Title
Essays in Dynamic Games and Information Economics

Author
Wei, Dong

Publication Date
2020

Peer reviewed|Thesis/dissertation

Essays in Dynamic Games and Information Economics

by

Dong Wei

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Chris Shannon, Co-chair
Professor Philipp Strack, Co-chair
Professor Brett Green

Fall 2020

Essays in Dynamic Games and Information Economics

Abstract

Essays in Dynamic Games and Information Economics

by

Dong Wei

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Chris Shannon, Co-chair

Professor Philipp Strack, Co-chair

This dissertation consists of three chapters, studying questions in dynamic games and information economics. These chapters represent a selection of my research conducted during the period of my PhD studies.

Chapter 1 is motivated by the "starting small" phenomena which are prevalent in long-term relationships. For example, in credit relationships, it is often observed that the credit limit granted to a borrower by a lender tends to increase over time conditional on satisfactory repayment histories. Why does this happen? To provide a novel explanation, we study a repeated lender-borrower game with anonymous re-matching; that is, once an ongoing relationship is terminated, players are rematched with new partners and prior histories are unobservable. We propose an equilibrium refinement based on two assumptions: (a) default implies termination of the current relationship; (b) in a given relationship, a better loan-repayment history implies weakly higher continuation values for both parties. This refinement captures the idea of "justifiable punishments" in repeated games. We show that if players are patient enough and re-matching is sufficiently likely, then the loan size is strictly increasing over time along the equilibrium path of *all* non-trivial equilibria. As such, this chapter helps explain gradualism in long-term relationships, especially credit relationships.

Chapter 2, based on the joint work with Elliot Lipnowski and Laurent Mathevet, concerns the optimal design of information transmission protocols when the information recipient is rationally inattentive. We develop a model of a well-intentioned principal who provides information to a rationally inattentive agent. Processing information is costly to the agent, but the principal does not internalize this cost. In a world with two states, it is shown that providing full information is universally optimal for the principal, even though the agent will typically not pay full attention. We then introduce a tractable specification with quadratic payoffs and study optimal information provision when full disclosure is not optimal. We characterize incentive compatible information policies, that is, those to which the agent willingly pays full attention. In a leading example with

three states, optimal disclosure involves information distortion at intermediate costs of attention. As the cost increases, optimal information abruptly changes from downplaying the state to exaggerating the state.

Chapter 3 naturally extends the theoretical framework introduced in Chapter 2 to a different but relevant setting where the decision interests between the sender and the receiver of information are misaligned. In our model, a Sender (seller) tries to persuade a rationally inattentive Receiver (buyer) to take a particular action (e.g., buying). Learning is costly for the Receiver who can choose to process strictly less information than what the sender provides. In a binary-action binary-state model, we show that optimal disclosure involves information distortion, but to a lesser extent than the case without learning costs; meanwhile, the Receiver processes less information than what he would under full disclosure. While the Sender is always worse off when facing a less attentive Receiver, the amount of information processed in equilibrium varies with learning costs in a non-monotone fashion. As such, this chapter sheds light on how to persuade a rationally inattentive decision maker.

To my beloved parents, Chunling Wei and Hong Ju.

Contents

List of Figures

Acknowledgments

I would not have been able to accomplish this dissertation without the tremendous help and support from my advisors, colleagues, and families.

The journey of my doctoral studies was a bit unusual. I started at NYU Stern in 2014 and transferred to UC Berkeley in 2016. Because of that, I have had the opportunity to learn from and collaborate with the fantastic people in both of these two great institutions.

First and foremost, I would like to thank my main advisors, Philipp Strack and Chris Shannon, for their invaluable support and advice throughout my time at UC Berkeley. Philipp is an extremely talented researcher and a wonderful advisor. Being one of the most productive economic theorists in the world, Philipp has never been short on time with advising and has always been able to provide timely and insightful feedback on my research projects. Likewise, Chris was equally generous with her time and has shared with me many of her experiences and insights. In particular, Chris provided me careful guidance that led to my first ever publication in a leading academic journal, and that same article also became the first chapter of my dissertation.

I would also like to thank Brett Green who is a member of my dissertation committee and a coauthor of mine. Brett is an extraordinary researcher with great insights. We have collaborated on another project, "(Reverse) Price Discrimination with Information Design," which for administrative reasons is not included in this dissertation but is an essential part of my doctoral research. I will always remember the many afternoons I spent at Brett's office where I discussed research ideas with him and listened to his generous feedback on various projects of mine.

I have always been extremely grateful for the time spent and knowledge learned at NYU. In particular, I want to thank Luís Cabral and Debraj Ray who would have been my advisors had I stayed at NYU to finish my degree. The first chapter of this dissertation originated from discussions with Luís in the summer of 2015. Later on, Debraj, who is one of the best experts in microeconomic theory, offered me extremely helpful feedback on this project while I was at NYU and after I left. Moreover, in the 2015-2016 academic year, I benefited tremendously from participating in the "New Research in Economic Theory" (NRET) workshop, a student-based economic theory seminar organized by Debraj, where I was able to sharpen my presentation skills and get exposed to cutting-edge theoretical research at an early stage.

During my PhD studies, I have had the honor to collaborate with a number of brilliant researchers in economic theory, among whom I would like to especially thank Laurent Mathevet and Elliot Lipnowski. Laurent has been a faculty member at NYU's Economics Department, and Elliot was a graduate student at NYU Stern while I was there. I took Laurent's class on Game Theory in the 2015 Fall semester. Laurent was an amazing instructor, and his course led me to the field of information design as well as our joint project with Elliot on "Attention Management." This project finally became the basis of the second and third chapters of my dissertation. Elliot is not only a colleague of mine, but also a great friend who has been extremely helpful to me with providing research feedback and job market advice.

Last but far from the least, I want to express my deepest gratitude to my parents, Chunling Wei and Hong Ju, to whom this dissertation is dedicated, and my wife, Rong Huang, for their unconditional support, constant encouragement and fullest trust over the past six years. I would not have carried through this journey without any of them.

Introduction

Information frictions are at the heart of many economic problems. For example, when there is moral hazard, repeated interactions are supposed to provide incentives for people to behave cooperatively, but a lack of information about a person's past records may undermine such incentives. In addition, learning or information processing is often costly, and as a result, individuals tend not to process every piece of information that is available to them. In this dissertation, we study several economic issues related to information frictions, and propose solutions to those situations where such frictions are present.

Chapter 1 is devoted to examining the cases where information frictions arise *exogenously* due to institutional constraints. In many developing countries, institutions such as credit bureaus are nonexistent. Consequently, it is very hard for creditors to acquire information about the past payment behavior of an individual borrower. How can trust be built in such scenarios? Motivated by the prevalent "starting small" phenomenon, we study a repeated lender-borrower game with anonymous re-matching; that is, once an ongoing relationship is terminated, players are rematched with new partners and prior histories are unobservable. The anonymity in re-matching captures the exogenous information frictions brought about by the lack of credit bureaus and the like. To deliver a sharp prediction, we propose an equilibrium refinement based on two assumptions: (a) default implies termination of the current relationship; (b) in a given relationship, a better loan-repayment history implies weakly higher continuation values for both parties. We show that if players are sufficiently patient and re-matching is highly likely, then the loan size is strictly increasing over time along the equilibrium path of all non-trivial equilibria. As such, this chapter illustrates how gradualism can facilitate trust over time in a society without access to credit histories and thus help people partially overcome such extreme information frictions. It also provides an novel explanation to the "starting small" phenomenon in long-term relationships, especially in credit relationships.

Chapters 2 and 3 turn to the scenarios where information frictions arise *endogenously* due to people's limited information processing capacity. Information is a gift that may not always be accepted and, hence, useful. Speaking to a toddler about grammar may not improve his linguistic abilities, just as an adult may learn less from a book, an email, or a contract that contains too much detail. Simon[41] foreshadows the potential hazards of detailed communication: "What information consumes is rather obvious: it consumes the attention of its recipients." Failures to recognize this fact can have counterproductive effects: consumers are frequently confused by nutritional labels; patients can be overwhelmed in parsing side effects of medications; and so on (see Ben-Shahar and Schneider[11]). How do such information frictions that arise from inattention affect people's communication? What are the effective ways to persuade an individual who has difficulties in processing too much or too complicated information? To shed light on these questions, we develop a model in Chapter 2 where a well-intentioned principal provides information to an agent for whom information is costly to process, but the principal does not internalize this cost. We show that full information is universally optimal if and only if the environment comprises one issue. With multiple issues, attention management becomes optimal: the principal restricts some information to induce the agent to pay attention to other aspects. In Chapter 3, we extend this theoretical framework to a different setting where the interests of the sender and the receiver of information are misaligned. In that model, a Sender (seller) tries to persuade a rationally inattentive Receiver (buyer) to take a particular action (e.g., buying). Learning is costly for the Receiver who can choose to process strictly less information than what the sender provides. In a binary-action binary-state model, we characterize optimal disclosure and provide a number of interesting results about the comparative statics with respect to learning costs.

Chapter 1

A Model of Trust Building with Anonymous Re-Matching

This chapter was published in the Journal of Economic Behavior and Organization.

1.1 Introduction

Gradualism, or "starting small," is often observed in long-term relationships. It is reflected in the increasing level of interactions over time between parties in a relationship. For example, the credit line granted by a credit card company to a specific customer normally increases over time as more and more on-time repayments are made; Antràs and Foley[5] document the pattern of financing terms of a U.S.-based exporter in the poultry industry, and find that the amount of trade credit granted by the exporter to its trading partners increases with the length of their relationships. More broadly, in large societies or informal credit markets of developing economies where personal history records are hard to obtain, people are usually cautious at the beginning of a new relationship, and put more at stake after satisfactory interactions.

This paper focuses on one specific driving force of gradualism: the re-matching opportunity. We study a repeated lender-borrower game with anonymous re-matching, in which a lender and a borrower interact in a society with a group of lenders and a group of borrowers. In addition to lending and repayment decisions, the lender and borrower can choose whether or not to continue their relationship at the end of each period. If a relationship is terminated, each party becomes unmatched; in the next period, an unmatched player will be anonymously matched with a new partner with some exogenous probability. The re-matching is *anonymous* in the sense that the players' actions in previous relationships are unobservable in the current relationship. This captures the idea that it is costly to acquire past history information of the other party in a new relationship. To highlight the importance of re-matching, the model abstracts away from incomplete information about players' type, and commitment power.

For this game, we focus on equilibria in which the same strategies are used in every relationship; that is, every new relationship is just a restart of the first relationship. We call such a strategy profile a *social equilibrium*. Such equilibria are studied by Datta[18] with a focus on the efficient ones, and by Ghosh and Ray[22, 23] in a setting with a simultaneous stage game and incomplete information about players' types. We propose a refinement called *orthodox social equilibria*, which are social equilibria such that: (i) an ongoing relationship is terminated by the lender on default; (ii) a better loan-repayment history is followed by weakly higher continuation values for the lender and the borrower. Our main result is that, if the discount factor and the re-matching probability are above certain thresholds, the size of loans along the equilibrium path of *any* (non-trivial) orthodox social equilibrium is strictly increasing over time.

In the definition of an orthodox social equilibrium, the first restriction that the relationship is terminated on default is standard in the literature (for example, see Datta[18], Kranton[28], Ghosh and Ray[22, 23]). The second restriction is motivated by the following idea. In credit relationships, repayment is viewed as a better action than default because default directly hurts the lender. In addition, a larger loan is better than a smaller one, in the sense that conditional on repayment, a larger loan leads to higher per-period payoffs for both parties. At the beginning of a given date t, if there are two histories of actions such that neither entails default and they differ only in the loan size in the last period, we may then view the history with a higher last-period loan as a result of the lender's deviation to offering an (unexpected) higher loan followed by the borrower's repayment. The lender's (and potentially the borrower's) deviation benefits both parties in this scenario, so the history with a higher last-period loan is arguably better. The second restriction in defining orthodox equilibrium says that, at any given t, the continuation values of the relationship following a better loan-repayment history should be weakly higher. Otherwise, if they are strictly lower, players can be viewed as punishing each other by using strictly lower continuation values, even though both of them have benefited from the deviation. But punishing the deviating player for taking better actions is not "morally justified" because she did not hurt anyone; moreover, it is not even "economically justified" if the punishment is costly to the non-deviating player (who is the "carrier" of the punishment).[1] The idea behind this refinement is related to the notion of "justifiable punishments" in normal-form repeated games (Aramendia and Wen[6]) which also puts a monotone restriction on continuation values following deviations that benefit the other party.

One important implication of our refinement is that, in any orthodox social equilibrium, the lender will make the borrower's no-deviation (i.e., no-default) constraint bind at all dates. This is because if it is not binding at some date, i.e., if the borrower strictly prefers repaying the loan, then

[1] To be more concrete, consider the following situation. Suppose that on an equilibrium path the lender should offer a loan of \$1,000 in the current period, but she deviates to lending \$2,000 and the borrower later repays. This leads them to an off-equilibrium history. At the end of that period, should the deviating player(s) be punished (for example, via terminating the relationship) just for the sake of the fact that some deviation has happened? More generally, should the value of a person's relationship with a bank be strictly reduced after the bank (unexpectedly) lends more and the person repays? Although subgame perfection does not offer an answer, the standard analysis of repeated games, which imposes the worst punishment to *any* deviation (Abreu[1]), answers "yes." However, in the context such as credit relationships where certain actions are arguably better, when deviations to better actions/histories occur, reducing the continuation value of the deviating or non-deviating player is hardly justified, morally or economically, as argued in the main text.

the lender can increase the loan size at that date by a little without inducing default. The reason why the borrower will not default on the slightly higher loan is as follows. If the borrower repays, she will enjoy a weakly higher continuation value than before because repaying a higher loan results in a better history than the one on path;[2] since her no-deviation constraint is not binding on path, the same constraint is still satisfied at a slightly higher loan followed by a weakly higher continuation value. But if a slightly higher loan does not induce default, by offering it the lender will also enjoy a weakly higher continuation value (together with a strictly higher current payoff) because a better loan-repayment history will be achieved. So in any orthodox social equilibrium, the lender has an incentive to increase the loan size as long as the borrower's no-deviation constraint is not binding, which implies that those constraints must be binding at all dates on path.

Given the above implication, we can explain the intuition for our main result, which shows that the loan sequence in any orthodox social equilibrium is strictly increasing. On the one hand, when past histories are unobservable, the high re-matching probability undermines the punishment power of the threat of terminating a relationship; so in order to induce repayment, there has to be some additional cost of starting a new relationship. This additional cost is reflected in the fact that the value of a relationship is increasing in its length, so that restarting a new relationship is worse than staying in the current relationship. On the other hand, the lender knows that the longer the borrower has stayed in a given relationship, the higher the value of this relationship to the borrower will be. Since the value of becoming unmatched is constant, the cost inflicted upon the borrower by terminating the current relationship becomes larger as time goes on; that is, the cost of defaulting increases over time. Then, since the lender has an incentive to exploit the borrower's no-deviation constraints, the loan size she offers also increases over time.

The remainder of this chapter is organized as follows. Section 1.2 sets up the model and states the main results. Section 1.3 discusses the intuition for our results, multiplicity of equilibria, and extension to mixed strategies. Section 1.4 concludes. All proofs for the results in this chapter are in Appendix A.

Related Literature

This paper contributes to the literature on trust building by showing that, in a repeated lender-borrower model, high likelihood of anonymous re-matching and sufficiently patient players—without assuming efficiency, multiple types and/or contractual commitment—are enough to deliver a unique prediction of strictly increasing loans over time in non-trivial equilibria.

The paper closest to ours is Datta[18]. He considers a special case of our model with linear payoffs and immediate re-matching, and shows that the *value of the relationship*—defined as the discounted sum of current and future loans—is nondecreasing in *efficient* social equilibria. This paper offers a sharper and more testable prediction in a more general setup where we allow for probabilistic re-matching and richer payoff structures. Compared to the result in Datta[18], our prediction of strictly increasing loans is sharper because nondecreasing values still permit non-monotone loans over time; it is also much easier to test because the value of a relationship is much

[2] Note that these two histories only differ in the last period and the new history has a larger loan in that period.

harder to observe than the level of interaction in each period.[3] Moreover, we propose an equilibrium refinement to address what should happen after an unexpected higher loan is repaid. It rules out unreasonable equilibria supported by "unjustifiable punishments." Unlike Datta[18], since this refinement does not impose efficiency, our result is robust to inefficiency, whereas Datta's weaker characterization hinges on the efficiency assumption. Kranton[28] studies a repeated game with re-matching where the stage game is of *simultaneous* moves with incentives similar to a prisoners' dilemma. She characterizes the cooperation levels in efficient equilibria to be "starting small" in the first period and reaching the efficient level from the second period on. Social equilibria of this type are ruled out by our refinement, and are arguably unreasonable in credit relationships where the stage game is of *sequential* moves (see Section 1.3 for a detailed discussion).

Two other explanations for gradualism can be found in the literature. One strand of explanations combines moral hazard with incomplete information by introducing multiple types (usually reflected in patience levels) to one side or both sides of a relationship. The reason for starting small in such environments is that, when the history of cooperation is longer, Bayes' rule implies that the probability of the other party being the "good" type is higher, so the optimal level of interaction increases over time. Ghosh and Ray[22] study a repeated game with both incomplete information (two types) and re-matching, where the stage game is of simultaneous moves with incentives similar to a prisoners' dilemma. Their characterization of the evolution of cooperation levels is similar to Kranton[28]: "starting small" only in the very first period. Watson[48, 49] studies a general model of long-term relationships with incomplete information, and characterizes the level of interactions to be increasing over time under certain refinements. Kartal et al.[27] studies a repeated trust game with two types of "receivers" (without re-matching), and finds a similar increasing pattern of trust levels. Applications of this idea, among others, include Rauch and Watson[39] in the context of trading between a supplier and a less informed buyer, as well as Araujo et al.[7] and Antràs and Foley[5] in the context of trade credit. The driving force of gradualism in these settings is the gradual learning of the other party's type.

Another strand of explanations shows up in the literature on self-enforcing contracts. Ray[40] studies a general repeated moral hazard problem without re-matching or incomplete information, and proves a similar result of "starting small" for *efficient* self-enforcing contracts. That is, in all efficient self-enforcing contracts, the continuation payoffs move over time in the direction of the agent who has an incentive to renege in the stage game. The main idea there is to apply the backloading argument up to a point: postponing higher rewards to the agent can keep the agent's current value the same while relaxing the agent's current incentive constraint (because future becomes more valuable), so that the agent can be incentivized to work harder or repay more today, which increases the principal's value and improves efficiency. Applications of this insight include, among others, Lazear[29] and Thomas and Worrall[47] in the context of wage contracts, Thomas and Worrall[46] in the context of foreign direct investment with threat of expropriation, and Albuquerque and Hopenhayn[2] in the context of credit relationships. The driving force of gradualism in these settings is the interaction between efficiency and self-enforcing constraints.

[3] For example, our result can be falsified if the *observed* history of loans is not monotone. However, this simple test does not work for Datta's result because non-monotone loans could still be consistent with monotone values.

1.2 The Model

Model Setup

Consider a lender and a borrower in a society with a group of lenders and a group of borrowers. Both are infinitely lived. Time is discrete and starts from 0. The lender's and the borrower's utility functions are the discounted sums of their expected period payoffs. Specifically, let δ be the common discount factor and let $\mathbf{y}^L = \{y_0^L, y_1^L, ...\}$ be the sequence of expected period payoffs to the lender. The utility function of the lender at time t is:

$$V_t^L(\mathbf{y}^L) = \Sigma_{i=0}^{\infty}\delta^i y_{t+i}^L. \tag{1.1}$$

Similarly, let $\mathbf{y}^B = \{y_0^B, y_1^B, ...\}$ be the sequence of expected period payoffs to the borrower, and define

$$V_t^B(\mathbf{y}^B) = \Sigma_{i=0}^{\infty}\delta^i y_{t+i}^B. \tag{1.2}$$

The stage game takes the following form. Each period t is divided into three subperiods. At t^0, the lender chooses the size of the loan, $L_t \in [0, L^*]$, granted to the borrower. At t^1, the borrower chooses whether to repay or default. At t^2, the lender and the borrower simultaneously choose whether or not to continue their relationship. The relationship continues if and only if both of them choose to do so. We assume away exogenous separation only for expository purposes.[4]

If the relationship continues, in the next period the players repeat the stage game as described; if it is terminated, each party enters the next period as an unmatched lender/borrower. In each period, an unmatched player will be anonymously matched with a new partner with an exogenous probability $\lambda \in [0, 1]$. If matched, she starts the stage game with the new partner in this period;[5] otherwise, she earns a payoff of 0 in this period and enters the next period as an unmatched player.

The history of players' actions in the *current* relationship is common knowledge to the borrower and the lender forming this relationship; histories of all past relationships of any party are unobservable,[6] so newly matched partners effectively restart from the very beginning of the game.

The payoffs in each period are determined by the loan size offered by the lender and the repay-

[4]As is standard in repeated games, all results still hold if Nature terminates the relationship with probability β in each period, with the only change being that the discount factor used in defining the thresholds in our main propositions should be $\delta(1-\beta)$, instead of δ. In that case, we refer to $\delta(1-\beta)$ as the "effective discount factor." The detailed analysis of the case with exogenous separations can be found in the Supplementary Material of the published version in the Journal of Economic Behavior and Organization.

[5]Once matched, a previously unmatched lender (borrower) still plays the role of a lender (borrower).

[6]This is equivalent to assuming that the histories of all past relationships of the *other* party are unobservable, and that the players also do not base their decisions on their own history of actions in past relationships which in principle are observable to themselves.

ment decision of the borrower. Specifically, for a relationship in its period t,

$$y_t^L = \begin{cases} R(L_t), & \text{if repayment happens;} \\ -L_t, & \text{if default happens;} \end{cases} \tag{1.3}$$

$$y_t^B = \begin{cases} C(L_t), & \text{if repayment happens;} \\ D(L_t), & \text{if default happens.} \end{cases} \tag{1.4}$$

To summarize, a lender-borrower game with anonymous re-matching depends on the following elements: the discount factor δ, the re-matching probability λ, the loan upper bound L^*, function R for the lender, and functions C and D for the borrower. For the sake of exposition, when no confusion arises we denote the game by $G(\delta, \lambda)$, omitting the dependence with respect to the R, C, D functions and the loan upper bound L^*.[7]

Assumption 1.

- $R(\cdot), C(\cdot), D(\cdot)$ *are continuous and strictly increasing, and* $R(0) = C(0) = D(0) = 0$*;*
- $\Delta(L) \equiv D(L) - C(L) > 0$ *for all* $L \in (0, L^*]$*, and is strictly increasing;*
- *There exist* $\underline{\alpha}, \bar{\alpha}$*, such that* $0 < \underline{\alpha} \leq \bar{\alpha} < 1$*, and* $\underline{\alpha} D(L) < C(L) < \bar{\alpha} D(L)$ *for all* $L \in (0, L^*]$.

Assumption 1 requires that the borrower's payoffs from both default and repayment increase with loan size; in addition, the gains from default, $\Delta(L)$, are positive and also increase with loan size, capturing the borrower's myopic incentive to default. For the lender, if repayment occurs, a larger loan generates a higher period payoff; meanwhile, if default occurs, the lender bears a cost that increases with loan size. The last part of Assumption 1 ensures that $\frac{C(L)}{D(L)}$ is always bounded away from 0 and 1.[8] Note that L_t in general can be viewed as the level of trust offered by the lender. The higher the level of trust, the higher the payoffs for both parties conditional on cooperation (repayment); however, a higher L_t also leads to a higher temptation to defect (default), as captured by the assumption that $\Delta(L) \equiv D(L) - C(L)$ is strictly increasing.

Example. Borrowing Capital for Production

Suppose that the borrower needs capital for production. Their payoffs can be modeled as follows:

$$y_t^L = \begin{cases} rL_t, & \text{if repayment happens;} \\ -L_t, & \text{if default happens;} \end{cases}$$

$$y_t^B = \begin{cases} F(L_t) - (1+r)L_t, & \text{if repayment happens;} \\ F(L_t), & \text{if default happens.} \end{cases}$$

[7] The results of this paper, Propositions 1 and 2, still hold even when there is no upper bound on loans.

[8] If C and D are differentiable at 0, the last part of Assumption 1 is equivalent to $0 < C'(0) < D'(0) < \infty$.

where F is a production function of capital, and r is the fixed interest rate. Assumption 1 is satisfied if $F(L) - (1+r)L$ is strictly increasing in L and $1 + r < F'(0) < \infty$.

Notice that in the unique subgame perfect Nash equilibrium of the stage game (without termination decision), the lender chooses a loan size of 0 and the borrower always defaults on any positive loan. In the repeated-game setting, we would like to focus on the equilibria in which the strategies depend only on the history of the current relationship, and the same strategy profile is played by all players in all relationships. That is, what players do in a new relationship is an exact repetition of any prior relationship. Given our focus, it is without loss to look at the game of a lender and a borrower in a single relationship while taking as given the continuation game (payoffs) on terminating the current relationship, which itself is determined in equilibrium.

Orthodox Social Equilibrium

We now define an equilibrium concept for our game. Let $L_t \in [0, L^*]$ be the loan size chosen at t^0; let $d_t \in \{0, 1\}$ denote the borrower's defaulting decision at t^1, such that $d_t = 1$ if and only if default happens at t; let $f_t \in \{0, 1\}$ and $g_t \in \{0, 1\}$ denote the lender's and the borrower's decisions on continuing the relationship at t^2, such that their relationship is terminated at t if and only if $f_t g_t = 0$. Denote by a_t the outcome of these decisions at period t; that is, $a_t = (L_t, d_t, f_t, g_t)$. The history at each node is denoted by:

$$\begin{aligned} h(t^0) &= \{a_0, a_1, ...a_{t-1}\}, \text{ where } h(0^0) = \emptyset; \\ h(t^1) &= h(t^0) \cup \{L_t\}; \\ h(t^2) &= h(t^1) \cup \{d_t\}. \end{aligned}$$

Let $H(t^i)$ be the collection of all possible histories at t^i, for $i = 0, 1, 2$.

A (pure) strategy of the lender $l = \{l_0, l_1, ...\}$ consists of a sequence of decision rules that maps each information set to her decision at that node. Specifically, $l_t = (\tilde{L}_{t^0}, \tilde{f}_{t^2})$, where $\tilde{L}_{t^0} : H(t^0) \to [0, L^*]$ and $\tilde{f}_{t^2} : H(t^2) \to \{0, 1\}$. Similarly, a (pure) strategy of the borrower $b = \{b_0, b_1, ...\}$ is defined as $b_t = (\tilde{d}_{t^1}, \tilde{g}_{t^2})$, where $\tilde{d}_{t^1} : H(t^1) \to \{0, 1\}$ and $\tilde{g}_{t^2} : H(t^2) \to \{0, 1\}$. Notice that given a strategy profile $\{l, b\}$, we are not yet able to compute the payoff of each player, if according to $\{l, b\}$ a relationship is terminated at some date. This is because the continuation values after termination of a relationship depend on equilibrium payoffs, but we have not solved for them; therefore, the equilibrium concept involves a fixed point between re-matching values and equilibrium payoffs.

Let $\bar{V}^L$ and $\bar{V}^B$ be the re-matching values (i.e., values of a newly-matched relationship) for the lender and the borrower, respectively. Note that the continuation values of an unmatched lender and borrower are given by $\lambda' \bar{V}^L$ and $\lambda' \bar{V}^B$, where $\lambda' = \frac{\lambda}{1-(1-\lambda)\delta}$.[9] Given a strategy profile (l, b), we are able to trace out a sequence of decisions on the equilibrium path, $\{a_t\}_t$, where $a_t = (L_t, d_t, f_t, g_t)$.

[9]To see this, let $\hat{V}^B$ be the continuation payoff of an unmatched borrower. We have $\hat{V}^B = \lambda \bar{V}^B + (1-\lambda)\delta \hat{V}^B$, so that $\hat{V}^B = \frac{\lambda}{1-(1-\lambda)\delta} \bar{V}^B$.

Let $T(l,b)$ be the date at which the relationship is terminated according to (l,b), i.e., $f_{T(l,b)}g_{T(l,b)} = 0$, and $f_t g_t = 1$ for all $t < T(l,b)$. From (1.1) and (1.2), we can write:

$$V_t^L(l,b,\bar{V}^L,\bar{V}^B) = \sum_{i=0}^{T(l,b)-t} \delta^i[(1-d_{t+i})R(L_{t+i}) - d_{t+i}L_{t+i}] + \delta^{T(l,b)-t+1}\lambda'\bar{V}^L, \tag{1.5}$$

$$V_t^B(l,b,\bar{V}^L,\bar{V}^B) = \sum_{i=0}^{T(l,b)-t} \delta^i[(1-d_{t+i})C(L_{t+i}) + d_{t+i}D(L_{t+i})] + \delta^{T(l,b)-t+1}\lambda'\bar{V}^B. \tag{1.6}$$

In (1.5), a lender's payoff is the discounted sum of her current and future period payoffs. At $t+i$, her period payoff is $R(L_{t+i})$ if repayment happens and $-L_{t+i}$ otherwise; in addition, at $T(l,b)$ when the relationship is terminated, she will also get the discounted continuation value of an unmatched borrower. In (1.6), a borrower's payoff has the same structure as that of a lender, where at $t+i$ the borrower gets $C(L_{t+i})$ if repayment happens and $D(L_{t+i})$ otherwise, plus the present value of her continuation payoff when the relationship is terminated.

Definition 1. *A **social equilibrium** consists of a strategy profile (l,b) and re-matching values $(\bar{V}^L,\bar{V}^B)$, such that*

(i) Gvien equilibrium payoffs $\bar{V}^L$ and $\bar{V}^B$, l and b are sequentially rational to each other;

(ii) $\bar{V}^L = V_0^L(l,b,\bar{V}^L,\bar{V}^B)$, $\bar{V}^B = V_0^B(l,b,\bar{V}^L,\bar{V}^B)$.

Part (i) of Definition 1 is just the standard requirement of subgame perfection, while part (ii) captures our fixed point requirement for re-matching values. We call it social equilibrium because part (ii) implicitly assumes that every pair in the society plays such a strategy profile in every relationship. We focus on pure-strategy equilibria in most of this paper, and discuss the extension to mixed strategies in Section 1.3.

We further restrict our attention to orthodox social equilibria, which are social equilibria such that: (i) an ongoing relationship is terminated on default; (ii) a better loan-repayment history is followed by weakly higher continuation values for the lender and the borrower. Formally, given a strategy profile (l,b), let $V_t^L : H(t^0) \to \mathbb{R}$ and $V_t^B : H(t^0) \to \mathbb{R}$ be the induced continuation value functions for the lender and the borrower at the beginning of each period. They specify the remaining values of the *current* relationship to the lender and the borrower following any history.

Definition 2. *A social equilibrium strategy profile (l,b) is **orthodox**, if*

(i) For any t and any $h(t^2) \in H(t^2)$ such that $d_\tau = 1$ for some $\tau \le t$, we have $\tilde{f}_{t^2}[h(t^2)] = 0$;

(ii) For any t and any $h(t^0), h'(t^0) \in H(t_0)$ such that for all $\tau < t$, $L'_\tau \ge L_\tau$ (with equality for all $\tau < t-1$), $d'_\tau = d_\tau = 0$ and $f'_\tau g'_\tau = f_\tau g_\tau = 1$, we have $V_t^L[h'(t^0)] \ge V_t^L[h(t^0)]$ and $V_t^B[h'(t^0)] \ge V_t^B[h(t^0)]$.[10]

[10] A weaker requirement that a better loan-repayment history leads to weakly higher continuation values to both players in the *rest of the game* (including both the remaining values from the current relationship, and the values from re-matching if the current relationship is terminated later) would deliver exactly the same results.

Part (i) of Definition 2 requires that the lender terminates the relationship if the borrower has defaulted before. This requirement is standard in the literature (see also Datta[18], Kranton[28], Ghosh and Ray[22, 23], etc.); it simplifies the analysis by imposing a fixed, in fact worst,[11] punishment on the borrower's default.

Part (ii) of Definition 2 needs more justifications. It says that at the beginning of any period t, if we consider two histories such that the borrower has made repayments at all dates in both of them, and the loan sizes only differ in the last period (i.e., $t-1$), then the continuation values for both players following the history with a higher last-period loan should be weakly higher.[12] This restriction is motivated by the following idea. One can view an equilibrium in a repeated game as an implicit contract between the players. Such a contract specifies each player's continuation value at every history. Let us fix a period t and look at two different histories at the beginning of period t—$h(t^0)$ and $h'(t^0)$—among which $h(t^0)$ is on the equilibrium path. To sustain $h(t^0)$ as part of the equilibrium path, any deviation that benefits the deviating player in the current period has to be prevented by sufficiently reducing her continuation value, which is achieved by changing the future course of play. But such a punishment imposed on the deviating player is not "morally justified" if the deviation unambiguously improved the other player's payoff. In the end, on what ground should a person be "punished" if her actual behavior—compared to what she was expected to do—strictly benefits another party? Moreover, it is not even "economically justified" if the punishment is costly to the non-deviating player (who is the "carrier" of the punishment). In the context of credit relationships, if the only difference between $h(t^0)$ and $h'(t^0)$ is that $h'(t^0)$ has a higher last-period loan (with all repayments made), it is fair to claim that both the lender and the borrower have unambiguously benefited from the deviation to $h'(t^0)$.[13] The previous argument then implies that neither player should receive a strictly lower continuation value after history $h'(t^0)$ than after $h(t^0)$, for otherwise it would be either morally unjustified or economically unjustified (or both).

The idea behind Part (ii) is related to the notion of "justifiable punishments" (Aramendia and Wen[6]), which is a refinement defined for normal-form repeated games with perfect monitoring. Using a reasoning similar to above, they also put a monotone restriction on the continuation values following deviations that benefit the other party.[14]

---

[11]The borrower cannot be made worse than being unmatched, as she can unilaterally terminate the relationship.

[12]To be sure, we only make comparisons at subperiod t^0, where a history $h(t^0)$ consists of a full set of actions (loan size, defaulting decision, continuation decisions) for each period from 0 to $t-1$. We do not make comparisons or impose such restriction at other subperiods. Moreover, it is allowed if one of the two histories has a *strictly* higher last-period loan but the continuation values following the two histories are the same. For example, this will be the case if the players simply ignore deviations to higher but repaid loans and keep the future course of play unchanged.

[13]The history $h'(t^0)$ may require deviations from both players, in which case both are deviating players.

[14]Our refinement is different from Aramendia and Wen[6] in several aspects. First, the monotonicity restriction in Aramendia and Wen[6] is only imposed on the non-deviating player's continuation value, while we require that the deviating player is also not punished if the deviation benefits the other player. Second, their refinement is only defined for normal-form repeated games with perfect monitoring, whereas in our context the stage game is of sequential moves. (Even though we can transform an extensive-form stage game into a normal-form one, the assumption of perfect monitoring would then become problematic because the strategy of the borrower in the stage game—a function mapping each loan size to a repayment decision—is not fully observed.) Moreover, we apply the logic of "justifiable punishments" only at the beginning of each period (i.e., only at t^0, not t^1 or t^2).

Finally, note that the definition of an orthodox social equilibrium does not require efficiency. So the main result of this paper, which establishes strictly increasing loans in all non-trivial orthodox social equilibria, holds for inefficient equilibria as well.

The Structure of Orthodox Social Equilibrium

Note first that a trivial (orthodox) social equilibrium always exists, in which the lender always offers a loan size of 0 and the borrower defaults on any positive loan. Note also that in any non-trivial social equilibrium, the relationship is never terminated on path by any player. This is because if a relationship is terminated at date t, then at that last period the borrower will default on any positive loan; as a result, the loan size at date t must be 0. But then, at the end of the second-to-last period $t-1$, each player has an incentive to terminate the relationship because both of them would prefer getting the values of an unmatched player right away (which is positive because the equilibrium is non-trivial), rather than waiting for another period with a payoff of 0 and then getting such values. But this is a contradiction to the optimality of continuing the relationship at the end of $t-1$. One immediate implication is that in any non-trivial orthodox social equilibrium, the borrower never defaults on path because the definition of such equilibria requires that any default is followed by termination.[15]

Now we state the main results of this paper.

Proposition 1. *Suppose that Assumption 1 holds. There exists a $\delta^* < 1$ such that the following holds: for all $\delta \in (\delta^*, 1)$, there exists a $\lambda^*_\delta < 1$ such that whenever $\lambda \in (\lambda^*_\delta, 1]$, the loan sequence $\{L_t\}_t$ is strictly increasing on the equilibrium path of any non-trivial orthodox social equilibrium of the game $G(\delta, \lambda)$.*

Proposition 2. *Suppose that Assumption 1 holds. Whenever $\delta \in (\delta^*, 1)$ and $\lambda \in (\lambda^*_\delta, 1]$, a non-trivial orthodox social equilibrium exists in the game $G(\delta, \lambda)$.*

Remark 1. Small Re-Matching Probability/Discount Factor
One may wonder about the structure of orthodox social equilibria when the re-matching probability or the discount factor is smaller than their thresholds. It turns out that if we assume linear payoff functions (i.e., $\frac{C(L)}{D(L)} \equiv \alpha \in (0, 1)$ for all L), we can characterize such equilibria for all parameter values.

As illustrated in Figure 1.1, when $\delta < 1 - \alpha$, only the trivial equilibrium ($L_t = 0$ for all t) exists. This feature often shows up in repeated games even without re-matching: when players do not care enough about the future, they can only play the unique SPE of the stage game. On the other hand, when the players are patient and the re-matching probability is low (i.e., the bottom right region of Figure 1.1), the threat of terminating a relationship is strong enough to sustain the maximum loan level from the beginning. This is because re-matching in this case is so unlikely

[15] Note that we impose the condition that default implies termination only in the definition of an orthodox social equilibrium. For a social equilibrium in general, although the relationship is never terminated by any player (as argued in the text), defaults may occur on the equilibrium path.

that losing the current relationship is nearly as bad as being thrown out of the market completely. Finally, notice that the threshold λ^*_δ increases with δ. This is intuitive: when players become more patient, it is easier to sustain $L_t \equiv L^*$ because the future cost of terminating a relationship is higher for a more patient borrower.

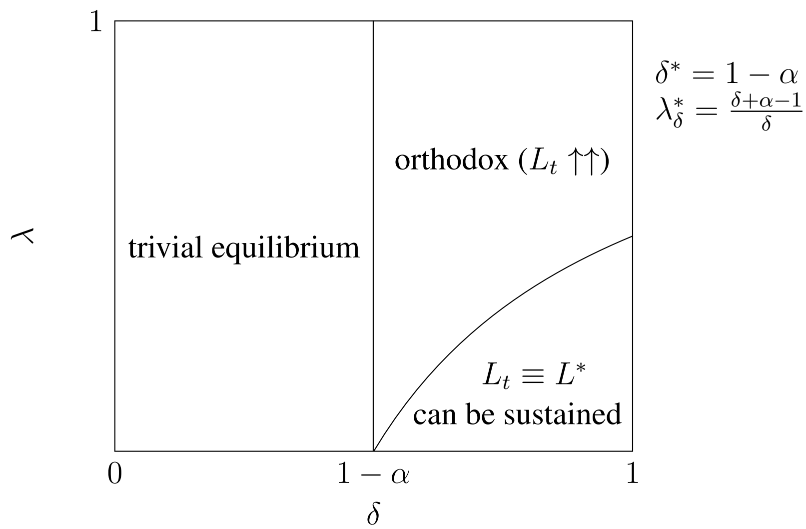

Figure 1.1: Equilibrium Characterization for Linear Payoff Functions

Therefore, an increasing sequence of loans will arise when the players are patient and the re-matching probability is high; the former ensures that the players are able to sustain some positive level interactions instead of only playing the trivial equilibrium, while the latter creates the need for “starting small” because a constant loan size over time will directly result in “hit and run.”

Remark 2. “Monotone Values” Restriction v.s. “Monotone Loans” Result
One important requirement in the definition of an orthodox social equilibrium is that a better loan-repayment history does not strictly decrease the continuation value of the relationship to either party. We motivated this requirement using an argument of “justifiable punishments.” But it is still fair to ask: does this monotone restriction on continuation values select equilibria with strictly increasing loan sizes by *assuming* it? The answer is “no” because the monotonicity in the restriction and in the result are across different dimensions. The “monotone values” restriction is a condition that compares continuation values across different histories at *fixed* t; in contrast, the “monotone loans” result is about the evolution of loan sizes *over time*. In addition, trivial equilibria in which loan size is always 0 are not ruled out by such a refinement. In fact, if the re-matching probability is too low (i.e., if the condition in Proposition 1 is violated), there are orthodox social equilibria in which the loan sizes are strictly *decreasing* over time.[16] These equibria are not ruled out by our

[16]For example, when “$\delta = 0.8, \alpha = 0.5, \boldsymbol{\lambda = 0.2}, L^* = 100$,” $\{L_t\}_t = \{93.8, 77.4, 67.1, 60.7, ...\}$ is sustained

refinement per se; however, they are impossible when the re-matching probability is high, because the borrower would want to "hit and run" if loans were decreasing and re-matching is easy.

As we will illustrate in Section 1.3: technically, what the "monotone values" restriction buys us is a sequence of *binding* no-default constraints; economically, the real driving forces of the gradualism result in this paper is the high re-matching probability and justifiable punishments.

1.3 Discussion

An Intuition for Proposition 1 with Linear Payoff Functions

To understand the intuition behind Proposition 1, consider the following case where payoff functions are linear in loan size:

$$y_t^L = \begin{cases} (1-\alpha)L_t, & \text{if repayment happens;} \\ -L_t, & \text{if default happens;} \end{cases}$$
$$y_t^B = \begin{cases} \alpha L_t, & \text{if repayment happens;} \\ L_t, & \text{if default happens.} \end{cases}$$

Let $\mathbf{L} = \{L_t\}_t$ be the sequence of loans on the equilibrium path of some non-trivial orthodox social equilibrium. As we observed earlier, a relationship is never terminated in any non-trivial social equilibrium. Therefore, we can write the lender's and the borrower's values at each date as:

$$V_t^L[(1-\alpha)\mathbf{L}] = \Sigma_{i=0}^{\infty}\delta^i(1-\alpha)L_{t+i},$$
$$V_t^B(\alpha\mathbf{L}) = \Sigma_{i=0}^{\infty}\delta^i\alpha L_{t+i}.$$

Note that the borrower's no-deviation (no-default) constraints are: for all t,

$$(1-\alpha)L_t \le \delta[V_{t+1}^B(\alpha\mathbf{L}) - \lambda' V_0^B(\alpha\mathbf{L})], \tag{1.7}$$

where the LHS is the current-period gains from default, and the RHS is the future cost of default which is the difference between the value of continuing the relationship and the value of terminating the relationship. (Recall that λ' is derived in footnote 9.)

One implication of part (ii) of Definition 2 is that in any orthodox social equilibrium, (1.7) must hold at equality for all t such that $L_t < L^*$. That is, the lender has an incentive to increase the

by an orthodox social equilibrium. In contrast to Proposition 1, our refinement cannot give a unique prediction on loan monotonicity when players are patient (high δ) and re-matching is unlikely (low λ). Such an equilibrium may not be very reasonable though, because under these parameters $L_t \equiv L^*$ can also be sustained as an orthodox equilibrium, and offering L^* in each period is the *first best* scenario that generates the highest possible payoffs to both parties in this model. Nonetheless, this example shows that the "monotone values" restriction itself does not select equilibria with increasing loans.

loan size as much as possible so that in each period, either the borrower's no-deviation constraint is binding, or the loan size reaches its maximum L^*. To see this, suppose that in some period t, (1.7) holds at strict inequality and $L_t < L^*$. Then in subperiod t^0 on the equilibrium path, the lender can consider deviating to $L'_t = L_t + \varepsilon < L^*$ such that

$$(1-\alpha)L'_t < \delta[V^B_{t+1}(\alpha\mathbf{L}) - \lambda' V^B_0(\alpha\mathbf{L})]. \tag{1.8}$$

Is this deviation profitable? Note first that such a deviation will not induce default in period t. This is because in subperiod t^1 when the borrower decides whether or not to repay, she will find:

$$(1-\alpha)L'_t < \delta[V'^B_{t+1} - \lambda' V^B_0(\alpha\mathbf{L})], \tag{1.9}$$

where V'^B_{t+1} is the borrower's continuation value at $(t+1)^0$ (i.e., the beginning of date $t+1$) following the new history with her repayment of a larger loan L'_t. (1.9) holds because of (1.8) and $V'^B_{t+1} \geq V^B_{t+1}(\alpha\mathbf{L})$, where the latter condition follows from the fact that the loan-repayment history with loans $\{L_0, L_1, ...L_{t-1}, L'_t\}$ and no default is better than that with $\{L_0, L_1, ...L_{t-1}, L_t\}$ and no default. As a result of (1.9), deviating to L'_t would not induce default.

Knowing this, the lender's total payoff from time t by offering L'_t will be $(1-\alpha)L'_t + V'^L_{t+1}$, which is strictly larger than $(1-\alpha)L_t + V^L_{t+1}[(1-\alpha)\mathbf{L}]$ because $L'_t > L_t$ and $V'^L_{t+1} \geq V^L_{t+1}[(1-\alpha)\mathbf{L}]$, where the second inequality is due to the fact that the lender's value following a better loan-repayment history is also weakly higher. But this just implies that deviating to L'_t is profitable for the lender, a contradiction to $\{L_t\}_t$ being on the equilibrium path of some (non-trivial) orthodox social equilibrium.

Therefore, when the deviation to a better loan-repayment history is not punished, the lender will have an incentive to increase the loan size as much as possible, so that in each period either the borrower's no-deviation constraint is binding, or L_t reaches its maximum L^*. It turns out that the latter case can be ruled out when λ (the re-matching probability) is high, so the borrower's no-deviation constraint is binding at each date.

Now we can explain the intuition for our main result, which says that when the discount factor and the re-matching probability are high, the equilibrium loan sequence in any orthodox social equilibrium is strictly increasing. On one hand, when past histories are unobservable, the high re-matching probability undermines the punishment power of the threat of terminating a relationship; so in order to induce repayment, there has to be some additional cost of starting a new relationship, which here is reflected in the fact that V^B_t is strictly increasing over time. On the other hand, because the value of becoming unmatched is constant while V^B_t grows, the cost inflicted upon the borrower by terminating the current relationship is increasing over time. This means that the cost of defaulting increases over time. Since the lender has an incentive to make the borrower's no-deviation constraint bind in each period, she will offer a strictly increasing sequence of loans.

Another way to understand the result, though more technical, is to directly inspect the loan sequence. First, high re-matching probability implies that a non-trivial equilibrium loan sequence cannot be constant (or decreasing) over time, for otherwise the borrower would simply default in the first period, run away, get rematched, and default again. As before, this calls for the value, as well as the loan size, to start small. But can it be that the loan size increases for a while and then stays constant starting from certain period $\tilde{t}$? This is the type of equilibrium showing up in Kranton[28] and Ghosh and Ray[22]. It is ruled out by our refinement because for any loan sequence of this type, the *future* interactions after period $\tilde{t}-1$ and after period $\tilde{t}$ are exactly the same (i.e., both consist of a constant loan size forever). That is to say, the borrower's costs of defaulting at $\tilde{t}-1$ and $\tilde{t}$, RHS of (1.7), are the same. But then, since by construction the loan size at $\tilde{t}-1$ is smaller than that at $\tilde{t}$, this implies that $L_{\tilde{t}-1} < L^*$ and the borrower's no-deviation constraint is not binding at $\tilde{t}-1$, a contradiction to what we have explained before (i.e., no-deviation constraint should be binding as long as $L_t < L^*$).[17] These arguments, together with the fact that we can rule out other possibilities (e.g., the sequence decreases once in a while, or it stays constant for a while and then increases, etc. See Lemma 6 in the Appendix), establishes that any non-trivial equilibrium loan sequence must be strictly increasing over time when the re-matching probability is high.

Equilibrium Non-Uniqueness

Though our refinement yields a unique prediction regarding loan monotonicity for interesting parameter values, it is unsurprising that equilibrium is not unique in this repeated game. We use the special case of linear payoff functions (as in the previous subsection) to illustrate one force that leads to the multiplicity of orthodox social equilibria. We then discuss the same multiplicity issue in Datta[18].

Consider the following parametrization:

$$\delta = 0.8, \alpha = 0.5, \lambda = 1, L^* = 100.$$

Under these parameters, the following two loan sequences can both be sustained by some orthodox social equilibria:

$$\{L_t\}_t = \{37.5, 60.9, 75.6, 84.7, 90.5, ...\};$$
$$\{L'_t\}_t = \{18.8, 30.5, 37.8, 42.4, 45.2, ...\}.$$

In our construction, $L'_t = \frac{1}{2}L_t$ for all t (before rounding up). In fact, with linear payoff functions, if $\{L_t\}_t$ is an equilibrium loan sequence, so is $\{\beta L_t\}_t$ for any $0 < \beta < 1$. This is because, moving from $\{L_t\}_t$ to $\{\beta L_t\}_t$, both sides of the borrower's no-default constraint (1.7) will be scaled

[17] We reiterate the reason why this type of social equilibria is unreasonable in the context of credit relationships. The reason lies in the strategy profile that supports such a loan sequence. We have seen that the borrower's no-deviation constraint is slack at $\tilde{t}-1$ on path. But the fact that it is an equilibrium implies that the lender does not want to increase the loan size at $\tilde{t}-1$ even by just a little. This means that either an increase in loan size causes default, or even after the repayment the lender is still worse off. In the former case, the future value of the relationship to the borrower must be lowered after she repays a higher loan; in the latter case, the future value of the relationship to the lender must be lowered after she offers a higher loan followed by repayment. In either case, had a better (off-equilibrium) loan-repayment history been reached, at least one player is "punished." This is the unreasonable feature suffered by *any* strategy profile that supports such a loan sequence which increases first and then stays constant.

by β; so it holds for the new sequence if and only if it holds for the original one.[18] Nevertheless, as Proposition 1 states, all these sequences are strictly increasing.

Datta[18] studies the linear payoff environment where the re-matching probability is 1. He focuses on the efficient equilibria within the class of social equilibria that entail no default on path, which he called *maximal social equilibria*. His main result is that the *value* sequences (of both parties) in any maximal social equilibrium are nondecreasing. While his prediction is weaker than ours, his refinement does not guarantee equilibrium uniqueness either. Under the same parameter values, both $\{L_t\}_t$ defined above and $\{L''_t\}_t$ defined below can be sustained by some maximal social equilibria:[19]

$$\{L''_t\}_t = \{35, 25, 100, 100, 100, ...\}.$$

Note that this loan sequence is non-monotone; in fact, Proposition 1 implies that most of the maximal social equilibria (except the orthodox ones) entail "unjustifiable punishments."

In the linear environment considered here, the sequence $\{37.5, 60.9, 75.6, 84.7, 90.5, ...\}$ is the unique equilibrium loan sequence that is both orthodox and maximal.[20] In nonlinear environments, even if we focus on *maximal orthodox* social equilibria, uniqueness is still not guaranteed. The main contribution of this paper is to find reasonable conditions that deliver a unique prediction of strictly increasing loans, despite multiplicity of equilibria.

Finally, it is worth pointing out that the key technical condition which drives our result is the binding no-default constraints. As discussed in the previous subsection, this condition is implied by the requirement of "justifiable punishments." This paper shows that such a requirement is enough to generate a unique prediction of strictly increasing loans, even without imposing efficiency. Datta[18] is unable to make a prediction regarding loan monotonicity because his analysis allows for "unjustifiable punishments" (and thus non-binding no-default constraints); and even after efficiency is imposed, only a weaker prediction regarding values can be made.

---

[18] With nonlinear payoff functions, this construction by scaling each loan size by β is no longer valid for creating another equilibrium loan sequence. But depending on the exact functional forms of C and D, reducing *all* loans by some (carefully chosen) amounts can still generate a new equilibrium loan sequence.

[19] To see that $\{L''_t\}_t = \{35, 25, 100, 100, 100, ...\}$ is maximal (i.e., efficient among the class of social equilibria that entail no default on path), note first that for any social equilibrium without default on path, $V_0^B \leq V_0^{*B}$, where V_0^{*B} solves $(1-\alpha)L^* = \delta(\frac{\alpha L^*}{1-\delta} - \lambda' V_0^{*B})$. This is because if there is a social equilibrium without default such that its loan sequence on path is $\{\tilde{L}_t\}_t$ and $\tilde{V}_0^B > V_0^{*B}$, without loss we have $\limsup \tilde{L}_t = L^*$ due to linearity of payoff functions. But then, for $\tilde{L}_t$ close enough to L^*, we have $(1-\alpha)\tilde{L}_t > \delta(\frac{\alpha \tilde{L}_t}{1-\delta} - \lambda' \tilde{V}_0^B)$, a contradiction to the requirement that the borrower does not default on path. It can be checked that $\{L''_t\}_t = \{35, 25, 100, 100, 100, ...\}$ achieves V_0^{*B} and satisfies the borrower's incentive constraints under our parametrization, so it is maximal.

[20] $\{L_t\}_t = \{37.5, 60.9, 75.6, 84.7, 90.5, ...\}$ is constructed using the algorithm proposed in the proof of Proposition 2 by setting the limit $L = L^*$. To see that $\{L_t\}_t$ is the unique loan sequence that is both orthodox and maximal, note first that by construction it can be supported by an orthodox social equilibrium using the strategy profile proposed in the proof of Proposition 2. Note also that by construction, $(1-\alpha)L_t = \delta(\frac{\alpha L_t}{1-\delta} - \lambda V_0^B)$, for all t. Since $L_t \to L^*$, we have $(1-\alpha)L^* = \delta(\frac{\alpha L^*}{1-\delta} - \lambda V_0^B)$; therefore $V_0^B = V_0^{*B}$, where V_0^{*B} is defined in footnote 19. This implies that such an orthodox social equilibrium is also maximal. Finally, the uniqueness of $\{L_t\}_t$ follows from the fact that it is the unique loan sequence that satisfies (A.1) and converges to L^* (see the proof of Proposition 2), and any other such sequence converging to $L < L^*$ generates a lower V_0^B.

Extension to Mixed Strategies

The preceding analysis has focused on pure-strategy equilibria, in which the loan amount, defaulting and termination decisions are deterministic in each period. In this subsection, we will first argue that in equilibria where the relationship is never terminated and default never occurs, randomizing over loan amounts is not compatible with the notion of orthodox social equilibrium. Next, we will show that if payoffs are linear in loan and if the borrower's defaulting decision breaks even in favor of the lender, then in any orthodox social equilibrium, indeed the relationship is never terminated and default never occurs, even if mixed strategies (in termination decisions) are allowed.

Note first that the definition of an orthodox social equilibrium can be readily extended to accommodate mixed strategies by interpreting the future values after a particular history as the *expected* future values. Then, let us take any orthodox social equilibrium that has no default and no termination on the equilibrium path. Suppose that at the beginning of period t with on-path history $h(t^0)$, equilibrium behavior involves mixing between two loan amounts, L_t and L'_t, such that $L_t < L'_t$. Since the lender is indifferent between L_t and L'_t, it means that

$$R(L_t) + V^L_{t+1}[h((t+1)^0)] = R(L'_t) + V^L_{t+1}[h'((t+1)^0)],$$

where the two histories at the beginning of $t+1$, $h((t+1)^0)$ and $h'((t+1)^0)$, only differ in their last-period (period t) loans. Since $L_t < L'_t$, our refinement requires that $V^L_{t+1}[h((t+1)^0)] \leq V^L_{t+1}[h'((t+1)^0)]$, but this is a contradiction to the indifference condition, as R is strictly increasing. This means that if default and termination never occur, randomizing over loan amounts is not compatible with the notion of orthodox social equilibrium.

Now suppose that:

- payoffs are linear in loan (as in the previous subsection);
- the borrower's defaulting decision breaks even in favor of the lender; i.e., she does not default when indifferent.[21]

In any mixed-strategy non-trivial orthodox social equilibrium, let us first show that default never happens on the equilibrium path. To see this, suppose (by contradiction) that default happens on path. Let t be the first period in which the borrower defaults on a positive loan. Since the relationship is terminated at the end of that period after default (as required by Definition 2), the lender can do strictly better by offering 0 loan in that period and then terminate the relationship. So a profitable deviation exists for the lender.

Next, we argue that the relationship is never terminated on path. Note that linearity and no default on path imply that the borrower's value is a fixed proportion $\frac{\alpha}{1-\alpha}$ of the lender's value. On the equilibrium path, let t be the first period in which the relationship is terminated with a positive probability. This implies that the borrower's value (on path) at the beginning of the next period

[21]This essentially assumes that the borrower's defaulting decision is deterministic, while allowing for mixing in loan sizes and termination decisions.

$t+1$ must be weakly lower than $\lambda' V_0^B$.[22] As a result, the borrower will default on any positive loan that leads to termination with a positive probability. So in this equilibrium, the history $h(t^2)$ on path after which termination is possible must have 0 loan in period t. But then, the lender should terminate the relationship one period earlier at the end of $t-1$, as she prefers getting the value of an unmatched lender right away (which is positive because the equilibrium is non-trivial) to waiting for another period with a payoff of 0 and then getting such a value. This is a contradiction to t being the first period in which the relationship is terminated with a positive probability.[23]

To summarize, we have argued that: (i) for any orthodox social equilibrium in mixed strategies where termination and default never occur on path, loan amounts must be deterministic; (ii) if payoffs are linear in loan and the borrower repays when indifferent, then all orthodox social equilibria in mixed strategies involve no termination and no default on path, and thus deterministic loan amounts. So Propositions 1 and 2 apply to these cases.

1.4 Conclusion

This paper studies a lender-borrower game in a pure moral hazard environment with anonymous re-matching. The main result states that as long as the discount factor and the re-matching probability are above certain thresholds, the size of loans along the equilibrium path of any orthodox social equilibrium is strictly increasing over time. This characterization gives a formal argument that qualifies the possibility of anonymous re-matching as a driving force of gradualism in long-term relationships. Certainly, re-matching is not the only reason for gradualism; gradual learning of another party's type, among other things, is also a reasonable and important driving force for such phenomena. However, given that in reality it is indeed costly to acquire the past history information of the other party in a new relationship, this paper provides insights from one specific aspect into understanding gradualism in long-term relationships, especially credit relationships.

This chapter is focused on the implications and remedies of anonymous re-matching which *exogenously* creates an extreme type of information frictions. Very often, however, information is left out not because it is not available, but because people choose not to process it due to their limited capacity/willingness to learn. In those cases, information frictions arise *endogenously*. The next two chapters turn to such scenarios where information frictions are caused by people's inattention. In particular, we will study optimal information transmission protocols when the information recipient is rationally inattentive.

[22] To see this, note that if the relationship is terminated (with a positive probability) because of the borrower's termination, we must have $V_{t+1}^B \leq \lambda' V_0^B$; if it is because of the lender's termination, we must have $V_{t+1}^L \leq \lambda' V_0^L$. In the latter case, since linearity implies that $V_{t+1}^B = \frac{\alpha}{1-\alpha} V_{t+1}^L$ and $V_0^B = \frac{\alpha}{1-\alpha} V_0^L$, we again have $V_{t+1}^B \leq \lambda' V_0^B$.

[23] More precisely, this is a contradiction unless $t=0$; but if the relationship is terminated with a positive probability at $t=0$, then the lender's initial value satisfies $V_0^L = 0 + \delta\lambda' V_0^L$, so that $V_0^L = 0$. But this can happen only in a trivial equilibrium.

Chapter 2

Attention Management

This chapter is joint work with Elliot Lipnowski and Laurent Mathevet. Part of this chapter was published in American Economic Review: Insights.

2.1 Introduction

Information is a gift that may not always be accepted and, hence, useful. Speaking to a toddler about grammar may not improve his linguistic abilities, just as an adult may learn less from a book, an email, or a contract that contains too much detail. Simon[41] offers an explanation for why less detailed communication may convey more: "What information consumes is rather obvious: it consumes the attention of its recipients." Failures to recognize this fact can have counterproductive effects: consumers are frequently confused by nutritional labels; patients can be overwhelmed in parsing side effects of medications; and so on (see Ben-shahar and Schneider[11]). As Simon[42, p.144] puts it: "The real design problem is not to provide more information to people ... but [to design] intelligent information-filtering systems."

This paper studies an information-filtering problem aimed at managing a receiver's attention. We study this problem in the context of a strong preference alignment: the sender and the receiver have the same material motives, but the receiver incurs an attention cost that the sender does not internalize. This is an instance of paternalistic, benevolent design, of the same nature as the teacher-student and doctor-patient interaction, as well as many others.

Consider a teacher preparing her students for an external evaluation. Each question in the test will have a correct answer $\theta \in \{0, 1\}$, where 0 means false and 1 means true, the two being correct in equal proportion. A student wants to answer questions correctly, and the teacher wants the same. For each question, the student must choose a number $a \in [0, 1]$, and his score on that question will be $1 - (a - \theta)^2$ (the teacher has the same objective). Depending on the course, the student will be more or less confident, and rightfully so, about the answers. For example, a completely ignorant student believes that both answers are equally likely for every question. Think of a course as an intended distribution $p \in \Delta[0, 1]$ of student belief $\nu = \mathbb{P}(\theta = 1)$ in the test. For example, for a fraction $p(\nu = 0.1)$ of questions, the student will be 10% sure the answer is true.

The teacher could cover all possible questions and their correct answers, so that a sufficiently attentive student will get a perfect grade. This corresponds to $p(1) = p(0) = 0.5$, because such a student would be sure that the answer is true when it is, and sure that it is false when it is, each occurring half of the time. Of course, a student may not be willing, or able, to process so much information. Instead of belief distribution p, he may experience a less well-informed distribution q. Suppose, for instance, a student incurs a cost $\kappa\mathbb{E}(\mu - \nu)^2$ of information processing to move his belief from prior μ to posterior ν. If $\kappa > 1$, any information the teacher provides is ignored, as the student deems it too costly.[24] In particular, then, a fully informative course is optimal (as is any other course). If $\kappa < 1$, a student wants as much information as possible, and so will be fully attentive. Full information is again optimal for the teacher.

Consider the same teacher, but now the test questions will have correct answers $\theta \in \{0, \frac{1}{2}, 1\}$, where $\frac{1}{2}$ means uncertain. Suppose that the exam has an equal fraction of all three answers; the scoring rule is the same as before; and a student's information cost is $\frac{1}{4}\mathbb{E}||\nu - \mu||^2$ where the distance is Euclidean. If the teacher covered everything, then a student would rationally pay partial attention, only learning to distinguish between true and false. That is, he would divide possible questions into two groups: those that are "probably true" (consisting of all true questions and some uncertain), and those that are "probably false" (consisting of all false questions and some uncertain).

This attention choice avoids large mistakes, such as answering true when the correct answer is false or vice versa, but it entails many small mistakes, as the student never gives a perfect answer. Given that the teacher and the students have aligned testing motives, it may seem surprising that withholding information could improve test scores. But, indeed, the teacher can strictly increase the average score by teaching students to answer -1 if and only if the correct answer is false. With this lesson plan, to which one can show the student will be fully attentive, the teacher eliminates all mistakes, both large and small, concerning false questions, and only leaves fine mistakes involved in discerning true from uncertain questions.

Justifying inattention is almost a matter of introspection: people and organizations have limited information processing capacity (Sims[43, 44]). So, it seems natural to ask how much information one should give such agents. In our model, a principal provides information to a rationally inattentive agent. The agent cares about his material benefit, but also about the cost of processing information. The principal, however, is only motivated by the agent's material benefit, such as the student's score or the fitness of the patient's medical decision. From a subjective cost-benefit analysis, the agent chooses how much of the principal's provided information to process. In other words, the agent decides how informed he wants to be, taking whatever information the principal makes available as an upper bound. Choosing that upper bound in an optimal way is our design problem.

Attention management gives a special role to information that is *incomparable* to what the agent would acquire given full information. Although more information is sufficient for better

[24] In this example, the student's indirect utility $U_A(\nu, \kappa) = -\nu(1-\nu) - \kappa(\mu - \nu)^2$, whose expectation he maximizes, is concave when $\kappa > 1$ and convex if $\kappa < 1$.

decision-making (Blackwell[12]), it is not necessary. A person with a different kind of information from that of another person, neither more nor less, can also make better decisions. In the teacher example, a course that focuses on false answers and keeps some uncertainty about true and uncertain is neither more nor less informative than one that focuses on true answers and keeps some uncertainty about false and uncertain. These two courses simply emphasize different angles of the same topic. Yet one course could result in higher average test scores than the other. Standard revealed preference reasoning tells us that our agent cannot be made better informed than what he pays attention to under full information (call that q). Therefore, any hope for improvement upon full disclosure must come from providing incomparable information, neither better nor worse (than q), but targeting the most pertinent dimensions for the task at hand. This manipulation prevents the agent from choosing q, while preserving some other policies that result in better decision-making.

When the state can take on only two values, as in the true/false teacher example, full information is always principal-optimal for *all* material motives and convex attention costs. This is our second theorem (our first theorem proves that a solution to the principal's problem exists, under general conditions, and can be found within a specific tractable class). Despite the binary characteristic of the state, there are still many incomparable information policies; however, if the agent willingly paid full attention to incomparable policies, then it can be shown that he would pay full attention to a third that is more informative than both. Hence, the principal would prefer to choose that one instead. In general, the principal sees no benefit from withholding information.

Following the above theorem, we study the simplest specification that departs from binary uncertainty: the Quadratic Model, defined by a motive to match the state, quadratic attention costs, and three evenly-spaced states. The second version of the teacher example is an instance of this. In contrast to the binary-state world, wherein the agent may also not pay full attention to full information, the principal may strictly prefer to withhold information—though her sole objective is to help the agent make good decisions. The Quadratic Model is the simplest available model with more than two states, because the agent's preferences over information are constant in a given "direction." For this reason, the model can isolate the new issues that come with additional states. In particular, it suggests two lessons on the nature of attention management. First, the principal should give as much information as possible, given its direction. This simplifying feature mutes the problem of choosing the amount of information, and isolates that of choosing *which aspects* of the world to emphasize. Second, an optimal policy should be as easy to pay attention to as possible, given its instrumental value. In practice, this efficiency concern garners a lot of structure. Based on these principles, an optimal policy can be derived explicitly (Proposition 5), taking a friendly form: the principal either downplays or exaggerates the state.

The remainder of this chapter is organized as follows. Below, we discuss related literature. Section 2.2 sets up a general model of attention management and describes a specialization with quadratic payoffs and three states. Section 2.3 presents our main results for the general model. Section 2.4 characterizes incentive compatibility and gives the full solution in the Quadratic Model. Section 2.5 discusses generalizations of our results, and 2.6 concludes. All proofs for the results in this chapter are in Appendix B.

Related Literature

Our paper lies at the interface of two literatures: persuasion of decision makers through flexible information (Kamenica and Gentzkow[26]; Aumann and Maschler[8]) and rational inattention (Sims[43, 44]).

Among other generalizations, the Bayesian persuasion framework has been extended to include costly information provision by the principal (Gentzkow and Kamenica[21]) and parallel costly information acquisition by the agent (Matysková[35]). Other works study persuasion games with departures from "classical" preferences, such as psychological preferences (Lipnowski and Mathevet[33]), ambiguity aversion (Beauchêne et al.[10]), and heterogenous beliefs (Alonso and Câmara[4]; Galperti[20]). A key feature of all aforementioned papers is that the receiver is a passive learner of provided information: he automatically processes whatever information is revealed by the sender. In our paper, the receiver will actively filter his own information to limit information processing costs. Our analysis illustrates that the belief-based approach commonly adopted to study persuasion games is still applicable in the face of this complication.

The rational inattention literature (Sims[43, 44]; Caplin and Dean[14]; Caplin and Martin[15]; Matějka and McKay[34], etc.) studies optimal decision making by agents who incur an attention cost (or face an attention constraint), and so decide which of the available information to process before acting. These models are the building blocks of our agent's problem, given the principal's disclosure choice.

One lesson from our paper is that attention management is fundamentally about choosing which aspects of the state to reveal: this is apparent from comparing Theorem 2 to Proposition 5. The literature on multidimensional cheap talk (Battaglini[9]; Levy and Razin[31]; Chakraborty and Harbaugh[16]) also focuses on revealing lower dimensional aspects of a state. While that literature focuses on trading off dimensions as a way to relax a sender's incentive constraints, lower-dimensional information is provided in our paper to restrict the receiver's latitude to pay partial attention.

Our paper also contributes to costly information acquisition under moral hazard through its main strategic tension. Previous work (e.g., Dewatripont and Tirole[19] and Li[32]) has identified various ways to provide better incentives for information acquisition. More directly pertinent, in a setting of delegated decision-making, Szalay[45] illustrates that eliminating "safe" actions from the agent's choice set can help align incentives to seek useful information, and that this may be valuable even if the principal never benefits from restricting the agent's behavior ex-post. Proposition 5 illustrates how limiting the information available to the agent *endogenously* eliminates such safe behavior.

The most related works, featuring information transmission under some form of inattention, are those of Bloedel and Segal[13] and Lester et al.[30].[25] Bloedel and Segal[13] study a problem

[25]Less related are Persson[37] and Hirshleifer et al.[24]. The former studies a model in which competing firms exploit consumers' limited attention through deliberate information overload. The latter shows, in a verifiable disclosure setting, how a simple form of inattention (specifically, some fraction of receivers being exogenously perfectly inattentive) can break standard equilibrium unraveling results.

with very similar motivation to ours, but with substantively different modeling assumptions. While there are other differences (they restrict attention to a binary-action world, with a possibly infinite state space), the most substantive is a qualitatively different cost specification. In information theory, the cost of information processing is often modeled as a cost of reducing uncertainty. But what is that uncertainty about? When information is provided by an intermediary (the principal), the receiver's uncertainty could be about the principal's realized message or about the underlying state of the world. Think, for example, about inviting a friend who has the bad habit of being late to social functions. The uncertainty will not be about what he says ("7pm", "8pm", "after dinner", or something else), but rather about his actual arrival time. In more complex communications, such as those by a technical expert who may employ jargon, the messages themselves can be difficult to parse. In Bloedel and Segal[13], a principal commits to an experiment, and the agent bears a (reduction of entropy) cost to learn what message the principal has sent; their cost is not directly related to reducing uncertainty about the state. In our model, the agent bears a (e.g., quadratic) cost to learn the state, the cost being unrelated to reducing uncertainty about which message the principal sent. Lester et al.[30] analyze a model of evidence exclusion in courts of law. In their model, a judge chooses which of finitely many pieces of evidence should be considered by the jury, who then choose a subset of those to examine at a cost. The authors provide examples in which evidence exclusion leads to fewer sentencing errors. Our paper studies this same basic tradeoff in a flexible information-choice framework, with a view toward giving general prescriptions about optimal attention management.

2.2 The Model

The General Model

Let Θ be a finite set of states and A be a compact metrizable space of actions. An agent must make a decision $a \in A$ in a world with uncertain state $\theta \in \Theta$ distributed according to (full-support) $\mu \in \Delta\Theta$. When the agent chooses a in state θ, his material payoff is given by $u(a, \theta)$, where $u : A \times \Theta \to \mathbb{R}$ is continuous. The principal's payoff is equal to the agent's material utility, u.

In addition to his material utility, the agent also incurs an attention cost. As in the rational inattention literature, this cost is interpreted as the utility loss from processing information. To define it, first let

$$\mathcal{R}(\mu) := \left\{ p \in \Delta\Delta\Theta : \int_{\Delta\Theta} \nu \, \mathrm{d}p(\nu) = \mu \right\}$$

be the set of **(information) policies**, which are the distributions over the agent's beliefs such that the mean equals the prior. It is well-known, for example from the work of Kamenica and Gentzkow[26], that signal structures and information policies are equivalent formalisms. For the purpose of this paper, an attention cost function is a mapping $C : \Delta\Delta\Theta \to \mathbb{R}_+$ such that for every policy p,

$$C(p) = \int_{\Delta\Theta} c \, \mathrm{d}p \tag{2.1}$$

for some convex continuous $c : \Delta\Theta \to \mathbb{R}_+$.[26] Jensen's inequality tells us that an agent who increases his attention, in the sense of obtaining a policy p that is more (Blackwell) informative than q, denoted $p \succeq_\mu^B q$,[27] incurs a higher cost for p than for q.

The timing of the game is as follows:

– The principal first commits to an information policy $p \in \mathcal{R}(\mu)$.

– The agent then decides to what extent he should pay attention to p: he chooses a policy $q \in \mathcal{R}(\mu)$ such that $q \preceq_\mu^B p$. Such a policy q is called an **(attention) outcome**.

– Finally, the agent's belief is drawn from q, at which point he takes an action $a \in A$. The agent's belief is his updated belief following reception of a message sent from the principal's signal structure.

We study principal-preferred subgame perfect equilibria.

It is convenient to work with the principal's indirect utility at $\nu \in \Delta\Theta$

$$U_P(\nu) = U(\nu) := \max_{a \in A} \int_\Theta u(a, \cdot)\, \mathrm{d}\nu,$$

and the agent's indirect utility

$$U_A(\nu) = U(\nu) - c(\nu).$$

Note that the attention cost does not affect the agent's optimal choice of a conditional on a given belief. The principal's problem can therefore be formalized as follows:

$$\begin{aligned} &\sup_{p,q} \int_{\Delta\Theta} U_P \,\mathrm{d}q \\ &\quad \text{s.t. } \; p \in \mathcal{R}(\mu) \;\text{ and }\; q \in G^*(p) \end{aligned} \tag{2.2}$$

where

$$G^*(p) := \operatorname*{argmax}_{q \in \mathcal{R}(\mu):\; q \preceq^B p} \left\{ \int_{\Delta\Theta} U \,\mathrm{d}q - C(q) \right\} = \operatorname*{argmax}_{q \in \mathcal{R}(\mu):\; q \preceq^B p} \int_{\Delta\Theta} U_A \,\mathrm{d}q$$

is the agent's optimal garbling correspondence. An information policy $p^* \in \mathcal{R}(\mu)$ is **(principal-) optimal** if (p^*, q^*) solves (2.2) for some outcome $q^* \in \Delta\Delta\Theta$. The corresponding q^* is an **optimal (attention) outcome**.

As formalized, it is clear that the principal's problem is one of delegation. The policy p chosen by the principal only appears in the constraint and does not directly affect any party's payoff. In effect, the principal makes available a menu of information policies, from which the agent picks his preferred one.

[26]One well-known example has $c(\nu) \propto H(\mu) - H(\nu)$, where H is Shannon entropy.

[27]For any $p, q \in \mathcal{R}(\mu)$, $p \succeq_\mu^B q$ (or simply $p \succeq^B q$) if p is a mean-preserving spread of q, that is, there is $r : \Delta\Theta \to \Delta\Delta\Theta$ such that (i) $p(S) = \int_{\Delta\Theta} r(S|\cdot)\, \mathrm{d}q, \forall$ Borel $S \subseteq \Delta\Theta$ and (ii) $r(\cdot|\nu) \in \mathcal{R}(\nu), \forall \nu \in \Delta\Theta$.

The Quadratic Model

Consider the following specification of our model:

1. Three (evenly spaced) states: $\Theta = \{-1, 0, 1\}$;
2. "Match the state" material preferences: $A = \text{co}\,(\Theta) \subseteq \mathbb{R}$ and $u(a,\theta) = -(a-\theta)^2$;
3. Quadratic information costs: $c(\nu) = \kappa||\nu - \mu||^2$, where $||\cdot||$ is the Euclidean norm and $\kappa > 0$ is a cost parameter.

We denote the common prior by $\mu = (\mu_{-1}, \mu_0, \mu_1) \in \text{int}(\Delta\Theta)$.

Under assumptions (2) and (3), the agent's attention preference is both translation- and scale-invariant. That is, his willingness to pay attention to binary information is unchanged if one translates both beliefs equally, or moves one belief toward or away from the other. By making information preferences especially simple "within a dimension" of the space of beliefs, this model isolates the richness imparted by multidimensional information. This three-state model is arguably the simplest model that escapes our full disclosure result in Section 2.3.

2.3 Optimal Disclosure

Simplifying Disclosure

Given the sequential nature of our game and the infinite number of alternatives in the agent's menu of attention policies, it is not immediate that a solution to the principal's problem exists. Our first result shows that it does indeed exist, and additionally shows that some optimum takes a special and convenient form. Say that an information policy $p \in \mathcal{R}(\mu)$ is **incentive compatible (IC)** if the agent finds it optimal to pay full attention to it, i.e., if $p \in G^*(p)$. Say that an information policy is **nonredundant** if supp(p) is affinely independent.

Theorem 1. *There exists a solution p^* to*

$$\begin{aligned} \sup_p \int_{\Delta\Theta} & U_P \, dp \\ \textit{s.t.} \quad & (i) \quad p \in \mathcal{R}(\mu) \\ & (ii) \quad p \textit{ is IC} \\ & (iii) \quad p \textit{ is nonredundant.} \end{aligned} \tag{2.3}$$

Moreover, any solution p^* to (2.3) is such that (p^*, p^*) is a solution to (2.2).

The theorem establishes existence of a solution to the general model and establishes that limiting attention to nonredundant, incentive compatible information policies is without loss. The existence result follows from a continuity argument, relying on the observation that the "garbling correspondence" is continuous. That IC policies are without loss, analogous to the revelation principle, relies on a revealed preference reasoning: if q is an optimal attention outcome, then it must be an optimal garbling of itself, $q \in G^*(q)$.

Finally, an important simplification is the sufficiency of nonredundant policies. A useful analogy is the optimality of posted price mechanisms in the sale of a good to a buyer with privately known valuation. There, revenue equivalence reduces the set of IC mechanisms to the set of increasing allocation rules; the principal's objective is a linear functional of this allocation rule; and the extreme points (of the set of increasing allocation rules) are simply the step functions, which correspond to posted price mechanisms. In the present setting, nonredundant policies are extreme in the set of all information policies.[28] Since Θ is finite, nonredundancy implies the need for fewer messages than there are states. Thus, just as the celebrated posted price result does, Theorem 1 reduces an infinite-dimensional optimization problem to a finite-dimensional one.

In addition to having a small support, a nonredundant information policy enjoys another technical convenience: its set of garblings is straightforward to characterize, which simplifies the task of checking whether it is IC. As proved in Lemma 8 in Appendix B, for all $p, q \in \mathcal{R}(\mu)$ with p nonredundant,

$$p \succeq^B q \quad \iff \quad \text{supp}(q) \subseteq \text{co}\left[\text{supp}(p)\right]. \tag{2.4}$$

Full Disclosure in Binary-State Environments

Let the full disclosure policy, $p^F \in \mathcal{R}(\mu)$, be defined as $p^F(\delta_\theta) = \mu(\theta)$ for all $\theta \in \Theta$, where $\delta_\theta \in \Delta\Theta$ is the belief that puts probability 1 on state θ.

Theorem 2. *If the state is binary, then full disclosure is a principal-optimal attention policy. That is, (p^F, q^*) is a solution to (2.2) for some $q^* \in \Delta\Delta\Theta$.*

The principal never has a reason to provide more or less information in the Blackwell sense than what the agent would acquire given full information.[29] More information would be ignored, and less would harm the principal. The only way to have the agent bear a greater cost of attention and make a better decision, is to provide a policy that is *incomparable* to the attention outcome from full information (represented by q^* (in red) in Figure 2.1). However, no such nonredundant policy is IC when the state is binary. For example, if the principal offered the policy $\tilde{q}$ (in blue) in Figure 2.1, which is incomparable to q^* by condition (2.4), then the agent would not pay full attention to $\tilde{q}$. Indeed, q^{**} (in purple) is a garbling of $\tilde{q}$ by (2.4), and it clearly gives the agent a strictly higher payoff than $\tilde{q}$ (since $\tilde{u}_A^* > \tilde{u}_A$). Therefore, $\tilde{q}$ is not IC. As the principal can only induce the agent to pay attention to q^* or to less informative policies than q^*, she finds it optimal to provide full information.

[28] Some care is required, as the set of IC information policies is not convex; but it is an extreme set, which enables an analogous argument.

[29] In knife-edge cases, the agent might have many best responses to full information; in that case, we take q^* to be the most informative one (which exists in the binary-state setting).

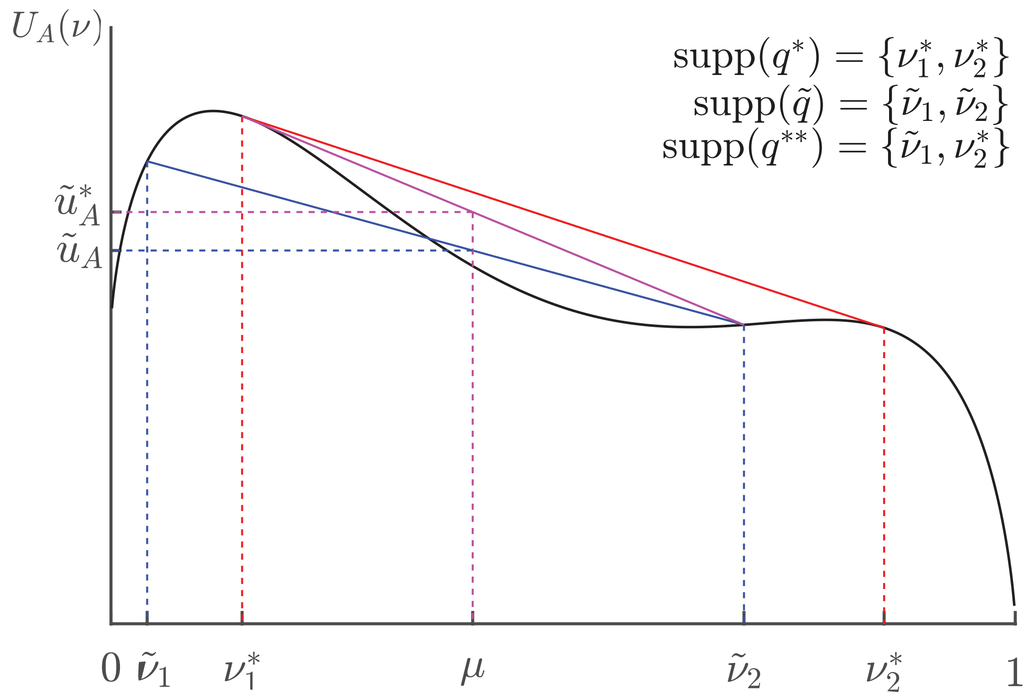

Figure 2.1: Blackwell-Incomparable Policies When $|\Theta| = 2$

With three or more states, the above reasoning breaks down, because there exist policies that are both IC and incomparable to the full-information outcome (q^*). The possibility of informing the agent along many distinct dimensions can enable attention outcomes that are *not* less informative than q^*. This type of partial information leaves room for attention management.

2.4 The Quadratic Model

We now study a parametric model with a richer state space, in which the agent's material payoff is the standard quadratic loss between the state and his action, and his attention cost is based on the Euclidean distance between the prior and his updated beliefs. In this model, as in the general model, the principal's problem can be seen as choosing the directional aspect and the scale aspect of an information policy, where the direction changes the agent's differential uncertainty between states—in much the same way as different textbooks can lead to different relative knowledge of the same topic—and scale refers to the "quantity" of information. Notice that information had no directional feature in the binary environment of the previous section: any (nonredundant) policy had a leftmost and a rightmost posterior, with information fully described by their distance from the prior. Therefore, Theorem 2 could be interpreted as the statement that, absent directional concerns, there is never any reason for the principal to restrict the scale of information. In the Quadratic

Model, while direction and scale are both relevant levers, the principal will never have any reason to restrict the scale of information once its direction is chosen. In this sense, the Quadratic Model is the simplest departure from binary uncertainty, because only the direction of optimal information matters, which is arguably the central aspect of information disclosure in a richer state space.

In this section, we characterize incentive compatibility, present principles of optimal attention management, and compute the resulting optimal policy for the Quadratic Model.

Payoffs and Incentives

Below we index statistical measures, such as expectation $\mathbb{E}$ and variance $\mathbb{V}$, by the distribution of the underlying random variable. For example, $\mathbb{E}_p\mathbb{V}_\nu\theta$ refers to $\mathbb{E}_{\nu\sim p}\left[\mathbb{V}_{\theta\sim\nu}(\theta)\right]$.

Principal's Payoff. The marginal distribution of actions is sufficient to compute the principal's expected payoffs in the Quadratic Model, due to the "match the state" motive.

The agent's optimal action at any posterior belief $\nu \in \Delta\Theta$ is

$$a^*(\nu) := \operatorname*{argmax}_{a\in A} u(a,\nu) = \mathbb{E}_\nu\theta.$$

The principal's value is then

$$U_P(\nu) := -\mathbb{E}_\nu[(\theta - a^*(\nu))^2] = -\mathbb{V}_\nu\theta, \tag{2.5}$$

which is strictly convex in ν.

Take any incentive compatible information policy $p \in \mathcal{R}(\mu)$ and note that

$$\int_{\Delta\Theta} U_P \,\mathrm{d}p = -\mathbb{E}_p\mathbb{V}_\nu\theta = \mathbb{V}_p\mathbb{E}_\nu\theta - \mathbb{V}_\mu(\theta) = \mathbb{V}_p[a^*(\nu)] - \mathbb{V}_\mu(\theta), \tag{2.6}$$

by the law of total variance. Therefore:

Observation 1. *In the Quadratic Model, for any IC policies p and p', the principal strictly prefers p to p' if and only if $\mathbb{V}_p[a^*(\nu)] > \mathbb{V}_{p'}[a^*(\nu)]$.*

Psychological vs. Material Incentives. Another advantage of the Quadratic Model, this one pertaining to the agent, is the simple characterization of IC by a sequence of comparisons between psychological cost and material benefit.

The agent's attention cost at ν is

$$c(\nu) = \kappa||\nu - \mu||^2 = \kappa\sum_\theta(\nu_\theta - \mu_\theta)^2,$$

so that his net indirect utility is

$$U_A(\nu) := U_P(\nu) - c(\nu) = \left[\left(\sum_\theta \nu_\theta\theta\right)^2 - \kappa\sum_\theta \nu_\theta^2\right] + h(\nu) \tag{2.7}$$

where $\nu_\theta := \mathbb{P}(\theta|\nu)$ and $h(\nu)$ is affine in ν.

To understand the driving force behind IC in this model, consider first a binary-support policy. Fix some ν, ν', and let $x_\theta := \nu_\theta - \nu'_\theta$ for all $\theta \in \Theta$. For $\epsilon \in [0,1]$,

$$\begin{aligned}\frac{1}{2}\frac{\mathrm{d}^2}{\mathrm{d}\epsilon^2}U_A(\nu + \epsilon(\nu' - \nu)) &= \left(\sum_\theta x_\theta \theta\right)^2 - \kappa \sum_\theta x_\theta^2 \\ &= |\mathbb{E}_\nu\theta - \mathbb{E}_{\nu'}\theta|^2 - \kappa||\nu - \nu'||^2.\end{aligned} \tag{2.8}$$

This derivative measures the curvature of U_A between ν and ν' and brings out two notions of distance between beliefs:

- The **choice distance**, $|\mathbb{E}_\nu\theta - \mathbb{E}_{\nu'}\theta|$, which describes the change in action caused by a change in beliefs.
- The **psychological distance**, $\sqrt{\kappa}||\nu - \nu'||$, which is a standard distance between beliefs (the scaled Euclidean metric, to be precise).

The psychological distance measures the marginal cost of extra attention in a given direction, while the choice distance measures the marginal benefit. A policy with binary support $\{\nu, \nu'\}$ is IC if and only if U_A is convex between ν and ν' (hence, if and only if the choice distance exceeds the psychological one).[30] The next proposition generalizes this fact to all nonredundant policies.

Proposition 3. *A nonredundant policy $p \in \mathcal{R}(\mu)$ is IC if and only if*

$$|a^{n+1} - a^n| \geq \sqrt{\kappa}||\nu^{n+1} - \nu^n|| \quad \forall n, \tag{2.9}$$

where supp$(p) = \{\nu^1, \ldots, \nu^N\}$ such that $a^n := \mathbb{E}_{\nu^n}(\theta)$ and $a^1 \leq \cdots \leq a^N$.

Condition (2.9) is much weaker than IC, only requiring incentive compatibility between any two messages that induce consecutive actions—that is, their choice distances should be weakly larger than their psychological distances. We refer to this property as **order-IC**.

In the Quadratic Model, the characterization of IC in Proposition 3 greatly simplifies the principal's problem. The first step for the principal is to understand the agent's (constrained) problem, but it typically depends on the shape of U_A over (the convex hull of) a policy's whole support. In general, it can be made difficult by boundary issues in the agent's problem. In the Quadratic Model, however, the agent's behavior can be understood only based on what happens on the edges of the policy's support, as formalized in (2.9).

[30] This follows from Jensen's inequality, together with the observation that $\epsilon \mapsto U_A(\nu + \epsilon(\nu' - \nu))$ is (being quadratic) either globally weakly convex or globally strictly concave.

Optimal Attention Management

Attention management is a form of damage control aimed at minimizing expected losses. Unless attention is so cheap that the agent willingly pays full attention to perfect information, the lack of attention will be the source of inevitable mistakes whose impact should be contained. Two main principles of good attention management emerge from our model and lay the foundations for the optimization that follows. The first one governs the quantity of information, and the second the directional aspect of information.

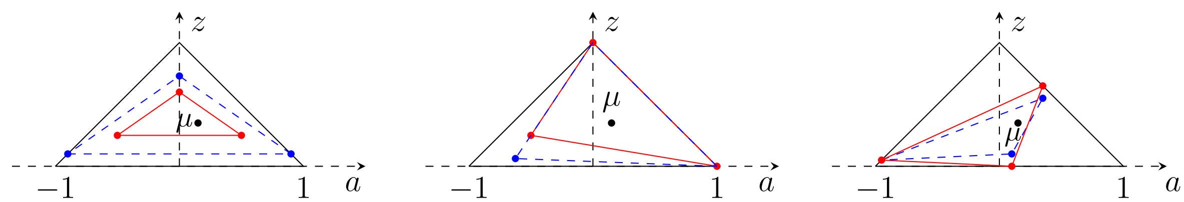

(a) Information Saturation (b) Information Saturation (c) Cost Efficiency

Figure 2.2: Principles of Attention Management

Information saturation says that information should be maximal once its (consecutive) direction is chosen. Given a policy with support $\{\nu^n\}_n$, let

$$\text{direction} := \left(\frac{\nu^{n+1} - \nu^n}{||\nu^{n+1} - \nu^n||} \right)_n . \tag{2.10}$$

Direction is the vector of normalized changes in uncertainty from one belief to the next. In the (a, z) space of Figure 2.2, let z be the agent's belief that θ is zero and a his optimal action.[31] By (2.10), direction is unaffected by any affine transformation of beliefs (as in Figure 2.2a) or changes in non-consecutive slopes (as in Figure 2.2b). Such transformations leave the comparisons in (2.9) (hence, IC) unchanged. Since IC is unchanged by making a policy more informative while holding its direction fixed, and since the principal wants the agent to pay more attention, information should be maximal once its direction is chosen.

Cost efficiency says that an optimal policy should be as easy to pay attention to as possible, given its instrumental value. Said differently, an optimal policy should induce an action distribution in a way that preserves IC for the largest set of cost parameters. A violation of cost efficiency means that the same behavior could be obtained from a less attentive agent. But then, intuitively, the principal should be able to induce even better-informed decisions from the more attentive agent. Rearranging (2.9), a policy p with supported beliefs $\nu^n = \nu_{(a^n, z^n)}$ is IC, if and only if $\kappa \leq 2$ and

$$\left| \frac{z^{n+1} - z^n}{a^{n+1} - a^n} \right| \leq s^*(\kappa) := \sqrt{\frac{2-\kappa}{3\kappa}} \quad \forall n. \tag{2.11}$$

[31] The associated belief is $\nu_{(a,z)} = z\delta_0 + \frac{1-z+a}{2}\delta_1 + \frac{1-z-a}{2}\delta_{-1}$, for $z \in [0, 1]$ and $a \in [z-1, 1-z]$.

Thus, a policy is ruled out by cost efficiency if the same action distribution can be implemented by another policy with smaller absolute slopes, for example the solid triangle in Figure 2.2c.

Now, we can combine Theorem 1 with the information saturation and cost efficiency principles to pare down the search for optimal information in the Quadratic Model. From here, we solve the residual (brute force) optimization problem over a simplified policy space.

Before stating the result, we introduce a language to speak about some special policies.

Say the principal **downplays** the state if she sends message:

$$m = \begin{cases} 0 & : \text{ with probability } 1, \text{ if } \theta = 0 \\ \theta & : \text{ with probability } 1-\pi_\theta, \text{ if } \theta \neq 0 \\ 0 & : \text{ with probability } \pi_\theta, \text{ if } \theta \neq 0. \end{cases}$$

The principal downplays the state if she misreports the truth only by occasionally saying 0 instead of an extreme state. Complete downplaying ($\pi_{-1} = \pi_1 = 1$) conveys no information; no downplaying ($\pi_{-1} = \pi_1 = 0$) fully discloses the state; other forms convey partial information. Let $\left\{p^D_{(\pi_{-1},\pi_1)} : (\pi_{-1},\pi_1) \in [0,1]^2\right\}$ denote the policies induced by downplaying the state. Say that the principal **exaggerates** the state if she sends message:

$$m = \begin{cases} \theta & : \text{ with probability } 1, \text{ if } \theta \neq 0 \\ 1 & : \text{ with probability } 1-\pi, \text{ if } \theta = 0 \\ -1 & : \text{ with probability } \pi, \text{ if } \theta = 0. \end{cases}$$

The principal exaggerates the state if she misreports the truth only by reporting extreme states instead of 0. Increasing π makes the agent more (less) certain that $\theta = m$ upon receiving $m = 1$ ($m = -1$). Let **maximal exaggeration** refer to $\pi \in \{0, 1\}$. Policies induced by exaggerating the state are denoted $\left\{p^E_\pi : \pi \in [0,1]\right\}$.

Let $a_\mu = \mu_1 - \mu_{-1}$ be the agent's optimal action at the prior, and define

$$\begin{aligned} \kappa_1 &:= \tfrac{1}{2} \\ \kappa_2 &:= \frac{2}{\frac{3}{4}\left(\frac{1-|a_\mu|+\mu_0}{1-|a_\mu|}\right)^2+1} \\ \kappa_3 &:= \frac{2}{3\left(\frac{\mu_0}{1-|a_\mu|}\right)^2+1} \\ \kappa_4 &:= 2. \end{aligned} \tag{2.12}$$

Proposition 4. *In the Quadratic Model, an optimal attention outcome:*

1. *fully reveals the state when $\kappa \in (0, \kappa_1]$;*
2. *downplays the state until (2.9) holds with equality when $\kappa \in (\kappa_1, \kappa_2]$;*
3. *maximally exaggerates the state when $\kappa \in [\kappa_2, \kappa_3]$;*[32]

[32] In particular, $\pi = 1$ ($\pi = 0$) is optimal if and only if $a_\mu \geq 0$ ($a_\mu \leq 0$).

4. *exaggerates the state until* (2.9) *holds with equality when* $\kappa \in (\kappa_3, \kappa_4]$*;*

5. *reveals no information when* $\kappa \in (\kappa_4, \infty)$.

Moreover, for generic κ*, every optimal attention outcome is of the above form.*

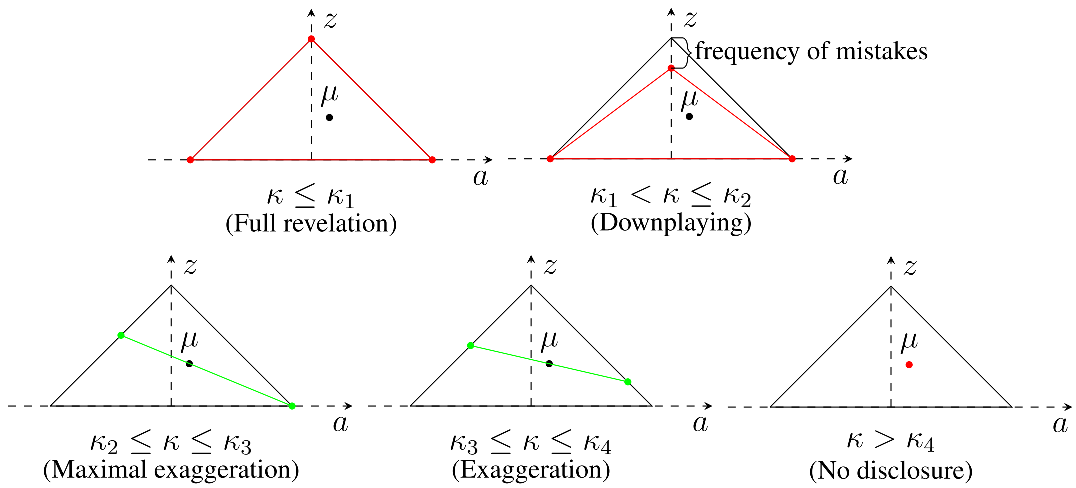

Figure 2.3: Optimal Attention Outcomes

Our result tells the story of a principal who optimally downplays the state when attention is cheap and, as it becomes costlier, reverses her strategy and exaggerates the state. Unless attention is nearly costless, a menu of mistakes becomes available to the principal: infrequent but (equally) large mistakes, a mix of very large and smaller mistakes, small but frequent mistakes, etc. As convex cost tend to discard the second kind, the main tradeoff is between size and frequency.

When attention comes at low cost, the principal keeps quiet (i.e., reports 0) on some extreme occasions and is truthful the rest of the time. That is, she downplays the state. As shown in Figure 2.3, this invites either inaction from the agent, who chooses 0, or an extreme reaction $\{-1, 1\}$. The latter are never mistaken, because they happen in precisely the principal's reported extreme state. However, inaction is harmful in extreme situations, so the agent makes *all* his mistakes through action 0. The frequency of mistakes, shown in Figure 2.3, is chosen so that the agent is barely willing to pay full attention. Under this type of information disclosure, mistakes are large, always of size 1, but they are kept *rare* as long as $\kappa \leq \kappa_2$.

As inattention grows more severe, the principal increasingly downplays the state to keep the agent's attention. In the process, mistakes become increasingly frequent. Eventually, the principal switches to extreme revelations that cause strong reactions: she exaggerates the state. This gets the agent's attention because of his desire to avoid the potentially harmful consequences of extreme behaviors. Although this results in smaller mistakes (than under downplaying, because their size is strictly less than 1), they occur *more* often. In our teacher example of the introduction, this corresponds to teaching students to answer -1 if and only if the correct answer is false, while pooling all uncertain questions with true ones. This eliminates all mistakes concerning false questions, and only leaves fine mistakes involved in discerning true from uncertain questions.

Finally, the following corollary follows naturally from the proposition. In the Quadratic Model, attention management is profitable *whenever* the agent strictly prefers to pay partial attention.

Corollary 1. *In the Quadratic Model, full disclosure is optimal if and only if* $\kappa \leq \kappa_1$ *or* $\kappa \geq \kappa_4$.

The exclusion of perverse garblings is the principal's main tool for managing attention. In Figure 2.4, the agent's attention response to full information (the flat line in green) is excluded from the menu defined by the (red) triangle policy. As a result, the agent pays full attention to that policy and enjoys greater material welfare. Restricting information limits the agent's latitude to respond to it, which, in turn, affects his incentives to listen to the information in the first place. This observation recalls principles from the delegation literature. Szalay[45] shows that restricting an agent's choices can incentivize him to seek useful information, and that this may be valuable even if interests are perfectly aligned conditional on said information. The intuition that eliminating moderate choices will provide greater information acquisition shows up here too: by downplaying the state, the principal makes it more difficult for the agent to hedge. Corollary 1 mirrors Szalay[45, Proposition 4] who shows that choice restrictions are profitable at intermediate costs of information. In our example, with such restrictions arising indirectly through information degradation, the same holds whenever an unrestricted agent would uniquely opt for partial information.

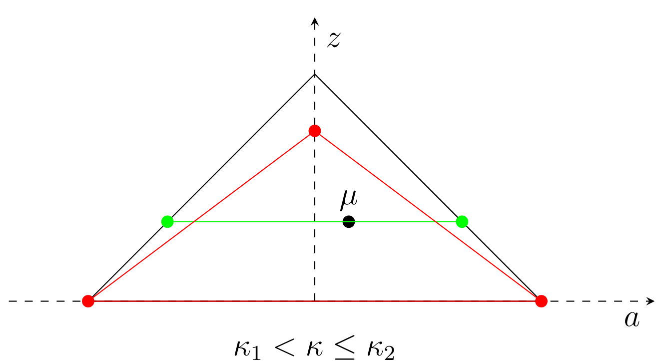

Figure 2.4: Excluding "Safe" Garblings

2.5 Generalizations

While we have focused on aligned interests, our model can also be formulated under conflicting interests. Given any continuous agent and principal objectives $u_A, u_P : A \times \Theta \times \Delta\Theta \to \mathbb{R}$, we can define the agent and principal indirect utilities via

$$U_A(\nu) := \max_{a \in A} \int_{\Delta\Theta} u_A(a, \cdot, \nu)\, \mathrm{d}\nu \tag{2.13}$$

$$U_P(\nu) := \max_{a \in a^*(\nu)} \int_{\Delta\Theta} u_P(a, \cdot, \nu)\, \mathrm{d}\nu, \tag{2.14}$$

where

$$a^*(\nu) := \arg\max_{a \in A} \int_{\Delta\Theta} u_A(a, \cdot, \nu)\, \mathrm{d}\nu.$$

As U_A is continuous and U_P upper semicontinuous, the proof of Theorem 1 then applies verbatim to this generalization.

Theorem 2 holds more generally: the proof applies if either U_P or $U_P - U_A$ is weakly convex. Within our story of paternalistic, benevolent design, a relevant extension here is one for which the principal cares about the agent's material welfare and *partially* internalizes the agent's cost of attention, placing a lower weight on it than does the agent.

The Quadratic Model plainly enjoys a lot of specific structure, but the most important simplification comes in the form of reducing the agent's willingness to pay attention to the study of binary policies. This particular simplification will hold for any model in which U_A is a quadratic function. Indeed, a weaker form of Proposition 3 holds in that case: a nonredundant information policy is incentive compatible if and only if the agent does not have a strict incentive to pool any two supported messages. This observation should facilitate a richer analysis of applied models of attention management.

2.6 Conclusion

We study the design problem of a well-intentioned principal who paternalistically seeks to help an inattentive agent make informed decisions. Even though the principal unequivocally wants the agent to be better informed, we find that withholding information can be optimal, helping guide the agent to make better decisions. A key takeaway from our analysis is that attention management is fundamentally about choosing the right "aspects" of the underlying state on which the agent should focus. We convey this point by showing that unidimensional information should never be withheld and by fully characterizing how to inform an inattentive agent in a simple multidimensional framework.

Chapter 3

Persuasion Under Costly Learning

3.1 Introduction

Learning or information processing can be costly. Due to such costs, a decision maker may not want to take in all available information. This motivated Simon[41, 42] to call for the design of "intelligent information-filtering systems."

The burgeoning literature on persuasion games following Kamenica and Gentzkow[26] studies the question of how to design optimal information-filtering systems. However, existing models almost always assume the information receiver to be a passive learner: he automatically processes whatever information revealed by the sender. Consequently, the design is not driven by learning costs or inattention, but by preference misalignments between the sender and the receiver.

How does inattention affect information disclosure and learning? To address this question, we study the information-filtering problem of a seller aimed at maximizing the purchase probability of a rationally inattentive buyer who finds it costly to learn. In our model, a Sender (seller) reveals information about an uncertain state of world (quality of the good) to persuade a Receiver (buyer) to take a particular action (buying). Processing information is costly for the Receiver, and he can learn strictly less information, in the sense of Blackwell, than what the Sender provides. Taking the Receiver's optimal learning into account, the Sender designs an information disclosure policy aiming to maximize the probability of the Receiver taking such an action. Other applications of the model include a biased advisor persuading a rationally inattentive politician to vote for a reform, a prosecutor persuading a jury to convict a defendant, or a pharmaceutical company convincing FDA to approve a new drug.

We study this problem in a binary-action binary-state environment, where we can explicitly characterize incentive compatible information policies, that is, those to which the buyer willingly pays full attention. To get a flavor of optimal disclosure, suppose that the buyer's optimal action absent any information is not buying. If information is *costless*, optimal disclosure is a recommendation rule that recommends the "buy" action whenever the quality is high and sometimes when the quality is low. The "buy" recommendation is obfuscated in such a way that the buyer finds it indifferent between buying and not buying upon receiving it. If learning is *costly*, however, such a disclosure rule becomes one of the worst because it fails to induce any attention from the buyer. Indeed, utility-wise the buyer does not benefit from paying any attention to such information, and thus he will ignore all of it without buying the good. To attract some attention, the seller now has to obfuscate the "buy" recommendation less often to increase the instrumental value of the available information. The seller does so by recommending the "buy" action less frequently when the quality is low, and as a result, the buyer will *strictly* prefer to buy when such a recommendation is made. Nevertheless, full disclosure (i.e., no obfuscation at all) is also suboptimal. Intuitively, starting from the buyer's best response to full disclosure, obfuscating the "buy" recommendation a little more has a first-order positive effect on the selling probability, while its (negative) effect through incentive compatibility is negligible due to an envelope-theorem argument. Indeed, optimal disclosure still involves information distortion, and in equilibrium the buyer will learn strictly less than what he would under full disclosure.

We also provide a set of comparative statics results with respect to learning costs. While the seller is always worse off when learning costs become higher, the amount of information processed by the buyer in equilibrium varies in a non-monotone fashion. Intuitively, when learning is more costly, in order to attract the buyer's attention, the seller has to increase the instrumental value of information *locally* by making the "buy" recommendation more informative about high quality. This effect tends to reduce her selling probability (as "buy" is recommended less frequently when quality is low) and increase the informativeness of the buyer's learning. On the other hand, a higher learning cost worsens the buyer's incentive to pay attention *globally*, which tends to shrink the set of incentive compatible policies and decrease the amount of information processed by the buyer. Both effects negatively impact the seller's payoff while generating opposing forces on the buyer's equilibrium learning, yielding our results.

The remainder of this chapter is organized as follows. Section 3.2 sets up the model. Section 3.3 presents the main results on optimal disclosure and comparative statics. Section 3.4 provides closed-form solution to an example with quadratic costs and illustrates explicitly the non-monotone variation of the buyer's equilibrium learning. Section 3.5 discusses some extensions beyond the binary-action binary-state case. Section 3.6 concludes. All proofs for the results in this chapter are in Appendix C. Ample discussion of the related literature has been provided in Chapter 2.

3.2 The Model

We build upon the general setting introduced in Chapter 2. Consider a persuasion game between a Sender (seller / prosecutor / pharmaceutical company; she) and a Receiver (buyer / jury / FDA; he). The Receiver must make a decision $a \in A = \{0, 1\}$ in a world with uncertain state $\theta \in \Theta = \{0, 1\}$ distributed according to $\mu \in \Delta\Theta$. To fix idea, we think of the Receiver as a buyer and the Sender as a seller. If $a = 1$ ($a = 0$), the buyer does (not) purchase the good. If $\theta = 1$ ($\theta = 0$), the quality of the good is high (low).

The seller always wants to sell the good regardless of its quality, and her state-independent payoff is given by

$$u_S(a, \theta) = a.$$

On the other hand, if the buyer chooses a in the state θ, his material payoff will be

$$u_R(a, \theta) = a(\theta - \lambda)$$

where λ is the fixed price paid by the buyer when he purchases the good. To make the problem nontrivial, we assume $0 < \mu < \lambda < 1$.[39] The fixed price may come from a company's resale price maintenance, and it is also a relevant assumption in other applications (e.g., prosecutor vs. jury, pharmaceutical company vs. FDA) where monetary transfer is not feasible.[40] In these applications, one can view $1 - \lambda > 0$ and $-\lambda < 0$ as the Receiver's state-dependent payoffs when taking action 1, while the payoff of taking action 0 is a normalized to 0.

In addition to his material utility, the buyer incurs a learning cost. As in the rational inattention literature, this cost is interpreted as the utility loss from processing information. To define it, let

$$\mathcal{R}(\mu) := \left\{ p \in \Delta\Delta\Theta : \int_{\Delta\Theta} \nu \, \mathrm{d}p(\nu) = \mu \right\}$$

be the set of **(information) policies**, which are the distributions over the buyer's beliefs such that the mean equals the prior. We define a learning cost function as a mapping $C : \mathcal{R}(\mu) \to \mathbb{R}_+$ such that for every policy p,

$$C(p) = \kappa \int_{\Delta\Theta} c \, \mathrm{d}p$$

for some strictly convex and differentiable $c : \Delta\Theta \to \mathbb{R}_+$ and some cost parameter $\kappa \in \mathbb{R}_+$. The strict convexity of c ensures that processing more information, in the sense of obtaining a policy p that is more (Blackwell) informative than q, denoted $p \succeq^B_\mu q$, incurs a higher cost.

The timing of the game is as follows:

– The seller first commits to an information policy $p \in \mathcal{R}(\mu)$.

– The buyer then decides to what extent he should pay attention to p: he chooses a policy $q \in \mathcal{R}(\mu)$ such that $q \preceq^B_\mu p$. Such a policy q is called an **(attention) outcome**.

– Finally, the buyer's belief is drawn from q, at which point he takes an action $a \in A$.

The second stage above captures the buyer's freedom to choose how to learn from the seller. Absent learning cost, the buyer will always pay full attention to the information provided by the seller even if he can learn strictly less, because more information always leads to (weakly) better decision-making. When learning is costly, the second stage becomes nontrivial because now the buyer must

[39] If $\mu \geq \lambda$, then the buyer is willing to buy absent any information. Since the seller always wants to sell the good, it is optimal to disclose no information to let the sale happen with probability 1.

[40] If the seller can set the price, she should set it to μ and reveal no information, in which way she will sell the good with probability 1 and extract the entire surplus μ.

make a choice on learning that best balances its cost and benefit. We study principal-preferred subgame perfect equilibria of this game.

Since $|\Theta| = 2$, we let the belief space $\Delta\Theta = [0, 1]$, where $\nu \in \Delta\Theta$ represents the probability of state 1 at that belief. We define $a^* : \Delta\Theta \to A$ by

$$a^*(\nu) = \mathbf{1}_{\{\nu \geq \lambda\}}$$

as the buyer's optimal action at belief ν, and define the players' indirect utility at ν by

$$\begin{aligned} U_S(\nu) &= a^*(\nu) = \mathbf{1}_{\{\nu \geq \lambda\}}, \\ U_R(\nu) &= \int_\Theta u_R(a^*(\nu), \cdot)\, \mathrm{d}\nu - \kappa c(\nu) = \max\{0, \nu - \lambda\} - \kappa c(\nu). \end{aligned}$$

The seller's problem can therefore be formalized as follows:

$$\begin{aligned} &\sup_{p,q} \int_{\Delta\Theta} U_S \, \mathrm{d}q \\ &\quad \text{s.t. } \ p \in \mathcal{R}(\mu) \ \text{ and } \ q \in G^*(p) \end{aligned} \tag{3.1}$$

where

$$G^*(p) := \operatorname*{argmax}_{q \in \mathcal{R}(\mu):\ q \preceq^B p} \int_{\Delta\Theta} U_R \, \mathrm{d}q$$

is the buyer's optimal garbling correspondence. An information policy $p^* \in \mathcal{R}(\mu)$ is **(sender-) optimal** if (p^*, q^*) solves (3.1) for some outcome $q^* \in \mathcal{R}(\mu)$. Such a q^* is an **optimal attention outcome**. In choosing what information to make available, the seller proposes a menu of information policies from which the buyer chooses his most preferred one.

Simplifying Disclosure

Say that an information policy $p \in \mathcal{R}(\mu)$ is **incentive compatible (IC)** if the buyer finds it optimal to pay full attention to it, i.e., if $p \in G^*(p)$. Analogous to the revelation principle, we can restrict attention to incentive compatible policies when searching for an optimal attention outcome.

With convex costs, the buyer will never pay full attention to any policy whose support contains two distinct posterior beliefs that induce the same optimal action. This is because if the buyer garbles these two messages into one, the same action is still optimal at the new belief, and thus this garbling generates the same material payoff as the original policy while reducing the cost of learning. Since the buyer is choosing between only two actions, the support of any IC policy will contain at most two beliefs. We call such policies with two supported beliefs **binary policies**.[41]

The following lemma, which follows directly from Theorem 1 of Chapter 2, formalizes the above reasoning.

[41] δ_μ (the Dirac measure of the prior belief, i.e., no information) is the unique information policy whose support is a singleton. To conserve language, we include it to the set of binary policies.

Lemma 1. *There exists a solution p^* to*

$$\begin{aligned} \sup_p \int_{\Delta\Theta} U_S \, dp & \\ \text{s.t.} \quad & \text{(i)} \quad p \in \mathcal{R}(\mu) \\ & \text{(ii)} \quad p \text{ is IC} \\ & \text{(iii)} \quad p \text{ is binary.} \end{aligned} \tag{3.2}$$

Moreover, p^* is a solution to (3.2) if and only if (r^*, p^*) is a solution to (3.1) for some $r^* \in \mathcal{R}(\mu)$.

A binary policy is uniquely pinned down by its support because there is a unique probability distribution over the support that satisfies Bayesian plausibility. For any binary policy, its set of garblings is straightforward to characterize. In fact, for $p, q \in \mathcal{R}(\mu)$, supp$(p) = \{\nu_1^p, \nu_2^p\}$ and supp$(q) = \{\nu_1^q, \nu_2^q\}$ with $\nu_1^p < \nu_2^p$ and $\nu_1^q < \nu_2^q$,

$$p \succeq^B q \quad \iff \quad \nu_1^p \leq \nu_1^q < \mu < \nu_2^q \leq \nu_2^p. \tag{3.3}$$

3.3 Optimal Disclosure

When learning is costless ($\kappa = 0$), every information policy is IC because the buyer always wants to learn as much as he can for best decision-making. In this case, the following recommendation rule is optimal for the seller:

$$\begin{aligned} \Pr(\text{recommend "buy"}|\theta = 1) &= 1, \\ \Pr(\text{recommend "not buy"}|\theta = 1) &= 0, \\ \Pr(\text{recommend "buy"}|\theta = 0) &= \frac{\mu}{1-\mu}\frac{1-\lambda}{\lambda}, \\ \Pr(\text{recommend "not buy"}|\theta = 0) &= \frac{\mu - \lambda}{\lambda(1-\mu)}. \end{aligned}$$

To maximize her selling probability, the seller adds as much noise as possible to the "buy" recommendation so long as the buyer still wants to follow it. Under this recommendation rule, conditional on full attention, the buyer knows for sure that the quality is low on hearing "not buy," while he is indifferent between buying and not on hearing "buy." It is easy to see that this recommendation rule induces an information policy with support $\{0, \lambda\}$.

When learning is costly ($\kappa > 0$), the above information policy supported on $\{0, \lambda\}$ is no longer IC. This is because the buyer's decision utility $\max\{0, \nu - \lambda\}$ does not change across the convex hull of the two supported beliefs, i.e., $[0, \lambda]$, as he always weakly prefers not to buy. This implies that whatever the buyer learns from the seller, his decision utility does not change, and thus he finds it optimal not to pay any attention. This is also reflected in U_R being strictly concave on $[0, \lambda]$, as shown in Figure 3.1. Hence, if the seller disclosed information as if the buyer had no learning cost, she would never sell the good because the buyer would optimally ignore all information without buying. To attract some attention, the seller has to be less aggressive in sending the "buy" signal, so that the buyer can get some strictly positive utility gain from paying attention. In particular, the buyer should *strictly* prefer to buy when such a recommendation is made.

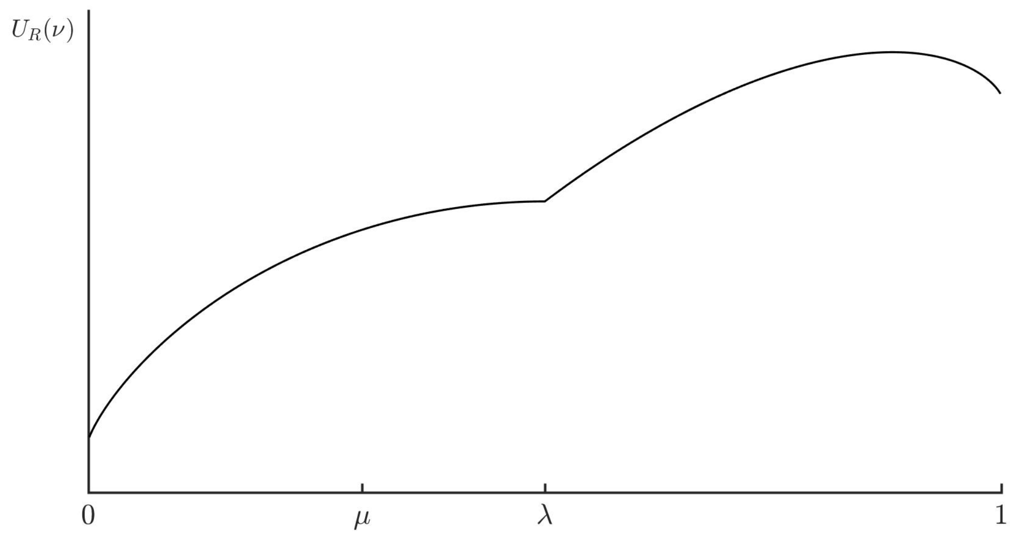

Figure 3.1: Buyer's Indirect Utility Function (With Entropy Learning Costs and $\kappa = 1$)

In what follows, we first analyze the buyer's best information choice q^* when the seller fully discloses the quality, and show that any optimal attention outcome must be less informative than q^*. We then present our characterization of optimal disclosure and its implementation. In the end, we provide some comparative statics results with respect to learning costs.

Full-Information Best Response

The buyer is best off when all information policies are available, that is, when the seller fully discloses the state. Let the full-disclosure policy, $p^F \in \mathcal{R}(\mu)$, be such that $\text{supp}(p^F) = \{0, 1\}$. We denote by q^* a best response to p^F.

Lemma 2. *The buyer has a unique best response to p^F and it is a binary policy.*

The above lemma, proved in Appendix C, is a direct implication of the strict convexity of c. Given this lemma, we let q^* be the buyer's unique best response to p^F and its support be $\{\nu_1^*, \nu_2^*\}$. If $|\text{supp}(q^*)| = 1$ (for example, when learning is very costly), the buyer finds it optimal to learn nothing even when all policies are available. Then, no matter what the seller chooses, the buyer will not learn anything, so the unique optimal attention outcome is no information and every policy is seller-optimal. The rest of the analysis focuses on the interesting case where q^* has binary support.

Assumption 2. *κ is small enough such that $|supp(q^*)| = 2$.*

If q^* has binary support, it is immediate that the recommendation rule described in the benchmark case without learning cost is no longer seller-optimal, because the full-disclosure policy p^F induces the buyer to choose q^* which leads to a strictly positive purchase probability. To find an optimal attention outcome, note first that any policy more informative than q^* will induce q^*, as the latter is the buyer's most preferred policy among *all* information policies and will be chosen whenever available. In Chapter 2, we have explained using Figure 2.1 that, with binary state, any policy Blackwell-incomparable to q^* is not IC. While in the aligned-interests case we can conclude from there that full disclosure is optimal, now with misaligned interests between the two players, we can only rule out those incomparable policies from optimal attention outcomes. The following lemma summarizes this observation.

Lemma 3. *If $\tilde{q}$ is Blackwell-incomparable to q^*, then $\tilde{q}$ is not incentive compatible. Hence, any optimal attention outcome is weakly less informative than q^*.*

Optimal Attention Outcome and Its Implementation

Given a policy p supported on $\{\nu_1, \nu_2\}$ such that $\nu_1 < \mu < \lambda < \nu_2$,[42] the seller's payoff is

$$\int_{\Delta\Theta} U_S \, \mathrm{d}p = p(\nu_2) = \frac{\mu - \nu_1}{\nu_2 - \nu_1}.$$

Note that the seller's payoff is decreasing in both ν_1 and ν_2. The following proposition characterizes optimal attention outcomes.

Proposition 5. *Suppose that Assumption 2 holds. There exists a solution $\{\nu_1^O, \nu_2^O\}$ to*

$$\begin{aligned} \sup_{\nu_1, \nu_2} \; & \frac{\mu - \nu_1}{\nu_2 - \nu_1} \\ \text{s.t.} \quad & \text{(i)} \quad \frac{U_R(\nu_2) - U_R(\nu_1)}{\nu_2 - \nu_1} = U_R'(\nu_1) \\ & \text{(ii)} \quad \nu_1^* \le \nu_1 < \mu < \nu_2 \le \nu_2^* \end{aligned} \tag{3.4}$$

Moreover, a binary policy p^* is an optimal attention outcome if and only if supp(p^*) solves (3.4).

To understand program (3.4) in Proposition 5, recall that Lemma 3 reduces the search to policies less informative than q^*, so that constraint (ii) follows from condition (3.3). Moreover, after the seller chooses a binary policy with support $\{\nu_1, \nu_2\}$, condition (3.3) implies that the buyer is restricted to choose from distributions over beliefs in $[\nu_1, \nu_2]$. As a result, a binary policy is IC if and only if the affine function connecting the points $(\nu_1, U_R(\nu_1))$ and $(\nu_2, U_R(\nu_2))$ lies above the values of U_R restricted to $[\nu_1, \nu_2]$. For a fixed ν_1, the smallest ν_2 such that this property holds is the

[42]If $\nu_1 < \nu_2 \le \lambda$, then p is not IC. See Figure 3.1.

one that makes said affine function tangent to U_R at ν_1 (see Figure 3.2). If ν_2 is lowered further, the buyer no longer wants to pay full attention. Intuitively, ν_2 cannot be too low, for otherwise the policy has too little instrumental value to attract the buyer's full attention. Finally, for any fixed ν_1 we can without loss focus on such lowest ν_2 because the seller's payoff is decreasing in ν_2, leading to constraint (i).

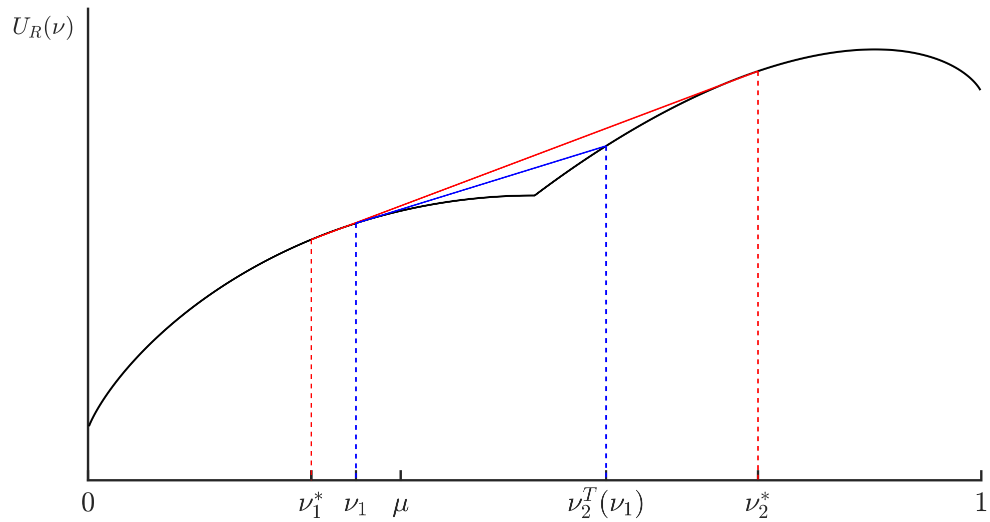

Figure 3.2: Lowest ν_2 Such That $\{\nu_1, \nu_2\}$ Is IC

Recall from Lemma 3 that any optimal attention outcome must be weakly less informative than the full-information best response, q^*. The next proposition shows that it is actually strictly less informative, so that full disclosure is *not* seller-optimal.

Proposition 6. *Suppose that Assumption 2 holds such that $\nu_1^* > 0$ and $\nu_2^* < 1$. Then, any optimal attention outcome p^* is strictly less informative than q^*, i.e., $p^* \prec^B q^*$.*

Remember that, to maximize her selling probability, the seller wants to lower both supported beliefs as much as possible. Nevertheless, to ensure incentive compatibility (captured by constraint (i) in (3.4)), a decrease in the higher supported belief must be associated with an increase in the lower one. Intuitively, when ν_2 is lowered while ν_1 is kept the same, the instrumental value of information can be outweighed by its cost. Consequently, to restore incentive compatibility, ν_1 has be to increased to reduce the cost (by making the policy even less informative).

Starting from q^*, decreasing the higher belief creates two effects on the seller's payoff: a direct positive effect, and an indirect negative effect through the IC constraint. Since q^* is the buyer's unconstrained optimal policy, the first order condition implies that the indirect effect is negligible, leaving only the positive effect present. Thus, the seller always wants to move away from q^* to some policy strictly less informative.

Implementation

When learning is costless, we have seen that the optimal policy can be implemented by a particular type of recommendation rule. Indeed, the seller always recommends "buy" whenever the quality is high, and she does so with some positive probability t when the quality is low. That is,

$$\begin{aligned} \Pr(\text{recommend "buy"}|\theta = 1) &= 1, \\ \Pr(\text{recommend "not buy"}|\theta = 1) &= 0, \\ \Pr(\text{recommend "buy"}|\theta = 0) &= t, \\ \Pr(\text{recommend "not buy"}|\theta = 0) &= 1 - t. \end{aligned}$$

We denote by r_t the information policy implied by such a recommendation rule. By Bayes' rule, supp$(r_t) = \left\{0, \frac{\mu}{\mu+(1-\mu)t}\right\}$. When learning is costly, any optimal attention outcome can also be induced by some recommendation rule in this class, with *partial* attention paid to it by the buyer.[43]

Proposition 7. *Let p^* be an optimal attention outcome such that supp$(p^*) = \{\nu_1^O, \nu_2^O\}$. Then p^* can be induced by r_{t^*} where $t^* = \frac{(1-\nu_2^O)\mu}{\nu_2^O(1-\mu)}$, i.e., $p^* \in G^*(r_{t^*})$. In other words, r_{t^*} is seller-optimal.*

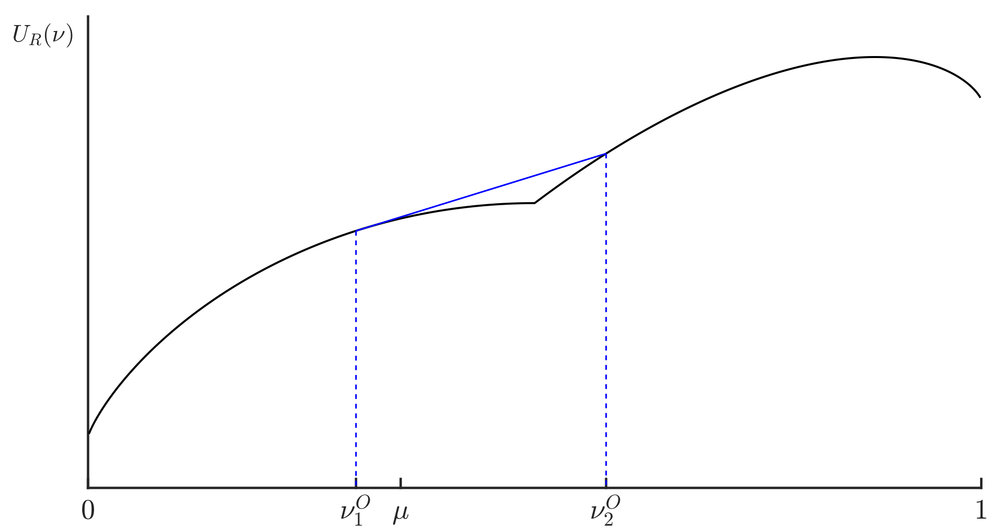

Figure 3.3: Inducing p^* With r_{t^*}

When learning is costly, disclosing information as if there were no learning cost attracts zero attention (thus zero purchase probability), because the available costly information does not help the buyer enough with his decision-making. Thus, the seller has to put more information content

[43]By saying "some policy/recommendation rule p *induces* an attention outcome q," we mean $q \in G^*(p)$.

into the "buy" signal, making the buyer strictly prefer to buy after seeing it (i.e., $\upsilon_2^O > \lambda$). The seller does so by sending the "buy" signal less frequently when the quality is low. This adjustment increases the (ex ante) instrumental value of the available information, making it worthwhile for the buyer to pay *some* attention. As illustrated in Figure 3.3, when r_{t^*} (supported on $\{0, \nu_2^O\}$) is provided, the buyer's unique best response is supported on $\{\nu_1^O, \nu_2^O\}$. In other words, r_{t^*} induces an optimal attention outcome p^*, and hence it is seller-optimal.

Comparative Statics

For general cost functions, we offer a comparative static result about the effect of learning costs (more specifically, κ) on the seller's payoff. We then argue that its effect on the buyer's equilibrium learning is non-monotone. In the next section, we solve an example with quadratic learning costs where we can see these effects explicitly.

Learning costs and seller's payoff

Proposition 8. *The seller is worse off when faced with a less attentive buyer (i.e., with a higher κ).*

With a less attentive buyer, full disclosure will induce less attention (i.e., $q^*(\bar{\kappa}) \preceq^B q^*(\underline{\kappa})$, if $\bar{\kappa} > \underline{\kappa}$). Since only policies less informative than q^* can be IC, this effectively narrows the constraint set of the seller. Moreover, if attention becomes more costly, the instrumental value of a "marginally" IC policy has to be increased for it to remain IC. To achieve so, the seller has to obfuscate the "buy" recommendation less frequently, which necessarily decreases the purchase probability. Both effects (reduction of constraint set, and decrease in purchase probability) lower the seller's payoff.

Learning costs and optimal attention outcome

Proposition 9. *The informativeness of optimal attention outcome $p^*(\kappa)$ is not monotone in κ.*

To see this, note first that when κ is equal (close) to 0, the support of the optimal attention outcome is (close to) $\{0, \lambda\}$. When κ is large enough such that $\nu_1^*(\kappa) > 0$, we denoted by $\{\nu_1^O, \nu_2^O\}$ the support of optimal attention outcome $p^*(\kappa)$. By Proposition 6, we know that $\nu_1^O \geq \nu_1^*(\kappa) > 0$. Moreover, since U_R is strictly concave on $[0, \lambda]$, we must have $\nu_2^O > \lambda$. That is, we have $0 < \upsilon_1^O < \lambda < \upsilon_2^O$, when κ is sufficiently large. Condition (3.3) then implies that $p^*(\kappa)$ is Blackwell-incomparable to $p^*(0)$, and consequently, $p^*(\cdot)$ cannot be monotone (under the Blackwell order).

In the next section, we show in an example with quadratic costs that, as κ increases, p^* first becomes more informative for small κ and then less informative for large κ, and we identify the driving forces of these variations.

3.4 An Example with Quadratic Learning Costs

In this section, we apply the general characterizations in Section 3.3 to a model with quadratic learning costs. Specifically, we assume that $c(\nu) = (\nu - \mu)^2$. Also, for simplicity we assume $\lambda = 1/2$. Under quadratic costs, the farther away a posterior belief is from the prior, the more costly it is to reach that belief. The buyer's indirect utility function becomes

$$U_R(\nu; \kappa) = \max\{\nu - 1/2, 0\} - \kappa(\nu - \mu)^2.$$

We first analyze the buyer's best response $q^*(\kappa)$ to full disclosure. Let $\{\nu_1^*(\kappa), \nu_2^*(\kappa)\}$ be set of points that supports $\text{conv}U_R(\nu; \kappa)$. If $\nu_1^*(\kappa) > 0$ and $\nu_2^*(\kappa) < 1$, the tangency condition of concavification requires that

$$-2\kappa(\nu_1^*(\kappa) - \mu) = \frac{\nu_2^*(\kappa) - 1/2 - \kappa[(\nu_2^*(\kappa) - \mu)^2 - (\nu_1^*(\kappa) - \mu)^2]}{\nu_2^*(\kappa) - \nu_1^*(\kappa)} = 1 - 2\kappa(\nu_2^*(\kappa) - \mu),$$

which gives us

$$\nu_1^*(\kappa) = \frac{1}{2} - \frac{1}{4\kappa},$$
$$\nu_2^*(\kappa) = \frac{1}{2} + \frac{1}{4\kappa}.$$

Therefore, when $\kappa \geq \frac{1}{2(1-2\mu)}$ (so that $\frac{1}{2} - \frac{1}{4\kappa} > \mu$), full disclosure induces no attention. When $1/2 \leq \kappa < \frac{1}{2(1-2\mu)}$, $q^*(\kappa)$ is supported on $\{\nu_1^*(\kappa), \nu_2^*(\kappa)\}$ defined above. When $0 \leq \kappa \leq 1/2$, full disclosure is IC (but not necessarily optimal).

Now we apply Proposition 5 to derive optimal attention outcomes for arbitrary cost parameters.

Proposition 10. *In this example with quadratic costs,*

- *When $0 \leq \kappa < \frac{1-2\mu}{2(1-\mu)^2}$, the unique optimal attention outcome $p^*(\kappa)$ is supported on the set $\left\{0, \frac{1-\sqrt{1-2\kappa}}{2\kappa}\right\}$, and the informativeness of $p^*(\kappa)$ increases in κ.*
- *When $\frac{1-2\mu}{2(1-\mu)^2} \leq \kappa < \frac{1}{2(1-2\mu)}$, the unique optimal attention outcome $p^*(\kappa)$ is supported on the set $\left\{1 - \mu - \sqrt{\frac{1-2\mu}{2\kappa}}, 1 - \mu\right\}$, and the informativeness of $p^*(\kappa)$ decreases in κ.*
- *When $\kappa \geq \frac{1}{2(1-2\mu)}$, the unique optimal attention outcome $p^*(\kappa)$ contains no information.*

Non-monotone variation of the optimal attention outcome with κ

In this example, we can see explicitly that the amount of information processed in equilibrium varies with κ in a non-monotone fashion. For $\kappa < (>) \frac{1-2\mu}{(1-\mu)^2}$, the optimal attention outcome becomes more (less) informative as κ increases. To understand the mechanisms at work, let us start with the optimal attention outcome for some κ, supported on $\{\nu_1^O, \nu_2^O\}$. As attention becomes more costly, it creates two effects. First, fixing ν_1^O, in order to attract the buyer's attention, the seller has to increase the instrumental value of information *locally* by making the "buy" signal more informative about high quality, leading to a higher ν_2 in the support. This effect tends to increase the informativeness of the optimal attention outcome, and it is dominant when κ is small. On the other hand, a higher attention cost reduces the buyer's incentive to pay attention *globally*, making it harder for any policy to attract attention. This effect tends to decrease the informativeness of IC policies, and it is dominant when κ is large.

3.5 Discussion

In this section, we discuss whether our results extend beyond the binary-action binary-state case and the potential difficulties for solving a more general model.

The binary-action assumption simplifies the seller's objective into maximizing the probability of one action. Nevertheless, given the binary-state assumption, a number of our characterizations still hold even with three or more actions. First, by Theorem 1 in Chapter 2, the restriction to binary policies (Lemma 1) is still without loss. Though the buyer's best response to full disclosure may no long be unique, there exists a most informative one (analogous to Lemma 2), and any optimal attention outcome will still be less informative than it (Lemma 3 and Proposition 6).

The binary-state assumption buys us a lot of technical convenience. In particular, since the belief space is unidimensional, there is no *multi-issue concern* in the design of information. In Chapter 2, we explore the design problem with richer uncertainty under aligned decision preferences, and illustrate that optimal disclosure can be hard to find. When decision preferences are misaligned, the multi-issue concern of information disclosure is combined with the sender's incentives to manipulate the receiver's action, which arguably makes the problem even harder. We leave it as an interesting direction for future research.

Nevertheless, given the binary-action assumption, the binary-state assumption is not as restrictive as it looks. Specifically, for an arbitrary state space Θ, let $\Theta_0 = \{\theta \in \Theta : u_R(1,\theta) < u_R(0,\theta)\}$ and $\Theta_1 = \{\theta \in \Theta : u_R(1,\theta) \geq u_R(0,\theta)\}$. The binary-state model in this paper is equivalent to the situation where the seller can only convey information about whether the true state is such that the seller should buy the good ($\theta \in \Theta_1$) or not ($\theta \in \Theta_0$).

3.6 Conclusion

This paper studies the information design problem of a Sender who tries to persuade a rationally inattentive Receiver to take a particular action. Due to learning costs, the Receiver may have an incentive to learn strictly less information than what the Sender provides. In a binary-action binary-state model, we show that the Receiver processes strictly less information than what he would under full disclosure. Moreover, optimal disclosure takes the same form as in the case without learning cost, but in order to attract attention, it has to bring more information content into the positive signal. While the Sender is always worse off when facing a less attentive Receiver, the Receiver's equilibrium learning varies with attention costs in a non-monotone way. As such, this paper sheds light on how to persuade a rationally inattentive decision maker.

Chapter 4

Concluding Remarks

Information frictions are central to many economic problems and phenomena. This dissertation studies various implications of and solutions to such frictions that can arise either *exogenously* due to institutional constraints or *endogenously* due to people's limited information processing capacity.

In Chapter 1, we analyze how trust can be built in a society where information about people's credit histories are not available to creditors, most likely due to the lack of credit agencies. Motivated by the prevalent "starting small" phenomenon, we study a repeated lender-borrower game with anonymous re-matching. The anonymity in re-matching captures the exogenous information frictions brought about by the lack of credit bureaus and the like. Under a reasonable equilibrium refinement, we show that if players are sufficiently patient and re-matching is highly likely, then the loan size is strictly increasing over time along the equilibrium path of all non-trivial equilibria. As such, our analysis illustrates how gradualism can facilitate trust over time in a society without access to credit histories and thus help people partially overcome such extreme information frictions. It also provides an novel explanation to the "starting small" phenomenon in long-term relationships, especially in credit relationships.

In Chapter 2, we turn to scenarios where information frictions arise from people's limited information processing capacity. In particular, we are interested in finding optimal information transmission protocols when the receiver of information is rationally inattentive. To this end, we develop a model where a well-intentioned principal provides information to an agent for whom information is costly to process, but the principal does not internalize this cost. We show that full information is universally optimal if and only if the environment comprises one issue. With multiple issues, attention management becomes optimal: the principal restricts some information to induce the agent to pay attention to other aspects.

In Chapter 3, we naturally extend the above theoretical framework to a different setting where the interests of the sender and the receiver of information are misaligned. In that model, a Sender (seller) tries to persuade a rationally inattentive Receiver (buyer) to take a particular action (e.g., buying). Learning is costly for the Receiver who can choose to process strictly less information than what the sender provides. In a binary-action binary-state model, we characterize optimal disclosure and provide a number of interesting results about the comparative statics with respect to learning costs. As such, this chapter sheds light on how to persuade a rationally inattentive decision maker.

Altogether, these three chapters present a holistic view of my doctoral studies on the implications of and solutions to information frictions in both static and dynamic settings. It is the author's sincere wish to continue investigating similar topics during the professional career and to further advance our knowledge and understanding of these economic issues.

Bibliography

[1] Dilip Abreu. "On the Theory of Infinitely Repeated Games with Discounting". In: *Econometrica* 56.2 (1988), pp. 383–96.

[2] Rui Albuquerque and Hugo A. Hopenhayn. "Optimal Lending Contracts and Firm Dynamics". In: *Review of Economic Studies* 71.2 (2004), pp. 285–315.

[3] Charalambos Aliprantis and Kim D. Border. *Infinite Dimensional Analysis*. Berlin: Springer, 2006.

[4] Ricardo Alonso and Odilon Câmara. "Bayesian Persuasion with Heterogeneous Priors". In: *Journal of Economic Theory* 165.C (2016), pp. 672–706.

[5] Pol Antràs and C. Fritz Foley. "Poultry in Motion: A Study of International Trade Finance Practices". In: *Journal of Political Economy* 123.4 (2015), pp. 853–901.

[6] Miguel Aramendia and Quan Wen. "Justifiable Punishments in Repeated Games". In: *Games and Economic Behavior* 88.C (2014), pp. 16–28.

[7] Luis Araujo, Giordano Mion, and Emanuel Ornelas. "Institutions and Export Dynamics". In: *Journal of International Economics* 98.C (2016), pp. 2–20.

[8] Robert J. Aumann and Michael B. Maschler. *Repeated Games with Incomplete Information*. Cambridge, MA: MIT Press, 1995.

[9] Marco Battaglini. "Multiple Referrals and Multidimensional Cheap Talk". In: *Econometrica* 70.4 (2002), pp. 1379–1401.

[10] Dorian Beauchêne, Jian Li, and Ming Li. "Ambiguous Persuasion". In: *Journal of Economic Theory* 179 (2019), pp. 312–365.

[11] O. Ben-Shahar and C. Schneider. "The Failure of Mandated Disclosure". In: *University of Pennsylvania Law Review* (2011), pp. 647–749.

[12] David Blackwell. "Equivalent Comparison of Experiments". In: *Annals of Mathematics and Statistics* 32 (1953), pp. 265–272.

[13] Alexander Bloedel and Ilya Segal. "Persuasion with Rational Inattention". Working Paper. 2018. URL: https://papers.ssrn.com/abstract_id=3164033.

[14] Andrew Caplin and Mark Dean. "Revealed Preference, Rational Inattention and Costly Information Acquisition". In: *American Economic Review* 105.7 (2015), pp. 2183–2203.

[15] Andrew Caplin and Daniel Martin. “A Testable Theory of Imperfect Perception”. In: *The Economic Journal* 125 (2015), pp. 184–202.

[16] Archishman Chakraborty and Rick Harbaugh. “Comparative Cheap Talk”. In: *Journal of Economic Theory* 132.1 (2007), pp. 70–94.

[17] Srishti Dhar Chatterji. “Martingales of Banach-Valued Random Variables”. In: *Bulletin of the American Mathematical Society* 66.5 (1960), pp. 395–398.

[18] Saikat Datta. *Building Trust*. Theoretical Economics Paper Series 305. STICERD, 1996.

[19] Mathias Dewatripont and Jean Tirole. “Advocates”. In: *Journal of Political Economy* 107.1 (1999), pp. 1–39.

[20] Simone Galperti. “Persuasion: The Art of Changing Worldviews”. In: *American Economic Review* 109.3 (2019), pp. 996–1031.

[21] Matthew Gentzkow and Emir Kamenica. “Costly Persuasion”. In: *American Economic Review Papers and Proceedings* 104.5 (2014), pp. 457–62.

[22] Parikshit Ghosh and Debraj Ray. “Cooperation in Community Interaction Without Information Flows”. In: *Review of Economic Studies* 63.3 (1996), pp. 491–519.

[23] Parikshit Ghosh and Debraj Ray. “Information and Enforcement in Informal Credit Markets”. In: *Economica* 83.329 (2016), pp. 59–90.

[24] David Hirshleifer, Sonya Seongyeon Lim, and Siew Hong Teoh. “Disclosure to an Audience with Limited Attention”. Working Paper. 2004. URL: https://papers.ssrn.com/abstract_id=604142.

[25] Olav Kallenberg. *Foundations of Modern Probability*. New York: Springer, 1997.

[26] Emir Kamenica and Matthew Gentzkow. “Bayesian Persuasion”. In: *American Economic Review* 101.6 (2011), pp. 2590–2615.

[27] Melis Kartal, Wieland Müller, and James Tremewan. “Gradualism in Repeated Games with Hidden Information: An Experimental Study”. Working Paper. 2015. URL: http://statweb.stanford.edu/~kjross/Stat220notes.pdf.

[28] Rachel E Kranton. “The Formation of Cooperative Relationships”. In: *Journal of Law, Economics and Organization* 12.1 (1996), pp. 214–33.

[29] Edward Lazear. “Agency, Earnings Profiles, Productivity, and Hours Restrictions”. In: *American Economic Review* 71.4 (1981), pp. 606–20.

[30] Benjamin Lester, Nicola Persico, and Ludo Visschers. “Information Acquisition and the Exclusion of Evidence in Trials”. In: *Journal of Law, Economics, and Organization* 28.1 (2012), pp. 163–182.

[31] Gilat Levy and Ronny Razin. “On the Limits of Communication in Multidimensional Cheap Talk: A Comment”. In: *Econometrica* 75.3 (2007), pp. 885–893.

[32] Hao Li. “A Theory of Conservatism”. In: *Journal of Political Economy* 109.3 (2001), pp. 617–636.

[33] Elliot Lipnowski and Laurent Mathevet. “Disclosure to a Psychological Audience”. In: *American Economic Journal: Microeconomics* 10.4 (2018), pp. 67–93.

[34] Filip Matějka and Alisdair McKay. “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model”. In: *American Economic Review* 105.1 (2015), pp. 272–298.

[35] Ludmila Matysková. “Bayesian Persuasion With Costly Information Acquisition”. Working Paper. 2018. URL: https://papers.ssrn.com/abstract_id=3161174.

[36] Richard C O’Brien. “On the Openness of the Barycentre Map”. In: *Mathematische Annalen* 223.3 (1976), pp. 207–212.

[37] Petra Persson. “Attention Manipulation and Information Overload”. In: *Behavioural Public Policy* 2.1 (2018), pp. 78–106.

[38] Robert R. Phelps. *Lectures on Choquet’s Theorem*. 2nd Ed. Berlin: Springer, 2001.

[39] James E. Rauch and Joel Watson. “Starting Small in an Unfamiliar Environment”. In: *International Journal of Industrial Organization* 21.7 (2003), pp. 1021–1042.

[40] Debraj Ray. “The Time Structure of Self-Enforcing Agreements”. In: *Econometrica* 70.2 (2002), pp. 547–582.

[41] Herbert A. Simon. “Designing Organizations for an Information Rich World”. In: *Computers, Communications, and the Public Interest*. Ed. by Martin Greenberger. Baltimore, 1971, pp. 37–72.

[42] Herbert A. Simon. *The Sciences of the Artificial*. Cambridge, MA: MIT Press, 1996.

[43] Christopher Sims. “Implications of Rational Inattention”. In: *Journal of Monetary Economics* 50.3 (2003), pp. 665–690.

[44] Christopher Sims. “Stickiness”. In: *Carnegie-Rochester Conference Series on Public Policy* 49.1 (1998), pp. 317–356.

[45] Dezsö Szalay. “The Economics of Clear Advice and Extreme Options”. In: *Review of Economic Studies* 72.4 (2005), pp. 1173–1198.

[46] Jonathan Thomas and Tim Worrall. “Foreign Direct Investment and the Risk of Expropriation”. In: *Review of Economic Studies* 61.1 (1994), pp. 81–108.

[47] Jonathan Thomas and Timothy Worrall. “Self-Enforcing Wage Contracts”. In: *Review of Economic Studies* 55.4 (1988), pp. 541–554.

[48] Joel Watson. “Starting Small and Commitment”. In: *Games and Economic Behavior* 38.1 (2002), pp. 176–199.

[49] Joel Watson. “Starting Small and Renegotiation”. In: *Journal of Economic Theory* 85.1 (1999), pp. 52–90.

[50] Wenhao Wu. “Sequential Bayesian Persuasion”. Working Paper. 2018. URL: https://drive.google.com/file/d/1gyFm7PVCn63-cVW8Ual6aAMfuKx6s9nF/view.

Appendix A

Proofs for Chapter 1

We present the proofs of Propositions 1 and 2 in this Appendix. Throughout Appendix A, we assume that Assumption 1 holds.

Lemma 4. *In any non-trivial social equilibrium, a relationship is never terminated by any player; in addition, in any non-trivial orthodox social equilibrium, there is no default on the equilibrium path.*

Proof. The proof follows the lines of the paragraph preceding Proposition 1. □

Lemma 5. *Let* (l, b) *be an orthodox social equilibrium, and* $\mathbf{L} = \{L_t\}_t$ *be the sequence of loan sizes on its equilibrium path. We have that for all* t*:*

$$\begin{aligned} & L_t = \max L \\ \textit{s.t. } & L \leq L^*, \\ & \Delta(L) \leq \delta[V^B_{t+1}(\mathbf{L}) - \lambda' V^B_0(\mathbf{L})]. \end{aligned} \tag{A.1}$$

where $V^B_t(\mathbf{L}) = \Sigma^{\infty}_{i=0}\delta^i C(L_{t+i})$.

Proof. By Lemma 4, in any orthodox social equilibrium, a relationship is never terminated by any player and there is no default on path; therefore, we can write the borrower's continuation value on path at the beginning of each period as $V^B_t(\mathbf{L}) = \Sigma^{\infty}_{i=0}\delta^i C(L_{t+i})$.

Note also that the borrower's no-default constraint at t can be written as

$$\Delta(L_t) \leq \delta\left[V^B_{t+1}(\mathbf{L}) - \lambda' V^B_0(\mathbf{L})\right].$$

where the LHS is the borrower's current gains from defaulting on L_t at t, and the RHS is the present value of her future cost from defaulting, i.e., the difference between the value of continuing the relationship by repaying and the value of terminating the relationship by defaulting.

Now we prove (A.1) by contradiction. Let (l, b) be an orthodox social equilibrium such that (A.1) does not hold. Let t be the first period that (A.1) fails. Since $\Delta(\cdot)$ is strictly increasing, either $\Delta(L_t) > \delta[V^B_{t+1}(\mathbf{L}) - \lambda' V^B_0(\mathbf{L})]$, or $L_t < L^*$ and $\Delta(L_t) < \delta[V^B_{t+1}(\mathbf{L}) - \lambda' V^B_0(\mathbf{L})]$.

If $\Delta(L_t) > \delta[V_{t+1}^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})]$, since there is no default on the equilibrium path and the no-default constraint is violated, the borrower will be better off by a one-shot deviation to default at time t. This is a contradiction to (l, b) being mutual perfect best responses.

If $L_t < L^*$ and $\Delta(L_t) < \delta[V_{t+1}^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})]$, choose L'_t such that $L_t < L'_t \leq L^*$ and $\Delta(L'_t) < \delta[V_{t+1}^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})]$. Denote by $V_{t+1}^{\prime B}$ and $V_{t+1}^{\prime L}$ of the continuation values for the borrower and the lender at $(t+1)^0$ following history $\{L_1, L_2, ..., L_{t-1}, L'_t\}$ (without default). By the fact that (l, b) is a social equilibrium, at time t with history $\{L_1, L_2, ..., L_{t-1}\}$ (without default), the lender should not find it profitable to offer L'_t. This implies that either L'_t does not induce default and the lender gets weakly worse off, i.e.,

$$R(L'_t) + \delta V_{t+1}^{\prime L} \leq R(L_t) + \delta V_{t+1}^L(\mathbf{L}), \tag{A.2}$$

or L'_t induces default, which means that from the borrower's viewpoint,

$$\Delta(L'_t) \geq \delta[V_{t+1}^{\prime B} - \lambda' V_0^B(\mathbf{L})]. \tag{A.3}$$

Since $L'_t > L_t$ and R is strictly increasing, we know that (A.2) implies $V_{t+1}^{\prime L} < V_{t+1}^L(\mathbf{L})$. Also, since by construction $\Delta(L'_t) < \delta[V_{t+1}^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})]$, we know that (A.3) implies $V_{t+1}^{\prime B} < V_{t+1}^B(\mathbf{L})$. Therefore, either $V_{t+1}^{\prime L} < V_{t+1}^L(\mathbf{L})$ or $V_{t+1}^{\prime B} < V_{t+1}^B(\mathbf{L})$. However, note that the loan sizes in history $\{L_1, L_2, ..., L_{t-1}, L'_t\}$ (without default) and in $\{L_1, L_2, ..., L_{t-1}, L_t\}$ (without default) only differ in the last period with $L'_t > L_t$, thus we have reached a contradiction to part (ii) of Definition 2 of orthodox social equilibrium.[44] □

Lemma 6. *Let (l, b) be an orthodox social equilibrium, and $\mathbf{L} = \{L_t\}_t$ be the sequence of loan sizes on its equilibrium path. If $\{L_t\}_t$ satisfies:*

$$\Delta(L_t) = \delta[V_{t+1}^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})], \textit{ for all } t \geq 0, \tag{A.4}$$

where $V_t^B(\mathbf{L}) = \Sigma_{i=0}^{\infty} \delta^i C(L_{t+i})$, then $\{L_t\}_t$ is strictly monotonic or constant.

Proof. We first show that $\{L_t\}_t$ is either weakly increasing or strictly decreasing. Take any $\{L_t\}_t$ satisfying (A.4). From (A.4), we have:

$$\Delta(L_t) - \Delta(L_{t-1}) = \delta \Sigma_{\tau=t}^{\infty} \delta^{\tau - t}[C(L_{\tau+1}) - C(L_\tau)], \text{ for all } t \geq 1. \tag{A.5}$$

Suppose $\{L_t\}_t$ is not weakly increasing. Then there exists a t such that $L_{t+1} < L_t$. Since Δ and C are strictly increasing, (A.5) implies there must be infinitely many t, such that $L_{t+1} < L_t$. We consider the following 2 cases.

Case 1: There exists a T, such that $L_{t+1} < L_t$ for all $t \geq T$. That is, eventually $\{L_t\}_t$ becomes strictly decreasing. We claim that in this case, $\{L_t\}_t$ must be a strictly decreasing sequence (from

[44] In fact, the contradiction is reached because the values from the rest of the game for both parties following any history (in particular, $\{L_1, L_2, ..., L_{t-1}, L'_t\}$ (without default)) are weakly higher than the remaining values of the current relationship following the same history (since the former includes both the latter and the values from the possibility of re-matching if the relationship is terminated at some date later).

time 0). To see this, notice that since $L_{t+1} < L_t$ for all $t \geq T$, we know from (A.5) that $L_T - L_{T-1} < 0$, i.e., $L_T < L_{T-1}$, because Δ and C are strictly increasing. Apply this step all the way back to $t = 1$ to obtain $L_t < L_{t-1}$ for all $t \geq 1$. So $\{L_t\}_t$ is a strictly decreasing sequence.

Case 2: There does *not* exist a T, such that $L_{t+1} < L_t$ for all $t \geq T$. That is, $\{L_t\}_t$ never becomes strictly decreasing after any T. Now let $t_1 \geq 1$ be the smallest time index, such that $L_{t_1} < L_{t_1-1}$. By the assumption in Case 2, there is a $t_2 \geq t_1$, such that $L_{t_2} < L_{t_2-1}$ but $L_{t_2+1} \geq L_{t_2}$. Now consider L_{t_2}. By (A.5), we have:

$$0 > \Delta(L_{t_2}) - \Delta(L_{t_2-1}) = \delta\{[C(L_{t_2+1}) - C(L_{t_2})] + \delta[C(L_{t_2+2}) - C(L_{t_2+1})] + ...\}, \tag{A.6}$$
$$0 \leq \Delta(L_{t_2+1}) - \Delta(L_{t_2}) = \delta\{[C(L_{t_2+2}) - C(L_{t_2+1})] + \delta[C(L_{t_2+3}) - C(L_{t_2+2})] + ...\}. \tag{A.7}$$

Let:

$$\begin{aligned} K &= (1-\delta)\{[C(L_{t_2+1}) - C(L_{t_2})] + \delta[C(L_{t_2+2}) - C(L_{t_2+1})] + ...\} \\ &= (1-\delta)\Sigma_{\tau=t_2}^{\infty}\delta^{\tau-t_2}[C(L_{\tau+1}) - C(L_\tau)] \\ &< 0, \end{aligned} \tag{A.8}$$

where the inequality follows directly from (A.6). We claim that $(1-\delta)\{[C(L_{t_2+2}) - C(L_{t_2+1})] + \delta[C(L_{t_2+3}) - C(L_{t_2+2})] + ...\} < K < 0$, which is a contradiction to (A.7). To see this, assume that $(1-\delta)\{[C(L_{t_2+2}) - C(L_{t_2+1})] + \delta[C(L_{t_2+3}) - C(L_{t_2+2})] + ...\} \geq K$. Then:

$$\begin{aligned} K &= (1-\delta)\{[C(L_{t_2+1}) - C(L_{t_2})] + \delta[C(L_{t_2+2}) - C(L_{t_2+1})] + ...\} \\ &\geq (1-\delta)\left\{[C(L_{t_2+1}) - C(L_{t_2})] + \frac{\delta K}{1-\delta}\right\} \\ &> (1-\delta)\left(K + \frac{\delta K}{1-\delta}\right) \\ &= K, \end{aligned}$$

where the second line follows from the assumption we just made and the third line follows from $C(L_{t_2+1}) - C(L_{t_2}) \geq 0 > K$. So Case 2 is not possible.

Combining Case 1 and 2, we conclude that any $\{L_t\}_t$ satisfying (A.4) must be either weakly increasing or strictly decreasing.

Now we show that in the case that $\{L_t\}_t$ is weakly increasing, it must be either strictly increasing or constant. Take any $\{L_t\}_t$ satisfying (A.4) and weakly increasing. If it is not strictly increasing, then there exists a smallest t, call it t_3, such that $L_{t_3} = L_{t_3-1}$. According to (A.5), since $\{L_t\}_t$ is weakly increasing, RHS of (A.5) at $t = t_3$ is 0 only if $L_{t+1} = L_t$ for all $t \geq t_3$. Therefore. $\{L_t\}_t$ is constant from $t_3 - 1$. Now we show that $t_3 = 1$. Assume not, i.e., $t_3 \geq 2$. Since $\{L_t\}_t$ is weakly increasing and t_3 is the smallest t, such that $L_t = L_{t-1}$, we have $L_{t_3-1} > L_{t_3-2}$. According to (A.5), we should have:

$$\Delta(L_{t_3-1}) - \Delta(L_{t_3-2}) = \delta\Sigma_{\tau=t_3-1}^{\infty}\delta^{\tau-(t_3-1)}[C(L_{\tau+1}) - C(L_\tau)]. \tag{A.9}$$

Notice that RHS of (A.9) is 0 because we have already proved that $\{L_t\}_t$ is constant from $t_3 - 1$, whereas LHS of (A.9) is positive. So (A.9) cannot hold, a contradiction. Therefore, $t_3 = 1$, which implies that $\{L_t\}_t$ has to be a constant sequence, if it is weakly increasing but not strictly increasing.

Therefore, we conclude that for any $\{L_t\}_t$ satisfying (A.4), it is strictly monotonic or constant. □

Lemma 7. *Let (l, b) be an orthodox social equilibrium, and $\mathbf{L} = \{L_t\}_t$ be the sequence of loan sizes on its equilibrium path. Then $\{L_t\}_t$ must satisfy one and only one of the following four properties:*
(i) it is strictly increasing;
(ii) it is constant;
(iii) it is strictly decreasing;
(iv) it is constant at L^* until some T and then becomes strictly decreasing.

Proof. By Lemma 5, we already know that $\{L_t\}_t$ satisfies (A.1). Consider the following 2 cases:

Case 1: $L_t < L^*$, for all $t \geq 0$, i.e., the 1st constraint in (A.1) never binds, which implies that $\{L_t\}_t$ satisfies (A.4). By Lemma 6, $\{L_t\}_t$ satiesfies (i), (ii) or (iii) in Lemma 7.

Case 2: $L_t = L^*$, for some $t \geq 0$. Note first that if $L_T = L^*$, then $L_{T-1} = L^*$, which implies $L_t = L^*$ for all $t \leq T$. To see this, consider time $T - 1$. The RHS of the 2nd constraint in (A.1) at time $T - 1$ is: $\delta[V_T^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})]$. Notice that:

$$\begin{aligned} V_T^B(\mathbf{L}) &= \Sigma_{i=0}^{\infty} \delta^i C(L_{T+i}) \\ &= C(L_T) + \delta V_{T+1}^B(\mathbf{L}) \\ &= C(L^*) + \delta V_{T+1}^B(\mathbf{L}) \\ &\geq V_{T+1}^B(\mathbf{L}), \end{aligned}$$

where the last line follows from $V_{T+1}^B(\mathbf{L}) \leq \frac{C(L^*)}{1-\delta}$ because C is strictly increasing. Then we have:

$$\delta[V_T^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})] \geq \delta[V_{T+1}^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})] \geq \Delta(L^*). \tag{A.10}$$

(A.10) implies that at time $T - 1$, L^* is the solution to (A.1). Apply this argument all the way back to $t = 0$ to obtain that $L_t = L^*$ for all $t \leq T$.

Now let t_4 be the period such that $L_t = L^*$ for all $t \leq t_4$, and $L_t < L^*$ for all $t \geq t_4 + 1$. If such a t_4 does not exist, then $\{L_t\}_t$ is a constant sequence at L^*, which satisfies (ii) in Lemma 7, so we're done. If t_4 exists, we claim that $L_{t+1} > L_t$, for all $t \geq t_4$.

To see this, notice first that the 2nd constraint in (A.1) is binding for all $t \geq t_4 + 1$ because $L_t < L^*$ for all these t's. Then by Lemma 6, $\{L_t\}_t$ from $t = t_4 + 1$ has to be strictly increasing, or strictly decreasing or constant. We now show that it must be strictly decreasing. Assume not, then it must be strictly increasing or constant from $t_4 + 1$. By definition of $V_t^B(\mathbf{L})$, we have $V_{t_4+2}^B(\mathbf{L}) \geq V_{t_4+1}^B(\mathbf{L})$, which implies:

$$\delta[V_{t_4+2}^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})] \geq \delta[V_{t_4+1}^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})] \geq \Delta(L^*), \tag{A.11}$$

where the last inequality follows from $L_{t_4} = L^*$. But this implies that at time $t_4 + 1$, L^* is the solution to (A.1), a contradiction to $\{L_t\}_t$ being on equilibrium path and $L_{t_4+1} < L^*$. Therefore, $\{L_t\}_t$ has to be strictly decreasing from $t_4 + 1$, meaning that it satisfies (iv) in Lemma 7. □

Proof of Proposition 1. Let $\bar{\alpha} = \sup_{L\in(0,L^*]} \frac{C(L)}{D(L)}$ and $\underline{\alpha} = \inf_{L\in(0,L^*]} \frac{C(L)}{D(L)}$. Under Assumption 1, we have $0 < \underline{\alpha} \leq \bar{\alpha} < 1$. Now define $\delta^* \equiv 1 - \underline{\alpha}$ and $\lambda^*_\delta \equiv \frac{\bar{\alpha}+\delta-1}{\delta}$. We will show that such δ^* and λ^*_δ work.

As we already defined, $\lambda' = \frac{\lambda}{1-(1-\lambda)\delta}$. It can be checked that $\lambda' > \frac{\bar{\alpha}+\delta-1}{\delta\bar{\alpha}}$, iff $\lambda > \frac{\bar{\alpha}+\delta-1}{\delta}$.

Call a sequence $\{L_t\}_t$ an equilibrium loan sequence if it satisfies (A.1) for all t. We first that show any non-trivial equilibrium loan sequence $\{L_t\}_t$ cannot be a constant sequence. Assume it were, i.e., $0 < L_t = \tilde{L} \leq L^*$ for all $t \geq 0$. Then we directly know that $V^B_t(\mathbf{L}) = \frac{C(\tilde{L})}{1-\delta}$ for all t. Now we check the 2nd constraint in (A.1):

$$\begin{aligned}
\delta[V_t^B(\mathbf{L}) - \lambda' V_0^B(\mathbf{L})] &= \delta\left[\frac{C(\tilde{L})}{1-\delta} - \lambda'\frac{C(\tilde{L})}{1-\delta}\right] \\
&< \delta\left(1 - \frac{\bar{\alpha}+\delta-1}{\delta\bar{\alpha}}\right)\frac{C(\tilde{L})}{1-\delta} \\
&= \delta\frac{(1-\bar{\alpha})(1-\delta)}{\delta\bar{\alpha}}\frac{C(\tilde{L})}{1-\delta} \\
&= \frac{(1-\bar{\alpha})}{\bar{\alpha}}C(\tilde{L}) \\
&\leq D(\tilde{L}) - C(\tilde{L}),
\end{aligned}$$

where the second line (strictly inequality) follows from the condition $\lambda' > \frac{\bar{\alpha}+\delta-1}{\delta\bar{\alpha}}$, and the last line follows from $\bar{\alpha} = \sup_{L\in[0,L^*]} \frac{C(L)}{D(L)} \geq \frac{C(\tilde{L})}{D(\tilde{L})}$. This implies that (A.1) is not satisfied (at any t), a contradiction to $\{L_t\}_t$ being an equilibrium loan sequence. Therefore $\{L_t\}_t$ cannot be a constant sequence when $\lambda' > \frac{\bar{\alpha}+\delta-1}{\delta\bar{\alpha}}$.

By Lemma 7 and the result above, we know that $\{L_t\}_t$ is convergent, and (A.4) is eventually satisfied by $\{L_t\}_t$ after some T; that is, there exists a T, such that (A.4) is satisfied for all $t \geq T$.[45] Let $\bar{L}$ be the limit of $\{L_t\}_t$. By definition of $V_t^B(\mathbf{L})$, it converges to $\frac{C(\bar{L})}{1-\delta}$. As (A.4) is eventually satisfied, we must have:

$$D(\bar{L}) - C(\bar{L}) = \delta[\frac{C(\bar{L})}{1-\delta} - \lambda' V_0^B(\mathbf{L})], \tag{A.12}$$

[45]This is because, the only case where (A.4) is not necessarily eventually satisfied (i.e., the 2nd constraint of (A.1) is not eventually binding) is that $\{L_t\}_t$ is constant at L^*, which has just been ruled out.

which gives us:

$$\begin{aligned} V_0^B(\mathbf{L}) &= \frac{\delta C(\bar{L}) - (1-\delta)[D(\bar{L}) - C(\bar{L})]}{(1-\delta)\delta\lambda'} \\ &\leq \frac{\delta - (1-\delta)(\frac{1}{\bar{\alpha}} - 1)}{(1-\delta)\delta\lambda'} C(\bar{L}) \\ &< \frac{C(\bar{L})}{1-\delta}, \end{aligned} \tag{A.13}$$

where the second line follows from $\frac{D(\bar{L})}{C(\bar{L})} \geq \frac{1}{\bar{\alpha}}$, and the third line is obtained by using $\lambda' > \frac{\bar{\alpha}+\delta-1}{\delta\bar{\alpha}}$.

Among the four possible properties in Lemma 7 one of which $\{L_t\}_t$ has to satisfy, only if $\{L_t\}_t$ is strictly increasing will (A.13) hold. Therefore we conclude that when $\lambda' > \frac{\bar{\alpha}+\delta-1}{\delta\bar{\alpha}}$, the loan sizes along the equilibrium path of any non-trivial orthodox social equilibrium must be strictly increasing. □

Proof of Proposition 2. We keep the definition of δ^* and λ^*_δ, i.e., $\delta^* \equiv 1 - \underline{\alpha}$ and $\lambda^*_\delta \equiv \frac{\bar{\alpha}+\delta-1}{\delta}$, where $\bar{\alpha} = \sup_{L\in(0,L^*]} \frac{C(L)}{D(L)}$ and $\underline{\alpha} = \inf_{L\in(0,L^*]} \frac{C(L)}{D(L)}$.

We first establish the existence of an equilibrium loan sequence, i.e., the sequence that satisfies (A.1). Define:

$$\hat{D}(L) = \begin{cases} D(L), & \text{if } L \in [0, L^*]; \\ D(L^*) + (L - L^*), & \text{if } L > L^*; \end{cases}$$

$$\hat{C}(L) = \begin{cases} C(L), & \text{if } L \in [0, L^*]; \\ C(L^*) + \frac{C(L^*)}{D(L^*)}(L - L^*), & \text{if } L > L^*. \end{cases}$$

Notice that by construction, $\hat{D}$ and $\hat{C}$ are continuous extensions of D and C, respectively; both are strictly increasing and unbounded above, satisfying $\inf_{L>0} \frac{\hat{C}(L)}{\hat{D}(L)} = \inf_{L\in[0,L^*]} \frac{C(L)}{D(L)} = \underline{\alpha}$ and $\sup_{L>0} \frac{\hat{C}(L)}{\hat{D}(L)} = \sup_{L\in[0,L^*]} \frac{C(L)}{D(L)} = \bar{\alpha}$.

As implied by Lemma 7, any equilibrium loan sequence converges. Given this property, pick any $\bar{L} \in [0, L^*]$, we can construct as follows a unique equilibrium loan sequence $\{L_t\}_t$ converging to $\bar{L}$:

$$V_0^B = \frac{\frac{C(\bar{L})}{1-\delta} - D(\bar{L})}{\delta\lambda'}, \tag{A.14}$$

$$L_t = \hat{D}^{-1}(V_t^B - \delta\lambda' V_0^B), \tag{A.15}$$

$$V_{t+1}^B = \frac{V_t^B - \hat{C}(L_t)}{\delta}. \tag{A.16}$$

We first show that $\{L_t, V_t^B\}_t$ constructed above is well-defined. It is enough to show that $V_t^B - \delta\lambda' V_0^B > 0$ for all t because then by the fact that $\hat{D}$ is continuous, strictly increasing and unbounded we know L_t is well-defined. To see $V_t^B - \delta\lambda' V_0^B > 0$, note that (A.15) implies that $\{V_t^B\}_t$ satisfies:

$$\hat{D}(L_t) + \delta\lambda' V_0^B = V_t^B.$$

Combining with (A.16) we have:

$$\frac{\hat{D}(L_t)}{\hat{C}(L_t)}(V_t^B - \delta V_{t+1}^B) + \delta\lambda' V_0^B = V_t^B.$$

Rearranging, we get:

$$V_{t+1}^B = \frac{1 - \frac{\hat{C}(L_t)}{\hat{D}(L_t)}}{\delta} V_t^B + \frac{\hat{C}(L_t)}{\hat{D}(L_t)}\lambda' V_0^B. \tag{A.17}$$

By (A.14) and the assumption that $\delta \geq 1 - \inf \frac{C(L)}{D(L)}$, we first have $V_0^B > 0$, so that $V_0^B - \lambda'\delta V_0^B > 0$. Then by induction on (A.17), we have $V_t^B > \delta\lambda' V_0^B$ for all t. Therefore $\{L_t, V_t^B\}_t$ constructed in (A.14) through (A.16) is well-defined.

Now we show that $\{V_t^B\}_t$ is bounded. Consider another sequence $\{\hat{V}_t\}_t$ such that $\hat{V}_0 = V_0^B$ and

$$\hat{V}_{t+1} = \frac{1 - \underline{\alpha}}{\delta}\hat{V}_t + \hat{V}_0. \tag{A.18}$$

Since for all t, $1 > \frac{\hat{C}(L_t)}{\hat{D}(L_t)} \geq \underline{\alpha} \equiv \inf \frac{\hat{C}(L)}{\hat{D}(L)}$, we know from (A.17) and (A.18) that $V_t^B \leq \hat{V}_t$ for all t. Note also that the solution to $\{\hat{V}_t\}_t$ is:

$$\hat{V}_t = \hat{V}_0 \left(\frac{\delta}{\underline{\alpha} + \delta - 1} - \frac{1 - \underline{\alpha}}{\underline{\alpha} + \delta - 1} (\frac{1 - \underline{\alpha}}{\delta})^t \right), \tag{A.19}$$

which is bounded because $\frac{1-\underline{\alpha}}{\delta} < 1$ as assumed. Then we know that $\{V_t^B\}_t$ is bounded above, and because $V_t^B - \delta\lambda' V_0^B > 0$ for all t and $V_0^B > 0$ as we have shown, it is also bounded below. Thus $\{V_t^B\}_t$ is bounded.

Now we claim that $\{L_t, V_t^B\}_t$ satisfies: for all $t \geq 0$,

$$\hat{D}(L_t) - \hat{C}(L_t) = \delta[V_{t+1}^B - \lambda' V_0^B], \tag{A.20}$$

$$V_t^B = \Sigma_{i=0}^{\infty} \delta^i \hat{C}(L_{t+i}). \tag{A.21}$$

It can be checked that (A.20) is obtained by substituting (A.16) into (A.15), and (A.21) is obtained by expanding (A.16) recursively and using the boundedness of $\{V_t^B\}_t$.

By applying Lemma 6 to (A.20) and (A.21),[46] we know that $\{L_t\}_t$ is monotonic. Since $\hat{C}$ are strictly increasing, V_t^B by (A.21) is monotonic. Because $\{V_t^B\}_t$ is also bounded, as we have just shown, $\{V_t^B\}_t$ is convergent. Then by (A.20) again, $\{L_t\}_t$ is also convergent. But then, by construction of V_0^B in (A.14), we have $L_t \to \bar{L}$.

[46]Note that the proof of Lemma 6 only use the conditions that D, C and $D - C$ are strictly increasing, which hold for $\hat{D}$, $\hat{C}$ and $\hat{D} - \hat{C}$ here.

Because $\{L_t\}_t$ can only be strictly increasing, strictly decreasing or constant (by Lemma 6), then using the condition $\lambda > \lambda^*_\delta \equiv \frac{\bar{\alpha}+\delta-1}{\delta}$ and by the same deduction as in (A.13), we know that L_t strictly increases to $\bar{L} \leq L^*$. This in turn implies that $L_t \in [0, L^*]$ for all t, and by definition of $\hat{D}$ and $\hat{C}$ as well as (A.20) and (A.21), we have:

$$\Delta(L_t) = \delta[V^B_{t+1} - \lambda' V^B_0], \tag{A.22}$$

$$V^B_t = \Sigma^{\infty}_{i=0} \delta^i C(L_{t+i}). \tag{A.23}$$

Therefore, $\{L_t\}_t$ satisfies (A.1), meaning that it is an equilibrium loan sequence converging to $\bar{L}$.

Finally we show the existence of an orthodox social equilibrium by construction. Let $\mathbf{L}^* = \{L^*_t\}_t$ be the (unique) equilibrium loan sequence converging to $\bar{L} = L^*$. We construct (l, b) as follows: $l = \{l_0, l_1, ...\}, b = \{b_0, b_1, ...\}$, where for all t, $l_t = (\tilde{L}_{t^0}, \tilde{f}_{t^2}), b_t = (\tilde{d}_{t^1}, \tilde{g}_{t^2})$, in which:

$$\tilde{L}_{t^0}[h(t^0)] = \begin{cases} L^*_t, & \text{if } d_\tau = 0, \text{ for all } \tau < t; \\ 0, & \text{otherwise}; \end{cases} \tag{A.24}$$

$$\tilde{d}_{t^1}[h(t^1)] = \begin{cases} 0, & \text{if } \Delta(L_t) \leq \delta[V^B_{t+1}(\mathbf{L}^*) - \lambda' V^B_0(\mathbf{L}^*)]; \text{ and for all } \tau < t, d_\tau = 0; \\ 1, & \text{otherwise}; \end{cases} \tag{A.25}$$

$$\tilde{f}_{t^2}[h(t^2)] = \begin{cases} 1, & \text{if for all } \tau \leq t, d_\tau = 0; \\ 0, & \text{otherwise}; \end{cases} \tag{A.26}$$

$$\tilde{g}_{t^2}[h(t^2)] = \begin{cases} 1, & \text{if for all } \tau \leq t, d_\tau = 0; \\ 0, & \text{otherwise}. \end{cases} \tag{A.27}$$

Recall that $\tilde{L}_{t^0}$ is the loan offered in period t, $\tilde{d}_{t^1}$ is the defaulting decision in period t, and $\tilde{f}_{t^2}$ and $\tilde{g}_{t^2}$ are continuation decisions in period t. Equations (A.24) through (A.27) describe the following strategy profile.

- For the lender, in period t, as long as default has never happened, she offers L^*_t and continues the relationship, *regardless of whether or not they are on the equilibrium path* (i.e., regardless of previous loan sizes);[47] if default has happened before, she offers 0 loan and terminates the relationship.
- For the borrower, in period t, she will repay as long as the current benefit from defaulting is less than its future cost and default has never happened; otherwise, she will default. The borrower continues the relationship if and only if she never defaulted before.

One particular feature of this construction is that, no matter L_t is on or off equilibrium path, the borrower's perception of her future value is always the one on path, i.e., $V^B_{t+1}(\mathbf{L}^*)$, as long as default never happened. This is a correct perception given the lender's strategy.

[47] In other words, as long as default never occurred, the lender always offers the on-path loan size L^*_t even if the previous loans were off equilibrium path.

Now we show that the constructed (l, b) is an orthodox social equilibrium. By construction of $\tilde{L}_{t^0}$ in (A.24), part (ii) of Definition 2 is satisfied because at the beginning of a given date the continuation values are always the same as long as there is no default. By construction of $\tilde{f}_{t^2}$ in (A.26), part (i) of Definition 2 are satisfied. It remains to be shown that (l, b) is a social equilibrium, i.e., given that the re-match values being V_0^L and V_0^B, (l, b) are sequentially rational with respect to each other. Since this is a game with complete information, it is sufficient to check for one-shot deviation at all possible histories.

First consider $h(t^0)$. At history $h(t^0)$ such that $d_\tau = 0$ for all $\tau < t$, i.e., there is no default before t, if the lender sets $L_t > L_t^*$, we know from (A.25) that this will induce default. But then by (A.26), this will result in the termination of the current relationship. So by setting $L_t > L_t^*$, the lender's payoff changes from $V_t^L(\mathbf{L}^*)$ to $\delta\lambda' V_0^L(\mathbf{L}^*)$, where $V_t^L(\mathbf{L}^*) = \Sigma_{i=0}^{\infty}\delta^i R(L_{t+i})$. As $\{L_t^*\}_t$ is strictly increasing, $V_t^L(\mathbf{L}^*) > V_0^L(\mathbf{L}^*) > \delta\lambda' V_0^L(\mathbf{L}^*)$, so this deviation is not profitable. If the lender sets $L_t < L_t^*$, based on (l, b) the borrower will not default and the loan sequence $\{L_t^*\}_t$ will be restored from next period; so this deviation just lowers the lender's current payoff from $R(L_t^*)$ to $R(L_t)$ while keeping future payoff constant, which is not profitable.

At history $h(t^0)$ such that $d_\tau = 1$ for some $\tau < t$, i.e., there exists default before t but the relationship is not yet terminated, we know from (A.25) and (A.26) that no matter what the lender offers, the borrower will default and the relationship will be terminated at the end of this period. Then offering anything larger than 0 will not be a profitable deviation for the lender.

Now consider $h(t^1)$. At history $h(t^1)$ such that $\Delta(L_t) \leq \delta[V_{t+1}^B(\mathbf{L}^*) - \lambda' V_0^B(\mathbf{L}^*)]$ and $d_\tau = 0$, for all $\tau < t$, a one-shot deviation to default will result in the termination of the current relationship, which is not profitable for the borrower exactly because $\Delta(L_t) \leq \delta[V_{t+1}^B(\mathbf{L}^*) - \lambda' V_0^B(\mathbf{L}^*)]$.

At history $h(t^1)$ such that $\Delta(L_t) > \delta[V_{t+1}^B(\mathbf{L}^*) - \lambda' V_0^B(\mathbf{L}^*)]$ or $d_\tau = 1$ for some $\tau < t$, i.e., either the extra payoff from default is higher than its cost, or there exists default record before t (or both) so that the current relationship will be terminated at the end of this period no matter whether she defaults. In both cases the borrower is better off by defaulting. So there is no profitable deviation.

Finally consider $h(t^2)$. At history $h(t^2)$ such that $d_\tau = 0$ for all $\tau \leq t$, i.e., there is no default at or before t, if the lender one-shot deviates to terminating the relationship, she will get $R(L_t^*) + \delta\lambda' V_0^L(\mathbf{L}^*)$ instead of $R(L_t^*) + \delta V_{t+1}^L(\mathbf{L}^*)$, which is not profitable because $\{V_t^L(\mathbf{L}^*)\}_t$ is increasing in t. Similarly, if the borrower one-shot deviates to terminating the relationship, she will get $C(L_t^*) + \delta\lambda' V_0^B(\mathbf{L}^*)$ instead of $C(L_t^*) + \delta V_{t+1}^B(\mathbf{L}^*)$, which is not profitable because $\{V_t^B(\mathbf{L}^*)\}_t$ is increasing in t.

At history $h(t^2)$ such that $d_\tau = 1$ for some $\tau < t$, i.e., default happens at or before t but the relationship is not yet terminated, if the lender or the borrower one-shot deviates to continuing the relationship, the relationship will still be terminated because the other party will do so according to her equilibrium strategy. So such a one-shot deviation will not change anything, which is not profitable.

Therefore, at no history can we find a profitable one-shot deviation for any player, so (l, b) are sequentially rational with respect to each other. $\square$

Appendix B

Proofs for Chapter 2

Toward the proof of Theorem 1

We first introduce some additional notation

For any $f : X \to \Delta Y, x \in X$ and Borel $B \subseteq Y$, let $f(B|x) := (f(x))(B)$.

Define the barycentre map $\beta_X : \Delta\Delta X \to \Delta X$ by $\beta_X(\hat{X}|m) := \int_{\Delta X} \gamma(\hat{X}) dm(\gamma)$, $\forall m \in \Delta\Delta X$, Borel $\hat{X} \subseteq X$. In other words, $\beta_X(m) = \mathbb{E}_{\nu\sim m}(\nu)$ for all $m \in \Delta\Delta X$. Note that $\mathcal{R}(\mu) = \beta_\Theta^{-1}(\mu)$, by definition.

Define $\Phi : \Delta\Delta\Delta\Theta \to (\Delta\Delta\Theta)^2$ by $\Phi(\mathbb{P}) = \left(\beta_{\Delta\Theta}(\mathbb{P}), \mathbb{P} \circ \beta_\Theta^{-1}\right)$. While we offer no specific interpretation to this map, it will be of use in deriving required properties of the Blackwell order.

Define the garbling correspondence $G : \Delta\Delta\Theta \rightrightarrows \Delta\Delta\Theta$ by

$$G(p) := \left\{q \in \Delta\Delta\Theta : p \succeq^B q\right\}.$$

We can view the principal's problem as a delegation problem in which she offers the agent a delegation set $\hat{G} \in \{G(p)\}_{p\in\mathcal{R}(\mu)}$, and the agent makes a selection $q \in \hat{G}$. Recall, the agent's optimal garbling correspondence $G^* : \Delta\Delta\Theta \rightrightarrows \Delta\Delta\Theta$ is given by

$$G^*(p) := \operatorname*{argmax}_{q\in G(p)} \int_{\Delta\Theta} U_A \, \mathrm{d}q.$$

Claim 1. *β_X is continuous for every compact metrizable space X.*

Proof. This follows from Phelps[38, Proposition 1.1]. □

Claim 2. *Φ is continuous.*

Proof. Suppose $\{\mathbb{P}_n\}_n \subseteq \Delta\Delta\Delta\Theta$ converges to $\mathbb{P}$. Since $\Delta\Theta$ is compact metrizable, $\beta_{\Delta\Theta}(\mathbb{P}_n) \to \beta_{\Delta\Theta}(\mathbb{P})$, by Claim 1. To show $\mathbb{P}_n \circ \beta_\Theta^{-1} \to \mathbb{P} \circ \beta_\Theta^{-1}$, take any continuous and bounded function

$f : \Delta \to \mathbb{R}$. Continuity of β_Θ implies that $f \circ \beta_\Theta$ is continuous. Then,

$$\begin{aligned}\int_{\Delta\Theta} f \, \mathrm{d}\left(\mathbb{P}_n \circ \beta_\Theta^{-1}\right) &= \int_{\Delta\Delta\Theta} f \circ \beta_\Theta \, \mathrm{d}\mathbb{P}_n \\ &\to \int_{\Delta\Delta\Theta} f \circ \beta_\Theta \, \mathrm{d}\mathbb{P} \\ &= \int_{\Delta\Theta} f \, \mathrm{d}\left(\mathbb{P} \circ \beta_\Theta^{-1}\right)\end{aligned}$$

where the second line follows from the weak convergence of $\mathbb{P}_n$ to $\mathbb{P}$. □

Claim 3. *The partial order $\succeq^B$ is given by $\succeq^B = \Phi(\Delta\Delta\Delta\Theta)$.*

Proof. First, take any $p \succeq^B q$ witnessed by mean-preserving spread $r : \Delta\Theta \to \Delta\Delta\Theta$ as in footnote 27. Define $\mathbb{P} := q \circ r^{-1} \in \Delta\Delta\Delta\Theta$. We now show that $\Phi(\mathbb{P}) = (p, q)$. Notice that $\mathcal{R}(\nu) \cap \mathcal{R}(\nu') = \emptyset$ for $\nu \neq \nu'$. Therefore, any $\nu \in \Delta\Theta$ satisfies $\beta_\Theta^{-1}(\nu) \cap r(\Delta\Theta) = r(\nu)$. As a result, for any Borel $S \subseteq \Delta\Theta$,

$$\mathbb{P} \circ \beta_\Theta^{-1}(S) = q \circ r^{-1}\left(\beta_\Theta^{-1}(S)\right) = q \circ r^{-1}(r(S)) = q(S),$$

and

$$\beta_{\Delta\Theta}(S|\mathbb{P}) = \int_{\Delta\Delta\Theta} \tilde{p}(S) \, \mathrm{d}\mathbb{P}(\tilde{p}) = \int_{\Delta\Delta\Theta} \tilde{p}(S) \, \mathrm{d}\left[q \circ r^{-1}\right](\tilde{p}) = \int_{\Delta\Theta} r(S|\tilde{p}) \, \mathrm{d}q(\tilde{p}) = p(S).$$

Therefore, $(p, q) = \Phi(\mathbb{P})$.

Next, take any $\mathbb{P} \in \Delta\Delta\Delta\Theta$ and let $(\bar{p}, \bar{q}) := \Phi(\mathbb{P})$. We want to show that $\bar{p} \succeq^B \bar{q}$. Notice that we can view β_Θ as a $(\Delta\Theta)$-valued random variable on the probability space $(\Delta\Delta\Theta, \mathcal{B}(\Delta\Delta\Theta), \mathbb{P})$. Let $\gamma : \Delta\Delta\Theta \to \Delta\Delta\Theta$ be a conditional expectation $\gamma = \mathbb{E}_{q\sim\mathbb{P}}[q|\beta_\Theta(q)]$, which exists by Chatterji[17, Theorem 1]. So γ is β_Θ-measurable, and $\forall$ Borel $S \subseteq \Delta\Theta$, we have

$$\int_{\Delta\Delta\Theta} q(S) \, \mathrm{d}\mathbb{P}(q) = \int_{\Delta\Delta\Theta} \gamma(S|\cdot) \, \mathrm{d}\mathbb{P}.$$

By Doob's theorem[25, Lemma 1.13], there exists a measurable $r : \Delta\Theta \to \Delta\Delta\Theta$ such that $\gamma = r \circ \beta_\Theta$. Then, $\forall$ Borel $S \subseteq \Delta\Theta$,

$$\begin{aligned}\int_{\Delta\Theta} r(S|\cdot) \, \mathrm{d}\bar{q} = \int_{\Delta\Delta\Theta} (r \circ \beta_\Theta)(S|\cdot) \, \mathrm{d}\mathbb{P} &= \int_{\Delta\Delta\Theta} \gamma(S|\cdot) \, \mathrm{d}\mathbb{P} \\ &= \int_{\Delta\Delta\Theta} q(S) \, \mathrm{d}\mathbb{P}(q) = \beta_{\Delta\Theta}(S|\mathbb{P}) = \bar{p}(S).\end{aligned}$$

Now, that β_Θ is affine and continuous implies

$$\beta_\Theta \circ \gamma = \mathbb{E}\left[\beta_\Theta \circ id_{\Delta\Delta\Theta} | \beta_\Theta\right],$$

which is $\mathbb{P}$-a.s. equal to β_Θ. That is, $\beta_\Theta \circ r \circ \beta_\Theta = id_{\Delta\Theta} \circ \beta_\Theta$, a.s.-$\mathbb{P}$. Equivalently, $\beta_\Theta \circ r = id_{\Delta\Theta}$, a.s.-$\bar{q}$. Then the measurable function

$$\begin{aligned} \bar{r} : \Delta\Theta &\to \Delta\Delta\Theta \\ \nu &\mapsto \begin{cases} r(\nu) & : \ r(\nu) \in \mathcal{R}(\nu) \\ \delta_\nu & : \ r(\nu) \notin \mathcal{R}(\nu) \end{cases} \end{aligned}$$

is $\bar{q}$-a.s. equal to r with $\beta_\Theta \circ \bar{r} = id_{\Delta\Theta}$, so $\bar{r}$ is a mean-preserving spread witnessing $\bar{p} \succeq^B \bar{q}$. □

Claim 4. *$\succeq^B$ is a continuous partial order, i.e., $\succeq^B \subseteq (\Delta\Delta\Theta)^2$ is closed.*

Proof. This follows from Claims 2 and 3, as the continuous image of a compact set is compact. □

Claim 5. *The garbling correspondence G is continuous and nonempty-compact-valued.*

Proof. It is nonempty-valued because $\succeq^B$ is reflexive, and upper hemicontinuous and compact-valued by Claim 4. To show that G is lower hemicontinuous, fix some open $D \subseteq \Delta\Delta\Theta$. Then,

$$\begin{aligned} \{p \in \Delta\Delta\Theta : G(p) \cap D \neq \emptyset\} &= \left\{p \in \Delta\Delta\Theta : \ p \succeq^B q, \ q \in D\right\} \\ &= \{p : \ (p, q) \in \Phi(\Delta\Delta\Delta\Theta), \ q \in D\} \\ &= \Phi_1 \circ \Phi_2^{-1}(D) \\ &= \beta_{\Delta\Theta}\left(\Phi_2^{-1}(D)\right) \end{aligned}$$

where the second line follows from Claim 3, and the last line follows from the definition of Φ_1. By Claim 2, since D is open, so is $\Phi_2^{-1}(D)$. In addition, $\beta_{\Delta\Theta}$ is an open map by O'Brien[36, Corollary 1]. So $\beta_{\Delta\Theta}\left(\Phi_2^{-1}(D)\right)$ is open, implying that G is lower hemicontinuous. □

Claim 6. *The optimal garbling correspondence G^* is upper hemicontinuous and nonempty-compact-valued.*

Proof. As the indirect utility function U_A is (by Berge's theorem) continuous, so is $q \mapsto \int_{\Delta\Theta} U_A \, \mathrm{d}q$. The result then follows from Claim 5 and Berge's theorem. □

Claim 7. *If $q^* \in \mathcal{R}(\mu)$ is such that (q^*, q^*) solves the principal's problem (3.1), then there is a set $\mathcal{P} \subseteq ext\,[\mathcal{R}(\mu)]$ such that $q^* \in \overline{co}\mathcal{P}$ and (p^*, p^*) solves the principal's problem for every $p^* \in \mathcal{P}$.*

Proof. By Choquet's theorem, $\exists \mathbb{Q} \in \Delta\,[\mathcal{R}(\mu)]$ such that:

$$\begin{aligned} \mathbb{Q}\,[\mathrm{ext}\mathcal{R}(\mu)] &= 1, \\ \beta_{\Delta\Theta}(\mathbb{Q}) &= q^*. \end{aligned}$$

By Claim 6 and the Kuratowski-Ryll-Nardzewski Selection Theorem [3, Theorem 18.13, which applies here by Theorem 18.10], there is some measurable selector g of G^*. The random posterior $q_g := \beta_{\Delta\Theta}(\mathbb{Q} \circ g^{-1})$ is then a garbling of q^*. Moreover, that $q^* \in G^*(q^*)$ implies

$$\begin{aligned} 0 &\leq \int_{\Delta\Theta} U_A \, \mathrm{d}q^* - \int_{\Delta\Theta} U_A \, \mathrm{d}q_g \\ &= \int_{\mathrm{ext}\mathcal{R}(\mu)} \left[\int_{\Delta\Theta} U_A \, \mathrm{d}q - \max_{\tilde{q} \in G(q)} \int_{\Delta\Theta} U_A \, \mathrm{d}\tilde{q} \right] \, \mathrm{d}\mathbb{Q}(q). \end{aligned}$$

Since the latter integrand is everywhere nonpositive and the integral is nonnegative, it must be that the integrand is almost everywhere zero. That is, $q \in G^*(q)$ for $\mathbb{Q}$-almost every q. Then, by Claim 6, $q \in G^*(q)$ for every $q \in \mathrm{supp}(\mathbb{Q})$. Therefore, $\mathcal{P} := \mathrm{supp}(\mathbb{Q}) \cap \mathrm{ext}\mathcal{R}(\mu)$ is as desired. □

Claim 8. *There is some $p^* \in ext\,[\mathcal{R}(\mu)]$ such that (p^*, p^*) solves the principal's problem in (2.2).*

Proof. The principal's objective can be formulated as a mapping $\mathrm{Graph}(G^*) \to \mathbb{R}$ with $(p, q) \mapsto \int_{\Delta\Theta} U_P \, \mathrm{d}q$. It is upper semicontinuous and, by Claim 6, has compact domain. Therefore, there is some solution $(\hat{p}, q^*)$ to (2.2). As $G(q^*) \subseteq G(\hat{p})$, it is immediate that $q^* \in G^*(q^*)$; that is, q^* is IC. Letting $\mathcal{P}$ be as delivered by Claim 7, and taking any $p^* \in \mathcal{P}$ completes the claim. □

Claim 9. *Given $|\Theta| < \infty$: $p \in ext\,[\mathcal{R}(\mu)]$ if and only if supp(p) is affinely independent.*[48]

Proof. First, we prove the "only if" direction. Take any $p \in \mathcal{R}(\mu)$. Then $\mu \in \overline{\mathrm{co}}[\mathrm{supp}(p)] = \mathrm{co}[\mathrm{supp}(p)]$, where the equality follows from Θ being finite. By Carathéodory's theorem, there exists an affinely independent $S \subseteq \mathrm{supp}(p)$ such that $\mu \in \mathrm{co}(S)$; without loss, let S be a smallest such set. Since Θ is finite, $S \subset \mathbb{R}^{|\Theta|}$, so affine independence implies that S is finite. Therefore, $\exists N : S \rightrightarrows \Delta\Theta$ such that, $\forall \nu \in S$, the set $N(\nu)$ is a closed convex neighborhood of ν with $S \cap N(\nu) = \{\nu\}$. Making $\{N(\nu)\}_{\nu \in S}$ smaller, we may assume for all selectors η of N, $\{\eta(\nu)\}_{\nu \in S}$ is affinely independent.

Now define a specific selector $\eta : S \to \Delta\Theta$ by:

$$\eta(\nu) = \beta_\Theta \left(\frac{p\,(N(\nu) \cap \cdot)}{p\,(N(\nu))} \right).$$[49]

Since $\mu \in \mathrm{co}(S)$, $\exists w \in \Delta S$ such that $\sum_{\nu \in S} w(\nu)\eta(\nu) = \mu$, and ($S$ being minimal) $w(\nu) > 0$ for all $\nu \in S$. Let

$$q := \sum_{\nu \in S} w(\nu) \frac{p\,(N(\nu) \cap \cdot)}{p\,(N(\nu))}$$

$$\varepsilon := \min_{\nu \in S} \frac{w(\nu)}{p(N(\nu))}$$

[48]Other than this claim, every step in the proof of Theorem 1 works verbatim for the case of a (possibly infinite) compact metrizable state space. Thus, the theorem holds in that more general setting if we replace nonredundancy of p with the condition that $p \in \mathrm{ext}\mathcal{R}(\mu)$.

[49]Note that $p(N(\nu)) > 0$ for every $\nu \in S \subseteq \mathrm{supp}(p)$, so that $\eta(\nu)$ is well-defined. That $N(\nu)$ is closed and convex for every $\nu \in S$ implies η is a selector of N.

Note that $q \in \mathcal{R}(\mu)$. Therefore, $\frac{p-\varepsilon q}{1-\varepsilon} \in \mathcal{R}(\mu)$ and $p \in \text{co}\left\{q, \frac{p-\varepsilon q}{1-\varepsilon}\right\}$.

Now, if $p \in \text{ext}\left[\mathcal{R}(\mu)\right]$, then it must be that $q = p$, even if we make each neighborhood in $\{N(\nu)\}_{\nu \in S}$ smaller, for otherwise $p \in \text{co}\left\{q, \frac{p-\varepsilon q}{1-\varepsilon}\right\}$ contradicts $p \in \text{ext}\left[\mathcal{R}(\mu)\right]$. But then, $\text{supp}(p) = S$, and since S is affinely independent, so is $\text{supp}(p)$.

Now, we prove the "if" direction. Suppose $p \in \mathcal{R}(\mu)$ has affinely independent support S. Suppose $q, q' \in \mathcal{R}(\mu)$ have $p = (1-\lambda)q + \lambda q'$ for some $\lambda \in (0,1)$. Then the support of q must be contained in S. However, q is Bayes-plausible:

$$\sum_{\nu \in S} q(\nu)\nu = \mu = \sum_{\nu \in S} p(\nu)\nu.$$

But S is affinely independent, implying that $q(\nu) = p(\nu)$ for all $\nu \in S$. That is, $q = p$. As q, q', λ were arbitrary, it must be that p is an extreme point. □

Proof of Theorem 1. By Claim 8, a solution to (2.2) exists. By Claims 8 and 9, (2.2) admits some optimal solution, (p^*, p^*), where $\text{supp}(p^*)$ is affinely independent. This implies that $p^* \in G^*(p^*)$. Finally, notice that the optimal value of the problem in (2.3) is no larger than that of (2.2). So (p^*, p^*) is also a solution to (2.3). □

Toward the proof of Theorem 2

We first prove a result that equivalently characterizes the Blackwell order, specialized to the case where the more informative information policy has affinely independent support. This characterization is important in proving both the binary-state result, and later results.

Lemma 8. *Suppose* $|\Theta| < \infty$. $\forall p, q \in \mathcal{R}(\mu)$ *with p nonredundant,*

$$p \succeq^B q \iff \mathit{supp}(q) \subseteq \mathit{co}\left[\mathit{supp}(p)\right].$$

The special case of this lemma with both information policies being finite support is the same as Wu[50, Theorem 5]. We include this slightly more general, nearly identical proof for the sake of completeness.

Proof. Take any $p, q \in \mathcal{R}(\mu)$ with p nonredundant. Since Θ is finite, affine independence implies that $\text{supp}(p)$ is finite.

"If" part: Suppose $\text{supp}(q) \subseteq \text{co}\left[\text{supp}(p)\right]$. Since $\text{supp}(p)$ is affinely independent, we can find a unique $r : \text{supp}(q) \to \Delta\left(\text{supp}(p)\right)$ such that $r(\cdot|\nu_q) \in \mathcal{R}(\nu_q)$ for all $\nu_q \in \text{supp}(q)$; that is,

$$\nu_q = \sum_{\nu_p \in \text{supp}(p)} \nu_p r(\nu_p|\nu_q), \forall \nu_q \in \text{supp}(q). \tag{B.1}$$

Moreover, we have

$$\begin{aligned}\sum_{\nu_p\in\text{supp}(p)} p(\nu_p)\nu_p = \mu \\ &= \int_{\Delta\Theta} \nu_q \,\mathrm{d}q(\nu_q) \\ &= \int_{\Delta\Theta} \left[\sum_{\nu_p\in\text{supp}(p)} r(\nu_p|\nu_q)\nu_p\right] \,\mathrm{d}q(\nu_q) \\ &= \sum_{\nu_p\in\text{supp}(p)} \left[\int_{\Delta\Theta} r(\nu_p|\cdot)\,\mathrm{d}q\right]\nu_p,\end{aligned}$$

where the first two equalities follow from $p, q \in \mathcal{R}(\mu)$, the third equality follows from (B.1), and the last equality comes from changing the order of summation. Since supp(p) is affinely independent, the weights under which the average of the supported beliefs is μ is unique. So

$$p(\nu_p) = \int_{\Delta\Theta} r(\nu_p|\cdot)\,\mathrm{d}q,\ \forall \nu_p \in \text{supp}(p). \tag{B.2}$$

From (B.1) and (B.2), we know that p is a mean-preserving spread of q (witnessed by r), thus $p \succeq^B q$.

"Only if" part: Suppose $p \succeq^B q$ is witnessed by $r : \text{supp}(q) \to \Delta(\text{supp}(p))$ such that (B.1) and (B.2) hold. By (B.1), we directly know that $\nu_q \in \text{co}\,[\text{supp}(p)]\,, \forall \nu_q \in \text{supp}(q)$, thus $\text{supp}(q) \subseteq \text{co}\,[\text{supp}(p)]$, as desired. □

We now consider binary environments.

Proof of Theorem 2. Let $p^F := \mu \circ (\delta_{(\cdot)})^{-1} \in \mathcal{R}(\mu)$, the full disclosure policy. Theorem 1 delivers an optimal IC policy $p^* \in \mathcal{R}(\mu)$ supported on at most two beliefs. The theorem is proved if we find some $q^* \in G^*(p^F)$ with $q^* \succeq^B p^*$. Indeed, convexity of U_P would imply that $\int_{\Delta\Theta} U_P\,\mathrm{d}q^* \geq \int_{\Delta\Theta} U_P\,\mathrm{d}p^*$; and optimality of (p^F, q^*) would then follow from optimality of (p^*, p^*).[50]

If $|\text{supp}(p^*)| = 1$, then any $q^* \in G^*(p^*)$ has $q^* \succeq^B p^*$. Now, focus on the complementary case, $|\text{supp}(p^*)| = 2$. Identifying $\Delta\Theta$ with $[0,1]$, say $\text{supp}(p^*) = \{\nu_0, \nu_1\}$, where $0 \leq \nu_0 < \mu < \nu_1 \leq 1$.

For any $\lambda \in (0,1)$, there is some $\epsilon \in (0,1)$ such that $\epsilon(1-\lambda,\ \lambda) \leq (p^*(\nu_0),\ p^*(\nu_1))$. Therefore, $p_\lambda := p^* - \epsilon\left[(1-\lambda)\delta_{\nu_0} + \lambda\delta_{\nu_1}\right] + \epsilon\delta_{(1-\lambda)\nu_0+\lambda\nu_1} \in \mathcal{R}(\mu)$. As $p^* \in G^*(p^*)$ and $p_\lambda \preceq^B p^*$, we have

$$0 \leq \int_{\Delta\Theta} U_A\,\mathrm{d}p^* - \int_{\Delta\Theta} U_A\,\mathrm{d}p_\lambda = \epsilon\left[(1-\lambda)U_A(\nu_0) + \lambda U_A(\nu_1) - U_A\left((1-\lambda)\nu_0 + \lambda\nu_1\right)\right].$$

[50]Under a different preference specification with U_P not convex, the same conclusion would obtain if $U_P - U_A$ were convex. Then we could deduce that

$$\int_{\Delta\Theta} U_P\,\mathrm{d}(q^* - p^*) \geq \int_{\Delta\Theta} U_A\,\mathrm{d}(q^* - p^*) \geq 0,$$

the first inequality following from Jensen's inequality, and the second following from $G(p^*) \subseteq G(p^F)$.

So, defining

$$\begin{aligned} r : \Delta\Theta &\to \Delta\Delta\Theta \\ \nu &\mapsto \begin{cases} (1-\lambda)\delta_{\nu_0} + \lambda\delta_{\nu_1} & : \ \nu = (1-\lambda)\nu_0 + \lambda\nu_1 \text{ for some } \lambda \in (0,1), \\ \delta_\nu & : \ \text{otherwise,} \end{cases} \end{aligned}$$

r is a mean-preserving spread with $\int_{\Delta\Theta} U_A \, \mathrm{d}r(\cdot|\nu) \geq U_A(\nu) \ \forall \nu \in \Delta\Theta$.

Now, take any $q^F \in G^*(p^F)$, and define $q^* := \int_{\Delta\Theta} r \, \mathrm{d}q^F \in \mathcal{R}(\mu)$. As

$$\int_{\Delta\Theta} U_A \, \mathrm{d}q^* - \int_{\Delta\Theta} U_A \, \mathrm{d}q^F = \int_{\Delta\Theta} \left[\int_{\Delta\Theta} U_A \, \mathrm{d}r(\cdot|\nu) - U_A(\nu) \right] \, \mathrm{d}q^F(\nu) \geq 0$$

and $q^F \in G^*(p^F)$, it follows that $q^* \in G^*(p^F)$ too. Also, by construction, $q^*\left([0,\nu_0] \cup [\nu_1, 1]\right) = 1$, so that $q^* \succeq^B p^*$. The theorem follows. □

Claim 10. *For any $p \in \mathcal{R}(\mu)$, the set $G^*(p)$ admits a $\succeq^B$-maximal element.*

Proof. Fix any $p \in \mathcal{R}(\mu)$. By Claim 6, $G^*(p)$ is nonempty and compact. By Claim 4, $\succeq^B$ is a continuous partial order. Therefore $MAX(G^*(p), \succeq^B) \neq \emptyset$. □

Proofs of results for the Quadratic Model

Lemma 9. *For any $\nu, \nu' \in \Delta\Theta$, $U_A|_{co\{\nu,\nu'\}}$ is convex (concave) if and only if $|\mathbb{E}_{\theta\sim\nu}(\theta) - \mathbb{E}_{\theta\sim\nu'}(\theta)| \geq (\leq)\sqrt{\kappa}||\nu - \nu'||$. Also, for any $\nu, \nu', \tilde{\nu}, \tilde{\nu}' \in \Delta\Theta$ such that $\nu - \nu' = k(\tilde{\nu} - \tilde{\nu}')$ for some $k \neq 0$, then $U_A|_{co\{\nu,\nu'\}}$ is convex (concave) if and only if $U_A|_{co\{\tilde{\nu},\tilde{\nu}'\}}$ is convex (concave).*

Proof. This follows directly from the computation in equation (2.8). □

Proof of Proposition 3. For the Quadratic Model, fixing any nonredundant $p \in \mathcal{R}(\mu)$, we will show that the following are equivalent:

1. p is IC.
2. $U_A|_{\text{co}\{\nu',\nu''\}}$ is weakly convex, $\forall \nu', \nu'' \in supp(p)$.
3. p is order-IC.

(1) ⇒ (2): Suppose condition (2) does not hold. As U_A is quadratic, there then exist $\nu', \nu'' \in \text{supp}(p)$ such that $U_A|_{\text{co}\{\nu',\nu''\}}$ is strictly concave. Define the finite-support random posterior $q \in G(p)$ by

$$q(\nu) = \begin{cases} 0 & : \ \nu \in \{\nu', \nu''\} \\ p(\nu') + p(\nu'') & : \ \nu = \frac{p(\nu')}{p(\nu')+p(\nu'')}\nu' + \frac{p(\nu'')}{p(\nu')+p(\nu'')}\nu'' \\ p(\nu) & : \ \text{otherwise.} \end{cases}$$

In words, q is a random posterior that replaces ν', ν'' in $\text{supp}(p)$ with their conditional mean. By construction, $p \succ^B q$, so that $q \in G(p)$. Also, since $U_A|_{\text{co}\{\nu',\nu''\}}$ is strictly concave, by Jensen's inequality,

$$\int_{\Delta\Theta} U_A \, \mathrm{d}q > \int_{\Delta\Theta} U_A \, \mathrm{d}p.$$

This implies that $p \notin G^*(p)$, i.e., (1) does not hold.

(2) $\Rightarrow$ (1): Suppose (2) holds. By Claim 10, there exists some Blackwell-maximal element $q^* \in G^*(p)$. Since Θ is finite and $\text{supp}(p)$ is affinely independent, we know that $S := \text{supp}(p)$ is finite, and that the map $\beta : \Delta S \to \text{co}(S)$, defined by $p \mapsto \sum_{\nu \in S} p(\nu)\nu$, is bijective.

Assume (toward a contradiction) that $\exists \nu^* \in \text{supp}(q^*) \setminus S$. As β is bijective and $\nu^* \notin S$, we know that $\beta^{-1}(\nu^*)$ is not a point mass. So $\exists \nu', \nu'' \in S$ and $\varepsilon > 0$ such that $\beta^{-1}(\nu'|\nu^*), \beta^{-1}(\nu''|\nu^*) > 2\varepsilon$. By continuity, there exists a neighborhood $N \subseteq \text{co}(S)$ of ν^* such that:

$$\beta^{-1}(\nu'|\nu), \beta^{-1}(\nu''|\nu) > \varepsilon, \forall \nu \in N.$$

Note that $q^*(N) > 0$ as $\nu^* \in \text{supp}(q^*)$.

Define $f_+, f_- : \Delta\Theta \to \Delta\Theta$ by $\nu \mapsto \nu + \varepsilon(\nu' - \nu'')\mathbf{1}_{\nu \in N}$ and $\nu \mapsto \nu - \varepsilon(\nu' - \nu'')\mathbf{1}_{\nu \in N}$, respectively.[51] Also, define

$$q' := \tfrac{1}{2} q^* \circ f_+^{-1} + \tfrac{1}{2} q^* \circ f_-^{-1} \in \Delta\Delta\Theta.$$

By construction, $q' \succ^B q^*$. Also, since $U_A|_{\text{co}\{\nu',\nu''\}}$ is convex, Lemma 9 tells us that $U_A|_{\text{co}\{\nu \pm \varepsilon(\nu'-\nu'')\}}$ is also convex, $\forall \nu \in N$. This implies

$$\int_{\Delta\Theta} U_A \, \mathrm{d}q' \geq \int_{\Delta\Theta} U_A \, \mathrm{d}q^*,$$

so that $q' \in G^*(p)$, contradicting maximality of q^*. Therefore, $\text{supp}(q^*) \subseteq S$, i.e., $q^* \in \Delta S$. But β is bijective and $\beta(q^*) = \mu = \beta(p)$, so that $q^* = p$. Hence, $p \in G^*(p)$, i.e., (1) holds.

(2) $\Rightarrow$ (3): This follows directly from the application of Lemma 9 to any consecutive pair of supported posteriors.

(3) $\Rightarrow$ (2): Suppose (3) holds. By order-IC, we then have $a^1 < \cdots < a^m$. Direct computation then shows that, for any $i, j \in \{1, \ldots, m\}$ with $i < j$:

$$\begin{aligned}
\frac{\nu^j - \nu^i}{a^j - a^i} &= \frac{1}{\sum_{\tilde{\ell}=i}^{j-1}\left(a^{\tilde{\ell}+1} - a^{\tilde{\ell}}\right)} \sum_{\ell=i}^{j-1} \left(\nu^{\ell+1} - \nu^{\ell}\right) \\
&= \sum_{\ell=i}^{j-1} \left[\frac{a^{\ell+1} - a^{\ell}}{\sum_{\tilde{\ell}=i}^{j-1}\left(a^{\tilde{\ell}+1} - a^{\tilde{\ell}}\right)} \left(\frac{\nu^{\ell+1} - \nu^{\ell}}{a^{\ell+1} - a^{\ell}} \right) \right] \\
&\in co \left\{ \frac{\nu^{\ell+1} - \nu^{\ell}}{a^{\ell+1} - a^{\ell}} \right\}_{\ell=i}^{j-1}.
\end{aligned}$$

[51] Since $\nu^* \notin S$, one can choose ε and N small enough to ensure f_+ and f_- are well-defined $\Delta\Theta$-valued maps.

As a norm is convex, and therefore quasiconvex, it follows that

$$\left\|\frac{\nu^j-\nu^i}{a^j-a^i}\right\| \le \max_{\ell\in\{i,\ldots,j-1\}} \left\|\frac{\nu^{\ell+1}-\nu^\ell}{a^{\ell+1}-a^\ell}\right\| \le \frac{1}{\sqrt{\kappa}}.$$

The result then follows from Lemma 9. □

Toward the proof of Proposition 5

Information policies with two supported messages are called **binary policies**, and those with three supported messages are called **ternary policies**. With three states, any nonredundant policy (except "no information") is either binary or ternary.

In what follows, we represent the belief space $\Delta\Theta$ parametrically. Let

$$\mathcal{B} = \left\{(a,z)\in\mathbb{R}^2 :\ z\in[0,1],\ a\in[z-1,1-z]\right\}.$$

For any $(a,z)\in\mathcal{B}$, let

$$\nu_{(a,z)} = z\delta_0 + \frac{1-z+a}{2}\delta_1 + \frac{1-z-a}{2}\delta_{-1}.$$

So $\nu_{(a,z)}\in\Delta\Theta$ has $a=\mathbb{E}_{\theta\sim\nu_{(a,z)}}(\theta)$ and $z=\nu_{(a,z)}(0)$. By construction, the map $(a,z)\mapsto\nu_{(a,z)}$ is an affine bijection between $\mathcal{B}$ and $\Delta\Theta$.

Under this representation, the Dirac measures in $\Delta\Theta$ are the extreme points of $\mathcal{B}$: $\delta_{-1}=\nu_{(-1,0)}$, $\delta_0=\nu_{(0,1)}$, and $\delta_1=\nu_{(1,0)}$. Also, the prior satisfies $\mu=\nu_{(a_\mu,\mu_0)}$, where $a_\mu=\mu_1-\mu_{-1}$ and $(a_\mu,\mu_0)\in\text{int}(\mathcal{B})$. Without loss of generality, *we focus hereafter on the case that* $a_\mu\ge 0$. Figure B.1 depicts the (a,z) representation of the belief space.

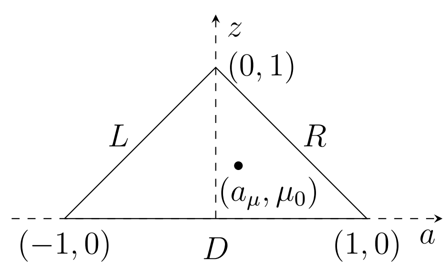

Figure B.1: (a,z)-Representation of $\Delta\Theta$

For any distinct beliefs $\nu=\nu_{(a,z)}$ and $\nu'=\nu_{(a',z')}$, let $(\Delta a,\Delta z)=(a-a',z-z')$. The derivation in the main text leading to (2.11) shows that

$$\begin{aligned} &|a-a'|\ge\sqrt{\kappa}||\nu-\nu'|| \\ \iff &\kappa\le 2,\ \Delta a\ne 0,\ \text{and}\ \left|\frac{\Delta z}{\Delta a}\right|\le\sqrt{\frac{2-\kappa}{3\kappa}}=:s^*(\kappa) \end{aligned} \tag{B.3}$$

By Proposition 3, a nonredundant $p \in \mathcal{R}(\mu)$ is IC, if and only if (B.3) holds for any two consecutive messages in supp(p).

Finally, we give the analytical expressions for the line segments $L := \text{co}\left(\{(0,1),(-1,0)\}\right)$, $R := \text{co}\left(\{(0,1),(1,0)\}\right)$ and $D := \text{co}\left(\{(-1,0),(1,0)\}\right)$.

$$L = \{(a,z):\ a \in [-1,0],\ z = 1+a\} \tag{B.4}$$
$$R = \{(a,z):\ a \in [0,1],\ z = 1-a\} \tag{B.5}$$
$$D = \{(a,z):\ a \in [-1,1],\ z = 0\} \tag{B.6}$$

Claim 11. *Suppose $p \in \mathcal{R}(\mu) \setminus \{\delta_\mu\}$ is nonredundant and IC, with support $\{\nu_{(a^n,z^n)}\}_{n=1}^m$ such that $a^1 < \cdots < a^m$.*[52] *If p is principal-optimal, then neither (a^1,z^1) nor (a^m,z^m) lies in int$(\mathcal{B})$.*

Proof. Suppose that $(a^1,z^1) \in \text{int}(\mathcal{B})$. For sufficiently small $\epsilon > 0$, interiority of (a^1,z^1) implies that

$$(\tilde{a},\tilde{z}) := (a^2,z^2) + (1+\epsilon)(a^1-a^2, z^1-z^2) \in \mathcal{B},$$

and $\{\nu_{(\tilde{a},\tilde{z})},\ \nu_{(a^2,z^2)},\ldots,\nu_{(a^m,z^m)}\}$ is still affinely independent. There is then a mean-preserving spread $\tilde{p}$ of p which is supported on $\{\nu_{(\tilde{a},\tilde{z})},\ \nu_{(a^2,z^2)},\ldots,\nu_{(a^m,z^m)}\}$.[53] As $\frac{z^2-\tilde{z}}{a^2-\tilde{a}} = \frac{z^2-z^1}{a^2-a^1}$ by construction, Proposition 3 tells us that $\tilde{p}$ is still IC. But the action distribution induced by $\tilde{p}$ is a strict mean-preserving spread of that induced by p. By Observation 1, p cannot be principal-optimal.

A symmetric argument proves that p is suboptimal if $(a^m,z^m) \in \text{int}(\mathcal{B})$ □

Below, we rule out ternary policies such that (when parametrized in $\mathcal{B}$) the middle message lies below the line segment between the two other messages.

Claim 12. *Take any nonredundant ternary policy $p \in \mathcal{R}(\mu)$ with support $\left\{\nu_{(a_1,z_1)},\nu_{(a_2,z_2)},\nu_{(a_3,z_3)}\right\}$ such that $a_1 < a_2 < a_3$. If p is IC and*

$$z_2 \le \tfrac{a_3-a_2}{a_3-a_1} z_1 + \tfrac{a_2-a_1}{a_3-a_1} z_3,$$

then p is not principal-optimal.

Proof. First, let us assume that $z_1 \le z_3$. Nonredundancy of p rules out the possibility that $z_2 = \frac{a_3-a_2}{a_3-a_1} z_1 + \frac{a_2-a_1}{a_3-a_1} z_3$; therefore,

$$z_2 < \tfrac{a_3-a_2}{a_3-a_1} z_1 + \tfrac{a_2-a_1}{a_3-a_1} z_3 \le z_3.$$

Let $s_\ell := \frac{z_2-z_1}{a_2-a_1}$ and $s_r := \frac{z_3-z_2}{a_3-a_2}$. Observe that

$$s_r - s_\ell = \tfrac{1}{a_3-a_2} z_3 + \tfrac{1}{a_2-a_1} z_1 - \left(\tfrac{1}{a_2-a_1} + \tfrac{1}{a_3-a_2}\right) z_2 = \tfrac{a_3-a_1}{(a_2-a_1)(a_3-a_2)} \left[\tfrac{a_3-a_2}{a_3-a_1} z_1 + \tfrac{a_2-a_1}{a_3-a_1} z_3 - z_2\right] > 0.$$

[52] Nonredundancy implies that $m \in \{2,3\}$.

[53] To construct it, take a dilation on $\Delta\Theta$ which fixes each of $\nu_{(a^2,z^2)},\ldots,\nu_{(a^m,z^m)}$ and splits $\nu_{(a^1,z^1)}$ to a measure with support $\{\nu_{(\tilde{a},\tilde{z})},\ \nu_{(a^2,z^2)}\}$.

This implies $s_r > 0$, because

$$(a_2 - a_1)s_\ell + (a_3 - a_2)s_r = z_3 - z_1 \geq 0.$$

Letting $p_i := p\left(\nu_{(a_i,z_i)}\right) \in (0,1)$ for $i = 1,2,3$, consider a small $\epsilon > 0$, and define the three vectors,

$$\begin{aligned}(\tilde{a}_1, \tilde{z}_1) &= (a_1, z_1)\\ (\tilde{a}_2, \tilde{z}_2) &= (a_2, z_2 + p_3\epsilon)\\ (\tilde{a}_3, \tilde{z}_3) &= (a_3, z_3 - p_2\epsilon)\,.\end{aligned}$$

All statements in what follows are taken to mean, *when $\epsilon > 0$ is sufficiently small*.

Applying Jensen's inequality to the concave function $1 - |\cdot|$,

$$z_2 < \tfrac{a_3-a_2}{a_3-a_1}z_1 + \tfrac{a_2-a_1}{a_3-a_1}z_3 \leq \tfrac{a_3-a_2}{a_3-a_1}(1-|a_1|) + \tfrac{a_2-a_1}{a_3-a_1}(1-|a_3|) \leq 1 - |a_2|.$$

Combining this with $z_3 > z_2 \geq 0$ tells us that $(\tilde{a}_2, \tilde{z}_2), (\tilde{a}_3, \tilde{z}_3) \in \text{int}(\mathcal{B})$. Next, define $\tilde{p} \in \Delta\Delta\Theta$ to be the measure which puts mass p_i on $\nu_{(\tilde{a}_i,\tilde{z}_i)}$ for $i = 1,2,3$. Direct computation shows that $\tilde{p} \in \mathcal{R}(\mu)$ and that $\tilde{p}$ generates the same action distribution as p. Observation 1 then tells us that $\tilde{p}$ yields the same principal payoff as does p. Now, to show that $\tilde{p}$ is IC, observe that $\frac{\tilde{z}_3-\tilde{z}_2}{\tilde{a}_3-\tilde{a}_2} = \frac{\tilde{z}_3-\tilde{z}_2}{a_3-a_2} \in [0, s_r]$ and $\frac{\tilde{z}_2-\tilde{z}_1}{\tilde{a}_2-\tilde{a}_1} = \frac{\tilde{z}_2-\tilde{z}_1}{a_2-a_1} \in [s_\ell, s_r]$. So $\tilde{p}$ is order-IC if p is. By Proposition 3, $\tilde{p}$ is IC because p is.

Since the principal is indifferent between p and $\tilde{p}$, $\tilde{p}$ is principal-optimal if p is. But Claim 11 tells us that (since $(\tilde{a}_3, \tilde{z}_3) \in \text{int}(\mathcal{B})$) $\tilde{p}$ is not principal-optimal, and so p is not principal-optimal.

A symmetric argument proves the same result in the case $z_1 \geq z_3$. ☐

Claim 13. *Suppose a nonredundant $p \in \mathcal{R}(\mu)$ is IC. If there exists $\nu_{(a,z)} \in$ supp(p) such that $(a,z) \in$ int(D), then p is not principal-optimal.*

Proof. Say p has support $\{\nu_{(a_1,z_1)}, \ldots, \nu_{(a_m,z_m)}\}$ such that $a_1 < \cdots < a_m$ and $m = 1,2,3$. First, let us assume $z_1 \leq z_m$.

If $m = 1$, the result follows from $\mu_0 > 0$, so assume otherwise.

If $m = 3$ and $z_2 \leq \frac{a_3-a_2}{a_3-a_1}z_1 + \frac{a_2-a_1}{a_3-a_1}z_3$, the result follows from Claim 12, so assume otherwise.

So we focus on the case that $m \in \{2,3\}$; $z_1 = 0$; and if $m = 3$, then $z_2 > \frac{a_2-a_1}{a_3-a_1}z_3$.

Letting $p_i := p\left(\nu_{(a_i,z_i)}\right) \in (0,1)$ for $i \in \{1,\ldots,m\}$, consider a small $\epsilon > 0$ and, define,

$$(\tilde{a}_i, \tilde{z}_i) = \begin{cases} (a_1, z_1 + p_2\epsilon) & i = 1\\ (a_2, z_2 - p_1\epsilon) & i = 2\\ (a_i, z_i) & \text{otherwise.}\end{cases}$$

All statements in what follows are taken to mean, *when $\epsilon > 0$ is sufficiently small*.

That $z_2 > 0$ and $(a_1, z_1) \in \text{int}(D)$ imply $(\tilde{a}_1, \tilde{z}_1), (\tilde{a}_2, \tilde{z}_2) \in \text{int}(\mathcal{B})$. Defining $\tilde{p} \in \Delta\Delta\Theta$ to be the measure which puts mass p_i on $\nu_{(\tilde{a}_i,\tilde{z}_i)}$ for $i = 1,\ldots,m$, observe that $\tilde{p} \in \mathcal{R}(\mu)$ and that $\tilde{p}$

generates exactly the same action distribution as p. Just as in the proof of Claim 12, an appeal to Claim 11 and Proposition 3 means we need only show that $\tilde{p}$ is order-IC.

But see that $\frac{\tilde{z}_2-\tilde{z}_1}{\tilde{a}_2-\tilde{a}_1} = \frac{z_2-z_1-(p_1+p_2)\epsilon}{a_2-a_1} \in \left[0, \frac{z_2-z_1}{a_2-a_1}\right]$. Moreover, in the ternary case, $\frac{z_2-z_1}{a_2-a_1} < \frac{z_3-z_2}{a_3-a_2}$ because $z_2 > \frac{a_2-a_1}{a_3-a_1}z_3$, so that $\frac{\tilde{z}_3-\tilde{z}_2}{\tilde{a}_3-\tilde{a}_2} = \frac{z_3-z_2+p_2\epsilon}{a_3-a_2} \in \left[\frac{z_2-z_1}{a_2-a_1}, \frac{z_3-z_2}{a_3-a_2}\right]$. The result follows.

A symmetric argument proves the same result in the case $z_1 \geq z_3$. □

Lemma 10. *Take any nonredundant binary policy $p \in \mathcal{R}(\mu)$ with supp$(p) = \left\{\nu_{(a_1,z_1)}, \nu_{(a_2,z_2)}\right\}$ such that $a_1 < a_2$. Suppose (p, p) is a solution to (2.2) (i.e., p is an optimal attention outcome). Then,*

1. $(a_1, z_1) \in L$, $(a_2, z_2) \in R$;
2. $\left|\frac{z_2-z_1}{a_2-a_1}\right| \leq s^*(\kappa)$

Proof. Claims 11 and 13 imply that $(a_1, z_1), (a_2, z_2) \in L \cup R$; that $p \in \mathcal{R}(\nu_{(a_\mu,\mu_0)})$ and $(a_\mu, \mu_0) \in$ int$(\mathcal{B})$ then imply part (1). Part (2) follows from Proposition 3 and equation (B.3). □

Lemma 11. *Suppose $\frac{1}{2} < \kappa < 2$. Take any nonredundant ternary policy $p \in \mathcal{R}(\mu)$ with supp$(p) = \left\{\nu_{(a_1,z_1)}, \nu_{(a_2,z_2)}, \nu_{(a_3,z_3)}\right\}$ such that $a_1 < a_2 < a_3$. Suppose (p, p) is a solution to (2.2) (i.e., p is an optimal attention outcome). Then,*

1. $(a_2, z_2) \in$ *int*$(\mathcal{B})$, $(a_1, z_1) \in L$, $(a_3, z_3) \in R$;
2. $\frac{z_2-z_1}{a_2-a_1} = s^*(\kappa)$, $\frac{z_3-z_2}{a_3-a_2} = -s^*(\kappa)$.

Proof. Since p is nonredundant and IC, Claims 11 and 13 imply that $(a_1, z_1), (a_3, z_3) \in L \cup R$. As $s^*(\kappa) < 1$ when $\kappa > \frac{1}{2}$, it cannot be (given Proposition 3) that L contains two distinct beliefs in the support of p; and similarly for R. This implies that $(a_1, z_1) \in L$, $(a_3, z_3) \in R$, and $(a_2, z_2) \in \mathcal{B} \setminus (L \cup R)$. But Claim 13 then rules out $(a_2, z_2) \in D$ as well. This delivers part (1).

Now we prove part (2). To start, it follows from Proposition 3 and equation (B.3) that we have $\left|\frac{z_2-z_1}{a_2-a_1}\right|, \left|\frac{z_3-z_2}{a_3-a_2}\right| \leq s^*(\kappa)$.

We now establish that $\left|\frac{z_2-z_1}{a_2-a_1}\right| = \left|\frac{z_3-z_2}{a_3-a_2}\right| = s^*(\kappa)$. Assume (toward a contradiction) that this does not hold. Let us first assume $\left|\frac{z_2-z_1}{a_2-a_1}\right| < s^*(\kappa)$. Consider a small $\epsilon > 0$, and define

$$\begin{aligned}
(\tilde{a}_1, \tilde{z}_1) &= (a_1, z_1) \\
(\tilde{a}_2, \tilde{z}_2) &= (1+\epsilon)(a_2, z_2) - \epsilon(a_3, z_3) \\
(\tilde{a}_3, \tilde{z}_3) &= (a_3, z_3).
\end{aligned}$$

Taking $\epsilon > 0$ to be sufficiently small, $\{(\tilde{a}_i, \tilde{z}_i)\}_{i=1}^3$ will be three affinely independent vectors, all in $\mathcal{B}$, whose convex hull contains $(0, \mu_0)$; and $\left|\frac{\tilde{z}_2-\tilde{z}_1}{\tilde{a}_2-\tilde{a}_1}\right| < s^*(\kappa)$. There is therefore a unique $\tilde{p} \in \mathcal{R}(\mu)$

whose support is $\{(\tilde{a}_i, \tilde{z}_i)\}_{i=1}^3$, and $\tilde{p}$ is order-IC, since $\left|\frac{\tilde{z}_3-\tilde{z}_2}{\tilde{a}_3-\tilde{a}_2}\right| = \left|\frac{z_3-z_2}{a_3-a_2}\right| \le s^*(\kappa)$. Proposition 3 guarantees that $\tilde{p}$ is IC. But, by construction, the action distribution induced by $\tilde{p}$ is a strict mean-preserving spread of that induced by p. Therefore, by Observation 1, p is not optimal, a contradiction.

A symmetric argument derives the same contradiction in the case $\left|\frac{z_3-z_2}{a_3-a_2}\right| < s^*(\kappa)$. Therefore, we have $\left|\frac{z_2-z_1}{a_2-a_1}\right| = \left|\frac{z_3-z_2}{a_3-a_2}\right| = s^*(\kappa)$.

Finally, nonredundancy implies that $\frac{z_2-z_1}{a_2-a_1} \ne \frac{z_3-z_2}{a_3-a_2}$, and Claim 12 rules out the possibility that $\frac{z_2-z_1}{a_2-a_1} < 0 < \frac{z_3-z_2}{a_3-a_2}$. Part (2) then follows. □

We define three special information policies that will be used in the coming proofs.

- The "full disclosure" policy p^F is such that $\text{supp}(p^F) = \{\nu_{(-1,0)}, \nu_{(0,1)}, \nu_{(1,0)}\}$
- The "no information" policy p^N is such that $\text{supp}(p^N) = \{\mu\}$
- The orthogonal policy p^O is such that $\text{supp}(p^O) = \{\nu_{(\mu_0-1,\mu_0)}, \nu_{(1-\mu_0,\mu_0)}\}$

The above lemmas can be combined into the following claim, which reduces our search for optimal attention outcomes to a single two-dimensional problem.

Claim 14. *Suppose $\frac{1}{2} < \kappa < 2$, and suppose the nonredundant policy $p \in \mathcal{R}(\mu)$ is an optimal attention outcome. Then there exist $(a_1, z_1), (a_2, z_2), (a_3, z_3) \in \mathcal{B}$ such that:*

- *$(a_1, z_1) \in L$, $(a_3, z_3) \in R$, and $a_1 \le a_2 \le a_3$;*
- *$z_2 - z_1 = s^*(\kappa)(a_2 - a_1)$, $z_3 - z_2 = -s^*(\kappa)(a_3 - a_2)$;*
- *$p\{\nu_{(a_1,z_1)}, \nu_{(a_2,z_2)}, \nu_{(a_3,z_3)}\} = 1$.*

Proof. As $\kappa \le 2$, the policy p^O is IC, generating strictly higher principal payoffs than no information. Therefore, p must be informative. Being nonredundant, it is either ternary or binary.

In the case that p is ternary, Lemma 11 delivers the result.

In the remaining case, p has binary support $\{\nu_{(a_1,z_1)}, \nu_{(a_3,z_3)}\}$, where $(a_1, z_1), (a_3, z_3) \in \mathcal{B}$ with $a_1 < a_3$. Lemma 10 then tells us that $(a_1, z_1) \in L$, $(a_3, z_3) \in R$, and $\left|\frac{z_3-z_1}{a_3-a_1}\right| \le s$, where $s := s^*(\kappa)$. Moreover, $s \in (0, 1)$ because $\frac{1}{2} < \kappa < 2$.

Now, the line of slope s through (a_1, z_1) and the line of slope $-s$ through (a_3, z_3) have different slopes (since $s > 0$), so that there they have a unique intersection point $(a_2, z_2) \in \mathbb{R}^2$. To prove the claim, then, all that remains is to show is that $a_1 \le a_2 \le a_3$ and $(a_2, z_2) \in \mathcal{B}$.

Define, for $\tilde{a}_2 \in \mathbb{R}$, the gap $g(\tilde{a}_2) = [z_3 - s(\tilde{a}_2 - a_3)] - [z_1 + s(\tilde{a}_2 - a_1)]$, which strictly decreases in $\tilde{a}_2$, so that $g^{-1}(0) = \{a_2\}$. Since $\left|\frac{z_3-z_1}{a_3-a_1}\right| \le s$, $g(a_3) \le 0 \le g(a_1)$. The intermediate value theorem then says that $a_2 \in [a_1, a_3]$.

As $s \leq 1$, $a_2 \geq a_1$, and (a_1, z_1) lies (weakly) below the line containing L, it follows that (a_2, z_2) lies below that line as well. As $s \leq 1$, $a_2 \leq a_3$, and (a_3, z_3) lies (weakly) below the line containing R, it follows that (a_2, z_2) lies below that line as well. As $s \geq 0$, $a_2 \geq a_1$, and (a_1, z_1) lies (weakly) above the line containing D, it follows that (a_2, z_2) lies above that line as well. Therefore $(a_2, z_2) \in \mathcal{B}$. $\square$

Lemma 12. *If $\frac{1}{2} < \kappa < 2$, then there exists an optimal attention outcome p such that one of the following holds*

- *p is binary: supp$(p) = \{\nu_{(a_1,z_1)},\ \nu_{(a_2,z_2)}\}$ for some $(a_1, z_1) \in L,\ (a_2, z_2) \in R$.*
- *p is ternary with critical slopes and support containing $\nu_{(1,0)}$:*[54]
  supp$(p) = \{\nu_{(a_1,z_1)},\ \nu_{(a_2,z_2)},\ \nu_{(1,0)}\}$ for some $(a_1, z_1) \in L,\ (a_2, z_2) \in$ int$(\mathcal{B})$ with $\frac{z_2 - z_1}{a_2 - a_1} = s = \frac{z_2}{1 - a_2}$.

Moreover, policies of the second form exist if and only if $s^*(\kappa) > \frac{\mu_0}{1 - a_\mu}$.

Proof. Here, $s := s^*(\kappa) \in (0, 1)$ because $\frac{1}{2} < \kappa < 2$.

Define the set $\mathcal{T} \subseteq \mathbb{R}^6$ of tuples $((a_i, z_i))_{i=1}^3 \subseteq B^3$ such that:

- $(a_1, z_1) \in L$, $(a_3, z_3) \in R$, and $a_1 \leq a_2 \leq a_3$;
- $z_2 - z_1 = s(a_2 - a_1)$, $z_3 - z_2 = -s(a_3 - a_2)$;
- $(a_\mu, \mu_0) \in \text{co}\{(a_1, z_1),\ (a_2, z_2),\ (a_3, z_3)\}$.

As $s > 0$, it follows that every element $t \in \mathcal{T}$ has $\{t_1, t_2, t_3\}$ an affinely independent subset of $\mathcal{B}$.[55] Given that $(a_\mu, \mu_0) \in \text{co}\{t_1, t_2, t_3\}$, it follows that there is a unique $p_t^{\mathcal{T}} \in \mathcal{R}(\mu) \cap \Delta\{\nu_{t_1}, \nu_{t_2}, \nu_{t_3}\}$. The policy $p_t^{\mathcal{T}}$ is IC by Proposition 3. Moreover, Claim 14 tells us that *any* nonredundant optimal attention outcome is of this form. Accordingly, we can reformulate the principal's problem as $\max_{t \in \mathcal{T}} \int_{\Delta\Theta} U_P \, \mathrm{d}p_t^{\mathcal{T}}$.

Toward a parametrization of $\mathcal{T}$, let $\mathcal{A} := \{(a_1, a_2) :\ t = (a_1, z_1, a_2, z_2, a_3, z_3) \in \mathcal{T}\}$.

Given $t = (a_1, z_1, a_2, z_2, a_3, z_3) \in \mathcal{T}$, we can infer the following:

- That $(a_1, z_1) \in L$ implies $z_1 = 1 + a_1$.
- Then, $z_2 = z_1 + s(a_2 - a_1) = 1 + (1 - s)a_1 + sa_2$.
- Finally, (a_3, z_3) belongs both to R (so that $z_3 = 1 - a_3$) and to the line of slope $-s$ through (a_2, z_2)—uniquely pinning it down as these two lines have different slopes. Direct computation then shows that $(a_3, z_3) = (-a_1 - \frac{2s}{1-s}a_2,\ 1 + a_1 + \frac{2s}{1-s}a_2)$.

[54] Recall our normalization that $a_\mu \geq 0$.

[55] In the extreme case where $a_2 = a_1$ (or $a_2 = a_3$), (a_2, z_2) and (a_1, z_1) (or (a_2, z_2) and (a_3, z_3)) collapse to one point, so that $\{t_1, t_3\}$ is still an affinely independent subset of $\mathcal{B}$.

So, for any $(a_1, a_2) \in \mathbb{R}^2$, let

$$t(a_1, a_2) := (a_1,\ 1 + a_1,\ a_2,\ 1 + (1-s)a_1 + sa_2,\ -a_1 - \tfrac{2s}{1-s}a_2,\ 1 + a_1 + \tfrac{2s}{1-s}a_2) \in \mathbb{R}^6.$$

Consistent with the previous notation, let

$$\begin{aligned} t_1(a_1, a_2) &= (a_1,\ 1 + a_1), \\ t_2(a_1, a_2) &= (a_2, 1 + (1-s)a_1 + sa_2), \\ t_3(a_1, a_2) &= (-a_1 - \tfrac{2s}{1-s}a_2,\ 1 + a_1 + \tfrac{2s}{1-s}a_2). \end{aligned}$$

The above derivations show that $\mathcal{T} = \{t(a_1, a_2)\}_{(a_1,a_2)\in\mathcal{A}}$. Hence, we can view the principal's problem as $\max_{(a_1,a_2)\in\mathcal{A}} \int_{\Delta\Theta} U_P \, \mathrm{d}p^{\mathcal{T}}_{t(a_1,a_2)}$.

Now we consider two cases separately.

- If $(a_1, a_2) \in \mathcal{A}$ and $t_2(a_1, a_2) \neq t_3(a_1, a_2)$ (i.e., $a_2 \neq -\frac{1-s}{1+s}a_1$), then

  $$a_1 (1-s) + a_2 (1+s) \neq 0.$$

  Moreover, $a_1 < a_2$.[56] Therefore, given $(a_1, a_2) \in \mathcal{A}$, the following three numbers are all well-defined:

  $$\begin{aligned} p_2(a_1, a_2) &:= \frac{a_1(a_1 + 1 - \mu_0)(1-s) + a_2(2a_1 + 1 - \mu_0 - a_\mu)s}{-s(a_2 - a_1)\left[a_1 (1-s) + a_2 (1+s)\right]}, \\ p_3(a_1, a_2) &:= \frac{(1-s)\left[sa_\mu + (1-s)\, a_1 + (1-\mu_0)\right]}{-2s\left[a_1 (1-s) + a_2 (1+s)\right]}, \\ p_1(a_1, a_2) &:= 1 - p_2(a) - p_3(a). \end{aligned}$$

  Observe that $p_1(a) + p_2(a) + p_3(a) = 1$ and $\sum_{i=1}^3 p_i(a)(a_i, z_i) = (a_\mu, \mu_0)$. The affine independence property of $\{t_i(a)\}_{i=1}^3$ then tells us that $p^{\mathcal{T}}_{t(a)} = \sum_{i=1}^3 p_i(a)\delta_{\nu_{t_i(a)}}$.

  Now, the principal objective can be expressed (by direct, tedious computation) as:[57]

  $$\begin{aligned} \int_{\Delta\Theta} U_P \, \mathrm{d}p^{\mathcal{T}}_{t(a_1,a_2)} + \mathbb{V}_{\theta\sim\mu}(\theta) + a_\mu^2 &= \mathbb{E}_{\nu\sim p^{\mathcal{T}}_{t(a_1,a_2)}}[a^*(\nu)^2] \\ &= p_1(a)a_1^2 + p_2(a)a_2^2 + p_3(a)\left(-a_1 - \tfrac{2s}{1-s}a_2\right)^2 \\ &= a_1^2 - \tfrac{a_1^2 + (1-\mu_0)a_1}{s} - \left(2a_1 + \tfrac{1-\mu_0}{1-s} + \tfrac{2s-1}{1-s}a_\mu\right) a_2. \quad \text{(B.7)} \end{aligned}$$

[56] This is because, if $a_1 = a_2$, then $(a_1, z_1) = (a_2, z_2) = (0, 1)$. But then, the fact that $(a_3, z_3) \in R$ implies that $(a_\mu, \mu_0) \in R$, contradicting $(a_\mu, \mu_0) \in \text{int}(\mathcal{B})$.

[57] The "$\mathbb{V}_{\theta\sim\mu}(\theta) + a_\mu^2$" term is an irrelevant constant, which we add for convenience. The first equality follows directly from Observation 1.

- If $(a_1, a_2) \in \mathcal{A}$ and $t_2(a_1, a_2) = t_3(a_1, a_2)$ (i.e., $a_2 = -\frac{1-s}{1+s}a_1$), that $(a_\mu, \mu_0) \in \text{co}\{t_1(a), t_2(a)\}$, $t_1(a) \in L$ and $t_2(a) \in R$ imply that

$$\begin{aligned}(a_1, a_2) &= \left(-\frac{1-\mu_0+sa_\mu}{1-s}, \frac{1-\mu_0+sa_\mu}{1+s}\right),\\ p^{\mathcal{T}}_{t(a_1,a_2)}(\nu_{t_1}) &= \frac{1-s}{2}\frac{1-\mu_0-a_\mu}{1-\mu_0+sa_\mu},\\ p^{\mathcal{T}}_{t(a_1,a_2)}(\nu_{t_2}) &= \frac{1+s}{2}\frac{1-\mu_0+a_\mu}{1-\mu_0+sa_\mu}.\end{aligned}$$

The value of the principal's objective is now

$$\begin{aligned}\int_{\Delta\Theta} U_P \, \mathrm{d}p^{\mathcal{T}}_{t\left(-\frac{1-\mu_0+sa_\mu}{1-s}, \frac{1-\mu_0+sa_\mu}{1+s}\right)} + \mathbb{V}_{\theta\sim\mu}(\theta) + a_\mu^2 &= p_1\left(-\tfrac{1-\mu_0+sa_\mu}{1-s}\right)^2 + p_2\left(\tfrac{1-\mu_0+sa_\mu}{1+s}\right)^2\\ &= \frac{(1-\mu_0)^2 - s^2a_\mu^2}{1-s^2},\end{aligned}$$

which one can directly verify is consistent with the value of equation (B.7) at $(a_1, a_2) = \left(-\frac{1-\mu_0+sa_\mu}{1-s}, \frac{1-\mu_0+sa_\mu}{1+s}\right)$.

Therefore, equation (B.7) summarizes the principal's payoff for all $(a_1, a_2) \in \mathcal{A}$. Observe that this objective is affine in a_2. But, $t(\cdot)$ being continuous and $\mathcal{T}$ being compact, the set of a_2 such that $(a_1, a_2) \in \mathcal{A}$ is (for fixed a_1) a compact set of real numbers. We may therefore find a principal-optimal attention outcome by restricting attention to the case that a_2 is the largest or smallest possible number for which $(a_1, a_2) \in \mathcal{A}$. Letting $\mathcal{A}^* \subseteq \mathcal{A}$ be the set of pairs with this property, we can view the principal's problem as

$$\max_{(a_1,a_2)\in\mathcal{A}^*} \int_{\Delta\Theta} U_P \, \mathrm{d}p^{\mathcal{T}}_{t(a_1,a_2)}.$$

What does $p^{\mathcal{T}}_{t(a)}$ look like if $(a_1, a_2) \in \mathcal{A}^*$? As $t(\cdot)$ is continuous, any $(a_1, z_1, a_2, z_2, a_3, z_3) \in \mathcal{T}$ with $(a_3, z_3) \in \text{int}(R)$, $(a_2, z_2) \in \text{int}\,(\mathcal{B})$ and $\mu_0 \in \text{int}\,(\text{co}\{(a_i, z_i)\}_{i=1}^3)$ *cannot* have $(a_1, a_2) \in \mathcal{A}^*$. The reason is that, from the definition of $\mathcal{T}$, it would then contain $t(a_1, a_2 \pm \epsilon)$ for sufficiently small ϵ. So, if $(a_1, a_2) \in \mathcal{A}^*$, then $(a_3, z_3) = (1, 0)$,[58] or (a_2, z_2) belongs to the boundary of $\mathcal{B}$, or μ_0 belongs to the boundary of $\text{co}\{(a_i, z_i)\}_{i=1}^3$.

If $(a_3, z_3) = (1, 0)$, then we have established part (ii) of this lemma.

If μ_0 belongs to the boundary of $\text{co}\{(a_i, z_i)\}_{i=1}^3$, it must be that $p_i(a_1, a_2) = 0$ for some i, then $p^{\mathcal{T}}_{t(a)}$ has binary support.

If (a_2, z_2) belongs to the boundary of $\mathcal{B}$, since $a_1 < a_2$ and $0 < s < 1$, it cannot be that $a_2 \in L \cup D$. Since $-s \neq -1$, it can only be that $(a_2, z_2) \in R$ if $a_2 = a_3 = -a_1 - \frac{2s}{1-s}a_2$, i.e., if $a_2 = -\frac{1-s}{1+s}a_1$. But then $p^{\mathcal{T}}_{t(a)}$ also has binary support.

[58] It cannot be that $(a_3, z_3) = (0, 1)$ because $a_\mu \geq 0$ and $(a_\mu, \mu_0) \in \text{int}(\mathcal{B})$.

So any $(a_1, a_2) \in \mathcal{A}^*$ has $p^{\mathcal{T}}_{t(a)}$ either binary, or ternary with $\nu_{(1,0)}$ in the support (and critical slopes), and the result follows.

Finally, for any $t \in \mathbb{R}^6$ such that $(a_3, z_3) = (1, 0)$, $(a_1, z_1) \in L$ and $\frac{z_2 - z_1}{a_2 - a_1} = s = \frac{z_2}{1 - a_2}$, we have $(a_\mu, \mu_0) \in \text{int}\left[\text{co}\{(a_1, z_1),\ (a_2, z_2),\ (a_3, z_3)\}\right]$ if and only if $s > \frac{\mu_0}{1 - a_\mu}$. So if $s \leq \frac{\mu_0}{1 - a_\mu}$, either $t \notin \mathcal{T}$ or $p^{\mathcal{T}}_t$ is binary; while if $s > \frac{\mu_0}{1 - a_\mu}$, $p^{\mathcal{T}}_t$ is of the second form in this lemma. (See Figures B.2 and B.3 for illustration.) The last part of Lemma 12 follows. $\square$

Graphically, when $s > \frac{\mu_0}{1 - a_\mu}$, for any given a_1, varying a_2 leads to policies depicted in Figure B.2 (following the notation of Lemma 12's proof). On the other hand, when $s \leq \frac{\mu_0}{1 - a_\mu}$, any $(a_1, a_2) \in \mathcal{A}^*$ leads to a binary policy, as depicted in Figure B.3.

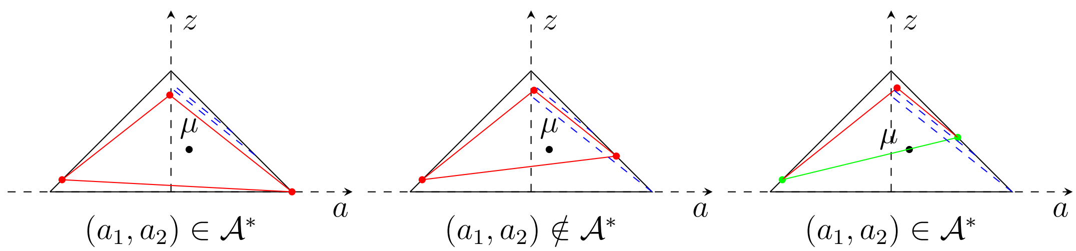

Figure B.2: When $s > \frac{\mu_0}{1 - a_\mu}$, for a Fixed a_1, Varying a_2 Such That $(a_1, a_2) \in \mathcal{A}$

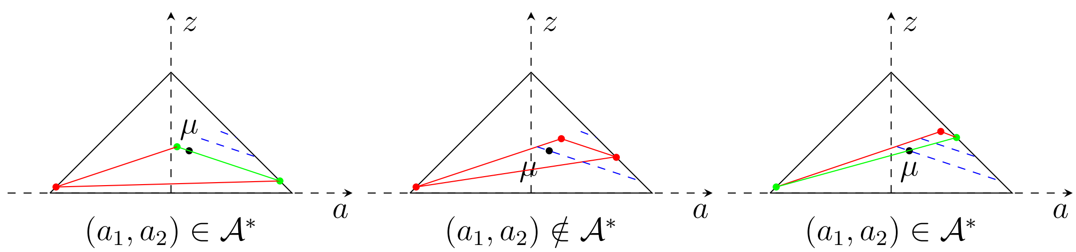

Figure B.3: When $s \leq \frac{\mu_0}{1 - a_\mu}$, for a Fixed a_1, Varying a_2 Such That $(a_1, a_2) \in \mathcal{A}$

Proof of Proposition 5. We will find a nonredundant optimal attention outcome, which exists by Theorem 1. Note that a nonredundant p is an optimal attention outcome, if and only if it is IC and principal-optimal, if and only if (p, p) is a solution to (2.2).

If $\kappa \leq \kappa_1 = \frac{1}{2}$, then $s^*(\kappa) \geq 1$. Note that the support $\text{supp}(p^F)$ of the full disclosure policy is $\{\nu_{(-1,0)}, \nu_{(0,1)}, \nu_{(1,0)}\}$. So by Proposition 3 and equations (2.11), (B.4), (B.5), p^F is IC. Since $p^F \succeq^B q$ for all $q \in \mathcal{R}(\mu)$ and $U_P(\nu)$ is convex, by Jensen's inequality $\left(p^F, p^F\right)$ is a solution to (2.2), so p^F is an optimal attention outcome, and part (1) of Proposition 5 follows.

If $\kappa > \kappa_4 = 2$, by equation (2.11) any nonredundant information policy with more than one messages is *not* IC. Therefore, the "no information" policy p^N is the only nonredundant IC policy, and so is the solution to (2.3). Then by Theorem 1, $\left(p^N, p^N\right)$ is a solution to (2.2), so p^N is an optimal attention outcome, and part (5) of Proposition 5 follows.

If $\kappa = \kappa_4 = 2$, then $s^*(\kappa) = 0$. Given equation (2.11), then, any nonredundant $p \in \mathcal{R}(\mu)$ must have $p\{\nu_{(a,z)} : \ (a,z) \in B, \ z = \mu_0\}$. But p^O is IC too, and $p^O \succeq^B p$ for such p. As U_P is convex, it follows that p^O is an optimal attention outcome. For this value of κ, the content of part (4) of Proposition 5 follows.

Henceforth, we consider the remaining case that $\kappa \in (\kappa_1, \kappa_4) = \left(\frac{1}{2}, 2\right)$, collectively covering parts (2), (3), and (4)—excluding the special case of $\kappa = \kappa_4$—of the proposition. Here, $s := s^*(\kappa) \in (0,1)$. In this case, Lemma 12 applies, telling us that there is an optimal attention outcome p^* that has one supported belief on each of L and R, and is either ternary with support containing $\nu_{(1,0)}$ (with slopes between adjacent beliefs equal to s) or binary.

- If p^* is binary, then p^* has support $\{\nu_{(a_1,z_1)}, \ \nu_{(a_2,z_2)}\}$ for some distinct $(a_1, z_1) \in L$ and $(a_2, z_2) \in R$. Let $\tilde{s} := \frac{z_2 - z_1}{a_2 - a_1}$. The fact that $(a_1, z_1) \in L$ and $(a_2, z_2) \in R$ implies that $\tilde{s} \in \left[-\frac{\mu_0}{1-a_\mu}, \frac{\mu_0}{1+a_\mu}\right]$. We also know that $\tilde{s} \in [-s, s]$ by Lemma 10.

  Given such $\tilde{s}$,

$$
\begin{aligned}
(a_1, z_1) &= \left(-\frac{1-\mu_0+\tilde{s}a_\mu}{1-\tilde{s}}, \frac{\mu_0 - \tilde{s}(1+a_\mu)}{1-\tilde{s}}\right) \text{ and} \\
(a_2, z_2) &= \left(\frac{1-\mu_0+\tilde{s}a_\mu}{1+\tilde{s}}, \frac{\mu_0 + \tilde{s}(1-a_\mu)}{1+\tilde{s}}\right)
\end{aligned}
$$

  are the unique intersections of L and R, respectively, with the line of slope $\tilde{s}$ through (a_μ, μ_0). Next, Bayes-plausibility tells us that

$$
\begin{aligned}
p^*\{\nu_{(a_1,z_1)}\} &= \frac{1-\tilde{s}}{2}\,\frac{1-\mu_0-a_\mu}{1-\mu_0+\tilde{s}a_\mu} \\
p^*\{\nu_{(a_2,z_2)}\} &= \frac{1+\tilde{s}}{2}\,\frac{1-\mu_0+a_\mu}{1-\mu_0+\tilde{s}a_\mu}
\end{aligned}
$$

  Therefore, the principal's objective can be written as (appealing to Observation 1)

$$
\begin{aligned}
\int_{\Delta\Theta} U_P \, \mathrm{d}p^* &= \mathbb{V}_{\mu\sim p^*}\left[a^*(\nu)\right] - \mathbb{V}_{\theta\sim\mu}(\theta) \\
&= \int_{\Delta\Theta} \left(a^*(\nu) - a_\mu\right)^2 \, \mathrm{d}p^*(\nu) - \mathbb{V}_{\theta\sim\mu}(\theta) \\
&= \tfrac{1-\tilde{s}}{2}\tfrac{1-\mu_0-a_\mu}{1-\mu_0+\tilde{s}a_\mu}\left(-\tfrac{1-\mu_0+a_\mu}{1-\tilde{s}}\right)^2 + \tfrac{1+\tilde{s}}{2}\tfrac{1-\mu_0+a_\mu}{1-\mu_0+\tilde{s}a_\mu}\left(\tfrac{1-\mu_0-a_\mu}{1+\tilde{s}}\right)^2 - \mathbb{V}_{\theta\sim\mu}(\theta) \\
&= \tfrac{(1-\mu_0+a_\mu)(1-\mu_0+a_\mu)}{2(1-\mu_0+\tilde{s}a_\mu)}\left(\tfrac{1-\mu_0+a_\mu}{1-\tilde{s}} + \tfrac{1-\mu_0-a_\mu}{1+\tilde{s}}\right) - \mathbb{V}_{\theta\sim\mu}(\theta) \\
&= \frac{(1-\mu_0)^2 - a_\mu^2}{1-\tilde{s}^2} - \mathbb{V}_{\theta\sim\mu}(\theta). \qquad \text{(B.8)}
\end{aligned}
$$

As $0 \leq a_\mu < 1-\mu_0$ (because $(a_\mu, \mu_0) \in \text{int}(\mathcal{B})$), the above objective is strictly increasing in $|\tilde{s}|$, so p^* being an optimal attention outcome implies that $|\tilde{s}| = \min\{s, \frac{\mu_0}{1-a_\mu}\}$.

- If there is no binary optimal attention outcome, then p^* is ternary and is supported on the set $\{\nu_{(a_1,z_1)},\ \nu_{(a_2,z_2)},\ \nu_{(1,0)}\}$ such that $(a_1, z_1) \in L$ and $\frac{z_2-z_1}{a_2-a_1} = s = \frac{z_2}{1-a_2}$. By the last part of Lemma 12, this is only possible when $s > \frac{\mu_0}{1-a_\mu}$. Recall that

$$\begin{aligned} p_2(a_1,a_2) &= \frac{a_1(a_1+1-\mu_0)(1-s)+a_2(2a_1+1-\mu_0-a_\mu)s}{-s(a_2-a_1)\left[a_1\left(1-s\right)+a_2\left(1+s\right)\right]}, \\ t_3(a_1,a_2) &= (-a_1 - \tfrac{2s}{1-s}a_2,\ 1+a_1+\tfrac{2s}{1-s}a_2). \end{aligned}$$

Hence, that $t_3 = (1,0)$ and $p_2(a_1,a_2) \geq 0$ imply that $a_2 = -\frac{1-s}{2s}(1+a_1)$ and $a_1 \in \left[-1, -\frac{1-\mu_0-a_\mu}{1-\mu_0+a_\mu}\right]$; and moreover, any such pair has $(a_1,a_2) \in \mathcal{A}$. Substituting into (B.7) yields

$$\begin{aligned} \int_{\Delta\Theta} U_P \,\mathrm{d}p^* + \mathbb{V}_{\theta\sim\mu}(\theta) &= -a_\mu^2 + a_1^2 - \tfrac{a_1^2+(1-\mu_0)a_1}{s} + \left(2a_1 + \tfrac{1-\mu_0}{1-s} + \tfrac{2s-1}{1-s}a_\mu\right)\frac{1-s}{2s}(1+a_1) \\ &= -a_\mu^2 + \tfrac{1}{2s}\left[a_1(1-a_\mu+\mu_0-2s(1-a_\mu))+1-\mu_0-a_\mu+2sa_\mu\right]. \end{aligned} \tag{B.9}$$

Note that the above objective function is affine in a_1. Note also that when $a_1 = -\frac{1-\mu_0-a_\mu}{1-\mu_0+a_\mu}$, $p_2(a_1,a_2) = 0$, so that the binary policy with slope $-\frac{\mu_0}{1-a_\mu}$ is obtained. Since, by hypothesis, no binary policy is an optimal attention outcome, a_1 must be optimally set to -1, i.e., $s > \frac{1-a_\mu+\mu_0}{2(1-a_\mu)}$.

Now, we consider various subcases for the value of κ, and find an optimal policy using payoffs computed in (B.8) and (B.9).

- First, suppose $\kappa \in (\kappa_1, \kappa_2]$ so that $\frac{1-a_\mu+\mu_0}{2(1-a_\mu)} \leq s < 1$. In (B.9), it is optimal to set $a_1 = -1$ because $s \geq \frac{1-a_\mu+\mu_0}{2(1-a_\mu)}$. But since setting $a_1 = -\frac{1-\mu_0-a_\mu}{1-\mu_0+a_\mu}$ results in an optimal binary policy (i.e., a binary policy that maximizes the principal's utility among all IC binary policies),[59] Lemma 12 then implies that the resulting ternary policy by setting $a_1 = -1$ is an optimal attention outcome. This proves part (2) of the proposition.

- Next, suppose $\kappa \in [\kappa_2, \kappa_3]$, so that $\frac{\mu_0}{1-a_\mu} \leq s \leq \frac{1-a_\mu+\mu_0}{2(1-a_\mu)}$. In (B.9), it is optimal to set $a_1 = -\frac{1-\mu_0-a_\mu}{1-\mu_0+a_\mu}$ because $s \leq \frac{1-a_\mu+\mu_0}{2(1-a_\mu)}$, which, as noted before, results in a binary policy. This means that any ternary policy with $\nu_{(1,0)}$ in the support (and critical slopes) is dominated by a binary policy. Then, by Lemma 12, we know that an optimal binary policy is optimal. That is, the binary policy with slope $|\tilde{s}| = \min\{s, \frac{\mu_0}{1-a_\mu}\} = \frac{\mu_0}{1-a_\mu}$ is an optimal attention outcome. This proves the part (3) of the proposition.

[59] Note that now $\min\{s, \frac{\mu_0}{1-a_\mu}\} = \frac{\mu_0}{1-a_\mu}$, and the resulting binary policy by setting $a_1 = -\frac{1-\mu_0-a_\mu}{1-\mu_0+a_\mu}$ has slope $-\frac{\mu_0}{1-a_\mu}$.

- Finally, suppose $\kappa \in (\kappa_3, \kappa_4)$ so that $0 < s < \frac{\mu_0}{1-a_\mu}$. Then by Lemma 12, we know that an optimal binary policy is optimal.[60] That is, the binary policy with slope $|\tilde{s}| = \min\{s, \frac{\mu_0}{1-a_\mu}\} = s$ is an optimal attention outcome. This proves the remainder of part (4) of the proposition.

All that remains is the uniqueness result. Let us show that there is a unique (up to reflection) optimal attention outcome whenever $\kappa \neq \kappa_2$.

If $\kappa \leq \frac{1}{2}$, then full information is IC, and it is the unique minimizer of $p \mapsto \int_{\Delta\Theta} U_P \, \mathrm{d}p$ over $\mathcal{R}(\mu)$, and so is the unique optimal attention outcome. If $\kappa > 2$, then no information is uniquely IC, and so is the unique optimal attention outcome. If $\kappa = 2$, then every IC information policy is supported on $\{\nu_{(a,z)} : \ (a, z) \in \mathcal{B}, \ z = \mu_0 \}$, a convex set over which U_P is strictly convex; p^O is therefore (being a unique $\succeq^B$-maximum over all such information policies) the unique optimal attention outcome.

Now, we focus on the remaining case of $\frac{1}{2} < \kappa < 2$, and let p^* be an arbitrary optimal attention outcome.

Letting $\mathcal{P}^*$ be the set of nonredundant optimal attention outcome, Claims 7 and 9 tell us that $p^* \in \overline{\mathrm{co}}(\mathcal{P}^*)$. As U_A is strictly concave on L (since $s < 1$) and p^* is IC, it must be that the positive measures $p^*(\cdot \cap L)$ and $p^*(\cdot \cap R)$ both have support of size at most one. But, by Lemma 12, any element of $\overline{\mathrm{co}}(\mathcal{P}^*) \setminus \mathcal{P}^*$ would violate this property. Therefore $p^* \in \mathcal{P}^*$, i.e., it is nonredundant.

We now show that p^* is unique (up to reflection) unless $\kappa = \kappa_2$. To this end, we need to rule out both multiplicity within the class of policies of the form guaranteed by Lemma 12, and existence of an optimal nonredundant policy outside of that class. The first part of the current proof already shows that within that class, the optimal attention outcome is unique unless $\kappa = \kappa_2$. We now argue that, as long as $\kappa \neq \kappa_2$, no nonredundant policy outside that class can be an optimal attention outcome. Following the notation of Lemma 12's proof, take any $(a_1, a_2) \in \mathcal{A} \setminus \mathcal{A}^*$, and assume, for a contradiction, that $p^{\mathcal{T}}_{t(a)}$ is an optimal attention outcome. As $\tilde{a}_2 \mapsto \int_{\Delta\Theta} U_P \, \mathrm{d}p^{\mathcal{T}}_{t(a_1,\tilde{a}_2)}$ was shown to be affine, this can only be true if the affine function is constant, so that $p^{\mathcal{T}}_{t(a)}$ generates the same payoff to the principal as $p^{\mathcal{T}}_{t(a_1,\underline{a}_2)}$ and $p^{\mathcal{T}}_{t(a_1,\bar{a}_2)}$ do, where $(a_1, \underline{a}_2), (a_1, \bar{a}_2) \in \mathcal{A}^*$ with $\underline{a}_2 < a_2 < \bar{a}_2$. In particular, both $p^{\mathcal{T}}_{t(a_1,\underline{a}_2)}$ and $p^{\mathcal{T}}_{t(a_1,\bar{a}_2)}$ are also optimal attention outcomes. Consider alternative cases of $\kappa \in (\frac{1}{2}, 2) \setminus \{\kappa_2\} = (\kappa_1, \kappa_2) \cup (\kappa_2, \kappa_3) \cup [\kappa_3, \kappa_4)$.

- If $\kappa \in (\kappa_1, \kappa_2) \cup (\kappa_2, \kappa_3)$, then $p^{\mathcal{T}}_{t(a_1,\underline{a}_2)}$ is ternary and $p^{\mathcal{T}}_{t(a_1,\bar{a}_2)}$ is binary. But, as we argued above, the optimal ternary attention outcome and the optimal binary attention outcome are strictly payoff ranked for $\kappa \neq \kappa_2$. This is a contradiction to $p^{\mathcal{T}}_{t(a_1,\underline{a}_2)}$ and $p^{\mathcal{T}}_{t(a_1,\bar{a}_2)}$ *both* being optimal attention outcomes.

- If $\kappa \in [\kappa_3, \kappa_4)$, then both $p^{\mathcal{T}}_{t(a_1,\underline{a}_2)}$ and $p^{\mathcal{T}}_{t(a_1,\bar{a}_2)}$ are binary, while one element in $\mathrm{supp}(p^{\mathcal{T}}_{t(a_1,\underline{a}_2)})$ is in $\mathrm{int}(\mathcal{B})$ (see the left panel of Figure B.3 for illustration). By Claim 11, $p^{\mathcal{T}}_{t(a_1,\underline{a}_2)}$ is not principal-optimal, a contradiction to $p^{\mathcal{T}}_{t(a_1,\underline{a}_2)}$ being an optimal attention outcome. □

[60] Note that the last part of Lemma 12 says that policies of the second form (ternary policies with $\nu_{(1,0)}$ in the support and critical slopes) do not exist when $s < \frac{\mu_0}{1-a_\mu}$, so we can only focus on binary policies.

Proof of Corollary 1. First, by Proposition 5, when $\kappa \le \kappa_1$, full disclosure policy p^F is a principal-optimal attention outcome. Next, by Proposition 5, when $\kappa > \kappa_4$, the policy p^N (i.e., no disclosure) is a principal-optimal attention outcome; as (p^F, q) has a weakly higher principal objective for any $q \in \mathcal{R}(\mu)$, it follows that full disclosure is a principal-optimal disclosure policy.

Now, suppose $\kappa \in (\kappa_1, \kappa_4] = (\frac{1}{2}, 2]$. For convenience, we work in the space $\mathcal{R}_{\mathcal{B}} = \{p \in \Delta\mathcal{B} : \int_B (a, z)\, \mathrm{d}p(a, z) = (a_\mu, \mu_0)\}$. Consider any $p_1 \in \mathcal{R}_{\mathcal{B}}$. The map $\eta : \mathcal{B} \to \Delta(L \cup R)$ given by

$$\eta(\cdot|a, z) := \begin{cases} \frac{1-z+a}{2(1-z)}\delta_{(1-z,z)} + \frac{1-z-a}{2(1-z)}\delta_{(z-1,z)} & : \ z \neq 1, \\ (a, z) & : \ z = 1, \end{cases}$$

is a mean-preserving spread. Therefore, $p_2 \in \mathcal{R}_{\mathcal{B}}$, where $p_2(\hat{B}) = \int_{\mathcal{B}} \eta(\hat{B}|\cdot)\, \mathrm{d}p_1$ for all Borel $\hat{B} \subseteq \mathcal{B}$. Notice now that, for $(a, z) \in \mathcal{B}$:

$$\begin{aligned} U_A(\nu_{(a,z)}) &= -\left[\mathbb{E}_{\theta\sim\nu_{(a,z)}}(\theta^2) - a^2\right] - \kappa||\nu_{(a,z)} - \mu||^2 \\ &= a^2 - (1-z) - \tfrac{\kappa}{2}\left[(a-a_\mu)^2 + 3(z-\mu_0)^2\right] \\ &= \left(1 - \tfrac{\kappa}{2}\right)a^2 + \kappa a_\mu a - \tfrac{\kappa}{2}a_\mu^2 - \tfrac{3\kappa}{2}(z-\mu_0)^2 - (1-z). \end{aligned}$$

First, $U_A(\nu_{(a,z)})$ is convex in a, strictly so if $\kappa \neq \kappa_4$. Thus, $\int_{\Delta\Theta} U_A(\nu_{(\cdot)})\, \mathrm{d}p_1 \le \int_{\Delta\Theta} U_A(\nu_{(\cdot)})\, \mathrm{d}p_2$, strictly so if $\kappa \neq \kappa_4$ and $p_1 \neq p_2$. This implies if $\kappa \in (\kappa_1, \kappa_4)$, the agent best response to full disclosure q^* must have its support contained in $L \cup R$. Next, for $z \in [0, 1]$,

$$\begin{aligned} \tfrac{\mathrm{d}^2}{\mathrm{d}z^2} U_A(\nu_{(1-z,z)}) &= \tfrac{\mathrm{d}^2}{\mathrm{d}z^2}\left[\left(1 - \tfrac{\kappa}{2}\right)(1-z)^2 + \kappa a_\mu(1-z) - \tfrac{\kappa}{2}a_\mu^2 - \tfrac{3\kappa}{2}(z-\mu_0)^2 - (1-z)\right] \\ &= (2-\kappa) - 3\kappa - 0 < 2 - 4\kappa_1 = 0. \end{aligned}$$

Therefore, $U_A(\nu_{(\cdot)})$ is strictly concave along R; the same is true along L. As $(a_\mu, \mu_0) \in \mathrm{int}(\mathcal{B})$, this implies that $|\mathrm{supp}(q^*) \cap L| = |\mathrm{supp}(q^*) \cap R| = 1$; that is, if $\kappa \in (\kappa_1, \kappa_4)$, the agent best response to full disclosure is a binary policy with one supported belief on L and the other on R.[61]

Now we consider the agent's optimization problem within this class of binary policies. Consider such a policy $q_{\tilde{s}} \in \mathcal{R}_{\mathcal{B}}$ with support $\{(a_1, z_1), (a_2, z_2)\}$ such that $a_1 < a_2$ and $z_2 - z_1 = \tilde{s}(a_2 - a_1)$. As derived before, given an $\tilde{s} \in \left[-\frac{\mu_0}{1-a_\mu}, \frac{\mu_0}{1+a_\mu}\right]$, the line of slope $\tilde{s}$ intersects L and R, respectively, at

$$\begin{aligned} (a_1, z_1) &= \left(-\frac{1-\mu_0+\tilde{s}a_\mu}{1-\tilde{s}}, \frac{\mu_0 - \tilde{s}(1+a_\mu)}{1-\tilde{s}}\right) \text{ and} \\ (a_2, z_2) &= \left(\frac{1-\mu_0+\tilde{s}a_\mu}{1+\tilde{s}}, \frac{\mu_0 + \tilde{s}(1-a_\mu)}{1+\tilde{s}}\right) \end{aligned}$$

[61] When $\kappa = \kappa_4$, there exists an agent best response to full disclosure with the same features, though some non-binary policy may also be optimal.

Next, Bayes-plausibility tells us that

$$q_{\tilde{s}}\{(a_1, z_1)\} = \frac{1-\tilde{s}}{2}\frac{1-\mu_0-a_\mu}{1-\mu_0+\tilde{s}a_\mu}$$
$$q_{\tilde{s}}\{(a_2, z_2)\} = \frac{1+\tilde{s}}{2}\frac{1-\mu_0+a_\mu}{1-\mu_0+\tilde{s}a_\mu}$$

The agent's attention cost from this binary policy $q_{\tilde{s}}$ is

$$\begin{aligned}\kappa \int_{\Delta\Theta} c(\nu_{(\cdot)})\, \mathrm{d}p_{\tilde{s}} &= \kappa \int_{\Delta\Theta} ||\nu_{(\cdot)} - \mu||^2\, \mathrm{d}p_{\tilde{s}} \\ &= \sum_{i=1}^{2} \frac{\kappa}{2}\left[(a_i - a_\mu)^2 + 3(z_i - \mu_0)^2\right] q_{\tilde{s}}\{(a_i, z_i)\} \\ &= \frac{\kappa(1+3\tilde{s}^2)}{2(1-\tilde{s}^2)}\left[(1-\mu_0)^2 - a_\mu^2\right] \end{aligned} \tag{B.10}$$

The agent's utility from this binary policy $q_{\tilde{s}}$ is

$$\begin{aligned}\int_{\Delta\Theta} U_A(\nu_{(\cdot)})\, \mathrm{d}p_{\tilde{s}} &= \int_{\Delta\Theta} U_P(\nu_{(\cdot)})\, \mathrm{d}p_{\tilde{s}} - \kappa \int_{\Delta\Theta} c(\nu_{(\cdot)})\, \mathrm{d}p_{\tilde{s}} \\ &= \frac{(1-\mu_0)^2 - a_\mu^2}{1-\tilde{s}^2} - \frac{\kappa(1+3\tilde{s}^2)}{2}\frac{(1-\mu_0)^2 - a_\mu^2}{1-\tilde{s}^2} - \mathbb{V}_{\theta\sim\mu}(\theta) \\ &= \left(\frac{3\kappa}{2} - \frac{2\kappa - 1}{1-\tilde{s}^2}\right)\left[(1-\mu_0)^2 - a_\mu^2\right] - \mathbb{V}_{\theta\sim\mu}(\theta) \end{aligned} \tag{B.11}$$

where the second line follows from equations (B.8) and (B.10). Since $\kappa > \kappa_1 = \frac{1}{2}$, the agent's objective is decreasing in $|\tilde{s}|$, so that the agent's unique optimal policy in this class is $q_0 = p^O$. Therefore, p^O is the unique agent best response to full disclosure if $\kappa \in (\kappa_1, \kappa_4)$.

In the case that $\kappa = \kappa_4$, same argument as above shows that p^O is an agent best response to full disclosure, while Proposition 5 shows p^O to be a principal-optimal attention outcome. Therefore, $\left(p^F, p^O\right)$ is also a solution to program (2.2), and p^F is principal-optimal.

Finally, in the case that $\kappa \in (\kappa_1, \kappa_4)$, we showed above that p^O is a unique agent best response to p^F. Direct computation then shows that the optimal attention outcome named in Proposition 5 yields a strictly higher principal value. Thus, providing full information is strictly suboptimal. □

Appendix C

Proofs for Chapter 3

Proofs for Section 3.3

Proof of Lemma 2. Since U_R is continuous, a best response q^* to full disclosure exists, and can be found by concavifying U_R. Moreover, since q^* is IC, we have $|\text{supp}(q^*)| \leq 2$ by the argument in the paragraph preceding Lemma 1. Take such a q^*. If $|\text{supp}(q^*)| = 1$, then "no information" is the unique best response to full disclosure.[62] If $|\text{supp}(q^*)| = 2$, let $\text{supp}(q^*) = \{\nu_1^*, \nu_2^*\}$. Since U_R is strictly concave on $[0, \lambda]$ and on $[\lambda, 1]$, we have $\nu_1^* < \lambda < \nu_2^*$. But then, by piecewise strict concavity of U_R again, we have $U_R(\nu) < \phi(\nu; \nu_1^*, \nu_2^*)$ for all $\nu \neq \nu_1^*, \nu_2^*$, where

$$\phi(\nu; \nu_1^*, \nu_2^*) \equiv U_R(\nu_1^*) + \frac{U_R(\nu_2^*) - U_R(\nu_1^*)}{\nu_2 - \nu_1}(\nu - \nu_1^*). \tag{C.1}$$

Hence, any policy other than q^* gives the buyer a strictly lower payoff than q^* does, and so his best response to full disclosure is unique. □

Proof of Lemma 3. This follows directly from the argument above Figure 2.1 in Chapter 2. □

To establish Proposition 5, we need the following lemma.

Lemma 13. *For any $\nu_1 \in [\nu_1^*, \mu]$, there is a unique $\nu_2 \in (\lambda, \nu_2^*]$ such that*

$$\frac{U_R(\nu_2) - U_R(\nu_1)}{\nu_2 - \nu_1} = U_R'(\nu_1). \tag{C.2}$$

Proof of Lemma 13. By definition of q^* and $\{\nu_1^*, \nu_2^*\}$, the concavification condition implies that

$$U_R'(\nu_1^*) \leq \frac{U_R(\nu_2^*) - U_R(\nu_1^*)}{\nu_2^* - \nu_1^*} \leq U_R'(\nu_2^*). \tag{C.3}$$

[62] If there is another different policy q that generates the same payoff to the buyer as q^*, then a convex combination of q and q^* is also a best response (as the buyer's payoff is linear in information policies). But such a convex combination has at least three supported beliefs, thus can be strictly improved upon, a contradiction.

The first (second) weak inequality holds with equality whenever $\nu_1^* > 0$ ($\nu_2^* < 1$). This implies that $\nu_1^* < \mu < \lambda < \nu_2^*$, for otherwise U_R would be strictly concave on $[\nu_1^*, \nu_2^*]$ in which case the buyer has an incentive to garble information further.

Consider the function $\eta(\nu) \equiv U_R(\nu_1) + U_R'(\nu_1)(\nu - \nu_1)$. By strict concavity of U_R on $[0, \lambda]$, we have $\eta(\lambda) > U_R(\lambda)$. In addition,

$$\begin{aligned}
\eta(\nu_2^*) &= U_R(\nu_1) + U_R'(\nu_1)(\nu_2^* - \nu_1) \\
&\le U_R(\nu_1) + U_R'(\nu_1^*)(\nu_2^* - \nu_1) \\
&\le [U_R(\nu_1^*) + U_R'(\nu_1^*)(\nu_1 - \nu_1^*)] + U_R'(\nu_1^*)(\nu_2^* - \nu_1) \\
&= U_R(\nu_1^*) + U_R'(\nu_1^*)(\nu_2^* - \nu_1^*) \\
&\le U_R(\nu_1^*) + \frac{U_R(\nu_2^*) - U_R(\nu_1^*)}{\nu_2^* - \nu_1^*}(\nu_2^* - \nu_1^*) \\
&= U_R(\nu_2^*)
\end{aligned}$$

where the second and third lines follow from the strict concavity of U_R on $[0, \lambda]$ and $\nu_1^* \le \nu_1 \le \mu(< \lambda)$, holding with strictly inequality whenever $\nu_1 > \nu_1^*$, and the fifth line follows from (C.3). Intermediate value theorem implies that there exists a $\nu_2 \in (\lambda, \nu_2^*]$ such that (C.2) holds. Moreover, for all $\nu \in (\lambda, \nu_2^*)$, piecewise strict concavity of U_R and condition (C.3) imply that $U_R'(\nu) > U_R'(\nu_2^*) \ge \frac{U_R(\nu_2^*) - U_R(\nu_1^*)}{\nu_2^* - \nu_1^*} \ge U_R'(\nu_1^*) \ge U_R'(\nu_1)$. So, $\eta(\nu) - U_R(\nu)$ is strictly monotonic on $(\lambda, \nu_2^*]$. Thus, for any given $\nu_1 \in [\nu_1^*, \mu]$, there is a unique $\nu_2 \in (\lambda, \nu_2^*]$ which satisfies (C.2). □

Lemma 13 allows us to define a mapping $\nu_2^T : [\nu_1^*, \mu] \to (\lambda, \nu_2^*]$ such that $\{\nu_1, \nu_2^T(\nu_1)\}$ satisfies (C.2). It is easy to see that ν_2^T is strictly decreasing in ν_1. We then define $\nu_1^T : [\nu_2^T(\mu), \nu_2^*] \to [\nu_1^*, \mu]$ as its inverse function, that is, $\nu_1^T = (\nu_2^T)^{-1}$. Note that $\nu_2^T(\nu_1^*) = \nu_2^*$ and $\nu_1^T(\nu_2^*) = \nu_1^*$, whenever $\nu_1^* > 0$.

Proof of Proposition 5. We first show that any $\{\nu_1, \nu_2\}$ that satisfies the constraints of program (3.4) renders an information policy that is IC and binary, i.e., such a policy satisfies the constraints in program (3.2). This will imply that the value of program (3.4) is weakly lower than the value of the seller's program. We then argue that these two values are equal.

Take any $\{\nu_1, \nu_2\}$ that satisfies constraints (i) and (ii) of program (3.4) and let p be the information policy supported on $\{\nu_1, \nu_2\}$. By Lemma 13, $\nu_2 = \nu_2^T(\nu_1)$, and so $\nu_1 < \lambda < \nu_2$. Apparently, p is binary, so we now show that p is IC. If the seller chooses p, condition (3.3) implies that the buyer is restricted to choose from distributions over posterior beliefs in $[\nu_1, \nu_2]$. For $\nu \in [0, 1]$, let

$$\begin{aligned}
\phi(\nu; \nu_1, \nu_2) &\equiv U_R(\nu_1) + \frac{U_R(\nu_2) - U_R(\nu_1)}{\nu_2 - \nu_1}(\nu - \nu_1) \\
&= U_R(\nu_1) + U_R'(\nu_1)(\nu - \nu_1)
\end{aligned} \tag{C.4}$$

be the expression for the affine function connecting $(\nu_1, U_R(\nu_1))$ and $(\nu_2, U_R(\nu_2))$, where the second line follows from constraint (i) in program (3.4). To establish that p is IC, it suffices to show

$$U_R(\nu) \le \phi(\nu; \nu_1, \nu_2), \forall \nu \in [\nu_1, \nu_2]$$

which will imply conv $\left(U_R|_{[\nu_1,\nu_2]}\right) = \phi(\cdot;\nu_1,\nu_2)$. We now show that the above inequality indeed holds. First, equation (C.4) and strict concavity of U_R on $[\nu_1,\lambda]$ imply that, for all $\nu \in (\nu_1,\lambda]$, $U_R(\nu) < \phi(\nu;\nu_1,\nu_2)$. Moreover, piecewise strict concavity of U_R and condition (C.3) imply that[63]

$$\phi'(\cdot;\nu_1,\nu_2) = U_R'(v_1) \leq U_R'(\nu_2). \tag{C.5}$$

But then, strict concavity of U_R on $[\lambda,\nu_2]$ imply that, for all $\nu \in [\lambda,\nu_2)$, $U_R(\nu) < \phi(\nu;\nu_1,\nu_2)$. Hence, for all $\nu \in [\nu_1,\nu_2]$, $U_R(\nu) \leq \phi(\nu;\nu_1,\nu_2)$, and so p is IC.

Since any $\{\nu_1,\nu_2\}$ satisfying constraints (i) and (ii) in program (3.4) delivers an IC and binary policy, the value of program (3.4) is weakly lower than the value of program (3.2). We now argue that the values of these two programs are actually equal. To this end, note first that by Lemma 3 and condition (3.3), constraint (ii) in program (3.4) is without loss. Note also that for a *fixed* $\nu_1 \in [\nu_1^*,\mu)$, the *smallest* ν_2 such that $\phi(.;\nu_1,\nu_2)$ majorizes $U_R|_{[\nu_1,\nu_2]}$ is $\nu_2^T(\nu_1)$.[64] Since the objective is decreasing in ν_2, constraint (i) in program (3.4) is also without loss. Therefore, the values of programs (3.2) and (3.4) are equal, and the existence of solution to program (3.4) follows directly from Lemma 1.

Finally, it is easy to see that the support of any optimal attention outcome must solve (3.4). □

Proof of Proposition 6. By Lemma 3, we only need to show that q^* is not an optimal outcome.

Since $\nu_1^* > 0$ and $\nu_2^* < 0$, we now have

$$U_R'(\nu_1^*) = \frac{U_R(\nu_2^*) - U_R(\nu_1^*)}{\nu_2^* - \nu_1^*} = U_R'(\nu_2^*). \tag{C.6}$$

By Lemma 13, program (3.4) is equivalent to

$$\max_{\nu_1 \in [\nu_1^*,\mu)} \frac{\mu - \nu_1}{\nu_2^T(\nu_1) - \nu_1}, \tag{C.7}$$

which (by Lemma 13 again) is equivalent to

$$\max_{\nu_2 \in (\nu_1^T(\mu),\nu_2^*]} \frac{\mu - \nu_1^T(\nu_2)}{\nu_2 - \nu_1^T(\nu_2)}. \tag{C.8}$$

Let $\pi(\nu_2) \equiv \frac{\mu - \nu_1^T(\nu_2)}{\nu_2 - \nu_1^T(\nu_2)}$. We have

$$\frac{d\pi}{d\nu_2} = -\frac{\mu - \nu_1^T(\nu_2)}{\left(\nu_2 - \nu_1^T(\nu_2)\right)^2} - \frac{\nu_2 - \mu}{\left(\nu_2 - \nu_1^T(\nu_2)\right)^2}\frac{d\nu_1^T}{d\nu_2}.$$

By (C.2), $\frac{d\nu_1^T}{d\nu_2} = \frac{U_R'(\nu_2) - U_R'(\nu_1^T(\nu_2))}{U_R''(\nu_1^T(\nu_2))(\nu_2 - \nu_1^T(\nu_2))}$. By (C.6), $\left.\frac{d\nu_1^T}{d\nu_2}\right|_{\nu_2=\nu_2^*} = 0$, and so $\left.\frac{d\pi}{d\nu_2}\right|_{\nu_2=\nu_2^*} = -\frac{\mu-\nu_1^*}{\left(\nu_2^*-\nu_1^*\right)^2} < 0$. Therefore, q^* is not an optimal attention outcome. Since the Blackwell order is a partial order (i.e., no indifference), any optimal attention outcome p^* is strictly less informative than q^*. □

[63] Specifically, $U_R'(v_1) \leq U_R'(v_1^*) \leq \frac{U_R(\nu_2^*) - U_R(\nu_1^*)}{\nu_2^* - \nu_1^*} \leq U_R'(\nu_2^*) \leq U_R'(\nu_2)$.

[64] For any $\tilde{\nu}_2 < \nu_2^T(\nu_1)$, we have $\phi(\nu;\nu_1,\tilde{\nu}_2) < U_R(\nu)$ for ν slightly greater than ν_1 (see Figure 3.2). Hence, any policy supported on $\{\nu_1,\tilde{\nu}_2\}$ with $\tilde{\nu}_2 < \nu_2^T(\nu_1)$ is not IC, while the one supported on $\{\nu_1,\nu_2^T(\nu_1)\}$ is IC (as shown).

Proof of Proposition 7. Since $\text{supp}(r_{t^*}) = \{0, \nu_2^O\}$, condition (3.3) implies that when r_t is chosen by the seller, the buyer is restricted to choose from distribution over beliefs in $[0, \nu_2^O]$. It is sufficient to show that $\phi(\cdot; \nu_1^O, \nu_2^O)$ (defined in (C.4)) majorizes U_R on $[0, \nu_2^O]$.

Since p^* is an optimal attention outcome (thus IC), we have $\phi(\nu; \nu_1^O, \nu_2^O) \geq U_R(\nu), \text{ for all } \nu \in [\nu_1^O, \nu_2^O]$. Moreover, since $\{\nu_1^O, \nu_2^O\}$ solves program (3.4) and U_R is strictly concave on $[0, \nu_1^O]$, we know that $\frac{d\phi}{d\nu} = U_R'(\nu_1^O) > U_R(\nu)$ for all $\nu \in [0, \nu_1^O)$, and so, $\phi(\nu; \nu_1^O, \nu_2^O) \geq U_R(\nu), \text{ for all } \nu \in [0, \nu_1^O]$. Hence, $\phi(\nu; \nu_1^O, \nu_2^O) \geq U_R(\nu)$ for all $\nu \in [0, \nu_2^O]$, as desired. $\square$

Proofs for the Results on Comparative Statics

Let $q^*(\kappa)$ be the buyer's best response to full information when the cost parameter is κ, and let $\text{supp}(q^*(\kappa)) = \{\nu_1^*(\kappa), \nu_2^*(\kappa)\}$. For any κ and any $\nu_1 \in [\nu_1^*(\kappa), \mu]$, let $\nu_2^T(\nu_1; \kappa)$ be the unique $\nu_2 \in (\lambda, \nu_2^*(\kappa)]$ such that (C.2) holds.

Lemma 14. *For all $\nu_1 \in (\nu_1^*(\kappa), \mu)$, $\frac{\partial \nu_2^T}{\partial \kappa} > 0$. Moreover, for all $\nu_2 \in (\lambda, \nu_2^*(\kappa))$, $\frac{\partial \nu_1^T}{\partial \kappa} > 0$.*

Proof of Lemma 14. Recall that $U_R(\nu) = \max\{0, \nu - \lambda\} - \kappa c(\nu)$. If condition (C.2) holds, we have

$$\nu_2 - \lambda - \kappa(c(\nu_2) - c(\nu_1)) = -(\nu_2 - \nu_1)\kappa c'(\nu_1).$$

Totally differentiating both sides with respect to (ν_1, ν_2, κ), we have

$$(\nu_2 - \nu_1)\kappa c''(\nu_1) d\nu_1 + [1 - \kappa(c'(\nu_2) - c'(\nu_1))]\, d\nu_2 + [(\nu_2 - \nu_1)c'(\nu_1) - (c(\nu_2) - c(\nu_1))]\, d\kappa = 0.$$

Therefore,

$$\frac{\partial \nu_2^T}{\partial \kappa} = \frac{c(\nu_2) - c(\nu_1) - c'(\nu_1)(\nu_2 - \nu_1)}{1 - \kappa(c'(\nu_2) - c'(\nu_1))} > 0,$$

where the numerator is greater than 0 because of the strict convexity of c, and the denominator is greater than 0 because $U_R'(\nu_1) < U_R'(\nu_2)$.[65] Similarly,

$$\frac{\partial \nu_1^T}{\partial \kappa} = \frac{c(\nu_2) - c(\nu_1) - c'(\nu_1)(\nu_2 - \nu_1)}{(\nu_2 - \nu_1)\kappa c''(\nu_1)} > 0. \qquad \square$$

Lemma 15. *Take any $\underline{\kappa}$ and $\bar{\kappa}$ with $\bar{\kappa} > \underline{\kappa}$. Then $\nu_1^*(\bar{\kappa}) \geq \nu_1^*(\underline{\kappa})$, strictly so if $\nu_1^*(\underline{\kappa}) > 0$.*

[65] See condition (C.5) and note that here $\nu_1 > \nu_1^*(\kappa)$ and $\nu_2 < \nu_2^*(\kappa)$.

Proof of Lemma 15. If $\nu_1^*(\underline{\kappa}) = 0$, the (weak) inequality automatically holds. Suppose $\nu_1^*(\underline{\kappa}) > 0$. We then have

$$
\begin{aligned}
& \frac{U_R(\nu_2^*(\underline{\kappa});\underline{\kappa}) - U_R(\nu_1^*(\underline{\kappa});\underline{\kappa})}{\nu_2^*(\underline{\kappa}) - \nu_1^*(\underline{\kappa})} = U_R'(\nu_1^*(\underline{\kappa});\underline{\kappa}) \\
\iff & \frac{\nu_2^*(\underline{\kappa}) - \lambda - \kappa[c(\nu_2^*(\underline{\kappa})) - c(\nu_2^*(\underline{\kappa}))]}{\nu_2^*(\underline{\kappa}) - \nu_1^*(\underline{\kappa})} = -\underline{\kappa}c'(\nu_1^*(\underline{\kappa})) \\
\iff & \frac{\nu_2^*(\underline{\kappa}) - \lambda}{\nu_2^*(\underline{\kappa}) - \nu_1^*(\underline{\kappa})} = \underline{\kappa}\left[\frac{c(\nu_2^*(\underline{\kappa})) - c(\nu_2^*(\underline{\kappa}))}{\nu_2^*(\underline{\kappa}) - \nu_1^*(\underline{\kappa})} - c'(\nu_1^*(\underline{\kappa}))\right] \\
\Rightarrow & \frac{\nu_2^*(\underline{\kappa}) - \lambda}{\nu_2^*(\underline{\kappa}) - \nu_1^*(\underline{\kappa})} < \bar{\kappa}\left[\frac{c(\nu_2^*(\underline{\kappa})) - c(\nu_2^*(\underline{\kappa}))}{\nu_2^*(\underline{\kappa}) - \nu_1^*(\underline{\kappa})} - c'(\nu_1^*(\underline{\kappa}))\right] \\
\iff & U_R(\nu_2^*(\underline{\kappa});\bar{\kappa}) < U_R(\nu_1^*(\underline{\kappa});\bar{\kappa}) + U_R'(\nu_1^*(\underline{\kappa});\bar{\kappa})(\nu_2^*(\underline{\kappa}) - \nu_1^*(\underline{\kappa}))
\end{aligned}
\tag{C.9}
$$

where the fourth line follows from the strict convexity of c and that $\bar{\kappa} > \underline{\kappa}$. Conditions (C.9) and piecewise concavity of U_R imply that the affine function η of ν, defined by

$$\eta(\nu;\nu_1^*(\underline{\kappa}),\bar{\kappa}) := U_R(\nu_1^*(\underline{\kappa});\bar{\kappa}) + U_R'(\nu_1^*(\underline{\kappa});\bar{\kappa})(\nu - \nu_1^*(\underline{\kappa})),$$

is greater than $U_R(\cdot;\bar{\kappa})$ on $[0,\lambda]$ and $[\nu_2^*(\underline{\kappa}),1]$. Now consider two cases:

- $\eta(\nu;\nu_1^*(\underline{\kappa}),\bar{\kappa}) > U_R(\nu;\bar{\kappa})$ for all $\nu \in (\lambda,\nu_2^*(\underline{\kappa}))$.
  In this case, since U_R is strictly concave on $[0,\lambda]$, the following holds for all $\nu_1 \in [0,\nu_1^*(\kappa)]$:
  $$\eta(\nu;\nu_1,\bar{\kappa}) > U_R(\nu;\bar{\kappa}) \text{ for all } \nu \neq \nu_1.$$
  Hence, the concavification of $U_R(\cdot;\bar{\kappa})$ cannot have any $\nu_1 \leq \nu_1^*(\underline{\kappa})$ in the support, i.e., $\nu_1^*(\bar{\kappa}) > \nu_1^*(\underline{\kappa})$.

- $\eta(\nu;\nu_1^*(\underline{\kappa}),\bar{\kappa}) \leq U_R(\nu;\bar{\kappa})$ for some $\nu \in (\lambda,\nu_2^*(\underline{\kappa}))$.
  In this case, by piecewise strict concavity of U_R, we know that $v_2^*(\bar{\kappa}) < \nu_2^*(\underline{\kappa})$. But then,
  $$
  \begin{aligned}
  \nu_1^*(\underline{\kappa}) &= \nu_1^T(\nu_2^*(\underline{\kappa});\underline{\kappa}) \\
  &< \nu_1^T(\nu_2^*(\bar{\kappa});\underline{\kappa}) \\
  &< \nu_1^T(\nu_2^*(\bar{\kappa});\bar{\kappa}) \\
  &= \nu_1^*(\bar{\kappa})
  \end{aligned}
  $$
  where the second line follows from $v_2^*(\bar{\kappa}) < \nu_2^*(\underline{\kappa})$ and ν_1^T being strictly decreasing in ν_2 for fixed κ, and the third line follows from Lemma 14 (i.e., $\frac{\partial v_1^T}{\partial \kappa} > 0$). □

Proof of Proposition 8. Recall that for a fixed κ, the seller's program is equivalent to (C.7):

$$\max_{\nu_1 \in [\nu_1^*(\kappa),\mu)} \frac{\mu - \nu_1}{\nu_2^T(\nu_1;\kappa) - \nu_1}.$$

By Lemma 15, $[\nu_1^*(\bar{\kappa}),\mu) \subseteq [\nu_1^*(\underline{\kappa}),\mu)$. By Lemma 14, for any fixed $\nu_1 \in [\nu_1^*(\bar{\kappa}),\mu)$, we have $\nu_2^T(\nu_1;\underline{\kappa}) < \nu_2^T(\nu_1;\bar{\kappa})$. Since the objective is strictly decreasing in ν_2^T, we conclude that the value of the seller's program when $\kappa = \bar{\kappa}$ is strictly lower than when $\kappa = \underline{\kappa}$ □

Solution to the Example with Quadratic Cost

Proof of Proposition 10. If $\kappa \geq \frac{1}{2(1-2\mu)}$, "no information" is the buyer's unique best response to full information, and thus the unique optimal attention outcome.

We now consider the case where $\kappa < \frac{1}{2(1-2\mu)}$. Recall that the seller's program is equivalent to (C.8). To characterize $\nu_1^T(\nu_2;\kappa)$, notice that condition (C.2) becomes:

$$-2\kappa(\nu_1-\mu) = \frac{\nu_2 - 1/2 - \kappa[(\nu_2-\mu)^2 - (\nu_1-\mu)^2]}{\nu_2-\nu_1},$$

which leads to

$$\nu_1^T(\nu_2;\kappa) = \nu_2 - \sqrt{\frac{2\nu_2-1}{2\kappa}}.$$

It can be checked that $\nu_1^T(\nu_2;\kappa)$ is strictly decreasing in ν_2 whenever $\nu_2 < \nu_2^*(\kappa)$. Let $\underline{\nu}_2(\kappa) \in (\lambda, \nu_2^*(\kappa))$ be the unique solution to $\nu_1^T(\nu_2;\kappa) = \mu$. Such a $\underline{\nu}_2$ exists whenever $\kappa < \frac{1}{2(1-2\mu)}$.

By Proposition 5 and the deduction leading to (C.8), the seller's value is equal to

$$\begin{aligned} \max_{\nu_2 \in (\underline{\nu}_2(\kappa), \nu_2^*(\kappa)]} \quad & 1 - \frac{\nu_2-\mu}{\sqrt{\frac{2\nu_2-1}{2\kappa}}} \\ s.t. \quad & \nu_2 - \sqrt{\frac{2\nu_2-1}{2\kappa}} \geq 0 \end{aligned}.$$

It is uniquely maximized at $\nu_2^O = 1-\mu$ if $\kappa \geq \frac{1-2\mu}{2(1-\mu)^2}$, while uniquely maximized at $\nu_2^O = \frac{1-\sqrt{1-2\kappa}}{2\kappa}$ if $\kappa < \frac{1-2\mu}{2(1-\mu)^2}$. Therefore, the policies proposed in the proposition is the unique optimal attention outcome for each range of κ. The monotonicity of the informativeness of $p^*(\kappa)$ in each range follows directly from condition (3.3) and Lemma 14. □