Signaling with Commitment

Raphael Boleslavsky Mehdi Shadmehr*

April 10, 2024

Abstract

We study the canonical signaling game, endowing the sender with commitment power: before learning the state, sender designs a strategy, which maps the state into a probability distribution over actions. We provide a geometric characterization of the sender's attainable payoffs, described by the topological join of the graphs of the interim payoff functions associated with different sender actions. We extend the sender's commitment power to the design of a communication protocol, characterizing whether and how sender benefits from revealing information about the state, beyond what is inferred from his action. We apply our results to the design of adjudication procedures, rating or grading systems, and sequencing algorithms for online platforms.

Keywords: Signaling, Information Design, Bayesian Persuasion, Belief-Based Approach

JEL: C72, D82, D83, G24, K4, L81, M3

*Boleslavsky: Department of Business Economics & Public Policy, Indiana University (e-mail: rabole@iu.edu); Shadmehr: Department of Public Policy, University of North Carolina at Chapel Hill (e-mail: mshadmeh@gmail.com). We thank Nageeb Ali, Nemanja Antic, Heski Bar-Isaac, Mike Baye, Dan Bernhardt, Ben Brooks, Steven Callander, Rick Harbaugh, Alexander Jakobsen, Navin Kartik, John Maxwell, Paula Onuchic, Marilyn Pease, Luciano Pomatto, Doron Ravid, Santanu Roy, Fedor Sandomirskiy, Mark Whitmeyer, Kun Zhang, and seminar audiences at the Kelley School of Business, New York University (Stern), SAET 2023, SEA 2023, and Southern Methodist University.

1 INTRODUCTION

Signaling games, including Spence signaling (Spence 1973b), reputation games (Kreps and Wilson 1982), communication with lying costs (Kartik 2009), and burning money (Austen-Smith and Banks 2000) arise naturally in various strategic environments, including labor economics, industrial organization (Nelson 1974, Milgrom and Roberts 1986, Bagwell and Riordan 1991), finance (Leland and Pyle 1977, Ross 1977, DeMarzo and Duffie 1999, DeMarzo 2005), macroeconomics (Angeletos et al. 2006), organizational economics (Bar-Isaac and Deb 2020), political economy (Banks 1991), and cultural economics (Spence 1973a, Camerer 1988, Bénabou and Tirole 2006). We study a standard class of signaling games, with the novel feature that the sender can commit to his strategy before learning the state. Depending on the application, such commitment power could arise from the design of institutions, formal contracts, reputation incentives, algorithms, or AI.[1] We call this strategic environment *signaling with commitment*.

In particular, consider the standard signaling game described in Fudenberg and Tirole (1991, pp. 324-5). There is a sender (leader) and a receiver (follower). There is a state of the world $\omega \in \Omega$, and a common prior, $\omega \sim \mu_0$ with full support. The sender observes the state of the world and chooses an action $s \in S$. The receiver observes the sender's action s and chooses an action $a \in A$. The players' payoffs depend on the state and actions: let sender's payoff be $v(s, a, \omega)$ and receiver's be $u(s, a, \omega)$. The sender's strategy Π is a set of probability distributions $\pi(\cdot|\omega)$ over actions $s \in S$, conditional on state $\omega \in \Omega$. Receiver's strategy is also a set of probability distributions $\sigma(\cdot|s)$ over actions $a \in A$, conditional on sender action $s \in S$. We build on the standard game by allowing sender to commit to his strategy Π before the state is realized.

In this paper, we study signaling with commitment using the "belief-based" approach. In particular, we do not study the sender's strategy directly; instead, we focus on the receiver's posterior beliefs, which are generated by Bayes' rule applied to the sender's strategy. In other words, we treat the sender's strategy $\Pi = \{\pi(s|\omega) \,|\, s \in S, \, \omega \in \Omega\}$ as a signal structure

[1] A number of papers provide justification for the assumption that sender can commit to a communication strategy in the Bayesian Persuasion framework, many of which have natural analogs in our setting. Best and Quigley (2020) derive conditions under which the desire to maintain influence in future interactions can generate commitment power for the sender. Deb et al. (2022) study how a formal contract allows sender to gain commitment power in the production of information. Moreover, as Shadmehr and Boleslavsky (2015) and Meng (2020) argue, by choosing the personnel composition of institutions, even autocrats can commit to different policies.

or statistical experiment, which conveys information about the state ω to the receiver. The sender's realized action s is treated as the realization of the signal or the outcome of the experiment. Thus, the problem of finding the sender's optimal strategy in a signaling game with commitment is replaced with the problem of designing the sender's optimal statistical experiment. Because the realization of the statistical experiment is his action, the sender's problem has two novel features relative to the information design literature (Kamenica and Gentzkow 2011, Bergemann and Morris 2019, Kamenica 2019, Taneva 2019). First, the signal realization is directly payoff relevant, potentially for both sender and receiver. Second, the sender does not have the freedom to create new signal realizations—signal realizations are actions and must be elements of S.

We provide a geometric characterization of the sender's optimal strategy and payoff. Applying the belief-based approach, we graph the sender's payoff in the space of posterior beliefs. Because the sender's payoff depends directly on the realization of the signal (his action) in addition to the posterior belief it induces, in the space of beliefs, the sender has a separate payoff function for each signal realization $s \in S$. We show that the set of feasible payoffs for the sender can be characterized geometrically as the *topological join*, "join" henceforth, of the graphs of these payoff functions.[2] Though it can be represented as a particular convex combination of the sender's payoffs, in general the join is a subset of the convex hull of the sender's payoffs, is not a convex set, and its upper boundary cannot be found via concavification, contrasting with information design.

In signaling with commitment, the receiver's only source of information about the state is the sender's action. For example, an institution that commits to procedures for adjudicating internal grievances may be barred from releasing additional information about individual cases. A financial firm can commit to a trading strategy via an algorithm or AI, but its public statements may be restricted or regulated. A university can commit to a grading policy, which allows it to reveals information about student ability, but a commitment to reveal information through other channels may not be plausible. Nevertheless, in some settings, the sender may be able to provide the receiver with additional information, beyond what is inferred from his action. In such settings, understanding the positive and normative implications of such additional communication is important (see Section 3.2, where we analyze online sales platforms).

[2]If A and B are subsets of $\mathbb{R}^n$, then the join of A and B is the set of all line-segments connecting a point in set A to a point in set B (Rourke and Sanderson 1982).

To address this possibility, we study a benchmark with "extended commitment," in which sender can provide the receiver with additional information. In particular, sender designs an arbitrary message space M, and commits to a joint distribution of message and action $(m, s) \in M \times S$, conditional on the state. Thus, sender can transmit information using the action, message, or any combination of the two. We show that in the extended commitment benchmark, sender can achieve any payoff in the convex hull of the union of the payoff graphs, and the sender's optimal payoff lies on its upper boundary. In order to achieve this payoff, commitment to actions is essential—in general, it cannot be achieved by designing a communication protocol alone. Furthermore, we characterize settings in which extended commitment results in the same payoff as signaling with commitment. In such settings, the ability to transmit additional information does not increase the sender's payoff.

Throughout the paper, we illustrate our findings in an example. An organization designs an adjudication procedure that maps valid (invalid) grievances into a probability distribution over resolutions, either dismissal or a costly remedy. As mentioned above, confidentiality restrictions prevent the organization from revealing information about the merits of a grievance, beyond what is inferred from the resolution. Outside stakeholders (donors, advertisers, advocacy groups) observe the adjudication procedure and the resolution, and then decide whether to retaliate against the organization (e.g., withhold donations or advertising, stage protests). Stakeholders retaliate if they believe that a grievance has been mishandled, but their power is limited: the organization would rather dismiss and incur retaliation than implement a remedy. If the organization cannot commit to a procedure, then it dismisses all grievances and the stakeholders always retaliate. We show that if the organization can commit to its procedures, then it continues to dismiss all grievances when the prior belief is high or low, but for intermediate priors, it remedies valid grievances with positive probability. Removing confidentiality restrictions allows the organization to transmit additional information about the merits of a grievance. With this ability, the organization again dismisses all cases, suggesting a rationale for confidentiality restrictions beyond privacy concerns.

We study two applications in more detail. First, we consider the design of a rating system, where the designer's payoff depends directly on the nominal rating, beyond the information it provides. This direct preference for nominal ratings introduces a signaling concern which distinguishes the problem from standard information design: the designer's payoff includes terms resembling Spence signaling (Spence 1973a) and "burned money" (Austen-Smith and

Banks 2000, Kartik 2007). A preference for nominal ratings arises naturally in applications. For example, when designing grading policies, schools may internalize professors' distaste for inflating grades of undeserving students. Schools may also internalize the cost of student effort or the value of additional skills that must be acquired to achieve a high nominal grade (Onuchic and Ray 2023). Certification intermediaries may have a career concern which rewards favorable ratings. The possibility of naive receivers, who interpret ratings according to their nominal meanings, also creates a preference for nominal ratings, as in Kartik et al. (2007) and Inderst and Ottaviani (2012a,b). The existence of such receivers has been documented empirically in financial and labor markets (De Franco et al. 2007, Malmendier and Shanthikumar 2007, Hansen et al. 2023).

To fix ideas, we focus on an environment in which an analyst designs a rating policy for a financial asset, with the goal of increasing investment. Sophisticated investors observe both the rating policy and the realized rating when updating their beliefs; naive investors update as if the rating policy is honest. Thus, a favorable rating leads to greater investment than is warranted given its true information content, and an unfavorable rating leads to less. In addition, we assume that exaggeration is costly: assigning the high rating to a bad asset may open the analyst to legal sanction or create psychological distress. The analyst's optimal rating policy with commitment differs significantly from the ratings that arise without commitment. Showing that the optimal rating policy depends on the global shape of the analyst's payoff, we characterize conditions under which the optimal rating *understates* the asset value, becomes *more* informative when there are more naive investors, and *inverts* the ratings' nominal meanings—a surprising result that is consistent with some puzzling empirical findings (De Franco et al. 2007).

Our second application is motivated by the idea that online "platforms often have a financial incentive to steer customers to particularly profitable products and can use the power of defaults and ordering to accomplish that effectively" (Scott Morton et al. 2019, p.51, quoted in Bar-Isaac and Shelegia (2022)). We develop a model in which a platform designs a sequencing algorithm that determines the ordering of products in a consumer's search results to maximize its expected commission. Consumers must pay a privately known cost to learn about or purchase a product that is lower in the sequence, which allows the platform to steer consumers toward products that are displayed more prominently (Dinerstein et al. 2018, Ursu 2018, Bar-Isaac and Shelegia 2022). Because the platform has access to extensive information about tastes and product characteristics, it knows the consumer's match-quality for each product, and this information is exploited by the sequencing algorithm. Thus, consumers learn about match-quality from a product's position in the sequence, as in Janssen et al. (2023) and Kaye (2024).[3]

We characterize conditions under which the platform chooses a positive (negative) sequencing algorithm, in which a more prominent position conveys good news (bad news) about match-quality. Positive sequencing leverages the cost to deter search, steering all consumers toward the profitable product. In contrast, negative sequencing steers high cost consumers toward the profitable product, but incentivizes search among low cost consumers. Consequently, an increase in search costs affects the algorithms differently: positive sorting becomes more profitable, negative sorting less profitable. With "extended commitment," the platform can commit to a recommendation system, which transmits (free) personalized product recommendations to consumers. In this setting, the platform always displays the more profitable product first, and uses the custom recommendations to reveal information, which influences the consumer's search and purchase decision. Critically, we show that the consumer's expected welfare is lower with the recommendation system than with the optimal sequencing algorithm. In other words, constraining the platform to reveal information with the product sequence rather than personalized reviews increases the consumer's expected welfare.

Our paper connects to the growing literature on Bayesian persuasion and information design, which can be interpreted equivalently as cheap talk with sender commitment, the design of statistical experiments, or as disclosure rules Kamenica (2019). In Kamenica and Gentzkow (2011), sender and receiver play a cheap talk game, and sender commits to his message strategy—or equivalently, to a statistical experiment that generates information about the state. They show that this problem can be analyzed using the belief-based approach, whereby a message strategy is identified with the distribution of posterior beliefs that it induces. In this characterization, the sender's optimal payoff is the concave envelope of the sender's payoff graph in the space of posterior beliefs. The belief-based approach has been used by Alonso and Câmara (2016) to study political persuasion, by Boleslavsky and Cotton (2015, 2018) and Au and Kawai (2020) to study competitive information production, and by Zhang and Zhou (2016) to study contest design. Aumann and Hart (2003) and Lipnowski and Ravid (2020) use this approach to study cheap talk without commitment. A number of

[3]Kaye (2024) focuses on estimating the consequences of a given algorithm (Expedia's). In Janssen et al. (2023), there is no steering, as product position does not directly affect search costs.

papers use this approach to study costly information production or rational inattention (Sims 2003, Gentzkow and Kamenica 2014, Caplin and Dean 2015, Bloedel and Segal 2021). In this literature, the cost of an information structure is posterior-separable (Denti 2022)—it can be expressed as an expectation of a function of the realized posterior. Pomatto et al. (2023) provide an axiomatic foundation for a class of such cost functionals. In contrast to these literatures, signaling with commitment can be viewed as an information design problem in which the realization of the experiment is directly payoff-relevant, beyond its effect on beliefs.

1.1 AN EXAMPLE

To communicate the key ideas, we begin with an example. An organization (sender, he) designs procedures or formal rules for addressing grievances. A grievance is either valid (e.g., a true violation of Title IX), which we denote $\omega = v$, or invalid, $\omega = f$ (e.g., a mistaken or false accusation). Both types of grievance arise exogenously within the organization at different rates. Consequently, a new grievance is believed to be invalid with probability μ_0. After learning the details of the grievance and determining its type, the organization has two available actions, $S = \{a, d\}$. By selecting a, the organization addresses the grievance and implements a costly remedy; by selecting d, it dismisses the grievance without redress. In either case, stakeholders—including potential donors, customers, advertisers, partners, or advocacy groups—observe the organization's decision and update their beliefs, $\mu_s = \Pr(\omega = f|s)$ for $s \in \{a, d\}$. Moreover, confidentiality restrictions prohibit the organization from communicating about the details of individual grievance cases. Thus, the organization's action is the stakeholders' only signal about the merits of the case. The stakeholders, as a unitary actor (receiver, she), then decide whether to retaliate against the organization, for example by withholding donations, boycotting, dissolving partnership, or protesting. In particular, stakeholders retaliate when they believe that the organization is likely to have mishandled the grievance. Thus, if the organization dismisses the grievance, stakeholders retaliate when they believe that the grievance is likely valid, $\mu_d < \theta_d$; if the organization addresses the grievance, then the stakeholders retaliate when they believe the case is likely to be invalid, $\mu_a > \theta_a$ (where θ_a, θ_d are exogenous). For concreteness, we focus on $\theta_a < \theta_d$, but the following analysis holds for all θ_a, including $\theta_a > 1$.

Both redress and retaliation are costly for the organization. The organization's payoff from redress ($s = a$) is 0, and its payoff from dismissal ($s = d$) is 1, provided that the stake-

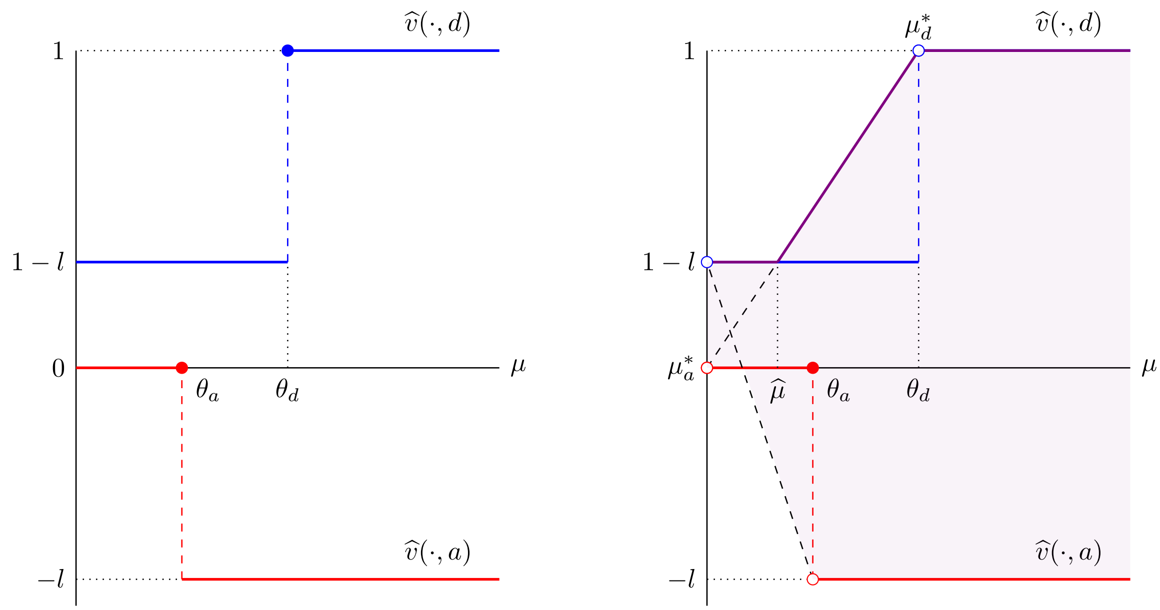

Figure 1: Motivating Example.

holders do not retaliate. Retaliation by stakeholders imposes a cost of $l \in (0,1)$ on the organization. With $l < 1$, stakeholders have limited sway over the organization: the organization would rather dismiss a grievance and face the stakeholders' punishment than address it.

The organization designs and commits to formal procedures for addressing grievances. These procedures determine the probability $\pi(s|\omega)$ that action $s \in \{a,d\}$ is taken in state $\omega \in \{v,f\}$. Consequently, they determine the stakeholders' posterior belief induced by each realized action, $\mu_s \equiv \Pr(\omega = f|s)$, $s \in \{a,d\}$, and the unconditional probability $\tau(\mu_s)$ of each posterior belief.[4] The law of iterated expectations implies $E_\tau[\mu_s] = \mu_0$. Furthermore, any combination of beliefs $\{\mu_a, \mu_d\}$ and probabilities $\{\tau(\mu_a), \tau(\mu_d)\}$ that satisfies the law of iterated expectations arises from some adjudication procedure $\pi(\cdot|\cdot)$.

Building on these observations, the organization's optimal commitment can be analyzed graphically using Figure 1. On the horizontal axis, we have μ, the stakeholder's posterior belief that the grievance is invalid after observing the organization's action. The blue step function is the graph of $\widehat{v}(\mu,d) = 1-\mathcal{I}(\mu < \theta_d)l$, which is the organization's expected payoff if it dismisses the case ($s = d$) and the stakeholder's updated belief is μ.[5] Similarly, the red step function, $\widehat{v}(\mu,a) = -\mathcal{I}(\mu > \theta_a)l$, is the organization's expected payoff, if it addresses

[4] In particular, $\tau(\mu_s) \equiv \mu_0\pi(s|f) + (1-\mu_0)\pi(s|v)$ and $\mu_s = \Pr(\omega = f|s) = \mu_0\pi(s|f)/\tau(\mu_s)$.

[5] The organization expects 1 if the stakeholders do not retaliate, (i.e., if $\mu \geq \theta_d$) and $1-l$ otherwise.

the case ($s = a$) and the stakeholder's updated belief is μ.[6]

As described above, the organization's strategy induces a distribution over posteriors $\{\mu_a, \mu_d\}$, with a mean equal to the prior μ_0. The organization's ex ante expected payoff of such a strategy, is then the expectation of $\widehat{v}(\mu_s, s)$ according to the same distribution over posteriors. Graphically, this payoff can be found by drawing a line segment connecting the point on the red curve above or below μ_a with the point on the blue curve above or below μ_d, and then evaluating the height of the line segment at μ_0. We then vary the organization's strategy over all possible $\{\mu_a, \mu_d\}$ to increase this height as much as possible, as illustrated in the right panel. Carrying out this procedure, we find that the largest payoff that the organization can attain with prior μ_0 is the height of the purple curve above μ_0. Furthermore, at any prior the shaded purple region represents all payoffs that the organization could attain by varying its procedures.

As Figure 1 illustrates, when the prior is high or low, $\mu_0 \geq \theta_d$ or $\mu_0 \leq \widehat{\mu}$, the organization always dismisses all grievances. In contrast, for moderate priors, $\mu_0 \in (\widehat{\mu}, \theta_d)$, the organization sometimes commits to address a valid grievance with positive probability, despite the fact that the worst payoff from dismissing a grievance is larger than the best payoff from addressing it—the highest value on the red curve is strictly below the lowest value on the blue one. This is evident in Figure 1, where moderate priors are optimally spread into posteriors $\{\mu_a^* = 0, \mu_d^* = \theta_d\}$.

Intuitively, if valid grievances are too unlikely under the prior, $\mu_0 \geq \theta_d$, then the organization does not need to reveal information in order to deter retaliation, and thus dismissing all grievances is optimal. In contrast, for $\mu_0 \leq \theta_d$, the organization has an incentive to reveal information in order to deter retaliation. However, in order to increase μ_d, the organization must commit to address valid grievances with positive probability. Because addressing a grievance is costly, deterring retaliation in this way is only worthwhile if valid grievances are relatively unlikely, $\mu_0 \in (\widehat{\mu}, \theta_d)$. If valid grievances are too likely under the prior, $\mu_0 < \widehat{\mu}$, it is better in expectation to dismiss all grievances and incur the wrath of stakeholders rather than paying the costs of redress.

In this example, the organization's ability to commit to its procedures expands stakeholders influence over the organization. Without commitment the organization dismisses all grievances (recall $1 - l > 0$). With commitment, however, valid grievances are addressed

[6] Here, the organization expects 0 if there is no retaliation, $\mu \leq \theta_a$ and $-l$ otherwise.

with positive probability at moderate priors. For this to occur, however, it is essential that stakeholders have at least some power initially: if $l = 0$, then the organization dismisses all grievances, even if it has commitment power. If retaliation does not harm the organization, then there is no reason deter it, and thus no reason for the organization to spread beliefs.

One key feature of this characterization is worth highlighting: the upper envelope representing the largest attainable payoff (in purple) is *not* the concave envelope of the organization's payoffs, as it would be in Bayesian Persuasion—indeed it is not even concave. In this example, the concave envelope connects two points on the blue curve (evidently $(0, 0)$ and $(\theta_d, 1)$). Because the only source of information for the stakeholders is the organization's action, and different actions correspond to different payoff functions, when constructing the line segments that allow us to evaluate the organization's payoff we can only connect a point on the blue curve to a point on the red curve. If a payoff constructed by connecting two points on the blue curve were feasible, then dismissing a grievance would generate two different posterior beliefs for the stakeholder. This is possible only if some additional source of information exists, beyond the organization's action. We consider an extension with this feature in Section 2.1. As we will see, when the stakeholders have limited sway ($l < 1$), relaxing privacy restrictions and allowing the organization to provide additional information to stakeholders leads to the dismissal of all grievances.

2 THEORY

Preliminaries. To keep the exposition as straightforward as possible, we focus on a finite state space $|\Omega| < \infty$ and sender action space $|S| < \infty$. Other than existence of a best response in the receiver problem, we make no assumptions about the receiver's action space or payoff. A belief $\mu \in \Delta(\Omega)$ is a probability distribution over states, and $\mu(\omega)$ is the probability of state ω under belief μ. By $\tau \in \Delta(\Delta(\Omega))$ we denote a probability distribution over beliefs, and $\tau(\mu)$ is the probability of belief μ in distribution τ. Players have a common prior μ_0 with full support on Ω.

Belief Systems and Sender Strategies. A belief system $\mathcal{B} \equiv \{\mu_s, \tau(\mu_s)\}_{s \in S}$, is a set of beliefs and associated probabilities, $\tau(\mu_s)$, indexed by the sender's set of available actions. The probabilities in a belief systems must be exhaustive, i.e., $\sum_{s \in S} \tau(\mu_s) = 1$, but $\tau(\mu_s) = 0$ is allowed. Following Kamenica and Gentzkow (2011), we say that sender strategy

$\Pi = \{\pi(s|\omega)\}_{(s,\omega)\in S\times\Omega}$ induces belief system $\mathcal{B}$ if and only if

$$\tau(\mu_s) = \sum_{\omega\in\Omega} \mu_0(\omega)\pi(s|\omega), \qquad \mu_s(\omega) = \frac{\mu_0(\omega)\pi(s|\omega)}{\tau(\mu_s)}.$$

If strategy Π induces belief system $\mathcal{B}$, then μ_s is the posterior belief associated with sender's action $s \in S$, and the probability $\tau(\mu_s)$ is the probability that action s is selected under the sender's strategy. In other words, $\tau(\mu_s)$ is the probability that the realization of the posterior belief is μ_s.

Remark 1 *If certain actions in S are never chosen, i.e. $\tau(\mu_s) = 0$, then the associated posterior is undefined. In this case, we allow μ_s to take on any value. Such realizations of the posterior have no effect on the sender's expected payoff, and, given the sender's commitment power, they play no role in the analysis.*

As is well-known, a belief system induced by a sender strategy must be Bayes-Plausible: $E_\tau[\mu] = \sum_{s\in S} \tau(\mu_s)\mu_s = \mu_0$. Following the argument of Kamenica and Gentzkow (2011), it is straightforward to show that any Bayes-Plausible belief system can be induced by a corresponding sender strategy. Therefore, rather than designing his or her strategy, we can consider the sender designing a belief system directly, with the constraint that it is Bayes-Plausible.

Proposition 1 *A belief system $\mathcal{B}$ is Bayes-Plausible if and only if it is induced by some sender strategy Π.*

Unlike Bayesian Persuasion, where the sender is able to design the message space as part of the signal structure, in signaling with commitment, the sender's messages must belong to S, which restricts the number of messages available.

Remark 2 *For Proposition 1, it is important that all sender actions $s \in S$ are available in all states, $\omega \in \Omega$. If sender cannot select certain actions in certain states, then some Bayes-Plausible belief systems cannot be induced by a sender strategy. While important to the analysis, the assumption that all actions are available in all states is without loss of generality. To analyze a setting in which action s is unavailable in ω, one simply imposes a large cost on the sender for such a choice. For a given prior distribution μ_0, a sufficiently large penalty ensures that selecting s in state ω with positive probability is suboptimal.*

Receiver's Response. The receiver selects his or her action after observing the sender's action s, which is associated with posterior belief μ_s. The receiver's optimal action solves

$$\widehat{a}(\mu_s, s) = \arg\max_{a \in A} E_{\mu_s}[u(s, a, \omega)].$$

If the receiver's problem has multiple solutions (over which the receiver must be indifferent), we focus on an equilibrium in which the sender's preferred action is chosen.[7]

Sender's Problem. Equilibrium with sender commitment requires that the belief system (and associated sender strategy) designed by the sender maximize the sender's ex ante payoff. To formulate the sender's problem, first consider the sender's utility at the interim stage, after his action has been realized, but before the receiver responds. At this stage, μ_s is the belief about the state and $\widehat{a}(\mu_s, s)$ is the receiver's response. The sender's interim utility is

$$\widehat{v}(\mu_s, s) = E_{\mu_s}[v(s, \widehat{a}(\mu_s, s), \omega)].$$

From this expression, we see that the signal realization (sender action) affects the sender's interim payoff in two ways, beyond its effect on the posterior belief. Keeping the posterior belief μ_s constant, changes in s can affect the interim utility directly through the sender's payoff (reflected in the first argument of $v(\cdot)$), or through its effect on the receiver's response (reflected in the second argument of $\widehat{a}(\cdot, \cdot)$).

The sender's problem is therefore to design a belief system, $\mathcal{B} = \{\mu_s, \tau(\mu_s)\}_{s \in S}$ in order to maximize his or her ex ante expected payoff, subject to the constraint that the belief system is Bayes-Plausible,

$$\text{(Sender's Problem)} \quad \max_{\mathcal{B}} E_\tau[\hat{v}(\mu_s, s)] \quad \text{subject to} \quad \sum_{s \in S} \mu_s \tau(\mu_s) = \mu_0.$$

The geometric characterization of the solution is based on the *topological join* ("join") of sets (Rourke and Sanderson 1982).

[7] In a generic environment, the sender can break the receiver's indifference between actions with a marginal adjustment in his strategy in such a way that receiver always selects sender's preferred action. Thus, equilibrium requires that receiver selects sender's preferred action when indifferent, lest the sender deviates. However in a non-generic setting, it may not be possible for sender to break the receiver's indifference in this manner.

Definition 1 *(Join of Sets). If $X_i \subset \mathbb{R}^n$ for $i = 1, ..., k$, then their join is*

$$join(X_1, ..., X_k) \equiv \Big\{ \sum_{i=1}^{k} \lambda_i x_i \,|\, x_i \in X_i,\ \lambda_i \geq 0,\ \sum_{i=1}^{k} \lambda_i = 1 \Big\}.$$

The join of the sets $(X_1, ..., X_k)$ is generated by forming all possible convex combinations, using at most one point from each set X_i. For two sets (X_1, X_2), it is the union of the sets themselves, along with all line segments connecting a point in X_1 to a point in X_2. With more than two sets, we can imagine first joining X_1 and X_2, then joining X_3 to the resulting set, and so on.[8] In Figure 1, the join of the graphs of the sender's payoff functions $\widehat{v}(\cdot, a)$ and $\widehat{v}(\cdot, d)$ is shaded purple.

In general, the $join(X_1, ..., X_k)$ is a subset of the convex hull of the union $\bigcup_{i=1}^{k} X_k$, which we denote by $con(X_1, ..., X_k)$. Indeed, the convex hull is formed by taking all possible convex combinations of points in $\bigcup_{i=1}^{k} X_i$, without the constraint that the convex combination includes at most one point from each set. These two coincide in the special case where the convex hull happens to satisfy this additional constraint. Furthermore, unlike the convex hull, which is a convex set by definition, the join is not necessarily a convex set. In Figure 1, the join of the graphs of the sender's payoffs is shaded purple. It is a strict subset of the convex hull of the union of the graphs, and it is not a convex set. We return to the connection between the convex hull of the union and the join in Section 2.1.

Definition 2 *(Join Envelope). If $X_i \subset \Delta(\Omega) \times \mathbb{R}$ for $i = 1, ..., k$, then their join envelope is a function of $\mu \in \Delta(\Omega)$*

$$V^{jo}(\mu | X_1, ..., X_k) \equiv \sup \Big\{ z \,|\, (\mu, z) \in join(X_1, ..., X_k) \Big\}.$$

The join envelope is the supremum of the join, in much the same way as the concave envelope, $V^{co}(\mu | X_1, ..., X_k)$, is the supremum of the convex hull. Although formally the join envelope and concave envelopes are defined as the supremum, in the context that we use them in this paper, the maximum is attained in the set.[9] When referring to the envelopes, we omit the sets from the notation where it does not create confusion. In Figure 1, the purple curve is the join envelope of the graphs of the expert's payoff function.

[8] The join is obviously associative and commutative, and thus the order is irrelevant.

[9] This is a consequence of our assumption that when the receiver has multiple best responses, the sender's preferred action is selected. In other words, $(\mu, V^{jo}(\mu | X_1, ..., X_k)) \in join(X_1, ..., X_k)$.

Definition 3 *(Interim Payoff Graphs). Sender's interim payoff graph for action s is the following subset of* $\Delta(\Omega) \times R$, $\widehat{v}_s \equiv \{(\mu, \widehat{v}(\mu, s)) \,|\, \mu \in \Delta(\Omega)\}$.

Figure 1 hints at the connection between the sender's problem, the interim payoff graphs, their join, and the join envelope, which we make precise in the following proposition.

Proposition 2 *(Signaling With Commitment). In sender's problem,*

(i) sender can attain payoff z *if and only if* $(\mu_0, z) \in join(\widehat{v}_s)_{s\in S}$.

(ii) for $(\mu_0, z) \in join(\widehat{v}_s)_{s\in S}$, *some* $\{\lambda_s\}_{s\in S}$ *exist such that* $\lambda_s \geq 0$, $\sum_{s\in S} \lambda_s = 1$, *and* $(\mu_0, z) = \sum_{s\in S} \lambda_s(\mu_s, \widehat{v}(\mu_s, s))$. *Sender attains* z *with belief system* $\{\mu_s, \tau(\mu_s) = \lambda_s\}_{s\in S}$.

(iii) sender's largest attainable payoff with prior μ_0 *is* $V^{jo}(\mu_0|(\widehat{v}_s)_{s\in S})$.

To simplify notation, we abbreviate $V^{jo}(\mu_0|(\widehat{v}_s)_{s\in S})$ to $V^{jo}(\mu_0)$.

To this point, we have focused on the "literal interpretation" of signaling with commitment—sender designs and commits to a strategy to maximize his expected payoff. We can also build on Proposition 2 to present a "metaphorical interpretation," in the spirit of Bergemann and Morris (2019). Imagine a signaling game in which sender may be a commitment type. With probability p, sender's action is the realization of some *exogenous* "commitment strategy," $\pi_C(\cdot|\omega)$. With complementary probability, the action is freely chosen by the sender after observing the state. Sender and receiver know the probability of the commitment type and the commitment strategy, but receiver cannot determine whether the sender's realized action was drawn from the commitment strategy or was sender's free choice. Imagine that an analyst observes this signaling game and is aware that sender may be a commitment type. However, the analyst cannot directly observe the probability p, or the commitment strategy. The analyst would like to determine the set of sender payoffs (beliefs, and strategies) that can arise in Perfect Bayesian Equilibria across all possible values of these unobservables. Thus, under the "metaphorical interpretation," a third-party introduces a commitment type into the game in pursuit of some unknown objective, and the analyst would like to provide a characterization of the Perfect Bayesian Equilibria that is robust to this possibility. With the advent of algorithms that offer product advice, trade in markets, or otherwise engage in signaling activity, such an exercise is particularly relevant.

A straightforward extension of Proposition 2 provides an answer. First, note that any equilibrium of the underlying signaling game with limited commitment induces some belief system. From Proposition 2, the equilibrium belief system cannot achieve a higher payoff than $V^{jo}(\cdot)$ at any prior belief. Furthermore, if the sender is a commitment type who plays the sender-optimal strategy with probability 1, then the beliefs, receiver responses, and sender's (commitment) strategy constitute a PBE. Thus, the upper bound of the set of attainable payoffs is the join envelope, as constructed in Proposition 2. To construct *all* of the achievable payoffs, we would need to redo the construction of the join, modifying the receiver's tie-breaking rule. In particular, if we wish to characterize all equilibria, then we can no longer impose that an indifferent receiver selects the sender's preferred action. Rather, we must include all possible best responses for receiver, any of which could be played in some equilibrium of the underlying game.[10] This modification transforms senders payoff graph into a correspondence (which complicates the notation), but it does not affect the essence of the construction or the characterization. Extending our previous argument, at prior μ_0, sender achieves payoff z in a PBE of the game for some commitment type $(p, \pi_C(\cdot|\cdot))$ if and only if (μ_0, z) belongs to the join of the sender's interim payoff graphs, modified to account for the full set of receiver best responses. Furthermore, based on Figure 1 (and others in the paper), we conjecture that with standard conditions on payoffs, the modified join is the closure of the original join, which was constructed under the sender-preferred tie-breaking rule.

Beer-Quiche Game. The findings of this section can be illustrated in the classic beer-quiche game. The tough sender's payoff is 1 if he chooses B (beer) and 0 if he chooses Q (quiche). The wimpy sender obtains a payoff of 1 from consuming quiche and a payoff of 0 from consuming beer. Furthermore, wimpy sender's payoff is reduced by $c \in (0, 1)$ if the receiver selects F (fight). Receiver's payoff is 0 if she chooses A (leave alone), is k if she chooses F and sender is wimpy, and is $-(1-k)$ if she chooses F and sender is tough.

Without sender commitment, only a separating strategy is sequentially rational for sender. The tough type always selects B which delivers 1 instead of 0 regardless of receiver's response; the wimpy type always selects Q since the worst payoff following this choice, $1-c$, is larger than the best payoff following B, which is 0.

To study the equilibrium with commitment, we first construct sender's interim payoffs.

[10] Suppose that under the prior belief, receiver is indifferent between actions a_1 and a_2 when sender chooses a particular action s. Suppose further that sender is known to be a commitment type ($p = 1$) who selects action s with probability 1 in all states. In this case all randomizations between a_1 and a_2 are optimal for the receiver, and any such randomization can arise in a PBE.

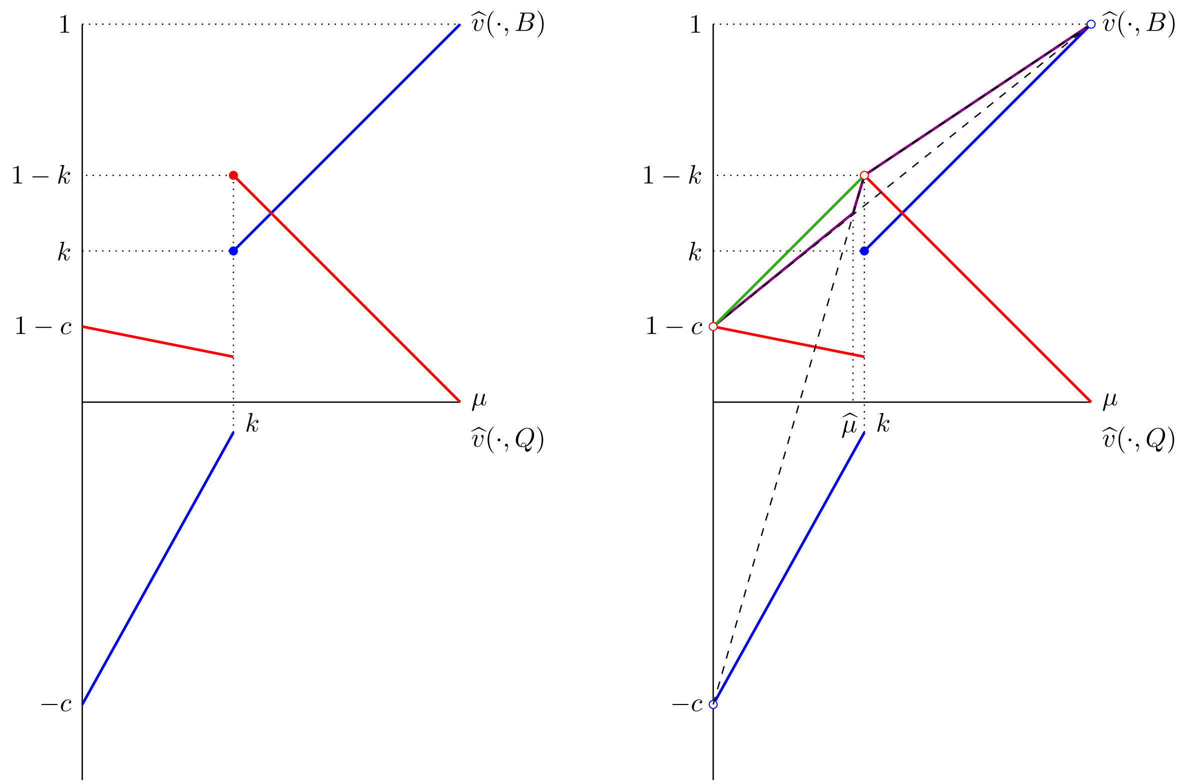

Figure 2: Beer-Quiche Game.

Let μ be the belief that the sender is tough. Receiver's expected payoff of F at belief μ is $(1-\mu)k - \mu(1-k) = k - \mu$, and thus receiver's optimal response is $\widehat{a}(\mu_s, s) = F$ if and only if $\mu < k$. Sender's interim payoffs are therefore,

$$\widehat{v}(\mu, B) = \mu - c(1-\mu)\mathcal{I}(\mu < k) \qquad \widehat{v}(\mu, Q) = (1-\mu)(1 - c\mathcal{I}(\mu < k)).$$

The interim payoff graphs (red, blue) and the join envelope (purple) appear in Figure 2 (green will be discussed later), where we have focused on the most interesting settings with $k \in (0, 1/2)$ and $c \in (\underline{c}, 1)$ for $\underline{c} \equiv k/(1-k)$. These conditions ensure that $\widehat{\mu}$ (the crossing between the dashed lines) exists. First, notice that the join envelope is convex on interval $[0, k]$ and thus does not coincide with the concave envelope. Next, note that for $\mu_0 \leq \widehat{\mu}$, sender's separating strategy without commitment ($B \iff$ tough), is also optimal with sender commitment. However, for larger prior beliefs this is no longer the case: for $\mu_0 \in (\widehat{\mu}, k)$ it is optimal for the tough sender to select Q and wimpy sender to mix, while for $\mu_0 \in (k, 1]$ it is optimal for tough sender to mix and wimpy sender to select Q.

2.1 EXTENDED COMMITMENT

In the main model, the sender can transmit information to the receiver only through his action. In the motivating example, the organization is legally barred from communicating with stakeholders about the merits of specific grievances; only the decision to dismiss or address the grievance is observed. To understand the implications of such restrictions, in this section we extend the scope of sender's commitment power: along with his strategy, sender commits to a communication protocol or signal structure that releases public information about the state. In particular, we allow sender to design an arbitrary message space M, and to commit to joint distributions $\pi_E(\cdot, \cdot|\omega)$ over messages and actions $(m, s) \in M \times S$ conditional on state $\omega \in \Omega$. Thus, sender can transmit information using his action, message, or any combination of the two. We refer to this environment as "extended commitment."

We apply the belief-based approach to characterize the set of attainable payoffs. First, note that each public realization (m, s) is associated with a posterior belief $\mu \in \Delta(\Omega)$. Thus, each sender strategy Π_E induces a finite joint distribution τ of the posterior belief and sender action over $G \equiv \Delta(\Omega) \times S$, and all strategies that induce the same joint distribution τ are outcome-equivalent. In the usual way, this joint distribution can be decomposed into a marginal for belief μ, denoted τ_m, and a distribution for the action $s \in S$ conditional on μ, denoted τ_c, i.e., $\tau(\mu, s) = \tau_m(\mu)\tau_c(s|\mu)$. From the law of iterated expectations, the marginal distribution of the posterior belief must be Bayes-Plausible: $\sum \tau_m(\mu)\mu = \mu_0$, where the summation is over the support of τ_m. Thus, Bayes-Plausibility of τ_m is a necessary condition for joint distribution τ to be induced by some sender strategy. Indeed, this is also sufficient, as we show in the following result.

Proposition 3 *(Inducible Joint Distribution). With extended commitment,*

(i) joint distribution τ is induced by some sender strategy if and only if the marginal distribution of the posterior belief τ_m is Bayes-Plausible.

(ii) if joint distribution τ can be induced by some sender strategy, then it can also be induced by a strategy in which observing the sender's action s, does not convey any additional information, beyond what is learned from the sender's message m.

Intuitively, whenever the marginal distribution of the belief, τ_m, is Bayes-Plausible, the joint distribution τ can be induced as the result of a two-step procedure. Design a sender strategy $\pi(\cdot|\omega)$ over some message space M that induces τ_m using messages alone; this can always be done if τ_m is Bayes-Plausible (Kamenica and Gentzkow 2011). In the first step, generate a realization from this strategy. If realized posterior belief μ is drawn in the first step, then select an action according to conditional distribution $\tau_c(\cdot|\mu)$ in the second step. By construction, the joint probability of a belief and action produced in this manner is $\tau_m(\mu)\tau_c(s|\mu) = \tau(\mu, s)$. Furthermore, the distribution of the action depends only on the realized posterior belief in the first stage, which itself depends only on the realized *message*. Conditional on the message, the action conveys no information about the state. In this sense, extended commitment allows information provision to be uncoupled from the choice of action.

Building on Proposition 3, part (i), with extended commitment, sender's problem is to design a joint distribution of belief and actions τ, that maximizes his expected payoff subject to the constraint that the marginal distribution of the belief τ_m is Bayes-Plausible.

$$\max_{\tau\in\Delta(G)} \sum_{(\mu,s)\in \mathrm{sp}[\tau]} \tau(\mu,s)\widehat{v}(\mu,s) \quad \text{subject to} \quad \sum_{\mu\in\mathrm{sp}[\tau_m]} \tau_m(\mu)\mu = \mu_0,$$

where sp[$\cdot$] denotes the support of the distribution.

To solve the sender's problem, note first that any attainable payoff for the sender at prior μ_0 can be written as a convex combination,

$$(\mu_0, V(\mu_0)) = \Big(\sum_{\mu\in\mathrm{sp}[\tau_m]} \tau_m(\mu)\mu, \sum_{(\mu,s)\in\mathrm{sp}[\tau]} \tau(\mu,s)\widehat{v}(\mu,s) \Big) = \sum_{\mu\in\mathrm{sp}[\tau_m]} \tau_m(\mu) \Big(\mu, \sum_{s\in S} \tau_c(s|\mu)\widehat{v}(\mu,s)\Big).$$

Next, note that $\sum_{s\in S}\tau_c(s|\mu)\widehat{v}(\mu,s)$ is a convex combination of the sender's interim payoff functions evaluated at posterior belief μ. Graphically, it is a convex combination of sender's payoff functions in the "vertical dimension," where the belief is fixed and the action s varies. From Proposition 3, sender can freely select the conditional distribution of actions $\tau_c(\cdot|\mu)$ at each posterior belief.[11] Thus, at a given posterior belief μ, varying $\tau_c(\cdot|\mu)$ allows the sender to achieve any payoff between the highest and lowest payoff graphs at that belief. Simultaneously, by varying the Bayes-Plausible distribution of the posteriors τ_m, sender can generate any convex combination of these "vertical payoffs" across posteriors (i.e., "horizontally").

[11]The decomposition of the joint distribution τ into a conditional $\tau_c(\cdot|\mu)$ and a marginal τ_m, and the expression for the sender's payoff in the text also apply in the main model. However, in the main model, sender faces additional constraints on the conditional distributions $\tau_c(\cdot|\mu)$. In particular, in signaling with commitment, each action s in sender's chosen belief system generates a single belief, μ_s. Thus, if $\tau_c(s|\mu) > 0$ for some μ, then $\tau_c(s|\mu') = 0$, for all $\mu' \neq \mu$.

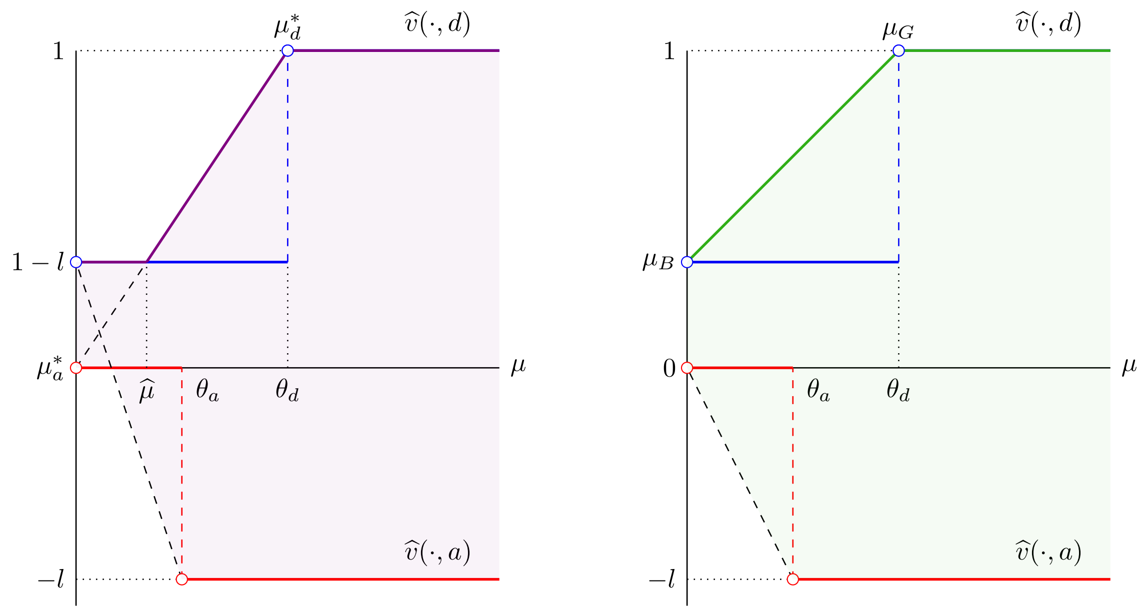

Figure 3: Motivating Example Continued.

By varying the distribution of the posterior τ_m and the conditional distribution of actions $\tau_c(\cdot|\mu)$ together, sender can achieve any payoff that is in the convex hull of the union of the payoff graphs. Obviously, under the sender's optimal strategy, for each realization of the posterior belief, sender selects the action(s) that generate the highest payoff at that belief; we refer to such actions as "belief-optimal."

Definition 4 *(Belief-Optimal Actions). Sender action $s \in S$ is belief-optimal at μ, if and only if $\widehat{v}(\mu, s) \geq \widehat{v}(\mu, s')$ for all $s' \in S$.*

In other words, a belief-optimal action at μ is the best available action for sender, if both he and receiver share posterior belief μ—crucially, a belief-optimal action at μ may not be optimal if sender knows the realized state ω.

We have thus proved the following proposition.

Proposition 4 *(Information and Actions). With extended commitment,*

(i) sender can achieve any payoff inside the convex hull of the union of his payoff graphs, $con((\widehat{v}_s)_{s\in S})$, and his maximum payoff at prior μ is their concave envelope $V^{co}(\mu)$.

(ii) in sender's optimal strategy, if $\tau(\mu, s) > 0$, then s is belief-optimal at μ.

Figure 3 illustrates. The left panel depicts the set of attainable payoffs in the motivating example with commitment to actions only. The right panel illustrates the set of attainable payoffs with extended commitment. From the preceding discussion, we may imagine that the sender designs a Bayes-Plausible distribution of posteriors, τ_m. For each realization of the posterior belief μ in the support, sender then selects the graph $\widehat{v}_s$ at which the payoff is evaluated, thereby assigning an action s to realization μ. By varying the distribution of posterior beliefs and the actions associated with each realization, the sender can achieve any payoff in the convex hull of the union of the sender's payoff graphs, shaded in green. The green curve is the concave envelope of the payoff graphs, evidently higher than the join envelope. While confidentiality restrictions that prevent organizations from communicating information about the details of individual grievances can be justified on the grounds of privacy protections, Figure 3 highlights an additional rationale for such restrictions. As is evident in the right panel, with extended commitment the organization always selects belief-optimal action d. In contrast, with confidentiality restrictions, the organization addresses valid grievances with positive probability at moderate prior beliefs. Thus, removing the restriction may reduce the probability that valid grievances are addressed. This is particularly concerning if the organization does not fully account for the benefits of addressing valid grievances.

Though it is a powerful instrument, in general, commitment to a communication protocol alone does not allow sender to attain the extended commitment payoff, $V^{co}(\cdot)$. By designing a communication protocol ahead of a standard signaling game, sender can generate any Bayes-Plausible distribution of the posterior belief τ_m. To achieve the extended commitment payoff $V^{co}(\cdot)$, however, the sender's action at each realized posterior μ must also be belief-optimal (Proposition 4, (ii)). In other words, sender must select an action that is optimal given only *public information*, summarized by the realized belief μ. However, in a standard signaling game, sender privately observes the *realized state* before selecting an action, and the belief-optimal action may not be optimal in certain states. Unless sender can commit to the belief-optimal action in such states, he cannot achieve $V^{co}(\cdot)$. For a concrete example, return to the Beer-Quiche game, illustrated in Figure 2. Focus on $\mu_0 \in (0, k]$. The concave envelope of the sender's payoff graphs is drawn in green, which evidently lies above the join envelope and interpolates points $(0, 1-c)$ and $(k, 1-k)$, both of which lie on $\widehat{v}(\cdot, Q)$. Thus, with extended commitment, the optimal distribution of posterior beliefs is supported on $\{0, k\}$, and sender always selects action Q. Because tough sender strictly prefers B regardless of the receiver's response, attaining payoff $V^{co}(\cdot)$ requires a commitment from sender to select Q, even when he is tough. By implication, the extended commitment payoff cannot be achieved by a communication protocol alone.

Although sender generally exploits both commitment to actions and the communication protocol in the extended commitment benchmark, in certain environments, sender can achieve the same payoff with commitment to actions only. In other words, given the prior belief and the players' payoff functions, signaling with commitment is enough for sender to achieve the extended commitment payoff. Intuitively, in signaling with commitment, the realized action plays two roles—it transmits information about the state, and it directly affects payoffs. As we have seen, these two roles are generally at odds, and the sender can benefit by uncoupling them. However, in a class of environments which we characterize below, both of these roles can be fulfilled simultaneously by realized actions without any adverse consequences for the sender.

To see how this works, it is helpful to think of sender choosing his strategy in the following way. Imagine that sender uses a public message to generate a Bayes-Plausible posterior belief distribution. Once the belief is realized, sender selects his optimal strategy in the "signaling with commitment" game, with the realized belief μ playing the role of the prior. Consistent with Proposition 2, sender can achieve any payoff in the join of the payoff graphs, and his highest attainable payoff is their join envelope $V^{jo}(\mu)$. Thus, when sender is designing the posterior belief distribution initially, he anticipates payoff function $V^{jo}(\cdot)$. From Kamenica and Gentzkow (2011), by designing a message protocol in the first stage, sender can achieve any payoff in the convex hull of the join, and his highest attainable payoff is the concave envelope of the join itself. In other words, with extended commitment, sender's optimal payoff can be found by concavifying the join of the sender's payoff graphs. We show this formally in Lemma 2 in the Appendix.

This perspective is helpful because it allows us to apply results from Kamenica and Gentzkow (2011) directly to characterize the extended commitment benchmark: Lemma 2 and Remark 1 of Kamenica and Gentzkow (2011) immediately imply part (i) of the following proposition.

Proposition 5 *(Extended Commitment vs. Signaling With Commitment).*

(i) At all prior beliefs, sender's equilibrium payoff in signaling with commitment is identical to his payoff with extended commitment, if and only if the join envelope of sender's payoff graphs is concave on $\Delta(\Omega)$.

(ii) If the sender's interim payoff function $\widehat{v}(\cdot, s)$ is concave on $\Delta(\Omega)$ for all $s \in S$, then the join envelope is concave on $\Delta(\Omega)$.

Part (i) of the Proposition provides a necessary and sufficient condition for the payoffs under extended commitment and signaling with commitment to coincide at *all* prior beliefs. In other words, if this condition holds, then a sender who can commit to actions (signaling with commitment) does not benefit from the ability to commit to a communication protocol (extended commitment), regardless of the prior information about the state. Part (ii) of the proposition provides a straightforward sufficient condition. Although (ii) requires that each interim payoff graph is concave, it does not rule out information provision by the sender, as illustrated in the application to nominal ratings (Proposition 6).

2.2 ADDITIONAL DISCUSSION

We highlight two points before we proceed to applications.

Information Design vs. Extended Commitment In the extended commitment benchmark, we focus on a communication protocol that transmits a single public message, which is observed by receiver along with sender's action. It is straightforward to show that sender cannot do better if he can design the game's information structure. In particular, suppose the sender designs message spaces M_σ, M_ρ and a joint distribution $\pi_I(\cdot, \cdot, \cdot|\omega)$ over messages and sender action $(m_\sigma, m_\rho, s) \in M_\sigma \times M_\rho \times S$, conditional on state ω. Next, the messages and sender action (m_σ, m_ρ, s) are realized. Message m_σ is privately observed by sender, m_ρ is privately observed by receiver, and s is observed publicly. Then receiver responds. With such strong commitment power, sender achieves the highest payoff possible in a signaling game: further increases require changes to the payoff functions, the prior, or the extensive form.

In the Appendix, we show that the sender achieves this payoff in the extended commitment benchmark (see Lemma 3). The idea is familiar: information that does not affect the distribution of player actions in each state also does not affect payoffs. Decompose the joint distribution of private messages and actions; imagine that (m_ρ, s) is drawn first conditional on ω, and then m_σ is drawn conditional on (m_ρ, s, ω), i.e., $\pi_I(m_\sigma, m_\rho, s|\omega) = \pi_{I1}(m_\rho, s|\omega)\pi_{I2}(m_\sigma|m_\rho, s, \omega)$. In this case, the sender's action and receiver's response in each state is determined solely by the first draw. Thus, the realization of sender's private message

m_σ does not affect the distribution of game outcomes in each state or the players' ex ante expected payoffs. By implication, we can assume that sender's message space is identical to receiver's and that their realized messages are perfectly correlated without changing sender's payoff. With these assumptions communication is public, as in the extended commitment benchmark.

Viewed from this perspective, Propositions 3 and 4 imply that in a signaling game, the highest payoff that sender can attain for fixed payoff functions and prior can be achieved using a combination of two familiar instruments, *persuasion* and *delegation*. Imagine that in a standard signaling game, sender has the capability to persuade: before learning the state, sender can design a statistical experiment that generates public information about the state, as in Kamenica and Gentzkow (2011). Given the optimal joint distribution in the extended commitment benchmark, τ, sender can induce the marginal distribution of the posterior τ_m, using the statistical experiment alone. Simultaneously, sender delegates authority over the action to an agent, whose payoffs are identical to his own. Crucially, the agent's only source of information about the state is the statistical experiment designed by the sender, which ensures that the agent chooses the belief-optimal action at each realized posterior. By generating the public belief and the action in this manner, sender can recreate the optimal joint distribution from the benchmark with extended commitment, attaining the same payoff as if he could design the information structure. Furthermore, when the conditions of Proposition 5 hold, the sender can achieve this highest possible payoff *without* designing the information structure and without revealing additional information—signaling with commitment is enough.

Signaling with Transparent Motives. In our motivating example, the sender's payoff is state-independent—the sender's motives are transparent. Lipnowski and Ravid (2020) study cheap talk with transparent motives (without sender commitment) using the belief-based approach, characterizing the set of attainable payoffs and the associated equilibrium beliefs across all equilibria. The key idea is that a Bayes-Plausible belief distribution is induced in some cheap talk equilibrium if and only if sender is indifferent between all beliefs in its support.[12] This result allows for an elegant geometric characterization. Consider the graph of the sender's payoff in the space of posterior beliefs. Select any point on this graph—

[12] If this were not so, then some belief in the support is strictly worse for the sender than some other belief. Because the sender does not have commitment power, she would never transmit the message that induces the worse belief, and the belief would not be in the support.

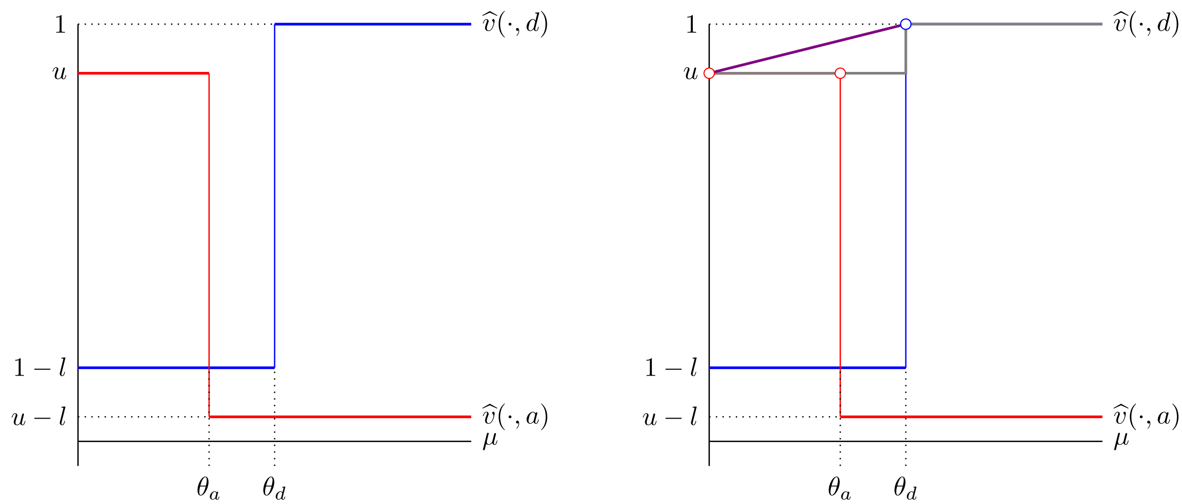

Figure 4: Motivating Example Without Sender Commitment

this payoff could always be attained in a babbling equilibrium. Are there other points on the graph at the same height (i.e., other posterior beliefs which generate the same sender payoff)? If so, connect the points with a line segment. Repeat until no new links are possible. The upper boundary of the resulting object is the *quasi-concave envelope* of the original graph, and it represents the sender's highest attainable payoff in any cheap talk equilibrium. Note that any line segments linking points on the underlying graph must be horizontal or "flat," because only points at the same height (sender value) are linked. This observation highlights the value of commitment power, which relaxes the restriction that the links are flat.

Assuming transparent motives, we can adapt insights and results of Lipnowski and Ravid (2020) to characterize the sender's highest equilibrium payoff in signaling games without commitment. In particular, following similar logic, the sender's maximum payoff in any equilibrium of the signaling game without commitment can be found by connecting the sender's interim payoff graphs using flat segments (as in Lipnowski and Ravid (2020)), but with the additional constraint that the flat segments must begin and end on different graphs (a "flat join"). The upper boundary of this set is the sender's highest possible equilibrium payoff without commitment (the "flat join envelope").[13]

[13] In our setting the characterization of the set of attainable sender payoffs is more demanding than in Lipnowski and Ravid (2020). Because they study cheap talk games, off-path messages are not a concern—setting the receiver's belief following an off-path message to be one of the beliefs inside the support of the equilibrium posterior belief distribution ensures that no off-path message is better for sender than the on-path messages. In contrast, in our setting, ensuring that no off-path action leads to a higher payoff than the on-path actions requires that for each off-path action a "punishing belief" exists at which the off-path action is worse than the on-path actions. This condition always holds for sender-optimal equilibria (on the flat join envelope).

To see how this works, return to the motivating example of adjudication. Modify the payoffs so that addressing a grievance is less onerous for the organization $u \in (1-l, 1)$, as illustrated in Figure 4. In the right panel, the flat join envelope is drawn in gray, while the join envelope and concave envelopes coincide and are drawn in purple. For prior beliefs outside interval (θ_a, θ_d), without commitment the sender can do no better than pooling on the belief-optimal action. For prior beliefs inside this interval, the flat join is strictly above the graph of the belief-optimal action. Thus, sender can do better than pooling, even without commitment. For prior beliefs below θ_d, sender can improve even more with commitment to actions, which allows him to join graphs at different heights. Extended commitment, which allows him to join payoff graphs of the same color, does not benefit the sender in this example.

3 APPLICATIONS

3.1 NOMINAL RATINGS

We study the design of an optimal certification, rating, or grading system, endowing the designer with a preference for the rating's *nominal* realization. We focus on two channels by which the nominal grade or rating matters, beyond the information it reveals. First, the designer benefits when the rating is favorable, an important feature in a variety of settings. For example, financial analysts who issue optimistic ratings are more likely to be promoted than their more-accurate peers (Hong and Kubik 2003). Similarly, professors may be able to improve their course evaluations by assigning higher grades, which could improve their career prospects (Johnson 2003). Certification intermediaries may similarly wish to present favorable findings to preserve business relationships with clients. The existence of naive receivers, who take ratings or grades at face-value, implies that a favorable rating can produce a stronger response than is warranted, based solely on its information content.[14] Second, exaggeration may also impose costs. Financial analysts who issue misleading reports risk legal sanction. Certifiers who recommend unsuitable or dangerous products risk litigation. Professors may have a distaste for inflating the grades of undeserving students. More broadly, the designer may have an aversion to lying, particularly in a self-serving manner.

[14] Inderst and Ottaviani (2012a,b) study models of financial advice with naive consumers. Kartik et al. (2007) incorporate credulity into a model of strategic communication. Comparing individual and institutional investors, Malmendier and Shanthikumar (2007) and De Franco et al. (2007) provide evidence that individual investors tend to interpret analyst recommendations naively. In the context of hiring, Hansen et al. (2023) find that some employers react credulously to changes in undergraduate GPA at the time of hiring.

To fix ideas, consider a financial analyst who rates an asset. The state $\omega \in \{b, g\}$ is the asset's quality, and the prior is $\mu_0 = \Pr(\omega = g)$. The analyst can assign two possible ratings to the asset, $s \in \{H, L\}$. Before learning the state, he designs an evaluation procedure or disclosure rule, $\pi(s|\omega)$, which specifies the probability that rating s is issued given state ω. We call $\pi(s|\omega)$ a rating policy. After setting the policy, the analyst privately observes quality and issues a public rating.

The rating is observed by a continuum of investors with total mass 1, each of whom has a single infinitesimal unit of capital.[15] Investors may be either sophisticated or naive. Sophisticated investors observe both the rating policy $\pi(\cdot|\cdot)$ and the realized rating when updating their beliefs. In contrast, naive investors interpret ratings literally: $s = H$ means $\omega = g$, and $s = L$ means $\omega = b$. The fraction of naive investors is $\nu \in (0, 1)$. After observing the rating and updating beliefs, each investor $i \in [0, 1]$ draws an independent outside option $\theta_i \sim F(\cdot)$, where F is twice continuously differentiable with support $[0, 1]$. Investors then decide whether to allocate capital to the asset or their outside options.[16] Therefore, if sophisticated investors believe $\mu = \Pr(\omega = g)$, then aggregate investment is $(1 - \nu)F(\mu) + \nu\mathcal{I}(s = H)$.

The analyst would like to increase investment in the asset. This incentivizes him to exaggerate the rating to exploit the naive investors. This incentive is tempered by the possibility of sanctions and psychological costs. In particular, when the high rating is assigned to an asset with low returns, the analyst pays a (expected) cost $k > 0$. Normalizing the analyst's payoffs by $1/(1 - \nu)$, we have the following interim payoffs,

$$\widehat{v}(\mu, s) = F(\mu) + \mathcal{I}(s = H)(b - c(1 - \mu)),$$

where $b \equiv \nu/(1 - \nu)$ and $c \equiv k/(1 - \nu)$. Formally, the payoffs with $b = 0$ and $c > 0$ correspond to a Spence (1973b) model, in which sender has a signaling action available whose cost depends on the state of the world. With $b > 0$ and $c = 0$, payoffs correspond to a money burning model, where one action is "purely dissipative" (Austen-Smith and Banks 2000, Kartik 2007). Note that identical payoffs would arise in the absence of naive investors ($\nu = 0$) if b was interpreted as a career concern.

[15] Allowing for investors who are not infinitesimal makes no difference to the analysis.

[16] One can interpret $E[\theta]$ as the asset price, and $\theta_i - E[\theta]$ as investor i's idiosyncratic benefit or loss from investing, e.g., tax motives, liquidity needs, or warm glow. To streamline the model and focus on the rating design, we abstract from endogenous price adjustments, though these could be incorporated within our framework. In some cases, (e.g., IPO) financial assets sell for fixed prices.

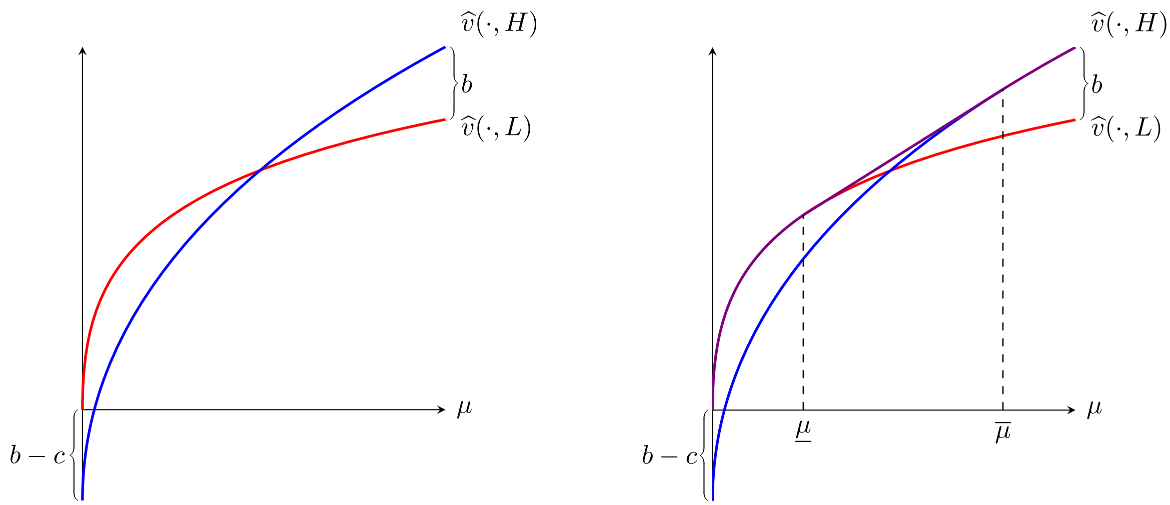

Figure 5: $F(\cdot)$ concave, $c > b$

With commitment, the analyst's optimal rating policy is determined by the interaction of two incentives. First, the concavity or convexity of $F(\cdot)$ determines whether the analyst wishes to reveal or conceal information from sophisticated investors. When $F(\cdot)$ is concave (convex), concealing (revealing) information increases the aggregate investment from the sophisticated segment of the market. Second, the analyst would like to manipulate the naive investors to boost their investment, but he must be wary of a high rating in the bad state. In particular, when $b > c$, the analyst benefits (on net) from boosting naive investment in the bad state even if it means he incurs the cost. In contrast, when $c > b$, the boost to naive investment in the bad state is not enough to offset the cost, and the analyst has an incentive not to exaggerate. Together, these incentives determine the shape and position of the analyst's interim payoff graphs. In particular, both $\widehat{v}(\cdot, H)$ and $\widehat{v}(\cdot, L)$ have identical concavity to $F(\cdot)$, and their relative positions are determined by the sign of $b - c$. When $b > c$, the graph of $\widehat{v}(\cdot, H)$ lies above $\widehat{v}(\cdot, L)$. In the reverse case, the graph of $\widehat{v}(\cdot, H)$ crosses $\widehat{v}(\cdot, L)$ once, and this crossing is from below—see Figure 5.

First, consider concave $F(\cdot)$. If $b > c$, then $\widehat{v}_H$ lies strictly above $\widehat{v}_L$. With concavity, it is optimal for the analyst to pool on the high rating. However, when $c > b$, the payoff graphs cross. Although each graph is concave when viewed in isolation, the crossing creates non-concavity around the point of intersection. Thus, the analyst benefits by joining the graphs near the crossing, as illustrated in Figure 5. The following Proposition follows from inspection of Figure 5.

Proposition 6 *(Financial Ratings, Concave F). In the model of financial ratings, if $F(\cdot)$ is concave and $c > b$, then belief thresholds $\underline{\mu}, \overline{\mu}$ exist, with $0 \leq \underline{\mu} < \overline{\mu} \leq 1$, such that,*

- *if $\mu_0 \leq \underline{\mu}$, then the optimal rating policy always assigns the low rating, $\mu_L = \mu_0$.*
- *if $\underline{\mu} < \mu_0 < \overline{\mu}$, then the optimal rating policy induces beliefs $\{\mu_L = \underline{\mu}, \mu_H = \overline{\mu}\}$. If $\underline{\mu} > 0$, then the optimal rating understates asset value with positive probability, $\pi(L|g) > 0$; if $\overline{\mu} < 1$, then it overstates with positive probability, $\pi(H|b) > 0$.*
- *if $\mu_0 \geq \overline{\mu}$, then the optimal rating policy always assigns the high rating, $\mu_H = \mu_0$.*

To see the intuition, note that when $c > b$, the analyst would like to avoid exaggeration in the bad state and boost naive investment in the good state, pushing him toward a separating strategy. The incentive to separate is tempered by the concavity of $F(\cdot)$—expected investment by the sophisticated types is maximized when the rating is uninformative. At moderate priors, both good state and bad state are relatively likely, giving the strongest incentives for separation. In this case, the optimal rating policy is informative, though it generally misstates asset value with positive probability. When the prior is high, avoiding exaggeration in the bad state is less of a concern because the bad state is unlikely. Thus, the analyst pools on H. Similarly, when the prior is low, boosting naive investment in the good state is less of a concern because the good state is unlikely, and the analyst pools on L.

The logic of the previous paragraph suggests that the analyst's incentive to separate is driven by the existence of naive investors. Here, we show that when the cost of manipulation is not too large, naive investors are critical for the rating system to be informative.

Corollary 1 *(Naive Investors and Informativeness). Suppose $k \in (0, 1)$. At each prior belief $\mu_0 \in (0, 1)$, there exist thresholds $0 < \nu_L(\mu_0) < \nu_M(\mu_0) < \nu_H(\mu_0) \leq 1$ such that (i) the optimal rating is uninformative if $\nu < \nu_L(\mu_0)$, and (ii) the optimal rating is informative if $\nu \in (\nu_M(\mu_0), \nu_H(\mu_0))$.*

Thus, an increase in the fraction of naive investors from a low level ($\nu < \nu_L$) to a moderate level ($\nu_M < \nu < \nu_H$) increases the informativeness of the optimal rating system.

Finally, we point out that both understatement (arising when $\underline{\mu} > 0$) and overstatement (arising when $\overline{\mu} < 1$) are generic features of the optimal rating system.

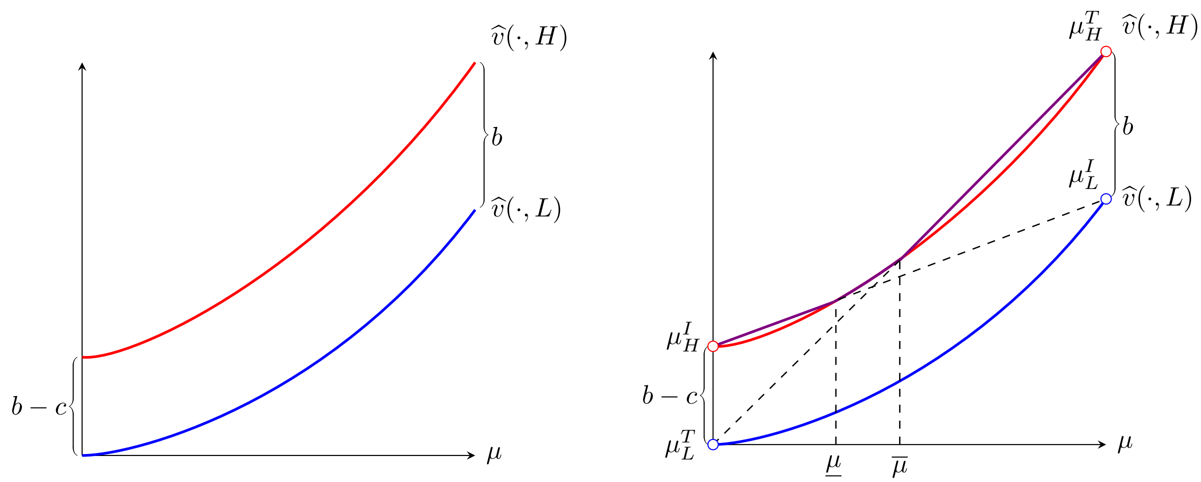

Figure 6: $F(\cdot)$ convex, $c < b$

Remark 3 *Consider a concave* $F(\cdot)$ *and any* (μ_1, μ_2) *with* $0 < \mu_1 < \mu_2 < 1$. *There exist parameters* $\{k, \nu\}$ *such that* $c > b$, *and* $(\underline{\mu}, \overline{\mu}) = (\mu_1, \mu_2)$ *in the optimal rating system described in Proposition 6.*

Next, consider convex $F(\cdot)$. If $c > b$, then both payoff graphs are convex, with $\widehat{v}_H$ crossing $\widehat{v}_L$ once from below. In this case, a truthful rating policy is optimal for the analyst, $\{\mu_H = 1, \mu_L = 0\}$. With convex $F(\cdot)$, full information maximizes aggregate investment from the sophisticated types. Furthermore, with $c > b$, it is too costly for the analyst to manipulate naive investors in the bad state. Both incentives reinforce each other, resulting in truthful ratings. With $c < b$, the incentives are more intricate. Here, the analyst would like to manipulate naive investors in the bad state, even if it means he pays the exaggeration cost. Thus, maximizing investment from the sophisticated types requires full separation, but from the naive types it requires full pooling on H. If a large fraction of investors are naive, pooling on H may be better than separation (even though each payoff graph is convex). Furthermore, even full separation is not as straightforward as it may appear. With a fully separating strategy, the naive types invest in only one state, the one where the asset is rated H. Thus, it is advantageous for the analyst to commit to assign the high rating in the state that is (relatively) more likely. In other words, if the prior is high, then the analyst wants to commit to generate the high rating in the good state, resulting in a truthful rating. But, if the prior is low, then the analyst wants to commit to generate the high rating in the bad state, resulting in an *inverted* rating. The following Proposition follows from inspection of Figure 6.

Proposition 7 *(Financial Ratings, Convex F). In the model of financial ratings, if $F(\cdot)$ is convex and $c < b$, then belief thresholds $\underline{\mu}, \overline{\mu}$ exist with $0 \leq \underline{\mu} \leq \overline{\mu} \leq 1$, such that*

- *if $\mu_0 \leq \underline{\mu}$, then the optimal rating policy assigns the high rating to the bad asset and the low rating to the good asset, $\{\mu_L = 1, \mu_H = 0\}$.*
- *if $\underline{\mu} < \mu_0 < \overline{\mu}$, then the optimal rating policy always assigns the high rating, $\mu_H = \mu_0$.*
- *if $\mu_0 \geq \overline{\mu}$, then the optimal rating policy assigns the high rating to the good asset and the low rating to the bad asset, $\{\mu_L = 0, \mu_H = 1\}$.*

The inverted rating result is surprising. Intuitively, with a convex $F(\cdot)$, sophisticated investors are likely to have good outside options, and their fraction is sufficiently large that the analyst would like to maximize their aggregate investment by revealing the asset's quality. With a low prior belief, however, asset quality is likely to be low, and the sophisticated investors do not invest most of the time. To generate some investment when the asset is bad, the analyst can manipulate the naive types to invest. When the (expected) penalty for such manipulation is low, doing so is advantageous. Therefore, under these conditions, the optimal rating policy inverts the nominal meaning—sophisticated types infer that a highly rated asset is bad and avoid it, while naive types infer that a highly rated is good and invest.

The inverted rating result is consistent with empirical findings in the literature on financial analysts. Using data from the *Global Research Analyst Settlement*, De Franco et al. (2007) find that in numerous instances analysts issued public guidance that was more positive than their private beliefs. The response to these public announcements differed between individual investors (who are more likely to be naive) and (sophisticated) institutional investors. Crucially, they document that institutions *sold* the highly rated assets, while individual investors bought. As we discuss further below, these findings are difficult to reconcile with a model in which analysts make public announcements without commitment. In such a model, sophisticated investors discount or ignore favorable announcements, but they do not interpret favorable announcements as bad news.

Formal Ratings vs. Informal Announcements We conclude this section by highlighting the effects of the analyst's commitment power. In particular, suppose that instead of a formal rating conducted via a specific procedure that is directly observed by sophisticated investors, the analyst simply announces his rating, $s \in \{H, L\}$, after observing asset quality. In other words, consider the identical game without sender commitment power. Perfect Bayesian Equilibria of this game can be classified in the standard taxonomy.

Remark 4 *With standard refinement, the Perfect Bayesian Equilibrium without analyst commitment is unique. In all refined equilibria, (i) a low rating fully reveals the bad state and a high rating is (weakly) good news, $\mu_L = 0$ and $\mu_H \geq \mu_0$; (ii), the rating strategy is determined exclusively by $\{F(\mu_0), \nu, k\}$ (the overall shape of $F(\cdot)$ is irrelevant); (iii) rating informativeness is decreasing in the fraction of naive investors ν.*

Intriguingly, all three of these features are different when the analyst has commitment power. First, recall that with concave $F(\cdot)$, the optimal rating system not only overstates asset value, it may understate it as well. Indeed, whenever $\underline{\mu} > 0$ and $\mu_0 < \overline{\mu}$, asset value is understated with positive probability ($\pi(L|g) > 0$), and thus, the low rating does not fully reveal the bad state, $\mu_L > 0$. Furthermore, with convex $F(\cdot)$, the analyst might use an inverted rating system in which the low rating reveals the good state and the high rating reveals the bad state $\mu_L = 1 > \mu_0 > \mu_H = 0$. Second, note that without commitment, the analyst's incentive to recommend H in the bad state depends only on the level of investment he expects and on the exaggeration cost. With commitment, the global shape of $F(\cdot)$ determines whether the analyst has an incentive to reveal or conceal information, playing a critical role in the characterization. Third, recall from Corollary 1 that with concave $F(\cdot)$ and $c > b$, an increase in the fraction of naive investors can switch the optimal rating from uninformative to informative.

3.2 PLATFORM DESIGN

Online platforms can customize search results based on their extensive information about consumers and products. Thus, a platform can steer consumers toward certain products and away from others by displaying products more or less prominently in its search results (Ursu 2018, Bar-Isaac and Shelegia 2022). By doing so, the platform also conveys some of what it knows to customers, indirectly influencing their search strategy (Janssen et al. 2023, Kaye 2024). In this application, we study the platform's design of an optimal algorithm for customizing search results, taking both of these features into account.

A platform (sender) carries two products, A and B. The consumer is familiar with B, but she is uncertain about A. Though she would rather buy either product than leave empty handed, product A may actually be better suited to her needs. Suppose that if she buys B, then the consumer's ex post payoff (inclusive of price) is $u \in (0,1)$; if she buys A, then it is $\omega \in \{0,1\}$. Thus, product A is a better match for the consumer's needs in state $\omega = 1$ and is a worse match in state $\omega = 0$. The consumer is uncertain whether A is a better match, and $\Pr(\omega = 1) = \mu_0$ under the prior; throughout, we use μ to denote the consumer's belief that $\omega = 1$. The platform earns a commission k_i when product $i \in \{A, B\}$ sells. We focus on the case where the commission on product A is larger, $k_A > k_B$. For convenience, we normalize $k_A = 1$ and $k_B = 0$; thus, the platform's expected commission is equal to the probability that the consumer buys A.

To purchase or learn more about a product, a consumer must access the product's "listing" on the platform. Product B's listing provides a link for the consumer to purchase it (because the consumer knows u, additional product information for B is irrelevant for the consumer). Product A's listing not only provides the link for the customer to purchase, it also reports product characteristics and reviews that may allow the consumer to determine her match quality for A. In particular, upon accessing A's listing, the consumer observes the realization of signal R, where

$$\Pr(R = 0|\omega = 0) = r \quad \Pr(R = \phi|\omega = 0) = 1 - r$$
$$\Pr(R = 1|\omega = 1) = r \quad \Pr(R = \phi|\omega = 1) = 1 - r$$

Thus, by observing the realization of R, the consumer learns the state perfectly with probability r, and with probability $1 - r$ she observes an uninformative "null realization." This "truth-or-noise" signal allows for a straightforward characterization.

The platform designs an algorithm that customizes the consumer's search results based on match-quality (ω). By changing the sequence in which listings are displayed, the platform makes it easier or more difficult for the consumer to purchase or learn about each product ("steering"). For example, one product may appear prominently at the top of the consumer's search results, making its listing easier to locate and access while the other is buried among other irrelevant listings or on later pages of search results. We say that the listing that is easier (more difficult) for the consumer to access is "positioned" or "sequenced" first (second). Furthermore, because the sequence is customized based on the consumer's match-quality, observing the sequence reveals information.

Formally, the platform commits to a strategy or algorithm, $\pi(s|\omega)$, which specifies the probability that product $s \in \{A, B\}$ is sequenced first in state $\omega \in \{0, 1\}$. In stage 1, the consumer observes the first product listing. The consumer then decides whether to purchase the first product or pay a privately known search cost.[17] With probability $h > 0$, the consumer's search cost is sufficiently high that she always purchases the first product in the sequence; with probability $1 - h$, the consumer has a low search cost, $c > 0$. If she purchases the first product, then the game ends. If she pays the search cost, then the game moves to stage 2, in which she accesses the second listing. She then decides whether to purchase the second product or return to purchase the first; either way, the game ends. We allow for free recall of the first product.[18] Note that upon accessing A's listing, the consumer observes a realization of R. This may occur during the first or second stage, depending on A's position in the sequence. To highlight the most interesting aspects of the model, we consider three parametric restrictions: $r < c/(u - u^2)$, $c < \min\{u, (1 - u)/2\}$, and $h < \bar{h} \equiv (1 - r)/(2 - r)$. Thus, the platform's private information about match quality is not easy for the consumer to glean from the product listing, the low search cost is relatively small, and the high search cost is relatively unlikely.

To apply our results, we construct the platform's interim payoff graphs associated with the two possible sequences, A-first (denoted A1), and B-first (B1). To do so, we must solve the consumer's search problem and derive the probability that she purchases A given the product sequence. The solution is standard and is derived in the Appendix.

Lemma 1 *(Search and Interim Payoffs). Let μ be the consumer's belief that $\omega = 1$ after observing the first product. Sender's interim payoffs, normalized by $(1 - h)^{-1}$ are*

$$\widehat{v}(\mu, A1) = \mathcal{I}(\mu \geq \theta_A)(1 - r) + r\mu + f$$
$$\widehat{v}(\mu, B1) = \mathcal{I}(\mu \geq \theta_B)(1 - r + r\mu),$$

where $f \equiv h/(1 - h)$. Furthermore, $0 < \theta_A < u < \theta_B < 1$, threshold θ_A is decreasing in c, and θ_B is increasing in c.

[17]Although match-quality ω is known to the platform, the consumer's search cost is private. The attributes that shape match-quality are deeper, persistent features of the consumer which the platform can learn from observing her past behavior; meanwhile search cost is affected by temporary or idiosyncratic shocks or circumstances that are not observed by the platform.

[18]With free recall of the first product, cost should be interpreted as the cost of identifying or locating the listing of the second product, not a cost for navigating between pages on the platform. With minimal changes, we could accommodate a cost of recalling the first product in stage 2, which would allow for such an interpretation.

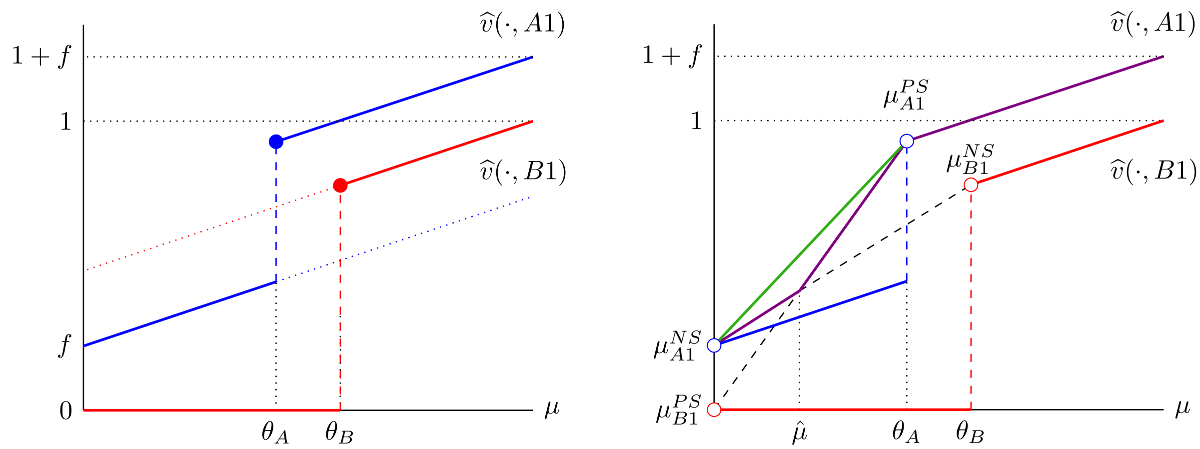

Figure 7: Optimal Sequencing Algorithms.

These graphs are illustrated in Figure 7. To understand their shape, note that the probability that A is bought is increasing in the belief that it matches, with a jump at thresholds $\{\theta_A, \theta_B\}$ arising from beneficial changes in the low cost consumer's search strategy. In particular, when A is sequenced first and $\mu < \theta_A$ the uninformed low cost consumer searches and buys B, but for $\mu \geq \theta_A$, she buys A instead. Similarly, when B is sequenced first, the low cost consumer buys B in the first stage when her belief is low, $\mu < \theta_B$, but she searches for A when her belief is high. Naturally, the cost of search biases the consumer's decision in favor of the first product, $\theta_A < u < \theta_B$, even when her search cost is low. Furthermore, by sequencing A first, the platform ensures that the high search cost consumer purchases it, increasing the platform's expected commission by h. With the normalization in Lemma 1, the purchase decision of the high cost consumer is reflected in the term f, which shifts $\widehat{v}(\cdot, A1)$ up.[19] Our assumption that $h \in (0, \bar{h})$ implies that $f < \bar{f} \equiv 1 - r$. Thus, $\widehat{v}(\cdot, A1)$ is not shifted up too far. Intuitively, this condition says that the platform is willing to lose the (normalized) commission from the high cost consumer (f), if doing so improves the low cost consumer's search behavior.[20]

Both payoff graphs are increasing, and A-first is belief optimal for all $\mu \in [0, 1]$. Thus, the flat-join envelope coincides with $\widehat{v}(\cdot, A1)$. Because the platform's payoff is state-independent,

[19] In a version of our model with $h = 0$, the term f could also be interpreted as a fee that is paid to the platform by marketers or promoters for sequencing A-first.

[20] When μ is high and A is first, the low cost consumer buys A unless she learns $\omega = 0$. Thus, she buys A with probability $r\mu + (1 - r)$. When μ is low and B is listed first, the low cost consumer never buys A. When $f < 1 - r$, the platform would give up f in order to switch the low cost consumer's search strategy, at all $\mu \in [0, 1]$.

we can conclude that without commitment, its highest equilibrium payoff is achieved in a pooling equilibrium with A-first. At low beliefs ($\mu_0 < \theta_A$), the platform would benefit if it could push the consumer's belief past thresholds, $\{\theta_A, \theta_B\}$, which would improve the low cost type's search strategy.

To hone in on the optimal algorithm with commitment, suppose first that the platform intends to use an algorithm in which the first position reveals good news about the product ($\mu_{B1} < \mu_0 < \mu_{A1}$). At low prior beliefs $\mu_0 < \theta_A$, the platform does not lose from reductions in μ_{B1}, and it does not gain much from increases in μ_{A1} past θ_A, the threshold at which the low cost consumer's search behavior becomes favorable. Therefore, among all algorithms with this structure the best one induces beliefs $\{\mu_{B1} = 0, \mu_{A1} = \theta_A\}$. We refer to this algorithm as *positive sequencing* (PS), reflecting the idea that the first position is good news. Conversely, the platform could use an algorithm in which the first position reveals bad news ($\mu_{A1} < \mu_0 < \mu_{B1}$). Following similar logic, among all such algorithms the best one induces beliefs $\{\mu_{A1} = 0, \mu_{B1} = \theta_B\}$, which we refer to as *negative sequencing* (NS). With $f \in (0, \bar{f})$, negative sequencing performs better than the belief-optimal algorithm at $\mu_0 < \theta_A$. What remains is to compare positive and negative sequencing to determine the optimal algorithm.

Proposition 8 *(Optimal Sequencing Algorithm). There is a $\widehat{\mu} \in (0, \theta_A)$ such that*

(i) if $\mu_0 < \widehat{\mu}$, then negative sequencing is the optimal algorithm, $\{\mu_{A1} = 0,\ \mu_{B1} = \theta_B\}$.

(ii) if $\mu_0 \in (\widehat{\mu}, \theta_A)$, then positive sequencing is the optimal algorithm, $\{\mu_{B1} = 0,\ \mu_{A1} = \theta_A\}$.

Furthermore, when c increases, expected profit from positive sequencing increases and expected profit from negative sequencing decreases. Moreover, for all $\mu_0 \in (0, \theta_A)$, the consumer's ex ante expected payoff with the optimal sequencing algorithm is strictly higher than her payoff when the platform always sequences A-first, as it would without commitment.

To see the intuition, suppose PS reveals good news about A by sequencing it first. In this case, the consumer's optimal search implies that the high cost consumer and the uninformed low cost consumer buy A in stage 1. Provided it generates good news, PS is better than NS, regardless of how NS positions the products. In contrast, if PS reveals bad news about A by sequencing it second, then the platform expects no sales of A, its worst possible outcome. Thus, PS is better if it is likely to produce good news and it is worse if it is likely to produce bad news. From Bayes-Plausibility, the probability that PS produces good news about A

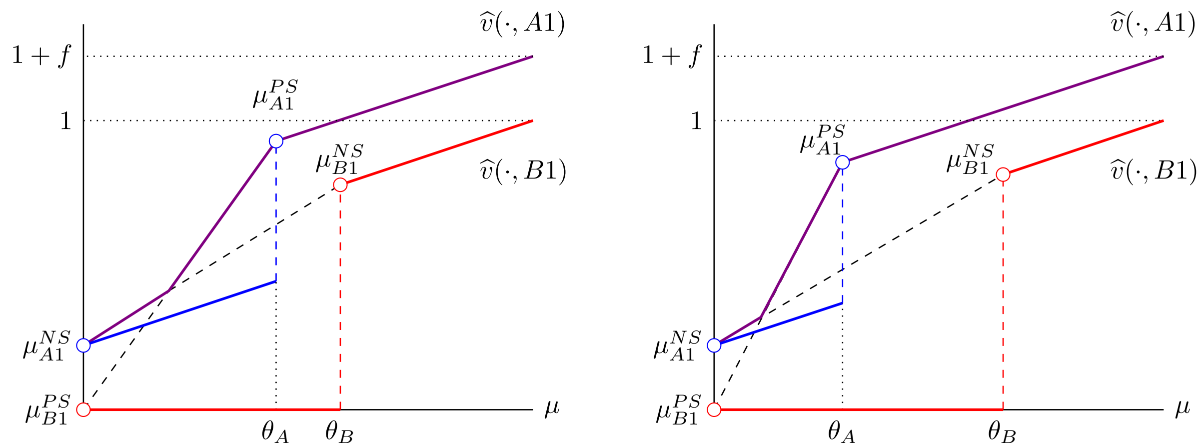

Figure 8: Effect of Search Cost. Increase in c shifts θ_A down and θ_B up, as in left panel.

is increasing in the prior belief μ_0; thus, positive sequencing does better at moderate priors, but negative sequencing does better at low priors.

These optimal algorithms lead to qualitatively different search behaviors from the consumer. With PS, the uninformed low cost consumer buys A in the first stage, even though she believes that B is probably a better choice ($\mu_{A1} = \theta_A < u$). Thus, PS leverages the search cost to deter search, steering the low cost consumer to A. With NS, this is not the case. When the low cost consumer searches for A, she believes that it is likely to be the better choice ($\mu_{B1} = \theta_B > u$). In other words, NS incentivizes search by the low cost consumer. Because of the different role of search cost in these algorithms, expected profit responds differently when it changes, as illustrated in Figure 8.

The optimal sequencing algorithm generates a higher ex ante expected payoff for the consumer than the algorithm which always positions A first. In other words, the platform's commitment to a steering algorithm leads to a Pareto improvement. Interestingly, with both the positive and negative sequencing algorithm, the low cost consumer's purchase decisions generate the same expected payoff as in the absence of commitment; the optimal sequencing algorithm reduces the low cost consumer's expected search cost. Furthermore, with positive sequencing the high cost consumer's purchase decision is more informed than steering to A, which benefits her. Under negative sequencing, however, the high cost consumer's purchase decision is distorted. Nevertheless, we show that under negative sequencing, the low cost consumer's expected gains outweigh the high cost consumer's expected losses.

The consumer's gains are undone when the platform has extended commitment, which allows it to transmit information to the consumer using messages, rather than the product sequence. In this case, the platform always positions A-first (it is belief-optimal) and commits to a recommendation system that transmits a message directly to the consumer. The optimal recommendation system either reveals that product A is a mismatch ($\mu_L = 0$), or it induces posterior belief $\mu_H = \theta_A$, leading to favorable search behavior (see the green line in Figure 7). Though the induced beliefs are identical to the positive sequencing algorithm, the recommendation system improves the platform's profit by selling A to the high cost consumer, even when it reveals that product B is the better match.[21] We show that the platform's gains come at the consumer's expense: the consumer's ex ante expected welfare is lower with the optimal recommendation system than with the optimal sequencing algorithm.

Proposition 9 *(Recommendation System vs. Sequencing) When the prior belief $\mu_0 < \theta_A$,*

(i) The consumer's ex ante expected payoff with the optimal recommendation system is strictly lower than with the optimal sequencing algorithm.

(ii) The platform's ex ante expected profit with the optimal recommendation system is strictly higher than with the optimal sequencing algorithm.

When the prior belief $\mu_0 \geq \theta_A$, consumer and platform are indifferent between the optimal recommendation system and optimal sequencing algorithm.

Together Propositions 8 and 9 present a nuanced perspective on the "power of defaults and ordering" in steering customers toward profitable products, highlighting a distinction between sequencing algorithms and recommendation systems. Further research could clarify this distinction, which may have significant policy implications.

4 CONCLUSION

In this paper, we study a standard signaling game, with the novel feature that the sender can commit to his strategy at the outset. We provide a geometric characterization of the sender's optimal payoff and strategy based on the topological join of the sender's interim payoff graphs. Extending our analysis to allow sender to commit to a communication protocol along with his action, we identify sender's gain from this additional commitment and

[21]Graphically, the platform's gains from the recommendation system is the vertical distance between the green and the purple lines in Figure 1.

characterize environments in which it is zero. We then apply our analysis to provide novel insights into three applications: the design of adjudication procedures, financial ratings, and sequencing algorithms for online platforms.

Multiple directions for future research stand out. One direction is to extend the signaling with commitment environment to incorporate dynamics (Kaya 2009), imperfect observability of sender actions (Jungbauer and Waldman 2023, 2024), receiver's private information (Feltovich et al. 2002), or some form of receiver commitment (Whitmeyer 2024). A second direction is to investigate robustness in the spirit of Bergemann et al. (2017), Bergemann and Morris (2019), and Taneva (2019). Suppose an outside observer knows that sender and receiver play a signaling game, but does not know the information of the sender or the receiver, or whether the sender has commitment power. It remains to characterize the possible outcomes from that observer's perspective and their relationship with the outcomes of the signaling with commitment environment. Methodologically, the robustness literature uses the "obedience approach" based on Bayes Correlated Equilibrium, rather than the belief-based approach that we use in our analysis. Understanding the connection between these approaches in signaling with commitment is a third direction for future research.

REFERENCES

Alonso, Ricardo and Odilon Câmara, "Persuading Voters," *American Economic Review*, 2016, *106* (11), 3590–3605.

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan, "Signaling in a Global Game: Coordination and Policy Traps," *Journal of Political Economy*, 2006, *114* (3), 452–484.

Au, Pak Hung and Keiichi Kawai, "Competitive Information Disclosure by Multiple Senders," *Games and Economic Behavior*, 2020, *119*, 56–78.

Aumann, Robert J. and Sergiu Hart, "Long Cheap Talk," *Econometrica*, 2003, *71* (6), 1619–1660.

Austen-Smith, David and Jeffrey S. Banks, "Cheap Talk and Burned Money," *Journal of Economic Theory*, 2000, *91* (1), 1–16.

Bagwell, Kyle and Michael H. Riordan, "High and Declining Prices Signal Product Quality," *American Economic Review*, 1991, *81* (1), 224–239.

Banks, Jeffrey S, *Signaling Games in Political Science*, New York: Routledge, 1991.

Bar-Isaac, Heski and Joyee Deb, "Reputation With Opportunities for Coasting," *Journal of the European Economic Association*, 06 2020, *19* (1), 200–236.

— **and Sandro Shelegia**, "Monetizing Steering," *CEPR Discussion Paper No. 17685*, 2022.

Bénabou, Roland and Jean Tirole, "Incentives and Prosocial Behavior," *American Economic Review*, December 2006, *96* (5), 1652–1678.

Bergemann, Dirk and Stephen Morris, "Information Design: A Unified Perspective," *Journal of Economic Literature*, March 2019, *57* (1), 44–95.

— **, Benjamin Brooks, and Stephen Morris**, "First-Price Auctions With General Information Structures: Implications for Bidding and Revenue," *Econometrica*, 2017, *85* (1), 107–143.

Best, James and Daniel Quigley, "Persuasion for the long run," *Mimeo*, 2020.

Bloedel, Alexander and Ilya Segal, "Persuading a Rationally Inattentive Agent," *Mimeo*, 2021.

Boleslavsky, Raphael and Christopher Cotton, "Grading standards and education quality," *American Economic Journal: Microeconomics*, 2015, *7* (2), 248–279.

Camerer, Colin, "Gifts as Economic Signals and Social Symbols," *American Journal of Sociology*, 01 1988, *94*, 180–214.

Caplin, Andrew and Mark Dean, "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 2015, *105* (7), 2183–2203.

De Franco, Gus, Hai Lu, and Florin P. Vasvari, "Wealth Transfer Effects of Analysts' Misleading Behavior," *Journal of Accounting Research*, 2007, *45* (1), 71–110.

Deb, Rahul, Mallesh M Pai, and Maher Said, "Indirect Persuasion," *Mimeo*, 2022.

DeMarzo, Peter and Darrell Duffie, "A Liquidity-Based Model of Security Design," *Econometrica*, 1999, *67* (1), 65–99.

DeMarzo, Peter M., "The Pooling and Tranching of Securities: A Model of Informed Intermediation," *Review of Financial Studies*, 2005, *18* (1), 1–35.

Denti, Tommaso, "Posterior Separable Cost of Information," *American Economic Review*, October 2022, *112* (10), 3215–59.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan, "Consumer Price Search and Platform Design in Internet Commerce," *American Economic Review*, July 2018, *108* (7), 1820–59.

Feltovich, Nick, Richmond Harbaugh, and Ted To, "Too Cool for School? Signalling and Countersignalling," *The RAND Journal of Economics*, 2002, *33* (4), 630–649.

Fudenberg, Drew and Jean Tirole, *Game theory*, New York: MIT press, 1991.

Gentzkow, Matthew and Emir Kamenica, "Costly persuasion," *American Economic Review*, 2014, *104* (5), 457–462.

Hansen, Anne Toft, Ulrik Hvidman, and Hans Sievertsen, "Grades and Employer Learning," *Journal of Labor Economics*, 2023.

Hong, Harrison and Jeffrey D. Kubik, "Analyzing the Analysts: Career Concerns and Biased Earnings Forecasts," *The Journal of Finance*, 2003, *58* (1), 313–351.

Inderst, Roman and Marco Ottaviani, "Financial Advice," *Journal of Economic Literature*, June 2012, *50* (2), 494–512.

_ and _ , "How (not) to pay for advice: A framework for consumer financial protection," *Journal of Financial Economics*, 2012, *105* (2), 393–411.

Janssen, Maarten, Thomas Jungbauer, Marcel Preuss, and Cole Williams, "Search Platforms: Big Data and Sponsored Positions," *CEPR Discussion Paper No. 18639*, 2023.

Johnson, Valen E., *Grade Inflation: A Crisis in College Education*, Springer, 2003.

Jungbauer, Thomas and Michael Waldman, "Self-Reported Signaling," *American Economic Journal: Microeconomics*, August 2023, *15* (3), 78–117.

_ **and** _ , "Actions and Signals," *Mimeo*, 2024.

Kamenica, Emir, "Bayesian Persuasion and Information Design," *Annual Review of Economics*, 2019, *11* (1), 249–272.

_ **and Matthew Gentzkow**, "Bayesian Persuasion," *American Economic Review*, October 2011, *101* (6), 2590–2615.

Kartik, Navin, "A note on cheap talk and burned money," *Journal of Economic Theory*, 2007, *136* (1), 749–758.

_ , "Strategic communication with lying costs," *The Review of Economic Studies*, 2009, *76* (4), 1359–1395.

_ **, Marco Ottaviani, and Francesco Squintani**, "Credulity, lies, and costly talk," *Journal of Economic theory*, 2007, *134* (1), 93–116.

Kaya, Ayça, "Repeated signaling games," *Games and economic behavior*, 2009, *66* (2), 841–854.

Kaye, Aaron P., "The Personalization Paradox: Welfare Effects of Personalized Recommendations in Two-Sided Digital Markets," *Mimeo*, 2024.

Kreps, David M and Robert Wilson, "Reputation and imperfect information," *Journal of Economic Theory*, 1982, *27* (2), 253–279.

Leland, Hayne E. and David H. Pyle, "Informational Asymmetries, Financial Structure, and Financial Intermediation," *The Journal of Finance*, 1977, *32* (2), 371–387.

Lipnowski, Elliot and Doron Ravid, "Cheap talk with transparent motives," *Econometrica*, 2020, *88* (4), 1631–1660.

Malmendier, Ulrike and Devin Shanthikumar, "Are small investors naive about incentives?," *Journal of Financial Economics*, 2007, *85* (2), 457–489.

Meng, Anne, *Constraining dictatorship: From personalized rule to institutionalized regimes*, Cambridge University Press, 2020.

Milgrom, Paul and John Roberts, "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 1986, *94* (4), 796–821.

Nelson, Phillip, "Advertising as information," *Journal of Political Economy*, 1974, *82* (4), 729–754.

Onuchic, Paula and Debraj Ray, "Conveying Value Via Categories," *Theoretical Economics*, 2023, *18* (4), 1407–1439.

Pomatto, Luciano, Philipp Strack, and Omer Tamuz, "The Cost of Information: The Case of Constant Marginal Costs," *American Economic Review*, 2023, *113* (5), 1360–93.

Ross, Stephen A., "The Determination of Financial Structure: The Incentive-Signalling Approach," *The Bell Journal of Economics*, 1977, *8* (1), 23–40.

Rourke, Colin P. and Brian J. Sanderson, *Introduction to Piecewise-Linear Topology*, Heidelberg: Springer-Verlag, 1982.

Scott Morton, Fiona, Pascal Bouvier, Ariel Ezrachi, Bruno Jullien, Roberta Katz, Gene Kimmelman, Douglas Melamed, and Jamie Morgenstern, "Report of the Market Structure and Antitrust Subcommittee, Digital Platforms Project," *George J Stigler Center for the Study of the Economy and the State, The University of Chicago Booth School of Business*, 2019.

Shadmehr, Mehdi and Raphael Boleslavsky, "Institutions, Repression, and the Spread of Protest," *Mimeo*, 2015.

Sims, Christopher A, "Implications of rational inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Spence, A. Michael, "Time and Communication in Economic and Social Interaction," *The Quarterly Journal of Economics*, 1973, *87* (4), 651–660.

Spence, Michael, "Job Market Signaling," *The Quarterly Journal of Economics*, 1973, *87* (3), 355–374.

Taneva, Ina, "Information Design," *American Economic Journal: Microeconomics*, November 2019, *11* (4), 151–85.

Ursu, Raluca M., "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions," *Marketing Science*, 2018, *37* (4), 530–552.

Whitmeyer, Mark, "Bayesian Elicitation," *Mimeo*, 2024.

Zhang, Jun and Junjie Zhou, "Information Disclosure in Contests: A Bayesian Persuasion Approach," *The Economic Journal*, 2016, *126* (597), 2197–2217.

5 APPENDIX: PROOFS

Proof of Proposition 1. The proof follows Kamenica and Gentzkow (2011). It is presented only for completeness.

First, we show that a belief system induced by strategy Π is Bayes-Plausible.

$$\sum_{s\in S}\tau(\mu_s)\mu_s(\omega)=\sum_{s\in S}\mu_0(\omega)\pi(s|\omega)=\mu_0(\omega)\sum_{s\in S}\pi(s|\omega)=\mu_0(\omega).$$

Next, we show that if a belief system is Bayes-Plausible, then it is induced by a strategy. Suppose belief system $\mathcal{B}\equiv\{\mu_s,\tau(\mu_s)\}_{s\in S}$ is Bayes-Plausible. Consider the following strategy:

$$\pi(s|\omega)=\frac{\tau(\mu_s)\mu_s(\omega)}{\mu_0(\omega)},$$

where the full support assumption ensures $\mu_0(\omega)>0$. Obviously, $\pi(s|\omega)\geq 0$. Moreover, Bayes-Plausibility implies $\sum_{s\in S}\tau(\mu_s)\mu_s(\omega)=\mu_0(\omega)$, and therefore $\sum_{s\in S}\pi(s|\omega)=1$. By construction, the proposed strategy induces belief system $\mathcal{B}$. ■

Proof of Proposition 2. *Proof of part (i) "if" and part (ii).* Suppose $(\mu_0,z)\in join(\widehat{v}_s)_{s\in S}$. By definition, $\{\lambda_s\}_{s\in S}$ exist such that $\lambda_s\geq 0$, $\sum_{s\in S}\lambda_s=1$, and $(\mu_0,z)=\sum_{s\in S}\lambda_s z_s$, where $z_s\in\widehat{v}_s$. From the definition of $\widehat{v}_s$, we have that for each $s\in S$, a $\mu_s\in\Delta(\Omega)$ exists such that $z_s=(\mu_s,\widehat{v}(\mu_s,s))$. Substituting, we have $(\mu_0,z)=\sum_{s\in S}\lambda_s(\mu_s,\widehat{v}(\mu_s,s))=(\sum_{s\in S}\lambda_s\mu_s,\sum_{s\in S}\lambda_s\widehat{v}(\mu_s,s))$. Thus, belief system $\{\mu_s,\tau(\mu_s)=\lambda_s\}_{s\in S}$ is Bayes-Plausible ($\sum_{s\in S}\tau(\mu_s)\mu_s=\mu_0$) and attains payoff z in sender's problem ($\sum_{s\in S}\tau(\mu_s)\widehat{v}(\mu_s,s)=z$).

Proof of part (i) "only if." We show that if a Bayes-Plausible belief system $\{\mu_s,\tau(\mu_s)\}_{s\in S}$ attains payoff z in sender's problem, then $(\mu_0,z)\in join(\widehat{v}_s)_{s\in S}$. Observe that $(\mu_0,z)=(\sum_{s\in S}\tau(\mu_s)\mu_s,\sum_{s\in S}\tau(\mu_s)\widehat{v}(\mu_s,s))=\sum_{s\in S}\tau(\mu_s)(\mu_s,\widehat{v}(\mu_s,s))$. Note that $(\mu_s,\widehat{v}(\mu_s,s))\in\widehat{v}_s$.

Furthermore, $\sum_{s\in S}\tau(\mu_s) = 1$ and $\tau(\mu_s) \geq 0$. Therefore, with $\lambda_s = \tau(\mu_s)$, we have $(\mu_0, z) = \sum_{s\in S}\lambda_s z_s$, where $z_s = (\mu_s, \widehat{v}(\mu_s, s)) \in \widehat{v}_s$. Therefore, $(\mu_0, z) \in join(\widehat{v}_s)_{s\in S}$.

Part (iii) follows immediately from the definition of join envelope. ■

5.1 EXTENDED COMMITMENT

Proof of Proposition 3. *Part (i).* First, we show that if $\tau(\mu, s)$ is induced by a sender strategy, then the marginal distribution of the posterior belief is Bayes-Plausible.

Consider a sender strategy, $\overline{\Pi} = \{\pi(m, s|\omega) \,|\, m \in M,\, s \in S,\, \omega \in \Omega\}$. Any realization from this strategy (m, s), generates a posterior belief about the state $\mu_{(m,s)}$ and arises with a certain probability $\hat{\tau}(m, s)$,

$$\hat{\tau}(m, s) = \sum_{\omega\in\Omega}\pi(m, s|\omega)\mu_0(\omega) \qquad \Pr(\omega|m, s) = \mu_{(m,s)}(\omega) = \frac{\pi(m, s|\omega)\mu_0(\omega)}{\hat{\tau}(m, s)}.$$

Consider a partition of the set of realizations (m, s) based on the posterior belief $\mu_{(m,s)}$ that they induce and the action s. Define a set $E(\mu, s) \equiv \{(m, s)|m \in M,\, \mu_{(m,s)} = \mu\}$, the set of all message/action pairs that generate posterior belief μ and have action s. By definition, joint probability $\tau(\mu, s) = \sum_{(m,s)\in E(\mu,s)}\hat{\tau}(m, s)$. Next, note that by the Law of Iterated Expectations (or direct substitution),

$$\sum_{(m,s)\in M\times S}\hat{\tau}(m, s)\mu_{(m,s)} = \mu_0. \tag{1}$$

Regrouping the sum according to the sets $E(\mu, s)$ yields

$$\begin{aligned}\sum_{(m,s)\in M\times S}\hat{\tau}(m, s)\mu_{(m,s)} &= \sum_{(\mu,s)\in\Delta(\Omega)\times S}\Big\{\sum_{(m,s)\in E(\mu,s)}\hat{\tau}(m, s)\mu_{(m,s)}\Big\}\\ &= \sum_{(\mu,s)\in\Delta(\Omega)\times S}\Big\{\mu\sum_{(m,s)\in E(\mu,s)}\hat{\tau}(m, s)\Big\}\\ &= \sum_{(\mu,s)\in\Delta(\Omega)\times S}\mu\ \tau(\mu, s). \qquad (2)\end{aligned}$$

Combining (1) and (2) yields

$$\sum_{(\mu,s)\in\Delta(\Omega)\times S}\tau(\mu, s)\mu = \mu_0.$$

Decomposing $\tau(\mu, s) = \tau(\mu)\tau(s|\mu)$ in the above equation yields

$$\sum_{(\mu,s)\in\Delta(\Omega)\times S} \tau(\mu)\tau(s|\mu)\mu = \sum_{\mu\in\Delta(\Omega)} \sum_{s\in S} \{\tau(\mu)\tau(s|\mu)\mu\} = \sum_{\mu\in\Delta(\Omega)} \tau(\mu)\mu = \mu_0.$$

Next, we show that if $\tau(\mu)$ is Bayes-Plausible, then it can be induced by a strategy such that $\Pr(\omega|m, s) = \Pr(\omega|m)$.

Consider joint distribution $\tau(\mu, s)$ such that the marginal distribution $\tau(\mu)$ is Bayes-Plausible. From Kamenica and Gentzkow (2011) Proposition 1, the sender can design a signal $\pi'(m|\omega)$ that induces $\tau(\mu)$ using the message space alone. Next, define a joint probability of messages and actions conditional on ω as $\pi'(m, s|\omega) = \Pr(m|\omega)\Pr(s|m, \omega) = \pi'(m|\omega)\tau(s|\mu = \mu_m)$, where μ_m is the posterior induced by message m from signal $\pi'(m|\omega)$. Clearly, this construction induces $\tau(\mu, s)$. Moreover, in this construction, $\Pr(s|m, \omega) = \tau(s|\mu = \mu_m)$, so that $\Pr(s|m, \omega) = Pr(s|m)$, and hence

$$\Pr(\omega|m, s) = \Pr(\omega|m)\frac{\Pr(s|m, \omega)}{\Pr(s|m)} = \Pr(\omega|m).$$

Part (ii). As part of the proof of part (i), we showed: if $\tau(\mu)$ is Bayes-Plausible, then it can be induced by a strategy such that $\Pr(\omega|m, s) = \Pr(\omega|m)$. This proves part (ii). ■

Lemma 2 *The convex hull of the union of the sender's payoff graphs is equal to the convex hull of the join of the sender's payoff graphs, $con((\widehat{v}_s)_{s\in S}) = con(join(\widehat{v}_s)_{s\in S})$.*

Proof of Lemma 2. First, we show that $con(\widehat{v}_s)_{s\in S} \subseteq con(join(\widehat{v}_s)_{s\in S})$. This is immediate, since the union of the payoff graphs is a weak subset of their join (by definition).

Next, we show that $con(join(\widehat{v}_s)_{s\in S}) \subseteq con(\widehat{v}_s)_{s\in S}$. Any point in $join(\widehat{v}_s)_{s\in S}$ can be expressed as a convex combination of some points belonging to the sender's payoff graphs. Thus, any point in $join(\widehat{v}_s)_{s\in S}$ is a convex combination of points in $\cup_{s\in S}\widehat{v}_s$. Therefore, any convex combination of points in $join(\widehat{v}_s)_{s\in S}$, is also a convex combination of points in $\cup_{s\in S}\widehat{v}_s$. The result follows. ■

Proof of Proposition 5. Part (i). From Lemma 2, we have that in the extended commitment benchmark, sender's optimal maximum payoff is the concave envelope of $V^{jo}(\cdot)$. The result follows from Jensen's inequality.

Part (ii). Suppose to the contrary that $\widehat{v}(\cdot, s)$ is concave for all $s \in S$, but $V^{jo}(\mu) \neq V^{co}(\mu)$ for some belief μ. It must be that $V^{co}(\mu)$ places positive weight on more than one point from $\widehat{v}_s$ for some $s \in S$. In particular, suppose $(\mu, V^{co}(\mu))$ includes two points, $x_{1s} = (\mu_1, \widehat{v}(\mu_1, s))$ and $x_{2s} = (\mu_2, \widehat{v}(\mu_2, s))$ from the same graph $\widehat{v}_s$ in the convex combination with strictly positive weights λ_1 and λ_2, respectively. In the convex combination, replace these two points with $\hat{\mu} = \alpha\mu_1 + (1-\alpha)\mu_2$, where $\alpha = \lambda_1/(\lambda_1 + \lambda_2)$, and the associated weight $\hat{\lambda} = \lambda_1 + \lambda_2$. By concavity, $\hat{\lambda}\, \widehat{v}(\hat{\mu}, s) \geq \lambda_1\, \widehat{v}(\mu_1, s) + \lambda_2\, \widehat{v}(\mu_2, s)$. This replacement results in a new convex combination with a higher value of the sender's payoff, contradicting the definition on $V^{co}(\mu)$.
■

5.2 ADDITIONAL DISCUSSION

Consider an *information design* benchmark. Sender designs message spaces M_σ, M_ρ and a joint distribution $\pi_I(\cdot,\cdot,\cdot|\omega)$ over private messages and a public action $(m_\sigma, m_\rho, s) \in M_\sigma \times M_\rho \times S$, conditional on state ω. Note that if we constrain sender to use a public message, then this benchmark reduces to extended commitment.

Lemma 3 *In the information design benchmark, any combination of sender and receiver payoffs that can be attained can also be attained with a public message.*

Proof of Lemma 3. Consider private communication. Sender's expected payoff is

$$\sum_{(m_\sigma, m_\rho, s, \omega)} \mu_0(\omega)\pi_I(m_\sigma, m_\rho, s|\omega)v(s, \widehat{a}(m_\rho, s), \omega),$$

and receiver's expected payoff is

$$\sum_{(m_\sigma, m_\rho, s, \omega)} \mu_0(\omega)\pi_I(m_\sigma, m_\rho, s|\omega)u(s, \widehat{a}(m_\rho, s), \omega),$$

where the summation is over all $(m_\sigma, m_\rho, s, \omega) \in M_\sigma \times M_\rho \times S \times \Omega$ and $\widehat{a}(m_\rho, s)$ denotes the receiver's optimal response at (m_ρ, s). Decompose the joint distribution $\pi_I(\cdot,\cdot,\cdot|\omega)$ to into two parts: $\pi_I(m_\sigma, m_\rho, s|\omega) = \pi_\rho(m_\rho, s|\omega)\pi_\sigma(m_\sigma|m_\rho, s, \omega)$. Substituting, we have that sender's payoff is

$$\sum_{(m_\rho,s,\omega)} \sum_{m_\sigma \in M_\sigma} \mu_0(\omega)\pi_\rho(m_\rho, s|\omega)\pi_\sigma(m_\sigma|m_\rho, s, \omega)v(s, \widehat{a}(m_\rho, s), \omega) = \sum_{(m_\rho,s,\omega)} \mu_0(\omega)\pi_\rho(m_\rho, s|\omega)v(s, \widehat{a}(m_\rho, s), \omega),$$

and receiver's is

$$\sum_{(m_\rho,s,\omega)} \mu_0(\omega)\pi_\rho(m_\rho, s|\omega)u(s, \widehat{a}(m_\rho, s), \omega).$$

Now, consider public communication with $M = M_\rho$ and a family of joint distributions $\Pi_E = \{\pi_\rho(m, s|\omega)|(m, s) \in M \times S, \omega \in \Omega\}$. Sender's expected payoff is

$$\sum_{(m_\rho,s,\omega)} \mu_0(\omega)\pi_\rho(m_\rho, s|\omega)v(s, \widehat{a}(m_\rho, s), \omega)$$

and receiver's is

$$\sum_{(m_\rho,s,\omega)} \mu_0(\omega)\pi_\rho(m_\rho, s|\omega)u(s, \widehat{a}(m_\rho, s), \omega),$$

which are evidently identical to the payoffs in the case of private communication. ■

5.3 NOMINAL RATINGS

Proof of Corollary 1. *Proof of Claim (i).*

Step 1. We show that there exists $\nu_L(\mu_0) > 0$ such that, if $\nu < \nu_L(\mu_0)$, then for all $\mu \in [0, 1]$,

$$F(\mu) + \frac{\nu - k}{1 - \nu} + \frac{k}{1 - \nu}\mu < F'(\mu_0)(\mu - \mu_0) + F(\mu_0).$$

Consider

$$Q(\nu) \equiv \min_{\mu \in [0,1]} F'(\mu_0)(\mu - \mu_0) + F(\mu_0) - F(\mu) - \Big(\frac{\nu - k}{1 - \nu} + \frac{k}{1 - \nu}\mu\Big),$$

and let $\mu^*(\nu)$ be the argminimum. By the Maximum Theorem, $Q(\cdot)$ is continuous. Therefore, it is enough to show $Q(0) > 0$.

$$Q(0) = F'(\mu_0)(\mu^*(0) - \mu_0) + F(\mu_0) - F(\mu^*(0)) + k(1 - \mu^*(0)).$$

First, suppose $\mu^*(0) < 1$. Concavity of $F(\cdot)$ implies $F'(\mu_0)(\mu^*(0)-\mu_0)+F(\mu_0)-F(\mu^*(0)) \geq 0$, and hence $Q(0) \geq k(1-\mu^*(0)) > 0$. Second, suppose $\mu^*(0) = 1$. Strict concavity of $F(\cdot)$ and $\mu_0 < 1$ imply $F'(\mu_0)(\mu^*(0)-\mu_0) + F(\mu_0) - F(\mu^*(0)) > 0$, and hence $Q(0) > 0$.

Step 2. We show that if $\nu < \nu_L(\mu_0)$, then any Bayes-Plausible belief system with $\tau(\mu_H) > 0$ is worse for sender than pooling on L, i.e., $\mu_L = \mu_0$ and $\tau(\mu_L) = 1$. Consider a prior belief $\mu_0 \in (0,1)$ and a Bayes-Plausible distribution supported on posteriors $\{\mu_L, \mu_H\}$ with $\tau(\mu_H) > 0$. Consider the analyst's payoff of such a belief system,

$$G(\mu_L, \mu_H) \equiv \tau(\mu_L)\widehat{v}(\mu_L, L) + \tau(\mu_H)\widehat{v}(\mu_H, H).$$

Using Step 1 and the concavity of $F(\cdot)$, we have

$$\begin{aligned}
\widehat{v}(\mu_H, H) &= F(\mu_H) + \frac{\nu - k}{1-\nu} + \frac{k}{1-\nu}\mu_H < F'(\mu_0)(\mu_H - \mu_0) + F(\mu_0), \\
\widehat{v}(\mu_L, L) &= F(\mu_L) \leq F'(\mu_0)(\mu_L - \mu_0) + F(\mu_0).
\end{aligned}$$

Using these bounds, along with $\tau(\mu_H) > 0$, we have

$$G(\mu_L, \mu_H) < \tau(\mu_L)(F'(\mu_0)(\mu_L - \mu_0) + F(\mu_0)) + \tau(\mu_H)(F'(\mu_0)(\mu_H - \mu_0) + F(\mu_0)) = F(\mu_0),$$

where the equality uses Bayes-Plausibility. Note that by pooling on L, the analyst's payoff is $\widehat{v}(\mu_0, L) = F(\mu_0)$. This completes the proof of Claim (i).

Proof of Claim (ii). Suppose $\mu_0 \in (0,1)$. Consider a Bayes-Plausible belief system $\{\mu_L = \mu_0 - \delta, \mu_H = \mu_0 + \delta\}$, $\tau(\mu_H) = \tau(\mu_L) = 1/2$. The payoff of such a belief system is

$$S(\delta, \nu) \equiv G(\mu_0 - \delta, \mu_0 + \delta) = \frac{1}{2}F(\mu_0 - \delta) + \frac{1}{2}\left(F(\mu_0 + \delta) + \frac{\nu - k + k\mu_0 + k\delta}{1-\nu}\right).$$

The highest possible payoff from an uninformative belief system is

$$P(\nu) = \max\{\widehat{v}(\mu_0, L), \widehat{v}(\mu_0, H)\} = \max\left\{F(\mu_0), F(\mu_0) + \frac{\nu - k + k\mu_0}{1-\nu}\right\}.$$

We show that there exists a $\nu^*(\mu_0) \in (0,1)$ and a $\delta^* > 0$ such that $S(\delta^*, \nu^*(\mu_0)) > P(\nu)$. Let $\nu^*(\mu_0) = k(1-\mu_0)$. Note that $k < 1 \Rightarrow \nu^*(\mu_0) \in (0,1)$. Note further,

$$\frac{\nu^*(\mu_0) - k + k\mu_0}{1-\nu} = 0 \Rightarrow P(\nu^*(\mu_0)) = F(\mu_0).$$

It follows that

$$\begin{aligned} S(\delta, \nu^*(\mu_0)) =& \frac{1}{2}F(\mu_0 - \delta) + \frac{1}{2}\left(F(\mu_0 + \delta) + \frac{\nu^*(\mu_0) - k + k(\mu_0 + \delta)}{1 - \nu^*(\mu_0)}\right) \\ =& \frac{1}{2}F(\mu_0 - \delta) + \frac{1}{2}\left(F(\mu_0 + \delta) + \frac{k\delta}{1 - \nu^*(\mu_0)}\right). \end{aligned}$$

Moreover, $S(0, \nu^*(\mu_0)) = F(\mu_0)$, and

$$\frac{\partial S(\delta, \nu^*(\mu_0))}{\partial \delta}\Big|_{\delta=0} = \frac{k}{1 - \nu^*(\mu_0))} > 0$$

Therefore, a $\delta^* > 0$ exists, such that $S(\delta^*, \nu^*(\mu_0)) > F(\mu_0) = P(\nu^*(\mu_0))$. Next, note that $S(\delta^*, \cdot)$ and $P(\cdot)$ are continuous at $\nu = \nu^*(\mu_0) < 1$. By implication, a non-empty interval $(\nu_M(\mu_0), \nu_H(\mu_0))$ exists such that, if $\nu \in (\nu_M(\mu_0), \nu_H(\mu_0))$, then $S(\delta^*, \nu) > P(\nu)$. Thus, the proposed informative belief system improves on an uninformative one. Therefore, the optimal rating must be informative. ■

Proof of Remark 3. Consider $0 < \mu_1 < \mu_2 < 1$. From Figure 5, a sufficient condition for $\underline{\mu} = \mu_1$ and $\overline{\mu} = \mu_2$ consists of two parts. (1) The crossing of the payoff graphs, $\mu^\dagger \equiv 1 - b/c$, lies between μ_1 and μ_2, i.e., $\mu_1 < \mu^\dagger < \mu_2$. (2) the tangent line to $\widehat{v}(\cdot, L)$ at $\mu = \mu_1$ is also tangent to $\widehat{v}(\cdot, H)$ at $\mu = \mu_2$. Condition (2) is equivalent to

$$\frac{d\widehat{v}(\mu, L)}{d\mu}\Big|_{\mu=\mu_1} = \frac{d\widehat{v}(\mu, H)}{d\mu}\Big|_{\mu=\mu_2} = \frac{\widehat{v}(\mu_2, H) - \widehat{v}(\mu_2, L)}{\mu_2 - \mu_1}. \tag{3}$$

Step 1. We show that for any $0 < \mu_1 < \mu_2 < 1$, *there exist parameters* $b > 0$ *and* $c > 0$ *such that (3) holds.* Substituting $\widehat{v}(\mu, L)$ and $\widehat{v}(\mu, H)$ into (3),

$$F'(\mu_1) = F'(\mu_2) + c = \frac{F(\mu_2) - F(\mu_1) + b - c(1 - \mu_2)}{\mu_2 - \mu_1}.$$

Solving, we have

$$\begin{aligned} c &= F'(\mu_1) - F'(\mu_2) \\ b &= (F'(\mu_1)(1 - \mu_1) + F(\mu_1)) - (F'(\mu_2)(1 - \mu_2) + F(\mu_2)). \end{aligned} \tag{4}$$

Because $F(\cdot)$ is concave, $c > 0$. Furthermore, note that

$$\frac{d}{dx}\{F'(x)(1-x) + F(x)\} = F''(x)(1-x) < 0,$$

where the last inequality follows from concavity of $F(\cdot)$. By implication $b > 0$.

Step 2. We show that $\mu_1 < \mu^\dagger < \mu_2$ for (b, c) in (4). By routine simplification,

$$1 - \frac{b}{c} - \mu_1 = \frac{\mu_2 - \mu_1}{F'(\mu_1) - F'(\mu_2)} \left(\frac{F(\mu_2) - F(\mu_1)}{\mu_2 - \mu_1} - F'(\mu_2) \right) > 0,$$
$$\mu_2 - \left(1 - \frac{b}{c}\right) = \frac{\mu_2 - \mu_1}{F'(\mu_1) - F'(\mu_2)} \left(F'(\mu_1) - \frac{F(\mu_2) - F(\mu_1)}{\mu_2 - \mu_1} \right) > 0,$$

where both inequalities follow from concavity of $F(\cdot)$.

To complete the proof, note that the values of the parameters (b, c) in (4) are equivalent to $k = c/(1+b) > 0$ and $\nu = b/(1+b) \in (0, 1)$. ■

5.4 PLATFORM DESIGN

Proof of Lemma 1. *A-first.* Suppose that product A is sequenced first. The high search cost consumer always purchases product A, probability h. The low cost consumer purchases A if she learns that $\omega = 1$, probability $(1-h)r\mu$, and she does not purchase A if she learns that $\omega = 0$. If the low search cost consumer is uninformed, probability $(1-h)(1-r)$, then she purchases A if and only if her expected payoff of buying A exceeds her payoff of searching and buying B, i.e., $\mu \geq \theta_A \equiv u - c$. Thus, the probability that the consumer buys A when it is sequenced first is $(1-h)[\mathcal{I}(\mu \geq \theta_A)(1-r) + r\mu] + h$. Multiplying by $(1-h)^{-1}$ delivers the expression in the text. Obviously, $\theta_A \in (0, u)$ and is decreasing in c.

B-first. Suppose that product B is sequenced first. The high search cost consumer buys B. Consider the decision of the low search cost consumer. If she buys B in stage 1, her payoff is $V_1 \equiv u$. If she searches and learns then she will buy the better product. If she remains uninformed, then she buys A if $\mu \geq u$ and B otherwise. Hence, her expected payoff of searching is

$$V_2(\mu) = r(\mu + (1-\mu)u) + (1-r)\max\{u, \mu\} - c.$$

$V_2(\cdot)$ is increasing. Thus, the consumer searches if and only if $\mu \geq \theta_B$, where $V_2(\theta_B) = u$. Furthermore, $V_2(u) = u + r(u - u^2) - c < u$ and $V_2(1) = 1 - c > u$. Therefore, $\theta_B \in (u, 1)$.

Using $\theta_B > u$, we have

$$r(\theta_B + (1-\theta_B)u) + (1-r)\theta_B - c = u \Rightarrow \theta_B = \frac{u(1-r)+c}{1-ru}.$$

Clearly θ_B is increasing in c. Thus, the low cost consumer searches if and only if $\mu \geq \theta_B$, and she buys A if she learns that $\omega = 1$ or if she is uninformed, probability $r\mu + (1-r)$. Hence, the probability of buying A is $(1-h)\mathcal{I}(\mu \geq \theta_B)(r\mu + (1-r))$. Multiplying by $(1-h)^{-1}$ gives the expression in the text. ■

Proof of Proposition 8. *Positive vs. Negative Sequencing.* We show that a $\widehat{\mu} \in (0, \theta_A)$ exists, such that $V_{NS}(\mu_0) > V_{PS}(\mu_0)$ if and only if $\mu_0 < \widehat{\mu}$. Note that with PS, $\tau(\theta_A) = \mu_0/\theta_A$, and with NS, $\tau(\theta_B) = \mu_0/\theta_B$. The platform's normalized expected profit with positive sequencing at belief μ_0 is

$$V_{PS}(\mu_0) = \tau(\theta_A)\ \widehat{v}(\theta_A, A1) + \tau(0)\ \widehat{v}(0, B1) = \mu_0 \frac{1-r+r\theta_A+f}{\theta_A} = \left(r + \frac{1-r+f}{\theta_A}\right)\mu_0,$$

and with negative sequencing,

$$V_{NS}(\mu_0) = f + \mu_0 \frac{1-r+r\theta_B-f}{\theta_B} = f + \mu_0\left(r + \frac{1-r-f}{\theta_B}\right).$$

Both $V_{PS}(\cdot)$ and $V_{NS}(\cdot)$ are linear in μ, and $V_{NS}(0) > V_{PS}(0)$. What remains is to show $V_{PS}(\theta_A) > V_{NS}(\theta_A)$. Recall that $\theta_A < \theta_B$. We have,

$$V_{PS}(\theta_A) - V_{NS}(\theta_A) = (1-r)\left(1 - \frac{\theta_A}{\theta_B}\right) + \frac{\theta_A}{\theta_B} f > 0.$$

Comparative Static. From Lemma 1, an increase in c reduces θ_A, which increases $V_{PS}(\mu_0)$. From Lemma 1, an increase in c increases θ_B, which reduces $V_{NS}(\mu_0)$.

Consumer Welfare. First consider the consumer's ex ante expected welfare when the platform always sequences A-first as it does without commitment,

$$U_{SA} = (1-h)\{r\mu_0 + r(1-\mu_0)(u-c) + (1-r)(u-c)\} + h\mu_0.$$

A high cost consumer always buys A; a low cost consumer buys A if she learns $\omega = 1$, buys B if she learns $\omega = 0$, and buys B if she does not learn ($\mu_0 < \theta_A = u - c$).

Second, consider the consumer's payoff when the platform uses positive sequencing,

$$U_{PS} = (1-h)\{\tau_{PS}(0)u + \tau_{PS}(\theta_A)(r\theta_A + r(1-\theta_A)(u-c) + (1-r)\theta_A)\}$$
$$+h\{\tau_{PS}(0)u + \tau_{PS}(\theta_A)\theta_A\}.$$

The high cost consumer buys B when $\mu_{B1} = 0$ is realized and buys A when $\mu_{A1} = \theta_A$ is realized. The low cost consumer buys B when $\mu_{B1} = 0$ is realized. When $\mu_{A1} = \theta_A$ is realized, she buys A if she learns it matches ($\omega = 1$) or if she is uninformed, and she buys B if she is informed that $\omega = 0$.

Third, consider the consumer's payoff when the platform uses negative sequencing,

$$U_{NS} = (1-h)\{\tau_{NS}(0)(u-c) + \tau_{NS}(\theta_B)(r\theta_B(1-c) + r(1-\theta_B)u + (1-r)u)\}$$
$$+h\{\tau_{NS}(\theta_B)u\}.$$

The high cost consumer buys A when $\mu_{A1} = 0$ is realized and buys B when $\mu_{B1} = \theta_B$ is realized. The low cost consumer buys B when $\mu_{A1} = 0$ is realized. When $\mu_{B1} = \theta_B$ is realized, she buys A if she learns it matches and B if she learns that A mismatches. If she is uninformed, then she searches. Recall that if B is sequenced first and the consumer's belief is θ_B, then she is indifferent between buying B and searching, and therefore, her payoff is u.

To compare U_{PS} and U_{SA}, recall that $\theta_A = u - c$ (see proof of Lemma 1). With a straightforward rearrangement,

$$U_{PS} = (1-h)\{\tau_{PS}(0)(u-c) + \tau_{PS}(\theta_A)(r\theta_A + r(1-\theta_A)(u-c) + (1-r)(u-c))\}$$
$$+(1-h)\tau_{PS}(0)c + h\{\tau_{PS}(0)u + \tau_{PS}(\theta_A)\theta_A\}.$$

By Bayes-Plausibility, $\tau_{PS}(\theta_A)\theta_A = \mu_0$. Moreover, $\tau_{PS}(0) + \tau_{PS}(\theta_A) = 1$, and hence $\tau_{PS}(0) + \tau_{PS}(\theta_A)(1-\theta_A) = 1 - \mu_0$. Simplifying,

$$U_{PS} = (1-h)\{r\mu_0 + r(1-\mu_0)(u-c) + (1-r)(u-c)\} + (1-h)\tau_{PS}(0)c + h\{\tau_{PS}(0)u + \mu_0\}$$
$$= U_{SA} + (1-h)\tau_{PS}(0)c + h\tau_{PS}(0)u.$$

To compare U_{NS} and U_{SA}, note that with a straightforward rearrangement,

$$U_{NS} = (1-h)\{\tau_{NS}(0)(u-c)+\tau_{NS}(\theta_B)(r\theta_B + r(1-\theta_B)(u-c) + (1-r)(u-c)\} + h\mu_0$$
$$+(1-h)\tau_{NS}(\theta_B)\{-r\theta_B c + r(1-\theta_B)c + (1-r)c)\} + h\{\tau_{NS}(\theta_B)u - \mu_0\}.$$

By Bayes-Plausibility, $\tau_{NS}(\theta_B)\theta_B = \mu_0$ and $\tau_{NS}(0) + \tau_{NS}(\theta_B)(1-\theta_B) = 1-\mu_0$, and furthermore $\tau_{NS}(0) + \tau_{NS}(\theta_B) = 1$. Simplifying,

$$U_{NS} = U_{SA} + \tau_{NS}(\theta_B)(1-h)\{c(1-2r\theta_B) - f(\theta_B - u)\}.$$

Details of both simplifications are available upon request. Recall from Lemma 1 that $\theta_B > u$. Therefore, the term in braces is decreasing in f. For $f < \bar{f} \equiv 1-r$, the term in braces is greater than

$$c(1-2r\theta_B) - (1-r)(\theta_B - u) = \frac{2r}{1-ru}(c + u(1-r))\left(\frac{1-u}{2} - c\right) > 0,$$

where we substituted for $\theta_B = (u(1-r)+c)/(1-ru)$ from Lemma 1. Hence, $U_{NS} > U_{SA}$. ■

Proof of Proposition 9. We show that the consumer's ex ante expected payoff is higher with the optimal sequencing algorithm than with the optimal recommendation system. Consider the consumer's payoff when the platform uses the optimal recommendation system,

$$U_R = (1-h)\{\tau_R(0)(u-c) + \tau_R(\theta_A)(r\theta_A + r(1-\theta_A)(u-c) + (1-r)\theta_A)\} + h\mu_0.$$

The high cost consumer always buys A, since it is always sequenced first, obtaining expected payoff μ_0. The low cost consumer buys B if belief $\mu_L = 0$ is realized. If belief $\mu_H = \theta_A$ is realized, then she buys A unless she learns that it is a mismatch. Recall that $\theta_A = u - c$,

$$U_R = (1-h)\{\tau_R(0)(u-c) + \tau_R(\theta_A)(r\theta_A + r(1-\theta_A)(u-c) + (1-r)(u-c))\} + h\mu_0.$$

By Bayes-Plausibility, $\tau_R(\theta_A)\theta_A = \mu_0$ and $\tau_R(0) + \tau_R(\theta_A)(1-\theta_A) = 1-\mu_0$, and furthermore $\tau_R(0) + \tau_R(\theta_A) = 1$. Substituting,

$$U_R = (1-h)\{r\mu_0 + r(1-\mu_0)(u-c) + (1-r)(u-c))\} + h\mu_0 = U_{SA}.$$

From Proposition 8, $U_{PS} > U_{SA}$ and $U_{NS} > U_{SA}$. ■