

DISCUSSION PAPER SERIES

DP17955

INDIRECT PERSUASION

Rahul Deb, Mallesh Pai and Maher Said

**INDUSTRIAL ORGANIZATION AND
ORGANIZATIONAL ECONOMICS**

CEPR

INDIRECT PERSUASION

Rahul Deb, Mallesh Pai and Maher Said

Discussion Paper DP17955
Published 04 March 2023
Submitted 30 November 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization
- Organizational Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Rahul Deb, Mallesh Pai and Maher Said

INDIRECT PERSUASION

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 beliefs 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.

JEL Classification: D82, D83, D86, G14, G18

Keywords: Mechanism design, Persuasion, Communication, Commitment

Rahul Deb - rahul.deb@utoronto.ca
University Of Toronto and CEPR

Mallesh Pai - mallesh.pai@rice.edu
Rice University and CEPR

Maher Said - maher.said@nyu.edu
New York University and CEPR

Acknowledgements

We are grateful to Anna Sanktjohanser for her thoughtful discussion of an early version of this paper, as well as to Heski Bar-Isaac, Laura Doval, Michael Grubb, Fei Li, Volker Nocke, Eduardo Perez-Richet, Ludovic Renou, Vasiliki Skreta, Ina Taneva, 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 continued and 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.

INDIRECT PERSUASION

Rahul Deb*
Malleesh M. Pai†
Maher Said‡

March 3, 2023

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 beliefs 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, University of Toronto, rahul.deb@utoronto.ca

†Department of Economics, Rice University, malleesh.pai@rice.edu

‡Stern School of Business, New York University, maher.said@nyu.edu

We are grateful to Anna Sanktjohanser for her thoughtful discussion of an early version of this paper, as well as Heski Bar-Isaac, Laura Doval, Michael Grubb, Fei Li, Volker Nocke, Eduardo Perez-Richet, Ludovic Renou, Vasiliki Skreta, Ina Taneva, and participants at various seminars and conferences for their many helpful comments and suggestions. We would also like to thank CEPR and the Studienzentrums Gerzensee, where much of this work was completed, for their hospitality. Deb thanks the SSHRC for their continued and 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 action. In these models, the sender communicates via *commitment* to an information structure (equivalently, to a statistical experiment or to a state-dependent, possibly mixed, action recommendation). 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.¹ 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 an informed sender with unfettered access to information structures or statistical experiments. Instead, we consider an *uninformed principal* and a *strategic agent* with private information. In particular, the agent has private information about both the state and their own underlying ability, and is purely concerned with the employment terms they receive from the principal. The principal’s payoff depends on both their employment relationship with the agent and the action of an uninformed *receiver*. Further, motivated by applications, we assume 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 need to design a contract in which the agent’s public messages permit the principal to optimally manage the employment relationship *and* persuade the receiver.

To contextualize our results, consider first a natural benchmark where the principal can freely intermediate between the agent and the receiver. In this case, the principal could employ a direct mechanism that incentivizes the agent to truthfully report her information privately to the principal. The mechanism specifies, based on this report, the agent’s contractual outcome; it also garbles the agent’s information and transmits it directly to the receiver. The principal can thus separately solve the agency and persuasion problems. In particular, this implies that if the principal had direct access to the agent’s private information (and did not need to elicit it), she would persuade the receiver in exactly the same way as in this benchmark.

Our main result shows that the outcome of the private communication benchmark described above can *always* be achieved, even when the principal is restricted to using mechanisms that specify contractual terms based only on *public communication* by the agent. This implies that optimal persuasion of the receiver is obtainable by commitment to a standard employment contract without also requiring commitment to arbitrary information structures. Instead, the optimal contract leads the agent to internalize the principal’s objective. This result allows us to interpret the commitment assumption in Bayesian persuasion as a reduced-form stand-in for an informed intermediary facing optimal incentives designed by

¹Some recent papers use ideas from repeated games and reputation to provide foundations for this commitment assumption: see, for example, [Best and Quigley \(2022\)](#) and [Mathevet, Pearce, and Stacchetti \(2022\)](#).

their employer.

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.² These analysts are employed by investment 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), but also wish to persuade the investors to take actions that may benefit the bank (via commissions, brokerage fees, and the like). Consequently, this industry is highly regulated: starting with NASD Rule 2711 in 2002 and culminating with MiFID II in 2018, direct interaction between banks' research and investment arms is prohibited. (See Section 2.1 for additional institutional detail.) This is precisely what is captured by our assumption of public communication by the agent (analyst): analysts are not allowed to communicate privately with the banks that employ them. An immediate consequence of our main result is that such regulation can be rendered ineffective by a bank that appropriately designs its analysts' employment contracts. Interestingly, we also show that the bank-optimal implementation takes a natural form in which the analyst's recommendations correspond to the commonly observed five-point asset rating scale (typically "strong buy," "buy," "hold," "sell," and "strong sell"). Finally, we provide conditions under which a regulatory intervention that eliminates uncertainty about the interpretation of analyst recommendations can strictly increase client welfare.

While our main application is to 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 preferences may not align with those of their employer; 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 to clients during site visits. 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 what the investigator finds and reports to the judge. Such strategic information intermediaries are ubiquitous.

1.1. SUMMARY OF MODEL AND RESULTS

Before proceeding, we first describe our model and the main result in slightly more detail, and provide some intuition.

In our model, there are three parties: a principal, an agent, and a receiver. The principal and the receiver have no private information. The agent has private information about the state and also about their underlying ability; in keeping with standard mechanism design practice, we refer to the totality of the agent's private information as simply their type.³ The agent's utility depends only on the underlying state,

²There 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.

³In our concluding remarks ([Section 4](#)), we also discuss an extension of our model where, instead of simply being endowed with it, the agent must take a costly action to acquire her private information.

their type, and their bilateral employment terms with the principal, while the receiver's utility depends on the state and their action. The principal's payoff depends on all the model parameters—that state, the agent's type, the employment terms, and the receiver's action—but it is separable in the variables that govern the interaction with the agent and the receiver respectively. Importantly, any communication by the agent must be publicly observed by both the principal and the receiver. In particular, this rules out secret messages from the agent to the principal.

As we described above, if the agent *could* secretly communicate with the principal, a standard revelation principle applies. The principal could commit to a direct mechanism in which the agent reports their type privately to the principal. Based on this report, the mechanism specifies both the employment terms between the principal and agent, and the distribution from which the action that the agent should recommend to the receiver realizes. Incentives to the agent are provided via the employment terms (for instance, bonus payments or whether or not the agent is promoted or fired). Because the principal's payoff is separable and the agent's payoff does not depend on the receiver's action, the principal can thus separately solve the agency and persuasion problems. In other words, if the principal had direct access to the agent's private information (without needing to elicit it), they would persuade the receiver in exactly the same way as in the optimal direct mechanism when this information is privately held by the agent.

Now let us consider our setting where any message from the agent must be public. We consider a game in which the principal is restricted to (indirect) public communication mechanisms. Formally, the game we study proceeds as follows. 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. Only the agent observes the selected contract—the receiver does not. The contract specifies the terms of the agent's employment 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 agent's message, the receiver updates their belief about the underlying state and, finally, chooses an action.

This class of public communication mechanisms has at least three desirable properties. First, the principal only commits to a standard employment contract and not to an information structure: as we argued above, the former is well-understood and -motivated, while the latter is less so. Second, in keeping with real-world regulations, there is no private communication from the agent to the principal as the agent's message is public. Third, as we discuss below, such mechanisms are commonly observed in practice.

Our main result shows that the outcome of any principal optimal direct mechanism (that requires private communication from the agent) can *always* be achieved as an equilibrium of the game in which the principal employs public communication mechanisms. In particular, this implies that optimal persuasion is obtainable by standard commitment to an employment contract without requiring commitment to an information structure. More broadly, our result permits us to reinterpret the commitment assumption in Bayesian persuasion as a reduced-form implementation of optimal employment incentives for an information intermediary.

To understand the tension, recall that when the message from the agent is private, the principal can solve the screening and persuasion portions of their problem “separately.” But when communication by the agent must be public, a potential conflict arises: the agent's public message influences both their employment terms *and* the receiver's action. Therefore, one might imagine that the principal would face a tradeoff 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.

Our construction leverages the fact that the receiver observes the agent’s public message but not the principal’s choice of mechanism. 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. We show that an appropriately chosen randomization over public communication mechanisms by the principal allows for both optimal screening of the agent and optimal persuasion of the receiver. Thus, the deliberate introduction of vague language is often necessary for indirect persuasion to be optimal.

1.2. RELATED LITERATURE

The closest papers within the Bayesian persuasion literature are [Lipnowski, Ravid, and Shishkin \(2022\)](#) and [Min \(2021\)](#).⁴ These papers examine a setting where the sender initially commits to an information structure but, with an exogenously given probability, is released from her commitment after observing the state and can send any other message. As in our setting, this implies that the receiver must interpret messages while accounting for the sender’s *partial* commitment. In this vein, [Nguyen and Tan \(2021\)](#) consider a sender who first commits to an information structure but then privately observes the message (generated by the information structure) which she can manipulate at a cost. [Bizzotto, Perez-Richet, and Vigier \(2021\)](#) also consider mediated communication, but in a setting where a principal commits to an information structure and uses monetary transfers to induce a third party to effectively communicate it to a receiver. We take a different view by assuming that direct commitment to an information structure is not possible but that standard organizational contracts can instead suffice for optimal persuasion.

The resulting optimal indirect mechanism we construct makes use of deliberately “vague” public communication to generate uncertainty and persuade the receiver. A similar idea appears in the literature on mechanism design and 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 \(2022\)](#)—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 stark contrast to our result, however, these works require commitment to the (randomization over the) information design.

In our model, the contracting terms offered by the principal serve dual functions: in addition to indirectly persuading the receiver, they also screen the agent.⁵ 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 \(2005\)](#), who analyzes a reputational cheap-talk model of sell-side analyst communication with career concerns—but sets aside the bank’s organizational design problem and simply takes 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

⁴The excellent surveys of [Bergemann and Morris \(2019\)](#) and [Kamenica \(2019\)](#) describe the broader literature.

⁵A 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.

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) may be of independent interest.

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 throughout the body of the paper wherever relevant.

1.3. STRUCTURE OF THE PAPER

We first present the application to financial analysts in Section 2. This serves the dual purpose of motivating the environment and highlighting the key intuition that underlies our main result. Then, in Section 3, we define the general model and present the main theorem. These sections are intended to be self-contained. A reader who is interested in the application but not in the full generality can proceed to the concluding remarks in Section 4 after reading Section 2; conversely, readers interested in the general result can skip straight to Section 3.

2. APPLICATION: REGULATING FINANCIAL ANALYSTS

We consider a game with three players: a bank (the principal), an analyst (the agent), and a representative client meant to capture the “market” (the receiver). These three players interact in the following stylized setting that captures key features of the market for sell-side investment research. Section 3 makes clear that our “punchline” carries over to substantially more general models.

The state: There is a binary and unobserved *state* $\omega \in \Omega := \{b, s\}$ of the world, either *buy* or *sell*. The prior probability that the state is ω is given by $\pi_\omega \in (0, 1)$ where $\pi_b + \pi_s = 1$.

The state captures the unknown value of an asset, with *b* referring to whether the asset is of high value (in which case the client should buy), and *s* low value (in which case the client should sell or go short). The state is publicly revealed after the client makes their trading decision.

Analyst ability: The analyst has a privately known *ability* $\theta \in \Theta := \{h, l\}$ which is either *high* or *low*; this type captures her skill. The likelihood that the analyst is ability θ is given by $\mu_\theta \in (0, 1)$ with $\mu_h + \mu_l = 1$.

Analyst information: The analyst learns about the state by observing an informative signal. We do not explicitly model this signal; instead we describe the analyst’s *information* via the distribution of the possible posterior beliefs that can arise by Bayesian updating upon observing signal realizations. Since the state is binary, the posterior belief can be summarized by a probability $p \in [0, 1]$ that the state is $\omega = b$.

We denote the (cumulative) distribution of posterior beliefs for the ability- θ analyst by F_θ , and we assume that this distribution has mean π_b (since the expectation of the posteriors must equal the prior) and admits a density f_θ . Note that we do not impose any additional structure on these distributions other than that

they differ, so that $F_h \neq F_l$. In particular, while it may be natural to assume that the high-ability analyst has “better” information (whether in the sense of Blackwell (1953) or some other stochastic order), our result does not require such an assumption.

Public communication mechanisms: A *public communication mechanism* is a pair (\mathcal{M}, x) consisting of a finite *public message space* \mathcal{M} for the analyst and a *retention rule* $x : \mathcal{M} \times \Omega \rightarrow [0, 1]$. The message space should be interpreted as the set of possible public recommendations (for instance, “buy,” “sell,” or “hold”) that the analyst can issue. The retention rule specifies the probability $x(m, \omega)$ with which the analyst is retained after sending message $m \in \mathcal{M}$ when state $\omega \in \Omega$ is realized.

Note that a public communication mechanism (\mathcal{M}, x) is an *indirect* mechanism and *not* a direct mechanism as might be employed upon invocation of the revelation principle. This is because (i) the public message space \mathcal{M} need not correspond to the analyst’s private information $\Theta \times [0, 1]$; (ii) messages $m \in \mathcal{M}$ are publicly observed, so the analyst cannot privately communicate information to the bank (consistent with the prohibitions imposed by conflict-of-interest regulations in the real world); and (iii) the mechanism only specifies contractual terms x with the analyst and not action recommendations provided by the bank to the client.

Bank strategy: The bank chooses a finite public message space \mathcal{M} , a finite set of public communication mechanisms $\mathbb{M} \subset [0, 1]^{\mathcal{M} \times \Omega}$, and a distribution $\rho \in \Delta(\mathbb{M})$ on this set.⁶ We assume that the message space \mathcal{M} and the set of mechanisms \mathbb{M} are publicly observed by both the analyst and the client, while the bank’s chosen distribution ρ is private and unobservable. In particular, the analyst only observes the realization of the bank’s (mixed) strategy. Thus, there is common knowledge of the set of possible contracts, but only the bank and analyst know which specific contract governs their relationship. In other words, this assumption amounts to the client knowing the set of possible contracts that might be offered to financial analysts that work at the bank but not the exact contract that is offered to any particular analyst.

Analyst preferences: The analyst wants to maximize the probability that she is retained. Formally, this implies that the analyst’s utility from choosing a message $m \in \mathcal{M}$ when state ω realizes is simply $x(m, \omega)$. Therefore, when her information is $p \in [0, 1]$, her expected utility from message m is

$$px(m, b) + (1 - p)x(m, s).$$

Note that this payoff only depends on the analyst’s information but *not* on her ability. Since trading commissions accrue to the bank, the analyst’s payoff also does not depend on the client’s actions.

Analyst strategy: After observing the set of possible mechanisms \mathbb{M} and realized public communication mechanism (\mathcal{M}, x) resulting from the bank’s strategy, the analyst chooses a message. The ability- θ analyst’s strategy or *recommendation* is denoted by $\sigma_\theta(x, p) \in \Delta(\mathcal{M})$. In words, when the bank selects the mechanism (\mathcal{M}, x) and the analyst’s posterior belief is p , the analyst chooses a (potentially mixed) recommendation $\sigma_\theta(x, p)$. (Note that we suppress the dependence of the analyst’s strategy on the bank’s choice of \mathbb{M} and \mathcal{M} to simplify notation.)

The client: The client updates their *belief* $q : \mathcal{M} \rightarrow [0, 1]$ about the state after observing the bank’s set

⁶We impose the restriction to finite messages and a finite set of mechanisms for two reasons. First, it makes our result starker: finite mechanisms suffice to implement the full-commitment solution even though the analyst’s private information is continuous. Second, it obviates the need for the additional measure-theoretic formalism required to define mixed strategies.

of possible mechanisms \mathbb{M} and the analyst's recommendation $m \in \mathcal{M}$. (As with the analyst strategy, we suppress the dependence on the former for ease of notation.) We denote the client's belief that the state is $\omega = b$ following a recommendation m by $q(m) \in [0, 1]$.

The client uses this belief to determine whether to *trade* and, if so, in which direction. Formally, the client picks an action a from the set $\mathcal{A} := \{\varphi, b, s\}$; here φ denotes no trade whereas b and s denote *buying* and *selling*, respectively. The notation overload with the state is deliberate; as we describe next, the client wants to match their action to the state.

The client's payoff from choosing an action a is given by

$$\mathbb{E}_{q(m)} \left[\mathbb{1}\{a = \omega\}(v - c) + \mathbb{1}\{a \neq \omega, a \neq \varphi\}(-v - c) \right].$$

In words, the client trades at a transaction cost $c > 0$. If the action a matches the state ω , they earn a payoff $v > c$. If the action a mismatches ω , they incur a loss $-v$. The payoff from not trading and taking action φ is 0. The expectation in the above expression is taken with respect to the client's updated belief $q(m)$ following a message m .

The client's action strategy is a mapping from each history to the set \mathcal{A} . The relevant history for the client is the message space \mathcal{M} and set of mechanisms \mathbb{M} chosen by the bank followed by the message $m \in \mathcal{M}$ reported by the analyst. For brevity, we suppress the dependence on the former and define the client's *action strategy*

$$\alpha : \mathcal{M} \rightarrow \mathcal{A}$$

as a mapping from the set of messages to the set of actions. Note that this definition restricts the client to follow a pure strategy, but this is without loss since we will examine the equilibrium that leads to the highest payoff for the bank. In any such equilibrium, we can simply break ties in favor of the bank whenever the client is indifferent.

Observe that, given these preferences, the client has a *cutoff belief* $\bar{q} \in (\frac{1}{2}, 1)$ such that it is optimal for them to choose b whenever $q(m) \geq \bar{q}$ and s whenever $q(m) \leq 1 - \bar{q}$. This captures the fact that the client only wants to trade when they are sufficiently confident that the asset's value will either appreciate or depreciate. Therefore, if the prior π_b lies in $(1 - \bar{q}, \bar{q})$, the client will not trade absent sufficiently informative analyst recommendations.

Bank preferences: The bank has dual objectives: it wants to determine the ability of the analyst so that only the high-ability type is retained, but it also wants the analyst to provide recommendations that induce trade and thereby generate commissions.

Formally, the payoff of the bank from offering a public communication mechanism (\mathcal{M}, x) is

$$\begin{aligned} \Pi(\mathcal{M}, x, \sigma, \alpha) &= \mu_h \int_0^1 \mathbb{E}_{\sigma_h(x,p)} [\kappa \mathbb{1}\{\alpha(m) \neq \varphi\} + (px(m, b) + (1 - p)x(m, s))] dF_h(p) \\ &\quad + \mu_l \int_0^1 \mathbb{E}_{\sigma_l(x,p)} [\kappa \mathbb{1}\{\alpha(m) \neq \varphi\} - \gamma (px(m, b) + (1 - p)x(m, s))] dF_l(p). \end{aligned}$$

The first $\mathbb{E}_{\sigma_\theta(x,p)}[\kappa \mathbb{1}\{\alpha(m) \neq \varphi\}]$ term in each integral captures the likelihood that the client trades based on the analyst's recommendation; the expectation is taken with respect to the analyst's (potentially

mixed) strategy that governs the distribution of messages m . Thus, this term reflects the volume of trade and the bank’s commissions from trading activity, where the constant parameter $\kappa > 0$ captures the weight placed on this objective. Equivalently, this term can be interpreted as the bank’s desire for informative analyst recommendations that push the client’s posterior outside the no-trade range $(1 - \bar{q}, \bar{q})$.

The second term in each integral reflects the fact that the bank only wants to retain the high-ability analyst. Formally, the bank receives a payoff of 1 and $-\gamma$ whenever the bank retains the high- and low-ability analyst, respectively, and a payoff of 0 if the analyst is fired.

Equilibrium: We analyze perfect Bayesian equilibria in this game. These consist of strategies for the client, analyst, and bank, along with consistent beliefs, that satisfy the following properties: (i) the client’s action strategy prescribes a best response to their beliefs which are, in turn, derived by Bayes’ rule for all on-path messages sent by the analyst; (ii) the analyst’s recommendation is a best response to the (on- or off-path) mechanism offered by the bank; and (iii) the bank’s strategy is a best response to the client’s and analyst’s strategies.

A *bank-optimal equilibrium* yields the bank its maximal payoff across all perfect Bayesian equilibria.

As mentioned above, perfect Bayesian equilibrium requires the analyst to best respond to any public communication mechanism offered by the bank. This does not, however, restrict the client’s beliefs after off-path play—for instance, after a deviation by the bank to a different message space or set of mechanisms, or by the analyst to a message that does not lie in the support of her strategy. As will shortly become clear, however, our main insight and results do not hinge on a particular or unrealistic choice of off-path beliefs for the client.

To make the timing explicit, [Figure 1](#) presents a flow chart describing the *public communication game*.

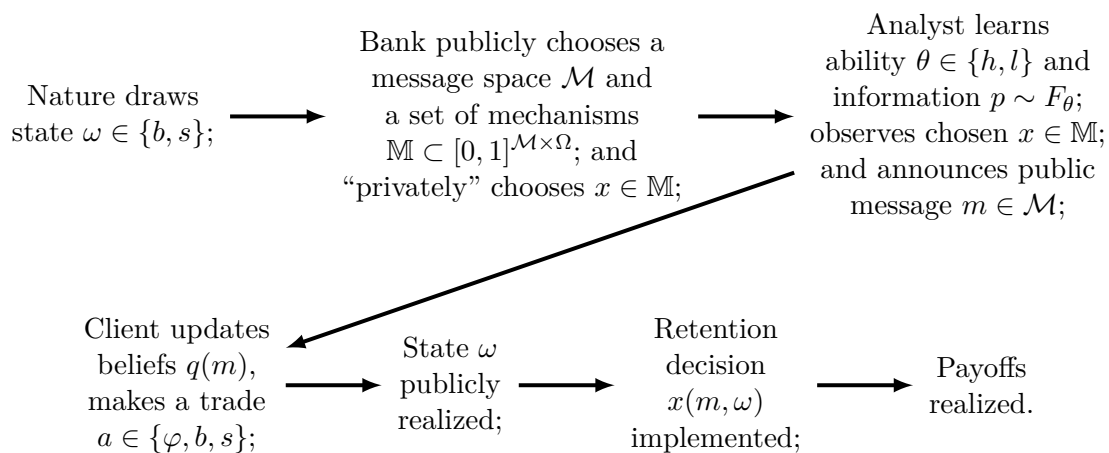


Figure 1: Timing of the public communication game.

2.1. MAPPING THE MODEL TO THE APPLICATION

Despite our (deliberate) simplifications, the main features of the model closely align with our application, and the analysis that follows is an independent contribution to the financial economics literature

on analysts. The bank-optimal equilibrium from this simple model has qualitative features that are observed in practice, and the results have important implications for the role of regulation in the industry for financial advice.

Recall that in the class of mechanisms we consider, there is no private communication from the analyst to the bank. This is motivated by several conflict-of-interest regulations enacted by the Securities and Exchange Commission starting with NASD Rule 2711 (Research Analysts and Research Reports) in 2002. The goal of these regulations is 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 investment banking business and commission revenues. In our model, the bank does not control the analyst's recommendations, and its retention policy only depends on the analyst's accuracy (captured by the relationship between the message and the realized state). The assumption that the retention rule is unobserved by the client corresponds to the privacy of employment terms and contracts in practice.

We assume that the bank incentivizes the analyst via her retention decision but not by tying financial compensation to the accuracy of the recommendation. This mirrors empirical evidence that shows that forecast accuracy is not driven by compensation incentives but is instead by the threat of termination. The recent survey of the literature on financial analysts by [Kothari, So, and Verdi \(2016\)](#) (with further relevant references found therein) concludes that "the evidence suggests that small deviations in accuracy have a minimal impact on analyst compensation, but large (negative) forecast inaccuracy can affect analyst wealth by increasing the probability of dismissal." That said, as we show in the next section, our main insight remains unchanged if we were to include financial compensation as part of the bank's mechanism.

The fact that analysts differ in ability (see [Crane and Crotty \(2020\)](#) for recent evidence) and that the bank only wants to retain those with higher abilities should be uncontroversial. This payoff captures, in a simple reduced-form way, the long-term value of increasing the organization's human capital. We assume that analysts generate revenue for the bank via trading fees, and so analyst advice has to be sufficiently informative to influence the market and increase the volume of trade. Therefore, investing in human capital within the organization by retaining and promoting higher-skilled analysts can improve the bank's prospects for future persuasion. 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.

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 \(2011\)](#). The simplification we employ is that the client only decides on the direction (long or short) of the trade but not its size. This can be easily generalized to allow the bank's payoff to be a more complex function of client beliefs that captures how the volume of trade depends on the degree of optimism or pessimism in the analyst's recommendation.

2.2. REGULATING COMMUNICATION IS INEFFECTIVE

If the principal had the ability to act as a Myersonian intermediary with full commitment to contracting decisions and information structures, the revelation principle would apply and it would be without loss to consider only incentive compatible direct revelation mechanisms. In terms of practical applicability, a direct mechanism (fully defined and described below) has two undesirable features. First, the analyst shares her private information with the bank, with the latter directly choosing its desired recommenda-

tions to the client; as discussed above, this is explicitly illegal. Second, a direct mechanism will typically require commitment to a mixed action recommendation as a function of the analyst’s information; as discussed in the introduction, this may require an unrealistic level of commitment power. Despite these potential shortcomings, the revelation principle implies that such mechanisms provide an upper bound on the principal’s attainable payoffs—which we will show can be achieved even with partial commitment only to contractual terms with the analyst that depend only on her public communications.

Formally, a *direct mechanism* is a pair (X, A) consisting of a *retention rule*

$$X : \Theta \times [0, 1] \times \Omega \rightarrow [0, 1]$$

and an *action recommendation*

$$A : \Theta \times [0, 1] \rightarrow \Delta(\mathcal{A}).$$

In words, the analyst reports her private ability and information to the bank. This determines both the retention decision X (which also conditions on the eventual observation of the realized state) and the (potentially stochastic) action recommendation A conveyed to the client.

If the bank could choose any direct mechanism, it would pick one that maximizes its payoff subject to *incentive compatibility* constraints for the analyst and the client. In particular, incentive compatibility requires the analyst to report her private information truthfully, so that

$$(\theta, p) \in \operatorname{argmax}_{\theta', p'} \{pX(\theta', p', b) + (1 - p)X(\theta', p', s)\} \quad (\text{IC-A})$$

for all $\theta \in \Theta$ and $p \in [0, 1]$. (Since the analyst only cares about the retention decision, the recommended action does not enter this constraint.) Likewise, incentive compatibility for the client requires them to optimally follow all recommended actions: a recommendation a in the support of A must be optimal with respect to the client’s updated posterior belief $q(a)$, so that

$$a \in \operatorname{argmax}_{a' \in \mathcal{A}} \left\{ \mathbb{E}_{q(a)} \left[\mathbb{1}\{a' = \omega\}(v - c) + \mathbb{1}\{a' \neq \omega, a' \neq \varphi\}(-v - c) \right] \right\}. \quad (\text{IC-C})$$

Thus, the *optimal direct mechanism* (X^*, A^*) solves the “full-commitment” problem

$$\max_{X, A} \left\{ \begin{array}{l} \int_0^1 \left[\mu_h (pX(h, p, b) + (1 - p)X(h, p, s)) dF_h(p) \right. \\ \quad \left. - \mu_l \gamma (pX(l, p, b) + (1 - p)X(l, p, s)) dF_l(p) \right] \\ \quad \left. + \kappa \int_0^1 \left[\mu_h \Pr[A(h, p) \neq \varphi] dF_h(p) + \mu_l \Pr[A(l, p) \neq \varphi] dF_l(p) \right] \right\} \quad (\text{FC})$$

subject to (IC-A) and (IC-C).

Here, the first integral represents the payoff from analyst retention while the second is that from trading commissions (where $\Pr[A(\theta, p) \neq \varphi]$ is the probability the mechanism recommends—and obediently induces via (IC-C)—trading). Observe that the bank’s problem is separable. Only the first integral in the bank’s payoff in (FC) depends on the retention decision X , so the analyst incentive compatibility

condition (IC-A) only constrains these terms. Likewise, the action recommendation A only enters the second integral, and so the client incentive compatibility constraint (IC-C) only constrains these terms of the bank's payoff. It is worth emphasizing that this implies that the optimal direct mechanism's action recommendation A^* is identical to what the bank would choose if it had direct access to the analyst's information p and its *only* goal was maximizing trading commissions.

Our first result characterizes the optimal direct mechanism (X^*, A^*) that solves the full-commitment problem (FC). To streamline the statement of the result, we use the shorthand notation $X^*(\theta, p, \cdot)$ to denote the vector $(X^*(\theta, p, b), X^*(\theta, p, s))$.

THEOREM 1. *The optimal retention rule X^* that solves problem (FC) is such that either*

- (i) *it is optimal to not screen and always retain the analyst, so $X^*(\theta, p, \cdot) = (1, 1)$ for all $p \in [0, 1]$ and $\theta \in \Theta$;*
- (ii) *it is optimal to not screen and always fire the analyst, so $X^*(\theta, p, \cdot) = (0, 0)$ for all $p \in [0, 1]$ and $\theta \in \Theta$; or*
- (iii) *it is optimal to nontrivially screen, in which case there exist cutoffs $\tilde{p}_* \leq 1/2 \leq \tilde{p}^*$ such that the retention rules take one of the following forms:*

$$\begin{aligned}
(a) \quad X^*(\theta, p, \cdot) &= \begin{cases} \left(0, \frac{\tilde{p}_*}{1-\tilde{p}_*}\right) & \text{if } p < \tilde{p}_*, \\ (1, 0) & \text{if } p \geq \tilde{p}_*; \end{cases} \\
(b) \quad X^*(\theta, p, \cdot) &= \begin{cases} (0, 1) & \text{if } p \leq \tilde{p}^*, \\ \left(\frac{1-\tilde{p}^*}{\tilde{p}^*}, 0\right) & \text{if } p > \tilde{p}^*; \end{cases} \text{ or} \\
(c) \quad X^*(\theta, p, \cdot) &= \begin{cases} (0, 1) & \text{if } p < \tilde{p}_*, \\ (\tilde{x}_b, \tilde{x}_s) & \text{if } p \in [\tilde{p}_*, \tilde{p}^*], \text{ where } \tilde{x}_b := \frac{(1-\tilde{p}_*)(2\tilde{p}^*-1)}{\tilde{p}^*-\tilde{p}_*} \text{ and } \tilde{x}_s := \frac{\tilde{p}^*(1-2\tilde{p}_*)}{\tilde{p}^*-\tilde{p}_*}. \\ (1, 0) & \text{if } p > \tilde{p}^*, \end{cases}
\end{aligned}$$

In addition, there are cutoffs $\tilde{q}^ \geq \tilde{q}_*$ such that the optimal action recommendation A^* solving (FC) is*

$$A^*(h, p) = A^*(l, p) = \begin{cases} b & \text{if } p > \tilde{q}^*, \\ s & \text{if } p < \tilde{q}_*, \\ \varphi & \text{otherwise.} \end{cases}$$

Before proceeding to the main result of this section—that regulation constraining communications between the bank and the analyst is ineffective—we briefly discuss and provide some intuition for the structure of the optimal direct mechanism.

The retention rule X^* either involves no screening as in cases (i) and (ii), or screening is nontrivial and takes a simple cutoff form as in case (iii). In case (i), the analyst is always retained regardless of private information; this can be optimal when the analyst is very likely to be high ability (so μ_h is close to 1) or the cost γ of hiring the low-ability analyst is small. Similarly, in case (ii), the analyst is never retained regardless of private information; this can be optimal when the analyst is very likely to be low ability

(so μ_h is close to 0) or the cost γ of hiring the low-ability analyst is large. Effective screening may also be impossible if the low-ability analyst has “better” information—a counterintuitive scenario that our model permits but which is unrealistic in the context of the application.

Even when nontrivial screening *is* beneficial for the bank, as in case (iii), the optimal retention rule does not depend on the analyst’s reported private ability. To see why, note that the analyst’s interim payoff—given her information p —does not depend on her ability θ . Therefore, any incentive compatible retention rule must offer the same probability of retention to an agent as a function of p , regardless of her ability. Thus, it is not possible to screen on the ability dimension of the analyst’s private information, and the principal must rely on the accuracy of the analyst’s recommendations in order to indirectly evaluate and screen ability.

More specifically, the principal must reward the analyst via retention for information that is sufficiently accurate while punishing inaccuracy with termination. The cutoff \tilde{p}^* is thus chosen to maximally separate types based on their relative likelihood of possessing information in favor of state b : reporting a posterior belief $p > \tilde{p}^*$ that is ultimately corroborated when state b is realized serves as evidence that the analyst is indeed accurate and of high ability, while a contradictory realization of state s is evidence that the analyst is inaccurate and of low ability. Likewise, the cutoff \tilde{p}_* is chosen to maximally separate types based on their relative likelihood of possessing information in favor of state s .

The optimal action recommendation A^* is designed to induce maximal trade. Recall that the client engages in trade only when they are sufficiently confident; that is, only when their posterior belief leaves the interval $(1 - \bar{q}, \bar{q})$. So suppose the bank were to recommend action b only when $p > \bar{q}$ (regardless of the analyst’s ability θ). This yields a posterior belief $q > \bar{q}$, 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 cutoff \tilde{q}^* is chosen to push the client’s posterior belief down as close to possible to \bar{q} , so that the client is *just* indifferent to buying when action b is recommended. A symmetric intuition applies to the cutoff \tilde{q}_* and s recommendations. (Note that, for some parameters, it is possible to have $\tilde{q}_* = \tilde{q}^*$, in which case the bank induces trade with probability one.) Note that both the optimal retention rule and optimal action recommendation are effectively finite; this feature appears in the outcome-equivalent equilibrium of the public communication game (in which the bank chooses a finite message space and a finite set of mechanisms).

The characterization of X^* may be of independent interest, as the bank’s screening problem constitutes a novel instance of a “mechanism design without transfers” setting.⁷ Conversely, the derivation of the action recommendation A^* in [Theorem 1](#) follows from standard Bayesian persuasion arguments.

We are now in a position to present the main result of this section: regulation is toothless, and the bank can generate exactly the same trading activity using a public communication mechanism as it could if it directly observed the analyst’s private information.

THEOREM 2. *Let (X^*, A^*) be an optimal direct mechanism that solves (FC). There is a perfect Bayesian equilibrium of the public communication game that is outcome-equivalent to (X^*, A^*) ; that is, there is a*

⁷Local incentive compatibility does not suffice to characterize the optimal retention rule, and so the argument employs a combination of ironing and Lagrangian methods to incorporate a global incentive constraint. See [Appendix A.1](#) for a self-contained analysis of this problem.

bank-optimal equilibrium in which, for every $(\theta, p) \in \Theta \times [0, 1]$, the action distribution chosen by the client is $A^*(\theta, p)$ and the probability of analyst retention in each state $\omega \in \Omega$ is $X^*(\theta, p, \omega)$.

When the optimal screening rule X^* is as in [Theorem 1](#)'s cases (i) or (ii), so there is no screening and the analyst is either always retained or always fired, she can be trivially induced to garble information on behalf of the bank (since she would be indifferent between all messages). The more surprising cases are when the optimal screening rule is as in case (iii) and the bank engages in nontrivial screening. The underlying intuition for [Theorem 2](#) is demonstrated most clearly by considering the special case where the optimal direct mechanism (X^*, A^*) is as described in case (iii).(c) of [Theorem 1](#): the optimal retention rule X^* is piecewise constant with cutoffs $\tilde{p}_* < \tilde{p}^*$, and the optimal action recommendation rule A^* is piecewise constant with cutoffs $\tilde{q}^* \geq \tilde{q}_*$.

The taxation principle implies that there is no loss in the screening problem from allowing the analyst to directly choose her preferred retention probabilities from the range of X^* ; by implication, there is no loss in coarsening the principal's information about the agent (relative to the full revelation of a direct mechanism) to the partition induced by the cutoffs \tilde{p}_* and \tilde{p}^* . Meanwhile, the client information structure generated by the action recommendation rule A^* is itself partitional. Thus, for a public communication mechanism to achieve the same outcomes and payoffs as the optimal direct mechanism, it must be the case that it generates an information partition that is mutually compatible with *both* the principal's and the client's partitions.

This mutual compatibility is relatively easy to achieve when $\tilde{q}_* < \tilde{p}_*$ and $\tilde{q}^* > \tilde{p}^*$ (so the cutoffs corresponding to (X^*, A^*) are as displayed in [Figure 2](#)): the meet of the principal's and the client's information partitions changes neither the retention probabilities nor the actions. To see this latter point, note that since A^* is optimal, standard Bayesian persuasion arguments imply that $\tilde{q}^* < \bar{q}$ and $\tilde{q}_* > 1 - \bar{q}$. Therefore, any refinement of the client's partition when the analyst's information lies in $(A^*)^{-1}(\varphi) = [\tilde{q}_*, \tilde{q}^*]$ will continue to yield posterior beliefs in the "no-trade" inaction region $(1 - \bar{q}, \bar{q})$.




Figure 2: Optimal direct mechanism when $\tilde{q}_* < \tilde{p}_*$ and $\tilde{q}^* > \tilde{p}^*$.

With this in mind, define the message space $\mathcal{M}^* = \{B, b, \varphi, s, S\}$ and a public communication mechanism $x^* : \mathcal{M}^* \times \Omega \rightarrow [0, 1]$ such that

- x^* pools messages in $\{B, b\}$ and treats them as reports of $p > \tilde{p}^*$;
- x^* treats the message φ as a report of $p \in [\tilde{p}_*, \tilde{p}^*]$; and
- x^* pools messages in $\{s, S\}$ and treats them as reports of $p < \tilde{p}_*$.

Since messages B and b are identical in terms of their implications for retention, the analyst is indifferent between these two reports; therefore, by publicly reporting B only when her information is $p > \tilde{q}^*$ and publicly reporting b when her information is $p \in (\tilde{p}^*, \tilde{q}^*]$, the analyst induces the client to trade only

on the more “extreme” recommendation B .⁸ A symmetric argument applies for the messages s and S , implying that this (indirect) public communication mechanism (depicted in Figure 3) exactly implements the optimal direct mechanism.




Figure 3: Bank-optimal equilibrium when $\tilde{q}_* < \tilde{p}_*$ and $\tilde{q}^* > \tilde{p}^*$.

Note that the message space \mathcal{M}^* above has more elements than just the three actions available to the client. 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, each message in \mathcal{M}^* has a natural interpretation: messages B and b can be interpreted as “strong buy” and “weak buy” recommendations, message φ is a “hold” or neutral recommendation; and messages s and S are “weak sell” and “strong sell” recommendations. Thus, optimal public communication in our environment corresponds directly to the traditional five-point rating scales employed by sell-side analysts.

Critical to our construction above is the fact that the client’s action is, in a sense, more sensitive than the principal’s screening to changes in the analyst’s information, and so it was possible to refine the principal’s information partition while maintaining the client’s incentives. But in other cases—for instance, when $\tilde{q}^* < \tilde{p}^* < \bar{q}$, as in Figure 4—the intuition above no longer applies. In order to implement the optimal




Figure 4: Optimal direct mechanism when $\tilde{q}_* < \tilde{p}_*$ and $\tilde{q}^* < \tilde{p}^* < \bar{q}$.

screening rule X^* , the principal must be able to identify when the analyst’s information p is just above or just below the screening cutoff \tilde{p}^* ; that is, for all sufficiently small $\varepsilon > 0$, the analyst’s message when $p = \tilde{p}^* + \varepsilon$ must be distinct from that when $p = \tilde{p}^* - \varepsilon$. But if the public messaging strategy is deterministic, the client is also able to distinguish between $p = \tilde{p}^* + \varepsilon$ and $p = \tilde{p}^* - \varepsilon$. This reduces the mass of analyst types above \tilde{q}^* that send messages resulting in trade, thereby decreasing the principal’s payoff from persuasion.

How then can the optimal direct mechanism be implemented with public communication? More specifically, how is it possible for the analyst to publicly report distinct messages when $p = \tilde{p}^* + \varepsilon$ and $p = \tilde{p}^* - \varepsilon$ *without* the client distinguishing between them?

⁸The analyst’s indifference is not critical here or in the argument that follows; her incentives can be made strict at an arbitrarily small payoff loss for the bank.

The key is to note that persuading the client to always trade when $p > \tilde{q}^*$ does *not* require the analyst to always report the same message when her information is sufficiently optimistic; it is sufficient for the analyst to report using the same *distribution* of messages when her information is above the threshold \tilde{q}^* . In particular, if the principal randomizes over private contracts, it is possible to induce the requisite analyst randomization over public messages.

So consider the message space $\mathcal{M}^* = \{B, b, \varphi, s, S\}$, and suppose the principal *privately* randomizes with equal probability between the two public communication mechanisms $x_1^* : \mathcal{M}^* \times \Omega \rightarrow [0, 1]$ and $x_2^* : \mathcal{M}^* \times \Omega \rightarrow [0, 1]$ such that

- both x_1^* and x_2^* pool messages in $\{S, s\}$ and treat them as reports of $p < \tilde{p}_*$;
- both x_1^* and x_2^* treat the message φ as a reports of $p \in [\tilde{p}_*, \tilde{q}^*]$;
- x_1^* treats the message b as a report of $p \in (\tilde{q}^*, \tilde{p}^*]$ and the message B as a report of $p > \tilde{p}^*$; and
- x_2^* treats the message b as a report of $p > \tilde{p}^*$ and the message B as a report of $p \in (\tilde{q}^*, \tilde{p}^*]$.

Since the analyst observes the realized retention rule, “sincere” reporting is optimal: it yields her the same mapping from private information to retention probabilities. Meanwhile, since the client does *not* observe the principal’s realized choice of contract, they are unable to distinguish between messages b and B , and so the two messages collectively generate the same posterior belief, and hence the same optimal action, as when $A^* = b$ (in Figure 4). Finally, it is clear that the bank has no incentive to deviate from its strategy since the analyst’s and client’s best responses yield the bank-optimal outcome; thus, the indirect mechanism (depicted in Figure 5) once again exactly implements the optimal direct mechanism.




Figure 5: Bank-optimal equilibrium when $\tilde{q}_* < \tilde{p}_*$ and $\tilde{q}^* < \tilde{p}^* < \bar{q}$.

The construction above for this more “tricky” case suggests that effective indirect persuasion may require the use of intentionally *vague* language: without uncertainty about the interpretation of messages, the principal must compromise on either screening, persuasion, or both. This uncertainty is not unnatural: in practice, there is a substantial lack of clarity about the interpretation of analyst ratings: what exactly is the difference 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? Importantly, however, this uncertainty about the interpretation of language is one-sided. To enable effective screening, the principal and agent *must* have a shared understanding of how public messages will be *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!

Theorem 2 deploys this general intuition beyond just the cases described above: the bank can *always* privately randomize over contractual terms in order to make the client uncertain about the meaning of public messages, thereby implementing its optimal outcome. The result also holds more generally than under the conditions of the stylized model above. In particular, we can allow for multiple states, multiple analyst abilities, mechanisms that include transfers, and multiple actions for the client. We describe a general abstract framework that encompasses these generalizations in [Section 3](#).

2.3. THE WELFARE CONSEQUENCES OF TRANSPARENCY

The above discussion highlights the important role of private bank–analyst contractual terms in permitting the bank-optimal equilibrium to implement the full-commitment outcome. Note that when the cutoffs for the optimal action recommendation satisfy $\tilde{q}_* < \tilde{q}^*$, the client’s expected utility is zero. It is therefore natural to ask whether the client can benefit if the contractual terms offered by the bank to the analyst are instead mandated to be public, and so all communication is transparent?

Clearly, making contractual terms public when $\tilde{q}_* < \tilde{q}^*$ cannot hurt the client, as they can always guarantee themselves the (equilibrium) payoff of zero by not trading. Indeed, the next result shows that public contracting can instead make the client strictly better off.

THEOREM 3. *Suppose $\tilde{q}_* < \tilde{q}^*$ and the cutoffs from every optimal retention rule X^* satisfy $0 < \tilde{p}_* < \tilde{q}_*$ or $\tilde{q}^* < \tilde{p}^* < 1$. Then, there is a cutoff $\bar{\kappa} > 0$ such that, for any weight on trading and commissions $\kappa < \bar{\kappa}$, the client earns a positive payoff in every bank-optimal equilibrium with public contracting.*

The intuition for this result is straightforward. As we illustrated in the previous subsection, when $\tilde{p}_* < \tilde{q}_*$ or $\tilde{q}^* < \tilde{p}^*$, the bank cannot obtain the full-commitment profits without mixing over contractual terms with the analyst.⁹ 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. When the weight κ on the latter is sufficiently low relative to the former, trading commissions are sacrificed. The proof of [Theorem 3](#) shows that, in this case, every bank-optimal equilibrium with public contracting features cutoffs $\bar{p}_s \leq \tilde{q}_* < \tilde{q}^* \leq \underline{p}^b$ (with at least one of the inequalities being strict) such that the client is induced to take action $a = s$ only when the analyst’s information type is $p \leq \bar{p}_s$, and likewise is induced to take action $a = b$ only when $p \geq \underline{p}^b$. Since \bar{p}_s and \underline{p}^b are more “extreme” than the bank-optimal cutoffs \tilde{q}_* and \tilde{q}^* , the client is no longer indifferent between trading and holding, and so must have a strictly positive payoff.

The case where $\tilde{q}_* < \tilde{q}^*$ is the case of more practical interest, as we generally should not expect analyst recommendations to generate trade with probability one. However, and perhaps surprisingly, making the bank–analyst contractual terms public can sometimes be detrimental to the client when $\tilde{q}_* = \tilde{q}^*$. We end this section with an example that demonstrates this.¹⁰

⁹When neither of these inequalities hold, mixing is not necessary and the bank simply chooses a single mechanism. A natural sufficient condition for this is that the information environment is symmetric (so F_h and F_l are both symmetric about their shared mean $\pi_b = \frac{1}{2}$) and F_h dominates F_l in the rotation order. A formal statement and proof are available on request.

¹⁰We will deliberately state this example in terms of its endogenous cutoffs for ease of exposition. Note, however, that it is easy to construct primitives (distributions, probabilities, and costs) that yield cutoffs with the requisite properties.

Suppose that the client’s cutoff belief for trading \bar{q} is sufficiently close to $\frac{1}{2}$ that they are “easily persuadable” and it is possible to induce trade with probability 1. When this is the case, there exists $\hat{q}^- < \frac{1}{2}$ and $\hat{q}^+ > \frac{1}{2}$ such that any action recommendation rule with cutoffs $\tilde{q}^* = \tilde{q}_* = \hat{q} \in [\hat{q}^-, \hat{q}^+]$ are optimal; assume further that $[\hat{q}^-, \hat{q}^+] \subset (1 - \bar{q}, \bar{q})$. Meanwhile, suppose that the optimal retention rule X^* corresponds to case (iii).(a) of [Theorem 1](#) with a single cutoff satisfying $\tilde{p}_* \in [\hat{q}^-, \frac{1}{2})$, so that the optimal retention rule only depends on whether the analyst’s information type is above or below \tilde{p}_* .

It is easy to see that, among all action recommendation rules that always induce trade, the symmetric one with $\tilde{q}^* = \tilde{q}_* = \frac{1}{2}$ is client-optimal. Moreover, [Theorem 2](#) implies that there is a bank-optimal equilibrium—featuring randomization over contracts that is unobserved by the client—that implements both the retention rule X^* and the (constrained) client-optimal action recommendation rule with $\tilde{q}^* = \tilde{q}_* = \frac{1}{2}$. Note, however, that since $\tilde{p}_* \in [\hat{q}^-, \frac{1}{2})$, there is also a bank-optimal equilibrium—with public contracting—that implements X^* and the action recommendation rule with $\tilde{q}^* = \tilde{q}_* = \tilde{p}_*$. This is, of course, strictly worse for the client than the symmetric recommendation rule implemented with private contracts.

The existence of this latter public-contracting equilibrium implies that any other bank-optimal public-contracting equilibrium must also implement X^* and trade with probability 1. But implementing X^* must involve the analyst sending different messages above and below \tilde{p}_* , while implementing any other action recommendation rule that yields trade with probability 1 involves the analyst sending different messages above and below $\hat{q} \neq \tilde{p}_*$, implying that analyst information types $p \in (\min\{\tilde{p}_*, \hat{q}\}, \max\{\tilde{p}_*, \hat{q}\})$ are distinguishable by the client from those outside that interval. But this implies that no trade occurs in this interval, as the client’s posterior belief will lie in the no-trade region $(1 - \bar{q}, \bar{q})$. Thus, the *only* bank-optimal public-contracting equilibrium is that with $\tilde{q}^* = \tilde{q}_* = \tilde{p}_*$, leaving the client worse off than with private contracting.

3. THE GENERAL MODEL

We now set aside the analyst framework and show that our result—a principal can indirectly implement the optimal direct mechanism while relying only on commitment to an employment contract and public communication—applies much more generally. The structure of the presentation mirrors that of the previous section. At the expense of repetition, this makes the mapping from the general model to the application clearer from the outset although we discuss this in [Section 3.1](#).

With this in mind, we consider a sequential game with three players: a principal, an agent, and a receiver. The principal wants to elicit the agent’s private information, but must also incentivize the agent to communicate payoff-relevant information about the economic environment to the receiver.

Note that we will generally restrict attention to finite sets, an assumption that is not necessary for our result but that greatly simplifies exposition and eliminates some technical complications.

The state: We denote by Ω the finite set of possible *states* of the world. These states are distributed according to a commonly known prior distribution $\pi \in \Delta(\Omega)$, where $\pi_\omega > 0$ is the prior probability of any given state $\omega \in \Omega$.

Agent’s private information: The agent’s *type* θ is drawn from a distribution $\mu(\cdot|\omega) \in \Delta(\Theta)$ that depends on the state $\omega \in \Omega$, where Θ is a finite set of possible types.

Once the agent observes her realized type θ , she forms a posterior belief about the state ω using the dis-

tribution μ . Thus the agent’s private type is not only information that determines her value to the firm (such as her ability or skill); it also implicitly contains information about the state that only the agent possesses.

Public communication mechanisms: A *public communication mechanism* is a pair (\mathcal{M}, x) consisting of a finite message space \mathcal{M} and a contract $x : \mathcal{M} \rightarrow \Delta(\mathcal{T})$, where \mathcal{T} is the set of *contractible decisions*. Elements of the set \mathcal{T} are the firm specific decisions that the principal can commit to; these might represent transfers, state-contingent retention probabilities, or some other general contractible outcomes.¹¹

As in Section 2, a public communication mechanism (\mathcal{M}, x) is an *indirect* mechanism: the public message space \mathcal{M} need not correspond to the agent’s private information Θ ; messages $m \in \mathcal{M}$ are publicly observed, so the agent cannot privately communicate information to the principal; and the mechanism only specifies contractual terms x with the agent and not action recommendations provided by the principal to the receiver.

Principal’s strategy: The principal chooses a public finite message space \mathcal{M} , a finite set of public communication mechanisms $\mathbb{M} \subset \Delta(\mathcal{T})^{\mathcal{M}}$, and a distribution $\rho \in \Delta(\mathbb{M})$ on this set. We assume that the message space \mathcal{M} and the set of mechanisms \mathbb{M} are 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 contracts, but only the principal and agent know which specific contract governs their relationship.

Agent’s strategy: After observing the set of possible mechanisms \mathbb{M} and the selected public communication mechanism (\mathcal{M}, x) and learning her type θ , the *agent’s strategy* is a distribution over possible reports in the message space \mathcal{M} .

Receiver: After observing the set of possible mechanisms \mathbb{M} and the agent’s message $m \in \mathcal{M}$, the receiver updates their belief to $q(m) \in \Delta(\Omega)$ and chooses an *action* $a \in \mathcal{A}$, possibly as the realization of a mixed strategy. (Note that the belief $q(m)$ may also depend on the set \mathbb{M} ; our notation suppresses this for convenience.)

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

The principal’s payoff depends on all parameters of the model: the state, the agent’s private information, the contracted decision, and the receiver’s action. The agent’s and the receiver’s utilities only depend on a subset of these parameters, however. The agent’s payoff is independent of the receiver’s action, while the receiver’s payoff is, as in most cheap talk and persuasion settings, a function only of the underlying state and their chosen action.

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

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

Recall that our primary goal is to situate a “standard” persuasion problem in an organizational context

¹¹The structure of the set \mathcal{T} of contractible decisions may incorporate feasibility or resource constraints; we do not, however, explicitly model such constraints, nor do we incorporate any participation constraints in our analysis. While these constraints influence the *form* of any optimal contracts, they are orthogonal to our central insight and do not affect our main result.

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 separability assumption, along with the structure of the agent’s and receiver’s preferences, ensure that a principal with full commitment power could maximize their payoff by separately solving the contracting problem with the agent and the persuasion problem with the receiver. In other words, in an optimal direct mechanism, the principal persuades the receiver in exactly the same way she would if she knew θ (and did not need to incentivize the agent to reveal it).

Suppose, instead, that the agent’s payoff depended on the receiver’s action. The optimal direct mechanism would then have to account for these incentives, distorting the principal’s action recommendation relative to an environment where the principal had direct access to the agent’s information. A similar tradeoff arises when the receiver’s payoff depends on the agent’s type. We therefore rule out such environments as they do not conform to the above mentioned primary goal.

Equilibrium: We analyze perfect Bayesian equilibria of the extensive form game described above. As usual, perfect Bayesian equilibrium requires that the receiver’s strategy maximize their payoff subject to beliefs that are derived using Bayes’ rule for on-path messages (and are unrestricted off-path); that the agent’s strategy maximizes her payoff in all public communication mechanisms (on- and off-path) offered by the principal; and that the principal’s strategy be a best response to those of the receiver and agent.

To summarize and make the timing explicit, [Figure 6](#) presents a flow chart depicting the more general public communication game we study in this section.

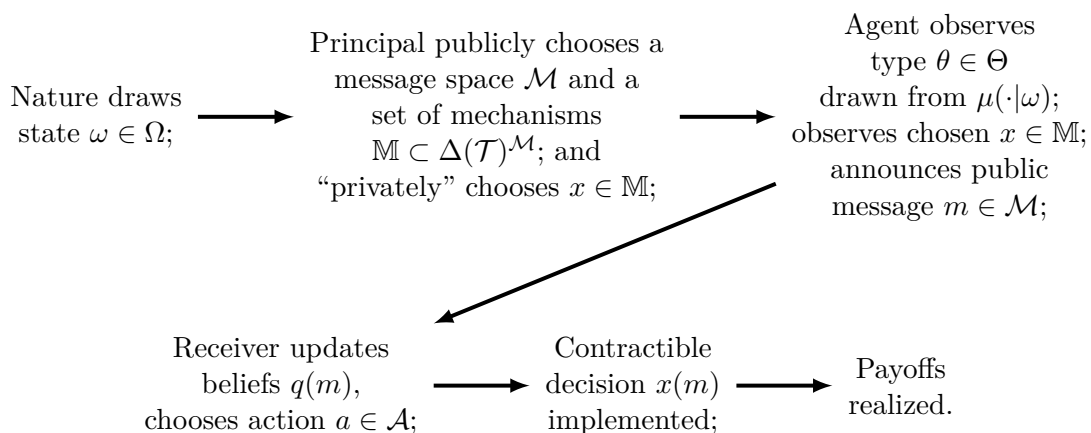


Figure 6: Timing of the general public communication game.

3.1. EMBEDDING THE APPLICATION IN THE GENERAL MODEL

Before describing the main result of the paper, we briefly map the analyst application from the previous section into our general framework. Readers who skipped [Section 2](#) can proceed directly to [Section 3.2](#).

In the application, the agent’s private information was two-dimensional, with one dimension representing her ability and the other her information about the underlying state. Relative to the definitions of this section, we employed slightly different notation in [Section 2](#) that was more natural for the application. In the present framework, the realization of the type θ determines, via Bayes’ rule and the distribution

μ , the agent's posterior over states. Since the type is private information, this posterior is also privately known. Moreover, conveying information about θ to the principal and the receiver allow those two players to similarly infer information about the state. One additional difference is that the state information p in the application to take continuous values in $[0, 1]$. This is not a substantive difference, however, and the restriction to finite types in this section is simply for ease of exposition.

The mapping to the set of contractible decisions is a little more subtle since the public communication mechanisms in Section 2 permit the principal to condition on the realized state, whereas we have not explicitly allowed for this here. It is possible, however, to embed that framework within the present environment. Suppose the set of contractible decisions is $\mathcal{T} = [0, 1] \times [0, 1]$, with the interpretation that $(t_b, t_s) \in \mathcal{T}$ represents the probability t_b of retaining the analyst when $\omega = b$ and the probability t_s of retaining her when $\omega = s$. The analyst's utility from any $t \in \mathcal{T}$ in state $\omega \in \{b, s\}$ is simply $u_A(\omega, \theta, t) = t_\omega$ and therefore, when her type θ has a posterior belief p for state b , her expected utility is $pt_b + (1 - p)t_s$. Thus, the generality of the set of contractible decisions \mathcal{T} allows us to incorporate the state dependence of the mechanism.

Finally, the principal and the client map directly from the application into the general framework.

3.2. THE OPTIMALITY OF PUBLIC COMMUNICATION MECHANISMS

Our main insight is that the principal can achieve their full-commitment payoff using (indirect) public communication mechanisms. To state this result, we need to define direct mechanisms.

A *direct mechanism* consists of a contracting rule $X : \Theta \rightarrow \Delta(\mathcal{T})$ and an action recommendation $A : \Theta \rightarrow \Delta(\mathcal{A})$. A direct mechanism (X, A) is *incentive compatible* if it satisfies the following two sets of constraints:

- For any recommendation $a \in \mathcal{A}$ that lies in the support of A , it is incentive compatible for the receiver to choose a over any alternative $a' \in \mathcal{A}$; that is, for all a in the support of A and all $a' \in \mathcal{A}$,

$$\mathbb{E}_{q(a)}[u_R(\omega, a)] \geq \mathbb{E}_{q(a)}[u_R(\omega, a')], \quad (\text{R-IC})$$

where the expectations are taken with respect to the receiver's posterior belief $q(a)$ (formed via Bayes' rule) after receiving action recommendation a .

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

$$\mathbb{E}_{\mu(\cdot|\theta), X(\theta)}[u_A(\omega, \theta, t)] \geq \mathbb{E}_{\mu(\cdot|\theta'), X(\theta')}[u_A(\omega, \theta, t)], \quad (\text{A-IC})$$

where the expectations are taken over the state $\omega \in \Omega$ (distributed according to the conditional distribution $\mu(\omega|\theta)$ derived from the primitives π and μ) and the contracting decision $t \in \mathcal{T}$ (distributed according to $X(\cdot)$, which in turn depends on the agent's report θ or θ').

An *optimal* direct mechanism (X^*, A^*) maximizes the principal's payoff subject to the incentive compatibility constraints (R-IC) and (A-IC). Note that the separability of the principal's payoff implies that they can separately choose X to maximize $u_P^A(\omega, \theta, t)$ subject to (A-IC) and A to maximize $u_P^B(\omega, a)$ subject to (R-IC). 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 .

We are now in a position to state our main result.

THEOREM 4. *Fix an optimal direct mechanism (X^*, A^*) . There is a perfect Bayesian equilibrium of the public communication game that is outcome-equivalent to this mechanism: that is, for all agent types $\theta \in \Theta$, the distribution of contractual outcomes is $X^*(\theta)$ and distribution of actions taken by the receiver is $A^*(\theta)$.*

The proof of this result is in [Appendix D](#), but its intuition is the same as that of the analogous result ([Theorem 2](#)) for our application.¹² In order to optimally screen according to X^* , the principal needs to elicit the agent’s type—but optimal persuasion according to A^* will not fully convey that information to the receiver. However, the principal can mix over public communication mechanisms that correspond to the same contracting rule while inducing uncertainty about the public meaning of each message. Because the agent observes the realized mechanism, she can always best respond, and optimal screening can be implemented. Since the receiver does not observe the realized mechanism, they do not learn which public communication mechanism the agent is best responding to. Appropriately mixing thus allows the principal to indirectly garble the agent’s private information and achieve optimal persuasion as in A^* .

4. 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 and a privately informed agent—and show that optimal persuasion can be achieved by partial commitment only to standard contractual terms (and not to information structures). This serves as an organizational economics micro-foundation for commitment in strategic communication.

Our main result, [Theorem 4](#), 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 $\mu(\omega, e) \in \Delta(\Theta)$. 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.) It is straightforward to show (under similar conditions on payoffs) that the principal’s optimal direct mechanism can be indirectly implemented as a perfect Bayesian equilibrium of a public communication game (defined in analogous fashion to that of [Section 3](#)). 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. This caveat aside, the same intuition applies even more broadly to environments with multiple receivers, many agents, and dynamically evolving private types.

¹²It is worth noting that, in order to simplify notation and exposition, the formal proof employs more randomization and messages than is strictly necessary.

A. PROOF OF THEOREM 1

We begin by noting that the principal's full-commitment problem (FC) is separable in X and A . Fixing any solution (X^*, A^*) to (FC), the retention rule X^* therefore solves

$$\max_X \left\{ \int_0^1 \left[\mu_h (pX(h, p, b) + (1-p)X(h, p, s)) dF_h(p) - \mu_l \gamma (pX(l, p, b) + (1-p)X(l, p, s)) dF_l(p) \right] \right\} \quad (\text{FC}_A)$$

subject to (IC-A),

while the recommendation rule A^* solves

$$\max_A \left\{ \kappa \int_0^1 \left[\mu_h \Pr[A(h, p) \neq \varphi] dF_h(p) + \mu_l \Pr[A(l, p) \neq \varphi] dF_l(p) \right] \right\} \quad (\text{FC}_C)$$

subject to (IC-C).

We can therefore characterize the optimal retention rule X^* (Theorem A.1) and the optimal persuasion rule A^* (Theorem A.2) independently; the two results jointly prove Theorem 1.

A.1. CHARACTERIZING THE OPTIMAL RETENTION RULE

Consider any direct revelation retention rule $X : \Theta \times [0, 1] \times \Omega \rightarrow [0, 1]$. We define the analyst's utility from reporting her type as (θ', p') when her true type is (θ, p) as

$$U(\theta', p' | \theta, p) := pX(\theta', p', b) + (1-p)X(\theta', p', s).$$

Analyst incentive compatibility (IC-A) therefore requires that $U(\theta, p | \theta, p) \geq U(\theta', p' | \theta, p)$ for all types $(\theta', p'), (\theta, p) \in \Theta \times [0, 1]$.

LEMMA A.1. *Suppose the screening rule $X : \Theta \times [0, 1] \times \Omega \rightarrow [0, 1]$ satisfies (IC-A). Then for almost all $p \in [0, 1]$, $X(h, p, \omega) = X(l, p, \omega)$ for all $\omega \in \Omega$.*

PROOF. Note first that (IC-A) implies that

$$U(h, p | h, p) \geq U(l, p | h, p) \text{ and } U(l, p | l, p) \geq U(h, p | l, p) \text{ for all } p \in [0, 1].$$

But notice that $U(\theta', p' | \theta, p)$ only depends on the analyst's reported type θ' and not on her true type θ , so that

$$U(h, p | h, p) = U(h, p | l, p) \text{ and } U(l, p | l, p) = U(l, p | h, p) \text{ for all } p \in [0, 1].$$

This of course implies that $U(h, p | h, p) = U(l, p | l, p)$ for all $p \in [0, 1]$: the expected utility of an analyst who reports her private information truthfully is the same regardless of her ability (whether true or reported) $\theta \in \Theta$.

Furthermore, since (IC-A) requires that $U(\theta, p | \theta, p) \geq U(\theta', p' | \theta, p)$ for all $p, p' \in [0, 1]$ and all $\theta \in \Theta$, standard arguments (see Milgrom and Segal (2002), for instance) imply that $U(\theta, p | \theta, p)$ is almost

everywhere differentiable in p , and that we can write

$$U(\theta, p|\theta, p) = U(\theta, 0|\theta, 0) + \int_0^p [X(\theta, z, b) - X(\theta, z, s)] dz \text{ for all } p \in [0, 1].$$

Thus, $X(h, p, b) - X(h, p, s) = X(l, p, b) - X(l, p, s)$ for almost all $p \in [0, 1]$. But we can write

$$\begin{aligned} X(h, p, s) &= U(h, p|h, p) - p[X(h, p, b) - X(h, p, s)] \\ &= U(l, p|l, p) - p[X(l, p, b) - X(l, p, s)] = X(l, p, s) \end{aligned}$$

for almost all $p \in [0, 1]$. This immediately implies that the retention rules are identical (almost everywhere) across the two abilities h and l . \blacksquare

As a consequence of [Lemma A.1](#) above, we can simplify our notation and define a retention rule to be a pair of functions

$$x_b : [0, 1] \rightarrow [0, 1] \text{ and } x_s : [0, 1] \rightarrow [0, 1],$$

where $x_\omega(p)$ is the probability that the analyst is retained when she reports information p and the realized state is ω . The incentive compatibility constraint (**IC-A**) rewritten with this simplified notation is then

$$U(p', p) := px_b(p') + (1 - p)x_s(p') \leq U(p, p) \text{ for all } p', p \in [0, 1]. \quad (\text{IC-}p)$$

Thus, we can rewrite the bank's screening problem (**FC_A**) as

$$\begin{aligned} \max_{x_b, x_s} \left\{ \int_0^1 U(p, p) d[F_h(p) - \beta F_l(p)] \right\} \\ \text{subject to } (\text{IC-}p) \text{ and } x_b(p), x_s(p) \in [0, 1], \end{aligned} \quad (\mathcal{P}_A)$$

where we have renormalized the principal's screening payoff by μ_h , and so

$$\beta := \frac{\mu_l}{\mu_h} \gamma.$$

We are now in a position to state the result characterizing the optimal retention rule.

THEOREM A.1. *The optimal retention rule solving (\mathcal{P}_A) is such that either*

- (i) *it is optimal to not screen and always retain the analyst, so $x_b^*(p) = x_s^*(p) = 1$ for all $p \in [0, 1]$;*
- (ii) *it is optimal to not screen and always fire the analyst, so $x_b^*(p) = x_s^*(p) = 0$ for all $p \in [0, 1]$; or*
- (iii) *it is optimal to nontrivially screen, and there exist cutoffs $\tilde{p}_* \leq 1/2 \leq \tilde{p}^*$ such that x_b^* and x_s^* take one of the following forms:*

$$(a) \quad x_b^*(p) = \begin{cases} 0 & \text{if } p < \tilde{p}_*, \\ 1 & \text{if } p \geq \tilde{p}_*, \end{cases} \text{ and } x_s^*(p) = \begin{cases} \frac{\tilde{p}_*}{1-\tilde{p}_*} & \text{if } p < \tilde{p}_*, \\ 0 & \text{if } p \geq \tilde{p}_*; \end{cases}$$

$$(b) \ x_b^*(p) = \begin{cases} 0 & \text{if } p \leq \tilde{p}^*, \\ \frac{1-p}{\tilde{p}^*} & \text{if } p > \tilde{p}^*, \end{cases} \text{ and } x_s^*(p) = \begin{cases} 1 & \text{if } p \leq \tilde{p}^*, \\ 0 & \text{if } p > \tilde{p}^*; \end{cases} \text{ or}$$

$$(c) \ x_b^*(p) = \begin{cases} 0 & \text{if } p < \tilde{p}_*, \\ \frac{(1-\tilde{p}_*)(2\tilde{p}^*-1)}{\tilde{p}^*-\tilde{p}_*} & \text{if } p \in [\tilde{p}_*, \tilde{p}^*], \\ 1 & \text{if } p > \tilde{p}^*, \end{cases} \text{ and } x_s^*(p) = \begin{cases} 1 & \text{if } p < \tilde{p}_*, \\ \frac{\tilde{p}^*(1-2\tilde{p}_*)}{\tilde{p}^*-\tilde{p}_*} & \text{if } p \in [\tilde{p}_*, \tilde{p}^*], \\ 0 & \text{if } p > \tilde{p}^*. \end{cases}$$

In the next subsection, we will prove this result. It is worth mentioning that, in the proof, we will provide sufficient conditions under which it is optimal for the bank not to screen. Moreover, we will be more precise about how the cutoffs \tilde{p}_* and \tilde{p}^* are defined. We elide these details here since the simplified statement above conveys the main qualitative features of the optimal retention rule.

A.2. PROOF OF THEOREM A.1

Standard arguments from quasilinear mechanism design with transfers imply that the incentive compatibility constraint (IC- p) is equivalent to an envelope and a monotonicity condition. Formally, these are

$$U(p, p) = px_b(p) + (1-p)x_s(p) = U(q, q) + \int_q^p \Delta(z)dz \text{ for all } p, q, \text{ and} \quad (\text{ENV})$$

$$\Delta(p) := x_b(p) - x_s(p) \text{ is nondecreasing in } p, \quad (\text{MON})$$

respectively. We now use condition (ENV) to reformulate the objective in the brokerage's problem (\mathcal{P}_A). Integration by parts yields

$$\begin{aligned} \int_0^1 U(p, p)d[F_h(p) - \beta F_l(p)] &= \left[U(p, p)[F_h(p) - \beta F_l(p)] \right]_{p=0}^1 - \int_0^1 \Delta(p)[F_h(p) - \beta F_l(p)]dp \\ &= \int_0^1 \Delta(p)[\beta F_l(p) - F_h(p)]dp + (1-\beta)x_b(1), \end{aligned}$$

where we have made use of the fact that $F_h(1) = F_l(1) = 1$ and $F_h(0) = F_l(0) = 0$. In addition, note that the integral condition in (ENV), evaluated at $p = 1$ and $q = 0$, becomes

$$\int_0^1 \Delta(z)dz = x_b(1) - x_s(0). \quad (\text{INT})$$

Finally, note that the feasibility constraints on x_b and x_s combine with (MON) to imply that

$$\Delta(p) \in [-x_s(0), x_b(1)] \text{ for all } p \in [0, 1] \text{ and } x_b(1), x_s(0) \in [0, 1]. \quad (\text{FEAS})$$

We now define the following *relaxed problem*:

$$\begin{aligned} \max_{\Delta, x_s(0), x_b(1)} \left\{ \int_0^1 \Delta(p)[\beta F_l(p) - F_h(p)]dp + (1-\beta)x_b(1) \right\} \\ \text{subject to (INT), (MON), and (FEAS).} \end{aligned} \quad (\mathcal{R}_A)$$

Notice that we have replaced the envelope constraint (ENV) (which must hold for all $p, q \in [0, 1]$) with

the weaker constraint (INT), which is simply the envelope condition (ENV) evaluated at a specific pair $p = 1$ and $q = 0$. We will solve problem (\mathcal{R}_A) and show that it admits a solution that satisfies the stronger constraint (ENV), and therefore also maximizes (\mathcal{P}_A).

Before proceeding, it is worth mentioning a minor detail left implicit in the above. Recall that we optimize over retention rules x_b and x_s in problem (\mathcal{P}_A), while optimization in the relaxed problem (\mathcal{R}_A) is over Δ , $x_s(0)$, and $x_b(1)$. However, the retention rule x_b and x_s corresponding to any Δ that satisfies the envelope condition (ENV) is fully determined given $x_s(0)$. This is because, for any $p \in [0, 1]$, we have

$$x_s(p) = U(p, p) - p\Delta(p) = U(0, 0) + \int_0^p \Delta(z)dz - p\Delta(p) = x_s(0) + \int_0^p \Delta(z)dz - p\Delta(p)$$

and

$$x_b(p) = x_s(p) + \Delta(p) = x_s(0) + \int_0^p \Delta(z)dz + (1 - p)\Delta(p).$$

Define the shorthand notation $u_0 := x_s(0)$ and $u_1 := x_b(1)$, and let

$$\mathcal{L}(\Delta|u_0, u_1, \lambda) := \int_0^1 [\Delta(p)[\beta F_l(p) - F_h(p)] + (1 - \beta)u_1 + \lambda(u_1 - u_0 - \Delta(p))] dp$$

be the Lagrangian corresponding to the objective function of (\mathcal{R}_A) and the integral constraint (INT), where $\lambda \in \mathbb{R}$ is the multiplier on that constraint.

Let

$$\Delta_{u_0, u_1, \lambda}^* \in \underset{\Delta}{\operatorname{argmax}} \{ \mathcal{L}(\Delta|u_0, u_1, \lambda) \} \text{ subject to (FEAS) and (MON);}$$

that is, the function $\Delta_{u_0, u_1, \lambda}^*$ maximizes $\mathcal{L}(\Delta|u_0, u_1, \lambda)$, subject to (FEAS) and (MON), for some *fixed values* of $u_0, u_1 \in [0, 1]$ and $\lambda \in \mathbb{R}$.

We will solve (\mathcal{R}_A) as follows. Consider values of u_0 , u_1 , and λ such that there is a maximizer $\Delta_{u_0, u_1, \lambda}^*$ of the above problem that additionally satisfies the integral constraint (INT). Among all such values of u_0 , u_1 , and λ , we then pick u_0^* , u_1^* , and λ^* that, jointly with their corresponding maximizer $\Delta_{u_0^*, u_1^*, \lambda^*}^*$ satisfying (INT), yield the highest value of the objective function in (\mathcal{R}_A). In other words, $\Delta_{u_0^*, u_1^*, \lambda^*}^*$, u_0^* and u_1^* solve the relaxed problem.¹³

We will employ the ironing procedure in [Toikka \(2011\)](#) to derive $\Delta_{u_0, u_1, \lambda}^*$. But first, we unavoidably need to define some additional notation.¹⁴ Let

$$h(p) := \beta F_l(p) - F_h(p) \text{ and } H(p) := \int_0^p h(z)dz,$$

and let

$$G := \operatorname{conv} H$$

¹³As this is a linear program, the validity of our Lagrangian approach follows from Theorems 8.3.1 and 8.4.1 of [Luenberger \(1969\)](#).

¹⁴An aside to the well-versed reader: note that the integral in the objective function in (\mathcal{R}_A) is taken with respect to the uniform measure (since it is the difference of expected utilities across the two analyst types). Therefore, there is no implicit change of variables in the definitions that follow.

be the lower convex envelope of H .¹⁵ Note that this implies

$$G(0) = H(0) = 0 \text{ and } G(1) = H(1) = (\beta - 1)\pi_s,$$

where the latter follows from the observation that $\int_0^1 F_\theta(p)dp = \pi_s$ for all $\theta \in \Theta$. G is continuously differentiable everywhere on $(0, 1)$, as it is the convex hull of a (by definition) differentiable function. So for all $p \in (0, 1)$, let

$$g(p) := G'(p),$$

and extend g to the rest of $[0, 1]$ by continuity; the convexity of G implies that g is nondecreasing.

Applying [Toikka \(2011\)](#)'s Theorem 3.7, the function $\Delta_{u_0, u_1, \lambda}^*$ solves

$$\sup_{\Delta} \{\mathcal{L}(\Delta | u_0, u_1, \lambda)\} \text{ subject to (FEAS) and (MON)}$$

if, and only if, $\Delta_{u_0, u_1, \lambda}^*$ is a pointwise optimizer of the Lagrangian, so

$$\Delta_{u_0, u_1, \lambda}^*(p) \in \operatorname{argmax}_{\Delta(p) \in [-u_0, u_1]} \{\Delta(p)g(p) + (1 - \beta)u_1 + \lambda(u_1 - u_0 - \Delta(p))\} \text{ almost everywhere;}$$

$\Delta_{u_0, u_1, \lambda}^*$ is monotone (which can be guaranteed since the objective has, by construction, increasing differences); and, finally, $\Delta_{u_0, u_1, \lambda}^*$ satisfies the “pooling property” (which requires $\Delta_{u_0, u_1, \lambda}^*$ be constant on all open intervals $I \subset [0, 1]$ for which $G(p) < H(p)$ for all $p \in I$).

Pointwise optimization of $[\Delta(p)g(p) + (1 - \beta)u_1 + \lambda(u_1 - u_0 - \Delta(p))]$ for each $p \in [0, 1]$ yields

$$\Delta_{u_0, u_1, \lambda}^*(p) \begin{cases} = -u_0 & \text{if } g(p) < \lambda, \\ \in [-u_0, u_1] & \text{if } g(p) = \lambda, \\ = u_1 & \text{if } g(p) > \lambda. \end{cases} \quad (\text{OPT})$$

The first lemma derives a simple sufficient condition under which it is optimal for the bank to *not* screen. We will see shortly that this condition is also necessary whenever $\beta \leq 1$; the condition is not tight, however, when $\beta > 1$.

LEMMA A.2. *Suppose $H(p) \geq pH(1)$ for all $p \in [0, 1]$. Then the optimal mechanism involves no screening, so either $x_b^*(p) = x_s^*(p) = 0$ for all $p \in [0, 1]$ or $x_b^*(p) = x_s^*(p) = 1$ for all $p \in [0, 1]$.*

PROOF. For any $p \in [0, 1]$ and any $(\zeta, p_1, p_2) \in [0, 1]^3$ such that $\zeta p_1 + (1 - \zeta)p_2 = p$, we have

$$\zeta H(p_1) + (1 - \zeta)H(p_2) \geq \zeta p_1 H(1) + (1 - \zeta)p_2 H(1) = pH(1).$$

This in turn implies

$$G(p) = \min\{\zeta H(p_1) + (1 - \zeta)H(p_2) | (\zeta, p_1, p_2) \in [0, 1]^3 \text{ and } \zeta p_1 + (1 - \zeta)p_2 = p\} \geq pH(1).$$

But letting $\zeta = p$, $p_1 = 1$, and $p_2 = 0$ in the above, we see that the lower bound $pH(1)$ is actually

¹⁵Formally, $G(p) := \min\{\zeta H(p_1) + (1 - \zeta)H(p_2) | (\zeta, p_1, p_2) \in [0, 1]^3 \text{ and } \zeta p_1 + (1 - \zeta)p_2 = p\}$ for all $p \in [0, 1]$.

achievable, and so $G(p) = pH(1)$. Thus, $H(p) \geq pH(1)$ for all $p \in [0, 1]$ implies that $G(p) = pH(1)$ and $g(p) = H(1)$ for all $p \in [0, 1]$.

We cannot, however, have $H(p) = pH(1)$ for all $p \in (0, 1)$ as well. If this were true, it would imply that

$$\beta F_l(p) - F_h(p) = h(p) = H(1) = (\beta - 1)\pi_s \text{ for all } p \in [0, 1],$$

which is clearly not possible when $\beta \neq 1$ since $F_l(0) = F_h(0) = 0$. If $\beta = 1$, then this means that $F_l(p) = F_h(p)$ for all p , contradicting the assumption that the two analyst types receive different information. Thus, we must have $H(p) \geq pH(1)$ for all $p \in [0, 1]$, with strict inequality for some p .

Now since $g(p)$ is a constant, we set $\lambda^* = g(p)$ and note that any constant function

$$\Delta_{u_0, u_1, \lambda^*}^*(p) = \bar{\Delta} \in [-u_0, u_1]$$

is a pointwise optimizer (OPT). Note that this $\Delta_{u_0, u_1, \lambda^*}^*$ is monotone and trivially satisfies the pooling property.

In order for this $\Delta_{u_0, u_1, \lambda^*}^*$ to satisfy the integral constraint (INT), we need that

$$\bar{\Delta} = \int_0^1 \Delta_{u_0, u_1, \lambda^*}^*(z) dz = u_1 - u_0,$$

and so the value of the objective function in (\mathcal{R}_A) is

$$\int_0^1 \bar{\Delta} h(p) dp + (1 - \beta)u_1 = (u_1 - u_0)H(1) + (1 - \beta)u_1 = (1 - \beta) [\pi_b u_1 + \pi_s u_0].$$

If $\beta < 1$, optimizing over u_0 and u_1 yields $u_0^* = u_1^* = 1$ (that, the analyst is always retained) and a payoff of $(1 - \beta) > 0$. Conversely, if $\beta > 1$, optimizing over u_0 and u_1 yields $u_0^* = u_1^* = 0$ (that is, the analyst is always fired) and a payoff of 0. Lastly, if $\beta = 1$, then the payoff is 0 for any values of u_0 and u_1 . ■

To summarize, when $H(p) \geq pH(1)$ for all $p \in [0, 1]$, the optimal mechanism never screens the analyst: it always retains her if $\beta < 1$ (and the odds and costs of retaining type l are low), it always fires her if $\beta > 1$ (and the odds and costs of retaining type l are high), and it is indifferent if $\beta = 1$. (In this latter case, $H(p) \geq pH(1) = 0$ for all p implies that F_l is a mean-preserving spread of F_h and so type l is better-informed than type h . This informational advantage implies type l cannot be screened out.)

Henceforth, we will therefore assume that

$$H(p) < pH(1) \text{ for some } p \in (0, 1),$$

and so effective screening might be possible.

LEMMA A.3. *Suppose $(\Delta_{u_0^*, u_1^*, \lambda^*}^*, u_0^*, u_1^*)$ solves problem (\mathcal{R}_A) and that $u_0^* + u_1^* > 0$. Then*

- (i) $\max\{u_0^*, u_1^*\} = 1$;
- (ii) $\lambda^* = g\left(\frac{u_0^*}{u_0^* + u_1^*}\right)$; and

(iii) the value of the objective is $-G\left(\frac{u_0^*}{u_0^*+u_1^*}\right)(u_0^*+u_1^*)+(1-\beta)\pi_b u_1^*$.

PROOF. Consider a solution $(\Delta_{u_0^*,u_1^*,\lambda^*}^*, u_0^*, u_1^*)$ to (\mathcal{R}_A) in which $u_0^*+u_1^* > 0$.

Suppose first that $\lambda^* > g(1) \geq g(p)$ for all $p \in [0, 1]$. Pointwise maximization and condition (OPT) imply that $\Delta_{u_0^*,u_1^*,\lambda^*}^*(p) = -u_0^*$ for all $p \in [0, 1]$. In addition, note that condition (INT) implies that

$$\int_0^1 (-u_0^*)dz = u_1^* - u_0^*, \text{ and so } u_0^* > u_1^* = 0.$$

Thus, the value of the Lagrangian is

$$\int_0^1 -u_0^*h(p)dp + (1-\beta)u_1^* = -u_0^*H(1) = (1-\beta)\pi_s u_0^*.$$

If $\beta > 1$, this value is negative—but the principal can guarantee a payoff of zero by always firing the agent. If $\beta < 1$, then this value is positive, but it is less than the payoff $(1-\beta)$ the principal can achieve by always retaining the agent. And finally, if $\beta = 1$, this value is zero—but it is easy to see that there are feasible candidate solutions with strictly positive objective values.¹⁶ Thus, for any β we have a contradiction of optimality, and so we must have $\lambda^* \leq g(1)$.

Similarly, suppose that $\lambda^* < g(0) \leq g(p)$ for all $p \in [0, 1]$. Pointwise maximization and condition (OPT) imply that $\Delta_{u_0^*,u_1^*,\lambda^*}^*(p) = u_1^*$ for all $p \in [0, 1]$. In addition, note that condition (INT) implies that

$$\int_0^1 (u_1^*)dz = u_1^* - u_0^*, \text{ and so } u_1^* > u_0^* = 0.$$

Thus, the value of the Lagrangian is

$$\int_0^1 u_1^*h(p)dp + (1-\beta)u_1^* = (\beta-1)(1-\pi_b)u_1^* + (1-\beta)u_1^* = (1-\beta)\pi_b u_1^*.$$

This again contradicts the supposed optimality. If $\beta > 1$, this value is negative—but the principal can guarantee a payoff of zero by always firing the agent. If $\beta < 1$, then this value is positive, but it is less than the payoff $(1-\beta)$ the principal can achieve by always retaining the agent. And finally, if $\beta = 1$, this value is zero and thus worse than the feasible candidate solutions (akin to the retention rule in footnote 16) with strictly positive objective values.

Therefore, we must have $\lambda^* \in [g(0), g(1)]$. Let

$$p_0 := \min\{p \mid g(p) = \lambda^*\} \text{ and } p_1 := \max\{p \mid g(p) = \lambda^*\}.$$

Pointwise maximization and condition (OPT) imply that $\Delta_{u_0^*,u_1^*,\lambda^*}^*(p) = -u_0^*$ for any $p < p_0$, and that

¹⁶Fix any $p \in (0, 1)$ with $H(p) < 0$. If $p < \frac{1}{2}$, the retention rule given by $\Delta(z) = -1$ below p and $\Delta(z) = \frac{p}{1-p}$ above is feasible, incentive compatible, and yields a payoff $-\frac{H(p)}{1-p} > 0$. If $p \geq \frac{1}{2}$, the retention rule given by $\Delta(z) = -\frac{1-p}{p}$ below p and $\Delta(z) = 1$ above is feasible, incentive compatible, and yields a payoff $-\frac{H(p)}{p} > 0$.

$\Delta_{u_0^*, u_1^*, \lambda^*}^*(p) = u_1^*$ for any $p > p_1$. The integral constraint (INT) then implies that

$$\begin{aligned} u_1^* - u_0^* &= \int_0^{p_0} (-u_0^*) dz + \int_{p_0}^{p_1} \Delta_{u_0^*, u_1^*, \lambda^*}^*(z) dz + \int_{p_1}^1 (u_1^*) dz \\ &\leq \int_0^{p_0} (-u_0^*) dz + \int_{p_0}^1 (u_1^*) dz = -p_0 u_0^* + (1 - p_0) u_1^*. \end{aligned}$$

Consequently, we must have $p_0 \leq \frac{u_0^*}{u_0^* + u_1^*}$. Likewise, we must also have $p_1 \geq \frac{u_0^*}{u_0^* + u_1^*}$, since

$$\begin{aligned} u_1^* - u_0^* &= \int_0^{p_0} (-u_0^*) dz + \int_{p_0}^{p_1} \Delta_{u_0^*, u_1^*, \lambda^*}^*(z) dz + \int_{p_1}^1 (u_1^*) dz \\ &\geq \int_0^{p_1} (-u_0^*) dz + \int_{p_1}^1 (u_1^*) dz = -p_1 u_0^* + (1 - p_1) u_1^*. \end{aligned}$$

Since g is continuous and monotone, we can conclude that $\lambda^* = g(\hat{p})$ in any solution, where $\hat{p} := \frac{u_0^*}{u_0^* + u_1^*}$; this also implies that $G(\cdot)$ is linear on $[p_0, p_1]$.

The integral constraint (INT) then implies that

$$\int_{p_0}^{p_1} \Delta_{u_0^*, u_1^*, \lambda^*}^*(z) dz = u_1^* - u_0^* - \int_0^{p_0} (-u_0^*) dz - \int_{p_1}^1 (u_1^*) dz = p_1 u_1^* - (1 - p_0) u_0^*.$$

Applying [Toikka \(2011\)](#)'s Theorem 4.4, the value of the Lagrangian $\mathcal{L}(\Delta_{u_0^*, u_1^*, \lambda^*}^* | u_0^*, u_1^*, \lambda^*)$ evaluated at this solution is equal to the value of its "ironed" counterpart where we replace h with g ; thus,

$$\begin{aligned} &\mathcal{L}(\Delta_{u_0^*, u_1^*, \lambda^*}^* | u_0^*, u_1^*, \lambda^*) \\ &= \int_0^{p_0} (-u_0^*) g(z) dz + \int_{p_0}^{p_1} \Delta_{u_0^*, u_1^*, \lambda^*}^*(z) g(z) dz + \int_{p_1}^1 (u_1^*) g(z) dz + (1 - \beta) u_1^* \\ &= -G(p_0) u_0^* + g(\hat{p})(p_1 u_1^* - (1 - p_0) u_0^*) + (G(1) - G(p_1)) u_1^* + (1 - \beta) u_1^* \\ &= -(G(\hat{p}) + g(\hat{p})(p_0 - \hat{p})) u_0^* + g(\hat{p})(p_1 u_1^* - (1 - p_0) u_0^*) \\ &\quad - (G(\hat{p}) + g(\hat{p})(p_1 - \hat{p})) u_1^* + (1 - \beta) \pi_b u_1^* \\ &= -G(\hat{p})(u_0^* + u_1^*) + g(\hat{p})(\hat{p} u_1^* - (1 - \hat{p}) u_0^*) + (1 - \beta) \pi_b u_1^* \\ &= -G(\hat{p})(u_0^* + u_1^*) + (1 - \beta) \pi_b u_1^*. \end{aligned}$$

Notice that in the first line, there is no term corresponding to the (INT) constraint due to complementary slackness; the second equality follows from the fact that $g(z) = \lambda^* = g(\hat{p})$ for all $z \in [p_0, p_1]$; the third equality follows by rewriting $G(p_0)$ and $G(p_1)$ using the linearity of G on $[p_0, p_1]$; and the remaining equalities follow by applying the definition of \hat{p} .

Finally, suppose that $\bar{u} := \max\{u_0^*, u_1^*\} < 1$. Then define $u'_0 := u_0^*/\bar{u}$ and $u'_1 := u_1^*/\bar{u}$, and define

$$\Delta'(p) := \frac{\Delta_{u_0^*, u_1^*, \lambda^*}^*(p)}{\bar{u}} \text{ for all } p \in [0, 1].$$

It is straightforward to note that Δ' , u'_0 , u'_1 and λ^* jointly satisfy the constraints (INT), (MON), and (FEAS) since Δ , u_0^* , and u_1^* were assumed to have done so (and have merely been scaled up by the same constant). Moreover, note that $\frac{u'_0}{u'_0+u'_1} = \hat{p}$, so that (OPT) is also satisfied; finally, the pooling property also continues to be satisfied since we have only made a scalar transformation. Thus, the value of the Lagrangian $\mathcal{L}(\Delta'|u'_0, u'_1, \lambda^*)$ evaluated at this new candidate solution is

$$\begin{aligned}\mathcal{L}(\Delta'|u'_0, u'_1, \lambda^*) &= \int_0^1 \Delta'(z)g(z)dz + (1 - \beta)u'_1 \\ &= \frac{\mathcal{L}(\Delta_{u_0^*, u_1^*, \lambda^*}^*|u_0^*, u_1^*, \lambda^*)}{\bar{u}} > \mathcal{L}(\Delta_{u_0^*, u_1^*, \lambda^*}^*|u_0^*, u_1^*, \lambda^*),\end{aligned}$$

contradicting the claimed optimality of the original solution. Thus, $\bar{u} := \max\{u_0^*, u_1^*\} = 1$. \blacksquare

Now consider the expression $-G\left(\frac{u_0}{u_0+u_1}\right)(u_0 + u_1) + (1 - \beta)\pi_b u_1$ from part (iii) of Lemma A.3, but evaluated at arbitrary (and not necessarily optimal) values $u_0, u_1 \in [0, 1]$. For any u_0 and u_1 such that $\max\{u_0, u_1\} = 1$, there is a unique corresponding $p = \frac{u_0}{u_0+u_1}$ (where $p \leq \frac{1}{2}$ corresponds to $u_0 \leq u_1 = 1$ and $p \geq \frac{1}{2}$ corresponds to $u_1 \leq u_0 = 1$). Thus, we can rewrite the expression above as a function $\xi : [0, 1] \rightarrow \mathbb{R}$ given by

$$\xi(p) := \begin{cases} -\frac{G(p)}{1-p} + (1 - \beta)\pi_b & \text{if } p \leq \frac{1}{2}, \\ -\frac{G(p)}{p} + (1 - \beta)\pi_b \frac{1-p}{p} & \text{if } p > \frac{1}{2}. \end{cases} \quad (\text{A.1})$$

The next lemma completes the proof of Theorem A.1 by showing that the maximizers of π correspond to solutions of the relaxed problem (\mathcal{R}_A), and moreover that these solutions are feasible in the original screening problem (\mathcal{P}_A). It also derives the qualitative properties of the solution as described in the theorem; indeed, the lemma is more detailed than the statement of theorem in that it characterizes the cutoffs and provides conditions under which each case corresponds to the optimal retention rule.

LEMMA A.4. *Suppose there exists $q \in [0, 1]$ such that $\xi(q) > 0$, and let $p^* := \max\{\text{argmax}_p\{\xi(p)\}\}$ and $p_* := \min\{\text{argmax}_p\{\xi(p)\}\}$.*

(i) *If $p^* < \frac{1}{2}$, then (\mathcal{P}_A) is solved by*

$$x_b^*(p) = \begin{cases} 0 & \text{if } p < p^*, \\ 1 & \text{if } p \geq p^*, \end{cases} \text{ and } x_s^*(p) = \begin{cases} \frac{p^*}{1-p^*} & \text{if } p < p^*, \\ 0 & \text{if } p \geq p^*. \end{cases} \quad (\text{A.2})$$

(ii) *If instead $p_* > \frac{1}{2}$, then (\mathcal{P}_A) is solved by*

$$x_b^*(p) = \begin{cases} 0 & \text{if } p \leq p_*, \\ \frac{1-p_*}{p_*} & \text{if } p > p_*, \end{cases} \text{ and } x_s^*(p) = \begin{cases} 1 & \text{if } p \leq p_*, \\ 0 & \text{if } p > p_*. \end{cases} \quad (\text{A.3})$$

(iii) Finally, if $p_* \leq \frac{1}{2} \leq p^*$, then $\frac{1}{2} \in \operatorname{argmax}_p \{\xi(p)\}$ and (\mathcal{P}_A) is solved by

$$x_b^*(p) = \begin{cases} 0 & \text{if } p < \bar{p}_*, \\ \frac{(1-\bar{p}_*)(2\bar{p}^*-1)}{\bar{p}^*-\bar{p}_*} & \text{if } p \in [\bar{p}_*, \bar{p}^*], \\ 1 & \text{if } p > \bar{p}^*, \end{cases} \text{ and } x_s^*(p) = \begin{cases} 1 & \text{if } p < \bar{p}_*, \\ \frac{\bar{p}^*(1-2\bar{p}_*)}{\bar{p}^*-\bar{p}_*} & \text{if } p \in [\bar{p}_*, \bar{p}^*], \\ 0 & \text{if } p > \bar{p}^*, \end{cases} \quad (\text{A.4})$$

where $\bar{p}^* := \min\{p \geq 1/2 \mid H(p) = G(p)\}$, $\bar{p}_* := \max\{p \leq 1/2 \mid H(p) = G(p)\}$, and we define $x_b^*(1/2) = 0$, $x_s^*(1/2) = 1$ when $\bar{p}^* = \bar{p}_* = 1/2$.

PROOF. We begin by considering the case where $p^* < \frac{1}{2}$. Suppose first that $G(p^*) < H(p^*)$, so G is linear on the interval $[p^* - \varepsilon, p^* + \varepsilon]$ for $\varepsilon > 0$ sufficiently small (in particular, we can choose ε such that $\varepsilon < \frac{1}{2} - p^*$). Then

$$\begin{aligned} \xi(p^*) - \xi(p^* + \varepsilon) &= \frac{G(p^* + \varepsilon)}{1 - p^* - \varepsilon} - \frac{G(p^*)}{1 - p^*} \\ &= \frac{(G(p^*) + \varepsilon g(p^*))(1 - p^*) - G(p^*)(1 - p^* - \varepsilon)}{(1 - p^*)(1 - p^* - \varepsilon)} \\ &= \varepsilon \frac{G(p^*) + (1 - p^*)g(p^*)}{(1 - p^*)(1 - p^* - \varepsilon)}. \end{aligned}$$

But note that since $p^* \in \operatorname{argmax}_p \{\xi(p)\}$ is interior, the first order condition $\xi'(p^*) = 0$ is satisfied; that is, since $p^* < \frac{1}{2}$,

$$\xi'(p^*) = G(p^*) + (1 - p^*)g(p^*) = 0.$$

This implies that $\xi(p^* + \varepsilon) = \xi(p^*)$, contradicting the definition of p^* as the largest maximizer of π . Thus, we must have $G(p^*) = H(p^*)$.

Now let $\lambda^* := g(p^*)$, $u_0^* := \frac{p^*}{1-p^*}$, $u_1^* := 1$, and $\Delta_{u_0^*, u_1^*, \lambda^*}^*(p) := x_b^*(p) - x_s^*(p)$, where x_b^* and x_s^* are as defined in (A.2). It is straightforward to see that this proposed solution satisfies constraints (FEAS) and (MON). In addition, note that

$$\int_0^1 \Delta_{u_0^*, u_1^*, \lambda^*}^*(z) dz = \int_0^{p^*} -\frac{p^*}{1-p^*} dz + \int_{p^*}^1 1 dz = \frac{-(p^*)^2 + (1-p^*)^2}{1-p^*} = \frac{1-2p^*}{1-p^*} = u_1^* - u_0^*,$$

implying that (INT) is satisfied. Finally, note that this proposed solution satisfies the ‘‘pooling property’’ since it is constant both above and below p^* (where we have shown that G and H coincide). Thus, as p^* is a maximizer of π , it solves problem (\mathcal{R}_A) .

Finally, note that we can write

$$U(p, p) = \begin{cases} \frac{p^*(1-p)}{1-p^*} & \text{if } p < p^*, \\ p & \text{if } p \geq p^*. \end{cases}$$

It is trivial to see at this point that the ‘‘full’’ envelope condition (ENV) is satisfied, and so the proposed solution is incentive compatible. Since it solves the relaxed problem (\mathcal{R}_A) , it must therefore also solve the original problem (\mathcal{P}_A) .

Now consider the case where $p_* > \frac{1}{2}$. Suppose first that $G(p_*) < H(p_*)$, so G is linear on some interval $[p_* - \varepsilon, p_* + \varepsilon]$ for $\varepsilon > 0$ sufficiently small (in particular, we can choose ε such that $\varepsilon < p_* - \frac{1}{2}$). Then

$$\begin{aligned}\xi(p^*) - \xi(p^* - \varepsilon) &= \frac{G(p_* - \varepsilon)}{p_* - \varepsilon} - \frac{G(p_*)}{p_*} + (1 - \beta)\pi_b \left(\frac{1 - p_*}{p_*} - \frac{1 - p_* + \varepsilon}{p_* - \varepsilon} \right) \\ &= \frac{(G(p_*) - \varepsilon g(p_*))p_* - G(p_*)(p_* - \varepsilon)}{p_*(p_* - \varepsilon)} \\ &\quad + (1 - \beta)\pi_b \frac{(1 - p_*)(p_* - \varepsilon) - (1 - p_* + \varepsilon)p_*}{p_*(p_* - \varepsilon)} \\ &= \varepsilon \frac{G(p_*) - p_*g(p_*) - (1 - \beta)\pi_b}{p_*(p_* - \varepsilon)}.\end{aligned}$$

But note that since $p_* \in \operatorname{argmax}_p \{\xi(p)\}$ is interior, the first order condition $\xi'(p_*) = 0$ is satisfied; that is, since $p_* > \frac{1}{2}$,

$$\xi'(p_*) = G(p_*) - p_*g(p_*) - (1 - \beta)\pi_b = 0.$$

This implies that $\xi(p_* - \varepsilon) = \xi(p_*)$, contradicting the definition of p_* as the smallest maximizer of $\xi(\cdot)$. Thus, we must have $G(p^*) = H(p^*)$.

Now let $\lambda^* := g(p_*)$, $u_0^* := 1$, $u_1^* := \frac{1-p_*}{p_*}$, and $\Delta_{u_0^*, u_1^*, \lambda^*}^*(p) := x_b(p) - x_s(p)$, where x_b and x_s are as defined in (A.3). It is straightforward to see that this proposed solution satisfies constraints (FEAS) and (MON). In addition, note that

$$\int_0^1 \Delta_{u_0^*, u_1^*, \lambda^*}^*(z) dz = \int_0^{p_*} -1 dz + \int_{p_*}^1 \frac{1 - p_*}{p_*} dz = \frac{-(p_*)^2 + (1 - p_*)^2}{p_*} = \frac{1 - 2p_*}{p_*} = u_1^* - u_0^*,$$

implying that (INT) is satisfied. Finally, note that this proposed solution satisfies the ‘‘pooling property’’ since it is constant both above and below p_* (where we have shown that G and H coincide). Thus, as p_* is a maximizer of π , it solves problem (\mathcal{R}_A).

Finally, note that we can write

$$U(p, p) = \begin{cases} 1 - p & \text{if } p < p_*, \\ \frac{p(1-p_*)}{p_*} & \text{if } p \geq p_*. \end{cases}$$

As before, the ‘‘full’’ envelope condition (ENV) is satisfied, implying that the proposed solution is incentive compatible. As it solves the relaxed problem (\mathcal{R}_A), it must also solve the original problem (\mathcal{P}_A).

Finally, we turn to the case where $p_* \leq 1/2 \leq p^*$. Suppose first that $1/2$ is *not* a maximizer of π , so both inequalities are strict. This of course implies that the first-order conditions $\xi'(p_*) = 0$ and $\xi'(p^*) = 0$ hold, implying that

$$g(p_*) = -\frac{G(p_*)}{1 - p_*} \text{ and } g(p^*) = \frac{G(p^*) - (1 - \beta)\pi_b}{p^*}.$$

Note, however, that we can write

$$\begin{aligned}\xi(p_*) + \xi(p^*) &= \left(-\frac{G(p_*)}{1-p_*} + (1-\beta)\pi_b \right) + \left(-\frac{G(p^*)}{p^*} + (1-\beta)\pi_b \frac{1-p^*}{p^*} \right) \\ &= g(p_*) - g(p^*) \leq 0,\end{aligned}$$

where the final inequality follows from the convexity of G . But of course this implies that $\max_p \{\xi(p)\} \leq 0$, contradicting the assumption that $\xi(q) > 0$ for some $q \in [0, 1]$. Thus, $1/2 \in \operatorname{argmax}_p \{\xi(p)\}$.

Now let $\lambda^* := g(1/2)$, $u_0^* := 1$, $u_1^* := 1$, and $\Delta^*(p) := x_b(p) - x_s(p)$, where x_b^* and x_s^* are as defined in (A.4). (Note that if $G(1/2) = H(1/2)$, then $\bar{p}^* = \bar{p}_* = 1/2$, and so by convention we can take $x_b^*(1/2) = 0$, $x_s^*(1/2) = 1$.) It is straightforward to see that this proposed solution satisfies constraints (FEAS) and (MON). In addition, note that

$$\begin{aligned}\int_0^1 \Delta_{u_0^*, u_1^*, \lambda^*}^*(z) dz &= \int_0^{\bar{p}_*} -1 dz + \int_{\bar{p}_*}^{\bar{p}^*} \frac{\bar{p}_* + \bar{p}^* - 1}{\bar{p}^* - \bar{p}_*} dz + \int_{\bar{p}^*}^1 1 dz \\ &= -\bar{p}_* + (\bar{p}_* + \bar{p}^* - 1) + (1 - \bar{p}^*) = 0 = u_1^* - u_0^*,\end{aligned}$$

implying that (INT) is satisfied. Finally, note that since G and H coincide (by definition) at both \bar{p}^* and \bar{p}_* and this proposed solution is constant on $[0, \bar{p}_*]$, $[\bar{p}_*, \bar{p}^*]$, and $(\bar{p}^*, 1]$, then it satisfies the pooling property. Thus, as $1/2$ is a maximizer of π , it solves problem (\mathcal{R}_A).

Finally, note that we can write

$$U(p, p) = \begin{cases} 1 - p & \text{if } p < \bar{p}_*, \\ \frac{(\bar{p}^* + \bar{p}_* - 1)p - (1 - 2\bar{p}_*)\bar{p}^*}{\bar{p}^* - \bar{p}_*} & \text{if } p \in [\bar{p}_*, \bar{p}^*], \\ p & \text{if } p > \bar{p}^*. \end{cases}$$

The “full” envelope condition (ENV) is clearly satisfied, implying that this proposed solution is indeed incentive compatible. Moreover, it must also be a solution to the original problem (\mathcal{P}_A) as it solves the relaxed problem (\mathcal{R}_A). \blacksquare

A.3. CHARACTERIZING OPTIMAL PERSUASION

In order to characterize the solution to the principal’s persuasion problem (FC_C), it is helpful to first define the thresholds

$$p^* := \min \left\{ p \in [0, 1] \left| \int_p^1 z \frac{\mu_h f_h(z) + \mu_l f_l(z)}{\mu_h(1 - F_h(p)) + \mu_l(1 - F_l(p))} dz \geq \bar{q} \right. \right\} \quad (\text{A.5})$$

and

$$p_* := \max \left\{ p \in [0, 1] \left| \int_0^p z \frac{\mu_h f_h(z) + \mu_l f_l(z)}{\mu_h F_h(p) + \mu_l F_l(p)} dz \leq 1 - \bar{q} \right. \right\}. \quad (\text{A.6})$$

We can interpret p^* as the lowest analyst posterior such that, when pooled with all $p > p^*$ regardless of the analyst’s ability, leads the client to buy; likewise, p_* is the largest analyst posterior such that, when pooled with all $p < p_*$ regardless of the analyst’s ability, leads the client to sell.

THEOREM A.2. *The principal's full-commitment persuasion problem (FC_C) is solved by the deterministic action recommendation rule*

$$A^*(h, p) = A^*(l, p) = \begin{cases} b & \text{if } p > \tilde{q}^*, \\ s & \text{if } p < \tilde{q}_*, \\ \varphi & \text{otherwise,} \end{cases}$$

where $\tilde{q}^* = p^*$ and $\tilde{q}_* = p_*$ if $p^* \geq p_*$, and $\tilde{q}^* = \tilde{q}_* = \bar{p}$ for arbitrary $\bar{p} \in [p^*, p_*]$ if $p^* < p_*$.

PROOF. We begin by rewriting the bank's persuasion problem (FC_C) in terms of the probabilities $B(\theta, p)$ and $S(\theta, p)$ with which the bank recommends the client chooses actions b and s , respectively. (It is without loss to consider only a single message recommending either b or s ; if multiple such messages induced the desired action, “merging” them would induce it as well.) Rescaling the objective by $1/\kappa > 0$, this problem can be written as

$$\begin{aligned} & \max_{B, S} \left\{ \sum_{\theta \in \{h, l\}} \mu_\theta \int_0^1 (B(\theta, p) + S(\theta, p)) dF_\theta(p) \right\} \\ & \text{subject to } \sum_{\theta \in \{h, l\}} \mu_\theta \int_0^1 (p - \bar{q}) B(\theta, p) dF_\theta(p) \geq 0, \\ & \sum_{\theta \in \{h, l\}} \mu_\theta \int_0^1 (1 - \bar{q} - p) S(\theta, p) dF_\theta(p) \geq 0, \\ & \mu_\theta f_\theta(p) B(\theta, p) \geq 0, \mu_\theta f_\theta(p) S(\theta, p) \geq 0, \text{ and} \\ & \mu_\theta f_\theta(p) (1 - B(\theta, p) - S(\theta, p)) \geq 0 \text{ for all } \theta \in \{h, l\} \text{ and } p \in [0, 1]. \end{aligned}$$

Notice that the first two constraints are simply the “obedience” constraints (IC-C) rewritten using Bayes’ rule (and multiplying through each by its denominator). The remaining constraints ensure feasibility of any potential recommendations, and have been normalized by $\mu_\theta f_\theta(p) \geq 0$ for all θ and p . (Note that if $\mu_\theta f_\theta(p) = 0$ for some type (θ, p) , that type’s contribution to the expected payoff is zero, and so the normalization is without loss.)

This is a well-defined linear optimization problem, and can therefore be solved using Lagrangian methods. Thus, letting $\lambda_B, \lambda_S, \eta_B(\theta, p), \eta_S(\theta, p)$, and $\eta_\varphi(\theta, p)$ denote the (nonnegative) multipliers on each of the constraints above, a solution is characterized by the first-order conditions

$$\begin{aligned} & \mu_\theta f_\theta(p) (1 + \lambda_B(p - \bar{q}) + \eta_B(\theta, p) - \eta_\varphi(\theta, p)) = 0 \text{ and} \\ & \mu_\theta f_\theta(p) (1 + \lambda_S(1 - \bar{q} - p) + \eta_S(\theta, p) - \eta_\varphi(\theta, p)) = 0, \end{aligned}$$

as well as the usual complementary slackness conditions. (See Theorems 8.3.1 and 8.4.1 in Luenberger (1969) establishing necessity and sufficiency, respectively.) With this in mind, let $B^*(\theta, p)$ and $S^*(\theta, p)$ correspond to the proposed solution in Theorem A.2; that is,

$$B^*(\theta, p) := \begin{cases} 0 & \text{if } p \leq \tilde{q}^*, \\ 1 & \text{if } p > \tilde{q}^*, \end{cases} \text{ and } S^*(\theta, p) := \begin{cases} 1 & \text{if } p < \tilde{q}_*, \\ 0 & \text{if } p \geq \tilde{q}_*, \end{cases}$$

where we let $\tilde{q}^* = p^*$ and $\tilde{q}_* = p_*$ if $p^* \geq p_*$, and $\tilde{q}^* = \tilde{q}_* = \bar{p}$ for arbitrary $\bar{p} \in [p^*, p_*]$ if $p^* < p_*$, with p^* and p_* as defined in (A.5) and (A.6).

Notice that if $p^* < p_*$, then the proposed solution (where $\tilde{q}^* = \tilde{q}_* = \bar{p}$ for an arbitrary $\bar{p} \in [p^*, p_*]$) yields client posterior beliefs that are always either above \bar{q} or below $1 - \bar{q}$, and so trade is induced with probability 1. By observation, this trivially achieves the optimum and solves (FC_C).

If, on the other hand, $\tilde{q}^* = p^* \geq p_* = \tilde{q}_*$, we can further define

$$\lambda_B^* := \frac{1}{\bar{q} - \tilde{q}^*} \text{ and } \lambda_S^* := \frac{1}{\tilde{q}_* - (1 - \bar{q})}$$

as the multipliers on the rewritten ‘‘obedience’’ constraints, and

$$\eta_B^*(\theta, p) := \begin{cases} \frac{\tilde{q}^* - p}{\bar{q} - \tilde{q}^*} + \frac{\tilde{q}_* - p}{\tilde{q}_* - (1 - \bar{q})} & \text{if } p < \tilde{q}_*, \\ \frac{\tilde{q}^* - p}{\bar{q} - \tilde{q}^*} & \text{if } p \in [\tilde{q}_*, \tilde{q}^*], \\ 0 & \text{if } p > \tilde{q}^*, \end{cases}$$

$$\eta_S^*(\theta, p) := \begin{cases} 0 & \text{if } p < \tilde{q}_*, \\ \frac{p - \tilde{q}_*}{\tilde{q}_* - (1 - \bar{q})} & \text{if } p \in [\tilde{q}_*, \tilde{q}^*], \\ \frac{p - \tilde{q}^*}{\bar{q} - \tilde{q}^*} + \frac{p - \tilde{q}_*}{\tilde{q}_* - (1 - \bar{q})} & \text{if } p > \tilde{q}^*, \end{cases}$$

$$\text{and } \eta_\varphi^*(\theta, p) := \begin{cases} \frac{\tilde{q}_* - p}{\tilde{q}_* - (1 - \bar{q})} & \text{if } p < \tilde{q}_*, \\ 0 & \text{if } p \in [\tilde{q}_*, \tilde{q}^*], \\ \frac{p - \tilde{q}^*}{\bar{q} - \tilde{q}^*} & \text{if } p > \tilde{q}^*, \end{cases}$$

as the multipliers on the feasibility constraints.

With these in hand, for all $\theta \in \{h, l\}$ and $p \in [0, 1]$, it is straightforward to verify that

$$1 + \lambda_B^*(p - \bar{q}) + \eta_B^*(\theta, p) - \eta_\varphi^*(\theta, p) = 0 \text{ and } 1 + \lambda_S^*(1 - \bar{q} - p) + \eta_S^*(\theta, p) - \eta_\varphi^*(\theta, p) = 0.$$

Therefore, the first-order conditions from above are satisfied, as are the complementary slackness conditions. Thus, the proposed solution is indeed a bank-optimal trading recommendation. ■

B. PROOF OF THEOREM 2

The result follows from the more general Theorem 4, making use of Section 3.1’s embedding of the analyst application into the general model.

The only caveat is that the type space $\Theta \times [0, 1]$ is infinite, whereas Theorem 4 is for a setting with finite types. Note, however, that the optimal direct retention rule X^* that solves (FC) has a finite range $X^*(\Theta \times [0, 1])$; its inverse therefore defines a finite partition Θ_X of $\Theta \times [0, 1]$ such that X^* is constant on each element of Θ_X . Likewise, the optimal direct action recommendation rule A^* that solves (FC) has a finite range $A^*(\Theta \times [0, 1])$; its inverse therefore defines a finite partition Θ_A of $\Theta \times [0, 1]$ such that A^* is constant on each element of Θ_A .

Let $\Lambda := \Theta_X \wedge \Theta_A$ be the meet of the two partitions Θ_X and Θ_A . By construction, Λ is finite, and

moreover X^* and A^* are both constant on each element of Λ —this partition of the type space $\Theta \times [0, 1]$ conveys all the information about the analyst’s type necessary to implement the optimal direct mechanism (X^*, A^*) . Treating Λ as the underlying (finite) type space, [Theorem 4](#) applies immediately. ■

C. PROOF OF THEOREM 3

Consider a bank-optimal equilibrium with public contracting in which the bank chooses a finite message space $\overline{\mathcal{M}}$ along with a retention rule \overline{x} ; the analyst chooses a reporting strategy $\overline{\sigma}_\theta$; and the client chooses an action strategy $\overline{\alpha}$.

For each client action $a = b, s$, we define

$$\mathcal{M}^a := \left\{ m \in \overline{\mathcal{M}} \left| \sum_{\theta \in \Theta} \mu_\theta \int_0^1 \overline{\sigma}_\theta(\overline{x}, p)(m) dF_\theta(p) > 0 \text{ and } \overline{\alpha}(m) = a \right. \right\}$$

to be the set of all messages reported by the analyst with positive probability such that the client chooses action a as a best response. Moreover, let

$$\begin{aligned} \overline{p}^a &:= \sup \left\{ p \left| \sum_{\theta \in \Theta} \mu_\theta \overline{\sigma}_\theta(\overline{x}, p)(m) > 0 \text{ for some } m \in \mathcal{M}^a \right. \right\} \text{ and} \\ \underline{p}^a &:= \inf \left\{ p \left| \sum_{\theta \in \Theta} \mu_\theta \overline{\sigma}_\theta(\overline{x}, p)(m) > 0 \text{ for some } m \in \mathcal{M}^a \right. \right\} \end{aligned}$$

be the largest and smallest analyst information types that announce a message in \mathcal{M}^a with positive probability.

The proof proceeds in two steps.

STEP 1: Consider a bank-optimal equilibrium in which the bank publicly chooses contractual terms $(\overline{\mathcal{M}}, \overline{x})$, the analyst best responds with strategy $\overline{\sigma}_\theta$ and the client’s expected payoff in this equilibrium is zero. Then, there exist cutoffs $\underline{p}^b \geq \tilde{q}^*$ and $\overline{p}^s \leq \tilde{q}^*$ such that, for both $\theta \in \Theta$, the analyst’s strategy satisfies $\overline{\sigma}_\theta(\overline{x}, p) \in \Delta(\mathcal{M}^b)$ and $p \in [\underline{p}^b, 1]$ and $\overline{\sigma}_\theta(\overline{x}, p) \in \Delta(\mathcal{M}^s)$ for $p \in [0, \overline{p}^s]$.

We prove the statement in this step via a series of lemmata.

LEMMA C.1. *For any $\overline{m}, \underline{m} \in \mathcal{M}^a$, $a \in \{b, s\}$, we have $\overline{x}(\overline{m}, b) = \overline{x}(\underline{m}, b)$ and $\overline{x}(\overline{m}, s) = \overline{x}(\underline{m}, s)$. Moreover, for $a \in \{b, s\}$ any $m \in \mathcal{M}^a$ is a best response for the analyst for all beliefs $p \in [\underline{p}^a, \overline{p}^a]$.*

PROOF. Consider any $\overline{m} \in \mathcal{M}^b$. Since the client’s expected payoff is zero, they must be indifferent between actions b and φ upon observing message \overline{m} , and so their posterior on observing \overline{m} is exactly \overline{q} . Because F_θ has no atoms and \overline{m} is reported by the analyst with positive probability, this implies that there exists \overline{p} and \underline{p} , with $\overline{p} > \overline{q} > \underline{p}$, such that $\sum_{\theta \in \Theta} \mu_\theta \overline{\sigma}_\theta(\overline{x}, \overline{p})(\overline{m}) > 0$ and $\sum_{\theta \in \Theta} \mu_\theta \overline{\sigma}_\theta(\overline{x}, \underline{p})(\overline{m}) > 0$.

Now fix any $p' \in (\underline{p}, \overline{p})$, and suppose that there exists some message $m' \in \overline{\mathcal{M}}$ that is strictly preferred by the analyst with information p' over message \overline{m} , implying that

$$p' \overline{x}(\overline{m}, b) + (1 - p') \overline{x}(\overline{m}, s) < p' \overline{x}(m', b) + (1 - p') \overline{x}(m', s). \quad (\text{C.1})$$

But since \bar{m} is an optimal message for information type \bar{p} , we have

$$\bar{p}\bar{x}(\bar{m}, b) + (1 - \bar{p})\bar{x}(\bar{m}, s) \geq \bar{p}x(m', b) + (1 - \bar{p})x(m', s).$$

Subtracting (C.1) from this inequality and dividing by $(\bar{p} - p') > 0$, we have

$$\bar{x}(\bar{m}, b) - \bar{x}(\bar{m}, s) > \bar{x}(m', b) - \bar{x}(m', s). \quad (\text{C.2})$$

Likewise, since \bar{m} is an optimal message for information type $\bar{\bar{p}}$, we have

$$\bar{\bar{p}}\bar{x}(\bar{m}, b) + (1 - \bar{\bar{p}})\bar{x}(\bar{m}, s) \geq \bar{\bar{p}}x(m', b) + (1 - \bar{\bar{p}})x(m', s).$$

Subtracting (C.1) from this inequality and dividing by $(\bar{\bar{p}} - p') < 0$, we have

$$\bar{x}(\bar{m}, b) - \bar{x}(\bar{m}, s) < \bar{x}(m', b) - \bar{x}(m', s). \quad (\text{C.3})$$

Since (C.2) and (C.3) cannot simultaneously hold, it must be the case that message \bar{m} is optimal for all $p' \in [\bar{\bar{p}}, \bar{p}]$.

Likewise, consider any other $\underline{m} \in \mathcal{M}^b$. Again, since the client's expected payoff is zero, they must be indifferent between actions b and φ upon observing message \underline{m} , and therefore their posterior on observing \underline{m} is exactly \bar{q} . This implies that there exists some $\underline{p} < \bar{q}$ with $\sum_{\theta \in \Theta} \mu_{\theta} \bar{\sigma}_{\theta}(\bar{x}, \underline{p})(\underline{m}) > 0$ and some other $\underline{\underline{p}} > \bar{q}$ with $\sum_{\theta \in \Theta} \mu_{\theta} \bar{\sigma}_{\theta}(\bar{x}, \underline{\underline{p}})(\underline{m}) > 0$. An identical argument to that above implies that message \underline{m} is optimal for all information types $p' \in [\underline{\underline{p}}, \underline{p}]$.

Note, however, that $\underline{\underline{p}} > \bar{q} > \bar{\bar{p}}$. Thus, information types $p \in (\bar{\bar{p}}, \underline{\underline{p}})$ must be indifferent between messages \bar{m} and \underline{m} . The indifference between \bar{m} and \underline{m} is possible throughout this interval if, and only if, $\bar{x}(\bar{m}, \omega) = \bar{x}(\underline{m}, \omega)$ for both $\omega \in \Omega$. Finally, the optimality of these messages for all $p \in [\underline{\underline{p}}^b, \bar{\bar{p}}^b]$ follows by choosing \bar{m} and \underline{m} so that $\bar{p} = \bar{\bar{p}}^b$ and $\underline{p} = \underline{\underline{p}}^b$.

A symmetric argument (where $1 - \bar{q}$ plays the role of \bar{q} in the above) applies to messages in \mathcal{M}^s that induce action s . ■

LEMMA C.2. *We must have $\max\{\bar{p}^b, \bar{p}^s\} = 1$ and $\min\{\underline{p}^b, \underline{p}^s\} = 0$.*

PROOF. Let $\bar{p} := \max\{\bar{p}^b, \bar{p}^s\}$ and first observe that $\bar{p} \geq \bar{p}^b \geq \bar{q}$. The latter inequality must hold as otherwise $\bar{\alpha}(m) = b$ cannot be a best response for $m \in \mathcal{M}^b$.

Suppose $\bar{p} < 1$. Let $m \in \bar{\mathcal{M}}$ be any message such that $\sum_{\theta \in \Theta} \mu_{\theta} \int_{\bar{p}}^1 \bar{\sigma}_{\theta}(\bar{x}, p)(m) dF_{\theta}(p) > 0$ or, in words, m is a message reported by a positive measure of types above \bar{p} . Clearly, since $\bar{p} \geq \bar{q}$, the client must best respond by choosing $\bar{\alpha}(m) = b$ which, in turn, implies $\bar{p}^b > \bar{p}$, a contradiction.

We can conclude that $\min\{\underline{p}^b, \underline{p}^s\} = 0$ using a symmetric argument. ■

LEMMA C.3. *We must have $\underline{p}^b \geq \bar{p}^s$.*

PROOF. Fix any $m^b \in \mathcal{M}^b$, and note that by Lemma C.1, m^b is optimal for all $p \in [\underline{p}^b, \bar{p}^b]$. Likewise, fix any $m^s \in \mathcal{M}^s$, and note that m^s is optimal for all $p \in [\underline{p}^s, \bar{p}^s]$.

Suppose, for purposes of contradiction, that $\underline{p}^b < \bar{p}^s$. Then both m^b and m^s are optimal for all $p \in [\underline{p}^b, \bar{p}^s]$; that is,

$$\begin{aligned}\underline{p}^b \bar{x}(m^b, b) + (1 - \underline{p}^b) \bar{x}(m^b, s) &= \underline{p}^b \bar{x}(m^s, b) + (1 - \underline{p}^b) \bar{x}(m^s, s) \text{ and} \\ \bar{p}^s \bar{x}(m^b, b) + (1 - \bar{p}^s) \bar{x}(m^b, s) &= \bar{p}^s \bar{x}(m^s, b) + (1 - \bar{p}^s) \bar{x}(m^s, s).\end{aligned}$$

Combining these two equations immediately implies that $\bar{x}(m^b, b) = \bar{x}(m^s, b)$ and $\bar{x}(m^b, s) = \bar{x}(m^s, s)$; therefore, the bank is (for screening purposes) pooling all information types in both $[\underline{p}^b, \bar{p}^b]$ and $[\underline{p}^s, \bar{p}^s]$ by offering them the same retention terms. Since, by assumption, these two intervals overlap, [Lemma C.2](#) implies that *all* analyst information types in $[0, 1]$ are offered the same retention terms.

But since all information types are indifferent between all messages (so the bank is not screening the analyst), the optimality of $(\bar{\mathcal{M}}, \bar{x})$ implies that the bank must be generating the highest possible trading commissions (obtained from the unique optimal action recommendation A^*). But this implies $\bar{p}^s = \tilde{q}_* < \tilde{q}^* = \underline{p}^b$ which provides the requisite contradiction. \blacksquare

LEMMA C.4. *We must have $\underline{p}^b \geq \tilde{q}^*$ and $\bar{p}^s \leq \tilde{q}_*$. Moreover, $\sum_{m \in \mathcal{M}^b} \bar{\sigma}_\theta(\bar{x}, p)(m) = 1$ for all $p \geq \underline{p}^b$ and $\sum_{m \in \mathcal{M}^s} \bar{\sigma}_\theta(\bar{x}, p)(m) = 1$ for all $p \leq \bar{p}^s$.*

PROOF. The probability with which the client chooses action b is given by

$$\sum_{m \in \mathcal{M}^b} \sum_{\theta \in \{h, l\}} \mu_\theta \int_0^1 \bar{\sigma}_\theta(\bar{x}, p)(m) dF_\theta(p) = \sum_{\theta \in \{h, l\}} \mu_\theta \int_{\underline{p}^b}^1 \sum_{m \in \mathcal{M}^b} \bar{\sigma}_\theta(\bar{x}, p)(m) dF_\theta(p).$$

Since $\bar{\sigma}_\theta(\bar{x}, p)(m) \leq 1$ for all $(\theta, p) \in \{h, l\} \times [0, 1]$ and all $m \in \bar{\mathcal{M}}$, there exists some $\bar{p} \geq \underline{p}^b$ such that

$$\sum_{\theta \in \{h, l\}} \mu_\theta \int_{\bar{p}}^1 1 dF_\theta(p) = \sum_{\theta \in \{h, l\}} \mu_\theta \int_{\underline{p}^b}^1 \sum_{m \in \mathcal{M}^b} \bar{\sigma}_\theta(\bar{x}, p)(m) dF_\theta(p).$$

Suppose that $\bar{p} > \underline{p}^b$; equivalently, suppose that there is a positive measure of analyst information types such that

$$0 < \sum_{\theta \in \{h, l\}} \mu_\theta \sum_{m \in \mathcal{M}^b} \bar{\sigma}_\theta(\bar{x}, p)(m) < 1.$$

We can then fix some $\bar{m} \in \mathcal{M}^b$ and consider the alternative mechanism (\mathcal{M}', x') , where $\mathcal{M}' := \bar{\mathcal{M}} \cup \{m', m''\}$ and x' is defined by

$$x'(m, \omega) := \begin{cases} \bar{x}(m, \omega) & \text{if } m \in \bar{\mathcal{M}}, \\ \bar{x}(\bar{m}, \omega) & \text{if } m \in \{m', m''\}. \end{cases}$$

Since [Lemma C.1](#) established that all types with $p \in [\underline{p}^b, 1]$ are indifferent between any $m \in \mathcal{M}^b$, it is clear that the analyst strategy given by

$$\sigma'_\theta(x', p)(m) := \bar{\sigma}_\theta(\bar{x}, p)(m) \text{ for all } m \in \bar{\mathcal{M}} \text{ and } p < \underline{p}^b, \text{ and}$$

$$\sigma'_\theta(x', p)(m) := \begin{cases} 1 & \text{if } m = m' \text{ and } p \geq \bar{p} - \varepsilon, \\ 1 & \text{if } m = m'' \text{ and } p \in [\underline{p}^b, \bar{p} - \varepsilon), \\ 0 & \text{otherwise} \end{cases}$$

is a best-response for any $\varepsilon \in [0, \bar{p} - \underline{p}^b]$.

When $\varepsilon = 0$, the client's posterior belief $q'(m')$ upon observing message m' (from reporting strategy σ'_θ) is strictly higher than their lowest posterior belief $\min_{m \in \mathcal{M}^b} \bar{q}(m)$ from any message in the set \mathcal{M}^b (from reporting strategy $\bar{\sigma}_\theta$). Therefore, for $\varepsilon > 0$ sufficiently small, $\alpha'(m') = b$ is a best response for the client to message m' (from reporting strategy σ'_θ). In other words, the mechanism (\mathcal{M}', x') and analyst reporting strategy σ'_θ induces the client to choose action b with strictly higher probability—without affecting the probability of action s or the bank's payoff from analyst retention. This contradicts the assumption that $(\bar{\mathcal{M}}, \bar{x})$ is bank-optimal and hence we must have $\bar{p} = \underline{p}^b$.

This in turn implies that $\underline{p}^b \geq \bar{q}^*$, as otherwise $\bar{\alpha}(m) = b$ would not be a best response for the client to messages $m \in \mathcal{M}^b$.

A symmetric argument immediately delivers the desired conclusions for messages inducing action s . ■

This last lemma completes the proof of Step 1.

STEP 2: There is a cutoff $\bar{\kappa} > 0$ such that, for all $\kappa < \bar{\kappa}$, the client achieves a positive payoff in every bank-optimal equilibrium with public contracting.

Consider the public communication mechanism (\mathcal{M}, x) where the message space $\mathcal{M} = \{b, \varphi, s\}$ has three messages, and a retention rule

$$x(m, \cdot) = X^*(\theta, p(m), \cdot), \text{ where } p(m) \begin{cases} > \tilde{p}^* & \text{if } m = b, \\ \in [\tilde{p}_*, \tilde{p}^*] & \text{if } m = \varphi, \\ < \tilde{p}_* & \text{if } m = s \end{cases}$$

that implements the optimal retention rule X^* . Given (\mathcal{M}, x) , it is easy to see that

$$\sigma_\theta(x, p)(m) = \begin{cases} 1 & \text{if } m = b \text{ and } p > \tilde{p}^*, \\ 1 & \text{if } m = \varphi \text{ and } p \in [\tilde{p}_*, \tilde{p}^*], \\ 1 & \text{if } m = s \text{ and } p < \tilde{p}_*, \\ 0 & \text{otherwise} \end{cases}$$

is a best response for the analyst.

Let Π_R^* be the bank's payoff from analyst retention given this reporting strategy. Note that, by construction, the bank cannot get a higher retention payoff from any public communication mechanism. Let the $\Pi_T \geq 0$ be the bank's payoff from trading commissions for some client best response α to the analyst's reporting strategy. Note that the bank's total payoff is given by

$$\Pi(\mathcal{M}, x, \sigma, \alpha) = \Pi_R^* + \kappa \Pi_T.$$

Now suppose, for the sake of contradiction, that for any $\bar{\kappa} > 0$, there exists a $\kappa \in (0, \bar{\kappa})$ such that there is a bank-optimal equilibrium in which the analyst's best response satisfies $\underline{p}^b = \tilde{q}^*$ and $\bar{p}^s = \tilde{q}_*$. Note that, from Step 1, only a bank-optimal equilibrium of this form results in a payoff of zero for the client. By assumption, this mechanism yields the same trading commission Π_T^* as that from the optimal action recommendation A^* . Moreover, there is an ε (that does not depend on κ) such that the payoff Π_R from analyst retention in this equilibrium satisfies $\Pi_R < \Pi_R^* - \varepsilon$. The latter inequality follows from the assumption that, either $0 < \tilde{p}_* < \tilde{q}_*$ or $\tilde{q}^* < \tilde{p}^* < 1$ and so the bank's payoff with public contracting is bounded away from the full-commitment payoff.

But there exists a $\bar{\kappa} > 0$ such that

$$\Pi_R^* > \Pi_R^* - \varepsilon + \kappa \Pi_T^*$$

for all $\kappa < \bar{\kappa}$. This in turn implies that, with public contracting and any $\kappa \in (0, \bar{\kappa})$, there cannot be a bank-optimal equilibrium with an analyst best response (satisfying $\underline{p}^b = \tilde{q}^*$ and $\bar{p}^s = \tilde{q}_*$) that induces the maximal trading commissions Π_T^* . This is because there would then be another equilibrium in which the bank offered (\mathcal{M}, x) (keeping off-path behavior of the analyst and client the same as the purported bank-optimal equilibrium) and get a strictly higher payoff $\Pi(\mathcal{M}, x, \sigma, \alpha)$.

This completes the proof of the theorem. ■

D. PROOF OF THEOREM 4

The proof is constructive. Fix an optimal direct mechanism $(X^*, A^*) \in (\Delta(\mathcal{T}) \times \Delta(\mathcal{A}))^\Theta$. By definition, it satisfies the incentive compatibility constraints (R-IC) and (A-IC).

Let $\text{supp}(A^*(\theta))$ denote the support of $A^*(\theta)$, 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 $\Theta := \{\theta_1, \dots, \theta_N\}$, where $N := |\Theta|$. We then define the message space $\mathcal{M} := \{m_{a,i}\}_{a \in \bar{A}, i=1, \dots, N}$, and for each $j = 1, \dots, N$, define the contract $x_j : \mathcal{M} \rightarrow \Delta(\mathcal{T})$ by

$$x_j(m_{a,i}) := X^*(\theta_k), \text{ where } k = 1 + ((i - 1) + (j - 1)) \bmod N.$$

The set of messages \mathcal{M} consists of N “publicly synonymous” copies of each action: for any $m_{a,i} \in \mathcal{M}$, the receiver interprets the message as a recommendation to take action a while the principal interprets the message as report of type θ_k , where k is the “cyclic” permutation defined above.

Now consider the following strategies for each player:

- the principal announces the message space \mathcal{M} and the set of public communication mechanisms $\mathbb{M} = \{(\mathcal{M}, x_j)\}_{j=1, \dots, N}$, randomizes uniformly over this set, and then privately communicates the realization $(\mathcal{M}, x_{\hat{j}})$ to the agent;
- in response, the agent of type $\theta_i \in \Theta$ publicly announces message $m_{a,k} \in \mathcal{M}$ with probability $A^*(\theta_i)[a]$, where $k = 1 + ((i - 1) + (\hat{j} - 1)) \bmod N$; and
- after observing the publicly realized message $m_{\hat{a},k}$, the receiver takes action \hat{a} .

Conversely, suppose the principal deviates and announces any other finite message space \mathcal{M}' along with a finite set of public communication mechanisms $\mathbb{M}' \subset (\mathcal{M}')^{\Delta(\mathcal{T})}$. We may arbitrarily choose any strate-

gies that constitute a sequential equilibrium for this 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 perfect Bayesian 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 $m_{\hat{a},k}$, the receiver's posterior belief that the state of the world is $\omega \in \Omega$ is, by Bayes' rule, given by

$$\frac{\pi(\omega) \sum_{\theta} \frac{1}{N} \mu(\theta|\omega) A^*(\theta)[\hat{a}]}{\sum_{\omega'} \pi(\omega') \sum_{\theta} \frac{1}{N} \mu(\theta|\omega') A^*(\theta)[\hat{a}]} = \frac{\pi(\omega) \sum_{\theta} \mu(\theta|\omega) A^*(\theta)[\hat{a}]}{\sum_{\omega'} \pi(\omega') \sum_{\theta} \mu(\theta|\omega') A^*(\theta)[\hat{a}]},$$

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 (R-IC), it must therefore be the case that taking action \hat{a} remains a best response in the public communication game.

Now consider the agent's message announcement decision, fixing an arbitrary type $\theta_i \in \Theta$ and a realization \hat{j} of the principal's randomization over public communication mechanisms $\{(\mathcal{M}, x_j)\}_{j=1,\dots,n}$. Since the image of $x_{\hat{j}}$ is the same as the image of X^* and the latter satisfies (A-IC), the agent's payoff is maximized by sending any message $m_{a,j} \in x_{\hat{j}}^{-1}(X^*(\theta_i))$. Given the definition of $x_{\hat{j}}$ (and the fact that it does not condition on the recommended action), this implies that all messages $m_{a,j}$ with $j = 1 + ((i - 1) + (\hat{j} - 1)) \bmod N$ are best responses for the agent. Moreover, since the agent is indifferent across all messages in $\{m_{a,j}\}_{a \in \bar{A}}$, mixing with the distribution $A^*(\theta_i)$ is optimal.

Finally, note that the principal is indifferent between offering the agent any public communication mechanism from the set $\mathbb{M} = \{(\mathcal{M}, x_j)\}_{j=1,\dots,N}$. 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.

We complete the proof by observing that the principal has no incentive to deviate since a perfect Bayesian equilibrium of any subgame (following the deviation) must deliver a lower payoff than that from full commitment. ■

REFERENCES

- 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 (2022): “Persuasion for the Long Run,” Unpublished manuscript, University of Oxford.
- BIZZOTTO, J., E. PEREZ-RICHET, AND A. VIGIER (2021): “Communication via Third Parties,” Unpublished manuscript, Sciences Po.
- BLACKWELL, D. (1953): “Equivalent Comparisons of Experiments,” *Annals of Mathematical Statistics*, 24(2), 265–272.
- 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.
- CRANE, A., AND K. CROTTY (2020): “How Skilled are Security Analysts?,” *Journal of Finance*, 75(3), 1629–1675.
- 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.
- ELLIOTT, G., AND A. TIMMERMANN (2016): “Forecasting in Economics and Finance,” *Annual Review of Economics*, 8(1), 81–110.
- INDERST, R., AND M. OTTAVIANI (2012): “Competition through Commissions and Kickbacks,” *American Economic Review*, 102(2), 780–809.
- IVANOV, M. (2022): “Perfect Robust Implementation by Private Information Design,” Unpublished manuscript, McMaster University.
- JACKSON, A. R. (2005): “Trade Generation, Reputation, and Sell-Side Analysts,” *Journal of Finance*, 60(2), 673–717.
- 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.
- 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*, pp. 863–894.
- LIPNOWSKI, E., D. RAVID, AND D. SHISHKIN (2022): “Persuasion via Weak Institutions,” *Journal of Political Economy*, 130(10), 2705–2730.
- LUENBERGER, D. G. (1969): *Optimization by Vector Space Methods*. John Wiley & Sons.
- 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, chap. 12, pp. 691–720. North-Holland.
- MATHEVET, L., D. PEARCE, AND E. STACCHETTI (2022): “Reputation for a Degree of Honesty,” Unpublished manuscript, European University Institute.
- MILGROM, P., AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70(2), 583–601.
- MIN, D. (2021): “Bayesian Persuasion under Partial Commitment,” *Economic Theory*, 72(3), 743–764.
- 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.
- TOIKKA, J. (2011): “Ironing without Control,” *Journal of Economic Theory*, 146(6), 2510–2526.