INDIRECT PERSUASION

Rahul Deb*
Mallesh M. Pai†
Maher Said‡

March 6, 2025

ABSTRACT

We provide an organizational economics foundation for commitment to information structures in persuasion. An uninformed principal faces a joint screening-and-persuasion problem: she wants to influence a receiver's belief about a payoff-relevant state using information elicited from a privately informed agent. The principal cannot act as an intermediary that commits to an optimal garbling of the agent's private communications; instead, the agent's messages are publicly observed by the receiver. We show that the principal can still (indirectly) implement the optimal unconstrained intermediation scheme. Commitment to an employment contract with the agent alone suffices for optimal persuasion of the receiver. We apply our result to the context of a brokerage contracting with a sell-side analyst, where private communication is constrained by conflict-of-interest regulations. We show that a public communication scheme—which closely corresponds to the investment ratings schemes observed in practice—can sidestep these regulations.

*Department of Economics, Boston College, rahul.deb@bc.edu
†Department of Economics, Rice University, mallesh.pai@rice.edu
‡Stern School of Business, New York University, maher.said@nyu.edu
We are grateful to Anna Sanktjohanser and Thomas Brzustowski for their thoughtful discussions of early versions of this paper, as well as to Heski Bar-Isaac, Joyee Deb, Laura Doval, Michael Grubb, Emir Kamenica, Fei Li, Volker Nocke, Eduardo Perez-Richet, Ludovic Renou, Vasiliki Skreta, Ina Taneva, four anonymous referees, and participants at various seminars and conferences for their many helpful comments and suggestions. We would also like to thank CEPR and the Studienzentrum Gerzensee, where much of this work was completed, for their hospitality. Deb thanks the SSHRC for their generous financial support. Pai gratefully acknowledges financial support from NSF grant CCF-1763349. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

1. INTRODUCTION

Bayesian persuasion problems consider a sender's manipulation of information to induce a receiver to take some desired actions. In these models, the sender communicates via commitment to an information structure. Similarly, in contract theory and mechanism design, it is typically assumed that the principal can commit to contractual terms—in moral hazard problems, commitment is to output-contingent payments, while in adverse selection problems, the principal commits to the actions that follow the revelation of private information. This commitment assumption is mostly uncontroversial in these latter contexts and is frequently motivated by mapping it to codified rules within an organization via, for example, its human resource policies. Unlike in contract theory and mechanism design however, commitment is harder to justify in the context of strategic communication and persuasion. Indeed, as Kamenica, Kim, and Zapechelnyuk (2021) observe: "Optimal information structures can be infeasible or difficult to implement in practice. A commitment to randomized messages is difficult to verify and enforce...."

In this paper, we provide an organizational economics microfoundation for the commitment assumption in communication. Unlike standard models of Bayesian persuasion that follow Kamenica and Gentzkow (2011), our setting is not one of a sender with unrestricted access to arbitrary information structures. Instead, our sender is an *uninformed principal* who aims to persuade a *receiver* by eliciting the private information of a *strategic agent*. The agent is purely concerned with the terms governing their contractual relationship with the principal. The principal's payoff depends on both the contractual relationship with the agent and the receiver's chosen action. We also assume, motivated by applications, that any communication by the agent must be public; that is, the agent's messages must be observed by both the principal and the receiver. The principal therefore faces a novel joint screening-and-persuasion problem: they must address both their employment relationship with the agent and their need to persuade the receiver using a contract that only elicits public messages from the agent. Our main result shows that the optimal persuasion of the receiver is obtainable by commitment to a standard employment contract without also requiring commitment to arbitrary information structures. This result allows us to reinterpret the commitment assumption in Bayesian persuasion as a reduced-form stand-in for an informed intermediary facing optimal incentives designed by their employer.

In our model, the agent has a privately known type whereas the principal and the receiver are uninformed. The payoffs of all three players depend on the agent's private type. In addition, the agent's utility depends on the contractual terms with the principal, while the receiver's utility depends on their action. The principal's payoff depends on both the contractual terms and the receiver's action, and it is additively separable across her interactions with the agent and the receiver.

We study the following extensive-form public communication game. The principal first publicly selects a message space and a set of contracts that are observed by both the agent and the receiver. The principal then chooses and commits to a contract from this set, possibly as the result of a mixed strategy. Only the agent observes the selected contract, while the receiver does not. The contract specifies the agent's contractual terms as a function of her public messages. The agent publicly announces a message from the message space; this message is observed by all players. After observing the message, the receiver updates her belief about the agent's type and, finally, chooses an action.

This class of (indirect) public communication mechanisms has several desirable properties. First, the principal need only commit to a standard employment contract with the agent and not to an arbitrary information structure. In addition, there is no private communication from the agent to the principal. As we highlight in our main application below, compliance regulations may prohibit such communication or it simply may not be feasible, and so such mechanisms are commonly observed in practice.

Our main result shows that there is a sequential equilibrium of this public communication game that achieves the same outcome as the full-commitment benchmark where the principal directly and privately intermediates between the agent and the receiver. In this benchmark, the agent truthfully reports their type privately to the principal, whose mechanism specifies, based on that report, both the agent's employment terms and a distribution over recommended receiver actions.

When communication by the agent is public, however, a conflict arises: the agent's public message directly influences *both* their employment terms *and* the receiver's action, potentially creating a tradeoff for the principal between screening and persuasion. For example, optimal screening might require the agent to publicly reveal "more" (or "more precise") information than is optimal for persuading the receiver. Despite this possibility, however, we show that the principal can still achieve both optimal screening and optimal persuasion—without direct access to private information and the ability to commit to an optimal information structure.

Our argument leverages the fact that the receiver observes the agent's public message but not the principal's choice of contract. In particular, when the principal follows a mixed strategy, the agent knows the realized mapping from messages to employment terms but the receiver does not. This uncertainty prevents the receiver from inverting the agent's strategy and essentially serves as a method of publicly garbling the agent's information while preserving its private meaning. Of course, not all garblings are feasible: any contract chosen with positive probability must maximize the principal's expected payoff. We show, however, that there exists an equilibrium randomization over contracts by the principal that allows for both optimal screening of the agent and optimal persuasion of the receiver (as in the benchmark). A consequence is that the deliberate introduction of vague language is often necessary for indirect persuasion to be optimal.

We apply this insight to the market for sell-side financial research. This application both motivates key features of our model and also delivers a surprising economic take away that we view to be of independent interest. Sell-side financial analysts are the preeminent financial market information intermediaries. They gather and analyze information, and then produce forecasts and recommendations for the investment community.[1] These analysts are employed by banks and brokerages who thus face a conflict of interest: they have an employment relationship with the analyst to manage (deciding, for instance, whether to promote or dismiss analysts based on ability or performance), but also wish to persuade investors to take actions that may benefit the bank (via commissions, brokerage fees, or the like).[2] Consequently, this industry is highly regulated: FINRA Rule 2241 in the US and MiFID II in the EU both prohibit direct interaction between banks' investment and research arms. This is what we aim to capture with our model's assumption of public communication by the agent: banks are not permitted to intermediate the communications between their analysts and their research clientele. An immediate consequence of our main result is that such regulation can be circumvented or rendered less effective by a bank that appropri-

[1]There is a vast and thorough empirical literature in finance analyzing various aspects of analysts. Bradshaw, Ertimur, and O'Brien (2017) is an excellent recent survey that describes what analysts do and how they have been affected by regulation, yet there is a paucity of theoretical work aimed at understanding how recommendations are influenced by career incentives and how banks that employ analysts provide these incentives.

[2]Generating trading fees is a well-documented role of analysts and indeed is often cited as a potential conflict of interest; see, for example, the discussion in the survey by Bradshaw (2011).

ately designs its analysts' employment contracts. Interestingly, our equilibrium takes a natural form in which the analyst's recommendations correspond to the commonly observed asset rating scale (typically "strong buy," "buy," "hold," and the like). Finally, we show that a regulatory intervention that forces the analyst's employment contract to be public can strictly increase client welfare.

While our application focuses on financial analysts, it is worth emphasizing that similar issues arise in any organization that hires experts to provide information and advise clients. As with financial analysts, the advice experts choose to provide is determined by their career incentives; these incentives may not align with their employers' preferences; and, critically, the employer may not be able to directly control the advice provided. For instance, consulting firms cannot directly control what information their consultants convey in response to clients' questions during meetings. A prosecutor wishing to convince a judge that a defendant is guilty (as in Kamenica and Gentzkow's (2011) canonical example) may hire an independent investigator to uncover evidence but then cannot control the investigator's findings and public testimony before the judge. Likewise, firms employing litigation consultants or expert witnesses face analogous conflicts; such strategic information intermediaries are ubiquitous.

1.1. RELATED LITERATURE

Our paper contributes to the literature studying commitment in Bayesian persuasion, where a receiver must interpret messages while accounting for the sender's *partial* commitment.[3] Several papers—most notably, Min (2021), Fréchette, Lizzeri, and Perego (2022), and Lipnowski, Ravid, and Shishkin (2022)—study a sender who initially commits to an information structure but, with an exogenously given probability, is released from her commitment and can send any other message. Others (Nguyen and Tan (2021) and Perez-Richet and Skreta (2022), for instance) instead permit the sender to deviate through costly manipulations of her information structure's inputs or outputs, while Lin and Liu (2024) limit these deviations to only those that leave the distribution of messages unchanged.

We take a different view by assuming that direct commitment to an information structure is not possible but instead show that standard organizational contracts can suffice for optimal persuasion.[4] In this way, we relate to papers that study persuasion with incentivized communication and intermediation. For example, Bizzotto, Perez-Richet, and Vigier (2021) study a setting where a principal uses monetary transfers to induce a third party to communicate the output of an information structure to a receiver. In Salamanca (2021) and Corrao and Dai (2024), a mediator designs a communication mechanism that elicits information from the sender to then transmit to the receiver. This work also shares features with the literature studying persuasion by an informed principal, including Perez-Richet (2014), Koessler and Skreta (2023), and Zapechelnyuk (2023). In contrast, we study an uninformed principal who is limited in the mechanisms she can use to *publicly* elicit and communicate information.[5]

The equilibrium we construct makes use of deliberately "vague" public communication to generate uncertainty and persuade the receiver.[6] A similar idea appears in the literature on mechanism design and

[3]The excellent surveys of Bergemann and Morris (2019) and Kamenica (2019) describe the broader literature.

[4]Some recent papers use ideas from repeated games and reputation to provide foundations for commitment: see, for example, Best and Quigley (2024) and Mathevet, Pearce, and Stacchetti (2024).

[5]Similar to our main result, Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017) also show an equivalence between public and private communication in persuasion games; in their environment, however, it is the *receiver* who is privately informed.

[6]The resulting one-sided uncertainty is similar to the "private communication in public" results in Antic, Chakraborty, and Harbaugh (2024); see also Strack and Yang (2024), who study the design of privacy-preserving signal structures.

communication with ambiguity-averse agents. For example, Bose and Renou (2014) show that the deliberate introduction of ambiguity into mediated communication can enlarge the set of implementable social choice functions. Beauchêne, Li, and Li (2019) also use "synonymous" messages to generate uncertainty and manipulate an ambiguity-averse receiver into taking a sender-preferred action. A recent sequence of papers—Krähmer (2020), Krähmer (2021), and Ivanov (2024)—also look at settings (without ambiguity aversion) where randomization over information structures can expand the set of outcomes in communication games. Similar to our work, private randomization whose realization is not observed by a relevant decision maker is a key ingredient; in contrast to our equilibrium construction, however, these works rely on commitment to that private randomization.

In our application, the contracting terms offered by the principal serve dual functions: in addition to indirectly persuading the receiver, they also screen the agent's skill.[7] The persuasion motive is akin to that in Inderst and Ottaviani (2012), who study how the design of commissions can influence financial advisors and steer their recommendations. These competing incentives also appear in Jackson (2005) and Beyer and Guttman (2011), who analyze models of trade generation by sell-side analysts with reputational concerns—but set aside the bank's organizational design problem and simply take analyst incentives to be exogenously fixed.

Meanwhile, the principal's screening motive connects our work to the comparatively small literature studying how strategic experts should be evaluated. There is an extensive literature on the statistical evaluation of forecasting models (see the work cited in Elliott and Timmermann (2016), for instance), but relatively less work examining the incentives faced by *strategic* experts who are potentially influenced by market or career incentives (see Marinovic, Ottaviani, and Sørensen (2013) for a survey of this literature). While many of the theoretical contributions in this latter area—Ottaviani and Sørensen (2006a,b,c) are particularly prominent examples—study environments where the experts' incentives are exogenously (and often suboptimally) given, we follow our earlier work in Deb, Pai, and Said (2018) and focus on the *design* problem faced by a principal that wants to separate a skilled from unskilled agent. Though this is a secondary goal of the paper, the characterization of the optimal screening contract in our application (Theorem 2) is of independent interest; see also Dasgupta (2023), who studies this question in the context of test design.

Lastly, our work builds on and contributes to the extensive finance literature studying sell-side analysts; the aforementioned Bradshaw, Ertimur, and O'Brien (2017) surveys much of this research. We will discuss the other related work in this literature when we present our application in Section 4.

2. MODEL

We study an extensive-form game with three players: a principal, an agent, and a receiver.

INFORMATIONAL ENVIRONMENT The agent's *type* θ is her private information, drawn from a distribution $\pi \in \Delta(\Theta)$, where Θ is a finite set of possible types. This type can, for instance, incorporate both information about an underlying state of the world and about the agent's ability or preferences (our application to financial analysts takes this form).

Neither the principal nor the receiver are endowed with any private information.

[7] A similar tension arises in the context of dynamic mechanism design without commitment; see Doval and Skreta (2022), who reinterpret sequential rationality as a principal's persuasion of their future self.

PUBLIC COMMUNICATION MECHANISMS A *public communication mechanism* is a pair $(\mathcal{M}, x)$ consisting of a finite *message space* $\mathcal{M}$ and a *contract* $x : \mathcal{M} \to \Delta(\mathcal{T})$, where $\mathcal{T}$ is the set of *contractible decisions*. Elements of the set $\mathcal{T}$ are the agent-specific contractible decisions that the principal can commit to; these might represent transfers, retention probabilities, or some other general contractible outcomes.

PRINCIPAL The principal chooses a finite message space $\mathcal{M}$ and a finite set of contracts $\mathbb{X} \subset \Delta(\mathcal{T})^{\mathcal{M}}$, yielding a finite set of public communication mechanisms $\mathbb{M} := \{(\mathcal{M}, x) \mid x \in \mathbb{X}\}$. In addition, the principal chooses a distribution $\rho \in \Delta(\mathbb{X})$ over the set of contracts. We assume that the set of contracts $\mathbb{X}$ is publicly observed by both the agent and the receiver, while the principal's chosen distribution ρ is private and unobservable. In particular, the agent only observes the realization of this distribution. Thus, there is common knowledge of the set of possible mechanisms, but only the principal and agent know which specific contract governs their relationship.

AGENT After learning her type $\theta \in \Theta$ and observing the selected public communication mechanism $(\mathcal{M}, x)$, the agent chooses to report a message $m \in \mathcal{M}$. The agent's strategy is a (θ-dependent) distribution over possible reports in the message space $\mathcal{M}$.

RECEIVER Upon observing the agent's reported message $m \in \mathcal{M}$, the receiver updates their *belief* to $q(m) \in \Delta(\Theta)$ about the agent's type θ; they then choose an *action* $a \in \mathcal{A}$. The receiver's strategy is thus a map from messages to distributions over actions.

PAYOFFS All players are expected utility maximizers, with *payoffs* given by $u_P : \Theta \times \mathcal{T} \times \mathcal{A} \to \mathbb{R}$ for the principal, $u_A : \Theta \times \mathcal{T} \to \mathbb{R}$ for the agent, and $u_R : \Theta \times \mathcal{A} \to \mathbb{R}$ for the receiver.

All three players' payoffs depend on the agent's type. The principal's payoff additionally depends on *both* the contracted decision and the receiver's action. The agent's payoffs also depend only on the contracted decision but not the receiver's action, while conversely the receiver's payoffs also depend only on her action but not the principal–agent contractual decision.

Importantly, we further assume that the principal's payoff function is additively separable across her interactions with the agent and the receiver, so that we can write

$$u_P(\theta, t, a) = u_P^A(\theta, t) + u_P^R(\theta, a).$$

Recall that our primary goal is to situate a "standard" persuasion problem in an organizational context and show that the principal can optimally persuade the receiver using only commitment to a standard employment contract with the agent (instead of commitment to arbitrary information structures). The assumption on payoffs is such that, were the principal to directly observe θ and have the ability to commit to an information structure, the resulting action choice by the receiver would correspond to the optimal persuasion outcome where the principal and receiver have payoffs $u_P^R(\theta, a)$ and $u_R(\theta, a)$ respectively.

EQUILIBRIUM We focus on the sequential equilibria of this public communication game. The *outcome* of a sequential equilibrium is the joint distribution over contractual decisions and receiver actions for each agent type $\theta \in \Theta$.

SUMMARY To summarize and make the timing explicit, Figure 1 presents a flow chart depicting the public communication game we study.

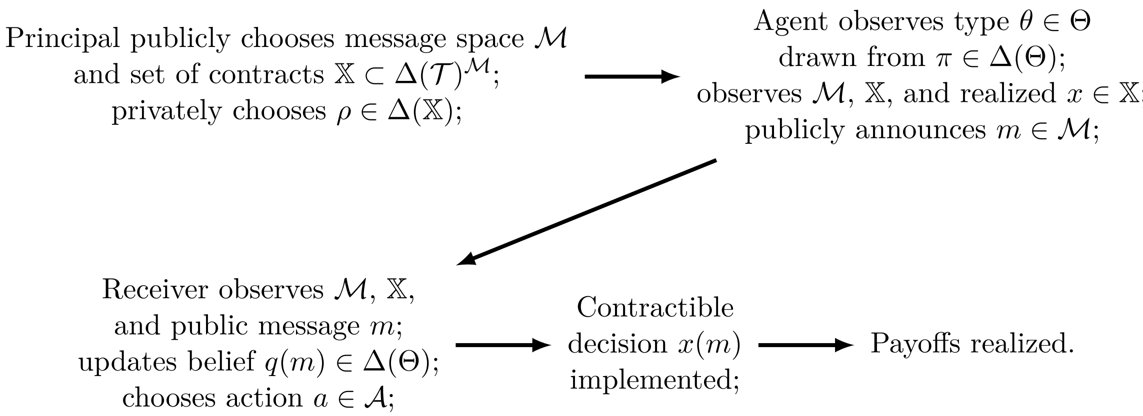

Figure 1: Timing of the public communication game.

3. ACHIEVING THE FULL-COMMITMENT OPTIMUM

In the public communication game described above, the agent's messages are observed by *both* the principal and the receiver. Thus, the principal's chosen public communication mechanism must balance her contracting goals with the agent and her desire to persuade the receiver. Contrast this with the case where the agent can communicate privately with a principal who has full commitment power: the principal could separately commit to contracting terms that maximize $u_P^A(\theta, t)$ and to an information structure that optimally garbles the agent's report to maximize her persuasion payoff $u_P^R(\theta, a)$. We now formally describe this full-commitment benchmark before presenting our main result showing that there is a sequential equilibrium of the public communication game that replicates this full-commitment outcome.

Following Myerson (1986), the revelation principle implies that it suffices to consider the class of incentive compatible direct mechanisms. A *direct mechanism* is a mapping $\chi : \Theta \to \Delta(\mathcal{T} \times \mathcal{A})$, which we decompose into its marginals: a *contracting rule* $X : \Theta \to \Delta(\mathcal{T})$ and an *action recommendation rule* $A : \Theta \to \Delta(\mathcal{A})$. In words, the joint distribution over contractible decisions and action recommendations is determined by the agent's reported type.

A direct mechanism χ is *incentive compatible* if it satisfies the following two sets of constraints:

- For any realized type $\theta \in \Theta$, it is optimal for the agent to report her type truthfully instead of misreporting it as some other type $\theta' \in \Theta$; that is,

$$\mathbb{E}_{X(\theta)}[u_A(\theta, t)] \geq \mathbb{E}_{X(\theta')}[u_A(\theta, t)] \text{ for all } \theta, \theta' \in \Theta, \quad \text{(IC-A)}$$

  where the expectations are taken over the contracting decision $t \in \mathcal{T}$, the distribution of which depends on the reported type via the contracting rule $X(\cdot) \in \Delta(\mathcal{T})$.

- It is optimal for the receiver to obediently choose any recommended action a that lies in the support of $A(\cdot) \in \Delta(\mathcal{A})$ for any $\theta \in \Theta$; that is,

$$\mathbb{E}_{q(a)}[u_R(\theta, a)] \geq \mathbb{E}_{q(a)}[u_R(\theta, a')] \text{ for all } a \text{ in the support of } A(\cdot) \text{ and all } a' \in \mathcal{A}, \quad \text{(IC-R)}$$

  where the expectations are taken over the agent's type $\theta \in \Theta$, given the receiver's posterior belief

$q(a) \in \Delta(\Theta)$ formed via Bayes' rule after observing action recommendation a.

Notice that, for any incentive compatible direct mechanism χ, every other direct mechanism χ' with the same marginals is also incentive compatible (since the constraints (IC-A) and (IC-R) are defined only in terms of marginals). The assumptions we impose on payoffs imply that any such χ' also delivers the same expected payoffs to all three players. This yields the following observation.

OBSERVATION. *It is without loss of generality to restrict attention to* product-separable *direct mechanisms that can be written as the product of their marginals; that is, direct mechanisms $\chi : \Theta \to \Delta(\mathcal{T} \times \mathcal{A})$ with marginals $X : \Theta \to \Delta(\mathcal{T})$ and $A : \Theta \to \Delta(\mathcal{A})$ such that $\chi(\theta) = X(\theta) \times A(\theta)$ for all $\theta \in \Theta$.*

Consequently, in searching for an *optimal* direct mechanism χ^* that maximizes the principal's payoff subject to the incentive compatibility constraints (IC-A) and (IC-R), we can separately and independently choose a contracting rule X^* that maximizes $u_P^A(\theta, t)$ subject to (IC-A) and an action recommendation rule A^* that maximizes $u_P^R(\theta, a)$ subject to (IC-R). In particular, the latter implies that the principal's optimal action recommendation A^* would remain optimal if the principal directly observed—without frictions—the agent's private information θ. In other words, A^* is the solution to the persuasion problem where the principal and receiver have payoffs $u_P^R(\theta, a)$ and $u_R(\theta, a)$ respectively. It is in this sense that our framework nests the standard persuasion setting of Kamenica and Gentzkow (2011).

3.1. THE MAIN RESULT

Direct mechanisms are convenient theoretical objects for deriving the full-commitment upper bound on the principal's attainable payoffs. In practice, however, a direct mechanism may have undesirable features. In settings like our application, for instance, there may be strong norms (legal or otherwise) precluding the principal's intermediation between the agent and the receiver. Moreover, in contexts where the agent's type encodes information about their innate abilities, they may be unwilling to explicitly announce to an employer that they are unskilled, preferring instead indirect mechanisms that may preserve some plausible uncertainty. Finally, a direct mechanism will typically require commitment to a mixed action recommendation as a function of the agent's information. As discussed in the introduction, such commitment (equivalently, to information structures) is frequently criticized on the grounds that it is a theoretical assumption with no real-world analogue.

By contrast, our main result shows that the principal can achieve their full-commitment payoff using *indirect* public communication mechanisms. These only require commitment to contractual terms (not information structures) and do not involve any private communication from the agent to the principal.

Given a direct mechanism χ with associated contracting rule X and action recommendation rule A, a sequential equilibrium of the public communication game *implements* χ if the outcome of this equilibrium is $X(\theta) \times A(\theta)$ for every agent type $\theta \in \Theta$. Note that this implementation criterion is stronger than payoff equivalence for all three players.[8]

[8]Note that the direct mechanism χ may, for some $\theta \in \Theta$, correlate the randomization in the contracting and action recommendation rules X and A. In this case, the implementing equilibrium features the same marginals X and A, but with independent distributions. For every $\theta \in \Theta$, the implementing equilibrium yields the same payoffs for all three players and, additionally, an identical outcome distribution whenever $\chi(\theta) = X(\theta) \times A(\theta)$.

THEOREM 1. *Every optimal direct mechanism χ^* is implementable by a sequential equilibrium of the public communication game.*

The formal proof of this result is in Appendix A, but we provide intuition here. In order to implement the optimal contracting rule X^*, the principal needs to elicit the agent's private information *without* conveying it perfectly to the receiver, as this may conflict with optimal persuasion according to A^*. To do so, the principal can choose a public message space $\mathcal{M}^*$ containing multiple "synonyms" for every recommended action in the support of A^*. By mixing over a set $\mathbb{X}^*$ of contracts that rely on permuted mappings from $\mathcal{M}^*$ to contractible decisions, synonymous messages can be garbled and endowed with a common *public* meaning that implements the recommendation rule A^*. But because the agent observes the realized choice of public communication mechanism, she and the principal have a shared *private* meaning for each synonym that simultaneously implements the contracting rule X^*.

Note that the message space $\mathcal{M}^*$ described above may contain more elements than just the set of actions recommended to the receiver. This is in contrast to "standard" Bayesian persuasion settings, where one typically requires only as many messages as on-path actions. But because the communication here is multivalent and directed towards multiple audiences, a richer language is necessary. Moreover, the construction above suggests that effective indirect persuasion may require the use of intentionally *vague* language: the principal would like the agent's message to reveal her type for screening purposes, but simultaneously needs the receiver to be uncertain about the mapping from messages to agent types for persuasion. This uncertainty about the interpretation of language is one-sided: the principal and the agent share, through their knowledge of the realized contract, an understanding of how public messages are *privately* interpreted. This shared understanding requires commitment to—and the privacy of—contractual terms. But, conversely, no ability whatsoever to commit to information structures is required.

It bears mention that our indirect implementation does not rely on implausible or pathological off-path beliefs. By publicly committing to a message space $\mathcal{M}$ and a set of contracts $\mathbb{X}$, the principal essentially acts as a mechanism designer whose payoff is bounded above by the optimal direct mechanism. But since it is precisely this optimum that is implemented on-path—regardless of the realization of the principal's mixed strategy—there is no incentive for the principal to deviate and attempt to implement any other (suboptimal) mechanism.[9] This of course relies on the observability of $\mathcal{M}$ and $\mathbb{X}$, which precludes an unobservable (to the receiver) deviation to a contract that, for instance, compensates the agent directly for inducing specific actions.

3.2. DISCUSSION

While Theorem 1 is a general result, it is worth discussing the key assumptions driving it. Recall that the principal's payoff $u_P^A(\theta, t) + u_P^R(\theta, a)$ is separable in the contractible decision t and the receiver's action a. As noted following the observation above, this payoff separability ensures that it is without loss to assume the optimal direct mechanism is product separable, with conditionally independent contractible decision and action recommendation rules. Note, however, that Theorem 1 continues to hold even for non-separable principal payoffs $u_P(\theta, t, a)$ as long as they admit a product-separable optimal direct mechanism. In other words, our payoff assumptions are natural conditions on primitives that are sufficient (but not necessary) for our result.

[9]Note that this argument also implies that the agent is indifferent between all contracts that might be chosen on-path; therefore, we can also implement the optimal mechanism by letting the *agent* select a contract from the principal's set $\mathbb{X}$.

To see why product separability is sufficient for our indirect implementation, recall that any mixing $X^*(\theta) \in \Delta(\mathcal{T})$ over contractible decisions is part of a principal–agent contract that can be committed to in a public communication mechanism. Likewise, any randomization $A^*(\theta) \in \Delta(\mathcal{A})$ in the receiver's action can be induced by the above-described mixing over contracts with synonymous messages. Since the agent's payoff $u_A(\theta, t)$ does not depend on a, this does not affect the agent's incentives in the public communication mechanism. Note, however, that this argument critically relies on the (conditional) independence of the randomization over actions a generated by the agent's strategy and the mixing over contractible decisions t generated by the contracts.

Consequently, Theorem 1 holds for any payoffs that admit an optimal direct mechanism with either a deterministic contracting or action recommendation rule (but not necessarily both), as any such mechanism is immediately product separable. While we do not have a general sufficient condition that ensures a deterministic optimum (in either $\mathcal{T}$ or $\mathcal{A}$), it is easy to construct examples where this is true.

To see why Theorem 1 need not hold when there is no product-separable optimal mechanism, consider a setting where the agent has two types $\Theta = \{\theta_1, \theta_2\}$, the receiver has two actions $\mathcal{A} = \{a, a'\}$, the principal has three contractible decisions $\mathcal{T} = \{t_1, t_1', t_2\}$, and the *unique* optimal direct mechanism is

$$\chi^*(\theta_1) = \begin{cases} (t_1, a) & \text{with probability } \frac{1}{2}, \\ (t_1', a') & \text{with probability } \frac{1}{2}, \end{cases} \quad \text{and } \chi^*(\theta_2) = (t_2, a') \text{ with probability } 1.$$

Suppose further that the agent's payoffs satisfy

$$u_A(\theta, t_1) > u_A(\theta, t_2) > u_A(\theta, t_1') \text{ for both } \theta \in \Theta,$$

so that both agent types prefer t_1 to t_2 to t_1', but that (as required for the incentive compatibility of χ^*),

$$\frac{1}{2}u_A(\theta_1, t_1) + \frac{1}{2}u_A(\theta_1, t_1') > u_A(\theta_1, t_2) \text{ and } u_A(\theta_2, t_2) > \frac{1}{2}u_A(\theta_2, t_1) + \frac{1}{2}u_A(\theta_2, t_1')$$

and the types differ in their preference for an equal-likelihood mixture of t_1 and t_1' relative to deterministically receiving t_2.[10]

Suppose that χ^* could be implemented as a sequential equilibrium of the public communication game. Then every contract x^* in the support of the principal's strategy must be such that $x^*(m_1) \in \{\delta_{t_1}, \delta_{t_1'}\}$ for every message m_1 in the support of θ_1's strategy and $x^*(m_2) = \delta_{t_2}$ for every message m_2 in the support of θ_2's strategy.[11] The latter condition follows immediately from the form of χ^*, while the former is true because if $x(m_1)$ was a mixture of t_1 and t_1', the receiver's action could not be perfectly correlated with the contractible decisions t_1 and t_1' as χ^* requires. Now observe that type θ_1 would never pick a message m_1 such that $x^*(m_1) = \delta_{t_1'}$ as they would prefer to pick m_2 for which $x^*(m_2) = \delta_{t_2}$. Conversely, type θ_2 would never pick a message m_2 such that $x^*(m_2) = \delta_{t_2}$ when there is another message m_1 for which $x^*(m_1) = \delta_{t_1}$. Thus, χ^* cannot be implemented as a sequential equilibrium of the public communication game.

[10] For brevity, we omit explicitly defining the principal's and receiver's payoffs as they are not required for the argument. It is straightforward, however, to specify payoffs for which the direct mechanism above is indeed the unique optimum.

[11] We use δ_t to denote the degenerate distribution that assigns probability one to the contractible decision $t \in \mathcal{T}$.

Theorem 1 also does not generalize to settings with a "biased" agent whose payoff $u_A(\theta, t, a)$ depends arbitrarily on the receiver's action. Recall that the principal effectively garbles the agent's information through randomization over contracts to ensure that a type-θ agent chooses different messages for different realizations of the contract. In a direct mechanism χ, the agent evaluates their expected utility $\mathbb{E}_{\chi(\theta)}[u_A(\theta, t, a)]$ based on the joint distribution over t and a induced by χ. Conversely, in any realized contract, each public message typically induces a unique receiver action (barring knife-edge cases where the receiver is indifferent over multiple actions). Thus, the principal may not be able to implement χ as they may not be able to guarantee agent incentives—even if each dimension of χ is independent. Nonetheless, Theorem 1 applies in cases where the agent's payoff depends on the receiver's action but there is an optimal direct mechanism whose corresponding action recommendations $A^*(\theta)$ are deterministic for all θ (even if the contracting rule X^* is not).[12]

Lastly, Theorem 1 also does not accommodate arbitrary receiver payoffs $u_R(\theta, t, a)$ that depend on the contractible decision. As with the case of nonseparable principal payoffs, optimal direct mechanisms may involve correlation between t and a. As discussed above, such correlation may not be achievable in the public communication game. But again, Theorem 1 extends to cases where the receiver payoff takes the form $u_R(\theta, t, a)$ yet there is a product-separable optimal direct mechanism such that each dimension of $\chi^*(\theta) = X^*(\theta) \times A^*(\theta)$ is independent for all $\theta \in \Theta$.

4. APPLICATION: REGULATING FINANCIAL ANALYSTS

In this section, we develop an application of Theorem 1 to the market for sell-side financial research. In addition to being of independent interest, this application serves several purposes. First, it helps motivate the public communication game we study by situating it in an important real-world context. It also permits a transparent presentation and interpretation of our result's underlying intuition, especially in light of the prevalence of vague communications in analyst ratings. Finally, the result has consequential implications for the effectiveness—or lack thereof—of conflict-of-interest regulations in this industry. We note, however, that the tensions appearing in this application are pervasive in the wide variety of environments that feature strategic information intermediaries.

As discussed in the introduction, sell-side financial analysts are important information intermediaries in financial markets, and strict compliance regulations govern their relationships with the institutions that employ them. These regulations aim to ensure that analysts' stock recommendations reflect their actual opinions about the investment potential of subject companies and are not influenced by incentives to generate trading activities, commissions, and investment banking business. These regulatory constraints correspond to our assumption of public communication and, as we will show, a consequence of Theorem 1 is that such regulation can be circumvented or rendered less effective by a bank that appropriately designs its analysts' employment contracts.

4.1. ENVIRONMENT

We now specialize the general public communication game in Section 2 to this setting. The three players are a bank (the principal), an analyst (the agent), and a representative client (the receiver) standing in for "the market" as a whole. In what follows, we deliberately choose the simplest possible version of

[12] This is, of course, subject to the caveat that all communication by the agent occurs publicly in the language specified by the principal; if the agent had a separate communication channel with the receiver, then any preference misalignment with the principal would distort the receiver's action.

the model required to illustrate our main insight. As long as the conditions of Theorem 1 are met, all assumptions can be generalized.

The analyst has a private type $\theta = (s, p) \in \Theta := \{h, l\} \times \{0, \frac{1}{n}, \frac{2}{n}, \ldots, 1\}$. The first dimension $s \in \{h, l\}$ is their *skill* or ability (either high or low). The second dimension represents their *information* about the quality of an investment opportunity: the analyst observes an (unmodeled) informative signal about the investment, and $p \in \{0, \frac{1}{n}, \frac{2}{n}, \ldots, 1\}$ is their posterior belief that it is profitable (an event we denote by b), while $1 - p$ is their complementary belief that it is unprofitable (an event we denote by ϕ).

The analyst's type is distributed according to $\pi(s, p) = \frac{1}{2} f_s(p)$, with $\sum_p f_s(p) = 1$; that is, the analyst is equally likely to be high- or low-skill, and, conditional on her skill s, her posterior p is distributed according to f_s. We assume that the distributions f_s are strictly positive and symmetric around $\frac{1}{2}$, so that $\sum_p p f_s(p) = \frac{1}{2}$ and the investment is ex ante equally likely to be profitable or not. Finally, we assume that f_h dominates f_l in Johnson and Myatt's (2006) rotation order, so that the high-skill analyst is better informed than a low-skill analyst.[13] Given our symmetry assumptions, this requires that the cumulative distributions satisfy $F_l(p) \leq F_h(p)$ whenever $p < \frac{1}{2}$, and likewise that $F_l(p) \geq F_h(p)$ whenever $p \geq \frac{1}{2}$.

The bank chooses a contractible decision $(t_b, t_\phi) \in \mathcal{T} := \{0, 1\} \times \{0, 1\}$. This captures whether the bank, after eventually observing the ex post outcome of the investment opportunity, retains the analyst when event b occurs (the first dimension t_b) or when event ϕ occurs (the second dimension t_ϕ).[14] The analyst cares solely about being retained, and so their payoff is given by

$$u_A((s, p), (t_b, t_\phi)) = p t_b + (1 - p) t_\phi.$$

We focus on the bank's retention decision, and not on monetary compensation or other contractual terms. This mirrors empirical evidence showing that analyst performance is primarily motivated by the threat of termination and not by compensation incentives. Kothari, So, and Verdi's (2016) survey of the literature on financial analysts (with further relevant references found therein) argues that banks "do not rely on forecast accuracy as a first-order determinant of annual compensation," but that there is a "strong relation between analysts' accuracy and other career outcomes"; in particular, "forecast inaccuracy can affect analyst wealth by increasing the probability of dismissal." Focusing on retention alone also permits a cleaner exposition of the indirect implementation result of Theorem 1 (which still holds even in environments with richer contracting instruments).

The client makes a trading decision $a \in \mathcal{A} := \{b, \phi\}$ to maximize her expected gains from (costly) trade

$$u_R((s, p), a) = \begin{cases} pv - c & \text{if } a = b, \\ 0 & \text{if } a = \phi. \end{cases}$$

If the client chooses $a = b$ and buys the asset, she incurs a transaction cost $c > 0$ and earns a payoff $v > c$ if the investment is profitable (the event b); if she chooses $a = \phi$ and does not buy, or if she does but the

[13]See Crane and Crotty (2020) for recent evidence on the heterogeneity in analysts' abilities to produce new information.

[14]The assumption that the ex post outcome is observable and potentially useful for evaluating the analyst is standard in models of forecasters and analysts; see, for instance, Marinovic, Ottaviani, and Sørensen's (2013) survey and the references therein.

investment is unprofitable (the event ϕ), she earns a payoff 0. This payoff structure yields a threshold

$$\bar{q} := \frac{c}{v}$$

such that the client finds it optimal to choose action b whenever the expected value she assigns to the analyst's information p is greater than, or equal to, $\bar{q}$ (we assume the client trades when she is indifferent). Otherwise she chooses action ϕ. We assume that $\bar{q} > \frac{1}{2}$ to rule out the uninteresting case where the client is willing to engage in trade even in the absence of any additional information.

Lastly, the bank has dual objectives: it earns trading fees and commissions linked to the volume of trade, and also wants to identify the analyst's ability so it only retains the high-skill analyst. Formally, the bank's preferences are given by

$$u_P((s,p),(t_b,t_\phi),a) = \kappa\mathbb{1}\{a=b\} + (pt_b + (1-p)t_\phi)(\mathbb{1}\{s=h\} - \mathbb{1}\{s=l\}).$$

The first term captures the trading fees: the bank receives a payoff of κ when the client chooses to buy the underlying financial asset. The second term reflects the bank's desire to retain high-skill analysts while firing low-skill analysts: the probability of retention is $pt_b + (1-p)t_\phi$, and the bank receives a payoff of 1 if it retains a high-skill analyst, a payoff of -1 if it retains a low-skill analyst, and a payoff of 0 if it fires the analyst. This captures, in a simple reduced-form way, the long-term value of increasing the organization's human capital. Since analysts generate revenue for the bank via their influence on the market, their advice must be sufficiently informative—and therefore their skill sufficiently high—to persuade the client and increase the volume of trade.[15] This is consistent with, for instance, Jackson's (2005) results showing that "analysts with better reputations generate significantly higher future trading volume" for their brokerages, and that these reputations are indeed consistently linked to forecast accuracy. Likewise, Lehmer, Lourie, and Shanthikumar (2022) provide evidence that "forecast and analyst quality are associated with a larger brokerage share of trading volume" and, moreover, that "brokerage houses are likely to consider trading volume generation to be a transferrable analyst skill"; therefore, investing in human capital within the organization by hiring, retaining, and promoting higher-skilled analysts can improve the bank's prospects for future persuasion.

4.2. OPTIMAL DIRECT MECHANISM

As mentioned above, the purpose of this application is twofold: to demonstrate that some conflict-of-interest regulations can be circumvented and rendered less effective in this market, and also to illustrate the operation of Theorem 1 in a specific context. The first of these goals will follow as an immediate consequence of Theorem 1, and so the main contribution of the next result is the explicit characterization of an optimal direct mechanism required for the latter goal.

Our payoff assumptions imply that there is an optimal direct mechanism $\chi^*(s,p) = X^*(s,p) \times A^*(s,p)$ for all $(s,p) \in \Theta$. Since there are only two actions, we can express the action recommendation rule $A^*(s,p) \in \Delta(\mathcal{A})$ as the probability $A^*(s,p) \in [0,1]$ with which the "buy" action $a = b$ is recommended. We can similarly express the contracting rule $X^*(s,p) \in \Delta(\mathcal{T})$ as the pair of probabilities $X^*(s,p) \in [0,1] \times [0,1]$ with which the analyst is retained when the events b and ϕ realize, respectively.

[15] It is easy to show that the probability of trade (given an optimal action recommendation rule) is increasing in the proportion of high-skill analysts; this follows from the assumption that a high-skill analyst is better-informed than a low-skill one.

(Given that both the bank and the analyst only care about the total retention probability, any correlation across the dimensions of $\Delta(\mathcal{T})$ can be ignored.) We abuse notation slightly in the remainder of this section and use $q^*(b)$ to denote the client's posterior belief that the investment is profitable when the buy action is recommended.[16]

THEOREM 2. *There is an optimal direct mechanism $\chi^*(s,p) = X^*(s,p) \times A^*(s,p)$ for all $(s,p) \in \Theta$, where*

$$X^*(s,p) := \begin{cases} (0,1) & \text{if } p < \frac{1}{2}, \\ (1,0) & \text{if } p \geq \frac{1}{2}, \end{cases} \quad \text{and } A^*(s,p) := \begin{cases} 0 & \text{if } p < \hat{p}, \\ \hat{\alpha} & \text{if } p = \hat{p}, \\ 1 & \text{if } p > \hat{p}, \end{cases}$$

and the threshold informational type $\hat{p} < \bar{q}$ and the buy recommendation probability $\hat{\alpha} \in (0,1]$ at this type are are such that $q^*(b) = \bar{q}$.

Moreover, there is a sequential equilibrium of the public communication game that implements χ^*.

The optimal contracting rule X^* is essentially a "prediction mechanism": the analyst is retained for sure if the event (b or ϕ) that they claim to be more likely actually occurs, and is fired for sure if it does not.[17] Note further that X^* does not depend on the analyst's reported skill s. To see why, note that the analyst's interim payoff—given their information p—does not depend on their skill s. Incentive compatibility (IC-A) thus implies that the probability of retention must be the same for both types (h,p) and (l,p) for all p. Moreover, there is no benefit to choosing different $X^*(h,p)$ and $X^*(l,p)$ that yield the same expected retention probability since the payoffs of both the bank and the analyst only depend on the latter.[18] Thus, as it is not possible to screen on the skill dimension of the analyst's private type, the bank must rely solely on the accuracy of the analyst's recommendations in order to indirectly evaluate and screen ability; note that this also implies that this mechanism remains optimal even if the analyst does not know their skill s and only observes the signal p about the underlying asset.

The derivation of the optimal action recommendation rule A^* in Theorem 2 follows from standard Bayesian persuasion arguments. Recall that the client engages in trade only when they are sufficiently confident; that is, only when their posterior $q^*(b)$ following a buy recommendation is greater than, or equal to $\overline{q} > \frac{1}{2}$. Suppose the bank were to recommend action b only when the analyst's information satisfied $p \geq \overline{q}$ (regardless of the analyst's skill s). Such a recommendation would yield a posterior belief $q > \overline{q}$, and it is clearly optimal for the client to obey this "naive" buy recommendation. But the bank can induce a greater volume of trade by lowering the threshold value of p above which buy recommendations are made. The cutoffs $\hat{p}$ and $\hat{\alpha}$ are chosen to exactly push the client's posterior belief down to $\overline{q}$, so that the client is *just* indifferent to buying when action b is recommended.[19]

[16] Since the client's payoff does not depend on the analyst's skill, we disregard her belief about that dimension of θ.

[17] A more general characterization of X^* that dispenses with our simplifying assumptions on the distribution $\pi \in \Delta(\Theta)$ is of independent interest, as the bank's screening problem is a novel mechanism design without transfers problem where local incentive compatibility constraints are not sufficient for optimality. An earlier version of this paper presented a characterization that relied on a combination of ironing and Lagrangian methods; Dasgupta (2023) uses the structure of the set of implementable mechanisms to also study this problem and derive interesting comparative statics.

[18] Other work has similarly noted the optimality of ignoring payoff-irrelevant (to the informed agent) private information; see, for instance, Jehiel and Moldovanu (2001), Che, Dessein, and Kartik (2013), or Dworczak, Kominers, and Akbarpour (2021).

[19] These cutoffs can be described directly in terms of the model's primitives; see equations (B.2) and (B.3). Also, note that the mixed recommendation when $p = \hat{p}$ is an artifact of the discrete type space.

4.3. PUBLIC COMMUNICATION IS OPTIMAL

Having provided some intuition for the form of the optimal direct mechanism, we now explain how it can be implemented as a sequential equilibrium of the public communication game. We begin with an explicit description of the bank's strategy.[20]

Define the message space $\mathcal{M}^* := \{B, \beta, \varphi\}$. First, suppose that the cutoff $\hat{p}$ from Theorem 2 satisfies $\hat{p} \geq \frac{1}{2}$. Consider the singleton set of contracts $\mathbb{X}^* := \{x^*_{B\beta}\}$, where $x^*_{B\beta}$ is given by

$$x^*_{B\beta}(m) := \begin{cases} (1,0) & \text{if } m \in \{B, \beta\}, \\ (0,1) & \text{if } m = \varphi. \end{cases}$$

Both messages B and β are interpreted by the bank as a prediction that the investment will be profitable while the message φ predicts unprofitability.

What is the type-(s,p) analyst's best response to the public communication mechanism $(\mathcal{M}^*, x^*_{B\beta})$? When $p > \frac{1}{2}$, they are indifferent between messages β and B but strictly prefer both to φ; when $p < \frac{1}{2}$, the analyst strictly prefers φ. Letting $\sigma(s,p|\mathcal{M},x) \in \Delta(\mathcal{M})$ denote the type-(s,p) analyst's strategy when faced with public communication mechanism $(\mathcal{M}, x)$, it is optimal for the analyst to respond to $(\mathcal{M}^*, x^*_{B\beta})$ with

$$\sigma^*(s,p|\mathcal{M}^*, x^*_{B\beta}) = \begin{cases} (0,0,1) & \text{if } p < \frac{1}{2}, \\ (1,0,0) & \text{if } \frac{1}{2} \leq p < \hat{p}, \\ (1-\hat{\alpha}, \hat{\alpha}, 0) & \text{if } p = \hat{p}, \\ (0,1,0) & \text{if } p > \hat{p}. \end{cases}$$

In the above strategy, the first, second, and third dimensions denote the probabilities with which the analyst reports B, β, and φ, respectively. Note that the above strategy implies that the ex ante likelihood of retention for each analyst type $(s,p) \in \Theta$ is precisely $X^*(s,p)$ from the optimal contracting rule.

Now observe that the client's posterior belief following a report β is, by construction, the same as her posterior $q^*(b)$ when the optimal action recommendation rule A^* suggests action b. Thus, the client optimally chooses action b after the message β. Conversely, the client's posterior following messages B or φ is lower than $\hat{p}$, which is in turn lower than $\bar{q}$. The client will hence optimally pick action ϕ after observing messages B or φ. In other words, the outcome is the same as that from the optimal action recommendation rule A^*.

Now suppose instead that the cutoff from Theorem 2 is such that $\hat{p} < \frac{1}{2}$. Indirect implementation of the optimal mechanism via a public communication is now more subtle: a single public message must separate analyst posteriors in $[\hat{p}, \frac{1}{2})$ from those in $[\frac{1}{2}, 1]$ (to implement the optimal screening rule X^*) while simultaneously pooling those posteriors (to implement the optimal recommendation rule A^*).

[20] Unsurprisingly, there are multiple sequential equilibria that implement the optimal direct mechanism. The proof of Theorem 1 constructs a "universal implementation" for all settings, including the present one. Here, we instead present a simpler implementation that is provably minimal (in number of messages required) and more natural for our application.

To achieve these dual objectives, consider the set of contracts $\mathbb{X}^* := \{x^*_B, x^*_\beta\}$, where

$$x^*_B(m) := \begin{cases} (1,0) & \text{if } m = B, \\ (0,1) & \text{if } m \neq B, \end{cases} \text{ and } x^*_\beta(m) := \begin{cases} (1,0) & \text{if } m = \beta, \\ (0,1) & \text{if } m \neq \beta. \end{cases}$$

Under x^*_B, the message B is interpreted by the bank as a prediction that the investment will be profitable while messages β and φ predict unprofitability. Under x^*_β, on the other hand, it is message β that is interpreted as a prediction of profitability while the remaining messages B and φ do not. Thus, the two contracts x^*_B and x^*_β permute the bank's interpretation of the message B and β.

When faced with public communication mechanism $(\mathcal{M}^*, x^*_B)$, it is clear that the type-(s,p) analyst strictly prefers to report message B whenever $p > \frac{1}{2}$. When instead $p < \frac{1}{2}$, she is indifferent between reporting β or φ but strictly prefers both to B. Therefore, a best response to $(\mathcal{M}^*, x^*_b)$ is the strategy

$$\sigma^*(s,p|\mathcal{M}^*, x^*_B) = \begin{cases} (0,0,1) & \text{if } p < \hat{p}, \\ (0,\hat{\alpha},1-\hat{\alpha}) & \text{if } p = \hat{p}, \\ (0,1,0) & \text{if } \hat{p} < p < \frac{1}{2}, \\ (1,0,0) & \text{if } p \geq \frac{1}{2}. \end{cases}$$

As the only difference between the public communication mechanisms $(\mathcal{M}^*, x^*_B)$ and $(\mathcal{M}^*, x^*_\beta)$ is the permutation of the messages B and β, a best response to $(\mathcal{M}^*, x^*_\beta)$ is the strategy

$$\sigma^*(s,p|\mathcal{M}^*, x^*_\beta) = \begin{cases} (0,0,1) & \text{if } p < \hat{p}, \\ (\hat{\alpha},0,1-\hat{\alpha}) & \text{if } p = \hat{p}, \\ (1,0,0) & \text{if } \hat{p} < p < \frac{1}{2}, \\ (0,1,0) & \text{if } p \geq \frac{1}{2}. \end{cases}$$

So suppose that the bank first commits to the set $\mathbb{X}^* := \{x^*_B, x^*_\beta\}$ of contracts, and then privately randomizes with equal probability over this set. Depending on the contract that realizes, the analyst best-responds according to the strategy σ^* above. By construction, the ex ante likelihood of retention for each analyst type $(s,p) \in \Theta$ is $X^*(s,p)$.

How does the client respond? It is easy to see that, from the client's perspective, messages B and β are indistinguishable: each is generated with equal likelihood by an identical distribution of analyst types. In other words, the client assigns an identical conditional distribution to the analyst's type when she observes either B or β. Moreover, by construction, this conditional distribution is identical to the posterior $q^*(b)$ corresponding to the client's belief when she receives the buy recommendation in the optimal direct mechanism. Similarly, the conditional distributions over the analyst's type when the client observes the message ϕ in the public communication game is the same as when she observes the recommendation $A^* = \phi$ in the optimal direct mechanism. Thus, the ex ante distribution over actions for each $(s,p) \in \Theta$ is the same in both the public communication game and the optimal direct mechanism.

Finally, note that the optimal direct mechanism is implemented regardless of the realization of the bank's randomization over the set $\mathbb{X}^*$ of contracts. Therefore, an equal likelihood mix over both public communication mechanisms is a best response for the bank—who also cannot improve its payoff by committing to a different set $\mathbb{X} \neq \mathbb{X}^*$ of contracts. Consequently, the strategies described are an equilibrium and implement the optimal direct mechanism as required.[21]

Moreover, the construction above highlights the importance of intentionally vague language for effective indirect persuasion: the client is unsure about the precise interpretation of message β relative to message B. This uncertainty is not unnatural as there is, in practice, a substantial lack of clarity about the interpretation of analyst ratings. For instance, what exactly is the threshold demarcating the boundary between a "strong buy" and a "weak buy" recommendation? And this vagueness is exacerbated when, as is often the case, analysts resort to even more opaque language: is an "outperform" rating better or worse than an "overweight" rating, and how do they both compare to a simple "buy" rating? Critically, however, this vagueness is one-sided, as the bank and analyst have a mutually shared understanding of how the public language is to be privately decoded.

While the specific construction yielding Theorem 2 is unique to the stylized example we describe here, the intuition holds far more generally. For example, it is straightforward (though notationally more cumbersome) to extend this argument to richer (and potentially asymmetric) environments with a "sell" state in which the financial asset depreciates in value.[22] Likewise, Theorem 1 applies in settings where the client can determine the volume of trade, the analyst has multiple abilities, and the contractible decisions include transfers.

Finally, it is worth noting that there is some evidence suggesting that other regulatory interventions may be effective in improving the informativeness of analyst ratings and recommendations. Chen and Chen (2009) and Chen, Novoselov, and Wang (2018) present evidence suggesting that NASD Rule 2711, the precursor to FINRA Rule 2241 that mandated firewalls between brokerages investment and research arms while also regulating analyst compensation, reduced the influence of conflicts of interest (including trading commissions) and increased the informativeness of analyst research. Meanwhile, Fang, Hope, Huang, and Moldovan (2020), Guo and Mota (2021), and Lang, Pinto, and Sul (2024) argue that MiFID II (the EU regulation unbundling sell-side research from transactions) has yielded improvements in research quality and analyst recommendation accuracy.

4.4. THE WELFARE CONSEQUENCES OF TRANSPARENCY

As should be clear from the above discussion, implementation of the optimal direct mechanism requires that the realized contract between the bank and the analyst is unobserved by the client. It is therefore natural to ask whether the client benefits if the bank's contractual terms are instead mandated to be public? In this case, the meaning of the message chosen by the analyst is common to all three players.

Recall that the action recommendation rule A^* in the optimal direct mechanism ensures that the client is indifferent between trading and not when she receives a buy recommendation, and so the client's ex ante payoff is zero. Making the contractual terms public therefore cannot hurt the client, as they can always guarantee themselves a payoff of zero by not trading. The next result shows that transparency and public contracting can instead make the client strictly better off.

[21] Of course, to fully describe the equilibrium requires a specification of off-path behavior should the bank deviate and offer a different set of contracts; these details are discussed in the proof of Theorem 1. We reiterate, however, that our construction does not rely on implausible or pathological off-path beliefs.

[22] This analysis (in an earlier version of this paper and available on request) yields the traditional five-point analyst rating scale.

THEOREM 3. *Suppose the optimal action recommendation A^* from Theorem 2 satisfies $\hat{p} < \frac{1}{2}$. Then, there is a weight on the bank's trading fees and commissions $\overline{\kappa} > 0$ such that, for any $\kappa < \overline{\kappa}$, the client earns a strictly positive payoff in every bank-optimal sequential equilibrium with public contracting.*

The intuition for this result is straightforward. As we illustrated in the previous subsection, when $\hat{p} < \frac{1}{2}$, the bank cannot obtain the full-commitment profits without mixing over contractual terms with the analyst.[23] Moreover, mixing is effective only when those contractual terms are unobserved by the client. Therefore, when contracting is public, the bank has to compromise on either the payoffs from retention or the payoffs from trading commissions—and when the weight κ on the latter is sufficiently low relative to the former, trading commissions are sacrificed.[24] The proof of Theorem 3 shows that, in this case, every bank-optimal equilibrium with public contracting features a cutoff $\tilde{p} > \hat{p}$ such that the client is only induced to buy the asset when the analyst's information satisfies $p \geq \tilde{p}$. Since $\tilde{p} > \hat{p}$, the client receives a strictly positive payoff from buying the asset.

5. CONCLUDING REMARKS

Bayesian persuasion is a powerful theoretical framework that relies on the ability of a sender to commit to arbitrary information structures, which may be an unrealistic assumption in practice. In this paper, we instead model the sender as *two* players—an uninformed principal that contracts with a privately informed agent—and show that optimal persuasion can be achieved by partial commitment only to standard contractual terms. Commitment to arbitrary information structures is thus not required to achieve optimal persuasion. This yields an organizational microfoundation for commitment in strategic communication.

Our main result, Theorem 1, applies to an environment with adverse selection. For certain applications, it might be more appropriate to instead consider a model where the agent must exert a costly private effort $e \in \mathcal{E}$ in order to learn her private type θ, which is drawn from a distribution $\pi(e) \in \Delta(\Theta)$. A direct mechanism here would need to provide incentives to ensure the agent obediently exerts the principal's desired effort and then reports her realized type truthfully. (Action recommendations to the receiver would be akin to those in the "pure" adverse selection case we study here.) It is straightforward to show (under similar conditions on payoffs) that the principal's optimal direct mechanism can be indirectly implemented as a sequential equilibrium of a public communication game (defined in analogous fashion to that of Section 2). The intuition remains the same: the principal can mix over public communication mechanisms, allowing a single public message from the agent to convey different meanings to the principal and the receiver. Of course, in such settings, the moral hazard problem introduces an additional friction, as a principal who "owns" the information acquisition process will typically choose a different level of effort than that required of the agent in the optimal direct mechanism.

[23] Recall that when $\hat{p} \geq \frac{1}{2}$, mixing is not necessary and the bank commits (deterministically) to a single contract $\mathbb{X}^* = \{x^*_{B\beta}\}$.

[24] This intuition extends to the case of private contracting with a positive probability that the contract is publicly revealed.

A. PROOF OF THEOREM 1

The proof is constructive. Fix an optimal direct mechanism χ^* with marginals X^* and A^*, and note that, by definition, X^* and A^* satisfy the incentive compatibility constraints (IC-A) and (IC-R).

Let $\text{supp}(A^*(\theta))$ denote the support of $A^*(\theta) \in \Delta(\mathcal{A})$, and let $\bar{A} := \cup_{\theta\in\Theta} \text{supp}(A(\theta))$ be the set of all action recommendations that might realize. Since Θ is finite, we can enumerate its elements by writing $\Theta := \{\theta_0, \ldots, \theta_{N-1}\}$, where $N := |\Theta|$. We then define the message space $\mathcal{M} := \bar{A} \times \{0, \ldots, N-1\}$, and for each $k = 0, \ldots, N-1$, define the contract $x_k : \mathcal{M} \to \Delta(\mathcal{T})$ by

$$x_k(a, i) := X^*(\theta_{(i+k) \bmod N}).$$

The set of messages $\mathcal{M}$ consists of N "publicly synonymous" copies of each action: for any $(a, i) \in \mathcal{M}$, the receiver interprets the message as a recommendation to take action a while the principal interprets the message as report of type θ_j, where $j := (i + k) \bmod N$ is a "cyclic" permutation of Θ.

Now consider the following "on-path" strategies for each player:

- the principal publicly announces the message space $\mathcal{M}$ and set of contracts $\mathbb{X} = \{x_k\}_{k=0,\ldots,N-1}$, and then randomizes uniformly over this set with $\rho(x_k) = \frac{1}{N}$ for all $k = 0, \ldots, N-1$;
- after privately observing the realized contract $x_{\hat{k}}$, the agent of type $\theta_i \in \Theta$ publicly announces message $(a, j) \in \mathcal{M}$ with probability $A^*(\theta_i)[a]$ if $j = (i + \hat{k}) \bmod N$, and with probability 0 otherwise; and
- after observing the publicly realized message $(\hat{a}, j)$, for any $j = 0, \ldots, N-1$, the receiver takes action $\hat{a}$.

If the principal announces any other finite message space $\mathcal{M}'$ or finite set of contracts $\mathbb{X}' \subset (\mathcal{M}')^{\Delta(\mathcal{T})}$, arbitrarily choose any continuation strategies that constitute a sequential equilibrium for this "off-path" subgame, where existence of such a selection is guaranteed because all players pick from finite action sets; see Kreps and Wilson (1982).

We will now establish that these strategies constitute a sequential equilibrium of the public communication game and that the equilibrium outcome coincides with that of the optimal direct mechanism.

We begin by examining the responses to the principal's on-path play starting with the receiver's action choice. After observing public message $(\hat{a}, j)$, the receiver's posterior belief that the agent's type is $\theta_i \in \Theta$ is, by Bayes' rule, given by

$$\frac{\pi(\theta_i)\rho(x_{(j-i) \bmod N})A^*(\theta_i)[\hat{a}]}{\sum_k \pi(\theta_k)\rho(x_{(j-k) \bmod N})A^*(\theta_k)[\hat{a}]} = \frac{\pi(\theta_i)A^*(\theta_i)[\hat{a}]}{\sum_k \pi(\theta_k)A^*(\theta_k)[\hat{a}]},$$

where the right-hand side is precisely the receiver's posterior belief after observing action recommendation $\hat{a}$ from the direct mechanism. Since that latter satisfies (IC-R), it must therefore be the case that taking action $\hat{a}$ remains a best response (on-path) in the public communication game.

Now consider the agent's message announcement decision, fixing an arbitrary type $\theta_i \in \Theta$ and a realization $x_{\hat{k}}$ of the principal's randomization over $\mathbb{X} = \{x_k\}_{k=0,\ldots,N-1}$. Since the image of $x_{\hat{k}}$ is the same as that of X^* and the latter satisfies (IC-A), type θ_i's payoff is maximized by sending any message

$(a,j) \in x_{\hat{k}}^{-1}(X^*(\theta_i))$. Given the definition of $x_{\hat{k}}$ (and the fact that it does not condition on the recommended action), this implies that *all* messages in $\{(a,(i+\hat{k}) \bmod N)\}_{a\in\bar{A}}$ are best responses for the type-θ_i agent; in particular, the mixture over this set that corresponds to the distribution $A^*(\theta_i) \in \Delta(\mathcal{A})$ is optimal for type θ_i.

Finally, note that the principal is indifferent between offering the agent any contract from the set $\mathbb{X} = \{x_k\}_{k=0,\ldots,N-1}$. In particular, the agent's contractual outcome does not depend on the realization; because the receiver does not observe the realization, their action is also independent of the realization. Thus, uniform randomization is optimal for the principal and on-path play results in the identical outcomes as the optimal direct mechanism χ^*.

We complete the proof by observing that the principal has no incentive to deviate by announcing a different message space $\mathcal{M}' \neq \mathcal{M}$ or set of contracts $\mathbb{X}' \neq \mathbb{X}$: any sequential equilibrium of any such off-path subgame must deliver (by definition) a lower payoff than the optimal direct mechanism χ^*. ■

B. PROOF OF THEOREM 2

Recall that a direct mechanism is a mapping $\chi : \Theta \to \Delta(\mathcal{T} \times \mathcal{A})$ with marginals $X : \Theta \to \Delta(\mathcal{T})$ and $A : \Theta \to \Delta(\mathcal{A})$. We will write $X(s,p) = (X_b(s,p), X_\phi(s,p))$, where $X_\omega(s,p)$ denotes the probability with which the analyst is retained when she reports $\theta = (s,p)$ and event $\omega \in \{b, \phi\}$ realizes. Likewise, we interpret $A(s,p)$ to be the probability with which the "buy" action $a = b$ is recommended when the analyst reports $\theta = (s,p)$.

The analyst's truthful reporting constraints (IC-A) can be written as

$$(s,p) \in \underset{(s',p')}{\operatorname{argmax}} \{pX_b(s',p') + (1-p)X_\phi(s',p')\} \text{ for all } (s,p) \in \Theta. \tag{AT}$$

For the client, on the other hand, the obedience constraints (IC-R) can be rewritten as

$$q(b) := \frac{\sum_{(s,p)} pA(s,p)\pi(s,p)}{\sum_{(s,p)} A(s,p)\pi(s,p)} \geq \bar{q} \text{ and } q(\phi) := \frac{\sum_{(s,p)} p(1-A(s,p))\pi(s,p)}{\sum_{(s,p)} (1-A(s,p))\pi(s,p)} \leq \bar{q}. \tag{CO}$$

These two conditions require that the client's posterior upon receiving a "buy" recommendation is sufficiently high that the recommendation is indeed optimally followed, and that her posterior upon receiving a "hold" recommendation is sufficiently low that she does not buy.

Thus, the bank's optimal direct mechanism χ^* solves the following "full commitment" problem

$$\max_{\chi} \left\{ \begin{aligned} &\sum_{(s,p)} (pX_b(s,p) + (1-p)X_\phi(s,p)) \left(\mathbb{1}\{s=h\} - \mathbb{1}\{s=l\}\right) \pi(s,p) \\ &\qquad\qquad\qquad\qquad\qquad\qquad + \kappa \sum_{(s,p)} A(s,p)\pi(s,p) \end{aligned} \right\} \tag{FC}$$
$$\text{s.t. (AT), (CO), and } X(s,p) \in [0,1]^2, A(s,p) \in [0,1] \text{ for all } (s,p) \in \Theta.$$

Problem (FC) is clearly separable—the analyst truthtelling constraint (AT) only affects the part of the objective function containing X, while the client obedience constraint (CO) only affects the part of the objective function containing A. In particular, since $\pi(s,p) = \frac{1}{2} f_s(p)$ for all $(s,p) \in \Theta$, we can rewrite problem (FC) as a stand-alone screening problem

$$\max_X \left\{ \sum_p \left((pX_b(h,p) + (1-p)X_\phi(h,p))\, f_h(p) - (pX_b(l,p) + (1-p)X_\phi(l,p))\, f_l(p) \right) \right\} \quad (S)$$
$$\text{s.t. (AT) and } X(s,p) \in [0,1]^2 \text{ for all } (s,p) \in \Theta,$$

and a stand-alone persuasion problem

$$\max_A \left\{ \sum_p \left(A(h,p) f_h(p) + A(l,p) f_l(p) \right) \right\} \quad (P)$$
$$\text{s.t. (CO) and } A(s,p) \in [0,1] \text{ for all } (s,p) \in \Theta.$$

Below, we characterize the optimal retention rule X^* (Lemma B.1) and the optimal persuasion rule A^* (Lemma B.2) independently. These two results jointly yield the characterization in Theorem 2; meanwhile, the discussion in the main text proves the indirect implementation result. ■

LEMMA B.1. *The optimal retention rule X^* that solves problem (S) is*

$$X^*(s,p) := \begin{cases} (0,1) & \text{if } p < \frac{1}{2}, \\ (1,0) & \text{if } p \geq \frac{1}{2}. \end{cases} \quad (S^*)$$

PROOF. Consider any direct retention rule $X : \Theta \to \Delta(\mathcal{T})$. We write the analyst's payoff from reporting (s',p') when her true type is (s,p) by

$$U(s',p'|s,p) := pX_b(s',p') + (1-p)X_\phi(s',p').$$

Note first that (AT) implies that

$$U(h,p|h,p) \geq U(l,p|h,p) \text{ and } U(l,p|l,p) \geq U(h,p|l,p) \text{ for all } p.$$

But $U(s',p'|s,p)$ does not depend on s, so $U(h,p|h,p) = U(h,p|l,p)$ and $U(l,p|l,p) = U(l,p|h,p)$ for all p. Together, these two sets of expressions imply that, arbitrarily choosing $s \in \{h,l\}$, we can define a "truthful" expected utility

$$U(p) := U(s,p|s,p) \text{ for all } p.$$

Furthermore, (AT) also implies that, for all s, p, and p',

$$\begin{aligned} U(p) \geq U(s,p'|s,p) &= p\left[X_b(s,p') - X_\phi(s,p')\right] + X_\phi(s,p') \\ &= U(p') + (p-p')\left[X_b(s,p') - X_\phi(s,p')\right], \text{ and} \\ U(p') \geq U(s,p|s,p') &= p'\left[X_b(s,p) - X_\phi(s,p)\right] + X_\phi(s,p) \\ &= U(p) + (p'-p)\left[X_b(s,p) - X_\phi(s,p)\right]. \end{aligned}$$

Adding these two inequalities, we obtain

$$(p-p')\left[X_b(s,p)-X_\phi(s,p)\right] \geq (p-p')\left[X_b(s,p')-X_\phi(s,p')\right] \text{ for all } s, p, \text{ and } p',$$

or, equivalently,

$$\Delta(s,p) := X_b(s,p)-X_\phi(s,p) \text{ is weakly increasing in } p, \text{ for all } s. \tag{MON}$$

Moreover, the two inequalities above can also be combined to yield a discrete "envelope condition":

$$(p-p')\Delta(s,p) \geq U(p)-U(p') \geq (p-p')\Delta(s,p') \text{ for all } s, p, \text{and } p'. \tag{ENV}$$

With these observations in hand, let us consider a relaxed version of the screening problem (S) where, instead of imposing (AT) directly, we instead impose the monotonicity and "envelope" conditions above:

$$\begin{aligned}&\max_X \left\{\sum_p U(p)(f_h(p)-f_l(p))\right\}\\ &\text{s.t. (MON), (ENV), and } X(s,p)\in[0,1]^2 \text{ for all } (s,p)\in\Theta.\end{aligned}$$

Note, however, that summation by parts yields

$$\begin{aligned}\sum_p U(p)f_s(p) = \sum_{k=0}^{n} U\left(\tfrac{k}{n}\right) f_s\left(\tfrac{k}{n}\right) &= \sum_{k=0}^{n} U\left(\tfrac{k}{n}\right)\left[F_s\left(\tfrac{k}{n}\right)-F_s\left(\tfrac{k-1}{n}\right)\right]\\ &= U(1)-\sum_{k=0}^{n-1}\left[U\left(\tfrac{k+1}{n}\right)-U\left(\tfrac{k}{n}\right)\right]F_s\left(\tfrac{k}{n}\right).\end{aligned}$$

Thus, the relaxed problem above can be rewritten as

$$\begin{aligned}&\max_X \left\{\sum_{k=0}^{n-1}\left[U\left(\tfrac{k+1}{n}\right)-U\left(\tfrac{k}{n}\right)\right]\left[F_l\left(\tfrac{k}{n}\right)-F_h\left(\tfrac{k}{n}\right)\right]\right\}\\ &\text{s.t. (MON), (ENV), and } X(s,p)\in[0,1]^2 \text{ for all } (s,p)\in\Theta.\end{aligned} \tag{S'}$$

Note that the rotation order on F_h and F_l, along with symmetry, implies that $F_l(p)\leq F_h(p)$ whenever $p<\frac{1}{2}$, and likewise that $F_l(p)\geq F_h(p)$ whenever $p\geq\frac{1}{2}$. With this in mind, it is helpful to rewrite the objective function in problem (S') as

$$\sum_{k<\frac{n}{2}}\left[U\left(\tfrac{k+1}{n}\right)-U\left(\tfrac{k}{n}\right)\right]\left[F_l\left(\tfrac{k}{n}\right)-F_h\left(\tfrac{k}{n}\right)\right]+\sum_{\frac{n}{2}\leq k<n}\left[U\left(\tfrac{k+1}{n}\right)-U\left(\tfrac{k}{n}\right)\right]\left[F_l\left(\tfrac{k}{n}\right)-F_h\left(\tfrac{k}{n}\right)\right].$$

Considering this objective pointwise (that is, for each $k=0,\ldots,n-1$) suggests that it is optimal to minimize the difference $U(\frac{k+1}{n})-U(\frac{k}{n})$ for $k<\frac{n}{2}$ and to maximize it for $k\geq\frac{n}{2}$.

To this end, we can apply the bounds from (ENV), which we rewrite here for adjacent analyst posteriors:

$$\frac{1}{n}\Delta(s, \tfrac{k+1}{n}) \geq U(\tfrac{k+1}{n}) - U(\tfrac{k}{n}) \geq \frac{1}{n}\Delta(s, \tfrac{k}{n}) \text{ for all } s \text{ and } k = 0, 1, \ldots, n-1.$$

These bounds can be tightened to

$$\frac{1}{n}\min_s\{\Delta(s, \tfrac{k+1}{n})\} \geq U(\tfrac{k+1}{n}) - U(\tfrac{k}{n}) \geq \frac{1}{n}\max_s\{\Delta(s, \tfrac{k}{n})\} \text{ for all } k = 0, 1, \ldots, n-1.$$

Thus, we can rewrite the objective function as

$$\frac{1}{n}\sum_{k<\frac{n}{2}} \max_s\{\Delta(s, \tfrac{k}{n})\}\left[F_l\left(\tfrac{k}{n}\right) - F_h\left(\tfrac{k}{n}\right)\right] + \frac{1}{n}\sum_{\frac{n}{2}\leq k<n} \min_s\{\Delta(s, \tfrac{k+1}{n})\}\left[F_l\left(\tfrac{k}{n}\right) - F_h\left(\tfrac{k}{n}\right)\right].$$

The feasibility constraints ($X(s,p) \in [0,1]^2$ for all $(s,p) \in \Theta$) imply that we must have $\Delta(s,p) \in [-1,1]$ for all $(s,p) \in \Theta$; this suggests the optimality of a bang-bang solution:

$$\Delta^*(s, \tfrac{k}{n}) = -1 \text{ for all } k < \frac{n}{2} \text{ and } \Delta^*(s, \tfrac{k+1}{n}) = 1 \text{ for all } k \geq \frac{n}{2}.$$

This is, of course, achievable by choosing X^* as in (S^*). Note that this "prediction mechanism" satisfies (AT), implying that it is also a solution to the original screening problem (S). ■

LEMMA B.2. *There exist $\hat{p} < \bar{q}$ and $\hat{\alpha} \in (0,1]$ such that the optimal action recommendation rule A^* that solves problem (P) is*

$$A^*(s,p) := \begin{cases} 0 & \text{if } p < \hat{p}, \\ \hat{\alpha} & \text{if } p = \hat{p}, \\ 1 & \text{if } p > \hat{p}. \end{cases} \tag{P^*}$$

PROOF. Because $\pi(s,p) = \frac{1}{2}f_s(p)$ for all $(s,p) \in \Theta$, we can rewrite the client's obedience constraints (CO) as

$$\frac{\sum_{(s,p)} pA(s,p)f_s(p)}{\sum_{(s,p)} A(s,p)f_s(p)} \geq \bar{q} \text{ and } \frac{\sum_{(s,p)} p(1-A(s,p))f_s(p)}{\sum_{(s,p)} (1-A(s,p))f_s(p)} \leq \bar{q}.$$

Multiplying these posteriors through by their denominators, we can can rewrite problem (P) as

$$\begin{aligned} &\max_A \left\{\sum_p (A(h,p)f_h(p) + A(l,p)f_l(p))\right\} \\ &\text{s.t. } \sum_s\sum_p (p-\bar{q})A(s,p)f_s(p) \geq 0, \\ &\qquad \sum_s\sum_p (\bar{q}-p)(1-A(s,p))f_s(p) \geq 0, \\ &\qquad f_s(p)A(s,p) \geq 0 \text{ for all } (s,p) \in \Theta, \text{ and} \\ &\qquad f_s(p)(1-A(s,p)) \geq 0 \text{ for all } (s,p) \in \Theta. \end{aligned} \tag{P'}$$

Note that the final two sets of constraints are the feasibility constraints (which require $A(s,p) \in [0,1]$ for all $(s,p) \in \Theta$) normalized by the (strictly positive) probability mass functions.

Problem (P') is just a linear program and can therefore be solved with Lagrangian methods. To this end, we let λ_b, λ_ϕ, $\eta_b(s,p)$, and $\eta_\phi(s,p)$ denote the (nonnegative) multipliers on each of the inequality constraints. The solution to (P') is therefore characterized by the following first-order conditions (along with the usual complementary slackness conditions):

$$f_s(p)\left[1 + \lambda_b(p - \bar{q}) - \lambda_\phi(\bar{q} - p) + \eta_b(s,p) - \eta_\phi(s,p)\right] = 0 \text{ for all } (s,p) \in \Theta. \tag{B.1}$$

Now define $\hat{p}$ by

$$\hat{p} := \min\left\{ p \in \{0, \tfrac{1}{n}, \tfrac{2}{n}, \ldots, 1\} \;\middle|\; \frac{\sum_{q>p} q(f_h(q) + f_l(q))}{\sum_{q>p} (f_h(q) + f_l(q))} > \bar{q} \right\}. \tag{B.2}$$

Thus, $\hat{p}$ is the lowest analyst posterior such that pooling all strictly higher posteriors, regardless of the analyst's skill, leads the client to strictly prefer to buy. (Note that $\hat{p}$ is well-defined, is strictly positive due to the symmetry of the distributions, and is necessarily smaller than $\bar{q}$.) Furthermore, we can define $\hat{a} \in (0,1]$ as the solution to

$$\frac{\alpha\hat{p}(f_h(\hat{p}) + f_l(\hat{p})) + \sum_{p>\hat{p}} p(f_h(p) + f_l(p))}{\alpha(f_h(\hat{p}) + f_l(\hat{p})) + \sum_{p>\hat{p}} (f_h(p) + f_l(p))} = \bar{q}. \tag{B.3}$$

It is easy to show that $\hat{\alpha}$ is well-defined: the expression on the left-hand side is strictly decreasing in α, and—by the definition of $\hat{p}$—a solution in $(0,1]$ exists. Moreover, note that $\hat{p} < \bar{q}$ implies that

$$\frac{\sum_{p<\hat{p}} p(f_h(p) + f_l(p)) + (1-\hat{\alpha})\hat{p}(f_h(\hat{p}) + f_l(\hat{p}))}{\sum_{p<\hat{p}} (f_h(p) + f_l(p)) + (1-\hat{\alpha})(f_h(\hat{p}) + f_l(\hat{p}))} < \bar{q}.$$

Consider, therefore, the following dual variables:

$$\lambda_b^* := \frac{1}{\bar{q} - \hat{p}},\, \lambda_\phi^* := 0,\, \eta_b^*(s,p) := \begin{cases} \frac{\hat{p}-p}{\bar{q}-\hat{p}} & \text{if } p < \hat{p}, \\ 0 & \text{if } p \geq \hat{p}, \end{cases} \text{ and } \eta_\phi^*(s,p) := \begin{cases} 0 & \text{if } p \leq \hat{p}, \\ \frac{p-\hat{p}}{\bar{q}-\hat{p}} & \text{if } p > \hat{p}. \end{cases}$$

Straightforward algebra verifies that A^* as defined in (P^*), along with these multipliers, satisfy the first-order conditions in (B.1), as well as the complementary slackness conditions, for all $(s,p) \in \Theta$. Thus, A^* is indeed a bank-optimal action recommendation rule. ■

C. PROOF OF THEOREM 3

We proceed in several steps. First, let $(\widetilde{\mathcal{M}}, \tilde{x})$ denote the public communication mechanism defined by

$$\widetilde{\mathcal{M}} := \{b, \phi\} \text{ and } \tilde{x}(m) := \begin{cases} (1,0) & \text{if } m = b, \\ (0,1) & \text{if } m = \phi. \end{cases}$$

Note that this public prediction mechanism implements the optimal retention rule X^* from Theorem 2; in addition, upon observing message $m = b$ (which is sent by the analyst whenever $p \geq \frac{1}{2} > \hat{p}$), the client believes that the "buy event" occurs with probability $\mathbb{E}[p \mid p \geq \frac{1}{2}] > \bar{q}$, and so has a strictly positive payoff from taking the recommended action.

Next, fix any bank-optimal equilibrium with public contracting in which the bank chooses a finite message space $\widehat{\mathcal{M}}$ along with a retention rule $\hat{x} : \widehat{\mathcal{M}} \to \Delta(\mathcal{T})$; the analyst chooses a reporting strategy $\hat{\sigma}(s,p) \in \Delta(\widehat{\mathcal{M}})$; and the client chooses an action strategy $\hat{\alpha} : \widehat{\mathcal{M}} \to \mathcal{A}$. It is without loss to assume that all messages $m \in \widehat{\mathcal{M}}$ are on-path, and we will denote by $q(m) \in [0,1]$ the client's posterior belief that the "buy event" will occur after observing message $m \in \widehat{\mathcal{M}}$.

STEP 1: ALL MESSAGES IN $\mathcal{M}^b$ HAVE 0 CLIENT PAYOFF. Notice that if there are any messages $m \in \widehat{\mathcal{M}}$ with $q(m) > \bar{q}$, the client receives a strictly positive payoff from $\hat{\alpha}(m) = b$ (and therefore also a strictly positive payoff in equilibrium). So assume instead that $q(m) \leq \bar{q}$ for all $m \in \widehat{\mathcal{M}}$, and let

$$\mathcal{M}^b := \left\{ m \in \widehat{\mathcal{M}} \,\middle|\, q(m) = \bar{q} \right\}$$

be the set of messages that make the client indifferent between actions. Note that in a bank-optimal equilibrium, the client breaks ties in favor of buying, and so $\hat{\alpha}(m) = b$ for all $m \in \mathcal{M}^b$. Note further that $|\mathcal{M}^b| > 0$, as otherwise this purported bank-optimal public mechanism yields a lower payoff than $(\widetilde{\mathcal{M}}, \tilde{x})$ (which achieves optimal screening along with some persuasion).

STEP 2: ANALYST TYPES WITH $p = 1$ MUST SEND A MESSAGE IN $\mathcal{M}^b$. Suppose that some type $(s', 1)$ never reports a message in $\mathcal{M}^b$; that is, suppose that $\hat{\sigma}(s',1)[\mathcal{M}^b] = 0$ for some $s' \in \{h, l\}$. We can then define an alternative public communication mechanism $(\mathcal{M}', x')$ by

$$\mathcal{M}' := \widehat{\mathcal{M}} \cup \{m'_{s',1}\} \text{ and } x'(m) := \begin{cases} \hat{x}(m) & \text{if } m \in \widehat{\mathcal{M}}, \\ \hat{x}(\hat{m}_{s',1}) & \text{if } m = m'_{s',1}, \end{cases}$$

for any $\hat{m}_{s',1} \in \text{supp}(\hat{\sigma}(s',1))$. Thus, type $(s',1)$ is assigned a new message $m'_{s',1}$ that the bank treats as equivalent to some message $m \in \mathcal{M}^b$ in the support of their reporting strategy in the original $\hat{x}$.

Since x' essentially replicates $\hat{x}$, it is a best response for type $(s',1)$ to choose $\sigma'(s',1)[m'_{s',1}] = 1$, and for all other types $(s,p) \in \Theta$ to choose $\sigma'(s,p) = \hat{\sigma}(s,p)$. This yields posterior belief

$$q'(m) \begin{cases} = 1 & \text{if } m = m'_{s',1}, \\ = \hat{q}(m) & \text{if } m \in \mathcal{M}^b, \\ \leq \hat{q}(m) & \text{otherwise.} \end{cases}$$

This yields a strictly greater probability of the client engaging in trade while achieving the same separation of type-h and type-l analysts, contradicting the optimality of $(\widehat{\mathcal{M}}, \hat{x})$.

Thus, we must have $\hat{\sigma}(s,1)[\mathcal{M}^b] > 0$ for both $s \in \{h, l\}$. Notice that this also implies that

$$\underline{p} := \min \left\{ p \,\middle|\, \hat{\sigma}(h,p)[\mathcal{M}^b] + \hat{\sigma}(l,p)[\mathcal{M}^b] > 0 \right\} < \bar{q}.$$

If this were not the case, then the posterior $q(m)$ after any message $m \in \text{supp}(\hat{\sigma}(h,1)) \cap \mathcal{M}^b$ would be strictly greater than $\bar{q}$, contradicting the assumption that the client is indifferent and achieves zero payoff.

STEP 3: THE BANK POOLS ALL MESSAGES IN $\mathcal{M}^b$. Now fix any $m \in \mathcal{M}^b$ with $\hat{\sigma}(s_1,p_1)[m] > 0$ and $\hat{\sigma}(s_2,p_2)[m] > 0$ for some $(s_1,p_1),(s_2,p_2) \in \Theta$ with $p_1 < \bar{q} < p_2$. (These types can be chosen to satisfy the latter set of inequalities due to the client's indifference upon observing message m.) We argue that, for any other $m' \in \widehat{\mathcal{M}}$ with $\hat{\sigma}(s,p)[m'] > 0$ for any $(s,p) \in \Theta$ with $p \in (p_1,p_2)$, we must have $\hat{x}(m) = \hat{x}(m')$. The optimality of reporting message m by types (s_1,p_1) and (s_2,p_2) and reporting m' by type (s,p) imply that

$$p_1\hat{x}_b(m) + (1-p_1)\hat{x}_\phi(m) \geq p_1\hat{x}_b(m') + (1-p_1)\hat{x}_\phi(m'), \tag{C.1}$$
$$p_2\hat{x}_b(m) + (1-p_2)\hat{x}_\phi(m) \geq p_2\hat{x}_b(m') + (1-p_2)\hat{x}_\phi(m'), \text{ and} \tag{C.2}$$
$$p\hat{x}_b(m') + (1-p)\hat{x}_\phi(m') \geq p\hat{x}_b(m) + (1-p)\hat{x}_\phi(m), \tag{C.3}$$

where write $\hat{x}(\cdot) = (\hat{x}_b(\cdot),\hat{x}_\phi(\cdot)) \in \Delta(\mathcal{T})$. Combining (C.1) and (C.3) and then dividing through by $(p_1 - p)$ yields

$$\hat{x}_b(m) - \hat{x}_\phi(m) \leq \hat{x}_b(m') - \hat{x}_\phi(m').$$

Likewise, combining (C.2) and (C.3) and then dividing through by $(p_2 - p)$ yields

$$\hat{x}_b(m) - \hat{x}_\phi(m) \geq \hat{x}_b(m') - \hat{x}_\phi(m').$$

Thus, we must have $\hat{x}_b(m) - \hat{x}_\phi(m) = \hat{x}_b(m') - \hat{x}_\phi(m')$. Of course, this then implies that we may rewrite (C.2) and (C.3) as

$$\hat{x}_\phi(m) \geq \hat{x}_\phi(m') \text{ and } \hat{x}_\phi(m') \geq \hat{x}_\phi(m),$$

jointly implying that $\hat{x}_\phi(m) = \hat{x}_\phi(m')$. Since we can repeat this argument with $p_1 = \underline{p}$, and again with $p_2 = 1$, this implies that the contract $\hat{x}$ treats all messages resulting in trade identically; that is,

$$\hat{x}(m) = \hat{x}(m') \text{ for all } m,m' \in \mathcal{M}^b.$$

Indeed, this pooling extends to all $m \in \widehat{\mathcal{M}}$ with $\hat{\sigma}(s,p)[m] > 0$ for any $p > \underline{p}$.

STEP 4: $(\widehat{\mathcal{M}},\hat{x})$ IMPLEMENTS THE OPTIMAL PERSUASION RULE A^* FROM THEOREM 2. With this in hand, we can proceed to argue that this optimal mechanism also pools the recommendations for all analyst types $(s,p) \in \Theta$ with $p > \underline{p}$; that is, $\hat{\sigma}(s,p)[\mathcal{M}^b] = 1$ for all $(s,p) \in \Theta$ with $p > \underline{p}$. Suppose that this were not the case, and consider $\bar{p} \geq \underline{p}$ and $\bar{\alpha} \in (0,1]$ such that

$$\sum_s \frac{1}{2}\left(f_s(\bar{p})\bar{\alpha} + \sum_{p>\bar{p}} f_s(p)\right) = \sum_s \frac{1}{2}\sum_p f_s(p)\hat{\sigma}(s,p)[\mathcal{M}^b] =: \widehat{T},$$

where $\widehat{T}$ is the volume of trade induced by $(\widehat{\mathcal{M}},\hat{x})$.

Fix any $\hat{m} \in \mathcal{M}^b$, and consider the new mechanism $(\mathcal{M}', x')$ defined by

$$\mathcal{M}' := \widehat{\mathcal{M}} \cup \{m', m''\} \text{ and } x'(m) := \begin{cases} \hat{x}(\hat{m}) & \text{if } m = m', m'', \\ \hat{x}(m) & \text{otherwise.} \end{cases}$$

Because of the contract pooling property described above, it is easy to see that, for sufficiently small $\varepsilon > 0$, it is optimal for the analyst to report according to

$$\sigma'(s,p)[m] := \begin{cases} 1 & \text{if } p > \bar{p}, m = m', \\ \bar{\alpha} + \varepsilon & \text{if } p = \bar{p}, m = m', \\ 1 - \bar{\alpha} - \varepsilon & \text{if } p = \bar{p}, m = m'', \\ 1 & \text{if } p \in [\underline{p}, \bar{p}), m = m'', \\ \hat{\sigma}(s,p)[m] & \text{if } p < \underline{p}, m \neq m', m'', \\ 0 & \text{otherwise.} \end{cases}$$

When $\varepsilon = 0$, the distribution of analyst types reporting m' under strategy σ' in mechanism $(\mathcal{M}', x')$ (strictly) first-order stochastically dominates that reporting any $m \in \mathcal{M}^b$ under strategy $\hat{\sigma}$ in mechanism $(\widehat{\mathcal{M}}, \hat{x})$. Therefore, the client's posterior belief is $q'(m') > \mathbb{E}[\hat{q}(\mathcal{M}^b)] = \bar{q}$, and the probability with which action b is induced is exactly $T' = \widehat{T}$. Then for sufficiently small $\varepsilon > 0$, action b is induced with strictly higher probability under this new mechanism than under $(\widehat{\mathcal{M}}, \hat{x})$, all the while leaving the bank's payoff from analyst retention unchanged. This contradicts the assumption that $(\widehat{\mathcal{M}}, \hat{x})$ is part of a bank-optimal equilibrium.

Therefore, in any bank-optimal equilibrium where the client's expected payoff is zero, the bank is indirectly implementing its optimal action recommendation rule A^* from Theorem 2.

STEP 5: $(\widehat{\mathcal{M}}, \hat{x})$ IS NOT OPTIMAL FOR κ SMALL. We can write the resulting bank payoff from using mechanism $(\widehat{\mathcal{M}}, \hat{x})$ as

$$\widehat{\Pi} := \widehat{\Pi}_S + \kappa \Pi_P^*,$$

where $\widehat{\Pi}_S$ is its resulting payoff from screening and Π_P^* is the (optimal) payoff from persuasion. On the other hand, the bank's payoff from using the "public prediction mechanism" $(\widetilde{\mathcal{M}}, \tilde{x})$ defined earlier is

$$\widetilde{\Pi} := \Pi_S^* + \kappa \widetilde{\Pi}_P,$$

where (by definition) Π_S^* is the (optimal) payoff from screening and $\widetilde{\Pi}_P$ is the induced payoff from persuasion. We then have $\widetilde{\Pi} > \widehat{\Pi}$ (and therefore a strictly positive payoff for the client) whenever

$$\kappa < \bar{\kappa} := \frac{\Pi_S^* - \widehat{\Pi}_S}{\Pi_P^* - \widetilde{\Pi}_P}.$$

Note that $\bar{\kappa}$ is well-defined and strictly positive: $\Pi_S^* > \widehat{\Pi}_S$ since $\underline{p} = \hat{p} < \frac{1}{2}$, and so the pooling under $(\widehat{\mathcal{M}}, \hat{x})$ is suboptimal, and likewise $\Pi_P^* > \widetilde{\Pi}_P$ since the persuasion under $(\widetilde{\mathcal{M}}, \tilde{x})$ is also suboptimal. ■

REFERENCES

Antic, N., A. Chakraborty, and R. Harbaugh (2024): "Subversive Conversations," *Journal of Political Economy*, forthcoming.

Beauchêne, D., J. Li, and M. Li (2019): "Ambiguous Persuasion," *Journal of Economic Theory*, 179, 312–365.

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

Best, J., and D. Quigley (2024): "Persuasion for the Long Run," *Journal of Political Economy*, 132(5), 1740–1791.

Beyer, A., and I. Guttman (2011): "The Effect of Trading Volume on Analysts' Forecast Bias," *Accounting Review*, 86(2), 451–481.

Bizzotto, J., E. Perez-Richet, and A. Vigier (2021): "Communication via Third Parties," Unpublished manuscript, Sciences Po.

Bose, S., and L. Renou (2014): "Mechanism Design with Ambiguous Communication Devices," *Econometrica*, 82(5), 1853–1872.

Bradshaw, M., Y. Ertimur, and P. O'Brien (2017): "Financial Analysts and Their Contribution to Well-Functioning Capital Markets," *Foundations and Trends in Accounting*, 11(3), 119–191.

Bradshaw, M. T. (2011): "Analysts' Forecasts: What Do We Know After Decades of Work?," Unpublished manuscript, Boston College.

Che, Y.-K., W. Dessein, and N. Kartik (2013): "Pandering to Persuade," *American Economic Review*, 103(1), 47–79.

Chen, C.-Y., and P. F. Chen (2009): "NASD Rule 2711 and Changes in Analysts' Independence in Making Stock Recommendations," *Accounting Review*, 84(4), 1041–1071.

Chen, P. F., K. E. Novoselov, and Y. Wang (2018): "Regulatory Effects on Analysts' Conflicts of Interest in Corporate Financing Activities: Evidence from NASD Rule 2711," *Journal of Corporate Finance*, 48, 658–679.

Corrao, R., and Y. Dai (2024): "The Bounds of Mediated Communication," Unpublished manuscript, Massachusetts Institute of Technology.

Crane, A., and K. Crotty (2020): "How Skilled are Security Analysts?," *Journal of Finance*, 75(3), 1629–1675.

Dasgupta, S. (2023): "Screening Knowledge," Unpublished manuscript, University of Chicago.

Deb, R., M. M. Pai, and M. Said (2018): "Evaluating Strategic Forecasters," *American Economic Review*, 108(10), 3057–3103.

Doval, L., and V. Skreta (2022): "Mechanism Design with Limited Commitment," *Econometrica*, 90(4), 1463–1500.

Dworczak, P., S. D. Kominers, and M. Akbarpour (2021): "Redistribution Through Markets," *Econometrica*, 89(4), 1665–1698.

Elliott, G., and A. Timmermann (2016): "Forecasting in Economics and Finance," *Annual Review* *of Economics*, 8(1), 81–110.

Fang, B., O.-K. Hope, Z. Huang, and R. Moldovan (2020): "The Effects of MiFID II on Sell-Side Analysts, Buy-Side Analysts, and Firms," *Review of Accounting Studies*, 25(3), 855–902.

Fréchette, G. R., A. Lizzeri, and J. Perego (2022): "Rules and Commitment in Communication: An Experimental Analysis," *Econometrica*, 90(5), 2283–2318.

Guo, Y., and L. Mota (2021): "Should Information be Sold Separately? Evidence from MiFID II," *Journal of Financial Economics*, 142(1), 97–126.

Inderst, R., and M. Ottaviani (2012): "Competition through Commissions and Kickbacks," *American Economic Review*, 102(2), 780–809.

Ivanov, M. (2024): "Perfect Robust Implementation by Private Information Design," *Economic Theory*, 78(3), 753–787.

Jackson, A. R. (2005): "Trade Generation, Reputation, and Sell-Side Analysts," *Journal of Finance*, 60(2), 673–717.

Jehiel, P., and B. Moldovanu (2001): "Efficient Design with Interdependent Valuations," *Econometrica*, 69(5), 1237–1259.

Johnson, J. P., and D. P. Myatt (2006): "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 96(3), 756–784.

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

Kamenica, E., and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101(6), 2590–2615.

Kamenica, E., K. Kim, and A. Zapechelnyuk (2021): "Bayesian Persuasion and Information Design: Perspectives and Open Issues," *Economic Theory*, 72, 701–704.

Koessler, F., and V. Skreta (2023): "Informed Information Design," *Journal of Political Economy*, 131(11), 3186–3232.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017): "Persuasion of a Privately Informed Receiver," *Econometrica*, 85(6), 1949–1964.

Kothari, S., E. So, and R. Verdi (2016): "Analysts' Forecasts and Asset Pricing: A Survey," *Annual Review of Financial Economics*, 8, 197–219.

Krähmer, D. (2020): "Information Disclosure and Full Surplus Extraction in Mechanism Design," *Journal of Economic Theory*, 187, 105020.

——— (2021): "Information Design and Strategic Communication," *American Economic Review: Insights*, 3(1), 51–66.

Kreps, D. M., and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50(4), 863–894.

Lang, M., J. Pinto, and E. Sul (2024): "MiFID II Unbundling and Sell-Side Analyst Research," *Journal of Accounting and Economics*, 77(1), 101617.

Lehmer, T., B. Lourie, and D. Shanthikumar (2022): "Brokerage Trading Volume and Analysts' Earnings Forecasts: A Conflict of Interest?," *Review of Accounting Studies*, 27(2), 441–476.

Lin, X., and C. Liu (2024): "Credible Persuasion," *Journal of Political Economy*, 132(7), 2228–2273.

Lipnowski, E., D. Ravid, and D. Shishkin (2022): "Persuasion via Weak Institutions," *Journal of Political Economy*, 130(10), 2705–2730.

Marinovic, I., M. Ottaviani, and P. N. Sørensen (2013): "Forecasters' Objectives and Strategies," in *Handbook of Economic Forecasting*, ed. by G. Elliott, and A. Timmermann, vol. 2B, pp. 691–720.

Mathevet, L., D. Pearce, and E. Stacchetti (2024): "Reputation and Information Design," Unpublished manuscript, New York University.

Min, D. (2021): "Bayesian Persuasion under Partial Commitment," *Economic Theory*, 72(3), 743–764.

Myerson, R. B. (1986): "Multistage Games with Communication," *Econometrica*, 54(2), 323–358.

Nguyen, A., and T. Y. Tan (2021): "Bayesian Persuasion with Costly Messages," *Journal of Economic Theory*, 193, 105212.

Ottaviani, M., and P. N. Sørensen (2006a): "Professional Advice," *Journal of Economic Theory*, 126(1), 120–142.

——— (2006b): "Reputational Cheap Talk," *RAND Journal of Economics*, 37(1), 155–175.

——— (2006c): "The Strategy of Professional Forecasting," *Journal of Financial Economics*, 81(2), 441–466.

Perez-Richet, E. (2014): "Interim Bayesian Persuasion: First Steps," *American Economic Review*, 104(5), 469–474.

Perez-Richet, E., and V. Skreta (2022): "Test Design Under Falsification," *Econometrica*, 90(3), 1109–1142.

Salamanca, A. (2021): "The Value of Mediated Communication," *Journal of Economic Theory*, 192, 105191.

Strack, P., and K. H. Yang (2024): "Privacy Preserving Signals," *Econometrica*, 92(6), 1907–1938.

Zapechelnyuk, A. (2023): "On the Equivalence of Information Design by Uninformed and Informed Principals," *Economic Theory*, 76(4), 1051–1067.