Delegated Cheap Talk: A Theory of Investment Banking*

Changyeop Lee† and Sungmin Park‡

March 31, 2024

Abstract

We study the role of investment banking as delegated cheap talk. In our model of initial public offering (IPO), two parties have conflicting interests: a seller wants to sell his firm, whereas a buyer wants to invest only in a good firm. All communication is cheap talk. We show that the seller can influence the buyer by contracting with an intermediary and delegating the communication. Any successful contract requires the intermediary to share the risk of loss with the buyer. A seller-optimal contract maximizes the intermediary's bias in the seller's favor while maintaining minimally sufficient alignment with the buyer's incentives.

Keywords: Investment banking, Initial public offering, Cheap talk, Costly information acquisition, Conflict of interest
JEL Codes: D83, G24

---

*We are grateful to Yaron Azrieli for his guidance and input throughout the paper's development. Our paper also benefited significantly from feedback by Sylvain Chassang, PJ Healy, Sejin Kang, Irfan Khan, Yonggyun Kim, Daehong Min, Junrok Park, Sumin Park, Jim Peck, René Stulz, Daisy Wang, Huanxing Yang, and Yuchi Yao. We are thankful for comments by participants of the OSU Theory/Experimental Reading Group, OSU Finance PhD Student Seminar, 2023 North American Summer Meeting of the Econometric Society, 34th Stony Brook International Conference on Game Theory, 2023 Asian School in Economic Theory, 2023 Midwest Economic Theory Conference, and 93rd Southern Economic Association Annual Meeting. We have no relevant or material financial interests that relate to the research described in this paper.

†Department of Economics, The Ohio State University. 1945 North High Street, Columbus, Ohio, United States 43210. Email: lee.8970@buckeyemail.osu.edu

‡Corresponding author. Department of Economics, The Ohio State University. 1945 North High Street, Columbus, Ohio, United States 43210. Email: park.2881@buckeyemail.osu.edu

1 Introduction

When Mark Zuckerberg announced his plan to sell the ownership shares of Facebook to the public in 2012, he chose Morgan Stanley as the company's investment bank, designating it as the lead underwriter that would help manage the company's initial public offering (IPO).[1] Among the underwriter's most crucial responsibilities were (a) *due diligence*—inspecting the company's business and finances, and (b) *roadshow*—meeting investors to convince them to buy its shares. For these services, Facebook's underwriters earned 1.1 percent of the total 16 billion US dollars raised. More typical, moderate-sized IPOs in the United States usually pay investment banks even more hefty fraction, of 7 percent (Chen and Ritter, 2000; Ritter, 2023). Conventional explanation behind such large compensation is that investment banks reduce the information asymmetry between the seller and the buyer through expertise (Baron and Holmström, 1980; Baron, 1982; Ramakrishnan and Thakor, 1984; Biais, Bossaerts, and Rochet, 2002) or reputation (Beaty and Ritter, 1986; Booth and Smith II, 1986; Carter and Manaster, 1990; Chemmanur and Fulghieri, 1994).

However, this view of investment banks as reputable experts seems at odds with what the general public thinks, what the popular media portrays, and what investment bankers say about themselves. For example, the protesters of the Occupy Wall Street movement in 2011 cried "We are the 99%," referring to investment banks as the undeservingly wealthy one percent that caused the 2008 financial crisis. Movies such as *Wall Street* (1987), *Margin Call* (2011), and *The Wolf of Wall Street* (2013) often depict investment bankers as sleazy, smooth-talking salesmen. In their memoir *Monkey Business: Swinging through the Wall Street Jungle*, former investment bankers Rolfe and Troob (2009) write:

> [Investment bankers] only want to say good things. The better they can make the company sound, the easier it will be for them to sell the securities. The easier it is for them to sell the securities, the more certain they'll be that the clients will be happy. That means fees. Fees are important.

[1] We focus on traditional IPOs, where investment banks help private companies go public by selling their ownership shares to large institutional investors and listing the shares on a stock market. They are not the only way companies go public, however. Businesses are going public increasingly through alternative arrangements, such as an auction, a direct listing, and a merger by a special purpose acquisition company (SPAC). See *Going Public* by Campbell (2022) on this recent development.

Motivated by the apparent contrast between the academic and popular views, this paper studies whether we can derive the value of investment banks not so much from their expertise or reputation but from being delegated talkers. Our model has one contract designer and two players: The contract designer is a seller (entrepreneur) owning a company with an uncertain future value; the two players are an intermediary (investment bank) and a buyer (investor). The model captures a few critical stages of a typical IPO, in which the investment bank conducts due diligence, generates private information about the firm, and communicates it to potential investors. First, in the *contracting stage*, the seller designs a contract that specifies the state-contingent transfers between the buyer and the intermediary. Next, in the *due diligence stage*, the intermediary publicly announces its due diligence structure: the structure of its costly information acquisition on the company's uncertain state. She then observes a private signal generated from that information structure. Finally, in the *roadshow stage*, the intermediary talks to the buyer, who decides afterward whether to invest in the company or not. The talk is cheap: the intermediary's message is costless, non-binding, and unverifiable.[2] The players' preferences are such that the seller wants to sell his firm, the intermediary wants to maximize her expected profits, and the buyer wants to invest in a firm with a high market value in the future.

In this setting, we show that a contract that implements a positive success rate requires that the intermediary shares the risk of loss with the buyer (Theorem 1). Namely, when the buyer experiences losses, the intermediary also incurs losses; when the buyer experiences gains, the intermediary also incurs gains. The intuition is simple. For the intermediary's communication to be credible to the buyer, the incentives of the intermediary and the buyer must be at least somewhat aligned. The result is similar to one in the canonical cheap talk game of Crawford and Sobel (1982), in which the conflict of interest between the sender and the receiver must be small enough for their communication to be informative. Our prediction is broadly con-

[2] In IPOs in the United States, the Securities and Exchange Commission (SEC) regulates investment banks' (and sellers') communication with investors to ensure complete and fair disclosure of all material information as required by law. In practice, investment banks have substantial room to tell the story of the seller's firm. For example, Facebook's registration statement for its IPO begins: "Our mission is to make the world more open and connected. People use Facebook to stay connected with their friends and family, to discover what is going on in the world around them, and to share and express what matters to them to the people they care about. [...] We believe that we are at the forefront of enabling faster, easier, and richer communication between people and that Facebook has become an integral part of many of our users' daily lives."

sistent with a well-known phenomenon: most IPOs in the United States are under "firm commitment" contracts—in which the investment bank takes on the risk of buying and reselling all of the shares—rather than "best-efforts" contracts, in which the investment bank receives a fixed fee and does not face such a risk.[3]

Building on the previous observation, we derive a property of any seller-optimal contract: an investment banking contract is seller-optimal only if it satisfies *minimally sufficient incentive alignment* (hereafter *minimal incentive alignment*) between the intermediary and the buyer (Theorem 2). By *incentive alignment*, we mean that the intermediary and the buyer agree on the preferred action in all probable realizations of the intermediary's due diligence. By *minimal* incentive alignment, we mean additionally that the buyer is indifferent between investing and not investing even when the due diligence produces good news. In other words, a contract satisfying minimal incentive alignment—a necessary condition for a seller-optimum—aligns the two parties' interests just enough to make the intermediary's communication credible.

This property of seller-optimal contracts points to a natural characterization: a contract is seller-optimal if and only if it maximizes the intermediary's bias in the seller's favor while maintaining minimal alignment with the buyer's incentives. (Theorem 3) Intuitively, if the intermediary's incentives were too distorted in the seller's favor, the buyer would stop believing in the intermediary's message. Therefore, the seller should align the intermediary's interests as close as possible to his own (i.e., incentivize selling) but not too close, so that the intermediary's talk remains credible to the buyer. To this end, the seller should promise a higher rate of return on investment (ROI) to the intermediary than to the buyer (Corollary 1). This result explains why, in observed IPOs, entrepreneurs sell their shares to investment banks at a significant discount relative to the prices offered to the public[4]: such contracts make the sales more profitable to the investment banks and bring their incentives closer to the sellers'.

[3]In contrast, it is not clear why best-efforts contract should not work as effectively as firm commitment contracts, if investment banks operated more like reputable experts than delegated talkers. See Ritter (2003) and Eckbo, Masulis, and Norli (2007) on the different types of IPO contracts and the prevalence of firm commitment contracts.

[4]For example, Facebook sold 421,233,615 shares in its IPO at an offering price of \$38 per share and an investment banking discount of 1.1 percent. This arrangement means that the investment banks led by Morgan Stanley bought these shares at a net offering price of \$37.582 per share and resold them to institutional investors at \$38, effectively earning 1.1 percent of the total value raised from investors as investment banking fees. Consequently, the investment banks earned $1.1\% \times \$38$ per share $\times\ 421,233,615$ shares $\approx \$176$ million.

Taken together, our results imply that delegated communication can make the seller better off relative to two benchmarks: (1) when the seller himself takes the investment bank's role and (2) when the buyer assumes that role. In the first scenario, when the seller serves as his investment bank, he is incentivized to recommend the buyer to invest regardless of his private information. Consequently, even if the seller has better expertise about his firm than anyone else, his talk cannot influence the buyer's decision. A crucial reason is that the seller cannot commit to state-contingent transfers, as he sells off the ownership shares of the firm "as is" without responsibility for future gains or losses. In the second scenario, when the buyer serves as the investment bank, he tends to acquire the firm's information more objectively than the seller prefers, resulting in a suboptimal success rate. In contrast to these benchmarks, using the intermediary allows the seller to design the intermediary's incentives to maximize the probability of a successful sale while keeping the intermediary's talk credible and informative.

Existing models in the literature customarily derive the value of investment banks from expertise or reputation. Their investment banks often have better prior information about the seller's firm (Baron and Holmström, 1980; Baron, 1982; Biais, Bossaerts, and Rochet, 2002) or collect the firm's information at a lower cost (Ramakrishnan and Thakor, 1984). Other models let investment banks face a repeated game (Beaty and Ritter, 1986; Booth and Smith II, 1986; Chemmanur and Fulghieri, 1994), making them concerned about their reputation to tell the truth. Empirical studies find mixed evidence on these models. Ritter and Welch (2002) argue that neither the expertise nor reputation of investment banks are a primary driver of their observed phenomena, whereas Fang (2005) and Brau and Fawcett (2006) argue otherwise. Nonetheless, recent works continue to broadly conform to either view (Eckbo, Masulis, and Norli, 2007; Ljungqvist, 2007; Ragupathy, 2011; Lee and Masulis, 2011; Katti and Phani, 2016; Lowry et al., 2017). In comparison, the investment bank in our model has no particular expertise or reputation, as she does not have better prior information or lower cost of acquiring information as others while facing a one-shot game.[5]

[5]More broadly, our paper adds to the theory of financial intermediation that studies the Principal-Agent relationship between financial intermediaries (agents) and their clients (principals). A canonical view in the literature is that of Diamond (1984). He argues that financial intermediaries serve as *delegated monitors*, pooling deposits from many customers and making loans to entrepreneurs. This view thus focuses on the relationship between the intermediary and the investor (depositor). In contrast, we focus on the relationship between the intermediary and the seller (entrepreneur) and argue that, to entrepreneurs, intermediaries serve as delegated talkers.

Although we use investment banking as the primary example, our model makes a more general point about Sender-Receiver games: We let the sender's credibility arise endogenously by allowing a contract designer to determine the conflict of interest between the seller and the buyer. In the classic setting by Crawford and Sobel (1982), the sender and the receiver have exogenously given biases and exogenously given private information. Many subsequent studies add endogenous information acquisition to this setting while retaining the exogenous conflict of interest (Austen-Smith, 1994; Pei, 2015; Argenziano, Severinov, and Squintani, 2016; Kreutzkamp, 2022; Lou, 2022; Lyu and Suen, 2022). A few papers allow endogenous conflicts of interest but do not have a separate designer; in Ivanov (2010), the receiver decides the conflict of interest between the sender and the receiver. In Antic and Persico (2020), both the sender and the receiver choose their incentive schemes—for example, by deciding their ownership shares in a company—before engaging in communication. In comparison, we let a distinct contract designer determine the conflict of interest, which consequently determines the sender's credibility. For this reason, our model differs from models of mediated communication (Goltsman et al., 2009; Ganguly and Ray, 2012; Ambrus, Azevedo, and Kamada, 2013)—which takes the conflicts of interest as given—and models of mediated persuasion (Arieli, Babichenko, and Sandomirskiy, 2022; Zapechelnyuk, 2022), which takes the sender's credibility or commitment power as given.

We organize the rest of our paper as follows. In Section 2, we define the model, the equilibrium, and the seller-optimal contract. In Section 3, we derive risk-sharing as a necessary condition for any contract with a positive success rate. In Section 4, we show that any seller-optimal contract satisfies minimal incentive alignment, arriving at a characterization. In Section 5, we discuss the seller's optimal outcome compared to those without an intermediary, concluding the paper. We include all proofs in the Appendix.

2 Model of Initial Public Offering (IPO)

An initial public offering (IPO) is an intricate process spanning several months to potentially years, with many rounds of planning, negotiation, and execution among entrepreneurs, investment bankers, financial analysts, lawyers, auditors, investors,

For more theories of financial intermediation, see Bhattacharya and Thakor (1993), Table 3 of Thakor (2020), and Section 2 of Clark, Houde, and Kastl (2021).

and regulators.[6] Our model focuses on the key actors and stages of IPOs that are relevant to our question.

There are Seller (S), Intermediary (I), and Buyer (B). The seller (he) is an entrepreneur whose firm has an uncertain future state $\omega \in \Omega = \{0, 1\}$, where ω is the firm's *opening price*, or the firm's market value when the ownership shares start trading on a stock exchange.[7] There is a publicly known, objective prior probability $p \in (0, 1)$ on the state $\omega = 1$. The intermediary (she) is an investment bank that can acquire private information about the firm's state and send a cheap message to the buyer. The buyer (he) is a consortium of investors whose final action is $a \in A = \{0, 1\}$ for *investing* (or *buying*, $a = 1$) and *not investing* (or *not buying*, $a = 0$) in the ownership share of the seller's firm. If the buyer invests, the seller gives up his firm and receives a *net offering price* of κ.[8] If the buyer does not invest, the seller keeps his firm and retains its objective value p. The net offering price is given as $\kappa \in (p, 1)$, interpreted as having been optimally chosen at the beginning.

The seller designs a mechanism between the intermediary and the buyer through the following arrangement. He first makes a take-it-or-leave-it offer[9] of a *contract* $t = (t_0, t_1) \in \mathbb{R}^2$ to the intermediary, where t_ω is the amount of transfer the intermediary receives from the buyer when the buyer invests ($a = 1$) and the realized state is ω.[10] If the intermediary rejects the offer, the seller retains his firm and other agents receive zero payoffs. If she accepts the offer, the intermediary and the buyer enter a

[6] Hall et al. (2016) is a good practical introduction to the detailed IPO process.

[7] We simplify our problem by assuming that there are only two possible values and normalize them to 0 and 1. For example, an IPO may either be over- or under-subscribed, depending on whether there is a greater or less quantity of shares demanded than offered by the seller through the intermediary. These two scenarios may lead to a higher or lower opening price than the net offering price.

[8] For example, when Facebook went public on May 18th, 2012, Mark Zuckerberg sold its ownership at \$37.582 per share net of investment banking fees (\$0.418 per share). On the same day, the shares started trading on the Nasdaq Stock Market at an opening price of \$42 per share. Assigning this outcome as the favorable state in our model ($\omega = 1$), Facebook's net offering price was $\kappa = 37.582/42 \approx 0.9$ after normalizing the opening price to 1

[9] By letting the seller make a take-or-leave-it offer, we assume that the seller has the full bargaining power between himself and the intermediary. This setup is reasonable when many investment banks compete to win the lead underwriter role for the seller's IPO. It is straightforward to relax the assumption by introducing a reservation payoff for the intermediary.

[10] In Facebook's example, t_1 represents the investment banking fee of \$0.42 per share in the good state, or $t_1 = 0.42/42 = 0.01$ after normalizing by the opening price of \$42 per share. In the same example, t_0 represents the investment bank's payoff in the bad state, such as when the IPO is under-subscribed and faces a smaller demand for shares than offered. As with most such contracts in the United States, Facebook's underwriting agreement specified that the intermediary would buy all shares first and resell them to the investors. That is, t_0 was negative, as the intermediary would face losses from unsold shares in an under-subscribed IPO.

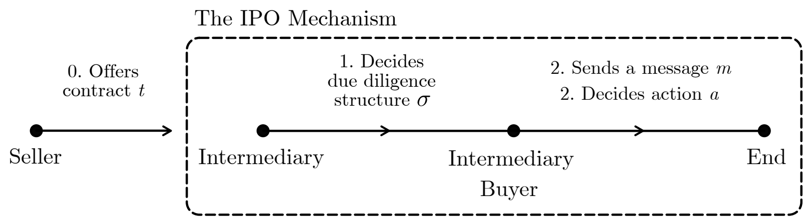

Figure 1: Summary of timing

two-player game called the *IPO mechanism*.

The IPO mechanism consists of two stages that involve only the intermediary and the buyer. Let $\mathcal{S} = \{s_0, s_1, \ldots, s_{J-1}\}$ be a fixed set of possible signals. In Stage 1 (the "due diligence stage"), the intermediary publicly chooses an *information structure* or a *due diligence structure*, a map $\sigma : \Omega \longrightarrow \Delta(\mathcal{S})$, where we write $\sigma(s|\omega)$ to represent the conditional probability of signal s given the state ω. We write the set of all possible due diligence structures as $\Sigma = (\Delta(\mathcal{S}))^{\Omega}$. After choosing a due diligence structure $\sigma \in \Sigma$, the intermediary privately observes a realized signal $s_j \in S$ from the probability distribution $\sigma(\cdot|\omega)$. Consequently, the probability of receiving a signal s_j under the due diligence structure σ given a prior belief p is

$$P_\sigma(s_j) = (1-p)\sigma(s_j|0) + p\sigma(s_j|1).$$

Consider any due diligence structure $\sigma \in \Sigma$. By Bayes' rule and consistency with the prior beliefs, the intermediary's private posterior belief on the state $\omega = 1$ given a realized signal $s_j \in \mathcal{S}$ is

$$q_j = \frac{p\sigma(s_j|1)}{P_\sigma(s_j)},$$

for all j such that $P_\sigma(s_j) > 0$. We say that the due diligence structure σ *induces* a posterior q_j with probability $P_\sigma(s_j)$. A posterior q_j is *probable* if $P_\sigma(s_j) > 0$.

In Stage 2 (the "roadshow stage"), the intermediary sends a message $m \in \mathcal{M} = \mathcal{S}$ to the buyer.[11] The message is cheap: it is (a) costless, (b) non-verifiable, and (c) non-binding. After receiving the message, the buyer chooses an action $a \in A = \{0, 1\}$. Finally, the state ω is publicly revealed. Figure 1 shows the summary of the game's timing.

[11] Letting $\mathcal{S} = \mathcal{M}$ is without loss of generality because, if $\mathcal{S} \neq \mathcal{M}$, we can redefine the signal space $\mathcal{S}'$ and message space $\mathcal{M}'$ as $\mathcal{S}' = \mathcal{M}' = \mathcal{S} \cup \mathcal{M}$.

After the IPO mechanism ends, the seller's payoff is

$$u^S(\omega, a) = \kappa a + p \cdot (1 - a).$$

The intermediary's payoff is

$$u^I_{t,\sigma}(\omega, a) = t_\omega a - C(\sigma),$$

where $t_\omega a$ is the intermediary's revenue and $C : \Sigma \longrightarrow \mathbb{R}$ is the *information cost function*.

Assumption 1. The information cost function C is posterior-separable.[12] That is, there exists a function $c : [0, 1] \longrightarrow \mathbb{R}$ with $c(p) = 0$ such that, for every due diligence structure $\sigma \in \Sigma$ that induces posteriors $q_0, q_1, \ldots, q_{J-1}$,

$$C(\sigma) = \sum_{j=0}^{J-1} P_\sigma(s_j) c(q_j).$$

In addition, the function c satisfies the properties P1–P3:

P1. (Smoothness) c is continuous on $[0, 1]$ and twice differentiable on $(0, 1)$,

P2. (Curvature) c is strictly convex, and

P3. (Steep boundaries) $-\lim_{q\to 0} c'(q) = \lim_{q\to 1} c'(q) = \infty$.

We can interpret the function $c(q)$ as representing a measure of relative certainty at the posterior q from the prior p, for example, as in the reduction in Shannon entropy. Note that the second property makes it more costly to choose a due diligence structure that is more informative in the sense of Blackwell (1953). The third property ensures that the intermediary choose a due diligence structure that induces probable posteriors in the interior of $[0, 1]$.

The buyer's payoff is

$$u^B_t(\omega, a) = (-\kappa + \omega - t_\omega)a.$$

In the IPO mechanism, the intermediary's strategy is $(\sigma, \boldsymbol{\mu})$, where $\boldsymbol{\mu} = \{\mu_\sigma\}_{\sigma\in\Sigma}$ is a collection of *message rules* $\mu_\sigma : \mathcal{S} \longrightarrow \Delta(\mathcal{M})$ that assigns each signal to a probability distribution over messages. We write $\mu_\sigma(m|s)$ to represent the conditional

[12] In the terminology by Caplin, Dean, and Leahy (2022), this cost function is also uniformly posterior-separable.

probability of the message m given the signal s, under the due diligence structure σ. The buyer's strategy is $\boldsymbol{\alpha} = \{\alpha_\sigma\}_{\sigma \in \Sigma}$, a collection of *action rules* $\alpha_\sigma : \mathcal{M} \longrightarrow A$. A *strategy profile* of the IPO mechanism is the triple $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$.

Given a due diligence structure σ, a message rule μ, and an action rule α, define $U^S(\sigma, \mu, \alpha)$, $U_t^I(\sigma, \mu, \alpha)$, and $U_t^B(\sigma, \mu, \alpha)$ as the expected payoffs of the seller, the intermediary, and the buyer, where the expectation is taken over all realizations of the states $\omega \in \Omega$, the signals $s \in \mathcal{S}$, and the messages $m \in \mathcal{M}$.

Definition 1 (Equilibrium)**.** A strategy profile $(\sigma^*, \boldsymbol{\mu}^*, \boldsymbol{\alpha}^*)$ is an *equilibrium* of the IPO mechanism under a contract $t = (t_0, t_1)$ if it satisfies conditions (a)–(c).

(a) ($\boldsymbol{\mu}^*$ and $\boldsymbol{\alpha}^*$ are mutual best responses) For every $\sigma \in \Sigma$, $\mu : \mathcal{S} \longrightarrow \Delta(\mathcal{M})$, and $\alpha : \mathcal{M} \longrightarrow A$,

$$U_t^I(\sigma, \mu_\sigma^*, \alpha_\sigma^*) \geq U_t^I(\sigma, \mu, \alpha_\sigma^*) \quad \text{and} \quad U_t^B(\sigma, \mu_\sigma^*, \alpha_\sigma^*) \geq U_t^B(\sigma, \mu_\sigma^*, \alpha).$$

(b) ($(\boldsymbol{\mu}^*, \boldsymbol{\alpha}^*)$ is not Pareto-dominated) There exists no pair $(\boldsymbol{\mu}, \boldsymbol{\alpha})$ that satisfies condition (a) and, for some $\sigma \in \Sigma$,

$$U_t^I(\sigma, \mu_\sigma^*, \alpha_\sigma^*) \leq U_t^I(\sigma, \mu_\sigma, \alpha_\sigma) \quad \text{and} \quad U_t^B(\sigma, \mu_\sigma^*, \alpha_\sigma^*) \leq U_t^B(\sigma, \mu_\sigma, \alpha_\sigma),$$

with at least one strict inequality.

(c) (σ^* is a best response to $(\boldsymbol{\mu}^*, \boldsymbol{\alpha}^*)$) For every $\sigma \in \Sigma$,

$$U_t^I(\sigma^*, \mu_{\sigma^*}^*, \alpha_{\sigma^*}^*) \geq U_t^I(\sigma, \mu_\sigma^*, \alpha_\sigma^*), \text{ and} \tag{1}$$

$$U_t^I(\sigma^*, \mu_{\sigma^*}^*, \alpha_{\sigma^*}^*) \geq 0. \tag{2}$$

This definition applies the sequential equilibrium (Kreps and Wilson, 1982) to our setting, with two additional restrictions. First, condition (b) demands that the Nash equilibria of the talking stage (Stage 2) subgame be Pareto-efficient. That is, we only consider cases in which the intermediary and buyer communicate efficiently in the talking stage and rule out, for example, babbling equilibria whenever there exists a mutually beneficial, informative one.[13] Second, the condition (c) includes

[13]This requirement shrinks the set of implementable outcomes to reasonable ones, thus preventing the designer from being too powerful. For example, without this requirement, the seller can

both the incentive compatibility and individual rationality constraints (1)–(2) for the intermediary in Stage 1.

We are ready to define a *seller-optimal investment banking contract* or simply *seller-optimal contract.*

Definition 2 (Seller-optimal contract)**.** Let $\mathcal{E}(t)$ denote the set of equilibria of the IPO mechanism under a contract $t = (t_0, t_1)$. A contract t^* is *seller-optimal* (or *optimal*) if

$$t^* \in \operatorname*{argmax}_{t\in\mathbb{R}^2} \left[\sup_{(\sigma,\boldsymbol{\mu},\boldsymbol{\alpha})\in\mathcal{E}(t)} U^S(\sigma, \mu_\sigma, \alpha_\sigma) \right].$$

In other words, a contract is seller-optimal if it maximizes seller's expected utility in the best equilibrium. This definition follows the standard approach in mechanism design, where the designer chooses the best-case outcome from the set of equilibria.

3 Successful investment banking contracts

Define the *success rate* of an equilibrium as the probability that the buyer invests. Let us say that a contract t *implements* a success rate ρ if there exists an equilibrium under t such that the success rate is ρ.

We note that it is impossible to achieve a 100-percent success rate.

Proposition 1. *There exists no contract that implements a success rate of 1.*

The crux of the proof is to apply Bayes-plausibility (Kamenica and Gentzkow, 2011), that the mean of the posteriors q equals the prior p. A success rate of 1 implies that the buyer is willing to invest at every probable posterior q. Bayes-plausibility implies that the buyer is willing to invest *ex ante*, contradicting our assumption on the prior p.

Because a sure success is impossible, we draw our attention to contracts with positive success rates, or *successful* contracts. A successful contract necessarily requires that the intermediary shares the risk of loss with the buyer.

Theorem 1. *If a contract $t = (t_0, t_1)$ implements a success rate $\rho > 0$, then $t_0 \in [-\kappa, 0)$ and $t_1 \in (0, 1 - \kappa]$.*

implement any due diligence structure σ in an equilibrium as long as it satisfies the participation constraint (2), by choosing the uninformative outcome for subgames with any other due diligence structures.

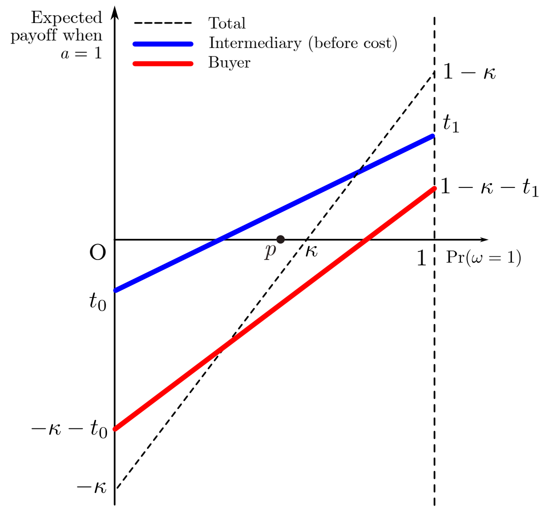

Figure 2: A successful contract requires a joint risk of loss

This result restricts the intermediary's and the buyer's payoffs from investing to have the same signs for any successful contract, as the buyer's payoff is $-\kappa - t_0 \leq 0$ in the bad state and $1 - \kappa - t_1 \geq 0$ in the good state. This restriction implies that a "best-efforts" contract—compensating the intermediary with a fixed transfer (*i.e.*, letting $t_0 = t_1$)—cannot be successful. Rather, a successful contract must make the intermediary incur losses whenever the buyer incurs losses, and make the intermediary incur gains whenever the buyer incurs gains. This result is broadly consistent with the fact that most observed IPOs in the United States are under "firm commitment" contracts, in which the investment bank faces a substantial risk of losses by agreeing to buy all offered shares from the seller and to resell them to investors. Therefore, although our model does not explicitly allow the underwriting process of buying and reselling, the result shows that it captures the essential phenomenon by subjecting the intermediary to a joint risk of losses.

Figure 2 illustrates a typical contract (t_0, t_1) satisfying this requirement. The left and right ends of the dashed line segment indicate the sum of the intermediary's and the buyer's expected payoffs in the bad and good states, respectively. The left ends of the solid line segments, indicating the payoffs in the bad state, are both below the x-axis. The right ends of the solid line segments, indicating the payoffs in the good state, are both above the x-axis.

The simple intuition behind the result is that the incentives of the intermediary and the buyer must be at least somewhat aligned for the cheap talk between them

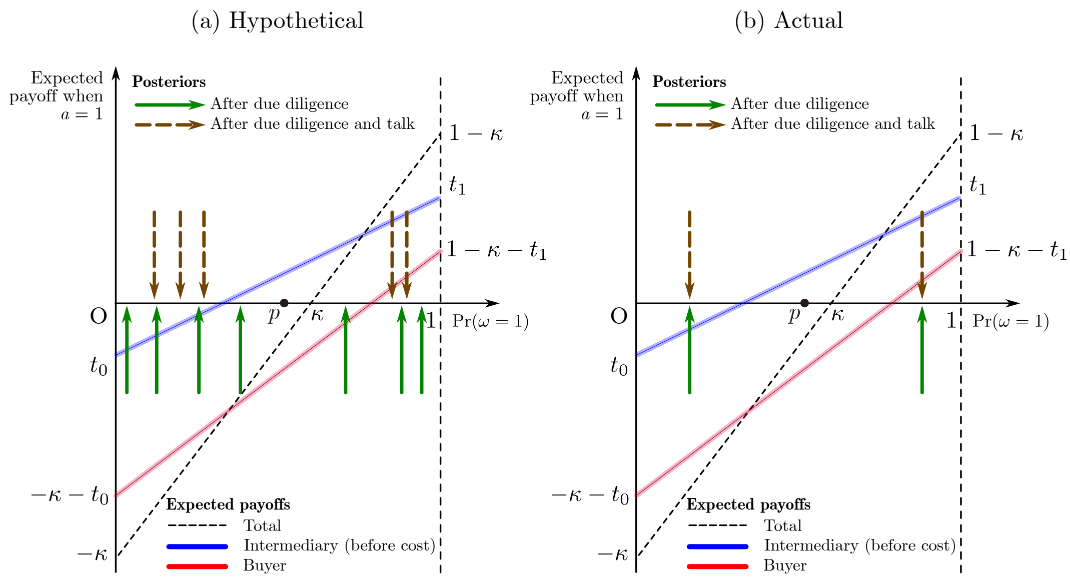

Figure 3: Distribution of the intermediary's and the buyer's posteriors

to be informative. The argument proceeds as follows. Suppose an equilibrium has a positive success rate. This implies that the distribution of the buyer's posteriors is such that the two prefer the same action (either to invest or not) regardless of the realized posterior. Otherwise, either the intermediary or the buyer would want to deviate from the supposed equilibrium, by sending a different message or taking a different action. Therefore, the intermediary's and the buyer's expected payoffs upon investing must have the same signs, for all probable posteriors of the buyer, as illustrated by the brown dashed arrows in Panel (a) of Figure 3. To achieve such incentive alignment, the contract must incur losses to both in the bad state and gains to both in the good state.

We make this idea more precise by introducing a few additional definitions and a key lemma. For any contract $t = (t_0, t_1)$, define functions V_t^I and V_t^B such that, for all $q \in [0, 1]$,

$$V_t^I(q) = (1-q)t_0 + qt_1, \text{ and}$$
$$V_t^B(q) = (1-q)(-\kappa - t_0) + q(1-\kappa-t_1).$$

The two functions represent the Stage-2 expected payoffs given a posterior q for the intermediary and the buyer, respectively, when the buyer chooses to invest ($a = 1$).

The solid line segments in Figure 2 illustrate the functions V_t^I and V_t^B for an example contract.

For every due diligence structure $\sigma \in \Sigma$ and every message rule $\mu : \mathcal{S} \longrightarrow \mathcal{M}$, define $(\mu \circ \sigma)(m|\omega)$ as

$$(\mu \circ \sigma)(m|\omega) = \sum_{s \in S} \mu(m|s)\sigma(s|\omega), \quad \text{for all } m \in \mathcal{M} \text{ and } \omega \in \Omega.$$

That is, $(\mu \circ \sigma)(m|\omega)$ is the induced distribution of messages $m \in \mathcal{M}$ in state $\omega \in \Omega$. Since $\mathcal{M} = \mathcal{S}$, the product $\mu \circ \sigma$ itself is an element of Σ. For example, in both panels of Figure 3, the solid green arrows indicate the intermediary's posteriors after the due diligence σ, whereas the dashed brown arrows indicate the buyer's posteriors after the due diligence and talk $\mu \circ \sigma$.

Definition 3. A due diligence structure $\sigma \in \Sigma$ is *incentive-aligned* under a contract t if, for all probable posteriors q induced by σ, $V_t^I(q)V_t^B(q) \geq 0$.

In other words, a due diligence structure σ is incentive-aligned if both players weakly prefer the same actions in all probable realizations of the signals s from σ.

The following lemma is the key step for the proof of Theorem 1.

Lemma 1. *Suppose a contract t implements a success rate $\rho \in (0, 1)$ with an equilibrium $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$. The equilibrium satisfies*

- **(a)** *Fully revealing messages: $\mu_\sigma \circ \sigma$ induces the same distribution of posteriors as σ,*
- **(b)** *Incentive alignment: σ is incentive-aligned, and*
- **(c)** *Binary due diligence: σ induces exactly two probable posteriors.*

Put differently, in any equilibrium with a positive success rate, the intermediary truthfully reports the result of her due diligence and prefers the same action as the buyer after seeing one of two probable signals. This result means that the distribution of posteriors for the intermediary and the buyer should resemble the arrows in Panel (b) rather than (a) in Figure 3. First, because the buyer's message rule is fully revealing, the points of the dashed brown arrows overlap exactly with those of the solid green arrows. Second, because the incentives of the intermediary and the buyer are aligned, their expected payoffs at those arrows have the same signs. Third, because the due diligence is binary, there are only two solid green arrows.

The proof of the lemma considers a hypothetical equilibrium with a positive success rate and shows that it should satisfy the three properties. The starting point is that the buyer's posteriors must satisfy incentive alignment for the cheap talk to be informative. Given this incentive alignment, the intermediary has no reason to garble her message; if she garbles, it is strictly less costly to obtain the same outcome by choosing a due diligence that directly induces the same distribution of posteriors and reporting truthfully. This result echoes Pei (2015) who finds that the sender in a cheap talk game with endogenous information acquisition always fully reveals his private information. Finally, given the incentive alignment and the truthful revelation, there is no reason for the intermediary to conduct a due diligence that produces more than two posteriors, because there are only two possible actions. This result is similar to Corollary 1 in Matějka and McKay (2015).

4 Optimal investment banking contracts

Building on our previous results about successful contracts, we narrow our attention to seller-optimal ones: Seller-optimal contracts satisfy minimal (or minimally sufficient) incentive alignment between the intermediary and the buyer.

4.1 Optimal contracts satisfy minimal incentive alignment

To define minimal incentive alignment, write $T = [-\kappa, 0) \times (0, 1-\kappa]$, the set of contracts that satisfy the necessary condition of a successful contract from Theorem 1. Without loss of generality, let $\widehat{\Sigma}$ denote the set of all $\sigma \in \Sigma$ such that, for all $\omega \in \Omega$ and for all $j = 2, 3, \ldots, J-1$, we have $\sigma(s_j|\omega) = 0$ and $\sigma(s_1|0) < \sigma(s_1|1)$. That is, $\widehat{\Sigma}$ is the set of all binary due diligence structures such that s_0 is bad news (with posterior less than p) and s_1 is good news (with posterior greater than p).

Definition 4 (Minimal incentive alignment)**.** A due diligence structure $\sigma \in \widehat{\Sigma}$ that induces a pair of posteriors $(\ell, r) \in [0, p) \times (p, 1]$ is *minimally incentive-aligned* under a contract $t \in T$ if it is incentive-aligned under t and it satisfies $V_t^B(r) = 0$.[14]

In other words, minimal incentive alignment means that the intermediary and the buyer agree on their interim weakly preferred actions just sufficiently, as the buyer is indifferent between investing and not investing even after hearing good news.

[14]Recall that V_t^I and V_t^B are functions defined as $V_t^I(q) = (1-q)t_0 + qt_1$ and $V_t^B(q) = (1-q)(-\kappa - t_0) + q(1-\kappa-t_1)$. A due diligence structure $\sigma \in \Sigma$ is *incentive-aligned* under a contract t if $V_t^I(q)V_t^B(q) \geq 0$ for all probable posteriors q induced by σ (Definition 3).

(a) Not minimal (b) Minimal

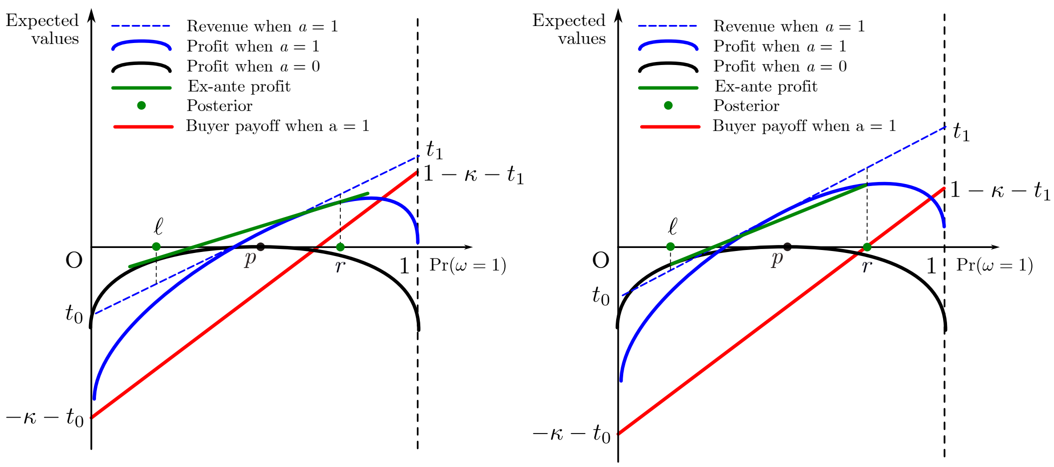

Figure 4: Incentive alignment may or may not be minimal in an equilibrium

Note: Profit refers to the intermediary's expected revenue ($V_t^I(q)$ if the buyer invests; 0 otherwise) minus the measure of relative certainty $c(q)$.

Theorem 2. *Suppose a contract $t^* \in T$ is seller-optimal, implementing a maximal success rate $\rho^* > 0$ with equilibrium $(\sigma^*, \boldsymbol{\mu}^*, \boldsymbol{\alpha}^*)$. The due diligence structure σ^* is minimally incentive-aligned under t^*.*

To better understand this statement, we illustrate in Figure 4 what it means for a binary due diligence structure to be minimally incentive-aligned. Panel (a) shows an equilibrium due diligence structure $\sigma \in \widehat{\Sigma}$ that is incentive-aligned under some contract $t \in T$. The points ℓ and r on the x-axis are the posteriors induced by σ. The expected payoffs V_t^I (dashed line) and V_t^B (solid red line) from investing are such that both the intermediary and the buyer prefer not investing ($a = 0$) when they see bad news (ℓ) and prefer investing ($a = 1$) when they see good news (r). This due diligence structure is not minimally incentive-aligned because the buyer strictly prefers investing when he sees good news. In contrast, the equilibrium due diligence structure shown in Panel (b) is minimally incentive-aligned, as the buyer is indifferent between investing and not investing when he sees good news (r). In essence, Theorem 2 says that any equilibrium due diligence structure under a seller-optimal contract must resemble Panel (b) rather than Panel (a).

The intuition behind the result is the following: Unless a contract already achieves minimal incentive alignment, the seller can always improve the success rate by adjusting the contract slightly to further incentivize selling. Specifically, the seller can bring the intermediary's incentives closer to his own by increasing the intermediary's revenues in both states (t_0 and t_1) upon a successful sale. In Panel (a) of Figure 4, this adjustment corresponds to shifting up the intermediary's revenue and profit curves upon success ($a = 1$). As the intermediary's profit from a successful sale shifts up, the intermediary adjusts its due diligence by lowering both posteriors (ℓ and r) to maximize its expected profit.[15] This change produces good news (r) with greater probability than before, increasing the success rate. In contrast, if a contract already achieves minimal incentive alignment as in Panel (b), such improvement is not possible. The reason is that increasing the intermediary's revenues (t_0 and t_1) further breaks the incentive alignment between the intermediary and the buyer, rendering the intermediary's communication incredible. It follows that any seller-optimal contract must satisfy minimal incentive alignment of its equilibrium due diligence structure. Our proof of the theorem reproduces this logic.

The same intuition makes it clear that even though our result may sound familiar from Bayesian persuasion and information design problems (Kamenica and Gentzkow, 2011; Kamenica, 2019; Bergemann and Morris, 2019), it is driven by a different force: the contractual design. Depending on the contract and the resulting ex-post payoffs, the information designer in our model (the intermediary) may happily choose a more balanced information structure, not a minimally aligned one. However, such a contract would be strictly suboptimal for the contract designer (the seller).

4.2 Characterization of the optimal contract

Theorem 2 is useful because it leads to a natural characterization of seller-optimal contracts: A contract is seller-optimal if and only if it maximizes the intermediary's bias in favor of the seller while maintaining minimal alignment with the buyer's incentives. To arrive at this characterization, we make an auxiliary definition.

[15] The intermediary's equilibrium choice of the due diligence follows the concavification method of Kamenica and Gentzkow (2011). That is, the intermediary finds a pair of posteriors such that the line passing through the corresponding points on the profit curves are tangent to both curves. This approach is valid in our setting as long as the posteriors satisfy strict incentive alignment ($V_t^I(q)V_t^B(w) > 0$ for all probable posteriors q) between the intermediary and the buyer, so that the intermediary's communication is strictly credible. The resulting solution is identical to that in a problem of rational inattention (Matějka and McKay, 2015; Caplin, Dean, and Leahy, 2019). Note that this method is no longer valid when the incentive alignment is minimal, as in Panel (b) of Figure 4.

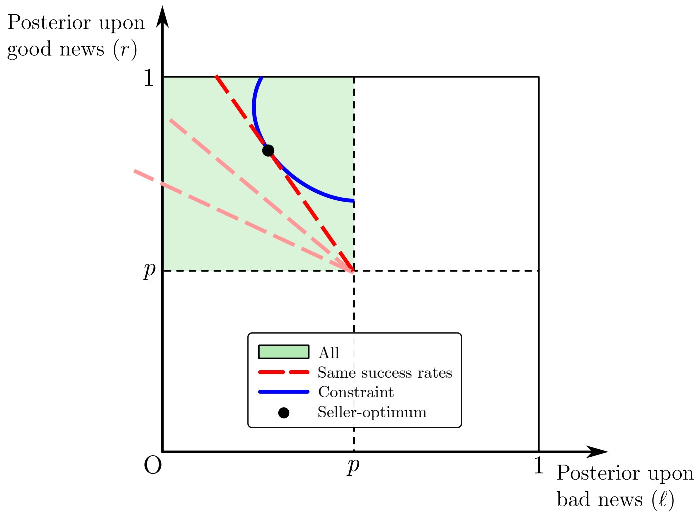

Figure 5: Finding a seller-optimal pair of posteriors

Definition 5. A pair of posteriors $(\ell^*, r^*) \in [0, p) \times (p, 1]$ is *seller-optimal* if

$$(\ell^*, r^*) \in \underset{(\ell,r)\in[0,p)\times(p,1]}{\text{argmax}} \frac{p-\ell}{r-\ell}$$

subject to the *implementability constraint*

$$c(r) - c(\ell) - (r-\ell)c'(\ell) = r - \kappa.$$

Note that the definition does not involve any specific contract t. To restate, a pair of posteriors is seller-optimal if it maximizes the probability of good news[16] subject to the implementability constraint. Intuitively, the constraint represents the set of all minimally incentive-aligned due diligence structures that the intermediary would choose as her best response under some contract. Figure 5 illustrates the implementability constraint and a seller-optimal pair of posteriors. Such a seller-optimal pair of posteriors exists if the cost of acquiring information is sufficiently small.[17]

Theorem 3. *Suppose a contract $t^* = (t_0^*, t_1^*) \in T$ implements a positive success rate.*

[16]The probability of good news is $P_\sigma(s_1) = \frac{p-\ell}{r-\ell}$, as $(1 - P_\sigma(s_1))\ell + P_\sigma(s_1)r = p$ by Bayes plausibility.

[17]If acquiring information is too costly, the set of posteriors satisfying the implementability constraint is empty, and no seller-optimal pair of posteriors exists. In Appendix B, we provide the precise necessary and sufficient conditions for the existence of a solution.

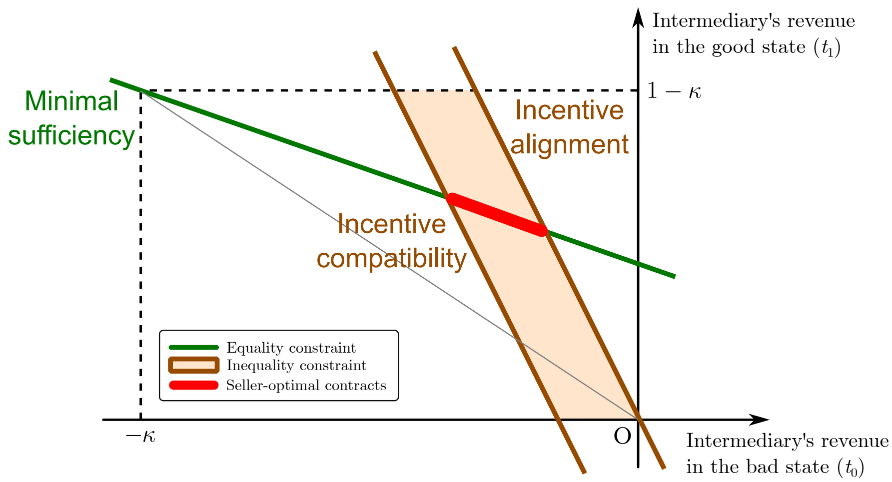

Figure 6: Seller-optimal contracts

The contract is seller-optimal if and only if

$$\text{(Minimal sufficiency)} \qquad (1-r^*)t_0^* + r^* t_1^* = r^* - \kappa, \tag{3}$$

$$\text{(Incentive alignment)} \qquad (1-\ell^*)t_0^* + \ell^* t_1^* \leq 0, \ \text{and} \tag{4}$$

$$\text{(Incentive compatibility)} \qquad (1-\ell^*)t_0^* + \ell^* t_1^* \geq c(r^*) - c(\ell^*) - (r^* - \ell^*)c'(r^*). \tag{5}$$

for a seller-optimal pair of posteriors $(\ell^*, r^*) \in [0, p) \times (p, 1]$.

In other words, a contract implements a seller-optimal pair of posteriors if and only if it satisfies the three conditions (3)–(5). Each condition has a concrete economic meaning. Condition (3) means that the contract makes the buyer indifferent between investing and not investing upon hearing good news. Condition (4) means that the contract makes the intermediary to prefer not investing upon seeing bad news, thus aligned with the buyer. Condition (5) means that the contract makes it incentive-compatible for the intermediary to choose the seller-optimal pair of posteriors over others. All in all, a contract is optimal if and only if it maximizes the success rate (implements the seller-optimal posteriors) while maintaining minimally sufficient incentive alignment (conditions 3–4) in equilibrium (condition 5). Because the right-hand side of (5) is strictly negative (by the strict convexity of c), the set of seller-optimal contracts is nonempty given (ℓ^*, r^*): It is the closed line segment defined by (3)–(5). Figure 6 illustrates the three conditions and shows the resulting set of seller-optimal contracts.

Many contracts implement the same seller-optimal pair of posteriors, as seen from Theorem 3 and Figure 6. This multiplicity is not an accident but a consequence of minimal incentive alignment. When the equilibrium due diligence structure is not minimally incentive-aligned (as in Figure 4a), it responds sensitively to changes in the contract. However, when the equilibrium due diligence structure is minimally incentive-aligned (as in Figure 4b), the intermediary does not lower the posterior upon good news (r) in response to certain changes in the contract, as it would break the incentive alignment and make her subsequent communication uninformative. This insensitivity implies that there are multiple contracts that induce the same pair of posteriors.

4.3 Implication for the return on investment (ROI)

An immediate corollary of Theorem 3 is that the intermediary earns a higher ex-post return relative to her risk than the buyer. Define the *(ex-post) return on investment* (ROI) as

$$\text{Return on investment (ROI)} = \frac{\text{Dollars earned in the good state}}{\text{Dollars lost in the bad state}},$$

which applies the conventional notion of ROI—the net value of investment divided by the cost of investment—to our setting. For example, the aggregate ROI for the intermediary and the buyer together is $\frac{1-\kappa}{\kappa}$, as the two agents altogether earn $1-\kappa$ (the net value of investment) in the good state and lose κ (the cost of investment) in the bad state. In comparison, the ROI for the intermediary is $\frac{t_1}{-t_0}$.

To focus on a salient case among the many seller-optimal contracts, define the *minimum-variance seller-optimal contract* as the contract that satisfies the condition (4) from Theorem 3 with equality. In Figure 6, it is the point closest to the origin on the line segment representing the optimal contracts. This terminology is adequate because it minimizes the variance of the transfer t_ω among all seller-optimal contracts. Let $(\ell^*, r^*) \in [0, p) \times (p, 1]$ denote a seller-optimal pair of posteriors.

Corollary 1. *Under any seller-optimal contract, the intermediary's return on investment (ROI) is higher than the buyer's ROI. In particular, under the minimum-variance seller-optimal contract,*

$$\textit{Intermediary's ROI} = \frac{1-\ell^*}{\ell^*} > \frac{1-\kappa}{\kappa} > \frac{1-r^*}{r^*} = \textit{Buyer's ROI.}$$

The intuition behind this corollary is our earlier characterization that seller-optimal contracts bring the intermediary's incentives closest to the seller's while maintaining minimal alignment with the buyer's. A necessary condition to achieve such a contract is to offer the intermediary a higher rate of return from investing than the overall return.[18]

4.4 Comparative statics

In addition to the immediate implications for the ROI, our full characterization of the seller-optimal contracts (Theorem 3) facilitates the following exercises in comparative statics. Our model has three exogenous variables: the initial probability p that the entrepreneur's firm will have high market value, the net offering price κ of the entrepreneur's firm, and the cost of information acquisition C by the investment bank. We summarize how changes in these underlying parameters affect the model's endogenous outcomes and include the precise results as Corollaries C.1–C.3 in the Appendix.

First, an increase in the ex-ante probability p raises the success rate ρ^* of selling the entrepreneur's firm to the buyer under optimal contracts. Although this result by itself is unsurprising, the changes in the due diligence structure and the contract that lead to this outcome are interesting: The intermediary collects (and communicates) less precise information than before, resulting in a reduced compensation (both t_0^* and t_1^*) relative to the firm's market value. The intuition is that the investment bank does not need to work as hard as before to sell the entrepreneur's firm to the buyer.

Second, an increase in the net offering price κ reduces the seller's success rate ρ^* and his expected revenue $\kappa\rho^*$ under optimal contracts. In other words, the buyer's demand ρ^* for the entrepreneur's firm through the investment bank is inelastic with respect to the firm's price κ. Intuitively, this result arises because of agency costs

[18] This implication is testable in principle by checking whether the condition $\frac{t_1}{-t_0} > \frac{1-\kappa}{\kappa}$ is satisfied, although data restrictions on t_0 make this test difficult in practice. For example, Facebook sold its shares at a net offering price of \$37.582 per share, and its opening price was \$42 per share on the Nasdaq Stock Market when the shares started trading there. After normalizing the opening price to 1, these values imply that $\kappa = 37.582/42 \approx 0.9$ and the overall ROI was $\frac{1-\kappa}{\kappa} \approx 0.12$. Meanwhile, the investment banks earned \$0.418 per share as fees or 1.1 percent of the gross offering price of \$38 per share. This value implies $t_1 = 0.418/42 \approx 0.01$ after normalization. Therefore, the required condition is satisfied as long as $t_0 > -t_1 \cdot \frac{\kappa}{1-\kappa} \approx -0.085$, *i.e.*, the investment bank in the bad state loses less than 8.5 percent of Facebook's potential market value. However, we do not observe t_0, because the bad state did not materialize, and the value is typically not disclosed in security filings in the United States.

or the implementability constraint: it is costly to induce the intermediary to collect and communicate information that extracts a greater surplus from the buyer.

Third, as acquiring information becomes less costly, the due diligence structure under optimal contracts approaches that of Bayesian persuasion by Kamenica and Gentzkow (2011). That is, it becomes close to the structure that the seller would choose if he could commit to revealing information truthfully after acquiring it costlessly. Unlike in their setting, however, the seller in our model achieves a similar outcome not by his commitment power but by appropriately designing the contract that aligns the intermediary's incentives with the buyer's. As a result, the success rate ρ^* under optimal contracts approaches p/κ as the cost of information scales down approaching 0.

5 Discussion

Our results make concrete predictions about investment banking contracts in a model in which an investment bank (intermediary) serves as a delegated communicator from an entrepreneur (seller) to an investor (buyer). Any successful contract must exhibit a shared risk of loss between the intermediary and the buyer (Theorem 1). A seller-optimal contract maximizes the intermediary's incentives in favor of the seller while maintaining a minimal alignment with the buyer's incentives (Theorems 2–3). Seller-optimal contracts provide a higher rate return relative to risk to the intermediary than to the buyer (Corollary 1).

It is straightforward to see from our analysis how this arrangement with an intermediary can strictly benefit the seller compared to scenarios without one. For instance, on the one hand, consider an alternative model in which the seller himself serves as the investment bank. The seller enters the IPO mechanism as the intermediary, with a critical restriction that $t_0 \geq 0$ and $t_1 \geq 0$—in other words, the entrepreneur cannot commit to any state-contingent transfers to the buyer. That is, the entrepreneur is selling the firm "as is" and cannot be held responsible for its future state. In this scenario, the seller's incentives cannot be aligned with the buyer's, as he gains a positive amount in both good and bad states when the buyer invests. It follows that there is no trade. This outcome remains the same even if the entrepreneur, with insider knowledge about his firm, has a significantly lower or zero cost of information than any investment bank. Therefore, intermediaries acting as delegated talkers with sufficient capital are valuable when the seller himself cannot credibly commit to a shared risk of loss.

On the other hand, consider another alternative model in which the investor serves as the investment bank: the original buyer plays both the intermediary and the buyer in the IPO mechanism. The contract is irrelevant in this case because it represents the amount of transfers from the buyer to the buyer himself. Regardless of the contract, the buyer gets the payoff $-\kappa$ in the bad state and $1-\kappa$ in the good state. As a result, the buyer collects information as the intermediary would in the original model under the contract $(t_0, t_1) = (-\kappa, 1-\kappa)$. Unlike the intermediary in the original model, the buyer is not constrained by incentive alignment because he communicates with himself. Despite this absence of friction in communication, this arrangement cannot be seller-optimal when acquiring information is sufficiently inexpensive. The success rate in the alternative model approaches p as the cost of information acquisition scales down approaching 0, whereas the success rate in the original model approaches a strictly higher value, p/κ.[19]

All in all, our model formalizes an under-studied function of investment banks in initial public offerings (IPOs) as delegated communicators. Without assuming their expertise or reputational concerns, we show that their risk-bearing capacity makes their cheap talk credible and informative. Such delegated talkers are especially valuable when the seller cannot be held liable for the ex-post losses to the buyer.

Although our paper has focused on investment banks and IPOs as a leading example, one can apply this model to similar settings where two parties have a substantial conflict of interest. For instance, consider a high school student seeking admission to a college and the college's admission committee that aims to accept only the most taleted students. Given the incentives of both, the student's application essay claiming his talent would not be credible. As a result, he needs an intermediary—his teacher or counselor—to evaluate him and send a recommendation letter. In this case, a student-optimal school policy would incentivize the teacher to be maximally biased in the student's favor while maintaining minimal alignment with the admission committee's interests.

[19] To see why the success rate approaches p/κ in the original model, see the related discussion in Subsection 4.4 (Comparative statics) or its precise statement in Corollary C.3 in the Appendix.

References

Ambrus, Attila, Eduardo M. Azevedo, and Yuichiro Kamada (2013) “Hierarchical cheap talk,” *Theoretical Economics*, 8 (1), 233–261.

Antic, Nemanja and Nicola Persico (2020) “Cheap talk with endogenous conflict of interest,” *Econometrica*, 88 (6), 2663–2695.

Argenziano, Rossella, Sergei Severinov, and Francesco Squintani (2016) “Strategic information acquisition and transmission,” *American Economic Journal: Microeconomics*, 8 (3), 119–155.

Arieli, Itai, Yakov Babichenko, and Fedor Sandomirskiy (2022) “Bayesian persuasion with mediators,” *arXiv preprint arXiv:2203.04285*.

Austen-Smith, David (1994) “Strategic transmission of costly information,” *Econometrica: Journal of the Econometric Society*, 955–963.

Baron, David P (1982) “A model of the demand for investment banking advising and distribution services for new issues,” *The Journal of Finance*, 37 (4), 955–976.

Baron, David P and Bengt Holmström (1980) “The investment banking contract for new issues under asymmetric information: Delegation and the incentive problem,” *The Journal of Finance*, 35 (5), 1115–1138.

Beaty, Randolph P and Jay Ritter (1986) “Investment banking, reputation, and the underpricing of initial public offering,” *Journal of Financial Economics*, 15 (1–2), 213–232.

Bergemann, Dirk and Stephen Morris (2019) “Information design: A unified perspective,” *Journal of Economic Literature*, 57 (1), 44–95.

Bhattacharya, Sudipto and Anjan V Thakor (1993) “Contemporary banking theory,” *Journal of financial Intermediation*, 3 (1), 2–50.

Biais, Bruno, Peter Bossaerts, and Jean-Charles Rochet (2002) “An optimal ipo mechanism,” *The Review of Economic Studies*, 69 (1), 117–146.

Blackwell, David (1953) “Equivalent comparisons of experiments,” *The annals of mathematical statistics*, 265–272.

Booth, James R and Richard L Smith II (1986) "Capital raising, underwriting and the certification hypothesis," *Journal of Financial Economics*, 15 (1-2), 261–281.

Brau, James C and Stanley E Fawcett (2006) "Initial public offerings: An analysis of theory and practice," *The Journal of Finance*, 61 (1), 399–436.

Campbell, Dakin (2022) "Going public: how Silicon Valley rebels loosened Wall Street's grip on the IPO and sparked a revolution."

Caplin, Andrew, Mark Dean, and John Leahy (2019) "Rational inattention, optimal consideration sets, and stochastic choice," *The Review of Economic Studies*, 86 (3), 1061–1094.

——— (2022) "Rationally inattentive behavior: Characterizing and generalizing shannon entropy," *Journal of Political Economy*, 130 (6), 1676–1715.

Carter, Richard and Steven Manaster (1990) "Initial public offerings and underwriter reputation," *the Journal of Finance*, 45 (4), 1045–1067.

Chemmanur, Thomas J. and Paolo Fulghieri (1994) "Investment bank reputation, information production, and financial intermediation," *The Journal of Finance*.

Chen, Hsuan-Chi and Jay R Ritter (2000) "The seven percent solution," *The Journal of Finance*, 55 (3), 1105–1131.

Clark, Robert, Jean-François Houde, and Jakub Kastl (2021) "The industrial organization of financial markets," in *Handbook of Industrial Organization*, 5, 427–520: Elsevier.

Crawford, Vincent P and Joel Sobel (1982) "Strategic information transmission," *Econometrica*, 50 (6), 1431–1451.

Diamond, Douglas W (1984) "Financial intermediation and delegated monitoring," *The Review of Economic Studies*, 51 (3), 393–414.

Eckbo, B Espen, Ronald W Masulis, and Øyvind Norli (2007) "Security offerings," *Handbook of empirical corporate finance*, 233–373.

Fang, Lily Hua (2005) "Investment bank reputation and the price and quality of underwriting services," *The Journal of Finance*, 60 (6), 2729–2761.

Ganguly, Chirantan and Indrajit Ray (2012) "Simple mediation in a cheap-talk game."

Goltsman, Maria, Johannes Hörner, Gregory Pavlov, and Francesco Squintani (2009) "Mediation, arbitration and negotiation," *Journal of Economic Theory*, 144 (4), 1397–1420.

Hall, Christopher, Sonny Allison, Justin Bastian, Eric DeJong, Nora Gibson, and David McShea (2016) *The IPO Handbook: A Guide for Entrepreneurs, Executives, Directors and Private Investors*: Smashwords Edition.

Ivanov, Maxim (2010) "Communication via a strategic mediator," *Journal of Economic Theory*, 145 (2), 869–884.

Kamenica, Emir (2019) "Bayesian persuasion and information design," *Annual Review of Economics*, 11, 249–272.

Kamenica, Emir and Matthew Gentzkow (2011) "Bayesian persuasion," *American Economic Review*, 101 (6), 2590–2615.

Katti, Supriya and BV Phani (2016) "Underpricing of initial public offerings: A literature review," *Universal Journal of Accounting and Finance*, 4 (2), 35–52.

Kreps, David M and Robert Wilson (1982) "Sequential equilibria," *Econometrica: Journal of the Econometric Society*, 863–894.

Kreutzkamp, Sophie (2022) "Endogenous information acquisition in cheap-talk games."

Lee, Gemma and Ronald W Masulis (2011) "Do more reputable financial institutions reduce earnings management by ipo issuers?" *Journal of Corporate Finance*, 17 (4), 982–1000.

Ljungqvist, Alexander (2007) "IPO underpricing," in *Handbook of Empirical Corporate Finance*, 375–422: Elsevier.

Lou, Yichuan (2022) "Sender-optimal learning and credible communication," *Available at SSRN 4298890*.

Lowry, Michelle, Roni Michaely, Ekaterina Volkova et al. (2017) "Initial public offerings: A synthesis of the literature and directions for future research," *Foundations and Trends® in Finance*, 11 (3-4), 154–320.

Lyu, Qianjun and Wing Suen (2022) "Information design in cheap talk," *arXiv preprint arXiv:2207.04929*.

Matějka, Filip and Alisdair McKay (2015) "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 105 (1), 272–298.

Pei, Harry Di (2015) "Communication with endogenous information acquisition," *Journal of Economic Theory*, 160, 132–149.

Ragupathy, MB (2011) "Initial public offering: A critical review of literature.," *IUP Journal of Behavioral Finance*, 8 (1).

Ramakrishnan, Ram TS and Anjan V Thakor (1984) "Information reliability and a theory of financial intermediation," *The Review of Economic Studies*, 51 (3), 415–432.

Ritter, Jay R (2003) "Investment banking and securities issuance," in *Handbook of the Economics of Finance*, 1, 255–306: Elsevier.

——— (2023) "Initial public offerings: updated statistics."

Ritter, Jay R and Ivo Welch (2002) "A review of ipo activity, pricing, and allocations," *The journal of Finance*, 57 (4), 1795–1828.

Rolfe, John and Peter Troob (2009) *Monkey business: Swinging through the Wall Street jungle*: Grand Central Publishing.

Thakor, Anjan V (2020) "Fintech and banking: What do we know?" *Journal of Financial Intermediation*, 41, 100833.

Zapechelnyuk, Andy (2022) "Sequential obfuscation and toxic argumentation," *Available at SSRN 4142663*.

A Proofs of statements in the main text

Proof of Proposition 1

On the contrary, suppose a contract t implements a success rate of 1 with equilibrium $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$. All posteriors q induced by $\mu_\sigma \circ \sigma$ with positive probability must satisfy $V_t^B(q) \geq 0$. It follows that $V_t^B(p) \geq 0$ because p is a weighted average of all posteriors induced by $\mu_\sigma \circ \sigma$. Observe that

$$V_t^I(p) + V_t^B(p) = (1-p)(-\kappa) + p(1-\kappa) < 0,$$

so $V_t^I(p) < 0$. We have

$$U_t^I(\sigma, \mu_\sigma, \alpha_\sigma) \leq \sum_{j=0}^{J-1} V_t^I(q_j) P_{\mu_\sigma \circ \sigma}(s_j) = V_t^I(p) < 0,$$

where $q_0, q_1, \ldots, q_{J-1}$ are the posteriors induced by $\mu_\sigma \circ \sigma$. The first inequality is from the definitions of U_t^I and V_t^I. The equality is due to the fact that V_t^I is affine. It follows that $U_t^I(\sigma, \mu_\sigma, \alpha_\sigma) < 0$. Thus, $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$ is not an equilibrium, a contradiction. ■

Proof of Lemma 1

From Proposition 1, we know that there exists no contract that implements a success rate of 1. Suppose t implements $\rho \in (0,1)$ with an equilibrium $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$ in an IPO mechanism. Let $\tilde{q}_0, \tilde{q}_1, \ldots, \tilde{q}_{J-1}$ denote the posteriors induced by $\tilde{\sigma} = \mu_\sigma \circ \sigma$. Define subsets of messages $M_0 = \{m \in \mathcal{M} : \alpha_\sigma(m) = 0\}$ and $M_1 = \{m \in \mathcal{M} : \alpha_\sigma(m) = 1\}$. Since $\rho \in (0,1)$, M_0 and M_1 are nonempty.

Step 1: Show that $\tilde{\sigma}$ is incentive-aligned. On the contrary, suppose that $\tilde{\sigma}$ is not incentive-aligned: there exists j such that $P_{\tilde{\sigma}}(s_j) > 0$ and $V_t^I(\tilde{q}_j) V_t^B(\tilde{q}_j) < 0$. First, suppose $V_t^I(\tilde{q}_j) < 0$ and $V_t^B(\tilde{q}_j) > 0$. Since α_σ is a best response to μ_σ, $\alpha_\sigma(s_j) = 1$. Then μ_σ is not a best response to α_σ, as decreasing $\mu_\sigma(s_j|s)$ and increasing $\mu_\sigma(s'|s)$ for some $s' \in M_0$ by the same amount makes the intermediary strictly better off, a contradiction. Second, suppose $V_t^I(\tilde{q}_j) > 0$ and $V_t^B(\tilde{q}_j) < 0$. Since α_σ is a best response to μ_σ, $\alpha_\sigma(s_j) = 0$. Then μ_σ is not a best response to α_σ, as decreasing $\mu_\sigma(s_j|s)$ and increasing $\mu_\sigma(s'|s)$ for some $s' \in M_1$ by the same amount makes the intermediary strictly better off, a contradiction.

Step 2: Construct a binary due diligence structure σ'. Let $\sigma' \in \Sigma$ such that, for all $\omega \in \Omega$,

$$\sigma'(s_j|\omega) = \sum_{s \in M_j} \tilde{\sigma}(s|\omega), \quad \text{for } j = 0, 1,$$

and $\sigma'(s_j|\omega) = 0$ for $j = 2, 3, \ldots, J - 1$. The constructed due diligence structure σ' induces two posteriors with positive probabilities: for $j = 0, 1$,

$$q'_j = \frac{p\sigma'(s_j|\omega)}{P_{\sigma'}(s_j)} = \frac{\sum_{s \in M_j} p\tilde{\sigma}(s|\omega)}{P_{\sigma'}(s_j)},$$

with probabilities

$$\begin{aligned} P_{\sigma'}(s_j) &= (1-p)\sigma'(s_j|0) + p\sigma'(s_j|1) \\ &= (1-p) \sum_{s \in M_j} \tilde{\sigma}(s|0) + p \sum_{s \in M_j} \tilde{\sigma}(s_j|1). \end{aligned}$$

Observe that, for all $j = 0, 1$, q'_j is a weighted average of posteriors $\tilde{q}_k$ whose signals belong to M_j. Namely, for all $j = 0, 1$,

$$q'_j = \sum_{s_k \in M_j} \tilde{q}_k \frac{P_{\sigma'}(s_k)}{P_{\sigma'}(s_j)}.$$

We show that the constructed posteriors q'_0 and q'_1 are distinct in Step 4.

Step 3: Show that σ' is incentive-aligned. Since $\tilde{\sigma}$ is incentive-aligned, $V_t^I(\tilde{q}_k)V_t^B(\tilde{q}_k) \geq 0$ for all posteriors $\tilde{q}_k$ with positive probabilities. Observe that, by construction, $\alpha_\sigma(s_k) = 0$ for all $s_k \in M_0$, so $V_t^B(\tilde{q}_k) \leq 0$ for all k such that $s_k \in M_0$. Then $V_t^I(\tilde{q}_k) \leq 0$ for all k such that $s_k \in M_0$. Similarly, by construction, $\alpha_\sigma(s_k) = 1$ for all $s_k \in M_1$, so $V_t^B(\tilde{q}_k) \geq 0$ for all k such that $s_k \in M_1$. It follows that $V_t^I(\tilde{q}_k) \geq 0$ for all k such that $s_k \in M_1$.

Observe that V_t^I and V_t^B are affine. Since q'_0 is a weighted average of posteriors $\tilde{q}_k$ for all k such that $s_k \in M_0$, $V_t^I(q'_0) \leq 0$ and $V_t^B(q'_0) \leq 0$. Similarly, since q'_1 is a weighted average of posteriors $\tilde{q}_k$ for all k such that $s_k \in M_1$, $V_t^I(q'_1) \geq 0$ and $V_t^B(q'_1) \geq 0$. Therefore, $V_t^I(q'_j)V_t^B(q'_j) \geq 0$ for all $j = 0, 1$.

Step 4: Show that $q'_0 \neq q'_1$. On the contrary, suppose $q'_0 = q'_1$. Recall from Step 3 that, for each $j = 0, 1$, q'_j is an weighted average of all $\tilde{q}_k$ such that $s_k \in M_j$. Also recall that $V_t^B(\tilde{q}_k) \leq 0$ for all $\tilde{q}_k$ such that $s_k \in M_0$. Similarly, $V_t^B(\tilde{q}_k) \geq 0$ for all $\tilde{q}_k$ such that $s_k \in M_1$. Since V_t^B is affine, these inequalities imply that $\tilde{q}_k = p$ for all $k = 0, 1, \ldots, J-1$. Then $\rho = 0$ or $\rho = 1$, a contradiction.

Step 5: Characterize $\mu_{\sigma'}$ and $\alpha_{\sigma'}$. Consider $\mu_{\sigma'}$ and $\alpha_{\sigma'}$. Because $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$ is an equilibrium, we know that $\mu_{\sigma'}$ and $\alpha_{\sigma'}$ are mutual best responses given σ' and that the pair is not Pareto-dominated by other such pairs of mutual best responses. Let $M_0' = \{m \in \mathcal{M} : \alpha_{\sigma'}(m) = 0\}$ and $M_1' = \{m \in \mathcal{M} : \alpha_{\sigma'}(m) = 1\}$. We claim that M_0' and M_1' are nonempty, and $\mu_{\sigma'}$ satisfies

$$\sum_{m \in M_0} \mu_{\sigma'}(m|s_0) = 1 \quad \text{and} \quad \sum_{m \in M_1} \mu_{\sigma'}(m|s_1) = 1, \tag{6}$$

which implies that $\mu_{\sigma'}$ is fully revealing. To show that M_0' and M_1' are nonempty, suppose that one of them is empty.

First, suppose M_0' is empty. Then $\alpha_{\sigma'}(m) = 1$ for all $m \in \mathcal{M}$, implying that $V_t^B(q_0') \geq 0$ and $V_t^B(q_1') \geq 0$. Let q'' denote the minimum of the two posteriors q_0' and q_1'. Then $q'' < p$ because $q_0' \neq q_1'$. From the definitions of V_t^I and V_t^B, $V_t^B(q'') \geq 0$ implies that

$$V_t^I(q'') \leq (1 - q'')(-\kappa) + q''(1 - \kappa).$$

The right-hand side of this inequality is strictly negative because $q'' < p < \kappa$. Then $V_t^I(q'') < 0$ whereas $V_t^B(q'') \geq 0$. Recall that σ' is incentive-aligned, so $V_t^I(q'')V_t^B(q'') \geq 0$ in particular. Then $V_t^B(q'') = 0$. However, $(\mu_{\sigma'}, \alpha_{\sigma'})$ is Pareto-dominated by a pair (μ', α') where

$$\begin{aligned} \mu'(s|s) &= 1, \quad \text{for all } s \in S, \\ \alpha'(m) &= 1, \quad \text{if and only if } m = s_1. \end{aligned}$$

Then $\alpha_{\sigma'}$ is not a best response, a contradiction. Second, suppose M_1' is empty. Then both the intermediary and the buyer earns zero payoffs at this stage, whereas (μ', α') earns the buyer a nonnegative payoff and the intermediary a positive payoff. Then $(\mu_{\sigma'}, \alpha_{\sigma'})$ is Pareto-dominated by (μ', α'), a contradiction.

To show equation (6), suppose that this equation is not true: there exists $m \notin M_j$ such that $\mu_{\sigma'}(m|s_j) > 0$ for some $j = \{0, 1\}$. Then $\mu_{\sigma'}$ is not a best response because the intermediary strictly gains by reducing $\mu_{\sigma'}(m|s_j) > 0$ and increasing $\mu_{\sigma'}(m'|s_d)$ for some $m' \in M_j$ by the same amount.

Conversely, let us show that if M_0' and M_1' are nonempty and equation (6) is true, then $\mu_{\sigma'}$ and $\alpha_{\sigma'}$ are mutual best responses and the pair is not Pareto-dominated. Since $\mu_{\sigma'}$ is fully revealing and σ' is incentive-aligned, $\mu_{\sigma'}$ and $\alpha_{\sigma'}$ are mutual best responses. Moreover, because $\mu_{\sigma'}$ is fully revealing, no other mutually best-responding pair gives a higher expected payoff to the intermediary or the buyer.

Step 6: Show that $U_t^I(\sigma', \mu_{\sigma'}, \alpha_{\sigma'}) = U_t^I(\sigma, \mu_\sigma, \alpha_\sigma)$. By the incentive alignment of σ and the affineness of V_t^I,

$$\sum_{j=0}^{1} P_{\sigma'}(s_j) V_t^I(q_j) = \sum_{j=0}^{J-1} P_{\tilde{\sigma}}(s_j) V_t^I(q_j) \tag{7}$$

Moreover, by construction, σ' is a garbling of $\tilde{\sigma}$ and $\tilde{\sigma}$ is a garbling of σ. Thus,

$$C(\sigma') \leq C(\tilde{\sigma}) \leq C(\sigma).$$

Observe that $U_t^I(\sigma', \mu_{\sigma'}, \alpha_{\sigma'})$ equals the left-hand side of (7) minus the cost $c(\sigma')$. Also, $U_t^I(\sigma, \mu_\sigma \alpha_{\sigma'})$ equals the right-hand side of (7) minus the cost $c(\sigma)$. Then $U_t^I(\sigma', \mu_{\sigma'}, \alpha_{\sigma'}) \geq U_t^I(\sigma, \mu_\sigma, \alpha_\sigma)$. Since σ is a best response to $(\boldsymbol{\mu}, \boldsymbol{\sigma})$, $U_t^I(\sigma', \mu_{\sigma'}, \alpha_{\sigma'}) \leq U_t^I(\sigma, \mu_\sigma, \alpha_\sigma)$. It follows that $U_t^I(\sigma', \mu_{\sigma'}, \alpha_{\sigma'}) = U_t^I(\sigma, \mu_\sigma, \alpha_\sigma)$.

Step 7: Show that $\sigma = \sigma'$. Observe that, if σ is binary, we have $\sigma = \sigma'$ by construction. Suppose σ is not binary, inducing more than two posteriors with positive probabilities. This implies that σ' is strictly less Blackwell-informative than σ, so $C(\sigma') < C(\sigma)$. This implies that $U_t^I(\sigma, \mu_\sigma, \alpha_\sigma) < U_t^I(\sigma', \mu_{\sigma'}, \alpha_{\sigma'})$, which is impossible. So $\sigma = \sigma'$. Observe from Step 3 that σ' is incentive-aligned, so σ is incentive-aligned. Observe also from Step 5 that $\mu_{\sigma'}$ is fully revealing, thus μ_σ is fully revealing. Therefore, the equilibrium $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$ satisfies fully revealing messages, incentive alignment, and binary due diligence. ■

Proof of Theorem 1

The following lemma establishes that the intermediary' incentives are balanced, in the sense that she prefers that the buyer does not invest when her posterior is pessimistic and prefers that the buyer does invest when her posterior is optimistic.

Lemma A.1. *Suppose a contract t implements a success rate $\rho \in (0, 1)$ with some equilibrium. Let $(\ell, r) \in [0, p) \times (p, 1]$ denote the pair of posteriors induced by the* *equilibrium due diligence. The intermediary's incentives are balanced:* $V_t^I(\ell) \leq 0$ *and* $V_t^I(r) > 0$*.*

Proof. Let $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$ denote the said equilibrium, and suppose that the intermediary's incentives are not balanced. This implies that $V_t^I(\ell)V_t^I(r) \geq 0$. Consider the following two cases. First, suppose $V_t^I(\ell) \leq 0$ and $V_t^I(r) \leq 0$. Then the intermediary's profit is negative, so $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$ cannot be an equilibrium, a contradiction. Second, suppose $V_t^I(\ell) > 0$. Then $V_t^B < 0$, because $V_t^I(\ell) + V_t^B(\ell) < V_t^I(p) + V_t^B(p) < 0$. Then σ is not aligned, a contradiction. Therefore, $V_t^I(\ell) \leq 0$ and $V_t^I(r) > 0$.

We now prove Theorem 1. Let $(\ell, r) \in [0, p) \times (p, 1]$ denote the pair of posteriors with positive probability induced by the equilibrium due diligence structure. By the definitions of V_t^I and V_t^B, for all $q \in [0, 1]$,

$$(t_1 - t_0)q \geq -t_0 \quad \text{if and only if} \quad V_t^I(q) \geq 0, \text{ and} \tag{8}$$

$$(1 - t_1 + t_0)q \geq \kappa + t_0 \quad \text{if and only if} \quad V_t^B(q) \geq 0, \tag{9}$$

which also hold with strict inequalities. By Lemma 1, (ℓ, r) satisfies $V_t^I(q)V_t^B(q) \geq 0$ for both $q \in \{\ell, r\}$ (*incentive alignment*). By Lemma A.1, (ℓ, r) satisfies $V_t^I(\ell) \leq 0$ and $V_t^I(r) > 0$ (*incentive balancedness*). We show that all of the following four cases violate these conditions.

Case 1: Suppose $t_0 < -\kappa$. First, if $1 - t_1 + t_0 \geq 0$, the strict inequality (9) implies that $V_t^B(q) > 0$ for all $q \in [0, 1]$. In particular, $V_t^B(\ell) > 0$ and $V_t^B(r) > 0$. Incentive alignment implies that $V_t^I(\ell) \geq 0$ and $V_t^I(r) \geq 0$, hence (ℓ, r) does not satisfy incentive balancedness. Second, if $1 - t_1 + t_0 < 0$, then $t_1 - t_0 > 1$. From (8)–(9), we have

$$q \geq \frac{-t_0}{t_1 - t_0} \quad \text{if and only if} \quad V_t^I(q) \geq 0, \text{ and}$$

$$q \leq \frac{\kappa + t_0}{1 - t_1 + t_0} \quad \text{if and only if} \quad V_t^B(q) \geq 0.$$

The incentive alignment requires that both posteriors ℓ and r lie in the interval $\left[\frac{\kappa+t_0}{1-t_1+t_0}, \frac{-t_0}{t_1-t_0}\right]$. This implies that $V_t^I(\ell) \geq 0$ and $V_t^I(r) \geq 0$, violating incentive balancedness.

Case 2: Suppose $t_0 \geq 0$. First, if $t_1 \geq t_0$, the inequality (8) implies that $V_t^I(q) \geq 0$ for all $q \in [0, 1]$. In particular, $V_t^I(\ell) \geq 0$ and $V_t^I(r) \geq 0$, violating the incentive balancedness. Second, if $t_1 < t_0$, we have $1 - t_1 + t_0 > 1$. From (8)–(9), we have

$$q \le \frac{-t_0}{t_1 - t_0} \quad \text{if and only if} \quad V_t^I(q) \ge 0, \text{ and}$$
$$q \ge \frac{\kappa + t_0}{1 - t_1 + t_0} \quad \text{if and only if} \quad V_t^B(q) \ge 0.$$

As in Case 1, incentive alignment requires that both posteriors ℓ and r lie on the interval $\left[\frac{\kappa+t_0}{1-t_1+t_0}, \frac{-t_0}{t_1-t_0}\right]$. Thus, $V_t^I(\ell) \le 0$ and $V_t^I(r) \le 0$, violating incentive balancedness.

Case 3: Suppose $t_1 \le 0$. We know from Cases 1–2 that $t_0 \in [-\kappa, 0)$, so $V_t^I(q) = (1-q)t_0 + qt_1 \le 0$ for all $q \in [0,1]$. This implies that $V_t^I(\ell) \le 0$ and $V_t^I(r) \le 0$, which violates incentive balancedness.

Case 4: Suppose $t_1 > 1-\kappa$. As before, we know from Cases 1–2 that $t_0 \in [-\kappa, 0)$. Consequently, $V_t^B(q) = (1-q)(-\kappa - t_0) + q(1-\kappa-t_1) < 0$ for all $q \in (0,1]$. In particular, $V_t^B(r) < 0$. Incentive alignment requires that $V_t^I(r) \le 0$, which violates incentive balancedness.

We know from Cases 1–2 that $t_0 \in [-\kappa, 0)$. In addition, we know from Cases 3–4 that $t_1 \in (0, 1-\kappa]$. These results imply that a contract that implements a success rate $\rho > 0$ satisfies $(t_0, t_1) \in [-\kappa, 0) \times (0, 1-\kappa]$. ■

Proof of Theorem 2

We introduce further notation and establish several lemmas. For any $\sigma \in \widehat{\Sigma}$, let σ_ω denote $\sigma(s_1|\omega)$, the conditional probability of the optimistic signal. Thus, $\sigma_0 = \sigma(s_1|0)$ and $\sigma_1 = \sigma(s_1|1)$. Consider a metric $d : \widehat{\Sigma} \times \widehat{\Sigma} \longrightarrow \mathbb{R}$ such that

$$d(\sigma, \sigma') = \sqrt{(\sigma_0 - \sigma_0')^2 + (\sigma_1 - \sigma_1')^2}, \tag{10}$$

which is the Euclidean distance between (σ_0, σ_1) and (σ_0', σ_1') in $\mathbb{R}^2$. For the rest of the proof, we use this metric for the distance between any two binary due diligence structures.

For any contract $t \in T$, define $\Lambda(t)$ as the set of all binary due diligence structures that are incentive-aligned under the contract t. Define $Q(t)$ as the set of all pairs of posteriors $(\ell, r) \in [0,p) \times (p,1]$ induced by some $\sigma \in \Lambda(t)$. The mapping from the set $\Lambda(t)$ of incentive-aligned due diligence structures and the set $Q(t)$ of their induced posteriors is invertible, resulting in the following lemma.

Lemma A.2. *Let $t \in T$ and $\sigma \in \Lambda(t)$ be given. Let $(\ell, r) \in [0,p) \times (p,1]$ denote the pair of posteriors induced by σ. Then $\sigma \in \operatorname{int}(\Lambda(t))$ if and only if $(\ell, r) \in \operatorname{int}(Q(t))$.*

Proof. Let us write $\sigma_\omega = \sigma(s_1|\omega)$ for both $\omega \in \Omega$. Because (ℓ, r) is the pair of posteriors induced by σ, it satisfies

$$\ell = \frac{(1-\sigma_1)p}{(1-\sigma_0)(1-p) + (1-\sigma_1)p} \quad \text{and} \quad r = \frac{\sigma_1 p}{\sigma_0(1-p) + \sigma_1 p}. \tag{11}$$

Equivalently, we have

$$\sigma_0 = \frac{1-r}{1-p} \cdot \frac{p-\ell}{r-\ell} \quad \text{and} \quad \sigma_1 = \frac{r}{p} \cdot \frac{p-\ell}{r-\ell}.$$

So there exists a continuous and invertible map $f : \Lambda(t) \longrightarrow Q(t)$ given by (11). As a result, $(\ell, r) \in \operatorname{int} Q(t)$ implies $f^{-1}(\ell, r) \in \operatorname{int} \Lambda(t)$. Conversely, $\sigma \in \operatorname{int}(\Lambda(t))$ implies $f(\sigma) \in (Q(t))$. This completes the proof of Lemma A.2.

Recall from Lemma A.1 that if $(\sigma, \boldsymbol{\mu}, \boldsymbol{\alpha})$ is an equilibrium with a positive success rate, σ is incentive-aligned. The following lemma generalizes this result to any binary due diligence structure σ not necessarily on the equilibrium path.

Lemma A.3. *Suppose a message rule μ and an action rule α are mutual best responses given $(t, \sigma) \in T \times \widehat{\Sigma}$. Then $U_t^I(\sigma, \mu, \alpha) \geq 0$ only if $\sigma \in \Lambda(t)$.*

Proof. Let $(\ell, r) \in [0,p) \times (p,1]$ denote the pair of posteriors induced by σ. Let $\tilde{\sigma} = \mu \circ \sigma$ so that, for all $\omega \in \Omega$ and $\tilde{s} \in \mathcal{S}$,

$$\tilde{\sigma}(\tilde{s}|\omega) = \sum_{s \in S} \mu(\tilde{s}|s)\sigma(s|\omega). \tag{12}$$

Let $q_0, q_1, \ldots, q_{J-1}$ denote the posteriors induced by $\tilde{\sigma}$. Let $M_0 = \{m \in \mathcal{M} : \alpha(m) = 0\}$ and $M_1 = \{m \in \mathcal{M} : \alpha(m) = 1\}$.

Step 1. Show that the success rate lies within the open interval $(0,1)$. Let $\rho = \frac{1}{\kappa} U_t^S(\sigma, \mu, \alpha)$ denote the success rate under σ, μ, and α. Observe that, because $\sigma \in \widehat{\Sigma}$ and thus $C(\sigma) > 0$, the intermediary's expected revenue is strictly positive. That is, $U_t^I(\sigma, \mu, \alpha) + C(\sigma) > 0$. First, suppose $\rho = 0$. Then $m \in M_0$ for all messages $m \in M$ such that $\pi(m|\tilde{\sigma}) > 0$. Then the intermediary's expected revenue is zero, a contradiction. Second, suppose $\rho = 1$. Since α is a best response, this implies that $V_t^B(q) \geq 0$ for all posteriors induced by $\tilde{\sigma}$ with positive probability. As in the proof of Proposition 1, this implies that $U_t^I(\sigma, \mu, \alpha) < 0$, a contradiction.

Step 2. Show that $\tilde{\sigma}$ is incentive-aligned. We proceed in the same way as in Step 1 of the proof of Lemma 1. Observe that neither M_0 nor M_1 are empty because $U_t^S(\sigma, \mu, \alpha) \in (0, 1)$. Suppose $\tilde{\sigma}$ is not incentive-aligned. Then there exists a posterior q_j induced by $\tilde{\sigma}$ with positive probability such that $V_t^I(q_j)V_t^B(q_j) < 0$. First, suppose $V_t^I(q_j) > 0$ and $V_t^B(q_j) < 0$. Because α is a best response, $\alpha(s_j) = 0$. It follows that μ is not a best response, a contradiction. Second, suppose $V_t^I(q_j) < 0$ and $V_t^B(q_j) > 0$. Because α is a best response, $\alpha(s_j) = 1$. Then μ is not a best response, a contradiction.

Step 3. Show that σ is incentive-aligned. Since $\tilde{\sigma}$ is a garbling of σ, $\ell \leq \min\{q_0, q_1, \ldots, q_{J-1}\}$ and $r \geq \max\{q_0, q_1, \ldots, q_{J-1}\}$. Since $\tilde{\sigma}$ is incentive-aligned and the intermediary's revenue is strictly positive, $V_t^I(q_j) \leq 0$ for some j and $V_t^I(q_k) > 0$ for some k. Since, from Theorem 1, V_t^I and V_t^B are both increasing functions, we have $V_t^I(\ell) < 0$, $V_t^I(r) > 0$, $V_t^B(\ell) \leq 0$, and $V_t^B(r) \geq 0$. Therefore, $V_t^I(\ell)V_t^B(\ell) \geq 0$ and $V_t^I(r)V_t^B(r) \geq 0$, meaning that $\sigma \in \Lambda(t)$. This complete the proof of Lemma A.3.

The next lemma shows that the intermediary' expected payoffs are smooth and strictly concave. A critical requirement for this result is that the collection of message rules and action rules are not Pareto-dominated. This condition ensures that the equilibrium of the talking-stage subgame is the informative one whenever it is available, leading to truthful revelation and mutually preferred actions.

Lemma A.4. *Let a contract $t \in T$ be given. Let $(\boldsymbol{\mu}, \boldsymbol{\alpha}) = \{(\mu_\sigma, \alpha_\sigma)\}_{\sigma \in \Sigma}$ be given such that it is a mutual best response and is not Pareto-dominated. The map $\sigma \mapsto U_t^I(\sigma, \mu_\sigma, \alpha_\sigma)$ with domain $\Lambda(t)$ is continuously differentiable and strictly concave.*

Proof. To streamline the notation, define $\pi(\omega)$ as the prior probability of the state ω, so that $\pi(0) = 1 - p$ and $\pi(1) = p$. Observe that

$$U_t^I(\sigma, \mu_\sigma, \alpha_\sigma) = \sum_{\omega \in \Omega} \sum_{s \in \mathcal{S}} \sum_{m \in \mathcal{M}} t_\omega \alpha_\sigma(m) \mu_\sigma(m|s) \sigma(s|\omega) \pi(\omega) - C(\sigma). \tag{13}$$

Let $M_j = \{m \in \mathcal{M} : \alpha_\sigma(m) = j\}$ for all $j = 0, 1$. As in Step 5 of Proof of Lemma 1, the fact that $\sigma \in \Lambda(t)$, $(\mu_\sigma, \alpha_\sigma)$ is a mutual best response and is not Pareto-dominated implies that $\mu(m|s_j) = 1$ for all $m \in M_j$ for all $j = 0, 1$. It follows that

$$U_t^I(\sigma, \mu_\sigma, \alpha_\sigma) = \sum_{\omega \in \Omega} \sum_{a \in A} t_\omega a \sigma(s_a|\omega)\pi(\omega) - C(\sigma). \tag{14}$$

The first term on the right-hand side of the above equation is affine in σ. The second term, $C(\sigma)$, is continuously differentiable and strictly convex in σ. Therefore, $U_t^I(\sigma, \mu_\sigma, \alpha_\sigma)$ is strictly concave in σ on the domain $\Lambda(t)$. This completes the proof of Lemma A.4.

Under any seller-optimal constraint, the intermediary's equilibrium choice of the due diligence structure is not an interior point, as she faces incentive alignment as a binding constraint:

Lemma A.5. *Suppose a contract $t^* \in T$ is seller-optimal, implementing a maximal success rate $\rho^* > 0$ with equilibrium $(\sigma^*, \boldsymbol{\mu}, \boldsymbol{\alpha})$. Then $\sigma^* \notin \operatorname{int} \Lambda(t^*)$.*

Proof. Let (ℓ^*, r^*) denote the pair of posteriors induced by σ^*. Since $(\sigma^*, \boldsymbol{\mu}, \boldsymbol{\alpha})$ is an equilibrium, (ℓ^*, r^*) maximizes the intermediary's expected payoffs given $(\boldsymbol{\mu}, \boldsymbol{\alpha})$. That is,

$$(\ell^*, r^*) \in \underset{(\ell, r) \in Q(t^*)}{\operatorname{argmax}} \frac{p-\ell}{r-\ell}[(1-r)t_0^* + rt_1^*] - \left(1 - \frac{p-\ell}{r-\ell}\right) c(\ell) - \frac{p-\ell}{r-\ell} c(r).$$

Contrary to the lemma's statement, suppose $\sigma^* \in \operatorname{int} \Lambda(t^*)$. Then by Lemma A.2, $(\ell^*, r^*) \in \operatorname{int} Q(t^*)$. This implies that $(\ell, r, t_0, t_1) = (\ell^*, r^*, t_0^*, t_1^*)$ satisfies the necessary first-order conditions

$$c'(r) - c'(\ell) - (t_1 - t_0) = 0, \text{ and} \tag{15}$$

$$(r-\ell)c'(r) - [c(r) - c(\ell)] - (r-\ell)(t_1 - t_0) + (1-r)t_0 + rt_1 = 0. \tag{16}$$

Let functions $f(\ell, r, t_0, t_1)$ and $g(\ell, r, t_0, t_1)$ denote the left-hand sides of the equations (15) and (16), respectively. Observe that the Jacobian matrix $\partial(f, g)/\partial(\ell, r)$ evaluated at $(\ell^*, r^*, t_0^*, t_1^*)$ yields

$$Y = \begin{bmatrix} \partial f/\partial \ell & \partial f/\partial r \\ \partial g/\partial \ell & \partial g/\partial r \end{bmatrix}_{(\ell^*, r^*, t_0^*, t_1^*)} = \begin{bmatrix} -c''(\ell^*) & c''(r^*) \\ 0 & (r^* - \ell^*)c''(r^*) \end{bmatrix}. \tag{17}$$

Since c is strictly convex, the Jacobian determinant $-(r^* - \ell^*)c''(\ell^*)c''(r^*)$ is nonzero. Observe also that another Jacobian matrix $\partial(f, g)/\partial(t_0, t_1)$ evaluated at the same point yields

$$Z = \begin{bmatrix} \partial f/\partial t_0 & \partial f/\partial t_1 \\ \partial g/\partial t_0 & \partial g/\partial t_1 \end{bmatrix}_{(\ell^*, r^*, t_0^*, t_1^*)} = \begin{bmatrix} 1 & -1 \\ 1-\ell^* & \ell^* \end{bmatrix}.$$

By the Implicit Function Theorem, there exists a unique pair $(\ell, r) = (\hat{\ell}(t_0, t_1), \hat{r}(t_0, t_1))$ for any (t_0, t_1) that satisfies equations (15)–(16) around a small neighborhood of (ℓ^*, r^*) and (t_0^*, t_1^*). Moreover, the Jacobian matrix $\partial(\hat{\ell}, \hat{r})/\partial(t_0, t_1)$ evaluated at $(\ell^*, r^*, t_0^*, t_1^*)$ is

$$-Y^{-1}Z = \begin{bmatrix} -\frac{1-\ell^*}{(r^*-\ell^*)c''(\ell^*)} + \frac{1}{c''(\ell^*)} & -\frac{\ell^*}{(r^*-\ell^*)c''(\ell^*)} - \frac{1}{c''(\ell^*)} \\ -\frac{1-\ell^*}{(r^*-\ell^*)c''(r^*)} & -\frac{\ell^*}{(r^*-\ell^*)c''(r^*)} \end{bmatrix}.$$

From the above Jacobian matrix $\partial(\hat{\ell}, \hat{r})/\partial(t_0, t_1)$, we have

$$\begin{aligned} \frac{\partial \hat{\ell}}{\partial t_0} + \frac{\partial \hat{\ell}}{\partial t_1} &= -\frac{1}{(r^* - \ell^*)c''(\ell^*)} < 0, \text{ and} \\ \frac{\partial \hat{r}}{\partial t_0} + \frac{\partial \hat{r}}{\partial t_1} &= -\frac{1}{(r^* - \ell^*)c''(r^*)} < 0. \end{aligned}$$

Thus, a small increase in both components of the contract (t_0^*, t_1^*) by the same size leads to strict decreases in both equilibrium posteriors ℓ^* and r^*, which strictly increases the success rate $\rho^* = \frac{p-\ell^*}{r^*-\ell^*}$. Therefore, t^* is not optimal, a contradiction. This completes the proof of Lemma A.5.

We now prove Theorem 2. Let $t^* = (t_0^*, t_1^*) \in T$ denote a seller-optimal contract implementing the maximal success rate ρ^* with equilibrium $(\sigma^*, \boldsymbol{\mu}, \boldsymbol{\alpha})$. Let $(\ell^*, r^*) \in [0, p) \times (p, 1]$ denote the pair of posteriors generated by σ^*. Because σ^* is the intermediary's best response,

$$\begin{aligned} (\ell^*, r^*) \in \underset{(\ell, r)}{\operatorname{argmax}} \; & \frac{p-\ell}{r-\ell} \cdot [(1-r)t_0^* + rt_1^*] - \left(1 - \frac{p-\ell}{r-\ell}\right) c(\ell) - \frac{p-\ell}{r-\ell} c(r) \\ \text{subject to} \quad & (\ell, r) \in [0, p) \times (p, 1], \\ & V_{t^*}^I(\ell) \leq 0, \text{ and} \\ & V_{t^*}^B(r) \geq 0. \end{aligned}$$

Contrary to the theorem's statement, suppose $V_{t^*}^B(r) > 0$. Because $c'(\ell) \longrightarrow -\infty$ as $\ell \longrightarrow 0$ and $c'(r) \longrightarrow \infty$ as $r \longrightarrow 1$, we know that $\ell^* \neq 0$ and $r^* \neq 1$. It follows that the only constraint that possibly binds in the above maximization problem is

$V_t^I(\ell) \leq 0$. Because (ℓ^*, r^*) is not an interior point of the constraint set (Lemma A.5), we have $V_{t^*}^I(\ell^*) = 0$. After letting the Lagrange multipliers be zero for the other constraints, the first-order necessary conditions for (ℓ^*, r^*) yield

$$(1-\ell^*)t_0^* + \ell^* t_1^* = 0, \text{ and}$$
$$(r^* - \ell^*)c'(r^*) - [c(r^*) - c(\ell^*)] - (r^* - \ell^*)(t_1^* - t_0^*) + (1 - r^*)t_0^* + r^* t_1^* = 0.$$

These imply that $(r^* - \ell^*)c'(r^*) = c(r^*) - c(\ell^*)$. Because c is strictly convex, it follows that $\ell^* = r^* = p^*$, a contradiction. ■

Proof of Theorem 3

Suppose a contract t^* implements a positive success rate ρ^* in an equilibrium $(\sigma^*, \boldsymbol{\mu}^*, \boldsymbol{\alpha}^*)$. From the definition of an optimal contract and Lemma 1, we know that the contract t^* is optimal if and only if

$$(t^*, \sigma^*) \in \underset{(t,\sigma) \in T \times \widehat{\Sigma}}{\operatorname{argmax}} \; U_t^S(\sigma, \mu_\sigma, \alpha_\sigma)$$

subject to

$$\sigma^* \in \underset{\sigma \in \Lambda(t)}{\operatorname{argmax}} \; U_t^I(\sigma, \mu_\sigma, \alpha_\sigma). \tag{18}$$

Equivalently, with a change of variables from σ^* to (ℓ^*, r^*), the contract t^* is optimal if and only if

$$(t^*, \ell^*, r^*) \in \underset{(t,\ell',r') \in T \times [0,p) \times (p,1]}{\operatorname{argmax}} \; \frac{p - \ell'}{r' - \ell'} \tag{19}$$

subject to

$$(\ell', r') \in \underset{(\ell,r) \in Q(t)}{\operatorname{argmax}} \; \frac{p - \ell}{r - \ell} \cdot [(1-r)t_0^* + r t_1^*] - \left(1 - \frac{p-\ell}{r-\ell}\right) c(\ell) - \frac{p-\ell}{r-\ell} c(r). \tag{20}$$

From Lemma A.4, $U_t^I(\sigma, \mu_\sigma, \alpha_\sigma)$ is continuously differentiable and strictly concave in σ. Thus, the first-order necessary conditions for the problem (18) are also sufficient. This implies that the first-order necessary conditions for the equivalent problem (20) are also sufficient.

Observe that the constraint $(\ell, r) \in Q(t)$ in (20) is equivalent to the combined condition

$$(\ell, r) \in [0, p) \times (p, 1], \tag{21}$$

$$(1-r)t_0 + rt_1 \geq r - \kappa, \text{ and} \tag{22}$$

$$(1-\ell)t_0 + \ell t_1 \leq 0. \tag{23}$$

From Theorem 2, the condition (22) holds with equality. With the constraints (21)–(23), the first-order conditions of (20) are

$$c(r) - c(\ell) - (r-\ell)c'(\ell) = (1-r)t_0 + rt_1, \text{ and} \tag{24}$$

$$c(r) - c(\ell) - (r-\ell)c'(r) \leq (1-\ell)t_0 + \ell t_1. \tag{25}$$

By substituting the condition (22) with equality into the right-hand side of (24), we obtain

$$c(r) - c(\ell) - (r-\ell)c'(\ell) = r - \kappa. \tag{26}$$

This result implies that the condition (19)–(20) is equivalent to the combined condition (22) with equality, (23), (25), and

$$(\ell^*, r^*) \in \underset{(\ell,r)\in[0,p)\times(p,1]}{\operatorname{argmax}} \frac{p-\ell}{r-\ell} \qquad \text{subject to (26)},$$

which is the desired result. ■

Proof of Corollary 1

Let $(t_0, t_1) \in T$ denote a seller-optimal contract that implements the maximal success rate. Let $(\ell, r) \in [0, p) \times (p, 1]$ denote the pair of seller-optimal posteriors.

From the definition, the intermediary's ROI is $-t_1/t_0$ whereas the buyer's ROI is $(1 - \kappa - t_1)/(\kappa + t_0)$. From Theorem 3, we have

$$(1-r)t_0 + rt_1 = r - \kappa, \tag{27}$$

which implies $(1-\kappa)t_0 + \kappa t_1 > 0$. It follows that

$$\frac{t_1}{-t_0} > \frac{1-\kappa-t_1}{\kappa+t_0}.$$

In addition, if the contract (t_0, t_1) is the minimum-variance seller-optimal contract, we have

$$(1-\ell)t_0 + \ell t_1 = 0. \tag{28}$$

Equation (28) implies that the intermediary's ROI is $\frac{1-\ell}{\ell}$. Equation (27) implies that the buyer's ROI is $\frac{1-r}{r}$. ■

B Existence of a seller-optimal pair of posteriors

In this section, we provide conditions for the existence of a seller-optimal pair of posteriors. Recall from Definition 5 that a seller-optimal pair of posteriors is a solution to the problem

$$\max_{(\ell,r)\in[0,p)\times(p,1]} \frac{p-\ell}{r-\ell} \quad \text{subject to} \quad c(r)-c(\ell)-(r-\ell)c'(\ell) = r-\kappa. \tag{29}$$

Let a function $h : [0,p] \times [p,1] \longrightarrow [-\infty,\infty)$ be defined as

$$h(\ell,r) = r - \kappa - [c(r)-c(\ell)] + (r-\ell)\lim_{q\to\ell} c'(q),$$

so that the implementability constraint in (29) is equivalent to $h(\ell,r) = 0$. Let

$$\overline{r} = (c')^{-1}\left(1 + c'(p)\right),$$

where the inverse $(c')^{-1}$ is well-defined as c is strictly convex.

The following provides a necessary and sufficient condition for the solution's existence.

Proposition B.1. *A seller-optimal pair of posteriors exists if and only if* $h(p,\overline{r}) > 0$.

Proof. To show the "if" part of the statement, suppose $h(p,\overline{r}) > 0$. Since $h(0,\overline{r}) = -\infty$ and h is continuous, the level set

$$L = \{(\ell,r) \in [0,p]\times[p,1] : h(\ell,r) = 0\}$$

is nonempty and closed. Observe that the objective function $f(\ell,r) = \frac{p-\ell}{r-\ell}$ is bounded everywhere except (p,p), which does not belong in the level set L. By Weirstrass' extreme value theorem, there exists $(\ell^*,r^*) \in L$ such that $f(\ell^*,r^*) = \sup_{(\ell,r)\in L} f(\ell,r)$. The maximizer (ℓ^*,r^*) belongs to the set $[0,p)\times(p,0]$, as $f(p,r) = 0$ for all $r \in [p,0]$ and $h(\ell,p) < 0$ for all $\ell \in [0,p]$.

Conversely, to show the "only if" part of the statement, suppose that a seller-optimal posterior exists, and suppose $h(p, \overline{r}) \leq 0$. Observe that $h(\ell, r)$ is strictly increasing in ℓ, as $h_\ell(\ell, r) = (r - \ell)c''(\ell)$. In addition, for all $\ell \in [0, p]$, $\overline{r} = \text{argmax}_{r \in [p,1]} h(\ell, r)$ as $h_r(\ell, \overline{r}) = 0$ and $h(\ell, r)$ is strictly concave in r. Consequently, for all $(\ell, r) \in [0, p) \times (p, 1]$, we have

$$h(\ell, r) < h(p, r) \leq h(p, \overline{r}) \leq 0.$$

It follows that the implementability constraint set is empty, a contradiction. ■

An intuitive sufficient condition for the existence is that the cost of acquiring information is small enough.

Corollary B.1. *Let the information cost function be scaled by a constant $\lambda > 0$; that is, let $\widetilde{C}(\cdot) = \lambda C(\cdot)$ be the new information cost function. There exists a threshold $\overline{\lambda} > 0$ such that a seller-optimal pair of posteriors exists if $\lambda \leq \overline{\lambda}$.*

Proof. Under the new information cost function, we have

$$\overline{r} = (c')^{-1}\left(1/\lambda + c(p)\right), \text{ and}$$
$$h(p, \overline{r}) = \overline{r} - \kappa - \lambda \cdot \left[c(\overline{r}) - (\overline{r} - p)c'(p)\right].$$

As λ approaches 0, $\overline{r}$ approaches 1 and $\overline{r} - \kappa$ approaches $1 - \kappa > 0$. Since the term $[c(\overline{r}) - (\overline{r} - p)c'(p)]$ is bounded, there exists some $\overline{\lambda}$ such that $h(p, \overline{r}) > 0$ for all $\lambda \leq \overline{\lambda}$. ■

If a seller-optimal pair of posteriors exists, it is generally unique except in pathological cases. Nonetheless, the following condition guarantees uniqueness.

Proposition B.2. *If c'' is nonincreasing on $[0, p]$, there exists at most one seller-optimal pair of posteriors.*

Proof. Suppose c'' is nonincreasing on $[0, p]$, and suppose there exist two distinct seller-optimal pairs of posteriors (ℓ, r) and (ℓ^*, r^*) that result in the same success rate. That is,

$$\frac{p - \ell}{r - \ell} = \frac{p^* - \ell^*}{r^* - \ell^*}.$$

Without loss of generality, suppose $\ell < \ell^* < p < r^* < r$. The first-order necessary conditions for the optima imply that

$$\frac{h_\ell(\ell, r)}{h_r(\ell, r)} = \frac{h_\ell(\ell^*, r^*)}{h_r(\ell^*, r^*)}.$$

where the denominators $h_r(\ell, r)$ and $h_r(\ell^*, r^*)$ are positive. We have

$$\frac{h_\ell(\ell, r)}{h_r(\ell, r)} = \frac{(r-\ell)c''(\ell)}{1 - c'(r) + c'(\ell)} > \frac{(r^* - \ell^*)c''(\ell^*)}{1 - c'(r^*) + c'(\ell^*)} = \frac{h_\ell(\ell^*, r^*)}{h_r(\ell^*, r^*)},$$

a contradiction. It follows that there cannot be two distinct seller-optimal pairs of posteriors. ■

C Comparative statics

Our model has three key exogenous components: the initial probability p that the entrepreneur's firm will have high market value, the net offering price κ of the entrepreneur's firm, and the cost of information acquisition C by the investment bank. We explore how changes in these underlying parameters affect the model's outcome variables.

To this end, we modify the assumption about the cost function to accommodate changes in p and scaling by a parameter $\lambda > 0$.

Assumption C.1. For every $p \in (0, 1)$ and $\lambda > 0$, the information cost function $C_{(p,\lambda)}$ is uniformly posterior-separable (Caplin, Dean, and Leahy, 2022). That is, there exists a function $\tilde{c} : [0, 1] \longrightarrow \mathbb{R}$ such that, for every due diligence structure $\sigma \in \Sigma$ that induces posteriors $q_0, q_1, \ldots, q_{J-1}$,

$$C_{(p,\lambda)}(\sigma) = \lambda \cdot \sum_{j=0}^{J-1} P_\sigma(s_j) \left[\tilde{c}(q_j) - \tilde{c}(p)\right].$$

In addition, the function $\tilde{c}$ satisfies the properties P1–P3 of Assumption 1.

Note that this modified cost function nests the original cost function as $C(\sigma) = C_{(p,\lambda)}(\sigma)$ with $\lambda = 1$ and $c(q) = \lambda \cdot [\tilde{c}(q) - \tilde{c}(p)]$. We maintain $c(q)$ for the base case of $\lambda = 1$ and introduce $\tilde{c}(q)$ only when we consider changes in λ.

Because there exist multiple seller-optimal contracts in general (Theorem 3), we narrow our focus on the responses of a specific optimum.

Definition C.1. Let (p, κ, λ) be given. For any $\theta \in [0,1]$, a *θ-seller-optimal contract* is a seller-optimal contract $t^* = (t_0^*, t_1^*)$ such that

$$(1-\ell^*)t_0^* + \ell^* t_1^* = \theta \left[c(r^*) - c(\ell^*) - (r^* - \ell^*)c'(r^*)\right],$$

where (ℓ^*, r^*) is a seller-optimal pair of posteriors.

Note that, with this definition, every seller-optimal contract qualifies as a θ-seller-optimal contract for some $\theta \in [0,1]$. For the remainder of this section, let $\theta \in [0,1]$ be fixed.

As in Appendix B, define the function $h : [0,p] \times [p,1] \longrightarrow [-\infty, \infty)$ as

$$h(\ell, r) = r - \kappa - [c(r) - c(\ell)] + (r - \ell)\lim_{q \to \ell} c'(q),$$

so that the implementability constraint in Definition 5 is equivalent to $h(\ell, r) = 0$. We assume that h is strictly quasiconcave, so that there exists at most one seller-optimal pair of posteriors given (p, κ, λ).

Corollary C.1 (Changes in p)**.** *Suppose $p' \in (p, \kappa)$, and let $((t_0, t_1), (\ell, r), \rho)$ and $((t_0', t_1'), (\ell', r'), \rho')$ denote the triples of θ-optimal contracts, the seller-optimal pair of posteriors, and the resulting success rates given ex-ante probabilities on the good state p and p', respectively. Then*

- **(a)** $t_0 > t_0'$ *and* $t_1 > t_1'$*,*
- **(b)** $\ell < \ell'$ *and* $r > r'$*, and*
- **(c)** $\rho < \rho'$.

In other words, as the ex-ante probability on the good state increases, the transfers to the intermediary in both states decrease, the due diligence structure becomes less informative, and the success rate increases. This shift occurs because, as the entrepreneur's firm is more likely to have high market value, the intermediary does not need to work as hard to sell the firm's shares to the buyer.

Proof of Corollary C.1. Observe that

$$\rho = \frac{p - \ell}{r - \ell} < \frac{p' - \ell}{r - \ell} \leq \frac{p' - \ell'}{r' - \ell'} = \rho',$$

where the first inequality is implied by $p < p'$ and the second inequality is due to (ℓ', r') being the seller-optimal pair of posteriors given p'. It follows that $\rho < \rho'$.

As (ℓ, r) and (ℓ', r') are the seller-optimal pairs of posteriors given p and p', their first-order conditions imply

$$\frac{1-\rho}{\rho} = \frac{h_\ell(\ell, r)}{h_r(\ell, r)} \quad \text{and} \quad \frac{1-\rho'}{\rho'} = \frac{h_\ell(\ell', r')}{h_r(\ell', r')}.$$

Because $\rho < \rho'$, the above implies that

$$\frac{h_\ell(\ell, r)}{h_r(\ell, r)} > \frac{h_\ell(\ell', r')}{h_r(\ell', r')}.$$

By the strict quasiconcavity of h, it follows that $\ell < \ell'$ and $r > r'$. By the definition of θ-seller-optimal contracts, these inequalities imply that $t_0 > t_0'$ and $t_1 > t_1'$, completing the proof. ■

Corollary C.2 (Changes in κ)**.** *Suppose $\kappa' \in (\kappa, 1)$, and let ρ and ρ' denote the maximal success rates given the net offering prices κ and κ', respectively. Then $\kappa\rho > \kappa'\rho'$.*

That is, raising the net offering price κ reduces the expected revenue $\kappa\rho$ of the entrepreneur (seller). In other words, the buyer's demand ρ for the seller's firm is inelastic with respect to its price κ. The simple intuition behind this result is the agency costs of contracting with the intermediary. It is costly for the seller to induce the intermediary to extract additional surplus from the buyer, as the seller must respect the intermediary's incentives. This cost is reflected in the implementability constraint for the seller-optimal pair of posteriors. In constrast, if there were no such constraint, the entrepreneur's expected revenue would be constant at p, as the maximal success rate is p/κ by letting the posteriors be $(\ell, r) = (0, \kappa)$.

Proof of Corollary C.2. Define $\hat{\rho}(\kappa)$ as the maximal success rate given a net offering price κ subject to the implementability constraint. That is,

$$\hat{\rho}(\kappa) = \max_{(\ell,r)\in[0,p)\times(p,1]} \frac{p-\ell}{r-\ell} \quad \text{subject to} \quad c(r) - c(\ell) - (r-\ell)c'(\ell) = r - \kappa.$$

Let $(\hat{\ell}(\kappa), \hat{r}(\kappa))$ denote the maximizer in the above problem. By Envelope Theorem, the derivative of $\hat{\rho}$ satisfies

$$\hat{\rho}'(\kappa) = \frac{\hat{\rho}(\kappa)}{\left[\hat{r}(\kappa) - \hat{\ell}(\kappa)\right] \cdot \left[1 - c'(\hat{r}(\kappa)) + c'(\hat{\ell}(\kappa))\right]}.$$

This derivative and the implementability constraint imply that $\kappa\hat{\rho}(\kappa)$ is strictly decreasing in κ. ■

Let $\overline{\lambda} > 0$ denote a threshold such that a seller-optimal pair of posteriors exists for every $\lambda \in (0, \overline{\lambda})$. From Corollary B.1, such threshold exists.

Corollary C.3 (Changes in C)**.** *Let* $((t_0(\lambda), t_1(\lambda)), (\ell(\lambda), r(\lambda)), \rho(\lambda))$ *denote the triple of* θ*-seller-optimal contract, the seller-optimal pair of posteriors, and the resulting success rate given the cost scaling parameter* $\lambda \in (0, \overline{\lambda})$*. As* λ *approaches* 0,

- **(a)** $t_0(\lambda) \longrightarrow 0$, $t_1(\lambda) \longrightarrow 0$,
- **(b)** $\ell(\lambda) \longrightarrow 0$, $r(\lambda) \longrightarrow \kappa$, *and*
- **(c)** $\rho(\lambda) \longrightarrow p/\kappa$.

A key insight from this result is that, as acquiring information becomes less costly, the seller-optimal pair of posteriors approaches $(0, \kappa)$, which is the same as the seller's choice if he had costless information with full commitment power as in Bayesian persuasion (Kamenica and Gentzkow, 2011). Rather than having such commitment power by assumption, however, the seller in our model gains credibility by designing an appropriate contract and delegating the communication to the intermediary. In this sense, our model offers a concrete microfoundation for the commitment assumption.

Proof of Corollary C.3. Observe that

$$(\ell(\lambda), r(\lambda)) \in \underset{(\ell,r)\in[0,p)\times(p,1]}{\operatorname{argmax}} \frac{p-\ell}{r-\ell}$$
$$\text{subject to} \quad \lambda \cdot \left[\tilde{c}(r) - \tilde{c}(\ell) - (r-\ell)\tilde{c}'(\ell)\right] = r - \kappa.$$

By the Maximum Theorem, $\ell(\lambda)$ and $r(\lambda)$ are continuous in λ.

Let $\varepsilon > 0$ be given. Define the ε-neighborhood of $(0, \kappa)$ as

$$B_\varepsilon(0, \kappa) = \big\{(\ell, r) : 0 \leq \ell < \epsilon \ \text{ and } \ \kappa \leq r < \kappa + \epsilon\big\}.$$

There exists a pair $(\ell_\varepsilon, r_\varepsilon) \in B_\varepsilon(0, \kappa)$ such that, for every $(\ell, r) \in [0, p) \times (p, 1]$, $f(\ell, r) \geq f(\ell_\varepsilon, r_\varepsilon)$ implies $(\ell, r) \in B_\varepsilon(0, \kappa)$. Define λ_ε so that $(\ell_\varepsilon, r_\varepsilon)$ satisfies the implementation constraint given λ_ε. That is, let

$$\lambda_\varepsilon = \frac{r_\varepsilon - \kappa}{\tilde{c}(r_\varepsilon) - \tilde{c}(\ell_\varepsilon) - (r_\varepsilon - \ell_\varepsilon)\tilde{c}'(\ell_\varepsilon)}.$$

Let $(\ell^*_\varepsilon, r^*_\varepsilon)$ denote the seller-optimal pair of posteriors given the cost scaling parameter λ_ε. By construction, $(\ell^*_\varepsilon, r^*_\varepsilon) \in B_\varepsilon(0, \kappa)$.

By the continuity of $\ell(\cdot)$ and $r(\cdot)$, there exists $\delta > 0$ such that $\lambda^* \in (\lambda_\varepsilon - \delta, \lambda_\varepsilon + \delta)$ implies $|(\ell(\lambda^*), r(\lambda^*)) - (\ell^*_\varepsilon, r^*_\varepsilon)| < \varepsilon$, where $|\cdot|$ denotes the Euclidean norm.

By triangle inequality, we have

$$\big|(\ell(\lambda^*), r(\lambda^*)) - (0, \kappa)\big| \leq \big|(\ell(\lambda^*), r(\lambda^*)) - (\ell^*_\varepsilon, r^*_\varepsilon)\big| + \big|(\ell^*_\varepsilon, r^*_\varepsilon) - (0, \kappa)\big| < 2\varepsilon.$$

This implies that, for every $\lambda \in (0, \lambda^*)$,

$$\frac{p - \ell(\lambda)}{r(\lambda) - \ell(\lambda)} = \rho(\lambda) \geq \rho(\lambda^*) > \frac{p - 2\varepsilon}{\kappa}.$$

As a result, we have

$$(\ell(\lambda), r(\lambda)) \in B_{\varepsilon'}(0, \kappa) \quad \text{where} \quad \varepsilon' = \max\left\{\frac{2\varepsilon\kappa}{p - 2\varepsilon}, \frac{2\varepsilon\kappa}{\kappa - (p - 2\varepsilon)}\right\}.$$

Since the choice of ε is arbitrary, the above implies that $(\ell(\lambda), r(\lambda)) \longrightarrow (0, \kappa)$ as $\lambda \longrightarrow 0$. Therefore, $\rho(\lambda) \longrightarrow p/\kappa$ as $\lambda \longrightarrow 0$.

Finally, observe that $t_0(\lambda)$ and $t_1(\lambda)$ satisfy

$$(1 - r(\lambda)) \cdot t_0(\lambda) + r(\lambda) \cdot t_1(\lambda) = r(\lambda) - \kappa, \text{ and}$$
$$(1 - \ell(\lambda)) \cdot t_0(\lambda) + \ell(\lambda) \cdot t_1(\lambda) = \lambda\theta\left[c(r(\lambda)) - c(\ell(\lambda)) - (r(\lambda) - \ell(\lambda))c'(r(\lambda))\right],$$

which imply that both $t_0(\lambda)$ and $t_1(\lambda)$ approach 0 as $\lambda \longrightarrow 0$. ■