since $\alpha = z^*$ is inefficient.[22] The next proposition considers the probability of eventual entry or fighting in relation to other parameters. First, fix an interval equilibrium and starting belief. As k increases, the fighting back frequency $\frac{1}{1+k}$ decreases, so conditional on $\theta = H$ the probability of eventual fighting decreases. On the other hand, player 2's indifference threshold $z^* = \ln \frac{1-k}{k}$ decreases, and L's entry frequency at the lower boundary, which is proportional to $\frac{k}{1-k}$, must rise. The net effect is that, conditional on $\theta = L$, the probability of fighting increases.

Next, consider increasing δ or decreasing α. One effect is that the belief is more likely to hit the upper boundary where entry and concession are certain. Another effect is that L enters less frequently upon reaching the lower boundary and is thus less likely to induce fighting. With respect to the probability of eventual fighting, these effects are reinforcing, and hence the probability of fighting decreases for both types. However, the above effects are opposing with respect to entry for type L. For an increase in δ, it turns out that the second effect dominates, and entry is less likely. For a decrease in α, the first effect dominates if and only if the starting belief is sufficiently high. Type H always enters at either boundary independent of its position, so her probability of entry is one.

Proposition 2.4.2 (Probability of Entry, Fighting)**.** *Consider any equilibrium* $\Xi(\delta, \alpha)$ *with* $\alpha > z^*$ *and starting belief* $z \in (\alpha - \delta, \alpha)$. *The probability of eventual entry conditional on type* H *is one in any such equilibrium. Conditional on type* L*, the*

[21] Here and in all other instances, uniform convergence refers to functions of z.

[22] That is, for any starting belief z and any interval equilibrium $\Xi(\delta, z^*)$, there exists an interval equilibrium $\Xi(\delta, \alpha)$ for some $\alpha > z^*$ with weakly higher payoffs for H, L, and player 2 at z.

probability of entry is increasing in k, decreasing in δ, and decreasing in α if and only if $2\alpha - \delta - z^* < z$. As k increases, the probability of fighting increases for L and decreases for H. For both types, the probability of fighting is increasing in α and decreasing in δ.

Proposition 2.4.3 addresses the payoff effects of changes in parameters for a fixed equilibrium. The earlier observation that higher news quality decreases delay also applies here. An increase in ϕ benefits both types of player 1 because she earns a nonnegative payoff at each barrier, and hence prefers to receive this payoff sooner. On the other hand, player 2 is made worse off as lump sum fighting costs and lost flow utility occur sooner. By the specification of payoffs, the game is constant-sum except for fighting costs, so total welfare also declines as the signal-to-noise ratio increases.[23] Now as k increases, type H is made worse off at the lower boundary and therefore at all states below α. Type L is unaffected by an increase in k because she is always indifferent at the lower boundary. Player 2 is made worse off because L enters more frequently at the lower boundary.

Proposition 2.4.3 (Value Functions within Equilibria)**.** *Consider any equilibrium $\Xi(\delta, \alpha)$ with $\alpha > z^*$ and starting belief $z \in (\alpha - \delta, \alpha)$. Both $V_H(z)$ and $V_L(z)$ are increasing in ϕ and decreasing in r, while $V_2(z)$ and $V(z)$ are decreasing in ϕ and increasing in r. $V_L(z)$ is independent of k, while $V_H(z)$, $V_2(z)$ and $V(z)$ are decreasing in k. Both $V_L(z)$ and $V(z)$ are increasing in z. As $\phi \to \infty$ or $r \to 0$,*

- $V_H(z) \to 1 - \frac{k}{k+1} \frac{e^{-z} - e^{-\alpha}}{e^{-(\alpha-\delta)} - e^{-\alpha}}$,
- $V_L(z) \to \frac{e^{z} - e^{\alpha-\delta}}{e^{\alpha} - e^{\alpha-\delta}}$, *and*
- $V_2(z) \to (1 - p(z)) \frac{e^{\alpha} - e^{z}}{e^{\alpha} - e^{\alpha-\delta}} (1 - e^{\alpha-\delta-z^*})$,

[23] I thank an anonymous referee for this observation.

where all convergence is uniform. In particular, in the limit, V_H is increasing and V_2 is decreasing in z.

There is a qualitative difference between the limiting payoffs as $\phi \to \infty$ and the payoffs of the complete information game. With complete information, type L always receives 0 and H receives 1, whereas with incomplete information and news, for starting beliefs in the delay region of a fixed equilibrium, there is a constant, positive probability of exit through the "wrong" boundary even as noise approaches 0. This same feature means that, compared to a game of incomplete information and no news, any player or type may be made better or worse off by the presence of news, depending on the starting belief. By contrast, given the presence of news, better news helps player 1 and hurts player 2 independently of the starting belief.

The following proposition describes how the parameters of the model affect the maximum length of the no-entry region in WBOP equilibria and the two cases of interval equilibria, $\alpha > z^*$ and $\alpha = z^*$. Intuitively, the temptation for H to deviate and enter inside the no-entry region is smallest when the fighting cost is large, when patience is high, and when news is highly informative, all of which incentivize type H to wait. When this temptation is low, larger no-entry regions can be sustained in equilibrium. In fact, arbitrarily large no-entry regions can be supported in the limit, with one notable exception. As $k \to 1$, δ^* is bounded above because $V_H(\alpha - \delta; \delta, \alpha)$ is bounded below by $\frac{1}{2}$, while values inside the delay region approach 0 as δ becomes large.

Proposition 2.4.4 (Delay region size)**.** *The values δ^*, δ^{IE} and $\delta^{IE}_{z^*}$ are increasing in k, increasing in ϕ, and decreasing in r. All three tend to ∞ as $\phi \to \infty$ or $r \to 0$, and to 0 as $k \to 0$, $\phi \to 0$, or $r \to \infty$. As $k \to 1$, δ^{IE} and $\delta^{IE}_{z^*}$ tend to ∞ while $\lim_{k \to 1} \delta^* < \infty$.*

There is an open set problem in finding optimal equilibria, so I characterize the suprema of value functions, which can be approximated uniformly.[24] Recall that $V_\theta(z;\delta,\alpha)$ is defined for all $(\delta,\alpha) \in \mathbb{R}_{++} \times \mathbb{R}$, even if there is no equilibrium equivalent to $\Xi(\delta,\alpha)$.

The final proposition of this section compares value functions across equilibria. As α increases, fighting happens more frequently and sooner, hurting H, L, and total welfare. The opposite is true of increases in δ for $\delta \leq \delta^*$. Below δ^*, an increase in δ causes player 1 value functions to rotate clockwise, as shown in Figure 2.3. For $\delta = \delta^*$, V_H satisfies a smooth-pasting condition at the lower boundary. As δ increases beyond δ^*, type H is made worse off as the wait for entry becomes longer, but total welfare and V_L remain increasing in δ. Intuitively, type L prefers to hold out as long as possible with the hope of entering at $z = \alpha$ with full concession, and in this sense she internalizes the social cost of fighting.

Proposition 2.4.5 (Value Functions across Equilibria)**.** *Consider any equilibrium $\Xi(\delta,\alpha)$ with $\alpha > z^*$ and starting belief $z \in (\alpha-\delta,\alpha)$. Both $V_L(z)$ and $V(z)$ are decreasing in α and increasing in δ. $V_H(z)$ is increasing in δ if and only if $\delta \leq \delta^*$, and is decreasing in α. For all $z \in \mathbb{R}$, the following hold:*

(i) $\sup_{\Xi \in \mathcal{X}} V_H(z;\Xi) = V_H(z;\delta^*,z^*)$,

(ii) $\sup_{\Xi \in \mathcal{X}} V_L(z;\Xi) = V_L(z;\delta^{IE},z^*)$,

(iii) $\sup_{\Xi \in \mathcal{X}^{NDVF}} V_\theta(z;\Xi) = V_\theta(z;\delta^*,z^*)$ *for* $\theta \in \{H,L\}$,

(iv) $\sup_{\Xi \in \mathcal{X}^{IE}} V(z;\Xi) = V(z;\delta^{IE},z^*)$,

(v) $\sup_{\Xi \in \mathcal{X}^{NDVF}} V(z;\Xi) = V(z;\delta^*,z^*)$,

(vi) $\min_{\Xi \in \mathcal{X}} V(z;\Xi) = V(z;\delta^{IE},z^*+\delta^{IE})$, *and*

(vii) $\min_{\Xi \in \mathcal{X}^{NDVF}} V(z;\Xi) = V(z;\delta^*,z^*+\delta^*)$.

[24] In Section 2.5.1, I describe a modified equilibrium concept which avoids this problem.

In addition, there exist a sequence of WBOP (and hence NDVF) equilibria $\Xi'(\delta^*, \alpha^n)$ *and a sequence of interval equilibria* $\Xi(\delta^{IE}, \alpha^n)$ *with value functions converging uniformly to the respective suprema in* $(i) - (v)$.

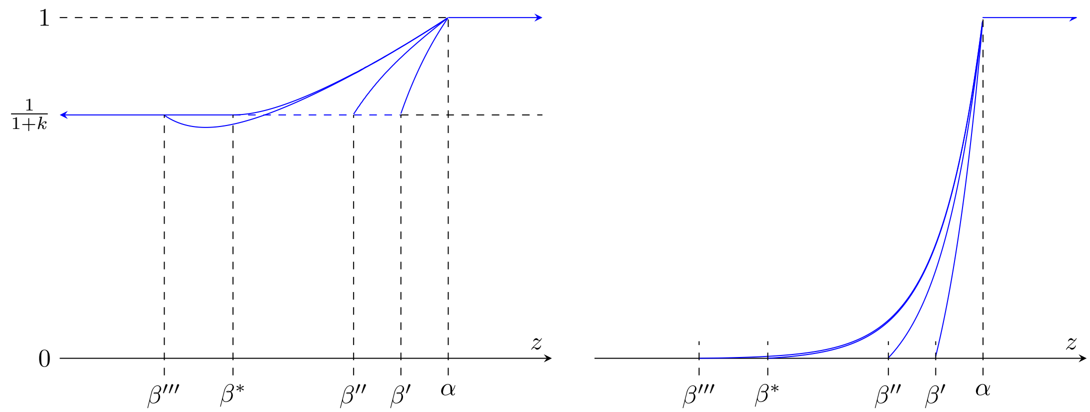

FIGURE 2.3: Value functions for equilibria with a fixed upper boundary $\alpha > z^*$ and delay regions of varying sizes, where $\alpha - \beta^* = \delta^*$. On the left, type H has increasing value in δ for $\delta < \delta^*$ and decreasing value in δ for $\delta > \delta^*$. On the right, type L always benefits as δ increases.

2.5 Discussion

2.5.1 Conclusions

In this paper I have analyzed the effect of news on an entry timing game of private information. News introduces delay as the high type has an incentive to wait for good news and the low type mimics. Under a novel equilibrium refinement, WBOP, equilibria share an appealing three-region structure and nondecreasing value functions. The length of the delay region is bounded above by the fact that a high type would prefer to enter early if the wait is too long otherwise. The D1 refinement produces the same equilibrium structure but places a tighter restriction on delay region size. As the signal-to-noise ratio or fighting cost increases or players become more patient, larger delay regions can occur in equilibrium. The presence of news can be beneficial or harmful depending on the starting belief; it can introduce delay when fighting would otherwise be immediate, but it can also introduce the possibility of fighting when entry and concession would otherwise be immediate. For a fixed equilibrium, a higher signal-to-noise ratio improves the welfare of both player 1 types, as it is equivalent to hastening the arrival of payoffs, but for the same reason it hurts player 2 and total welfare. In a richer model, the efficiency of delay would depend on the sign and magnitude of the game's externalities.

There are several directions available for future work. One could extend the current paper in a purely theoretical direction and explore the implications of WBOP-like refinements for equilibrium selection in more general dynamic games of private information. One could introduce pre-game investment by player 1. This may reveal more information about player 1, and may also reduce multiplicity of equilibria. A natural extension of the two-sided entry game would be introducing private information on the side of player 2; this may re-introduce nondegenerate interval equilibria because player 2 reveals information by waiting. Alternatively, one could allow players the more general choice of competition intensities as continuous variables as opposed to the binary decisions about entry and fighting.

2.5.2 Variations of the Model

The definition of equilibrium in this paper requires player 2's fighting probability to be a function of the post-entry belief. Alternatively, one could loosen this requirement using the pre-entry belief as the state variable. This gives player 2 the ability to mix with different probabilities after entry from beliefs z and z' as long as the post-entry belief is z^*. This alternative eliminates the open set problem, allowing WBOP equilibria with $\alpha = z^*$, because player 2 can fight back with probability 0 when entry occurs from z^*, so player 1's value at the boundary α can be 1. However, this flexibility increases the multiplicity of equilibria in a nontrivial manner because the fight back frequency then can be modified at arbitrary points inside the delay region, allowing for unnatural equilibria with a disconnected delay region of the form $(\alpha - \delta, z^*) \cup (z^*, \alpha)$ and raising the need for WBOP-like refinements for fighting frequencies in relation to prior beliefs.

In the supplementary appendix I explore three variations of the model, each retaining the original information structure. In the first variation, I allow only player 2 to enter; that is, player 1 is an incumbent and player 2 is an entrant learning about his strength relative to player 1. As player 1's behavior is trivial, the model reduces to a standard optimal stopping problem, with player 2 waiting until the belief falls to a threshold below the myopic threshold, the latter denoted z^{**}. In the second variation, I allow both players the option to enter. In equilibrium, the delay region cannot include points below z^{**} because player 2 would like to preempt entry at the lower boundary to save fighting costs against L. Furthermore, all delay is eliminated in interval equilibria. Hence, allowing player 2 to enter is similar to eliminating news, and the welfare effect depends on the starting belief. In the third variation, player 2 cannot enter, but at the beginning of each period, I allow player 2 to force entry by player 1 and then respond by fighting back or conceding.[25] I assume that player 2 learns before responding whether player 1 chose to enter ("unforced") or not ("forced"). I show that player 2 would never strictly benefit from forcing entry; by doing so, he introduces fighting costs against L in situations where L might have otherwise refrained from entry forever. In a formal sense, the class of equilibria is equivalent to that of the original game.

2.6 Appendix A: Proofs of Main Results

2.6.1 Proof of Proposition 2.3.2

I prove that Proposition 2.3.2 describes an equilibrium by satisfying all of the conditions in Definition 2.3.1. The belief process, updating rule and strategy profile

[25] I thank an anonymous referee for suggesting this variation.

are expressed formally as $\Xi(0, z^*)$.

1. Strategies for player 1 are trivial if $z = \pm\infty$, so assume $z \in \mathbb{R}$. Under the prescribed strategies, for $z > z^*$, both types enter and get the maximum payoff and thus have no strictly profitable deviation. For $z \leq z^*$, H gets $1 - k\frac{1}{k+1}$ from entering, but by waiting, the reputation drops to $-\infty$, so player 2 would fight back against any later entry, thus H obtains at most $1 - k$ by waiting, which is lower. L gets $1 - (1 + k)\frac{1}{1+k} = 0$ from entering. If she waits, she will never enter at reputation $-\infty$ and thus gets 0 as well, so she is indifferent.

2. See Lemma 2.3.2.

3. The premise only applies when $Z_{t_0} < z^*$, in which case $\ln \frac{1-A_{t_1-}^{H,t_0}}{1-A_{t_1-}^{L,t_0}} = -\infty$ and thus the belief process matches (2.3.1).

4. At $z \geq z^*$, Bayes' rule yields $j(z) = z$. At $z < z^*$, Bayes rule yields $j(z) = z^*$.

5. For all $u \geq 0$, the strategies yield $g^{H,u}(z) = 1$ for all $z \in \mathbb{R}$, $g^{L,u} = 1$ for $z \geq z^*$, and $g^{L,u} = \exp(z - z^*)$ for $z < z^*$.

6. See Lemma 2.3.2.

7. The state transitions to $-\infty$ from any state, establishing time-homogeneity, and the only time at which history is relevant is time 0, in which case the full history is simply z_0, establishing the Markov property. □

2.6.2 Proof of Theorem 2.3.1

That any interval equilibrium is equivalent to $\Xi(\delta, \alpha)$ for some $(\delta, \alpha) \in S^{IE}$ follows from the necessary condition (2.3.13) combined with Lemmas 2.7.9 and 2.7.10. The next step is to verify that for any $(\delta, \alpha) \in S^{IE}$, $\Xi(\delta, \alpha)$ is an interval equilibrium. Proposition 2.3.2 covers the only possibility for $\delta = 0$, so consider $\delta > 0$. First, I verify that $\Xi(\delta, \alpha)$ is itself an interval equilibrium when $\alpha > z^*$.

Case I: $\alpha > z^*$

1. For $z \geq \alpha$, both H and L earn the maximum payoff of 1, so entry is clearly optimal. For $z \leq \alpha - \delta$, L is indifferent, and H strictly prefers the immediate payoff $\frac{1}{1+k}$ to deviating to wait and causing the reputation to fall to $-\infty$ and getting at most $1-k$. For $z \in (\alpha - \delta, \alpha)$, the off-path beliefs are such that by entry H would get $1-k$ and L would get $-k$. Since the value of waiting to each type is W_θ, Lemma 2.7.9 implies that when $\delta \leq \delta^{IE}$, H weakly prefers to wait. L strictly prefers to wait because she has strictly positive value at z.

2. See Lemma 2.3.2.

3. When the premise holds, either the belief has remained in the no-entry region and is driven solely by news, or it has done so until exiting through the lower boundary $\alpha - \delta$, so that the belief jumps to $-\infty$ by Bayes' rule. Thus the belief process is consistent with (2.3.1).

4. For $z \leq \alpha - \delta$, Bayes' rule applies and yields $j(z) = z^*$. For $z \geq \alpha$, Bayes' rule applies and gives $j(z) = z$. Inside $(\alpha - \delta, \alpha)$, Bayes' rule does not apply.

5. This condition follows from the definition of the strategies combined with condition 7.

6. See Lemma 2.3.2.

7. Z_t is driven purely by the time-homogeneous, $\mathcal{F}_t$-Markov process Y_t until it exits $(\alpha - \delta, \alpha)$ and jumps to the absorbing state $-\infty$, so Z_t is also a time-homogeneous, $\mathcal{F}_t$-Markov process.

Case II: $\alpha = z^*$

Note that inside the delay region, the value function for H coincides with $\hat{W}_H$ instead of W_H, and L has the same value function as in Proposition 2.3.2. The above steps can be repeated exactly to show that for $\delta \leq \delta_{z^*}^{IE}$, $\Xi(\delta, z^*)$ is itself an interval equilibrium with the following modifications to the first step when $\alpha = z^*$:

- For $z = \alpha = z^*$, H prefers entry because entry yields $\frac{1}{1+k}$ whereas delay yields at most $1 - k$. L is indifferent.
- For $z \in (\alpha - \delta, \alpha) = (z^* - \delta, z^*)$, Lemma 2.7.10 and the fact that L earns 0 from entering imply that both types at least weakly prefer to wait. □

2.6.3 Proof of Theorem 2.3.2

By Lemma 2.7.11, every NDVF equilibrium is an interval equilibrium equivalent to some $\Xi(\delta, \alpha)$ with $\delta = 0$ or $\alpha > z^*$. By Lemma 2.7.7, $\delta \leq \delta^*$, establishing the first part of the theorem. In the second part of the theorem, the fact that $S^{NDVF} \subset S^{IE}$ and the second part of Theorem 2.3.1 ensure equilibrium, and Lemmas 2.7.7 and 2.7.12 ensure NDVF. □

2.6.4 Proof of Theorem 2.3.3

Together, Lemma 2.7.15 and Corollary 2.7.1 show that that every WBOP equilibrium is equivalent to $\Xi'(\delta, \alpha)$ for some $(\delta, \alpha) \in S^{NDVF}$. To prove the rest of Theorem 2.3.3, I show that for every $(\delta, \alpha) \in S^{NDVF}$, $\Xi'(\delta, \alpha)$ itself is an equilibrium and satisfies WBOP.

First, consider $\delta > 0$. The equilibrium conditions follow the same proofs as those of Theorem 2.3.1 with one minor exception. For condition 1 of Definition 2.3.1, for $z \in (\alpha - \delta, \alpha)$, the off-path beliefs are such that by entry H would get $\frac{1}{1+k}$ and L would get 0, which equal their respective boundary values at $z = \alpha - \delta$. Since the value functions are nondecreasing, both types prefer to wait. Since $\delta > 0$ implies

$\alpha > z^*$ for $(\delta, \alpha) \in S^{NDVF}$, player 1 optimality for $z \geq \alpha$ is the same as in Case I of the Proof of Theorem 2.3.1, as it is for $z \leq \alpha - \delta$.

Next, I verify WBOP conditions when $\delta > 0$:

1. By construction, $j(z)$ is a nondecreasing function, so this condition is immediate.

2. If $Z_s(\omega') \leq \alpha - \delta$ for any $s \in [0, t]$, then by construction $Z_{t+}(\omega') = -\infty$. If not, then $Z_t(\omega') = Y_t(\omega')$ and since Brownian motion is continuous with probability one, $Z_{t+}(\omega') = Y_{t+}(\omega') = Y_t(\omega') = Z_t(\omega')$ almost surely.

3. If $\exists s \in [0, t_2]$ such that $Y_s(\omega'_2) \leq \alpha - \delta$, then $Z_{t_2+}(\omega'_2) = -\infty$ and the claim is trivial. If $Y_s(\omega'_2) > \alpha - \delta$ for all $s \in [0, t_2]$, then the premise implies $Z_{t_1+}(\omega'_1) = Y_{t_1}(\omega'_1) \geq Y_{t_2}(\omega'_2) = Z_{t_2+}(\omega'_2) > \alpha - \delta$.

For the case $\delta = 0$, note that $\alpha = z^*$ since $(\delta, \alpha) \in S^{NDVF}$. The WBOP conditions are argued exactly as above. The equilibrium conditions are argued exactly as above with the following exception:

1. At $z = \alpha = z^*$, H and L earn $\frac{1}{1+k}$ and 0 from entry, respectively. By right-continuity condition of strategies, any deviation must require strictly positive delay, and by the properties of Brownian motion, the belief must enter $(-\infty, z^*)$ immediately and jump to $-\infty$, giving H at most $1 - k$ and L at most $-k$, so both types prefer entry. □

2.7 Appendix B: Supporting Results

2.7.1 Lemmas for All Equilibria

Lemma 2.7.1. *In any equilibrium, for any z, if L enters with probability one, then so does H.*

Proof. Suppose this were not true in some equilibrium for some z. If L gets value 1 from entry at z, it must be that player 2 concedes after updating his belief. H would strictly prefer to enter and receive 1 because otherwise, by right-continuity of the entry strategy CDF, there would exist some $\epsilon > 0$ such that with positive probability, she does not enter by time ϵ; and then by discounting, her expected payoff would be strictly less than 1. If L gets value less than 1 from entry, then she can deviate by waiting and mimic type H to force the belief to $+\infty$, and then enter after an arbitrarily small delay for an expected payoff arbitrarily close to 1, also a contradiction. □

Lemma 2.7.2. *In any equilibrium, if entry is in the support of at least one type's strategy at $z \in \mathbb{R}$, then $j(z) \geq z^*$.*

Proof. If $j(z) < z^*$, player 2 fights back, so L would strictly prefer to wait rather than receive a negative payoff. Hence entry is not in the support of L's strategy at z, and by Lemma 2.3.1 it is not in the support of H's strategy either. □

Lemma 2.7.3. *In any equilibrium, if $z \in \mathbb{R}$, entry is in the support of L's strategy, and $g^{L,0}(z) < 1$, then $j(z) = z^*$.*

Proof. By the previous lemma, the belief updates to at least z^*. If the updated belief exceeds z^*, player 2 would strictly prefer to concede following entry, and L would strictly prefer to enter at z, a contradiction. □

Beliefs in an equilibrium can be partitioned according to the strategies of H and L. In the table below, "Enter" in column θ means that type θ enters with probability 1, $g^{\theta,0}(z) = 1$. Use "Mix$^+$" to mean that $g^{\theta,0}(z) \in (0, 1)$ and "Mix^0" to mean $g^{\theta,0}(z) = 0$ but entering is in the support of θ's strategy. Use "Wait" to mean that entering is not in the support of θ's strategy. These labels are mutually exclusive and produce $4^2 = 16$ ordered pairs. Listing H's strategy first, by No Separation on Entry, one can rule out categories (Enter, Wait), (Enter, Mix^0), (Mix$^{+/0}$, Wait), (Mix$^+$, Mix^0), (Mix^0, Mix$^+$), (Wait, Enter) and (Wait, Mix$^{+/0}$); and by Lemma 2.7.1, one can rule out (Mix$^{+/0}$, Enter). Hence the following table is exhaustive of the possible configurations of strategies at beliefs $z \in \mathbb{R}$ that can occur in equilibrium.[26]

Table 2.2: Strategy Categories

Category	H	L
A	Enter	Enter
B	Enter	Mix$^+$
C	Mix$^+$	Mix$^+$
	Mix^0	Mix^0
D	Wait	Wait

Let $A, B, C, D \subseteq \mathbb{R}$ be the sets of beliefs that belong to the respective categories A, B, C and D. Next, I define a partial order over subsets of $\mathbb{R}$. Let $M, N \subseteq \mathbb{R}$. Then define $\succeq$ by

$$M \succeq N \iff z \succeq z' \qquad \forall z \in M, z' \in N$$

Note that if either M or N is empty, the claim trivially holds.

Lemma 2.7.4. *In any equilibrium, B is nonempty, $\{z^*\} \succeq B$ and $z^* \notin B$.*

Proof. Suppose instead that B is empty. If H enters w.p. 1 at any point $z < z^*$, by Lemmas 2.3.1 and 2.7.3, $z \in B$, a contradiction. Hence H weakly prefers to wait at every point below z^*. As in the proof of Proposition 2.3.1, this leads to a contradiction, so B is nonempty. Now suppose $z \geq z^*$ for some $z \in B$. Since H is entering w.p. 1 and L w.p. strictly less than 1, $j(z) > z \geq z^*$ by Bayes' rule. Then player 2 concedes w.p. 1 and L would strictly prefer to enter at z, a contradiction. □

[26] Trivially, both types enter at $z = +\infty$ while H enters and L waits at $z = -\infty$, so I consider only real-valued beliefs.

Proof of Lemma 2.3.2. Since B is nonempty by Lemma 2.7.4, there exists some belief at which L is indifferent between revealing herself to receive continuation value 0, or entering. The fighting frequency that makes her indifferent is the value $x = \frac{1}{1+k}$ that solves $0 = x(-k) + (1-x)1$. Condition 6 of Definition 2.3.1 requires the fighting frequency to be the same for any prior beliefs such that the post-entry belief is z^*. By the definition of z^*, the described strategy solves (FP_s). $\square$

Lemma 2.7.5. *In any equilibrium, $A \setminus \{z^*\}$ is nonempty and $A \succeq \{z^*\}$.*

Proof. If $A \setminus \{z^*\}$ is empty, it contains no points strictly above z^*, so by Lemma 2.7.4, all points above z^* belong to either C or D. Thus H either mixes or waits at beliefs strictly above z^*, so she weakly prefers to wait. Hence her expected payoff at these points can be calculated assuming that mixing always results in waiting everywhere in this region. An upper bound on the discounted value of entry is the present value of inducing full concession at the first time the belief falls to z^* or below, which is arbitrarily close to 0 for large z, so that H strictly prefers to enter at z for payoff $1 - k > 0$, a contradiction. Thus $A \setminus \{z^*\}$ is nonempty. Now for $z \in A$, both types enter at z, so the posterior is $j(z) = z$, and by Lemma 2.7.2, $z \geq z^*$, so $A \succeq \{z^*\}$. $\square$

Lemma 2.7.6. *In any equilibrium, C is empty.*

Proof. Suppose not, and consider any $z_c \in C$. I analyze three cases.

Case I: $z_c < z^*$: By the almost sure continuity of Brownian motion, the belief process is continuous in states in D w.p. 1. From states in C, entry causes the belief to jump to z^* by Lemma 2.7.3, so waiting must cause a downward jump. From states in B, waiting causes the belief to jump to $-\infty$. Now, by indifference, the payoff to H can be computed as if any randomization results in delay. Thus starting from z_c, w.p. 1, when H enters she enters from B for a payoff of $\frac{1}{1+k}$. Since waiting requires positive delay by right-continuity of entry strategy CDFs, the expected discounted value of

H's payoff starting at z_c is strictly less than $\frac{1}{1+k}$, which she could receive immediately by entering, contradicting the optimality of her strategy at z_c in equilibrium.

Case II: $z_c > z^*$: By similar logic to the above, the belief process starting from z_c evolves continuously or with upward jumps with probability one, so by Lemma 2.7.4 the game ends in B with probability zero. Given type H, by waiting, it must reach A in finite time with positive probability because otherwise, by her weak preference for waiting at points in C and D, H would get 0 payoff and would strictly prefer immediate entry. This implies that there is a positive measure of news paths such that the same is true for the belief process given type L. Thus L gets strictly positive payoff by waiting for the belief to reach A, but by immediate entry she gets exactly 0, so she strictly prefers to wait, a contradiction.

Case III: $z_c = z^*$: Lemma 2.7.4 and the right-continuity condition on strategies imply that $\sup\{z : z \in B\} < z^*$. By Cases I and II with Lemma 2.7.5, this implies that by waiting, the belief process starting from z_c evolves continuously for either type until it reaches B or A, so it must reach A with positive probability. Thus L has strictly positive continuation value at z_c, which contradicts her indifference. □

2.7.2 Interval Equilibria

If an interval equilibrium is equivalent to $\Xi(\delta, \alpha)$ for $\delta = 0$, then the necessary condition (2.3.4) implies $\alpha = z^*$. The proof of Proposition 2.3.2 confirms that this is indeed an equilibrium. The rest of this section considers $\delta > 0$. Here I consider only the case $z^* > \alpha$; for the case $z^* = \alpha$, which has different boundary value conditions, see the supplementary appendix.

When $\alpha > z^*$, both players receive value 1 at α. Solving the system of equations

(2.3.7)–(2.3.12) gives the following values for the constants:

$$C_1^H = M_H \left[\exp\left(\frac{-\lambda - 1}{2}(\alpha - \delta) \right) - \frac{1}{1+k} \exp\left(\frac{-\lambda - 1}{2}\alpha \right) \right] \tag{2.7.1}$$

$$C_2^H = M_H \left[\exp\left(\frac{\lambda - 1}{2}\alpha \right) \frac{1}{1+k} - \exp\left(\frac{\lambda - 1}{2}(\alpha - \delta) \right) \right] \tag{2.7.2}$$

$$C_1^L = M_L \exp\left(\frac{-\lambda + 1}{2}(\alpha - \delta) \right) \tag{2.7.3}$$

$$C_2^L = -M_L \exp\left(\frac{\lambda + 1}{2}(\alpha - \delta) \right), \tag{2.7.4}$$

where $M_H = \frac{\exp(\frac{\lambda+1}{2}(\alpha+(\alpha-\delta)))}{\exp(\lambda\alpha)-\exp(\lambda(\alpha-\delta))} > 0$ and $M_L = \frac{\exp(\frac{\lambda-1}{2}(\alpha+(\alpha-\delta)))}{\exp(\lambda\alpha)-\exp(\lambda(\alpha-\delta))} > 0$.

The proofs of the next three lemmas are mainly algebraic and are contained in the supplementary appendix.

Lemma 2.7.7. *There exists a unique* $\delta^* \in \mathbb{R}_{++}$ *such that* $W_H(z; \delta, \delta)$ *is strictly increasing in* z *on* $[0, \delta]$ *if and only if* $\delta \leq \delta^*$.

Lemma 2.7.8. *If* $z > 0$ *and* $\delta > 0$, *then* $W_H(z; \delta, \delta)$ *is strictly decreasing in* δ.

Lemma 2.7.9. *There exists* $\delta^{IE} \in (\delta^*, \infty)$ *such that* $W_H(z; \delta, \delta) \geq 1 - k$ *for all* $z \in [0, \delta]$ *if and only if* $\delta \leq \delta^{IE}$.

For the case $\alpha = z^*$, similar arguments in the supplementary appendix yield the following:

Lemma 2.7.10. *There exists a unique threshold* $\delta_{z^*}^{IE}$ *such that* $\hat{W}_H(z; \delta, z^*) \geq 1 - k$ *for all* z *if and only if* $\delta \leq \delta_{z^*}^{IE}$.

2.7.3 NDVF Equilibria

Lemma 2.7.11. *Every NDVF equilibrium is equivalent to an interval equilibrium* $\Xi(\delta, \alpha)$ *with* $\delta = 0$ *or* $\alpha > z^*$.

Proof. Consider any equilibrium. Recall the definitions of belief categories A, B, C and D. By Lemma 2.7.6, the belief space must consist of only regions A, B and D. If D is empty, then the equilibrium is an interval equilibrium equivalent to $\Xi(0, z^*)$, so suppose D is nonempty. Consider any $z \in D$, and define $\bar{z}(z) := \inf\{z' \in [z, \infty) \cap (A \cup B)\}$ and $\underline{z}(z) := \sup\{z' \in (-\infty, z] \cap (A \cup B)\}$. Note that $\bar{z}, \underline{z} \in A \cup B$ by right-continuity of the entry CDF, so $z \in (\underline{z}, \bar{z})$ and V_H in $(\underline{z}, \bar{z})$ is the solution to (2.3.5) with some boundary conditions at $\underline{z}$ and $\bar{z}$. Since entry is in the support of L's strategy at $\bar{z}$ and $\underline{z}$, the posterior after entry from either of these points must be at least z^* by Lemma 2.7.2, so player 2 fights back with probability 0 or with probability $\frac{1}{1+k}$; that is, $f(j(\bar{z})), f(j(\underline{z})) \in \{0, \frac{1}{1+k}\}$. If $f(j(\bar{z})) = f(j(\underline{z})) = 0$, then V_H and V_L are strictly convex functions inside $(\underline{z}, \bar{z})$ taking value 1 at the end points, and thus are strictly decreasing near $\underline{z}$. Similarly, if $f(j(\bar{z})) = f(j(\underline{z})) = \frac{1}{1+k}$, then V_H is strictly convex and takes value $\frac{1}{1+k}$ at each end point, so again it is strictly decreasing near $\underline{z}$. Thus exactly one of the boundaries must lie in A and strictly above z^*, and the other must lie weakly below z^*. If D is not an interval, then there exists a $z' \in D$ with $z' \neq z$ such that there exists $z'' \in A \cup B$ with $z'' \in (\min\{z, z'\}, \max\{z, z'\})$. Hence it must be that either (i) $\underline{z}(z') \geq \bar{z}(z) > z^*$ or (ii) $\bar{z}(z') \leq \underline{z}(z) \leq z^*$, and in either case there is a contradiction. Thus D is an interval and by Lemmas 2.7.4 and 2.7.5, the equilibrium is an interval equilibrium. Since A does not intersect z^* when D is nonempty, the equilibrium must be an interval equilibrium equivalent to $\Xi(\delta, \alpha)$ with $\alpha > z^*$. ☐

Lemma 2.7.12. *In any interval equilibrium, V_L is nondecreasing.*

Proof. If the equilibrium is equivalent to $\Xi(\delta, \alpha)$ with $\alpha = z^*$, then $V_L = 0$ for all $z \in (\alpha - \delta, \alpha)$ and thus V_L is the same as under $\Xi(0, z^*)$. If instead $\alpha > z^*$, then V_L is defined by (2.3.8), (2.7.3) and (2.7.4). By inspection, W_L is strictly increasing on $[\alpha - \delta, \alpha]$, and by construction it takes values 0 and 1 at the lower and upper boundaries, so V_L is again nondecreasing. □

2.7.4 WBOP Equilibria

Lemma 2.7.13. *In any WBOP equilibrium, $D \succeq B$.*

Proof. By Lemma 2.7.4, B is nonempty. If D is empty, the claim is trivial, so assume D is nonempty and suppose there exist $z_1 \in B$, $z_2 \in D$ s.t. $z_1 > z_2$. By the definition of D, both players wait at z_2 and the right-limit of the posterior is almost surely z_2. Since H enters at $-\infty$, $z_2 > -\infty$. Then for almost all $(\omega_1', \omega_2') \in \Omega' \times \Omega'$ and $t_1, t_2 \geq 0$ such that $Z_{t_1}(\omega_1') = z_1$ and $Z_{t_2}(\omega_2') = z_2$, $Z_{t_2+}(\omega_2') = z_2 > -\infty = Z_{t_1+}(\omega_1')$, which fails WBOP. □

Lemma 2.7.14. *In any WBOP equilibrium, $A \setminus \{z^*\} \succeq D$.*

Proof. If D is empty, the claim is trivial, so assume it is not. By Lemma 2.7.5, $A \setminus \{z^*\}$ is nonempty. If $A \setminus \{z^*\} \not\succeq D$, there exists $z \in A \setminus \{z^*\}$ and $z' \in D$ such that $z' > z > z^*$. L weakly prefers to wait at points in D. If $z' \in D$, then WBOP implies $j(z') \geq j(z) = z > z^*$, so player 2 would concede after entry at z' and both players would strictly prefer to enter at z', a contradiction. □

In what follows, define $\alpha := \inf\{z \in A : z > z^*\}$ and $\beta := \sup\{z \in B : z < z^*\}$.

Lemma 2.7.15. *All WBOP equilibria are interval equilibria.*

Proof. In any WBOP equilibrium, C is empty by Lemma 2.7.6 so the only possible regions are A, B, and D. If $\alpha = \beta$, then by definition both must equal z^*. In this case, as argued in the proof of Lemma 2.7.6, $z^* \notin D$, so by combining this with Lemma 2.7.4, $z^* \in A$. This gives an equilibrium that entails $A = [z^*, \infty)$ and $B = (-\infty, z^*)$ which is of the desired structure. Now consider $\alpha > \beta$. As argued in the proof of Lemma 2.7.6, this implies $(\beta, \alpha) = D$. Recall that $\alpha \geq z^*$. By Lemma 2.7.4, $\alpha \notin B$, so $\alpha \in A$. It must be that $\alpha > z^*$ because otherwise H gets $\frac{1}{1+k}$ at both α and β, and by WBOP, at every $z \in (\beta, \alpha)$ as well; she would then strictly prefer entry to waiting inside D, contradicting the optimality of her strategy in D. By Lemma 2.7.4, either $z^* \in A$ or $z^* \in D$. In the first case, by Lemmas 2.7.13 and 2.7.14, the only possibility is $B = (-\infty, z^*)$ and $A = \{z^*\} \cup [\alpha, \infty)$ which is of the interval structure. In the second case, the only possibility is $A = [\alpha, \infty)$, and $B = (-\infty, \beta]$, which is also of the interval structure. □

Lemma 2.7.16. *In any WBOP equilibrium, $V_H(z) \geq \frac{1}{1+k}$ for all $z \in \mathbb{R}$.*

Proof. By Lemma 2.7.15, any WBOP equilibrium has the same strategy profile as some $\Xi(\delta, \alpha)$. If $\delta = 0$ the claim is trivial, so assume $\delta > 0$. The proof of Lemma 2.7.15 rules out $\alpha = z^*$, so consider $\alpha > z^*$. Then $V_H(z) = 1$ for all $z \in [\alpha, \infty]$ and $V_H(z) = \frac{1}{1+k}$ for all $z \in [-\infty, \alpha - \delta]$. If $z \in (\alpha - \delta, \alpha)$, H must receive at least $\frac{1}{1+k}$ by waiting because if she were to enter, $j(z) = z^*$ and thus player 2 would fight back with probability $\frac{1}{1+k}$ and giving H a value of $\frac{1}{1+k}$ in expectation. □

Corollary 2.7.1. *In any WBOP equilibrium, $\delta \leq \delta^*$, and the value functions V_θ are nondecreasing.*

Proof. Recall Lemma 2.7.12. If $\delta \leq \delta^{C_2^H}$, then the result follows from Lemma 2.7.7. If $\delta > \delta^{C_2^H}$ then V_H is strictly convex on $[\alpha - \delta, \alpha]$, and since Lemma 2.7.16 implies that $V_H'(\alpha - \delta) \geq 0$, V_H is nondecreasing. □

3

A Delegation-Based Theory of Expertise

3.1 Introduction

There are many situations in which a principal lacks the knowledge and expertise to perform a certain task, and therefore has to delegate the job to a qualified expert. Examples include a candidate running for office who has to hire an expert to work out her economic agenda, or the CEO of a pharmaceutical company who must delegate building a research and development division to a scientist. Further complicating the principal's situation is that experts tend to have systemic biases, preferring suboptimal actions from the principal's perspective.

In this paper we investigate a model in which a principal has to delegate a task to one of two imperfectly informed experts. The need to delegate differentiates our model from models of expertise in which experts send cheap talk recommendations to the principal, such as Krishna and Morgan (2001b). In particular, we consider the following game. First experts receive noisy and conditionally independent signals of a single dimensional state variable. The principal's ideal action is equal to the state, but each expert has a constant bias (either positive or negative) and a resulting ideal point different from the sender's. Next, the experts simultaneously propose actions. A proposal is assumed to bind the expert to perform the given action whenever the principal delegates the task to him.[1] The principal then chooses one of the two offers, and the corresponding action is taken by the given expert. We later extend the model so that experts benefit from being selected, on top of caring about the implemented action (the latter being independent of who takes the action). In particular, we allow for a bonus to the chosen expert, either as a monetary payment or as a non-monetary benefit, such as increased prestige in his profession. We investigate two cases, with the bonus amount given exogenously in one case and optimally chosen by the principal in the other.

Our model best applies to situations in which the principal lacks the knowledge to implement or initiate changes in the proposed actions, so all she can do is solicit different proposals and choose one of them. The assumption that proposals commit the experts correspond to common law, according to which an offer is a statement of terms on which the offeror is willing to be bound, and it shall become binding as soon as it is accepted by the person to whom it is addressed.[2] A different application for our model is political competition: starting from Downs (1957), most papers on political competition in a Hotelling (1929) framework assume that candidates are committed to the policies they announce in the campaign, and the electorate can only choose between the policies announced by the candidates.[3] In this context the bonus corresponds to the rents from being in office. Yet another application for our model is a setting where a legislative body (floor) seeks legislative proposals for the same bill from multiple committees, using a modified rule (see Gilligan and Krehbiel (1989),

[1]Even if the principal might not have the knowledge to verify whether the expert indeed chose the action that he proposed, outside experts might be able to verify if that was the case and hence penalties can be imposed on experts deviating from their proposals.

[2]See Treitel (1999), p8.

[3]For theoretical motivations for this assumption, and empirical relevance in the political competition context, see Pétry and Collette (2009), Kartik et al. (2015), and papers cited therein.

Krishna and Morgan (2001a)), meaning that the floor cannot amend the proposals and can only accept one of the proposed bills without modification, conforming to the basic assumptions of our model.[4]

We focus our investigation on situations in which the principal has limited knowledge of the decision environment. Besides imposing the constraint that she cannot perform the task herself, and hence she needs to delegate, we make two additional assumptions along these lines. The first one is that the principal's prior is improper uniform over a state space represented by the real line. The second one is that the principal lacks the ability to measure the difference between two proposed actions, and she can only ordinally compare them. This leaves only a few simple strategies available to the principal: always choose expert 1's offer (effectively delegating the action choice to expert 1), always choose expert 2's offer, always choose the minimum of the two offers, and always choose the maximum of the two offers. This type of coarse information processing is considered in the context of consumer choice problems in Kamenica (2008). The main question we address is how much a principal with such minimal knowledge of the environment can benefit from the presence of multiple experts.

We note that even though we restrict players' strategies in our game to the above simple strategies, in all equilibria we characterize, no player has a unilateral incentive to deviate, at any information set. That is, even if a player was allowed to play any strategy, she would stick to playing a simple strategy as long as the other players play simple strategies. Hence, the equilibria we characterize would remain equilibria in any sensible extension of the game in which players were allowed to choose from larger sets of strategies. In particular, this holds for the game in

[4]Gilligan and Krehbiel (1989) analyze this situation with an additional option to the floor, in the form of not accepting either of the proposals and opting for a status quo outcome. As opposed to our model, Gilligan and Krehbiel (1989) assume perfectly informed experts (committees), which fundamentally changes the strategic interaction.

which the principal's strategy is not restricted to be based on coarse information (the binary relationship between the two proposed actions): the equilibria in our model would remain equilibria with a more sophisticated principal, but particularly simple equilibria in which the principal only uses coarse information.[5]

In our game there always exists a trivial pure strategy equilibrium, in the form of delegating the task to one of the experts. Formally, one expert always proposing his ideal action conditional on the signal he observes (equal to the signal plus his bias), the other expert proposing his signal minus the first expert's bias, and the principal always delegating the task to the first expert constitutes a (Bayesian Nash) equilibrium.[6] The question is whether there exist other equilibria of the game, in which the principal either always chooses the minimum or always chooses the maximum of the two proposals (hence her choice depends nontrivially on the proposals). We assume without loss of generality that the sum of the biases of the two experts is nonnegative.

In our baseline model, there always exists an equilibrium in which the principal chooses the minimum of the two offers and the experts apply markups above their biases; if both biases are positive, this means that both experts exaggerate their biases. This result is contrary to a naive intuition that experts in competition should move toward the center. We call this equilibrium "upward," as the experts on average exaggerate their signals upwards in their proposals. This result extends to the case where an exogenous bonus is given to the selected expert, as long as the bonus is not too large. To illustrate, suppose that experts have the same biases and the bonus is small. Then in this upward equilibrium both experts propose actions strictly above their ideal actions based purely on their private signals. This is because, similarly

[5]Of course in the extended game there can be other, more complicated equilibria.

[6]There are other Bayesian Nash equilibria on mixed strategies with the same outcome, in which one expert always proposes his ideal point, the other expert "babbles" (randomizes over possible messages he can send), and the principal always delegates the task to the first expert.

to the winner's curse phenomenon in common value auctions, being selected by the principal contains information on the other expert's signal (namely that his signal is higher), changing the optimal action of the expert. In equilibrium, proposals have to be optimal conditional on the event that the other expert's action proposal is higher.

We also show that if the experts' signals are noisy enough then, for a subset of the range of bonuses for which an upward equilibrium exists, there also exists a "downward" equilibrium in which experts propose actions on average *below* their signals, and the maximum of the two proposals is selected by the principal. The strategic forces are similar to those in upward equilibrium: the fact that the maximum of the two offers is selected pushes proposals downwards, and for noisy signals on average experts modify their proposals downward relative to their signals.[7] This type of equilibrium does not exist when the signals are very precise, because then the information conveyed from being selected does not shift the optimal proposals of the experts enough to make markups negative on average. We show that even when the downward equilibrium exists, the principal prefers the upward equilibrium to it, and thus for further welfare comparisons we need only consider upward equilibrium and simple delegation to the less-biased expert.

The feature of the above equilibria that similarly biased experts exaggerate in a particular direction (and hence the principal should choose the proposal least in the direction of the exaggeration) is in line with empirical evidence. For example, Zitzewitz (2001), Bernhardt et al. (2006) and Chen and Jiang (2006) find that financial analysts systematically exaggerate their forecasts relative to unbiased forecasts based on the analysts' information sets, while Iezzoni et al. (2012) report that 55% of doctors in a survey said that in the previous year they had been more positive about patients' prognoses than their medical histories warranted.

[7] There can also be an equilibrium in which the principal mixes with a particular probability between accepting the lower or the higher proposal. We provide a partial characterization of such mixed equilibria in the Supplementary Appendix.

We compare the principal's welfare between upward equilibrium and simple delegation in order to find the principal-optimal equilibrium. In general, the comparison is complicated and can go either way, but for several focal cases of parameter settings, the optimum is indeed the upward equilibrium. When one expert's bias is positive and the other's is zero, the principal prefers the upward equilibrium to simple delegation; this is despite the fact that under simple delegation the unbiased expert's incentives are perfectly aligned with the principal's. The intuition for the result is that in the upward equilibrium, the principal extracts some additional information the second expert, which reduces the variance of the chosen action, and this benefit always outweighs any cost associated with higher markups. In addition, the principal is always better off in the upward equilibrium when the experts have the same biases (as in settings where experts have similar agendas), or when the experts have exactly opposite biases. This result holds even when the bonus is zero and hence there is no competition among experts for being selected.

When the experts have opposite biases of sufficiently large magnitude, we show that the principal can improve upon the upward equilibrium by commitment to an "element of surprise," introducing a small probability of choosing the *higher* offer. In the upward equilibrium, the expert with the positive bias faces a very large winner's curse, and applies a markup well above his bias. The other expert faces an almost negligible winner's curse and applies a markup just slightly above his bias. Now by threatening to choose the higher offer some of the time, the principal induces the first expert to reduce his markup drastically and the second expert raises his markup; both markups move closer to zero. The cost of this deviation to the principal lies in mistakenly choosing the higher offer, but when the magnitude of the biases is sufficiently large, the benefit outweighs the cost.

Applied to a political setting, the results in the two preceding paragraphs offer an alternative explanation for known phenomena. First, the comparison between upward equilibrium and simple delegation helps explain what voters may otherwise perceive as corruption – a politician may want to seek a second opinion from a biased expert, even if it is common knowledge that she already has access to an unbiased expert. Second, that the principal can benefit from committing to a mixed strategy in certain situations is consistent with an observed pattern of regulatory uncertainty, which voters may otherwise perceive as erratic behavior. Ederer et al. (2014) show similarly that commitment to an opaque reward scheme reduces temptation to game the system in a principal-agent environment.

We also compare the principal's payoffs when experts have equal versus opposite biases, and our result contrasts sharply with some of the existing literature. In our model, assuming the upward equilibrium is played, having two experts with identical biases yields a higher payoff than having two antagonist experts with opposite biases. In general, the expected bias of the implemented action is smaller with antagonist experts than with experts having the same bias, but this benefit is outweighed by a higher variance of the implemented action that arises because the expert with the lower bias is selected most of the time, and so the information from the other expert's signal is only utilized to a limited extent. This result contrasts models of competition in persuasion (Milgrom and Roberts (1986), Gentzkow and Kamenica (2015)), in which antagonist experts benefit the principal by pressing each other to reveal more information,[8] and with the multi-sender cheap talk model of Krishna and Morgan (2001b), in which having a second sender with the same bias does not benefit the receiver.[9]

The principal's expected payoff depends in a complicated way on the noise in

[8]Experts with identical agendas can be better for the principal than experts with opposing agendas in the persuasion model of Bhattacharya and Mukherjee (2013). The mechanism is rather different than in our paper, though: with similar experts an undesirable default action can provide strong incentives for both experts to reveal information.

[9]See also Shin (1998) and Dewatripont and Tirole (1999) for different types of models making the case for adversarial procedures.

the experts' signals, and on the amount of the bonus. Hence, for these comparative statics we focus on the case of equally biased experts. Even in this case, the effect of the variance of the experts' signals is ambiguous. An increased precision of experts' signals reduces the variance of the implemented action conditional on the state. For small bonuses, this unambiguously increases the principal's expected payoff. However, for larger bonuses, it might benefit the principal in the upward equilibrium if the experts increase their markups,[10] which can result from increasing the variance of the signals. We provide an exact characterization (for equally biased experts) for when a decrease in the variance of experts' signals benefits the principal.

Increasing the bonus reduces the absolute values of the experts' markups, hence bringing their proposals closer to truthful reporting, both in the upward and downward equilibria. Intuitively, a higher bonus increases competition among experts, leading them to decrease their proposals in the upward equilibrium and increase their proposals in the downward equilibrium. In the upward equilibrium, this initially improves the principal's expected payoff by decreasing the expected bias of the implemented action. There is a threshold level of bonus though at which the expected bias of the implemented action becomes zero, and increasing the bonus above this threshold decreases the principal's payoff. When the bonus comes from exogenous sources, the optimal bonus from the principal's perspective is always strictly positive, and is on the interior of the interval of bonuses for which the upward equilibrium exists. When the bonus is paid by the principal, the optimal bonus amount is always strictly smaller than in the previous case, and depending on the parameters it can be either strictly positive or zero.

In the political competition application of the model the result implies that a

[10] This is related to the chunkiness of the principal's possible choices in our model: for certain parameter values sticking with choosing the minimum of two proposed actions is still optimal for the principal, even though it leads to the implemented action being negatively biased. This can happen if choosing the maximum offer would lead to an even larger positive bias. These are the cases when an increase in the expectation of the minimum offer benefits the principal.

small amount of office-seeking motivation can be beneficial for voters, but at higher levels a further increase in office-seek motivation can adversely affect voters' welfare.

We consider two extensions of our model, for equally biased experts. In the first one we allow the principal to commit ex ante to any mixture of simple strategies, and show that for bonuses that are not too large, such commitment leads to the same outcome as in the upward equilibrium of the original game, hence the ability to commit does not improve the principal's welfare. In the second extension we drop the dependence of the unselected expert's payoff on the implemented action, and instead assume that the expert gets a fixed outside option payoff. This variant of the model is more realistic in market transaction situations, such as when experts are car mechanics or doctors. A car mechanic might be biased towards larger repairs than necessary, but typically he does not care about what type of repair is chosen in case a different mechanic is selected to do the job. The analysis of this version of the model is more involved, but we show that under some parameter restrictions similar upward and downward equilibria exist as in the baseline model. The fact that the unselected expert gets a fixed outside payment increases the experts' proposals in upward equilibrium, and decreases them in downward equilibrium.

Finally, a methodological contribution of our paper is that, in contrast with existing game theoretical models of improper prior (Friedman (1991), Klemperer (1999), Morris and Shin (2002, 2003), Myatt and Wallace (2014)), we provide an alternative to the standard approach of leaving expected payoffs of the game undefined and only interpreting payoffs conditional on signal realizations. In particular, we characterize the maximal set of expert strategies in our game in which the concept of ex ante expected payoffs can be extended in a rigorous sense, despite the improper prior on the state space: strategies for which the limit payoffs of the players for any possible sequence of proper priors diffusing (converging in a formal sense to the improper prior) are always the same. We can then assign these limit payoffs as ex ante expected payoffs under an improper uniform prior, and as in standard games, have well-defined expected payoffs for every strategy profile. In our game the largest set of expert strategies for which the limit expected payoffs are well-defined, given the set of simple strategies of the principal, are equivalent to strategies in which the difference between the proposed action and the observed signal, which we label as *markup*, is constant over all possible signals.[11] We define the set of strategies for the experts to be such constant markup strategies.

The literature on delegation so far mainly focused on either the question of delegating the action choice versus retaining the right to take the action (Dessein (2002), Li and Suen (2004)) or on optimally constraining the action choices of a particular expert (Holmström (1977), Melumad and Shibano (1991), Alonso and Matouschek (2008)). Krishna and Morgan (2008) investigate how monetary incentives can be used optimally in delegation to a single agent. The first paper in a delegation framework with asymmetric information we are aware of that considers competing agents is Kartik et al. (2015), in the context of the classic Downs (1957) model of political competition. Similarly to our setting, politicians receive independent private signals about the state of the world and hence the optimal policy from the electorate's point of view. The main difference between the models is that the politicians in Kartik et al. (2015) do not have policy preferences, and they are purely office-motivated. For this reason neither the own private information nor the rival's private information directly affects their expected payoffs, and the candidates play a zero-sum game. In contrast, in our model the experts' signals are directly payoff-relevant for them, and their interests are partially aligned, as in higher states they would both like to induce higher actions.[12] This leads to different equilibrium dynamics than in Kartik

[11] More precisely, mixtures of such constant markup strategies.

[12] For classic papers in political science literature assuming that politicians have mixed motivation (having both policy preferences and wanting to win), as in our model, see Wittman (1983) and Calvert (1985).

et al. (2015), and to distinct conclusions: in particular, in their model the electorate can never strictly benefit from the presence of a second candidate (relative to just a single one).[13]

Outside the delegation literature, Gerardi et al. (2009) investigates aggregation of expert opinions through a particular mechanism that approximates the first best outcome if signals are very accurate. Pesendorfer and Wolinsky (2003) investigate the effects of being able to solicit a second opinion from a different expert, in a dynamic model in which experts are not biased but it is costly for them to gather information. Another line of literature investigates multi-sender extensions of the cheap talk model of Crawford and Sobel (1982), and finds that under certain conditions there can be equilibria in which the receiver can extract full or almost full information from the senders (Gilligan and Krehbiel (1989), Austen-Smith (1993), Wolinsky (2002), Battaglini (2002, 2004), Ambrus and Takahashi (2008), Ambrus and Lu (2014)). As opposed to the above papers, we investigate settings in which the principal cannot solicit information from experts and then take the action choice herself. Lastly, Ottaviani and Sørensen (2006) consider a model with multiple experts with reputational concerns reporting sequentially on privately observed signals. The issues they focus on (potential herding behavior of experts) are very different than in the current paper.

The coarse information processing on the part of the principal in our model can be considered as a particular instance of analogy classes, proposed in Jehiel (2005) and

[13] Our model approximates the model in Kartik et al. (2015) when the bonus payment is very large and so the agents mainly care about being selected. We find that for very large bonuses the only equilibria in our model involve delegating the action choice to a single agent, which is in line with the result on maximum informativeness of political competition in Kartik et al. (2015). Correspondingly, Kartik et al. (2015) discuss an extension of their model in which they show that allowing a small amount of ideological motivation for the candidates, and assuming that they are close to unbiased from the electorate's point of view implies that in equilibrium one candidate must be winning ex ante with probability close to 1. These results suggest that there is no discontinuity between no policy preference versus a small amount of policy preference for the agents. Our paper mainly focuses on cases in which agents' policy preferences are relatively important.

Jehiel and Koessler (2008).[14] In particular, the experts' proposals can be assigned to three analogy classes: expert 1's proposed action is higher, expert 2's proposed action is higher, and the two proposed actions are equal. As opposed to the examples these papers highlight, the equilibria we characterize in our model would remain equilibria even if the principal could condition her choice on the exact proposals.[15] Hence the outcomes we characterize can also be interpreted as equilibria in which the principal processes information from the proposals in a coarse way, even though she is not restricted to do so.

3.2 Model

We consider the following multi-stage game with incomplete information. There are three players: a principal and two experts. The set of states of the world is $\mathbb{R}$, and we assume that the common prior distribution of states is diffuse (improper uniform).

In stage 0 state $\theta \in \mathbb{R}$ realizes. In stage 1 each expert $i = 1, 2$ receives a noisy private signal about the state of the world $s_i = \theta + \epsilon_i$, where $\epsilon_i \sim N(0, \sigma^2)$, and ϵ_1 and ϵ_2 are independent. In stage 2 each expert i proposes an action $a_i \in \mathbb{R}$ to the principal. In stage 3 the principal chooses one of the two experts, who then implements the action he proposed in stage 2.

Let real-valued functions $a_1(s_1)$ and $a_2(s_2)$ denote the strategies of expert 1 and expert 2 respectively, while $C(a_1, a_2)$ is the principal's choice strategy. If action $a = a_i$ is taken then the principal receives payoff $V(a, \theta) = -(a - \theta)^2$, expert $j = 1, 2$ receives $-(a - \theta - b_j)^2$. In Section 3.4, we extend the model to include a bonus payment B to the chosen expert. We call b_i the *bias* of expert i. Without loss of

[14] Similar concepts have been introduced in Eyster and Rabin (2005) and in Esponda (2008).

[15] With finer information processing on the part of the principal, there could be other pure strategy equilibria, besides the simple ones characterized in this paper.

generality, we assume that $b_1 + b_2 \geq 0$ and $b_1 \geq b_2$. We further assume that all parameters of the game are common knowledge.

We are interested in situations in which the principal is unable to comprehend/measure the exact values of a_1 and a_2, she can only make a binary comparison and tell which proposed action is higher. Hence, we restrict the principal's strategies to the following simple strategies: (i) $C(a_1, a_2) = 1$ for any $a_1, a_2 \in \mathbb{R}$ (always choose expert 1); (ii) $C(a_1, a_2) = 2$ for any $a_1, a_2 \in \mathbb{R}$ (always choose expert 2); (iii) $C(a_1, a_2) = 1$ if $a_1 < a_2$, $C(a_1, a_2) = 2$ if $a_2 < a_1$, and $C(a_1, a_2) \in \{1, 2\}$ if $a_1 = a_2$ (always choose the expert with the smaller proposal); and (iv) $C(a_1, a_2) = 1$ if $a_1 > a_2$, $C(a_1, a_2) = 2$ if $a_2 > a_1$, and $C(a_1, a_2) \in \{1, 2\}$ if $a_1 = a_2$ (always choose expert with the higher proposal).

Define a markup of an expert's strategy as a difference between his proposed action and his signal: $k_i(s_i) = a_i(s_i) - s_i$. We restrict the experts' strategy sets, to constant markup/markdown strategies: $a_i(s_i) = s_i + k_i$, where $k_1, k_2 \in \mathbb{R}$. We motivate this restriction in the next section, where we show that given the set of strategies of the expert defined above, constant markup/markdown strategies are essentially the only strategies in our game for which the concept of ex ante expected payoff can be extended to in a formal way. With slight abuse of notation, we use simply $(k_1, k_2, C(a_1, a_2))$ to denote a strategy profile with constant markup strategies.

We use Bayesian Nash equilibrium (from now on referred to as equilibrium) as our solution concept, and in the main text focus attention on pure strategy equilibria. Given the restrictions we impose on strategies, all the equilibria we characterize in the paper are also perfect Bayesian Nash equilibria. We also note that in all equilibria we characterize, the action choice prescribed by the equilibrium strategy is sequentially rational at any history, even if the restrictions on the set of possible strategies is lifted. That is, the characterized equilibria remain perfect Bayesian Nash equilibria even when strategies are not restricted to be simple (for the principal) and constant markup/markdown (for the experts).

3.3 Equilibria in Pure Strategies

In this section we characterize all possible pure strategy equilibria in the delegation game. We start with determining the best response of the principal to all possible pairs of constant markup strategies by the experts, and then investigate the candidate pure strategy equilibria consistent with this best response behavior.

3.3.1 Best Response to Constant Markup Strategies

Here we analyze the principal's best response to constant markup strategies, and find that it only depends on the sum of markups. As the principal has a quadratic loss function, her expected payoff can be decomposed into losses from the uncertainty about the true state (which is independent of her action) and the losses from the expected difference between the chosen action and the true state. Therefore the principal prefers the offer which is closer to her posterior expectation of the true state. After observing the offers, the principal's expectation about the true state is lower (higher) than the average of the experts' offers if and only if the sum of the markups is positive (negative). Figure 3.1 illustrates a case where the markups have positive sum.

Let $\arg\min\{a_1, a_2\}$ be defined as $\{1\}$ if $a_1 < a_2$, $\{2\}$ if $a_1 > a_2$, and $\{1, 2\}$ if $a_1 = a_2$. Similarly, let $\arg\max\{a_1, a_2\}$ be defined as $\{1\}$ if $a_1 > a_2$, $\{2\}$ if $a_1 < a_2$, and $\{1, 2\}$ if $a_1 = a_2$.

Theorem 3.3.1. *If experts follow constant markup strategies $a_i(s_i) = s_i + k_i$, then*

- *if $k_1 + k_2 > 0$, the principal strictly prefers the lower offer, and $C(a_1, a_2) \in \arg\min\{a_1, a_2\}$;*

- *if* $k_1 + k_2 < 0$, *the principal strictly prefers the higher offer, and* $C(a_1, a_2) \in \arg\max\{a_1, a_2\}$*;*

- *if* $k_1 + k_2 = 0$, *the principal is indifferent between the offers.*

Proof. After observing both offers, the principal updates her belief: $\theta | a_1, a_2 \sim N(\frac{1}{2}(a_1+a_2) - \frac{1}{2}(k_1+k_2), \frac{\sigma^2}{2})$. Therefore the principal's expected utility from choosing offer a_1 or a_2 is:

$$V(a_1) = \mathbb{E}[-(\theta - a_1)^2] = -\text{Var}(\theta) - (\mathbb{E}[\theta] - a_1)^2 = -\frac{\sigma^2}{2} - \left[\frac{1}{2}(a_2 - a_1) - \frac{1}{2}(k_1 + k_2)\right]^2$$

$$V(a_2) = \mathbb{E}[-(\theta - a_2)^2] = -\text{Var}(\theta) - (\mathbb{E}[\theta] - a_2)^2 = -\frac{\sigma^2}{2} - \left[\frac{1}{2}(a_1 - a_2) - \frac{1}{2}(k_1 + k_2)\right]^2.$$

Hence, $V(a_1) - V(a_2) = (a_2 - a_1)(k_1 + k_2)$, which immediately implies the statements in the theorem. □

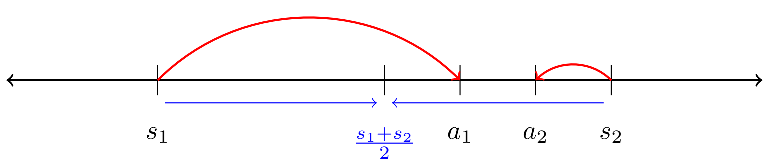

FIGURE 3.1: $k_1 > 0, k_2 < 0, k_1 + k_2 > 0$. The principal chooses the lower offer a_1, which lies closer to her expectation $\frac{s_1+s_2}{2}$.

An equilibrium $(k_1, k_2, C(a_1, a_2))$ is said to be an *upward equilibrium*, if on average experts adjust their signals upwards and the lower proposal is accepted: $k_1 + k_2 \geq 0$ and $C(a_1, a_2) \in \arg\min(a_1, a_2)$. In upward equilibrium, the principal's updated expectation of the state of the world is lower than the average of the two offers. Her best response is to choose the lower offer, which is closer to her expectation, as demonstrated in Figure 1. Likewise, an equilibrium $(k_1, k_2, C(a_1, a_2))$ is said to be a *downward equilibrium* if $k_1 + k_2 \leq 0$ and $C(a_1, a_2) \in \arg\max(a_1, a_2)$. Note that when $k_1 + k_2 = 0$ then the principal's posterior expectation of θ is exactly the average of the two offers, and she is indifferent between the two. This raises the possibility of equilibria in which the experts play constant markup strategies and the expert mixes between the lower and the higher offer with some fixed probability. We investigate such equilibria in the Supplementary Appendix.

3.3.2 Simple delegation

In our game there always exist simple pure strategy equilibria in which the principal always chooses the same expert, independently of two offers, in effect delegating the decision to her. In particular, Theorem 3.3.1 implies that if expert i chooses constant markup b_i and the other expert chooses constant markup $-b_i$ then the principal is always indifferent between the two offers, and she might as well always choose expert i. Given this strategy of the principal, expert i's best response is choosing exactly markup b_i, which in expectation implements his ideal action. The other expert has no profitable deviation since his proposal is never accepted. While such an equilibrium exists for each of the two experts, it is more natural to consider the one in which the principal always chooses the expert with the smaller absolute bias, who is expert 2 by convention. These observations are summarized in the next proposition.

Proposition 3.3.1. *For $i \in \{1, 2\}$, an equilibrium exists in which the principal always chooses expert i and markups are $k_i = b_i$ and $k_j = -b_i$ for $j \neq i$. The principal's expected payoff in this equilibrium is $\mathbb{E}[-(s + b_i)^2] = -\sigma^2 - b_i^2$, where $s \sim N(0, \sigma^2)$.*

In the rest of the section we examine equilibria in which the principal's choice between the experts depends in a nontrivial way on the pair of offers proposed.

3.3.3 *Upward equilibrium*

Here we investigate strategy profiles $\{(k_1, k_2, a \in \arg\min\{a_1, a_2\}) : k_1 + k_2 \geq 0\}$. We start by computing players' payoffs under such strategy profiles. Let $b(k_1, k_2, L) = \mathbb{E}(a-\theta)$ denote the expected bias of the chosen offer and $Var(k_1, k_2, L) = \text{Var}(a-\theta)$ denote the variance of the chosen offer. Then the principal's utility is: $V(k_1, k_2, L) = -\mathbb{E}(a-\theta)^2 = -b^2(k_1, k_2, L) - Var(k_1, k_2, L)$. Proposition 3.3.2 below provides the expanded forms of these expressions, which are useful to our analysis. Here and throughout the rest of the paper, let f and F denote the PDF and the CDF of the distribution $N(0, 2\sigma^2)$ and let $z = k_1 - k_2$. Note that the expected bias $b(k_1, k_2, L)$ is strictly less than the expected value of the selected markup, $k_1(1-F(z)) + k_2 F(z)$; this is because the lower offer is associated with a noise term which is normally distributed but truncated above and thus has negative expectation.

Proposition 3.3.2. *If both experts follow constant markup strategies $a_j(s_j) = s_j + k_j$ and the principal always chooses the lower offer, then*

$$\begin{aligned}
b(k_1, k_2, L) &= -2\sigma^2 f(z) + k_1(1-F(z)) + k_2 F(z);\\
Var(k_1, k_2, L) &= \sigma^2 - 4\sigma^4 f^2(z) - 2\sigma^2 z f(z)(2F(z)-1) + z^2 F(z)(1-F(z));\\
V(k_1, k_2, L) &= -\sigma^2 + 2(k_1+k_2)\sigma^2 f(z) - k_1^2(1-F(z)) - k_2^2 F(z);\\
U_i(k_1, k_2, L) &= -\sigma^2 + 2\sigma^2(k_i + k_j - 2b_i) f(z) - (k_j - b_i)^2 F(k_i - k_j)\\
&\quad - (k_i - b_i)^2 F(k_j - k_i).
\end{aligned}$$

Denote the hazard rate $\frac{f(x)}{1-F(x)}$ by $v(x)$ and let $w(x) := \frac{f(x)}{F(x)}$. The hazard rate plays an important role in our analysis. It represents the instantaneous probability that experts' signals differ by x, conditional on differing by at least x.

In upward equilibrium, the principal's choice of the lower offer implies that offers affected by negative realizations of noise are accepted more frequently. Therefore conditional on being chosen, an expert must revise his belief about θ upwards. This induces experts to increase their markups, similarly to bidders shading their bids downwards in a common value auction environment.[16]

Let z^* denote the unique solution[17] to the equation

$$z - \sigma^* \left[v(z) - w(z)\right] = b_1 - b_2. \tag{3.3.1}$$

Theorem 3.3.2. *There exists a unique upward equilibrium, characterized by markups $k_1^U = b_1 + \sigma^2 v(z^*)$ and $k_2^U = b_2 + \sigma^2 w(z^*)$ with $k_1^U - k_2^U = z^* \geq b_1 - b_2 \geq 0$.*

Notice that the equilibrium markup difference z^* as well as $\text{Var}(a - \theta)$ depend on the biases of the experts only through $b_1 - b_2$. In upward equilibrium expert 1 wins the bonus with probability $1 - F(z^*)$, which is reflected in the expression for b_U. The bonus affects the expected bias in two ways: through the change in experts' probabilities of winning and through the change in markups. In the appendix (Corollary 3.8.1), we give the full expansion of the players' utilities in upward equilibrium.

3.3.4 Downward equilibrium

Recall that an equilibrium is a downward equilibrium when it belongs to $\{(k_1, k_2, a \in \arg\max(a_1, a_2)) : k_1 + k_2 \leq 0\}$. Proposition 3.3.3 is analogous to Proposition 3.3.2. Note that $b(k_1, k_2, H)$ now exceeds the expected chosen markup because it is associated with a normally distributed noise term truncated from below.

Proposition 3.3.3. *If both experts follow constant markup strategies $a_j(s_j) = s_j + k_j$*

[16] A similar shading behavior emerges in the welfare-maximizing (first best) outcome in Kartik et al. (2015), as interestingly the latter outcome implies always selecting the politician with the larger signal in absolute terms.

[17] That there is a unique solution is shown in the proof of Theorem 3.3.2*.

and the principal always chooses the higher offer, then

$$\begin{aligned} b(k_1,k_2,H) &= k_1+k_2-b(k_1,k_2,L); \\ Var(k_1,k_2,H) &= Var(k_1,k_2,L); \\ V(k_1,k_2,H) &= V(-k_1,-k_2,L); \\ U_i(k_1,k_2,H) &= -\sigma^2-2\sigma^2(k_i+k_j-2b_i)f(z)-(k_i-b_i)^2F(k_i-k_j) \\ &\quad -(k_j-b_i)^2F(k_j-k_i). \end{aligned}$$

In downward equilibrium, the principal's strategy of choosing the higher offer implies that, conditional on being chosen, an expert must revise his belief and his markup downward. This downward force must be sufficiently large to ensure that the sum of markups is negative, so that the principal's choice of the higher offer is a best response. Hence, noise must be sufficiently large for downward equilibrium to exist.

Theorem 3.3.3. *A downward equilibrium exists if and only if* $b_1+b_2 \le \sigma^2(v(z^*)+w(z^*))$. *When it exists, it is unique and characterized by* $k_1^D = b_1-\sigma^2 w(z^*)$ *and* $k_2^D = b_2-\sigma^2 v(z^*)$, *with* $k_1^D-k_2^D = z^* \ge b_1-b_2$.

3.3.5 Principal-Optimal Equilibrium

In this section we compare the principal's expected utility in upward and downward equilibria, and in the case of simple delegation. We also investigate how the principal's expected payoff in equilibrium depends on the biases of the experts.

Note that the equilibrium markup difference z^* and $\text{Var}(a-\theta)$ stay the same as in upward equilibrium. In downward equilibrium expert 1 wins the bonus with probability $F(z^*)$, which is higher than $1-F(z^*)$ in upward equilibrium, and this shifts the expected bias higher. Furthermore, in downward equilibrium the bonus motivates experts to increase their markups, as opposed to upward equilibrium in which the bonus motivates experts to decrease their markups. As a result of the above effects, the expected bias is higher in downward equilibrium. In Corollary 3.8.2 we provide expressions for the players' utilities.

Proposition 3.3.4. *The principal prefers upward equilibrium to downward equilibrium whenever both exist.*

Given the above result, we now compare the principal's utility in upward equilibrium and simple delegation to the expert with the smaller absolute bias. In general this comparison is complicated, but in the next proposition we show that for upward equilibrium to be better for the principal, it has to be the case that either both experts have high enough biases, or that they are equally biased.

Proposition 3.3.5. *Parameterize biases as $b_1 = b + x$ and $b_2 = b - x$ for some $x, b \geq 0$. For all $x > 0$, there exists a threshold $\bar{b} > 0$ such that the principal prefers simple delegation to the upward equilibrium if and only if $b_1 > \bar{b}$. The principal always prefers upward equilibrium in the following cases: (i) biases are equal, i.e. $x = 0$, and (ii) $b_2 = 0$.*

For intuition behind the above result, first consider the case $x > 0$. The variance of the chosen action in upward equilibrium is always lower than that in simple delegation, and both are independent of b. As b increases, the expected bias $b-(2F(z^*)-1)x$ in upward equilibrium and $b - x$ in downward equilibrium increase at the same rate. Since the former is higher than the latter, and losses are quadratic, this increase hurts more in upward equilibrium than in simple delegation. Once b is high enough, this disadvantage outweighs the initial advantage of lower variance.

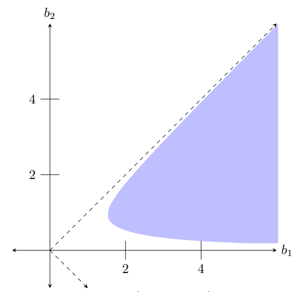

Figure 3.2: Depending on the biases (b_1, b_2), the principal's strategy under commitment is either simple delegation to expert 2 (shaded region) or choosing the lower offer (unshaded).

3.4 Extensions

3.4.1 Bonus Payments

In this extension, the chosen expert receives a bonus payment $B > 0$ in addition to his quadratic loss. Experts now have two contradictory incentives. To illustrate, recall the upward equilibrium, in which the principal always chooses the lower offer and the resulting winner's curse phenomenon exerts an upward force on experts' offers, leading them to apply markups above their biases. Introducing a positive bonus $B > 0$ induces experts to decrease their markups in order to be selected more frequently. Hence, if the bonus is small, the first force prevails, and the expert sets a markup higher than her bias. If the bonus is large then the second force prevails, and the expert sets a markup lower than his bias. These conclusions are confirmed by formulas in Theorem 3.3.2*. If $B < 2\sigma^2$, then the markups are higher than the corresponding biases, and if $B > 2\sigma^2$, the situation reverses. When the bonus becomes sufficiently high, markups become lower than the biases to an extent that the sum of markups is negative. Then the principal prefers to choose the higher offer. Therefore, for very high bonuses upward equilibria do not exist.

Let $\rho := \sigma^2 - \frac{B}{2}$. Equation (3.3.1) extends to

$$z - \sigma^* \left[v(z) - w(z) \right] = b_1 - b_2. \tag{3.4.1}$$

Theorem 3.3.2*. *There exists a threshold* $B_U > 0$ *such that an upward equilibrium exists if and only if* $B \leq B_U$. *When it exists, it is unique and characterized by markups* $k_1^U = b_1 + \rho v(z^*)$ *and* $k_2^U = b_2 + \rho w(z^*)$ *with* $k_1^U - k_2^U = z^* \geq 0$. *Moreover,* B_U *lies in the interval* $[2\sigma^2 + 2\sqrt{\pi}\sigma \max(0, b_2), 2\sigma^2 + \sqrt{\pi}\sigma(b_1 + b_2)]$. *For* $B \leq B_U$,

- $z^* \geq b_1 - b_2 \iff B \leq 2\sigma^2 \iff \rho \geq 0$;
- $b(k_1^U, k_2^U, L) = b_U := b_1(1 - F(z^*)) + b_2 F(z^*) - B f(z^*)$;
- $Var(k_1^U, k_2^U, L) = \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1) + (z^*)^2 F(z^*)(1 - F(z^*))$.

In the case of equally biased experts, we can obtain a closed form solution for strategies in upward equilibrium.

Proposition 3.4.1. *Consider* $b_1 = b_2 = b > 0$. *Then* $B_U = 2\sigma^2 + 2\sqrt{\pi}\sigma b$, *and in upward equilibrium, we have the following:*

- $k_1^U = k_2^U = k_U = b + \frac{\rho}{\sigma\sqrt{\pi}}$;
- $b(k_U, k_U, L) = b - \frac{B}{2\sqrt{\pi}\sigma}$;
- $Var(k_U, k_U, L) = \left(1 - \frac{1}{\pi}\right)\sigma^2$;
- $V(k_U, k_U, L) = -\left(b - \frac{B}{2\sqrt{\pi}\sigma}\right)^2 - \sigma^2 + \frac{\sigma^2}{\pi}$;
- $U_i(k_U, k_U, L) = -\rho + \frac{\sigma^2}{\pi} - \frac{B^2}{4\pi\sigma^2}$ *for* $i = 1, 2$.

The next theorem characterizes downward equilibrium. Note that if $B_D < 0$, no downward equilibrium exists. Recall the definition of z^* from (3.4.1).

Theorem 3.3.3*. *There exists a threshold* B_D *such that a downward equilibrium exists if and only if* $B \leq B_D$. *When it exists, it is unique and characterized by*

$k_1^D = b_1 - \rho w(z^*)$ *and* $k_2^D = b_2 - \rho v(z^*)$*, with* $k_1^D - k_2^D = z^* \geq b_1 - b_2$ *and* $B_D \in [2\sigma^2 - \sqrt{\pi}\sigma(b_1 + b_2), 2\sigma^2 - 2\sqrt{\pi}\sigma \max(0, b_2)]$*. For* $B \leq B_D$*,*

- $b(k_1^D, k_2^D, H) = b_D = b_1 F(z^*) + b_2(1 - F(z^*)) + Bf(z^*)$*;*
- $Var(k_1^D, k_2^D, H) = \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1) + (z^*)^2 F(z^*)(1 - F(z^*))$*.*

The following corollaries are immediate from Theorems 3.3.2* and 3.3.3*.

Corollary 3.4.1. *A downward equilibrium exists only if an upward equilibrium exists; that is,* $B_D \leq B_U$*.*

Corollary 3.4.2. *For* $B \leq B_D$*,* $k_1^U - b_1 = b_2 - k_2^D$ *and* $k_2^U - b_2 = b_1 - k_1^D$*.*

As before, we can obtain a closed form solution for the case of equally biased experts.

Proposition 3.4.2. *Consider* $b_1 = b_2 = b > 0$*. Then* $B_D = 2\sigma^2 - 2\sqrt{\pi}\sigma b$*, and in downward equilibrium, we have the following:*

- $k_1^D = k_2^D = k_D = b - \frac{\rho}{\sigma\sqrt{\pi}}$*;*
- $b(k_D, k_D, H) = b + \frac{B}{2\sqrt{\pi}\sigma}$*;*
- $Var(k_D, k_D, H) = \left(1 - \frac{1}{\pi}\right)\sigma^2$*;*
- $V(k_D, k_D, H) = -\left(b + \frac{B}{2\sqrt{\pi}\sigma}\right)^2 - \sigma^2 + \frac{\sigma^2}{\pi}$*;*
- $U_i(k_D, k_D, H) = -\rho + \frac{\sigma^2}{\pi} - \frac{B^2}{4\pi\sigma^2}$ *for* $i = 1, 2$*.*

First we establish that whenever both upward and downward equilibria exist, the principal always prefers the former.

Proposition 3.3.4*. *For any fixed $B \geq 0$, the principal prefers upward equilibrium to downward equilibrium whenever both exist.*

The intuition for the above result can be summarized as follows. As we pointed out earlier, in any state θ variances of the expected offer $\text{Var}(a-\theta)$ in upward and downward equilibria coincide, but the expected bias $\mathbb{E}(a-\theta)$ is higher in downward equilibrium: $b_D = b_1 F(z^*) + b_2(1-F(z^*)) + Bf(z^*) \geq b_1(1-F(z^*)) + b_2 F(z^*) - Bf(z^*) = b_U$. Hence, to conclude that the principal is better off in upward equilibrium it is enough to show that $|b_D| \geq |b_U|$ or, taking into account above, $b_U + b_D \geq 0$. Markup differences in upward and downward equilibria coincide, but the choice rule is opposite, therefore for both experts the probabilities of winning in upward and downward equilibria are complementary. The direct effects of the bonus on b_U and b_D are opposite and equal in absolute value. Therefore $b_U + b_D = b_1 + b_2 \geq 0$, implying that the principal is better off in upward equilibrium.

In the Supplementary Appendix we show that in case the two experts are not equally biased, the expert with the lower bias also prefers upward equilibrium to downward equilibrium, while the expert with the higher bias has the opposite preferences. This is both because the expected action is closer to expert 2's ideal point in upward equilibrium, and closer to expert 1's ideal point in downward equilibrium, and because expert 2 is chosen (and hence receives the bonus) with higher probability in upward equilibrium, and expert 1 is chosen with higher probability in downward equilibrium.

For equal biases, the intuition for why upward equilibrium, when it exists, yields a higher expected payoff to the principal than simple delegation generalizes for $B \geq 0$ as follows. With simple delegation, the expected utility of the principal is $V = -b^2 - \sigma^2$. In contrast, in upward equilibrium the expected bias of the action is $b - \frac{B}{2\sqrt{\pi}\sigma} < b$, and the variance is $\sigma^2 - \frac{\sigma^2}{\pi} < \sigma^2$. The decreased variance of the action implies that the principal prefers upward equilibrium to simple delegation, even when $B = 0$ and therefore the expected action is the same in upward equilibrium as in simple delegation.

3.4.2 Ex Ante Commitment by the Principal

The analysis in the previous sections assumes that the principal plays a best response to the strategies of the experts. Alternatively, we can consider situations in which the principal can ex ante commit to a simple strategy. In particular, assume that the principal can credibly commit to choose the lower offer with any probability $p \in [0, 1]$ and the higher offer with probability $1 - p$. In this section we provide two results. When biases are equal, as long as the bonus is small, commitment power does not help the principal; the principal would choose $p = 1$, which is already an (upward) equilibrium without commitment. When biases are opposite and sufficiently large in magnitude, the principal benefits from commitment to an interior p.

For the moment, restrict attention to the case of equally biased experts: $b_1 = b_2 = b$. While for very large bonuses such a commitment opportunity might be beneficial for the principal, here we show that for $B \leq 2\sigma^2$ commitment does not improve the principal's payoff, relative to the expected payoff in upward equilibrium in the game without commitment.

If the principal commits to choosing the lower offer with probability p, the experts' FOCs are

$$k_1 = b + \rho \frac{f(z)(2p-1)}{p(1-F(z)) + (1-p)F(z)}$$

$$k_2 = b + \rho \frac{f(z)(2p-1)}{pF(z) + (1-p)(1-F(z))},$$

where $z = k_1 - k_2$.

As $B \leq 2\sigma^2$, by Lemma $S.1$ from Supplementary Appendix, the only equilibrium of the game between the two experts, given the pre-committed strategy of the principal, involves $k_1 = k_2 = k = b + (2p-1)\frac{2\sigma^2 - B}{2\sqrt{\pi}\sigma}$. The principal's utility is

$$
\begin{aligned}
V &= p\left[-\sigma^2 - k^2 + \frac{2\sigma}{\sqrt{\pi}}k\right] + (1-p)\left[-\sigma^2 - k^2 - \frac{2\sigma}{\sqrt{\pi}}k\right] \\
&= \frac{\rho(2\sigma^2 + B)}{2\pi\sigma^2}(2p-1)^2 + \frac{bB}{\sqrt{\pi}\sigma}(2p-1) - \sigma^2 - b^2.
\end{aligned}
$$

The next proposition follows from inspection of the expression above.

Proposition 3.4.3. *For $B \leq 2\sigma^2$ and $b_1 = b_2$, the optimal strategy of the principal under commitment is $p = 1$.*

In particular, commitment by the principal results in the same outcome as in the upward equilibrium of the game without commitment.

Next, consider oppositely biased experts and a bonus of $B = 0$. As the common magnitude b of the biases increases, expert 1's winner's curse in the upward equilibrium becomes more severe, and his markup very large. By introducing a small probability of choosing the higher offer, expert 1 is incentivized to reduce his markup, benefiting the principal. The cost to the principal of doing so is in choosing the wrong offer. When b is large, the markup reduction is large and outweighs the cost, and thus the principal can profitably deviate from $p = 1$ to an interior p.

Proposition 3.4.4. *Let $b_1 = b > 0$, $b_2 = -b$, and $B = 0$. For sufficiently large b, the optimal p under commitment is strictly less than 1.*

3.4.3 Unselected Expert Indifferent over Actions

In the baseline model we assumed that an expert whose offer is not selected is still affected by principal's action. While this is a reasonable assumption in some contexts, in other situations it is more realistic to assume that the expert not selected by the principal receives an outside payoff that is independent of the state and the implemented action. For instance, a car mechanic is unlikely to care what kind of maintenance is done if he is not the one selected for the job. In this extension we assume that if expert i is chosen then expert j's realized payoff is normalized to be 0. We restrict attention the case of equally biased experts: $b_1 = b_2 = b > 0$

In this version of the model we assume that B is large enough that an expert's expected payoff under simple delegation to that expert is nonnegative; that is, he prefers simple delegation to not being selected at all. Under this condition, the same simple delegation equilibria exist in this version of the model as in the baseline model. Below we show that under some conditions there also exist symmetric pure strategy equilibria that are similar to the ones characterized in the baseline model. For this to be the case, the bonus payment must be neither too low nor too large.

First we examine the conditions for the existence of upward equilibrium. Using the same notation as before, we investigate strategy profiles $\{(k_1, k_2, C(a_1, a_2) \in \arg\min\{a_1, a_2\}) : k_1 + k_2 \geq 0\}$. While for any such profile the principal's payoff does not change, the experts' expected payoffs should be recalculated:

$$U_i(k_i, k_j, L) = \int_{k_i - k_j}^{\infty} \left[B - \left(k_i - b - \frac{t}{2} \right)^2 - \frac{\sigma^2}{2} \right] f(t)\, dt$$

$$= \left[B - \sigma^2 - (k_i - b)^2 \right] (1 - F(k_i - k_j))$$

$$+ \left[2\sigma^2 (k_i - b) - \frac{1}{2}\sigma^2 (k_i - k_j) \right] f(k_i - k_j).$$

Notice that for a fixed constant markup strategy of the other expert, an expert can choose arbitrarily high constant markup and guarantee an expected payoff arbitrarily close to 0.[18] Hence, 0 is a lower bound for experts' equilibrium payoffs.

In what follows, define $\beta := \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}}$.

[18]For this reason, we do not introduce an explicit participation constraint in this version of the model, even though such a constraint would be natural in many applications.

Proposition 3.4.5. *A symmetric upward equilibrium* $k_1^U = k_2^U = k_U$ *exists if and only if* $B \in \left[\left(\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2}\right)\sigma^2, \frac{5}{2}\sigma^2 + 2\sqrt{\pi}b\sigma + b^2\right]$. *When it exists, it is characterized by:*

- $k_1^U = k_2^U = k_U = b + (\sqrt{\pi} - \beta)\,\sigma$;
- $b(k_U, k_U, L) = b + \left(\sqrt{\pi} - \beta - \frac{1}{\sqrt{\pi}}\right)\sigma$;
- $Var(k_U, k_U, L) = \left(1 - \frac{1}{\pi}\right)\sigma^2$;
- $V(k_U, k_U, L) = -\left[b + \left(\sqrt{\pi} - \beta - \frac{1}{\sqrt{\pi}}\right)\sigma\right]^2 - \sigma^2 + \frac{\sigma^2}{\pi}$;
- $U_i(k_U, k_U, L) = \left[\frac{\pi-1}{\sqrt{\pi}}\beta + \frac{7}{4} - \pi\right]\sigma^2$ *for* $i = 1, 2$.

In the appendix we show that $k_U = b + (\sqrt{\pi} - \beta)\sigma > k_U^{bas.} = b + (1 - \frac{B}{2\sigma^2})\frac{\sigma}{\sqrt{\pi}}$, hence in this version of the model experts select higher markups in upward equilibrium than in the baseline model (for parameter values for which upward equilibrium exists in both model versions). The intuition behind this result is that in this alternative version of the model, the relative gain from being selected is reduced by the policy loss (that is not imposed on the expert if not selected). Since we consider B sufficiently large that expected payoffs are nonnegative, the resulting "net bonus" is still nonnegative; being selected is still preferable, conditional on having made the lower offer. It follows that an expert's equilibrium offer in either version is lower than what is ex-post optimal for that expert – that is, optimal after conditioning on both the expert's signal and having the lower offer. The smaller net bonus in the alternative version reduces the expert's incentive to marginally lower his offer in order to more frequently earn the net bonus. This reduction must be met by an offsetting reduction in his incentive to raise his offer, which is enforced through his bidding higher and thus closer to his ex-post optimum; due to quadratic losses, marginal movements toward the ex-post optimum have decreasing marginal benefits.

We note that the qualitative comparison between this extension and the baseline model is dependent upon the modeling of preferences over policy outcomes through losses. Such a model is appropriate for applications where an expert would prefer not to be associated with the project if his action would be sufficiently far from the true state; for example, this would be the case if the expert has a reputation at stake. Alternatively, one could model preferences through gains, using some single-peaked, nonnegative utility function of the distance between the action and the true state. In that model, the comparison above would be reversed, as being selected enhances the bonus and thus experts compete more aggressively by lowering their offers.

Next we turn attention to characterizing the conditions under which a downward equilibrium exists in which the principal always chooses the higher offer. Experts' expected payoffs can be calculated as:

$$
\begin{aligned}
U_i(k_i, k_j, H) &= \int_{-\infty}^{k_i - k_j} \left[B - \left(k_i - b - \frac{t}{2} \right)^2 - \frac{\sigma^2}{2} \right] f(t)\, dt \\
&= \left[B - \sigma^2 - (k_i - b)^2 \right] F(k_i - k_j) \\
&\quad - \left[2\sigma^2 (k_i - b) - \frac{1}{2}\sigma^2 (k_i - k_j) \right] f(k_i - k_j)
\end{aligned}
$$

As in upward equilibrium, 0 is a lower bound for experts' equilibrium payoffs.

Proposition 3.4.6. *Consider $b_1 = b_2 = b > 0$. If $\frac{b}{\sigma} > \frac{3\sqrt{\pi}}{4(\pi-1)}$, then no symmetric downward equilibrium exists. If $\frac{b}{\sigma} \leq \frac{3\sqrt{\pi}}{4(\pi-1)}$, then a symmetric downward equilibrium $k_1^D = k_2^D = k_D$ exists if and only if $B \in \left[\left(\frac{5}{2} - \frac{3\pi(8\pi - 11)}{16(\pi-1)^2} \right) \sigma^2, \frac{5}{2}\sigma^2 + 2\sqrt{\pi} b\sigma + b^2 \right]$. When it exists, it is characterized by:*

- $k_1^D = k_2^D = k_D = b - (\sqrt{\pi} - \beta)\, \sigma$;

- $b(k_D, k_D, H) = b - \left(\sqrt{\pi} - \beta - \frac{1}{\sqrt{\pi}}\right)\sigma;$

- $Var(k_D, k_D, H) = (1 - \frac{1}{\pi})\sigma^2;$

- $V(k_D, k_D, H) = -\left[b - \left(\sqrt{\pi} - \beta - \frac{1}{\sqrt{\pi}}\right)\sigma\right]^2 - \sigma^2 + \frac{\sigma^2}{\pi};$

- $U_i(k_D, k_D, H) = \left[\frac{\pi - 1}{\sqrt{\pi}}\beta + \frac{7}{4} - \pi\right]\sigma^2 \text{ for } i = 1, 2.$

For downward equilibrium, the difference relative to the baseline model is the mirror image of the difference described earlier for upward equilibrium. Again, the bonus is reduced by quadratic losses, but in downward equilibrium this causes markups to decrease, as experts compete less aggressively to make the higher offer.

3.5 Comparative Statics

We start this section with comparing the expected payoff the principal can achieve with two equally biased experts, $b_1 = b_2 = b$, to the expected payoff she can achieve with two oppositely biased experts, $b_1 = -b_2 = b$. Theorem 3.3.2* implies that the expected bias of the action in the case of equally biased experts is equal to b. In the case of oppositely biased experts the expected bias of the action is $b(1 - 2F(z^*)) - Bf(z^*)$, simplifying to $b(1 - 2F(z^*))$ when $B = 0$. Hence, with $B = 0$ the absolute value of the expected bias in the case of oppositely biased experts is lower than in the case of equally biased experts. However, the next Proposition shows that the variance of the action is lower in the symmetric case, and in fact this effect dominates, resulting in the principal preferring to have two equally biased experts. Let $V_{symm}(b)$ be the principal's expected payoff in upward equilibrium when $b_1 = b_2 = b$, let $V_{opp}(b)$ be the principal's expected payoff in upward equilibrium when $b_1 = -b_2 = b$, and let $V_{sim}(b) = -\sigma^2 - b^2$ be the principal's expected payoff in case of simple delegation to an expert with absolute bias b.

Proposition 3.5.1. *For any $b > 0$, $V_{symm}(b) \geq V_{opp}(b) \geq V_{sim}(b)$.*

Proof. Upward equilibrium for opposite biases exists only if $B \leq B_U = 2\sigma^2$. From Theorem 3.3.2*

$$V_{symm}(b) = -\sigma^2 - b^2 + 2Bbf(0) + (4\sigma^4 - B^2)f^2(0)$$

$$V_{opp}(b) = -\sigma^2 - b^2 + \left(\sigma^4 - \frac{B^2}{4}\right)\frac{f^2(z^*_{opp})}{F(z^*_{opp})(1 - F(z^*_{opp}))} \geq -\sigma^2 - b^2 = V_{sim}(b),$$

where z^*_{opp} is the upward equilibrium markup difference in the case of oppositely biased experts. As $2Bbf(0) \geq 0$ and $\frac{f^2(z)}{F(z)(1-F(z))}$ reaches its maximum at $z = 0$, we get $V_{symm}(b) \geq V_{opp}(b)$ □

3.5.1 Equally Biased Experts

For the remainder of this section we focus on the case when the experts are equally biased: $b_1 = b_2 = b > 0$. In this setting we can derive closed form solutions for the marginal affects of various parameter values on the principal's expected payoff, greatly simplifying comparative statics exercises.

Recall from Proposition 3.3.5 that in the case of equal biases, the principal always prefers upward equilibrium to simple delegation to one of the biased experts. Moreover, using the formula derived in Proposition 3.5.1 for the principal's expected utility in upward equilibrium, we can exactly characterize when it is the case that the principal prefers upward equilibrium with two equally biased experts to simple delegation to an unbiased expert.

Corollary 3.5.1. *The principal prefers upward equilibrium with equally biased experts to unconstrained delegation to an unbiased expert if and only if $B \in [2\sqrt{\pi}\sigma b - 2\sigma^2, B_u]$. For $B = 0$ this condition is equivalent to bias-to-noise ratio being low* *enough:*

$$\frac{b}{\sigma} \leq \frac{1}{\sqrt{\pi}}.$$

Proof. It is enough to compare $V_{symm} = -(b - \frac{B}{2\sqrt{\pi}\sigma})^2 - \sigma^2 + \frac{\sigma^2}{\pi}$ with utility from single delegation to unbiased expert $-\sigma^2$. □

From here on we investigate comparative statics of the principal's expected payoff in upward equilibrium. First we look at how the principal's expected payoff depends on b, the common bias of the experts. Recall that upward equilibrium exists only for $B \leq B_u = 2\sigma^2 + 2\sqrt{\pi}\sigma b$. The principal's expected payoff in upward equilibrium is $V(k_U, k_U, L) = -(b - \frac{B}{2\sqrt{\pi}\sigma})^2 - (1 - \frac{1}{\pi})\sigma^2$. This expression is maximized at $b^* = \frac{B}{2\sqrt{\pi}\sigma}$, taking the value $-(1 - \frac{1}{\pi})\sigma^2$. This implies that for $B = 0$, it is optimal for the principal to have nonbiased experts, but for $B > 0$, the optimal bias level is strictly positive and increasing in B.

Next we consider how the principal's expected payoff depends on the precision of the experts' signals. The expected bias of the action $b_U = b - \frac{B}{2\sqrt{\pi}\sigma}$ is increasing in σ. Nevertheless, the equilibrium utility of the principal, $V(k_U, k_U, L) = -(b - \frac{B}{2\sqrt{\pi}\sigma})^2 - \sigma^2 + \frac{\sigma^2}{\pi}$, is nonmonotonic in σ, for the following reason. If σ increases, then the expected variance of the action $(1 - \frac{1}{\pi})\sigma^2$ increases, leading to a decrease in the principal's expected payoff. At the same time when the variance of noise σ is small, the value of the expected bias $b - \frac{B}{2\sqrt{\pi}\sigma}$ stays negative and increases, leading to a decrease in absolute value of the expected bias, which increases the principal's expected payoff. After σ reaches the value when B is optimal in exogenous case ($B = 2\sqrt{\pi}\sigma b$) the value of expected bias turns positive and further increase in σ leads to decrease in utility. Hence, optimal value for principal lies on the interval $(0, \frac{B}{2\sqrt{\pi}b})$. The next proposition provides the formal comparative statics of the principal's utility in the precision of the experts' signals.

Let $B > 0$ and σ^* denote the unique positive solution to the equation $(4\pi - 4)\sigma^4 + 2\sqrt{\pi} b B \sigma = B^2$.

Proposition 3.5.2. *Consider $b_1 = b_2 = b > 0$.*

If $0 < B < \frac{2(\pi-1)\pi}{(\pi-2)^2} b^2$, there exists a non-empty interval $\sigma \in [\frac{\sqrt{\pi b^2 + 2B} - \sqrt{\pi} b}{2}, \sigma^*)$, where upward equilibrium exists and the principal's expected payoff is increasing in σ; when $\sigma > \sigma^*$, the principal's expected payoff is decreasing in σ.

If $B = 0$ or $B \geq \frac{2(\pi-1)\pi}{(\pi-2)^2} b^2$, the principal's expected payoff is always decreasing in σ.

For a concrete example when the principal's expected payoff increases in σ, consider $b_1 = b_2 = b = 2$ and $B = 9 < \frac{8(\pi-1)\pi}{(\pi-2)^2}$. Then $V(\sigma)$ is increasing on an interval containing $(0.992, 1.085)$.

Lastly, we address the question of how the principal's expected payoff depends on the bonus payment. We start with the case when the bonus payment comes from exogenous sources, and therefore only indirectly affecting the principal's payoff, through influencing experts' strategies in upward equilibrium. While the variance of the action $Var(k_U, k_U, L) = \sigma^2 - \frac{\sigma^2}{\pi}$ does not depend on B, the expected bias $b_U = b - \frac{B}{2\sqrt{\pi}\sigma}$ is decreasing in B. The principal prefers the expected bias to be as close to 0 as possible, so her expected payoff in upward equilibrium is maximized at $B = 2\sqrt{\pi}\sigma b$, where it is equal to $-\sigma^2[1 - \frac{1}{\pi}]$. At $B = 0$, $b_U = b - \frac{B}{2\sqrt{\pi}\sigma} = b \geq 0$ and a small increase in bonus decreases the experts' markup and benefits the principal. However, at $B = B_u$ $b_U = b - \frac{B}{2\sqrt{\pi}\sigma} = -\frac{\sigma}{\sqrt{\pi}} < 0$ and principal prefers to increase markups and correspondingly lower bonus. As a consequence, an intermediate point $B = 2\sqrt{\pi}\sigma b$ is optimal.

Next we consider the situation, when the bonus is paid from the principal's pocket. For this investigation, we append our game (described in Section 2) with a stage 0, preceding stage 1, in which the principal chooses B. We assume that the principal's

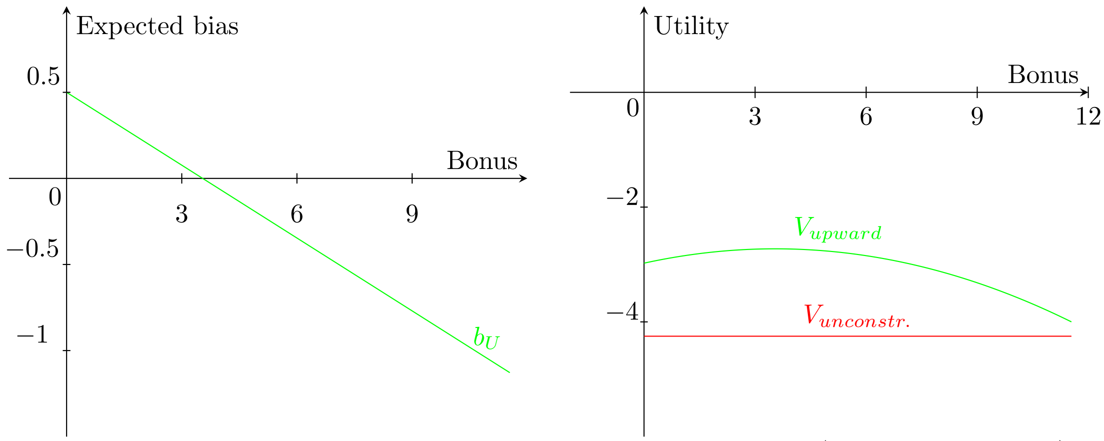

FIGURE 3.3: The Expected Bias and the Principal's Payoff (exogenous bonus) for $b = 0.5, \sigma = 2$

choice of B becomes public knowledge by stage 1. We also modify the principal's payoff to $V(a, \theta) = -(a-\theta)^2 - B$, corresponding to the assumption that the principal has to pay the bonus.

The next proposition characterizes the optimal bonus choice of the principal in this case.

Proposition 3.5.3. *Suppose that bonus B is paid by the principal. The principal's optimal choice of B depends on the bias-to-noise ratio.*
If $\frac{b}{\sigma} \leq \sqrt{\pi}$, then the principal pays no bonus: $B = 0$.
If $\frac{b}{\sigma} > \sqrt{\pi}$, then the principal chooses bonus $B = 2\sqrt{\pi}\sigma^2 \left[\frac{b}{\sigma} - \sqrt{\pi}\right]$. The increase in the principal's expected payoff, relative to when the bonus is restricted to be 0, is equal to $(b - \sqrt{\pi}\sigma)^2$.

Therefore, if the bias-to-noise ratio is high enough, it is optimal for the principal to offer a positive bonus.

At $B = B_u$: $b_U = k_U - \frac{\sigma}{\sqrt{\pi}} = -\frac{\sigma}{\sqrt{\pi}} < 0$ and the principal would like to to decrease the bonus, in order to increase experts' markups and drive the expected bias closer to 0.

At $B = 0$: $b_U = b \geq 0$, and if the principal increases the bonus then the expected bias decreases. Therefore a marginal increase of B from 0 improves the principal's expected payoff whenever the marginal gain from the decrease in expected bias ($\frac{b}{\sqrt{\pi}\sigma}$) exceeds the marginal expense from increasing the bonus (1).

Proposition 3.5.3 is illustrated on Figure 3: the principal chooses a strictly positive bonus if the decrease in expected bias exceeds the marginal disutility from increasing bonus: $\frac{b}{\sqrt{\pi}\sigma} > 1$. This for example holds for $b = 2$, $\sigma = 0.5$, but not for $b = 1$, $\sigma = 1$.

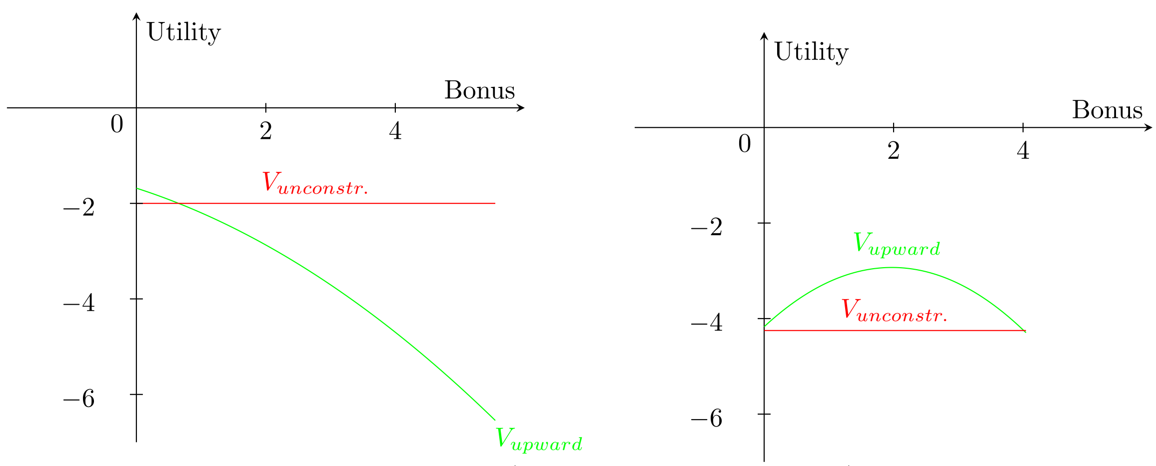

FIGURE 3.4: The Principal's Payoff (bonus paid by Principal) for $b = 1$, $\sigma = 1$ on the left; for $b = 2$, $\sigma = 0.5$ on the right

3.6 Conclusion

We proposed a model in which a principal can choose between two imperfectly informed experts, introducing the possibility of competition in a delegation framework. We showed that a principal with limited knowledge of the decision environment can benefit from the presence of two experts, relative to a simple unconstrained delegation to one of them, even if the experts have exactly the same bias. The main reason is that in equilibria in which the selection of the expert depends nontrivially on the experts' proposals, information is utilized from both experts' private signals. The option of offering a bonus payment to the selected expert can improve the principal's payoff, by inducing the experts to report more truthfully, but only to a certain point. The effects of more precise signals to the experts on the principal's payoff is in general ambiguous, but they are always beneficial in the absence of bonus payments.

We also find that in the equilibria we characterize, the principal being severely restricted in her choice of strategies, due to limited knowledge of the environment, might not hurt her. In particular, given the equilibrium strategies of the experts, the principal would not be able to do better without the restrictions to simple strategies. Moreover, for small bonuses the best equilibrium of the game achieves the payoff the principal could get by ex ante committing to a strategy. Hence, if there is competition between experts to be selected, sophistication on the part of the experts can lead to the same outcome no matter of the principal's knowledge about the environment.

As this is the first step in investigating the benefits of multiple choices of experts in a delegation problem, there are many avenues of future research. One is examining multi-dimensional environments, in which different experts differ in their dimensions of specialization. Another direction would be investigating the problem of choosing an expert to delegate a task to with a more general mechanism design approach.

3.7 Appendix A: Microfoundations for Ex Ante Expected Payoffs with a Diffuse Prior

In this section we provide a formal interpretation of a uniform improper prior as a limit construction, and characterize strategies for which ex ante expected payoffs are well-defined in a formal sense.

3.7.1 Overview

While the uninformative prior is often interpreted as a uniform distribution on the real line, it is not properly defined. We capture the main features of the uninformative prior by using a sequence of (proper) measures that *diffuse* in a formal sense (Definition 3.7.3). In this section, we provide two definitions of strategy admissibility and derive a main result for each. We say that a strategy is *weakly admissible* (Definition 3.7.4) when ex ante expected payoffs, taken along any diffusing sequence of priors, have a well-defined limit that does not depend on the particular sequence. Similarly, a strategy is *strongly admissible* (Definition 3.7.6) if the joint distribution over noise and markups it induces has such a limit. Our first result (Theorem 3.7.1) states roughly that in any class of weakly admissible strategies that includes constant markup strategies, every strategy is *nearly stationary* in a particular sense (Definition 3.7.5). Our second result (Theorem 3.7.2) states that any strongly admissible strategy has a stronger property: it is payoff-equivalent to some mixture of constant markup strategies.

Since we start from strategy sets that are not restricted to be mixtures of constant markup strategies, we need to provide a more careful formal definition of strategies. Since the state space is uncountable, it is not possible to define strategies as products of state-dependent distributions over markups. The key tension is that desirable topologies should be both rich enough so that payoff functions are continuous in strategies, but also coarse enough so that the strategy space is compact. We follow the approach of Milgrom and Weber (1985) in using *distributional strategies*, which are measures μ over the product space of signals and markups with bounded support with respect to the latter. A distributional strategy induces a conditional distribution on markups for any given signal, and by integrating over signals, it induces a conditional distribution on actions for any given θ.

Below we first provide outlines of the main results detailed in subsections 3.7.2 and 3.7.3.

Weak Admissibility

Given the form of the experts' loss functions, it is useful to describe strategies in terms of distributions over offers, net of the true state. For every expert strategy μ and state θ, we define $G_\mu(\cdot;\theta)$ as this distribution conditional on θ. We show that $G_\mu(\cdot;\theta)$ is continuous in both μ and θ, and the space of all such distributions, denoted $\mathcal{G}^M$, is compact (Lemmas 3.7.1 and 3.7.2).[19] We say that μ is *near* a given distribution $Q \in \mathcal{G}^M$, and hence *nearly stationary*, if the following property holds: given any $\eta > 0$, $G_\mu(\cdot, \cdot;\theta)$ is within η distance of Q for all θ except for a set of finite Lebesgue measure (Definition 3.7.2).

We begin the proof of Theorem 3.7.1 by establishing the existence of some distribution $Q \in \mathcal{G}^M$ with the following property: for all $\eta > 0$, $G_\mu(\cdot;\theta)$ is within η of Q for an infinite measure set of θ (Lemma 3.7.3). We then call Q an *attraction*. A necessary condition for μ to be near Q is that Q is an attraction.

Next, we argue that there can be at most one such attraction for any strategy. The proof of this claim is by contradiction and relies on several steps. We suppose that Q and Q' are two distinct attractions. We then show that there must exist some constant markup strategy of the rival, κ_m, against which these distributions yield distinct expected payoffs, denoted $u_{Q,\kappa_m}, u_{Q',\kappa_m} \in \mathbb{R}$. Given $\eta > 0$, we can construct a sequence of measures $(\mathbb{P}^1_n)_{n\in\mathbb{N}}$ (resp. $(\mathbb{P}^2_n)_{n\in\mathbb{N}}$) that places increasing mass on θ such that $G_\mu(\cdot;\theta)$ is within η of Q (resp. Q'). By continuity, the limit of expected payoffs must be at most η from u_{Q,κ_m} (resp. u_{Q',κ_m}). As these are assumed distinct, this is a contradiction.

---

[19]We use the topology of weak convergence for measures and distributions, which is metrized by the Prokhorov distance, and we use the usual topology for $\mathbb{R}$; continuity and compactness are with respect to these topologies.

Finally, given the unique attraction Q, we show that μ is near Q. If μ is near any distribution, that distribution must be an attraction, so Q is the only candidate. We show that if μ is not near Q, then there is a compact set of distributions that does not contain Q but contains $G_\mu(\cdot;\theta)$ for infinitely many θ. This compact set itself contains an attraction, and this contradicts the uniqueness of Q, completing the proof.

Strong Admissibility

For any expert strategy μ and state θ, we define $H_\mu(\cdot,\cdot;\theta)$ as the distribution of noise and markups conditional on θ. Strong admissibility is essentially the requirement that these distributions have well-defined expectations in the limit (Definition 3.7.6), and it implies weak admissibility. As before, we show that $H_\mu(\cdot,\cdot;\theta)$ is continuous in μ and in θ, and the space of all such distributions, denoted $\mathcal{H}^M$, is compact (Lemmas 3.7.4 and 3.7.5).

The proof of Theorem 3.7.2 begins in the same way as that of Theorem 3.7.1. Given a strongly admissible strategy μ, we first show the existence of some distribution $Q \in \mathcal{H}^M$ that is an attraction for μ. Next, we argue by contradiction that there can be at most one such attraction for any strategy. Given the unique attraction Q, we then show that μ is near Q.

To complete the proof, we define K_Q as the marginal distribution over markups from Q; hence K_Q characterizes a mixture over constant markup strategies. We show that expected payoffs along any diffusing sequence of priors must converge to the same limit under μ as under K_Q, and thus the two strategies are payoff equivalent.

3.7.2 Formal Analysis of Weak Admissibility

We begin by providing a formal definition of an expert's strategy. As mentioned previously, in games of incomplete information with uncountable type spaces, certain topological problems arise from interpreting strategies as mixed strategies (i.e., distributions over pure strategies) or behavioral strategies (i.e., products of history-contingent distributions over actions). As an alternative, Milgrom and Weber (1985) introduce the notion of *distributional strategies*, which we adapt in Definition 3.7.1.

We let $S_i = \mathbb{R}$ denote the signal space for expert i, and $X_i = \mathbb{R}$ the space of markups. For $M \in \mathbb{R}$, we use the notation $X_i^M := [-M, M]$. Let λ denote the Lebesgue measure and let γ denote the measure corresponding to $N(0, \sigma^2)$. We let $h(s|\theta)$ denote the PDF of $N(\theta, \sigma^2)$ and abbreviate $h(s|0)$ to $h(s)$.

Definition 3.7.1. *A strategy for expert $i \in \{1, 2\}$ is a probability measure μ_i on $S_i \times X_i$ such that:*

- *for all measurable $T \subset S_i$, $\mu_i(T \times X_i) = \gamma(T)$, and*
- *there exists $M \in \mathbb{R}$ such that* $\text{supp}(\mu(\cdot|s_i)) \subset X_i^M$ *for all $s_i \in S_i$,*

where $\mu(\cdot|s_i)$ is the regular conditional distribution.

Since $S_i \times X_i$ is complete and separable, the space of probability measures $\Delta(S_i \times X_i)$, under the topology of weak convergence, is metrized by the Prokhorov distance.[20] We let $\mathcal{M} \subset \Delta(S_i \times X_i)$ denote the space of all strategies, and $\mathcal{M}^M := \{\mu \in \mathcal{M} : \text{supp}(\mu(\cdot|s_i)) \subset X_i^M \, \forall s_i \in S_i\}$. Note that $\mathcal{M} = \cup_{M \in \mathbb{R}} \mathcal{M}^M$.

At face value, a strategy as defined above gives a distribution over signals and markups conditional on $\theta = 0$, since the marginal distribution γ is $N(0, \sigma^2)$. However, by multiplication of a factor $\frac{h(s_i|\theta)}{h(s_i|0)}$, the strategy uniquely determines these distributions conditional on any $\theta \in \mathbb{R}$.

For any strategy $\mu \in \mathcal{M}$ and $\theta \in \mathbb{R}$ we define a CDF

$$G_\mu(x; \theta) := \int_{-\infty}^{\infty} \mu((-\infty, x - (s_i - \theta)]|s_i) \frac{h(s_i|\theta)}{h(s_i|0)} \gamma(ds_i) \tag{3.7.1}$$

$$= \int_{-\infty}^{\infty} \mu((-\infty, x - (s_i - \theta)]|s_i) h(s_i|\theta) \lambda(ds_i).$$

[20] See Billingsley (2009, Theorem 6.8).

We interpret $G_\mu(x;\theta)$ as the joint probability, conditional on θ, that the expert makes an offer that exceeds θ by at most x. We define

$$\mathcal{G}^M := \{P \in \Delta(\mathbb{R}) : P \text{ is induced by the CDF } G_\mu(\cdot;\theta) \text{ for some } \theta \in \mathbb{R}, \mu \in \mathcal{M}^M\}.$$

By inspection of (3.7.1), it is not difficult to see that for all $Q \in \mathcal{G}^M$,

$$F_{-M} \preceq_{FSD} Q \preceq_{FSD} F_{+M}, \tag{3.7.2}$$

where $\preceq_{FSD}$ denotes first-order stochastic dominance and F_k denotes the cumulative distribution of $N(k,\sigma^2)$. Note also that for any $\theta \in \mathbb{R}$ and $\mu \in \mathcal{M}^M$, there exists $\mu' \in \mathcal{M}^M$, defined by $\mu'(Y|s_i) = \mu(Y|s_i + \theta)$ for all $Y \subset X^i$ and all $s \in S_i$, which satisfies $G_{\mu'}(\cdot;0) = G_\mu(\cdot;\theta)$, and thus $\mathcal{G}^M$ is the image of the set $\mathcal{M}^M$ under the map $G : \mu \mapsto G_\mu(\cdot;0)$. Henceforth, we abbreviate $G_\mu(\cdot;0)$ to G_μ. We let $\mathcal{G} := \cup_{M\in\mathbb{R}} G^M$. As with $\mathcal{M}^M$ and $\mathcal{M}$, we endow $\mathcal{G}^M$ and $\mathcal{G}$ with the topology of weak convergence, which is metrized by the Prokhorov distance.

The next two lemmas establish continuity and compactness properties, to be used toward the main result.

Lemma 3.7.1. *For any $\mu \in \mathcal{M}$, the mapping $G_\mu : \theta \mapsto G_\mu(\cdot;\theta)$ from $\mathbb{R}$ to $\mathcal{G}$ is continuous in θ.*

Lemma 3.7.2. *For all $M \in \mathbb{R}$, $\mathcal{M}^M$ and $\mathcal{G}^M$ are compact metric spaces.*

Definition 3.7.2. *For any $M \in \mathbb{R}$ and any $\mu \in \mathcal{M}^M$, we say that μ is* near *a distribution $Q \in \mathcal{G}^M$ if for all $\eta > 0$, $\lambda\{\theta : d_P(G_\mu(\cdot;\theta), Q) \geq \eta\} < \infty$, where d_P is the Prokhorov metric. We then classify μ as* nearly stationary.

It is immediate from the definition that a strategy can be near to at most one distribution in $\mathcal{G}^M$. However, a strategy need not be near any distribution; consider, for example, a strategy that assigns markup 1 for all signals $s \geq 0$ and markup 0 for all signals $s < 0$.

Given a strategy profile and bias, we define expert 1's state-dependent expected utility[21]

$$u(\theta; \mu_1, \mu_2, C, b) := -\int_{\mathbb{R}^2} \mathbb{1}\{C(x_1+\theta, x_2+\theta) = 1\}(x_1 - b)^2 \\ + \mathbb{1}\{C(x_1+\theta, x_2+\theta) = 2\}(x_2 - b)^2 \nu(d(x_1, x_2)),$$

where $\nu = G_{\mu_1}(\cdot;\theta) \otimes G_{\mu_2}(\cdot;\theta)$ denotes the product measure. When there is no risk of confusion, we write $u(\theta)$ for brevity.

We formalize the uniform improper prior through *diffusing* sequences of proper measures, defined next. The definition incorporates the intuition that these sequences should increasingly spread mass over the real line and approach uniformity.

Definition 3.7.3. *Consider a sequence $(\mathbb{P}_n)_{n\in\mathbb{N}}$ of Borel probability measures on $\mathbb{R}$. We say that this sequence is* diffusing *if for any nonempty and bounded measurable set X and any $\eta > 0$, there exists $N \in \mathbb{N}$ such that for all $n \geq N$, $\mathbb{P}_n(X) > 0$ and for any measurable $Y \subseteq X$, $\left|\frac{\mathbb{P}_n(Y)}{\mathbb{P}_n(X)} - \frac{\lambda(Y)}{\lambda(X)}\right| < \eta$.*

It follows from this definition that if $(\mathbb{P}_n)$ is diffusing, then for any bounded and measurable set Y, $\lim_{n\to\infty} \mathbb{P}_n(Y) = 0$. To see this, take any such set Y and any $\eta > 0$. Pick any M such that $Y \subseteq X := [-M, M]$ and $\lambda(Y) < 2M\eta = \lambda(X)\eta$. Then for sufficiently large N, $n \geq N$ implies $\mathbb{P}_n(X) > 0$ and $|\frac{\mathbb{P}_n(Y)}{\mathbb{P}_n(X)} - \frac{\lambda(Y)}{\lambda(X)}| < \eta$, which implies that $\mathbb{P}_n(Y) < \frac{\mathbb{P}_n(Y)}{\mathbb{P}_n(X)} < \eta + \frac{\lambda(Y)}{\lambda(X)} < 2\eta$.

For illustrative purposes, we highlight two specific diffusing sequences.

Example 1. *Both $(\mathbb{P}_n^1)$ given by the density $\frac{1}{2n}\mathbb{1}\{[-n, n]\}$ and $(\mathbb{P}_n^2)$ given by $N(0, n)$ are diffusing.*

Proof. If X is bounded and nonempty, then $X \subset [-n, n]$ for some n, and then $\mathbb{P}_n^1(X) = \lambda(X)/(2n) > 0$. Moreover, if $Y \subset X$, then $\mathbb{P}_n^1(Y) = \lambda(Y)/(2n)$ and thus

[21] It is without loss of generality to focus on expert 1, as the labeling of experts is arbitrary.

$\left|\frac{\mathbb{P}_n^1(Y)}{\mathbb{P}_n^1(X)} - \frac{\lambda(Y)}{\lambda(X)}\right| = 0$. Now $\mathbb{P}_n^2(X) > 0$ for all $n \in \mathbb{N}$. Let h_n denote the density of $N(0, n)$. Note that for all $\eta \in (0, 1)$, for sufficiently large n, $\frac{\inf_{x \in X} h_n(x)}{\sup_{x \in X} h_n(x)} > 1 - \eta$. Hence $(1-\eta)^2 \frac{\lambda(Y)}{\lambda(X)} < \frac{\mathbb{P}_n^2(Y)}{\mathbb{P}_n^2(X)} < (1-\eta)^{-2} \frac{\lambda(Y)}{\lambda(X)}$, and by algebra the result follows. □

Next, we define formally the key property featured in Theorem 3.7.1.

Definition 3.7.4. *A class $\mathcal{M}_0 \subset \mathcal{M}$ of expert strategies is said to be* weakly admissible *if for any $\mu_1, \mu_2 \in \mathcal{M}_0$, principal strategy C, and $b \in \mathbb{R}$:*

- *For any diffusing sequence $(\mathbb{P}_n)$ and sufficiently large n, $u(\cdot)$ is $\mathbb{P}_n$-integrable and has finite expectation.*
- *There exists $u^* \in \mathbb{R}$ such that for any diffusing sequence $(\mathbb{P}_n)$,*

$$\lim_{n \to \infty} \int_{\theta \in \mathbb{R}} u(\theta)\, d\mathbb{P}_n(\theta) = u^*.$$

For any $m \in \mathbb{R}$, a constant markup strategy, denoted κ_m, is the unique strategy such that for all $s_i \in S_i$, $\kappa_m(\{m\}|s_i) = 1$. Let $\mathcal{K} := \{\kappa_m : m \in \mathbb{R}\}$ denote the class of constant markup strategies. Note that for constant markup strategies and their mixtures, the function $u(\cdot)$ defined above is constant in θ, and hence the sequence of integrals in Definition 3.7.4 trivially converges to the same limit for every diffusing sequence.

Example 2. *Constant markup strategies and mixtures of constant markup strategies are weakly admissible.*

For the proof of the main result, we make use of a weaker property than that of Definition 3.7.2.

Lemma 3.7.3. *Let (Y, d) be a compact metric space, $X \subseteq \mathbb{R}$ with $\lambda X = \infty$, and $\pi : X \to Y$ a Lebesgue measurable function. There exists $y \in Y$ with the property* *that for all* $\eta > 0$, $\lambda\{x \in X : d(\pi(x), y) < \eta\} = \infty$. *We say that any such* y *is an* attraction.

We are now ready to state the main result for weak admissibility. The full proof can be found in the appendix.

Theorem 3.7.1. *If* $\mathcal{M}_0$ *is weakly admissible and* $\mathcal{K} \subset \mathcal{M}_0$*, then every* $\mu \in \mathcal{M}_0$ *is nearly stationary.*

3.7.3 Formal Analysis of Strong Admissibility

In order to characterize admissible strategies in a more transparent way, it is useful to define joint distributions over noise and markups given θ. For any strategy $\mu \in \mathcal{M}$ and $\theta \in \mathbb{R}$ we define a CDF

$$H_\mu(\epsilon, x; \theta) := \int_{-\infty}^{\epsilon} \mu((-\infty, x]|\hat{\epsilon}_i + \theta) h(\hat{\epsilon}_i + \theta|\theta) \lambda(d\hat{\epsilon}_i) \tag{3.7.3}$$
$$= \int_{-\infty}^{\epsilon} \mu((-\infty, x]|\hat{\epsilon}_i + \theta) h(\hat{\epsilon}_i) \lambda(d\hat{\epsilon}_i).$$

The interpretation of $H_\mu(\epsilon, x; \theta)$ is the probability, conditional on θ, that the noise realization is at most ϵ and the markup is at most x. We make several observations analogous to those following (3.7.1) in the previous section. Note that for any $\theta \in \mathbb{R}$ and $\mu \in \mathcal{M}^M$, there exists $\mu' \in \mathcal{M}^M$ with behavior conditional on state 0 matching that of μ conditional on state θ, and thus in some instances we can assume $\theta = 0$ without loss of generality. That is, set $\mu'(Y|s_i) = \mu(Y|s_i + \theta)$ for all measurable $Y \subset X_i$ and all $s_i \in S_i$, which satisfies $H_{\mu'}(\cdot, \cdot; 0) = H_\mu(\cdot, \cdot; \theta)$, and thus $\mathcal{H}^M$, defined as the image of the set $\mathcal{M}^M$ under the map $H : \mu \mapsto H_\mu(\cdot, \cdot; 0)$, is $\mathcal{M}^M$ itself. We let $\mathcal{H} := \cup_{M \in \mathbb{R}} \mathcal{H}^M$. As with $\mathcal{M}^M$ and $\mathcal{M}$, we endow $\mathcal{H}^M$ and $\mathcal{H}$ with the topology of weak convergence and the Prokhorov metric.

Lemma 3.7.4. *For any $\mu \in \mathcal{M}$, the mapping $H_\mu : \theta \mapsto H_\mu(\cdot, \cdot; \theta)$ from $\mathbb{R}$ to $\mathcal{H}$ is continuous in θ.*

The next lemma establishes compactness for the spaces we are analyzing. The compactness of $\mathcal{M}^M$ is a restatement of Lemma 3.7.2. The compactness of $\mathcal{H}^M = \Delta([-M, M])$ is standard but also follows from Lemma 3.7.4 together with the compactness of $\mathcal{M}^M$.

Lemma 3.7.5. *For all $M \in \mathbb{R}$, $\mathcal{M}^M$ and hence $\mathcal{H}^M$ are compact metric spaces.*

The next definition is analogous to Definition 3.7.2.

Definition 3.7.5. *For any $M \in \mathbb{R}$ and any $\mu \in \mathcal{M}^M$, we say that μ is* near *a distribution $Q \in \mathcal{H}^M$ if for all $\eta > 0$, $\lambda\{\theta : d_P(H_\mu(\cdot, \cdot; \theta), Q) \geq \eta\} < \infty$, where d_P is the Prokhorov metric. We then classify μ as* nearly stationary.

For a strategy to be strongly admissible, we require that it admits a well-defined expected distributional strategy on the product space of noise and markups; that is, the expected joint distribution over noise and markups, when taken along any diffusing sequence, converges to a fixed limiting distribution that is independent of the particular sequence. This condition ensures that for any strategy profile and conditional on any measurable set of signals, expected payoffs are well-defined in that for all diffusing sequences, the vector of expected payoffs converges to the same payoff vector.[22] Our definition uses the fact that convergence in distribution is equivalent to convergence of CDFs at points of continuity (Billingsley, 2009, Theorem 2.1 and Example 2.3).

[22]For any strategy μ and state θ, let $P_\mu(y; \theta) := \int_{-\infty}^{\infty} H_\mu(\epsilon, y - \epsilon; \theta)\, d\epsilon$ denote the probability that an offer exceeds θ by at most y. Strong admissibility guarantees that for any diffusing sequence, these distributions have a well-defined limiting expectation, $P^*_\mu := \int_{-\infty}^{\infty} Q(\epsilon, y - \epsilon)\, d\epsilon$, where Q is defined as in Definition 3.7.6. For any $\mathbb{P}_n$ and subset of signals, expected payoffs are calculated by integrating (on the outside) with respect to θ and (on the inside) with respect to y. By changing the order of integration and taking limits inside the integral, limiting expected payoffs are thus well-defined as single integrals over y according to the distribution P^*_μ. For payoffs conditional on a subset of signals, the logic is similar.

Definition 3.7.6. *An expert strategy $\mu \in \mathcal{M}^M$ is said to be* strongly admissible *if there exists a CDF Q over $S_i \times X_i^M$ such that for all (ϵ, x) at which Q is continuous and for all diffusing sequences $(\mathbb{P}_n)$,*

$$\lim_{n\to\infty} \int_{\theta\in\mathbb{R}} H_\mu(\epsilon, x|\theta)\, d\mathbb{P}_n(\theta) = Q(\epsilon, x).$$

For any $m \in \mathbb{R}$, a constant markup strategy, denoted κ_m, is the unique strategy such that for all $s_i \in S_i$, $\kappa_m(\{m\}|s_i) = 1$. Let $\mathcal{K} := \{\kappa_m : m \in \mathbb{R}\}$ denote the class of constant markup strategies. Note that for constant markup strategies, the function $H_\mu(\epsilon, x|\cdot)$ defined above is constant in θ, and hence the sequence of integrals in Definition 3.7.6 trivially converges to the same limit for every diffusing sequence.

Example 3. *All constant markup strategies and their mixtures are strongly admissible.*

Given a strongly admissible strategy μ, we use Lemma 3.7.3 to show the existence of an attraction for μ. We show that this attraction, denoted Q, is unique by way of contradiction; otherwise, there would exist diffusing sequences under which the limiting distributions of Definition 3.7.6 differ. Now μ is near to at most one strategy (and possibly none). We show that μ must be near some strategy, so it is near Q, and hence it is nearly stationary. Moreover, Q defines a mixture of constant markup strategies, K_Q. Finally, we show that μ is payoff equivalent to this mixture in that for any strongly admissible rival strategy, bias, and simple principal strategy, the limiting expected utility under μ is the same as under Q.

Theorem 3.7.2. *If μ is strongly admissible, then it is nearly stationary and payoff equivalent to some mixture of constant markup strategies.*

Motivated by Theorem 3.7.2, we hereafter restrict experts' strategies to be (mixtures over) constant markup strategies.

3.7.4 Proofs for Section 3.7.2

Proof of Lemma 3.7.1. The proof makes use of two facts: (i) for mappings on metric spaces, continuity is equivalent to sequential continuity,[23] and (ii) weak convergence of probability distributions is equivalent to pointwise convergence of the CDFs at all points of continuity.

To show that the mapping G_μ is continuous in θ, for any θ, choose any x at which $G_\mu(x;\theta)$ is continuous in x, and consider any sequence $\theta_n \to \theta$. We claim that $\mu((-\infty, x-(s_i-\theta)]|s_i)$ is continuous in x, and hence θ, for almost all $s_i \in S_i$. Since CDFs are right-continuous, it suffices to show left-continuity. If left-continuity fails for all $s_i \in D$ for some $D \subset S_i$ of positive measure, then for all $s_i \in D$ and all sequences $(x_n) \nearrow x$, $\mu((-\infty, x_n-(s_i-\theta)]|s_i) \leq \mu((-\infty, x-(s_i-\theta))|s_i) < \mu((-\infty, x-(s_i-\theta)]|s_i)$, and thus $\lim_{n\to\infty} G_\mu(x_n;\theta) < G_\mu(x;\theta)$, contradicting continuity of $G_\mu(x;\theta)$ in x. Thus for almost all $s_i \in S_i$, the integrand in (3.7.1) is continuous in θ. Thus $\mu((-\infty, x-(s_i-\theta_n)]|s_i)f(s_i|\theta_n) \to \mu((-\infty, x-(s_i-\theta)]|s_i f(s_i|\theta))$ pointwise almost everywhere. Since the $\mu(\cdot|s_i)$ are probability measures, we have $\mu((-\infty, x-(s_i-\theta_n)]|s_i)f(s_i|\theta_n) \leq f(s_i|\theta_n)$. Note that for all s_i, $f(s_i|\theta_n) \to f(s_i|\theta)$ and $1 = \int_{-\infty}^{\infty} f(s_i|\theta_n)\lambda(ds_i) \to \int_{-\infty}^{\infty} f(s_i|\theta)\lambda(ds_i) = 1$. By application of the generalized Lebesgue Convergence Theorem,[24] we have $G_\mu(x;\theta_n) = \int_{-\infty}^{\infty} \mu((-\infty, x-(s_i-\theta_n)]|s_i)f(s_i|\theta_n)\lambda(ds_i) \to \int_{-\infty}^{\infty} \mu((-\infty, x-(s_i-\theta)]|s_i)f(s_i|\theta)\lambda(ds_i) = G_\mu(x;\theta)$. Hence $G_\mu(\cdot;\theta_n)$ converges weakly to $G_\mu(\cdot;\theta)$, as desired. □

Proof of Lemma 3.7.2. We have already noted that these are metric spaces. We first show that $\mathcal{M}^M$ is a relatively compact. Given any $\eta > 0$, choose any compact $T \subset S_i$ such that $\gamma(T) > 1-\eta$. Since X_i^M is compact, $T \times X_i^M$ is compact and $\mu(T \times X_i^M) = \gamma(T) > 1-\eta$ for all $\mu \in \mathcal{M}^M$. By definition, $\mathcal{M}^M$ is thus a tight

[23] A mapping $f: X \to Y$ is sequentially continuous if for all $x \in X$, $x_n \to x$ implies $f(x_n) \to f(x)$. In general, sequential continuity is weaker than continuity.

[24] See Royden and Fitzpatrick (1988, Theorem 4.17).

family of measures. By Prokhorov's Theorem,[25] $\mathcal{M}^M$ is relatively compact. To see that $\mathcal{M}^M$ is closed, take any sequence $(\mu_n)_{n\in\mathbb{N}}$ of measures in $\mathcal{M}^M$ converging weakly to $\mu \in \mathcal{P}$; we show that $\mu \in \mathcal{M}^M$. For any closed $T \subset S_i$, the set $T \times X_i^M$ is closed and thus weak convergence implies[26] $\mu(T \times X_i^M) \geq \limsup_n \mu_n(T \times X_i^M) = \gamma(T)$. Suppose $\mu(T \times X_i^M) > \gamma(T)$ for some such T. If so, then $\mu(T \times X_i^M) + \mu(T^c \times X_i^M) = 1 = \gamma(T) + \gamma(T^c) \implies \mu(T^c \times X_i^M) < \gamma(T^c) - \eta$ for some $\eta > 0$. There exists a closed set $T' \subset T^c$ such that $\gamma(T') > \gamma(T^c) - \eta$, and thus $\mu(T' \times X_i^M) < \gamma(T')$, a contradiction. Hence $\mu(T \times X_i^M) = \gamma(T)$ for all closed $T \subset S_i$, which implies that $\mu \in \mathcal{M}$. Being both closed and relatively compact, $\mathcal{M}$ is compact.

That $\mathcal{G}^M$ is compact follows immediately from the compactness of $\mathcal{M}^M$ combined with the continuity of $G : \mu \mapsto G_\mu$ from Lemma 3.7.1. □

Proof of Lemma 3.7.3. Suppose on the contrary that for each $y \in Y$, there exists $\eta_y > 0$ such that $\lambda\{x \in \mathbb{R} : d(\pi(x), y) < \eta_y\} < \infty$. The collection $\{N_{\eta_y}(y) : y \in Y\}$ is an open covering of Y, and by compactness, it has a finite subcovering denoted $\{N_{\eta_i}(y_i)\}_{i=1}^n$ for some $n \in \mathbb{N}$. It follows that $\mathbb{R} \subseteq \cup_{i=1}^n \pi^{-1}(N_{\eta_i}(y_i))$ and thus $\lambda X \leq \sum_{i=1}^n \lambda\pi^{-1}(N_{\eta_i}(y_i)) < \infty$, a contradiction. □

Proof of Theorem 3.7.1. Given $\mu \in \mathcal{M}_0$, $\mu \in \mathcal{M}^M$ for some $M \in \mathbb{N}$. First, we show that μ is nearly stationary, and then we show payoff equivalence. Suppose otherwise that μ is not nearly stationary. By Lemma 3.7.1, $G_\mu : \theta \mapsto G_\mu(\cdot; \theta)$ is a continuous, and hence measurable, function from $\mathbb{R}$ to $\mathcal{G}^M$, which is compact by Lemma 3.7.2. By Lemma 3.7.3, there exists an attraction $Q \in \mathcal{G}^M$ for μ.

Next, given the existence of an attraction Q, we establish uniqueness. Suppose there also exists an attraction $Q' \in \mathcal{G}^M$ for G_μ with $Q' \neq Q$. We show that there exists a (constant markup) strategy of the rival and a simple princi-

[25]See, for example, Billingsley (2009, Theorem 5.1).

[26]See Billingsley (2009, Theorem 2.1).

pal strategy such that if expert 1 plays μ, she does not have well-defined payoff in the limit. For concreteness, suppose the following: the principal plays C defined by $C(x_1, x_2) = 1$ if $x_1 \leq x_2$ and 0 otherwise; expert 2 plays a constant markup strategy κ_m for some $m \in \mathbb{R}$; and $b_1 = 0$. Let $\nu_n := G_\mu(\cdot;\theta_n) \otimes G_{\kappa_m}$ and $\nu := Q \otimes G_{\kappa_m}$.[27] We show that if (θ_n) is a sequence such that $G_\mu(\cdot;\theta_n) \Rightarrow Q$, then $u(\theta_n) \to u_{Q,\kappa_m} := -\int_{\mathbb{R}^2}(\min\{x_1, x_2\})^2 d\nu(x_1, x_2)$. To prove the convergence, recall that $u(\theta_n) = -\int_{\mathbb{R}^2}(\min\{x_1, x_2\})^2 d\nu_n(x_1, x_2)$. Since $G_\mu(\cdot;\theta_n) \to Q$, we have $\nu_n \Rightarrow \nu$. Moreover, we have $u(\theta_n) \to u_{Q,\kappa_m} < \infty$.[28] Likewise, if (θ_n) is such that $G_\mu(\cdot;\theta_n) \Rightarrow Q'$, $u(\theta_n)$ tends to u_{Q',κ_m}, defined analogously to that above. We claim that $u_{Q,\kappa_m} \neq u_{Q',\kappa_m}$ for some $m \in \mathbb{R}$. Suppose not. Note that for any $Q \in \mathcal{G}_M$, $u_{Q,\kappa_m} = -\int_{\mathbb{R}^2}(\min\{x_1, s_2 + m\})^2 (Q \otimes H)(d(x_1, s_2))$, where H is the CDF of $N(0, \sigma^2)$. Next, we show that we can differentiate w.r.t. m under the integral. Write $\mathbf{x} := (x_1, s_2)$, $g(\mathbf{x}, m) := (\min\{x_1, s_2 + m\})^2$, and $\nu = Q \otimes H$. Note that for all $\mathbf{x} \in \mathbb{R}^2$, $g(\mathbf{x}, m)$ is absolutely continuous in m on bounded intervals. We have

$$a(m) := \int_{\mathbb{R}^2} g(\mathbf{x}, m)\nu(d\mathbf{x}) \tag{3.7.4}$$

$$= \int_{\mathbb{R}^2} \left[g(\mathbf{x}, m_0) + \int_{m_0}^{m} g_m(\mathbf{x}, z)dz \right] \nu(d\mathbf{x}) \tag{3.7.5}$$

$$= \int_{\mathbb{R}^2} g(\mathbf{x}, m_0)\nu(d\mathbf{x}) + \int_{m_0}^{m} \int_{\mathbb{R}^2} g_m(\mathbf{x}, z)\nu(d\mathbf{x})dz \tag{3.7.6}$$

$$\implies a'(m) := \int_{\mathbb{R}^2} g_m(\mathbf{x}, m)\nu(d\mathbf{x}) \qquad \text{a.e. } m, \tag{3.7.7}$$

where $m_0 < m$ can be chosen arbitrarily. To obtain (3.7.5) we have used the fact that absolutely continuous functions are the integral of their derivatives.[29] Fubini's The-

[27] Here we use the fact that $G_{\kappa_m}(\cdot;\theta)$ is constant in θ for any constant markup strategy κ_m.

[28] This follows from the fact that $(\min\{x_1, x_2\})^2 \leq x_1^2 + x_2^2$, the bounds (3.7.2), and the finiteness of second moments for normal distributions.

[29] See Royden and Fitzpatrick (1988, Corollary 5.15).

orem is used to obtain (3.7.6). Differentiability-a.e. of the integral yields (3.7.7).[30] By the same arguments, we obtain $a''(m) = \int g_{mm}(\mathbf{x}, m)\nu(d\mathbf{x}) = 2\int_{\{\mathbf{x}: x_1 > s_2 + m\}} \nu(d\mathbf{x})$ a.e. m. Applying this to $\nu = Q \otimes H$ and $\nu = Q' \otimes H$, it follows that if $u_{Q,\kappa_m} = u_{Q',\kappa_m}$ for all $m \in \mathbb{R}$, then $\int_{\{\mathbf{x}: x_1 > s_2 + m\}} (Q \otimes H)(d(x_1, s_2)) = \int_{\{\mathbf{x}: x_1 > s_2 + m\}} (Q' \otimes H)(d(x_1, s_2))$ a.e. m. Rearranging, we have

$$\int_{S_2} (1 - Q(s_2 + m))a(s_2)ds_2 = \int_{S_2} (1 - Q'(s_2 + m))a(s_2)ds_2$$

$$\implies 0 = \int_{S_2} (Q(m - s_2) - Q'(m - s_2))a(s_2)ds_2, \qquad (3.7.8)$$

by the evenness of the normal distribution. Letting $K := Q - Q'$, we have that the convolution $[K * f](m) = 0$ a.e. m. Let $\mathcal{F}$ denote the normalized Fourier transform, $\mathcal{F}(g)(z) := \int_{-\infty}^{\infty} e^{-2\pi i z x} g(x)dx$. Both K and h are Lebesgue integrable functions,[31] and thus by the convolution theorem,[32] $\mathcal{F}(K * h) = \mathcal{F}(K) \cdot \mathcal{F}(h)$. But since $K * h \equiv 0$, $\mathcal{F}(K * h) \equiv 0$. Since $\mathcal{F}(h)(z) = e^{-2(\pi\sigma z)^2} > 0$ for all z, we must have $\mathcal{F}(K) \equiv 0$. Applying the inverse Fourier transform, we have $K = 0$ almost everywhere. This contradicts the assumption that $Q \neq Q'$, so it must be that $u_{Q,\kappa_m} \neq u_{Q',\kappa_m}$ for some m.

Next, we show that $u_{Q,\kappa_m} \neq u_{Q',\kappa_m}$ contradicts weak admissibility. By continuity of $G_\mu(\cdot, \theta)$ in θ, for each $\eta > 0$, we there exists $\delta > 0$ such that $d_P(G_\mu(\cdot; \theta), Q) < \delta \implies |u(\theta) - u_{Q,\kappa_m}| < \eta$ and $d_P(G_\mu(\cdot; \theta), Q') < \delta \implies |u(\theta) - u_{Q',\kappa_m}| < \eta$.[33] Recall that by weak admissibility, there is some u^* such that $\lim_{n\to\infty} \int u(\theta)d\mathbb{P}_n = u^*$ for all diffusing sequences $(\mathbb{P}_n)$. For the contradiction, we construct two sequence of measures $(\mathbb{P}_n^1)$ and $(\mathbb{P}_n^2)$ along which the limits differ. Since Q is an attrac-

[30] See Royden and Fitzpatrick (1988, Theorem 5.10).

[31] The integrability of H follows from the bounds (3.7.2) applied to Q and Q', and the fact that for all $M > 0$, $\int_{-\infty}^{\infty} (H(x + M) - H(x - M))dx = 2M < \infty$.

[32] See Reed and Simon (1980, Theorem IX.3b).

[33] Here we remind the reader that Q, Q', m, and μ are fixed; θ is arbitrary.

tion and $\lambda\{\theta : d_P(G_\mu(\cdot,\theta),Q) < \delta\} = \infty$, for each $n \in \mathbb{N}$, there exists $C_n^1 \subset \{\theta : d_P(G_\mu(\cdot,\theta),Q) < \delta\} \setminus [-n,n]$ with $\lambda C_n^1 = 2n^2$. Define $B_n^1 := [-n,n] \cup C_n^1$, and define $\mathbb{P}_n^1(\theta) := \mathbf{1}_{B_n^1}(\theta)/\lambda(B_n^1)$. By construction, $(\mathbb{P}_n^1)_{n\in\mathbb{N}}$ is a diffusing sequence of measures, and by the assumption that μ is weakly admissible, we must have $\lim_{n\to\infty}\int u(\theta)d\mathbb{P}_n^1 = u^*$. Now $u(\theta)$ as a function of θ is bounded in magnitude by some $C > 0$. We have $|\int u(\theta)d\mathbb{P}_n^1 - u_{Q,\kappa_m}| \leq \int |u(\theta) - u_{Q,\kappa_m}|d\mathbb{P}_n^1 \leq \frac{\eta\cdot 2n^2+2C\cdot 2n}{2n^2+2n} \to \eta$ and thus $u^* \in [u_{Q,\kappa_m} - \eta, u_{Q,\kappa_m} + \eta]$. Likewise, Q' is an attraction, so $\lambda\{\theta : d_P(G_\mu(\cdot,\theta),Q') < \delta\} = \infty$, and we define $C_n^2 \subset \{\theta : d_P(G_\mu(\cdot,\theta),Q') < \delta\} \setminus [-n,n]$ with $\lambda C_n^2 = 2n^2$, $B_n^2 = [-n,n] \cup C_n^2$, and $\mathbb{P}_n^2(\theta) := \mathbf{1}_{B_n^2}(\theta)/\lambda(B_n^2)$, such that $|\int u(\theta)d\mathbb{P}_n^2 - u_{Q',\kappa_m}| \to \eta$, and thus $u^* \in [u_{Q',\kappa_m} - \eta, u_{Q',\kappa_m} + \eta]$. Since η is arbitrary, we choose $\eta < \frac{|u_{Q,\kappa_m} - u_{Q',\kappa_m}|}{2}$ and obtain a contradiction. Hence we conclude that Q is unique.

In the final step, we prove that μ is nearly stationary, and in particular μ is nearly Q, by showing that for all $\eta > 0$, $\lambda\{\theta : d_P(G_\mu(\cdot,\theta),Q) \geq \eta\} < \infty$. We derive a contradiction by showing that otherwise, the uniqueness result above would be violated. Let $\mathcal{E}$ be the set of all such η; by way of contradiction, we suppose $\bar{\eta} = \inf\{\eta \in \mathcal{E}\} > 0$.[34] Pick $\eta^* \in (0,\bar{\eta})$ and let $Y_{\eta^*} := \{Q \in \mathcal{G}^M : d_P(Q,Q) \geq \eta^*\}$; clearly $Q_0 \notin Y_{\eta^*}$. We have $\lambda\{\theta : G_\mu(\cdot,\theta) \in Y_{\eta^*}\} = \infty$, and Y_{η^*} is a compact subspace of $\mathcal{G}^M$. By Lemma 3.7.3, there exists an attraction $Q' \in Y_{\eta^*}$. This contradicts the uniqueness of the attraction Q, so μ is near Q, completing the proof. □

3.7.5 Proofs for Section 3.7.3

Proof of Lemma 3.7.4. To show that the mapping H_μ is continuous in θ, for any θ, choose any x at which $H_\mu(\epsilon,x;\theta)$ is continuous in (ϵ,x), and consider any sequence $\theta_n \to \theta$, which implies that $\mu((-\infty,x]|s)h(s|\theta_n) \to \mu((-\infty,x]|s)h(s|\theta))$ pointwise in

[34] The metric definition ensures that $\eta^* \leq 1$.

s. Since the $\mu(\cdot|s)$ are probability measures, we have

$$\mu((-\infty, x]|s)h(s|\theta_n) \leq h(s|\theta_n).$$

Note that $1 = \int_{-\infty}^{\infty} h(s_i|\theta_n)\lambda(ds_i) \to \int_{-\infty}^{\infty} h(s_i|\theta)\lambda(ds_i) = 1$. By application of the generalized Lebesgue Convergence Theorem[35] and continuity of the integral,

$$\begin{aligned} H_\mu(\epsilon, x; \theta_n) &= \int_{-\infty}^{\epsilon+\theta_n} \mu((-\infty, x]|s)h(s|\theta_n)\lambda(ds) \\ &\to \int_{-\infty}^{\epsilon+\theta} \mu((-\infty, x]|s)h(s|\theta)\lambda(ds) = H_\mu(\epsilon, x; \theta). \end{aligned}$$

Hence $H_\mu(\cdot, \cdot; \theta_n)$ converges weakly to $H_\mu(\cdot, \cdot; \theta)$, as desired. □

Proof of Theorem 3.7.2. Given an expert strategy μ, we have $\mu \in \mathcal{M}^M$ for some $M \in \mathbb{N}$. Suppose that μ is not nearly stationary. By a repeat of arguments in the proof of Theorem 3.7.1, there exists an attraction $Q \in \mathbb{H}^M$ for H_μ.

To prove uniqueness, we suppose both Q and $Q' \neq Q$ are attractions for μ. Then there exists a pair (ϵ, x) at which both Q' and Q are continuous but $Q'(\epsilon, x) \neq Q(\epsilon, x)$. As in the proof of Theorem 3.7.1, we can construct two sequence of measures $(\mathbb{P}_n^1)$ and $(\mathbb{P}_n^2)$ along which the limits converge to neighborhoods of $Q(\epsilon, x)$ and $Q'(\epsilon, x)$, respectively. This contradicts strong admissibility. Since the notational details are almost identical to those in the proof of Theorem 3.7.1, we do not repeat them here. The proof that μ is near Q is also a repeat of previous arguments and is therefore omitted.

Next, we show that Q is induced by some mixture of constant markup strategies. To do this, we show that there exists a regular conditional probability for Q, interpreted as a CDF $F_{Q,X}(\cdot|\epsilon)$ over markups conditional on ϵ, which is constant in ϵ. Suppose otherwise that for any such $F_{Q,X}(\cdot|\cdot)$, there exists an $x \in X_i^M$ and $\epsilon, \epsilon' \in S_i$

[35] See Royden and Fitzpatrick (1988, Theorem 4.17).

such that $F_{Q,X}(x|t)$ is continuous in x at x for $t = \epsilon, \epsilon'$ but $F_{Q,X}(x|\epsilon) \neq F_{Q,X}(x|\epsilon')$. Since μ is near Q, for any diffusing sequence $(\mathbb{P}_n)$, we have

$$\int_{-\infty}^{\infty} \mu(x|\epsilon + \theta) d\mathbb{P}_n(\theta) \to F_{Q,X}(x|\epsilon).$$

Now consider a change of variables $\theta' = \theta + \epsilon - \epsilon'$:

$$\int_{-\infty}^{\infty} \mu(x|\theta' + \epsilon') d\mathbb{P}_n(\theta') \to F_{Q,X}(x|\epsilon).$$

But $\int_{-\infty}^{\infty} \mu(x|\theta' + \epsilon') d\mathbb{P}_n(\theta') \to F_{Q,X}(x|\epsilon')$, so $F_{Q,X}(x|\epsilon) = F_{Q,X}(x|\epsilon')$, a contradiction. Given the existence of this $F_{Q,X}(\cdot|\cdot)$ constant in ϵ, we define a CDF over markups K_Q by $K_Q(x) := F_{Q,X}(x|\epsilon)$ for all $x \in X_i^M, \epsilon \in S_i$.

The last step of the proof is to show payoff equivalence between μ and Q (and hence K_Q). Take any bias b, strongly admissible rival expert strategy μ_2, and principal strategy C, and fixed θ. Let the variables $Q', Q_2' \in \mathcal{H}^M$ denote an arbitrary joint distribution over noise and markups. Define the expression

$$u(Q', Q_2') := \int \int [- \mathbb{1}\{C(a_1, a_2) = 1\}(a_1 - b)^2 \tag{3.7.9}$$

$$- \mathbb{1}\{C(a_1, a_2) = 2\}(a_2 - b)^2] dQ'(\epsilon_2, x_2) dQ'(\epsilon_1, x_1), \tag{3.7.10}$$

where $a_i = \epsilon_i + x_i$.[36] Note that by familiar convergence arguments, the function u in (3.7.10) is continuous in its arguments and bounded. Evaluating u at $Q' = H_\mu(\cdot, \cdot; \theta)$, $Q_2' = H_{\mu_2}(\cdot, \cdot; \theta)$ gives the expected payoff conditional on θ for expert strategy profile (μ, μ_2); we denote this as $u(\theta) := u(H_\mu(\cdot, \cdot; \theta), H_{\mu_2}(\cdot, \cdot; \theta))$ with some abuse of notation. Now μ is near Q and if μ_2 is strongly admissible, it is near some Q_2. Let $u^* := u(Q, Q_2)$. Now for any $\eta > 0$, by continuity there exists $\delta > 0$ such that if $\max\{d_P(H_\mu(\cdot, \cdot; \theta), Q), d_P(H_{\mu_2}(\cdot, \cdot; \theta), Q_2)\} < \delta$, then $|u(\theta) - u^*| < \eta$. By nearness

[36] Since the principal is restricted to simple strategies, the choice is translation invariant, and we thus drop θ.

to Q and Q_2, $\max\{d_P(H_\mu(\cdot,\cdot;\theta),Q), d_P(H_{\mu_2}(\cdot,\cdot;\theta),Q_2)\} < \delta$ holds for all but a finite measure of θ. Thus for any $\eta > 0$, $|u(\theta) - u^*| < \eta$ for all but a finite measure of θ. Integrating over any diffusing sequence $(\mathbb{P}_n)$, we have $\int u(\theta)d\mathbb{P}_n(\theta) \to u^*$. The same logic shows that $\int u(Q, H_{\mu_2}(\cdot,\cdot;\theta))d\mathbb{P}_n(\theta) \to u^*$, establishing payoff equivalence between μ and Q. □

3.8 Appendix B: Proofs for Sections 3.3-3.5

3.8.1 Proofs for Sections 3.3 and 3.4.1

We first provide an auxiliary lemma, which is proved in Supplementary Appendix.

Lemma 3.8.1. *Let $\epsilon_1, \epsilon_2 \sim N(0,\sigma^2)$, where ϵ_1 and ϵ_2 are independent, and define*

$$\xi(k_1,k_2) := \min(\epsilon_1 + k_1, \epsilon_2 + k_2),\ \eta(k_1,k_2) := \max(\epsilon_1 + k_1, \epsilon_2 + k_2).\ \textit{Then}$$

$$\mathbb{E}\xi(k_1,k_2) = -2\sigma^2 f(k_1-k_2) + k_1(1-F(k_1-k_2)) + k_2F(k_1-k_2);$$

$$\mathbb{E}\eta(k_1,k_2) = 2\sigma^2 f(k_1-k_2) + k_1F(k_1-k_2) + k_2(1-F(k_1-k_2));$$

$$\mathbb{E}\xi^2(k_1,k_2) = \sigma^2 - 2(k_1+k_2)\sigma^2 f(k_1-k_2) + k_1^2(1-F(k_1-k_2)) + k_2^2F(k_1-k_2);$$

$$\mathbb{E}\eta^2(k_1,k_2) = \sigma^2 + 2(k_1+k_2)\sigma^2 f(k_1-k_2) + k_1^2F(k_1-k_2) + k_2^2(1-F(k_1-k_2)).$$

Proof of Proposition 3.3.2. After observing the signal s_i, expert i does a Bayesian update of his beliefs:

$$\theta|s_i \sim N(s_i,\sigma^2) \text{ and } s_j|s_i \sim N(s_i, 2\sigma^2).$$

Since the principal chooses the lower offer, she accepts a_i iff $s_j > s_i + k_i - k_j$. Denote by g the PDF of $N(s_i, 2\sigma^2)$. Hence, the expected utility of expert i

$$U_i(k_1,k_2,L) = \int_{s_i+k_i-k_j}^{\infty} \mathbb{E}\left[B - (a_i - \theta - b_i)^2 | s_i, s_j\right] g(s_j)\, ds_j$$

$$+ \int_{-\infty}^{s_i+k_i-k_j} \mathbb{E}\left[-(a_j - \theta - b_i)^2 | s_i, s_j\right] g(s_j)\, ds_j$$

As $(\theta|s_i, s_j) \sim N(\frac{s_j+s_j}{2}, \frac{\sigma^2}{2})$, $a_i = s_i + k_i$, $a_j = s_j + k_j$, we obtain

$$U_i(k_1, k_2, L) = \int\limits_{s_i+k_i-k_j}^{\infty} \left[B - \left(k_i - b_i - \frac{s_j - s_i}{2}\right)^2 - \frac{\sigma^2}{2}\right] g(s_j)\, ds_j$$

$$+ \int\limits_{-\infty}^{s_i+k_i-k_j} \left[-\left(k_j - b_i + \frac{s_j - s_i}{2}\right)^2 - \frac{\sigma^2}{2}\right] g(s_j)\, ds_j.$$

Now make a substitution $t = s_j - s_i$ and denote by f and F the PDF and CDF of $N(0, 2\sigma^2)$.

$$U_i(k_1, k_2, L) = \int\limits_{k_i-k_j}^{\infty} \left[B - \left(k_i - b_i - \frac{t}{2}\right)^2 - \frac{\sigma^2}{2}\right] f(t)\, dt$$

$$+ \int\limits_{-\infty}^{k_i-k_j} \left[-\left(k_j - b_i + \frac{t}{2}\right)^2 - \frac{\sigma^2}{2}\right] f(t)\, dt.$$

Note that $U_i(k_1, k_2, L)$ does not depend on signal s_i, which is intuitive for the improper prior. As $\int\limits_a^{\infty} t f(t)\, dt = 2\sigma^2 f(a)$ and $\int\limits_{-\infty}^{\infty} t^2 f(t)\, dt = 2\sigma^2$, we get the expression for $U_i(k_i, k_j, L)$.

Now in state θ, the principal's action a is distributed as $\theta + \xi$, where $\xi = \min(\epsilon_1 + k_1, \epsilon_2 + k_2)$; $\epsilon_1, \epsilon_2 \sim N(0, \sigma^2)$, ϵ_1 and ϵ_2 are independent.

Therefore, from Lemma 3.8.1 the expected bias of the accepted offer is

$$b(k_1, k_2, L) = \mathbb{E}\xi(k_1, k_2) = -2\sigma^2 f(k_1 - k_2) + k_2 F(k_1 - k_2) + k_1(1 - F(k_1 - k_2))$$

and the expected utility of the principal is

$$V(k_1, k_2, L) = -\mathbb{E}(a - \theta)^2 = -\mathbb{E}(\theta + \xi - \theta)^2 = -\mathbb{E}\xi^2(k_1, k_2)$$
$$= -\sigma^2 + 2(k_1 + k_2)\sigma^2 f(k_1 - k_2) - k_2^2 - (k_1^2 - k_2^2)(1 - F(k_1 - k_2)).$$

Finally, the variance of the chosen offer is

$$
\begin{aligned}
Var(k_1, k_2, L) =& - V(k_1, k_2, L) - b^2(k_1, k_2, L) \\
=& \sigma^2 - 4\sigma^4 f^2(z) - 2\sigma^2 z f(z)(2F(z) - 1) + z^2 F(z)(1 - F(z)).
\end{aligned}
$$

□

The following lemma provides several useful bounds. They are the immediate corollaries of Sampford (1953).

Lemma 3.8.2. *The following inequalities hold for all $x \in \mathbb{R}$:*

- $0 < v'(x) < \frac{1}{2\sigma^2}$;
- $0 > w'(x) > -\frac{1}{2\sigma^2}$;
- $v''(x) > 0$.

Proof of Theorem 3.3.2.* We start by showing that $U_i(k_1, k_2, L)$ is a single-peaked function of k_i. Taking a derivative w.r.t. k_i yields

$$
\begin{aligned}
U_i'(k_i) &= -2\left[(k_i - b_i)(1 - F(k_i - k_j)) - \rho f(k_i - k_j)\right] \\
&= -2(1 - F(k_i - k_j))\left[k_i - b_i - \rho\frac{f(k_i - k_j)}{1 - F(k_i - k_j)}\right].
\end{aligned}
$$

Let $g(k_i)$ denote the term in square brackets above. Lemma 3.8.2 implies that $g'(k_i) = 1 - \rho\lambda(k_i - k_j) \geq 1 - \sigma^2\lambda(k_i - k_j) > 0$. Additionally, we have $\lim_{x\to\pm\infty} g(x) = \pm\infty$. Combining these facts, U_i has a unique critical point, which is a global maximum.

We now look for upward equilibria. The FOCs for the experts are equivalent to:

$$
k_1 - b_1 - \rho\frac{f(k_1 - k_2)}{1 - F(k_1 - k_2)} = 0 \tag{3.8.1}
$$

$$
k_2 - b_2 - \rho\frac{f(k_1 - k_2)}{F(k_1 - k_2)} = 0. \tag{3.8.2}
$$

Subtracting (3.8.2) from (3.8.1) substituting $z = k_1 - k_2$, we get

$$z - \rho \left[\frac{f(z)}{1 - F(z)} - \frac{f(z)}{F(z)} \right] = b_1 - b_2. \tag{3.8.3}$$

Denote $l(z) = -\rho \left[\frac{f(z)}{1-F(z)} - \frac{f(z)}{F(z)} \right] + z$. Using a) from Lemma 3.8.2, we obtain

$$l'(z) = 1 - \rho \left[v'(z) - w'(z) \right] + 1 \geq 1 - \sigma^2 [v'(z) - w'(z)] > 0$$

Now $l(z)$ is continuous, strictly increasing on $\mathbb{R}$, and ranges from $-\infty$ to $+\infty$. Therefore (3.4.1) has a unique solution, z^*; we use $z(B)$ to denote explicitly the dependence on B.

Using this solution, we get (k_1^U, k_2^U) as the only critical point and check that this point satisfies both initial FOCs. As $U_i(k_i, k_j, L)$ is a single-peaked function of k_i, (k_1^U, k_2^U) is a pair of best responses.

As it was shown in Theorem 3.3.1, choosing the lower offer is the BR strategy for the principal iff $k_1 + k_2 \geq 0$, or equivalently

$$b_1 + b_2 - \left[\frac{f(z^*)}{1 - F(z^*)} + \frac{f(z^*)}{F(z^*)} \right] \geq 0.$$

Also the LHS of (3.4.1) is equal to 0 at $z = 0$, and therefore $z^* \geq 0$ and $k_1^U - k_2^U \geq 0$.

Define a function $m(B) = b_1 + b_2 + \rho[v(z(B)) + w(z(B)]$; the upward equilibrium exists if and only if $m(B) \geq 0$.

1) For $B \leq 2\sigma^2$: $m(B) \geq 0$, therefore the upward equilibrium exists.

2) Next, we show that $m(B)$ is decreasing in B in the region $B \geq 2\sigma^2$.

$$m'(B) = -\frac{1}{2}[v(z(B)) + w(z(B))] + \rho[\lambda(z(B)) - \lambda(-z(B))]z'(B). \tag{3.8.4}$$

Differentiating equation (3.4.1) at point B, we get:

$$z'(B) - \rho[\lambda(z(B)) + \lambda(-z(B))]z'(B) + \frac{1}{2}[v(z(B)) - w(z(B))] = 0. \tag{3.8.5}$$

By substituting (3.8.5), the second term of (3.8.4) becomes

$$+\,\rho[\lambda(z(B)) - \lambda(-z(B))]\frac{-\frac{1}{2}[v(z(B)) - w(z(B))]}{1-\rho\,[\lambda(z(B)) + \lambda(-z(B))]}$$

$$\leq -\rho[\lambda(z(B)) - \lambda(-z(B))]\frac{\frac{1}{2}[v(z(B)) - w(z(B))]}{-\rho\,[\lambda(z(B)) + \lambda(-z(B))]}$$

$$= \frac{1}{2}\frac{\lambda(z(B)) - \lambda(-z(B))}{\lambda(z(B)) + \lambda(-z(B))}[v(z(B)) - w(z(B))]$$

$$\implies m'(B) \leq -\frac{1}{2}[v(z(B)) + w(z(B))] + \frac{1}{2}[v(z(B)) - w(z(B))] = -w(z(B)) < 0.$$

3) From Lemma 3.8.2 the hazard rate v is convex, so for any real x, $v(x) + w(x) = v(x) + v(-x) \geq 2v(0) > 0$, and $m(B)$ tends to $-\infty$ as B tends to ∞.

From 1)-3) follows that there exists $B_U : m(B) \geq 0$ iff $B \leq B_U$. Also

$$\left(\frac{B_U}{2} - \sigma^2\right)[v(z(B_U)) + w(z(B_U))] = b_1 + b_2. \tag{3.8.6}$$

As $z(B_U)$ satisfies equation (3.4.1), we have:

$$\left(\frac{B_U}{2} - \sigma^2\right)[v(z(B_U)) - w(z(B_U))] + z(B_U) = b_1 - b_2. \tag{3.8.7}$$

From the previous discussion and (3.8.6) we have $B_U \geq 2\sigma^2$. Also, (3.8.6) and the inequality $v(x) + w(x) = v(x) + v(-x) \geq 2v(0) = \frac{2}{\sqrt{\pi}\sigma}$ give an upper bound on B_U:

$$\left(\frac{B_U}{2} - \sigma^2\right)\frac{2}{\sqrt{\pi}\sigma} \leq b_1 + b_2.$$

Subtracting (3.8.7) from (3.8.6), we get a lower bound on B_U:

$$2b_2 = (B_U - 2\sigma^2)w(z(B_U)) - z(B_U) \leq (B_U - 2\sigma^2)w(0) = (B_U - 2\sigma^2)\frac{1}{\sqrt{\pi}\sigma}.$$

Finally, we calculate the expected bias of the chosen offer, its variance, and players' utilities:

$$\begin{aligned}
b(k_1^U, k_2^U, L) &= -2\sigma^2 f(z^*) + k_2^U F(z^*) + k_1^U(1 - F(z^*)) \\
&= -2\sigma^2 f(z^*) + b_2 F(z^*) + \rho f(z^*) + b_1(1 - F(z^*)) + \rho f(z^*) \\
&= b_1(1 - F(z^*)) + b_2 F(z^*) - Bf(z^*); \\
Var(k_1^U, k_2^U, L) &= \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1) + (z^*)^2 F(z^*)(1 - F(z^*)).
\end{aligned}$$

□

Corollary 3.8.1. *In upward equilibrium,*

- $V(k_1^U, k_2^U, L) = -\sigma^2 - b_1^2(1 - F(z^*)) - b_2^2 F(z^*) + B(b_1 + b_2)f(z^*)$ $+ \left(\sigma^4 - \frac{B^2}{4}\right) \frac{f^2(z^*)}{F(z^*)(1-F(z^*))}$;
- $U_1(k_1^U, k_2^U, L) = -\sigma^2 - (b_1 - b_2)^2 F(z^*) + B(1 - F(z^*)) - B(b_1 - b_2)f(z^*) + \left(\sigma^4 - \frac{B^2}{4}\right) \frac{f^2(z^*)}{F(z^*)(1-F(z^*))}$;
- $U_2(k_1^U, k_2^U, L) = -\sigma^2 - (b_1 - b_2)^2(1 - F(z^*)) + BF(z^*) + B(b_1 - b_2)f(z^*) + \left(\sigma^4 - \frac{B^2}{4}\right) \frac{f^2(z^*)}{F(z^*)(1-F(z^*))}$.

Proof. Immediate from applying Proposition 3.3.2 to the markups given by Theorem 3.3.2*. □

Proof of Proposition 3.3.3. Since the principal chooses the highest offer, she chooses a_i iff $s_j < s_i + k_i - k_j$. Using arguments similar to used in Proposition 3.3.2, we find the expected utility of expert i:

$$\begin{aligned}U_i(k_1,k_2,H) =& \int_{k_i-k_j}^{\infty} \mathbb{E}[-(a_j-\theta-b_i)^2|s_j]f(s_j)\,ds_j \\ &+\int_{-\infty}^{k_i-k_j} \mathbb{E}[B-(a_i-\theta-b_i)^2|s_j]f(s_j)\,ds_j \\ =& \int_{k_i-k_j}^{\infty}\left[-\left(k_j-b_i+\frac{s_j}{2}\right)^2-\frac{\sigma^2}{2}\right]f(s_j)\,ds_j \\ &+\int_{-\infty}^{k_i-k_j}\left[B-\left(k_i-b_i-\frac{s_j}{2}\right)^2-\frac{\sigma^2}{2}\right]f(s_j)\,ds_j \\ =&(B-(k_i-b_i)^2)F(k_i-k_j)-\sigma^2-(k_j-b_i)^2[1-F(k_i-k_j)] \\ &-2\sigma^2(k_i+k_j-2b_i)f(k_i-k_j).\end{aligned}$$

In state θ the principal's action a is distributed as $\theta+\eta$, where $\eta \sim \max(\epsilon_1+k_1, \epsilon_2+k_2)$; $\epsilon_1, \epsilon_2 \sim N(0,\sigma^2)$, ϵ_1 and ϵ_2 are independent.

From Lemma 3.8.1 the expected bias of the accepted offer is

$$b(k_1,k_2,H) = \mathbb{E}\eta(k_1,k_2) = 2\sigma^2 f(k_1-k_2)+k_2(1-F(k_1-k_2))+k_1F(k_1-k_2).$$

The expected utility of the principal is

$$\begin{aligned}V(k_1,k_2,H) &= -\mathbb{E}(a-\theta)^2 = -\mathbb{E}\eta^2(k_1,k_2) \\ &= -\sigma^2-2(k_1+k_2)\sigma^2 f(k_1-k_2)-k_2^2-(k_1^2-k_2^2)F(k_1-k_2).\end{aligned}$$

The variance of the chosen offer is

$$\begin{aligned}Var(k_1,k_2,H) =& -V(k_1,k_2,H)-b^2(k_1,k_2,H) \\ =&\sigma^2-4\sigma^4 f^2(z)-2\sigma^2 z f(z)(2F(z)-1)+z^2F(z)(1-F(z)).\end{aligned}$$

□

Proof of Theorem 3.3.3.* The proof is analogous to that of Theorem 3.3.2*. The

FOCs for experts are now:

$$k_1 - b_1 + \rho \frac{f(k_1 - k_2)}{F(k_1 - k_2)} = 0 \tag{3.8.8}$$

$$k_2 - b_2 + \rho \frac{f(k_1 - k_2)}{1 - F(k_1 - k_2)} = 0. \tag{3.8.9}$$

Subtracting equation (3.8.8) from equation (3.8.9) yields (3.8.3). Principal optimality holds if and only if $k_1 + k_2 \leq 0$, or equivalently

$$n(B) := b_1 + b_2 + \left[\frac{f(z^*)}{1 - F(z^*)} + \frac{f(z^*)}{F(z^*)} \right] \leq 0, \tag{3.8.10}$$

where z(B) is given by equation (3.4.1). For $B > 2\sigma^2$, we have $n(B) > 0$, and thus a downward equilibrium does not exist. Observe further that $n(2\sigma^2) = b_1 + b_2 \geq 0$. Since $m(B) + n(B) = 2(b_1 + b_2)$ and $m'(B) < 0$, we have $n'(B) > 0$. It follows that if $n(0) \leq 0$, then there exists $B_D \in [0, 2\sigma^2]$ such that $n(B) \leq 0$ iff $B \leq B_D$. Therefore

$$\left(\frac{B_U}{2} - \sigma^2 \right) [v(z(B_D)) + w(z(B_D))] = -(b_1 + b_2). \tag{3.8.11}$$

Also $z(B_D)$ satisfies equation (3.4.1), and therefore

$$\left(\frac{B_U}{2} - \sigma^2 \right) [v(z(B_D)) - w(z(B_D))] + z(B_D) = b_1 - b_2. \tag{3.8.12}$$

From the previous discussion and (3.8.11) we have $B_D \leq 2\sigma^2$. Also (3.8.11) and the inequality $v(x) + w(x) \geq 2v(0) = \frac{2}{\sqrt{\pi}\sigma}$ give the lower bound

$$\left(\frac{B_D}{2} - \sigma^2 \right) \frac{2}{\sqrt{\pi}\sigma} \geq -(b_1 + b_2).$$

Summing (3.8.12) and (3.8.11), we get the upper bound

$$-2b_2 = \left(\frac{B_U}{2} - \sigma^2 \right) v(z(B_D)) + z(B_D) \geq (B_D - 2\sigma^2) 2v(0).$$

Finally, we compute the following:

$$\begin{aligned} b(k_1^D, k_2^D, H) &= 2\sigma^2 f(z^*) + k_1^D F(z^*) + k_2^D (1 - F(z^*)) \\ &= 2\sigma^2 f(z^*) + b_1 F(z^*) - \rho f(z^*) + b_2(1 - F(z^*)) - \rho f(z^*) \\ &= b_1 F(z^*) + b_2(1 - F(z^*)) + B f(z^*); \\ Var(k_1^D, k_2^D, H) &= \sigma^2 - 4\sigma^4 f^2(z^*) - 2\sigma^2 z^* f(z^*)(2F(z^*) - 1) + (z^*)^2 F(z^*)(1 - F(z^*)). \end{aligned}$$

□

Corollary 3.8.2. *In downward equilibrium,*

$$V(k_1^D, k_2^D, H) = -\sigma^2 - b_1^2 F(z^*) - b_2^2(1 - F(z^*)) - B(b_1 + b_2) f(z^*) + P(z),$$

$$U_1(k_1^D, k_2^D, H) = -\sigma^2 - (b_1 - b_2)^2 (1 - F(z^*)) + BF(z^*) + B(b_1 - b_2) f(z^*) + P(z),$$

$$U_2(k_1^D, k_2^D, H) = -\sigma^2 - (b_1 - b_2)^2 F(z^*) + B(1 - F(z^*)) - B(b_1 - b_2) f(z^*) + P(z),$$

where $P(z) := \left(\sigma^4 - \frac{B^2}{4}\right) \frac{f^2(z^*)}{F(z^*)(1-F(z^*))}$.

Proof. Proposition 3.3.3 applied to Theorem 3.3.3*. □

Proof of Proposition 3.3.4.* From Theorems 3.3.2* and 3.3.3*

$$V_{upw.} - V_{downw.} = (2F(z^*) - 1)(b_1^2 - b_2^2) + 2B(b_1 + b_2) f(z^*) \geq 0,$$

with equality if and only if either $b_1 + b_2 = 0$ or both $B = 0$ and $b_1 = b_2$. □

Proof of Proposition 3.3.5. Delegation to expert 2 alone yields principal utility $V^S(b, x) = -\sigma^2 - (b - x)^2$, while upward equilibrium yields

$$V^U(b, x) = -\sigma^2 - b^2 - x^2 + (2F(z^*) - 1)2bx + (k_1 - b - x)(k_2 - b + x).$$

The difference between these is

$$\begin{aligned} V^U(b, x) - V^S(b, x) &= -2(1 - F(z^*))bx + (k_1 - b - x)(k_2 - b + x) \\ &= -2(1 - F(z^*))bx + \sigma^4 v(z^*) w(z^*). \end{aligned}$$

Recall that z^* is independent of b. If $x = 0$, then the above expression is always positive. For fixed $x > 0$ then the existence of $\bar{b}$ follows from the fact that this expression is decreasing linearly in b and positive for $b = 0$.

Next, consider $b_2 = 0$ and $b_1 = b > 0$. The principal's utility in upward equilibrium is

$$V = -\sigma^2 + 2(k_1 + k_2)\sigma^2 f(z) - k_1^2(1 - F(z)) - k_2^2 F(z),$$

which we aim to show is greater than $-\sigma^2$ as under simple delegation to expert two. Using the expressions $k_1 = b + \sigma^2 \frac{f(z)}{1-F(z)}$ and $k_2 = \sigma^2 \frac{f(z)}{F(z)}$, this is true if and only if

$$2\left(b + \sigma^2 \frac{f}{1-F} + \sigma^2 \frac{f}{F}\right)\sigma^2 f > \left(b^2 + 2b\sigma^2 \frac{f}{1-F} + \sigma^4 \left(\frac{f}{1-F}\right)^2\right)(1-F)$$
$$+ \sigma^4 \left(\frac{f}{F}\right)^2 F.$$

Using $b = z - \sigma^2 \frac{f(2F-1)}{F(1-F)}$ and simplifying, this is equivalent to

$$\sigma^4(4F - 1) > z(zF(1-F) - 2\sigma^2 f(2F-1)).$$

As $z > 0$, the left hand side is positive; we now show that the right hand side is negative. Let $h(z) := 2\sigma^2 f(2F-1) - zF(1-F)$, which we aim to show is positive. Then $h'(z) = 2f^2 - F(1-F)$. As shown in Sampford (1953), $k(z) := \frac{f^2}{F(1-F)}$ is decreasing for $z \geq 0$. It is easy to verify that $2k(0) > 1$ and that $\lim_{z\to+\infty} k(z) = 0$. It follows that there is a unique positive solution to $h'(z) = 0$. It is also easy to check that $h'(0) > 0$ and that $\lim_{z\to+\infty} h(z) = 0$. Together these facts imply that $h(z) > 0$ for all $z > 0$, as desired. □

Proof of Proposition 3.4.1. If $b_1 = b_2 = b > 0$, then upper and lower bounds on B_U coincide, and thus $B_U = 2\sigma^2 + 2\sqrt{\pi}\sigma b$. From Theorem 3.3.2*, the experts' markups are $k_1^U = k_2^U = k_U = b + \frac{\rho\sigma}{\sigma\sqrt{\pi}}$ and $z^* = 0$. The other results follow immediately. □

Proof of Proposition 3.4.2. If $b_1 = b_2 = b > 0$, then upper and lower bounds on B_D from Theorem 3.3.3* coincide, so $B_D = 2\sigma^2 + 2\sqrt{\pi}\sigma b$. The experts' markups are $k_1^D = k_2^D = k_D = b - \frac{\rho}{\sigma\sqrt{\pi}}$ and $z^* = 0$, which implies all other results. □

3.8.2 Proofs for Sections 3.4.2 and 3.4.3

Proof of Proposition 3.4.4. Let $V(p)$ denote the principal's utility from commitment to p. It suffices to show that $V'(1) < 0$. Given p, the markups satisfy

$$k_1(p) = b + \frac{(2p-1)f(z(p))}{W(p)}$$

$$k_2(p) = -b + \frac{(2p-1)f(z(p))}{1-W(p)}$$

$$\implies z(p) = 2b + \frac{(2p-1)f(z)(1-2W(p))}{W(p)(1-W(p))},$$

where $W(p) := p(1-F(z(p))) + (1-p)F(z(p))$. Differentiating with respect to p and solving for $z'(1)$ yields

$$z'(1) = \frac{f(4F+1)(2F-1)}{F^2(1-F)^2 - F(1-F)(f'(2F-1)+2f^2) - f^2(2F-1)^2}.$$

The principal's utility is

$$V(p) = 2(2p-1)f(k_1+k_2) - k_1^2 W - k_2^2(1-W).$$

By differentiating with respect to p, evaluating at $p=1$, substituting in the above expression for $z'(1)$ and simplifying, we obtain

$$V'(1) = \frac{f^2}{F^2(1-F)^2}\left(1 + \frac{g_1}{g_2}\right), \text{ where}$$

$$g_1 := (2F-1)(2F(1-F)f' + f^2(2F-1)),$$

$$g_2 := F^2(1-F)^2 - f^2(2F^2-2F+1) - f'(2F-1)F(1-F).$$

We claim that $g_2 > 0$, and thus it suffices to show that $g_2 < -g_1$ for sufficiently large z. To see this, note that by Lemma 3.8.2, $f' < \frac{1-F}{2} - \frac{f^2}{1-F}$, and thus

$$g_2 > F(1-F)^2 \left[\frac{1}{2} - \frac{f^2}{F(1-F)} \right].$$

It is easy to verify that $\frac{f^2}{F(1-F)} < \frac{1}{2}$ holds globally, and thus $g_2 > 0$ as desired. To see that $g_2 < -g_1$ for sufficiently large z, note that by algebra this comparison is equivalent to

$$F(1-F) + f'(2F-1) \leq 2f^2. \tag{3.8.13}$$

Using Lemma 3.8.2 again and simplifying, a sufficient condition for (3.8.13) is $2F - \frac{1}{2} < \frac{f^2}{(1-F)^2}$. The left hand side is bounded above by $\frac{3}{2}$, while the right hand side is increasing and unbounded above; thus for sufficiently large z, (3.8.13) holds. Finally, since z is increasing and unbounded above as a function of b, the proposition holds.

□

Proof of Proposition 3.4.5. First, we calculate marginal utilities:

$$U_i'(k_i) = -2(k_i - b)[1 - F(k_i - k_j)] + \left[\frac{1}{2}\sigma^2 + 2\rho + \frac{1}{4}(k_i + k_j - 2b)^2 \right] f(k_i - k_j).$$

Here, setting $U_i'(k) = 0$ gives two critical points:

$$k = b + \sqrt{\pi}\sigma - \sqrt{\left(\pi - \tfrac{5}{2}\right)\sigma^2 + B} \text{ and } k = b + \sqrt{\pi}\sigma + \sqrt{\left(\pi - \tfrac{5}{2}\right)\sigma^2 + B}.$$

The second derivative is:

$$U_i''(k_i) = -2[1 - F(k_i - k_j)] + f(k_i - k_j) \times \ldots$$

$$\ldots \left[2(k_i - b) + \frac{1}{2}(k_i + k_j - 2b) + \left(\frac{B}{2\sigma^2} - \frac{5}{4} \right)(k_i - k_j) - \frac{1}{8\sigma^2}(k_i - k_j)(k_i + k_j - 2b)^2 \right].$$

We get that only $k^* = b + \sqrt{\pi}\sigma - \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$ is a local maximum of experts' utility functions.

Optimality for the principal holds if and only if $k^* \geq 0$ or, equivalently, $B \leq b^2 + 2\sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2$.

Calculating, we get that $U_i(k^*, k^*, L) = \frac{(\pi-1)\sigma}{\sqrt{\pi}}\sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B} - (\pi - \frac{7}{4})\sigma^2$. As we noted earlier, a necessary condition for equilibrium is $U_i(k^*, k^*, L) \geq 0$ or, equivalently, $B \geq (\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2})\sigma^2$.

Therefore, upward equilibrium may exist only if

$$B \in \left[\left(\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2}\right)\sigma^2, \frac{5}{2}\sigma^2 + 2\sqrt{\pi}b\sigma + b^2\right].$$

To finish the proof, we show that if B lies on this interval, then $k = k^*$ is a global maximum of $U_1(k, k^*, L)$.

Denote $g(k) = -2(k-b) + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] v(k - k^*)$.[37]

Then $U_1'(k) = -2(k-b)[1 - F(k-k^*)] + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] f(k-k^*) = [1 - F(k-k^*)]g(k)$ and $sign(U_1'(k)) = sign(g(k))$.

The first and second derivatives of g are

$$g'(k) = -2 + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] v'(k-k^*) + \frac{1}{2}(k + k^* - 2b)v(k-k^*)$$

$$g''(k) = \left[\frac{\sigma^2}{2} + 2\rho - B + \frac{1}{4}(k + k^* - 2b)^2\right] v''(k-k^*) + (k + k^* - 2b)v'(k-k^*) + \frac{1}{2}v(k-k^*).$$

Consider two cases.

1. $B \in \left[\left(\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2}\right)\sigma^2, \frac{5}{2}\sigma^2\right]$. Here $k^* \geq b$.

a) On the interval $k < b$, $U_1'(k) > 0$ and hence there is no point of maximum there.

[37] Recall that $v(k-k^*) = \frac{f(k-k^*)}{1-F(k-k^*)})$.

b) On the interval $k \geq b$ we also have that $k + k^* - 2b \geq 0$. As all v, v' and v'' are strictly positive functions, $g''(k) > 0$. As $g'(k^*) < 0$ and $g'(+\infty) > 0$, hence there exists $k^{**} > k^*$: for $k < k^{**}$ $g(k)$ is decreasing; for $k > k^{**}$, $g(k)$ is increasing. As $g(b) > 0$, $g(k^*) = 0$, $g(k^{**}) < 0$ and $g(+\infty) > 0$, then there exists $k_0 > k^{**} : g(k_0) = 0$. In summary, $g(k)$ is negative only on (k^*, k_0). Consequently, $U_1(k)$ is increasing on $[b, k^*)$, decreasing on (k^*, k_0), increasing for $k > k_0$. Hence, to show that k^* is a maximum on the interval $k \geq b$ it is sufficient to verify that $U_1(k^*) \geq U_1(+\infty) = 0$, which was already done.

2. $B \in \left[\frac{5}{2}\sigma^2, \frac{5}{2}\sigma^2 + 2\sqrt{\pi}b\sigma + b^2\right]$. Here $k^* \leq b$.

a) On the interval $k < k^*$:

$$\begin{aligned} U_1'(k) >& 2(b - k^*)[1 - F(k^* - k^*)] + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k^* + k^* - 2b)^2\right] f(k - k^*) \\ =& b - k^* - \frac{b - k^*}{f(0)} f(k - k^*) \end{aligned}$$

as k^* is a solution of $\frac{5}{2}\sigma^2 - B + (k^* - b)^2 = \frac{k^* - b}{f(0)}$. Therefore $U_1'(k) > b - k^* - \frac{b-k^*}{f(0)} f(k - k^*) \geq b - k^* - \frac{b-k^*}{f(0)} f(0) = 0$.

b) On the interval $k \in (k^*, b]$:

$$\begin{aligned} \frac{U_1'(k)}{f(k - k^*)} =& -2(k - b)\frac{1 - F(k - k^*)}{f(k - k^*)} + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] \\ <& 2(b - k^*)\frac{1 - F(k^* - k^*)}{f(k^* - k^*)} + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k^* + k^* - 2b)^2\right] \\ =& \frac{b - k^*}{f(0)} - \frac{b - k^*}{f(0)} = 0, \end{aligned}$$

and hence $U_1'(k) < 0$.

c) On the interval $k \in \left(b, 2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}\right]$ we also have

$$\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2 \leq 0.$$

Then

$$U_1'(k) = -2(k-b)[1 - F(k-k^*)] + \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k+k^*-2b)^2\right] f(k-k^*)$$
$$<0.$$

d) On the interval $k > 2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}$ we also have $k + k^* - 2b > 0$. Hence, on this interval $g''(k) > 0$. Also notice that $g(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) < 0$. Then two cases are possible: (i) $g'(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) \geq 0$. Then on the whole interval $g'(k) > 0$ and $g(k)$ is increasing. As $g(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) < 0$ and $g(+\infty) > 0$, there exists k_0: $g(k) < 0$ for $k < k_0$ and $g(k) > 0$ for $k > k_0$. Now $U_1(k)$ is decreasing for $k < k_0$ and increasing for $k > k_0$. Hence, to show that k^* is a global maximum it is enough to check that $U_1(k^*) \geq U_1(+\infty) = 0$, which was already done.

(ii) $g'(2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}) < 0$. Then there exists $k^{**} > 2b - k^* + 2\sqrt{B - \frac{5}{2}\sigma^2}$: for $k < k^{**}$ $g(k)$ is decreasing; for $k > k^{**}$ $g(k)$ is increasing. As $g(+\infty) > 0$, there exists k_0: $g(k) < 0$ for $k < k_0$ and $g(k) > 0$ for $k > k_0$. Then $U_1(k)$ is decreasing for $k < k_0$ and increasing for $k > k_0$ and k^* is a global maximum as $U_1(k^*) \geq U_1(+\infty) = 0$.

We now verify that $k_U = k^* > k_U^{bas.}$:

$$k_U = b + \left(\sqrt{\pi} - \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}}\right)\sigma > k_U^{bas.} = b + \left(1 - \frac{B}{2\sigma^2}\right)\frac{\sigma}{\sqrt{\pi}}$$

$$\iff \sqrt{\pi} - \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}} > \left(1 - \frac{B}{2\sigma^2}\right)\frac{1}{\sqrt{\pi}}$$

$$\iff \sqrt{\pi} + \frac{1}{\sqrt{\pi}}\left(\frac{B}{2\sigma^2} - 1\right) > \sqrt{\pi + \frac{B}{\sigma^2} - \frac{5}{2}}$$

$$\iff \pi + \frac{B}{\sigma^2} - 2 + \frac{1}{\pi}\left(\frac{B}{2\sigma^2} - 1\right)^2 > \pi + \frac{B}{\sigma^2} - \frac{5}{2}$$

$$\iff \frac{1}{2} + \left(\frac{B}{2\sigma^2} - 1\right)^2 > 0.$$

□

Proof of Proposition 3.4.6. Start with calculation of marginal utilities:

$$U_i'(k_i) = -2(k_i - b_i)F(k_i - k_j) - \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k_i + k_j - 2b)^2\right] f(k_i - k_j)$$

Consider the symmetric case: $k_1 = k_2 = k$. The FOCs give two critical points:

$k = b - \sqrt{\pi}\sigma - \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$ and $k = b - \sqrt{\pi}\sigma + \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$.

Next, calculate second derivatives:

$$U_i''(k_i) = -2F(k_i - k_j) - f(k_i - k_j)\times$$

$$\Bigg[2(k_i - b_i) + \frac{1}{2}(k_i + k_j - 2b_i)\dots$$

$$\dots + \left(\frac{B}{2\sigma^2} - \frac{5}{4}\right)(k_i - k_j) - \frac{1}{8\sigma^2}(k_i - k_j)(k_i + k_j - 2b_i)^2\Bigg].$$

We get that only $k^* = b - \sqrt{\pi}\sigma + \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B}$ satisfies SOCs.

In order to satisfy principal optimality we need $k^* \le 0$ or, equivalently, both

$\frac{b}{\sigma} \leq \sqrt{\pi}$ and $B \leq b^2 - \sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2$.

Calculating, we get that $U_i(k^*, k^*, H) = \frac{(\pi-1)\sigma}{\sqrt{\pi}}\sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B} - (\pi - \frac{7}{4})\sigma^2$ (the same as in upward equilibrium). As in upward equilibrium case, a necessary condition is $B \geq \left(\frac{5}{2} - \frac{3\pi(8\pi-11)}{16(\pi-1)^2}\right)\sigma^2$.

From previous arguments downward equilibrium may exist only if

$$B \in \left[(\frac{5}{2} - \frac{3\pi(8\pi - 11)}{16(\pi - 1)^2}\right)\sigma^2, b^2 - 2\sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2)$$

and $\frac{b}{\sigma} \leq \sqrt{\pi}$. This interval is non-empty if and only if $\frac{b}{\sigma} \leq \frac{3\sqrt{\pi}}{4(\pi-1)}$. Note also that $B \leq b^2 - 2\sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2 \leq \frac{5}{2}\sigma^2$.

To finish the proof we show that if B lies on this interval, then $k = k^*$ is not only a local, but also a global maximum of $U_1(k, k^*, L)$.

Denote $r(k) = -2(k - b) - \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] w(k - k^*)$ (recall that $w(k - k^*) = \frac{f(k-k^*)}{F(k-k^*)}$).

Then $U_1'(k) = -2(k - b)F(k - k^*) - \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] f(k - k^*) = F(k - k^*)r(k)$ and $sign(U_1'(k)) = sign(r(k))$.

First and second derivatives of $r(k)$ are:

$r'(k) = -2 - \left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] w'(k - k^*) - \frac{1}{2}(k + k^* - 2b)w(k - k^*)$;

$r''(k) = -\left[\frac{\sigma^2}{2} + 2\rho + \frac{1}{4}(k + k^* - 2b)^2\right] w''(k - k^*) - (k + k^* - 2b)w'(k - k^*) - \frac{1}{2}w(k - k^*)$.

Notice that as $\frac{b}{\sigma} \leq \sqrt{\pi}$, $B \leq b^2 - \sqrt{\pi}b\sigma + \frac{5}{2}\sigma^2 \leq \frac{5}{2}\sigma^2$ and $k^* = b - \sqrt{\pi}\sigma + \sqrt{\left(\pi - \frac{5}{2}\right)\sigma^2 + B} \leq b$.

a) On interval $k > b$ $U'(k) < 0$, so there is no candidate for maximum there.

b) On interval $k < b$ we also have $k + k^* - 2b \leq 0$. As $w > 0$, $w' < 0$, $w'' > 0$, we have $r''(k) < 0$.

As $r'(-\infty) > 0$ and $r'(k^*) < 0$, there exists $k^{**} < k^* < b$: $r(k)$ is increasing for $k < k^{**}$, $r(k)$ is decreasing for $k > k^{**}$. As also $r(-\infty) < 0$, $r(k^* - 0) > 0$ and $r(k^* + 0) < 0$, there exists k_0: $r(k) > 0$ only on (k_0, k^*). Therefore, $U_1(k)$ is decreasing on $k < k_0$, increasing on (k_0, k^*), decreasing on (k^*, b). Hence, k^* is a global maximum if $U_1(k^*) \geq U_1(-\infty) = 0$, which has already been shown. □

3.8.3 Proofs for Section 3.5

Proof of Proposition 3.5.2. From Proposition 3.4.1, the symmetric upward equilibrium exists if $B \leq 2\sigma^2 + 2\sqrt{\pi}b\sigma$ or, equivalently, $\sigma \geq \frac{\sqrt{\pi b^2+2B}-\sqrt{\pi}b}{2}$. The principal's expected payoff in upward equilibrium is equal to $V = -(b - \frac{B}{2\sqrt{\pi}\sigma})^2 - (1 - \frac{1}{\pi})\sigma^2$. Then $V'(\sigma) = \frac{B}{\sqrt{\pi}\sigma^2}(\frac{B}{2\sqrt{\pi}\sigma} - b) - 2(1 - \frac{1}{\pi})\sigma = -\frac{1}{2\pi\sigma^3}[(4\pi - 4)\sigma^4 + 2\sqrt{\pi}bB\sigma - B^2] > 0$ if and only if $\sigma < \sigma^*$.

If $B = 0$, $V'(\sigma) < 0$. Otherwise, denote $\sigma_0 = \frac{\sqrt{\pi b^2+2B}-\sqrt{\pi}b}{2} > 0$.

The interval $[\sigma_0, \sigma^*)$ is non-empty if and only if $0 > (4\pi-4)\sigma_0^4+2\sqrt{\pi}bB\sigma_0-B^2 = (4\pi - 4)\sigma_0^4 - 2B\sigma_0^2 + B[2\sigma_0^2 + 2\sqrt{\pi}b\sigma_0 - B] = (4\pi - 4)\sigma_0^4 - 2B\sigma_0^2$ or, equivalently, $0 > (2\pi-2)\sigma_0^2-B = (2\pi-2)\sigma_0^2-(2\sigma_0^2+2\sqrt{\pi}b\sigma_0) = (2\pi-4)\sigma_0^2-2\sqrt{\pi}b\sigma_0$. The latter holds if and only if $\sigma_0 < \frac{\sqrt{\pi}b}{\pi-2}$ or, equivalently, $B = 2\sigma_0^2 + 2\sqrt{\pi}b\sigma_0 < \frac{2(\pi-1)\pi}{(\pi-2)^2}b^2$. □

Proof of Proposition 3.5.3. By Proposition 3.3.4, only upward equilibrium should be considered. Therefore, we seek to maximize the quadratic function $V(k_U, k_U, L) - B = -\left(b - \frac{B}{2\sqrt{\pi}\sigma}\right)^2 - \sigma^2 + \frac{\sigma^2}{\pi} - B$ on the interval $B \in [0, 2\sigma^2 + 2\sqrt{\pi}\sigma b]$.

a) First consider $\frac{b}{\sigma} \leq \sqrt{\pi}$. In this case the maximum is achieved at $B = 0$ and is equal to $V(k_U, k_U, L|B = 0) = -b^2 - \left(1 - \frac{1}{\pi}\right)\sigma^2$.

Hence, if $\frac{b}{\sigma} < \sqrt{\pi}$, then the principal pays no bonus.

b) Now consider $\frac{b}{\sigma} \geq \sqrt{\pi}$. In this case the principal achieves maximum in upward equilibrium at $B_R = 2\sqrt{\pi}\sigma\left(b - \sqrt{\pi}\sigma\right)$ (upward equilibrium exists for this point) and is equal to $V(k_U, k_U, L|B = B_R) = (\pi - 1 + \frac{1}{\pi})\sigma^2 - 2\sqrt{\pi}\sigma b$.

Her gains comparatively to $B = 0$ (if she is legally restricted from paying bonuses) are equal to:

$$V(k_U, k_U, L|B = B_R) - V(k_U, k_U, L|B = 0) = \left(b - \sqrt{\pi}\sigma\right)^2.$$

□

Bibliography

Akerlof, G. (1970), "The market for "lemons": Quality uncertainty and the market mechanism," *The quarterly journal of economics*, pp. 488–500.

Alonso, R. and Matouschek, N. (2008), "Optimal delegation," *The Review of Economic Studies*, 75, 259–293.

Ambrus, A. and Lu, S. (2014), "Almost fully revealing cheap talk with imperfectly informed senders," *Games and Economic Behavior*, 88, 174–189.

Ambrus, A. and Takahashi, S. (2008), "Multi-sender cheap talk with restricted state spaces," *Theoretical Economics*, 3, 1–27.

Ambrus, A., Burns, J., and Ishii, Y. (2013), "Gradual bidding in ebay-like auctions," Working paper.

Appuhamillage, T., Bokil, V., Thomann, E., Waymire, E., and Wood, B. (2010), "Solute transport across an interface: A Fickian theory for skewness in breakthrough curves," *Water Resources Research*, 46.

Appuhamillage, T., Bokil, V., Thomann, E., Waymire, E., and Wood, B. (2011), "Occupation and local times for skew Brownian motion with applications to dispersion across an interface," *The Annals of Applied Probability*, 21, 183–214.

Aumann, R. J. and Maschler, M. (1995), *Repeated games with incomplete information*, MIT press.

Austen-Smith, D. (1993), "Interested experts and policy advice: Multiple referrals under open rule," *Games and Economic Behavior*, 5, 3–43.

Banks, J. S. and Sobel, J. (1987), "Equilibrium selection in signaling games," *Econometrica: Journal of the Econometric Society*, pp. 647–661.

Bar-Isaac, H. (2003), "Reputation and Survival: learning in a dynamic signalling model," *The Review of Economic Studies*, 70, 231–251.

Battaglini, M. (2002), "Multiple referrals and multidimensional cheap talk," *Econometrica*, 70, 1379–1401.

Battaglini, M. (2004), "Policy advice with imperfectly informed experts," *Advances in Theoretical Economics*, 4, 1–32.

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

Bergemann, D. and Hege, U. (2005), "The financing of innovation: Learning and stopping," *RAND Journal of Economics*, pp. 719–752.

Bergemann, D. and Välimäki, J. (2002), "Entry and vertical differentiation," *Journal of Economic Theory*, 106, 91–125.

Berger, J., Sorensen, A. T., and Rasmussen, S. J. (2010), "Positive effects of negative publicity: When negative reviews increase sales," *Marketing Science*, 29, 815–827.

Bernhardt, D., Campello, M., and Kutsoati, E. (2006), "Who herds?" *Journal of Financial Economics*, 80, 657–675.

Bhattacharya, S. and Mukherjee, A. (2013), "Strategic information revelation when experts compete to influence," *The RAND Journal of Economics*, 44, 522–544.

Billingsley, P. (2009), *Convergence of probability measures*, vol. 493, John Wiley & Sons.

Board, S. and Meyer-ter Vehn, M. (2013), "Reputation for quality," *Econometrica*, 81, 2381–2462.

Board, S. and Meyer-ter Vehn, M. (2014), "A reputational theory of firm dynamics," Tech. rep., working paper, UCLA.

Bohren, A. (2012), "Stochastic Games in Continuous Time: Persistent Actions in Long-Run Relationships," Tech. rep.

Bolton, P. (1987), "The principle of maximum deterrence revisited," Tech. rep., working paper, University of California, Berkeley.

Bonatti, A. and Hörner, J. (2011), "Collaborating," *The American Economic Review*, 101, 632–663.

Calvert, R. L. (1985), "Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence," *American Journal of Political Science*, 29, 69–95.

Cantrell, R. and Cosner, C. (1999), "Diffusion models for population dynamics incorporating individual behavior at boundaries: applications to refuge design," *Theoretical Population Biology*, 55, 189–207.

Chen, Q. and Jiang, W. (2006), “Analysts’ weighting of private and public information,” *Review of Financial Studies*, 19, 319–355.

Cho, I.-K. and Kreps, D. M. (1987), “Signaling games and stable equilibria,” *The Quarterly Journal of Economics*, 102, 179–221.

Cisternas, G. (2012), “Two-sided Learning and Moral Hazard,” Tech. rep.

Corns, T. and Satchell, S. (2007), “Skew Brownian motion and pricing European options,” *The European Journal of Finance*, 13, 523–544.

Crawford, V. and Sobel, J. (1982), “Strategic information transmission,” *Econometrica: Journal of the Econometric Society*, pp. 1431–1451.

Daley, B. and Green, B. (2012), “Waiting for News in the Market for Lemons,” *Econometrica*, 80, 1433–1504.

Daley, B. and Green, B. (2015), “Bargaining with the Arrival of News,” Tech. rep.

Dessein, W. (2002), “Authority and communication in organizations,” *The Review of Economic Studies*, 69, 811–838.

Dewatripont, M. and Tirole, J. (1999), “Advocates,” *Journal of Political Economy*, 107, 1–39.

Dilmé, F. (2014), “Reputation Building through Costly Adjustment,” .

Dixit, A. (1989), “Entry and exit decisions under uncertainty,” *Journal of political Economy*, pp. 620–638.

Dixit, A. (1993), *Art of Smooth Pasting*, vol. 55, Routledge.

Downs, A. (1957), *An economic theory of democracy*, New York: Harper & Row.

Dumas, B. (1991), “Super contact and related optimality conditions,” *Journal of Economic Dynamics and Control*, 15, 675–685.

Durrett, R. (2010), *Probability: theory and examples*, Cambridge University Press.

Dynkin, E. (1969), “Game variant of a problem on optimal stopping,” in *Soviet Math. Dokl.*, vol. 10, pp. 270–274.

Ederer, F., Holden, R., and Meyer, M. A. (2014), “Gaming and strategic opacity in incentive provision,” .

Esponda, I. (2008), “Behavioral equilibrium in economies with adverse selection,” *The American Economic Review*, 98, 1269–1291.

Eyster, E. and Rabin, M. (2005), "Cursed equilibrium," *Econometrica*, 73, 1623–1672.

Faingold, E. and Sannikov, Y. (2011), "Reputation in Continuous-Time Games," *Econometrica*, 79, 773–876.

Feltovich, N., Harbaugh, R., and To, T. (2002), "Too Cool for School? Signalling and Countersignalling," *RAND Journal of Economics*, 33, 630–649.

Fontana, R. and Nesta, L. (2009), "Product innovation and survival in a high-tech industry," *Review of Industrial Organization*, 34, 287–306.

Frick, M. and Ishii, Y. (2015), "Innovation Adoption by Forward-Looking Social Learners," Tech. rep., working paper.

Friedman, D. (1991), "A simple testable model of double auction markets," *Journal of Economic Behavior and Organization*, 15, 47–70.

Gentzkow, M. and Kamenica, E. (2015), "Competition in persuasion," Working paper.

Georgiadis, G., Lippman, S. A., and Tang, C. S. (2014), "Project design with limited commitment and teams," *The RAND Journal of Economics*, 45, 598–623.

Gerardi, D., McLean, R., and Postlewaite, A. (2009), "Aggregation of expert opinions," *Games and Economic Behavior*, 65, 339–371.

Gilligan, T. and Krehbiel, K. (1989), "Asymmetric information and legislative rules with a heterogeneous committee," *American Journal of Political Science*, pp. 459–490.

Grigg, O. and Farewell, V. (2004), "An overview of risk-adjusted charts," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 167, 523–539.

Gryglewicz, S. (2009), "Signaling in a stochastic environment and dynamic limit pricing," Tech. rep., mimeo, Tilburg University.

Gul, F. and Pesendorfer, W. (2012), "The war of information," *The Review of Economic Studies*, 79, 707–734.

Harrison, J. M. (2013), *Brownian Models of Performance and Control*, Cambridge University Press.

Harrison, J. M. and Shepp, L. A. (1981), "On skew Brownian motion," *The Annals of probability*, pp. 309–313.

Hauser, J., Tellis, G. J., and Griffin, A. (2006), "Research on innovation: A review and agenda for marketing science," *Marketing science*, 25, 687–717.

Heinsalu, S. (2014), "Noisy signalling over time," Working paper.

Helfat, C. E. and Lieberman, M. B. (2002), "The birth of capabilities: market entry and the importance of pre-history," *Industrial and Corporate Change*, 11, 725–760.

Holmström, B. (1977), *On incentives and control in organizations*, Stanford University.

Hörner, J., Rosenberg, D., Solan, E., and Vieille, N. (2010), "On a Markov game with one-sided information," *Operations research*, 58, 1107–1115.

Hotelling, H. (1929), "Stability in competition," *The Economic Journal*, 39, 41–57.

Iezzoni, L., Rao, S., DesRoches, C., Vogeli, C., and Campbell, E. (2012), "Survey shows that at least some physicians are not always open or honest with patients," *Health Affairs*, 31, 383–391.

Ito, K. and McKean, H. (1965), "Diffusion Processes and Their Sample Paths," .

Jehiel, P. (2005), "Analogy-based expectation equilibrium," *Journal of Economic theory*, 123, 81–104.

Jehiel, P. and Koessler, F. (2008), "Revisiting games of incomplete information with analogy-based expectations," *Games and Economic Behavior*, 62, 533–557.

Jovanovic, B. (1981), "Entry with private information," *The Bell Journal of Economics*, pp. 649–660.

Kamenica, E. (2008), "Contextual inference in markets: On the informational content of product lines," *The American Economic Review*, 98, 2127–2149.

Kartik, N., Squintani, F., and Tinn, K. (2015), "Information revelation and pandering in elections," Tech. rep., mimeo, Columbia University.

Klemperer, P. (1999), "Auction theory: A guide to the literature," *Journal of Economic Surveys*, 13, 227–286.

Kolb, A. (2015), "Optimal entry timing," *Journal of Economic Theory*, 157, 973–1000.

Krishna, V. and Morgan, J. (2001a), "Asymmetric information and legislative rules: Some ammendments," *American Political Science Review*, 95, 435–452.

Krishna, V. and Morgan, J. (2001b), "A Model of Expertise," *Quarterly Journal of Economics*, 116, 747–775.

Krishna, V. and Morgan, J. (2008), "Contracting for information under imperfect commitment," *The RAND Journal of Economics*, 39, 905–925.

Le Gall, J. (1984), "One-dimensional stochastic differential equations involving the local times of the unknown process," in *Stochastic analysis and applications*, pp. 51–82, Springer.

Li, H. and Suen, W. (2004), "Delegating decisions to experts," *Journal of Political Economy*, 112, 311–335.

Loury, G. C. (1979), "Market structure and innovation," *The Quarterly Journal of Economics*, pp. 395–410.

Melumad, N. and Shibano, T. (1991), "Communication in settings with no transfers," *The RAND Journal of Economics*, pp. 173–198.

Milgrom, P. and Roberts, J. (1986), "Relying on the information of interested parties," *The RAND Journal of Economics*, pp. 18–32.

Milgrom, P. and Weber, R. (1985), "Distributional strategies for games with incomplete information," *Mathematics of Operations Research*, 10, 619–632.

Mirrlees, J. (1974), "Notes on welfare economics, information and uncertainty," *Essays on economic behavior under uncertainty*, pp. 243–261.

Mitchell, W. (1989), "Whether and when? Probability and timing of incumbents' entry into emerging industrial subfields," *Administrative Science Quarterly*, pp. 208–230.

Morris, S. and Shin, H. (2002), "The social value of private information," *The American Economic Review*, 92, 1521–1534.

Morris, S. and Shin, H. (2003), "Global games: Theory and applications," Econometrics Society Monographs.

Moscarini, G. and Squintani, F. (2010), "Competitive experimentation with private information: The survivor's curse," *Journal of Economic Theory*, 145, 639–660.

Murto, P. and Välimäki, J. (2011), "Learning and information aggregation in an exit game," *The Review of Economic Studies*, 78, 1426–1461.

Myatt, D. and Wallace, C. (2014), "Central bank communication design in a Lucas-Phelps economy," *Journal of Monetary Economics*, 63, 64–79.

Ortner, J. (2013), "Durable Goods Monopoly with Stochastic Costs," Working paper.

Ottaviani, M. and Sørensen, P. (2006), "Professional advice," *Journal of Economic Theory*, 126, 120–142.

Pesendorfer, W. and Wolinsky, A. (2003), "Second opinions and price competition: Inefficiency in the market for expert advice," *The Review of Economic Studies*, 70, 417–437.

Pétry, F. and Collette, B. (2009), "Measuring how political parties keep their promises: A positive perspective from political science," vol. 15 of *Studies in Public Choice*, pp. 65–80.

Reed, M. and Simon, B. (1980), *Methods of modern mathematical physics: Functional analysis*, vol. 1, Gulf Professional Publishing.

Royden, H. and Fitzpatrick, P. (1988), *Real analysis*, vol. 4, Prentice Hall New York.

Sampford, M. (1953), "Some inequalities on Mill's ratio and related functions," *The Annals of Mathematical Statistics*, pp. 130–132.

Schoenecker, T. S. and Cooper, A. C. (1998), "The role of firm resources and organizational attributes in determining entry timing: a cross-industry study," *Strategic Management Journal*, 19, 1127–1143.

Shin, H. (1998), "Adversarial and inquisitorial procedures in arbitration," *The RAND Journal of Economics*, pp. 378–405.

Sircar, R. and Xiong, W. (2007), "A general framework for evaluating executive stock options," *Journal of Economic Dynamics and Control*, 31, 2317–2349.

Taylor, C. (1999), "Time-on-the-Market as a Sign of Quality," *Review of Economic Studies*, 66, 555–578.

Tirole, J. (2014), "Cognitive Games and Cognitive Traps," Tech. rep., mimeo Toulouse School of Economics.

Touzi, N. and Vieille, N. (2002), "Continuous-time Dynkin games with mixed strategies," *SIAM Journal on Control and Optimization*, 41, 1073–1088.

Treitel, G. (1999), *The law of contract*, Sweet & Maxwell, 10 edn.

Wittman, D. (1983), "Candidate motivation: A synthesis of alternative theories," *American Political Science Review*, 77, 142–157.

Wolinsky, A. (2002), "Eliciting information from multiple experts," *Games and Economic Behavior*, 41, 141–160.

Zhang, M. (2000), "Calculation of diffusive shock acceleration of charged particles by skew Brownian motion," *The Astrophysical Journal*, 541, 428–435.

Zitzewitz, E. (2001), "Measuring herding and exaggeration by equity analysts and other opinion sellers," Tech. rep., mimeo, Stanford University.

Biography

Aaron Michael Kolb was born on December 17, 1984 in Columbus, Mississippi to Karen Kay Kolb and Michael Anthony Kolb. A graduate of St. Xavier High School in Cincinnati, Ohio, he earned a BA from Yale University in 2007 with a double major in Mathematics and Economics. He attended Duke University for graduate school, where he received his MA and PhD in Economics in 2011 and 2016, respectively.