Selling Information for Bilateral Trade*

Robert Evans
University of Cambridge

In-Uck Park
University of Bristol

June 25, 2026

We study the design and pricing of information by independent information providers for a buyer engaged in a trading relationship with a seller, and examine their welfare implications. While competitive information firms would sell full information about the value of the good, a monopolist supplies limited information in a particularly simple form—often resulting in higher total surplus than under full information, by influencing the seller's pricing decision. If only a single information structure can be offered, it takes a binary threshold form. If the provider can offer a menu of priced information structures before the seller sets a price, he offers a continuum of thresholds, inducing a unit-elastic demand function for the seller. In both scenarios, the information firm increases total welfare when the seller's production cost is high but reduces it when the cost is low. (*JEL* Codes: D42, D61, D82, D83, L15)

Keywords: Information Sale, Mechanism Design, Information Design

1. Introduction

Market efficiency and outcomes are critically affected by the information possessed by the transacting parties, who sometimes obtain information from a third-party firm specializing in information provision and advice. Information trade is more likely if estimating the transaction value requires detailed knowledge both about the buyer's characteristics and about the nature of the product. Prominent examples include investment banks advising about potential takeovers, and product testing websites selling detailed test results for various aspects of consumer products by subscription or per item (e.g., consumerreports.org, shop.oekotest.de, which.co.uk). The market for such paid-for advice can be expected to expand in various forms in

*We are grateful to an Editor and three anonymous referees for their advice. We thank Olivier Compte, Matt Jackson, Stephen Morris, Arina Nikandrova, Alessandro Pavan, Anne-Katrin Roesler, Larry Samuelson, Lily Ling Yang and Joel Watson for helpful comments and suggestions, as well as seminar and conference audiences at AEA Winter Meeting 2016, Paris School of Economics, Seoul National University, Toulouse School of Economics, the Universities of Cambridge, Kyoto, Mannheim, Osaka, and Oxford. The usual disclaimer applies. The paper supersedes our earlier working paper, entitled "Third-Party Sale of Information."

Emails: rae1@cam.ac.uk; i.park@bristol.ac.uk

the future as a result of the development of data science and artificial intelligence, and the accumulation of detailed knowledge not readily available to individual sellers or buyers. This calls for furthering our understanding of the ways in which economic outcomes and welfare depend on the interplay between information markets and goods markets in the context of independent information providers. In this paper we study the optimal design and pricing of information, and their welfare effects, by information firms who sell, to the buyer, information about the value of the match which neither the buyer nor the seller knows.

When an information firm designs and offers information to the buyer, it also shapes the demand function for the seller's good because demand is determined by whether the buyer obtains information from the information firm and, if so, what kind. As a result, the information firm's offer—by shaping the buyer's incentives—affects the seller's pricing decision, which in turn influences whether the buyer acquires information. By analysing this three-way strategic interaction among information firm, buyer, and seller, which has not been explored in previous studies, we provide new insights into the information firm's optimal policy and its welfare implications.

We consider a buyer (he) and a seller (she) who are symmetrically informed about the buyer's value for the seller's indivisible good (a random variable v), and an information firm, the "advisor" (he), who has access to more precise information about v. The advisor publicly offers a contract consisting of a signal on v and a fee; the seller announces a price for her good. The buyer then decides whether to accept the advisor's contract and whether to buy the good. If he accepts the contract, he uses the signal realization in his purchase decision for the good. We characterize equilibrium and welfare properties for various forms of information market, deriving guidelines on potential welfare-improving policies. Notably, the specific features of information markets give some implications at odds with conventional wisdom; for example, a monopoly information market may deliver higher welfare than a competitive one.

If the seller moves first and commits, publicly, to her price then it is easy to see that the monopoly outcome prevails, as if the buyer knows v precisely, because the advisor will inform the buyer whether v is above or below the seller's price (see Section 6). Hence, we focus on the other scenario, in which the advisor publicly offers a contract before the seller sets her price, where the three-way interaction mentioned above is pronounced. This case is particularly appropriate if the seller's price is unobservable to the advisor because, say, it may be privately negotiated by the trading parties (e.g., in takeover deals) or subject to secret discounts, but the analysis is also valid when price is observable.[1] We show that key equilibrium features and welfare effects are similar in the two scenarios. In particular, information is supplied in a simple, binary structure (albeit by a different logic in each case), and the advisor enhances total welfare if the seller's production cost is high but otherwise reduces it.

By offering a contract, the advisor designs a game between seller and buyer. If the seller's price is low enough, the buyer bypasses the contract offer and buys the good outright because the signal is unlikely to change his mind from buying the good. If the price is very high, he buys neither information nor the good. So, he purchases the signal when the good's price lies in an intermediate interval, where it is more likely to help his purchase decision. Hence, the seller's optimal price maximizes her profit either (i) subject to inducing the buyer to buy the signal, or (ii) subject to inducing the buyer to bypass the advisor's offer and buy outright, whichever gives higher profit.

The advisor's objective is to design a signal that commands the maximal fee, subject to inducing the seller to price so that the buyer buys the signal. It is far from obvious *ex ante* what form of signal achieves this dual objective. Perhaps surprisingly, we show that a simple, binary signal is optimal—the advisor offers to inform the buyer whether or not v exceeds a particular threshold level; subsequently, the seller sets the highest price at which the buyer would opt to purchase the offered signal (and then buy the good if and only if v is above the threshold).

The underlying intuition for this result is that a threshold signal has a twofold advantage. Firstly, for a given probability of trade of the good, consumer surplus is maximized if trade occurs when v exceeds an appropriate threshold, and the advisor can extract this surplus as fee. Secondly, the threshold structure reduces the seller's informational disadvantage since she can increase price, without the probability of trade falling, up to the point at which the buyer's surplus, net of information fee, is zero. This increases her incentive to induce information trade, allowing the advisor to increase his fee (the constraint on the fee being the threat of bypass). We stress that

[1] In the case in which the seller's price is observable and verifiable the information firm can do strictly better because it has additional ways of influencing the price—see footnote 7.

the binary nature of the optimal signal does not follow from the Revelation Principle, because simply giving the buyer a binary recommendation, either to buy the good or not, would not necessarily be flexible enough to induce the seller's price as desired while maximizing extractible surplus. As we discuss later, the signal structure and the logic of the argument are also very different from the case when the buyer can choose a signal costlessly himself to induce the seller's pricing in his favor (Roesler and Szentes, 2017).

We characterize the unique equilibrium threshold, the advisor fee and seller price (Proposition 1). We also derive clear welfare effects of an information firm. First, with no advisor, the good trades at a price equal to the mean valuation μ of the good so long as the seller's constant production cost is below this. Hence, the outcome is efficient when cost is zero but becomes less efficient as cost rises toward μ. The advisor, on the other hand, sets the equilibrium threshold above production cost, because setting it below it reduces the value of information for the buyer (hence the fee extractible), given that seller price will exceed cost. This inefficiency dissipates and eventually disappears as production cost reaches the mean valuation. As a result, the advisor reduces total welfare when cost is low and increases it when cost is high, the underlying reason being that information is more valuable for high cost goods.

However, the advisor does not benefit the trading parties: the seller's profit is lower if the advisor is present and the buyer's surplus is fully extracted either way. Hence, sellers may have an incentive to lobby to ban a third-party information firm, but our results above show that such a ban may be welfare-reducing.

Next, we extend our analysis to a more general environment in which the advisor offers a menu of contracts, from which the buyer may choose one. Importantly, the buyer can condition his selection from the menu on the seller's price. This more general analysis essentially reaffirms key equilibrium features and welfare effects obtained for the case that the advisor offers a single contract, but illuminates how extra contracts can boost the advisor's fee.

To see this, note that the advisor can raise the fee by lowering the bypass price, and hence the seller's incentive to bypass. Suppose that the advisor offers, in addition to the optimal contract, a second contract which the buyer strictly prefers to the initial contract at the bypass price but not at the initial optimal price. This lowers the bypass price, opening room for raising the fee. Thus, offering a menu of two contracts increases the advisor's payoff. This effect strengthens as more contracts are added. In Section 5 we characterize the optimal menu and show that it consists only of threshold contracts without loss of generality, and contains all thresholds from zero (the lowest value of v) up to some maximum level. But, in addition, the menu, and corresponding demand function (mapping seller price to trade probability), must take a particular form: the fees for the different threshold contracts must be chosen so that the seller is indifferent between all prices which induce information purchase (i.e., the implied demand function is unit-elastic) and, in equilibrium, the seller charges the highest such price.

An intuition for this result is as follows. For a given menu the buyer's indirect utility rises as the price of the good falls, at a rate equal to the trade probability at any given price.[2] To reduce the bypass price, therefore, the optimal menu is designed to increase the probability of trade, hence buyer's utility from information, as much as possible for prices below the equilibrium price, subject to seller profit not exceeding equilibrium profit. Thus, the seller is indifferent between all prices at which the buyer would buy information. Using these equilibrium features, we formulate and solve an optimization problem that characterizes the equilibrium outcome.

As noted above, if the seller first commits to a price, the monopoly outcome results, as if the buyer knows his value for the good precisely. It is straightforward to see that the same outcome may prevail if multiple identical advisors competitively offer contracts to the buyer, because the seller foresees that, owing to this competition, the buyer will obtain full information for free (hence, the buyer keeps his surplus, unlike in the previous case where the advisor extracts it). Relative to the case without advisor(s), in line with our main analysis, this outcome gives higher total welfare when production cost is high but lower welfare when cost is low. Moreover, in a wide range of cases (in particular, if production cost is relatively high or if the distribution of the buyer's valuation is not too concentrated) total welfare is higher in our main model than in the monopoly outcome. In other words, in many cases it may be more efficient to have a monopoly advisor who commits publicly to an information

[2] This corresponds to a familiar result in contract theory. In our setting the buyer's type is the seller price and the lowest type is the (high) equilibrium price.

policy than to have competitive information suppliers, although buyers are better off with competitive advisors. This suggests that caution may be needed in applying conventional wisdom when regulating information markets.

The paper is organized as follows. We discuss related literature next. We set up the model in Section 2 and in Section 3 we illustrate in an example the main strategic considerations facing the players. We analyze the single-contract game in Section 4, then the menu game in Section 5. In Section 6 we discuss two variant models which give rise to full information, and, in Section 7, the relation of our analysis to analyses of buyer-optimal signals and hard information. All proofs are in the Appendix.

Related Literature. Among early contributions on information design and sale, Admati and Pfleiderer (1986) study a monopoly seller of financial information but their focus is on selling independent noisy information to buyers in a competitive asset market. Rather than a third-party seller, Lewis and Sappington (1994) study a monopoly seller of a good who can provide a signal to a buyer for the purpose of price discrimination. Lizzeri (1999) (see also Albano and Lizzeri, 2001) studies a third-party intermediary but, unlike in our paper, it is paid by the (informed) seller to inform the buyer.

Our model differs in multiple respects from the early literature on Bayesian persuasion, pioneered by Kamenica and Gentzkow (2011). Firstly, the information designer designs a game to be played by two other players, buyer and seller, rather than for one receiver. Secondly, both information and a product are sold, so that prices are crucial strategic variables. In the language of Kamenica and Gentzkow (2011), we combine two ways in which an agent can be induced to do something, by pricing and by changing beliefs. That is, our paper is in the mechanism design tradition, in that the designer can manipulate outcomes (including the information fee) as well as the information structure. Bergemann and Morris (2019) and Kamenica (2019) survey the extensive literature on Bayesian persuasion and information design.

Of papers which study mechanism design combined with information design, Bergemann and Pesendorfer (2007), Eso and Szentes (2007), Li and Shi (2017), Wei and Green (2024) and Guo, Li and Shi (2024) are concerned with a seller of a good who designs, for one or more buyers, both an information mechanism and a trading mechanism; this involves different considerations from those of our setting, in which a third-party information firm designs a mechanism for a buyer.[3] In Bergemann, Bonatti and Smolin (2018) a principal designs a mechanism for a privately informed buyer (e.g. of credit scores) to acquire incremental information before taking an action. Again, this differs from our context, in which the designer influences a buyer-seller trading relationship. Hörner and Skrzypacz (2016) study information sale, but their focus is on gradual release of information by an informed agent, to mitigate a holdup problem.

Closer to our paper, because they concern a third party selling information to players engaged in a trading relationship, are Yang (2019, 2022) and Lee (2021), but they differ from our paper in several respects. In Yang (2019) the intermediary is a platform between consumers and the monopoly firm who can only contract via the platform. In Yang (2022) the intermediary sells information (about market segmentation) to the monopolist seller, rather than to the buyer as in our paper. In Lee (2021) too, the informed party deals with the seller, in the sense that it collects payments from sellers for recommendations to buyers (see also Inderst and Ottaviani, 2012). Bergemann and Bonatti (2019) review a number of papers which study markets for data.

Our paper also relates to the literature on sequential common agency; Calzolari and Pavan (2006) study sequential contracting of two principals with an agent and the conditions under which it is optimal for the first principal to sell information revealed in the first contracting stage to the second principal. Our focus is different since our buyer initially has no private information, the two principals choose mechanisms before the agent acts and the first principal sells information to the agent.

In Section 7 we compare our results to those of recent papers on related models: one in which a buyer designs their own signal before the seller sets her price (Roesler and Szentes, 2017; Ravid, Roesler and Szentes, 2022), and another in which a monopoly certifier sells hard information to sellers about the value of their good (Ali, Haghpanah, Lin and Siegel, 2022; Asseyer and Weksler, 2024; Celik and Strausz, 2025).

2. Model

There is a single seller (S) of an indivisible object/good and a single potential buyer (B). The value of the good to B, denoted by v, is distributed according to a

[3] See also Gershkov, Moldovanu, Strack and Zhang (2021) and Mensch (2022), in which the seller of the good designs the trading mechanism in the knowledge that the buyer(s) will acquire information, or take other actions which make their types endogenous, before accepting the mechanism.

CDF F with support $V \equiv [0,1]$, continuous density $F'(v)$ and mean μ. Neither S nor B knows the value of v; for each of them their subjective belief about v is given by F and this is common knowledge. There is also a third party, A (for 'advisor'), who can find out, and trade, more precise information about v as specified below.

The advisor A maximizes his payoff by selling information about v to B. A can offer any signal for any fee t, where a *signal* (aka experiment), ψ, maps V to the set of random variables on a signal realization space Σ. If B buys the signal ψ, A informs B of its realization. For example, depending on A's chosen ψ, he could reveal the true value of v, or reveal a partition element that contains it, or provide a stochastic signal which is imperfectly informative about the value of v. B can use the posterior expectation of v implied by the signal realization in his decision whether to buy the good (i.e., when it exceeds the price of the good) or not.

Our aim is to establish A's optimal selling scheme; in particular, what form the signal should take, and how much to charge for it. Since the trade outcome depends only on the marginal distribution of B's posterior mean of v we may without loss redefine the realized signal as the implied posterior mean, i.e., $E(v|s) = s$, where $s \in \Sigma$ is the signal realization. Furthermore, a distribution H is such a marginal distribution, for some signal, if and only if H is a mean-preserving contraction of F.[4] Therefore, we define signals as follows (cf. Kamenica, 2019).

Definition. *A signal ψ is a joint distribution over $V \times \Sigma \equiv [0,1] \times [0,1]$, such that (i) the marginal over V is F, i.e., $\psi([0,v] \times [0,1]) = F(v)$ for $v \in V$; and (ii) the marginal over Σ, denoted by H^{ψ}, is a mean-preserving contraction of F.*

We denote the set of signals by Ψ. Particularly useful in the sequel is the class of signals which reveal whether or not v exceeds a certain threshold $\theta \in V$. We refer to these as 'single-threshold' signals and denote them by T_θ. T_θ assigns probability 1 to the set $([0,\theta) \times \{E(v|v<\theta)\}) \bigcup ([\theta,1] \times \{E(v|v \geq \theta)\})$ and so the distribution H^{T_θ} assigns probability $F(\theta)$ to $E(v|v<\theta)$ and $1-F(\theta)$ to $E(v|v \geq \theta)$.

We denote by $\mathcal{C}$ the set of feasible contracts[5] which A may offer, where

[4] Strassen (1965), Theorem 8.

[5] Note that allowing the fee t to depend on the signal realization, or on B's action, would introduce moral hazard on the part of A at the stage of providing the signal realization.

$$\mathcal{C} \;\equiv\; \{(\psi, t) \,|\, \psi \in \Psi, t \in \mathbb{R}\}.$$

In the general game that we analyze (Section 5) A announces a menu of contracts from which B may select one. This *menu game*, denoted by Γ_M, is defined as follows.

(1) A publicly announces a menu of contracts, i.e., a subset $M \subseteq \mathcal{C}$.

(2) S announces price $p \in \mathbb{R}_+$; B observes p.

(3) B either selects one contract in M or none (i.e., rejects).

(4) If B selects contract $(\psi, t) \in M$: B pays t to A; A observes and supplies to B the realized signal as specified by ψ; B then decides either to buy S's good for price p, or not.

(5) If B rejected: B decides either to buy S's good for price p, or not.

However, before analyzing the menu game we study, in Section 4, the *single-contract game*, denoted Γ_1, in which the advisor may only offer one contract, i.e., $\#(M) = 1$. This case is of independent interest since there may be situations in which the more general menu case is infeasible; in addition, the single-contract case serves as a benchmark and is helpful in developing intuitions.

All parties are risk-neutral expected utility maximizers and have quasi-linear utility for money. Thus, if the good is traded at price p and B pays t to A, then A's payoff is t, B's payoff is $v - p - t$, and S's payoff is $p - c$, where $c \in [0, 1)$ is the cost of production, which is common knowledge.

We study *perfect Bayesian equilibrium.* It is characterized by backward induction in Γ_M because the belief on v at any information set is unambiguous[6] and every move is observed by all parties yet to make a decision. The *outcome* of an equilibrium refers to A's fee, S's price and the mapping from v to trading probability, on the equilibrium path. These determine equilibrium welfare and each player's utility.

We now discuss various aspects of the model. Modeling A as offering a menu for B to choose from is particularly appropriate if S's price p is unobservable[7] by A.

[6]We assume a 'no-signaling-what-you-don't-know' condition: the belief is F at all information sets except at those of B after he receives a signal realization from A, where it is the Bayes-updated posterior on v, corresponding to the chosen signal. This is implied by the *plain consistency* condition on beliefs defined in Watson (2024).

[7]Though the analysis remains valid if A can observe p. In the case in which the seller's price is

This is a natural assumption in many situations. Even if there is a publicly quoted price, S may be able to adjust the price in ways observable only by B.[8] In such a case the true price is effectively private information to B and it is natural, and without loss of generality, to assume that A proposes a menu for B to choose from, his choice depending on p. The same is true if the price is observable to A but not verifiable.

We model A as the first mover. This means that part of A's objective when designing the signal is to influence S's price in such a way as to increase the value to the buyer of the information contained in the signal. This move order is more appropriate than the reverse one, in which A reacts to S's price, when, as just discussed, p is unobservable by A. In addition, in many settings it is natural to think of the advisor as able to move first and commit to a strategy, if only because it is easier to change a price than to change the information policy. Consider, for example, a setting in which A is a consultant who provides information to a sequence of clients (buyers). A would like to set at the outset an information policy which maximizes his long-run payoff. Potential clients may observe, in a statistical sense, the outcomes of the consultant's previous advice, but only with a lag. Supposing that the consultant lacks commitment power, could he gain by deviating from this policy, for example by negotiating with a given buyer a higher fee in exchange for a different signal, after the seller has set her price? Such a deviation can only damage his future reputation and, since the buyer has no way of knowing whether the information supplied is indeed drawn from a different signal, this short-run renegotiation would not be credible.[9] A plausible way to represent such a situation is a three-player game in which A moves first, and commits to a strategy.[10]

observable and verifiable the information firm can do strictly better as there are additional ways to influence the price. For off-path prices, the contract can specify a price-contingent signal, free to the buyer, which minimizes seller profit. The results of the analysis of this case, which we omit, have broad qualitative similarities to those of this paper; for example, threshold signals are always optimal.

[8]For example, the seller may offer secret discounts, personalized prices or other kinds of side-payments. A contract in which (ψ, t) is contingent on a verifiable list price named by the seller would be vulnerable to such discounts, agreed collusively with the buyer.

[9]Note that, absent renegotiation, the consultant/advisor has no incentive to deviate from the announced information policy despite the fact that B cannot observe whether he has done so, as we assume no costs of learning or communicating information for A.

[10]An alternative possibility is that A and S move simultaneously or, equivalently, A offers privately to B. As noted above, there are many settings in which sequential moves are more plausible, but a further argument against this way of representing the interaction is that equilibrium may not exist in the game in which A and S move simultaneously. For example, it can be shown that there is no

In Section 6 we also consider two other scenarios for a comprehensive analysis of independent firms selling information to B. First, a game with the opposite order of moves, i.e., S first commits to a price, which A and B both observe, and A then offers a contract to B (a menu would be redundant in this situation). Second, multiple advisors competitively offer contracts to B. In both scenarios, we establish an equilibrium trade outcome which is equivalent to that when a monopoly seller faces a fully informed buyer, and provide a welfare comparison with our main model.

We abstract from the possibility that the information firm be paid by the seller to provide information to the buyer because of a credibility problem: in many situations, the buyer may not trust soft information[11] supplied by an agent of the other party. Furthermore, in some cases it is illegal for the buyer's advisor to take payment from the seller. For example, since 2012 independent financial advisors in the UK have been forbidden to take commissions from providers of certain investment products.[12] Since the buyer must trust the information provider's advice, it is important that the advice is seen to be unbiased, hence that the seller does not pay for it.[13]

Lastly, we assume that A does not incur any costs of learning or communicating information. However, the equilibrium in our model should be robust to small costs that A may incur to supply signals, so long as the cost is increasing in complexity of the signal, because single-threshold signals should be among the least costly ones.

3. An Illustrative Example

In this section we illustrate the optimal single contract for the special case in which the buyer's value v is uniformly distributed on $[0, 1]$ and production cost is zero. We show how, starting from an offered signal which gives full information, A may modify it to an optimal binary threshold signal.

Bayesian Nash equilibrium, pure or mixed, in this game when $c = 0$, F is uniform and A may only offer a single contract. For details, see our working paper, Evans and Park (2025).

[11] Soft information is the leading example which we have in mind, but if the information is hard then our results apply unchanged. In this case there is no credibility problem arising from contracting with the seller, but it may be that contracting with the buyer is more profitable.

[12] See UK Financial Services Authority PS10/6 (www.fca.org.uk/publication/policy/fsa-ps10-06.pdf)

[13] Luca, Wu, Couvidat and Frank (2015) provide evidence that Google's practice of prominently displaying Google content, for example local business reviews, in its search pages, at the expense of independent third-party content, reduces consumer welfare. This suggests that, for important purchase decisions, buyers should be willing to pay for unbiased, rather than self-interested, advice.

Suppose A offers full information for fee $t \geq 0$. If the buyer (B) buys information, he learns the true value v and buys the good if and only if v exceeds the price p of the good, thereby obtaining an expected payoff $(1-p)^2/2 - t$. B will indeed buy information if this exceeds the payoff from rejecting A's offer and buying the good outright (i.e., $0.5 - p$) and that from not buying at all (0), in other words, if

$$(1-p)^2/2 - t \;\geq\; \max\{0.5-p, 0\}, \quad \text{i.e.,} \quad p \in \left[\sqrt{2t}, 1-\sqrt{2t}\,\right].$$

Note that this is possible (i.e., $[\sqrt{2t}, 1-\sqrt{2t}\,] \neq \emptyset$) only if $t \leq 1/8$. B bypasses A's offer and buys the good outright if $p < \sqrt{2t}$, and does not buy at all if $p > 1-\sqrt{2t}$.

For $t \leq 1/8$, therefore, S obtains an expected profit of $(1-p)p$ by setting $p \in [\sqrt{2t}, 1-\sqrt{2t}\,]$, which is maximized at the standard monopoly price $p^m = 1/2$, for the monopoly profit of $\pi^m = 1/4$. Alternatively, S could sell outright by setting $p < \sqrt{2t}$, the maximum (supremum) profit from which is $\sqrt{2t}$. Since S must not prefer the latter for information transaction to take place (i.e., $\sqrt{2t} \leq 1/4$), the highest fee A can achieve is $t = 1/32$. If A offers full information for this fee, S sets $p^m = 1/2$, and B buys the information and buys the good if $v \geq 1/2$, obtaining an expected payoff (net information rent) of $(1/8) - (1/32) = 3/32$.

We now show that A can charge a strictly higher fee for a binary signal which only informs B whether v is above or below threshold 1/2. If B buys the signal, he learns that the good is worth $v_h = 3/4$ (i.e., $v \geq 1/2$) or $v_l = 1/4$ (i.e., $v < 1/2$) with equal probability. Hence, the signal can be useful only if $p \in (1/4, 3/4)$, in which case B's expected payoff from purchasing the signal is $(v_h - p)/2 - t$ as he would buy the good only if the good is worth v_h. As before, B will buy the signal if this exceeds the payoff from buying the good outright and that from not buying at all, in other words, if

$$\frac{3}{8} - \frac{p}{2} - t \;\geq\; \max\left\{\frac{1}{2} - p, 0\right\} \quad \text{i.e.,} \quad p \in \left[\frac{1}{4} + 2t,\, \frac{3}{4} - 2t\right]. \tag{1}$$

B buys outright if p is below this interval, and does not buy at all if p is above.

Thus, S may set a price in the interval in (1), inducing B's signal purchase and selling the good with probability 1/2. S's maximal profit from doing so is obtained with the highest price in the interval, $p = (3/4) - 2t \geq 1/2$, for a profit of $(3/8) - t$. Note that B's net payoff is 0 in this case (as he might as well not buy at all at that price). Alternatively, S may sell outright by setting any price below $p = (1/4) + 2t$ (the *bypass price*), and obtain a profit up to $(1/4)+2t$. Since S should not choose the latter if information transaction is to happen, the maximal fee A may achieve with the binary signal is $t = 1/24$ (from $(3/8) - t \geq (1/4) + 2t$). This is higher than the maximal fee obtainable with full information derived above, $1/32$.

Part of the reason that A does better with the binary signal is that it generates inelastic demand, allowing S to fully extract B's surplus from trading, net of A's fee, leaving no information rent for B. This in turn relaxes the constraint that S should not induce bypass, opening room for A to increase the fee.

As a matter of fact, the fee can be raised further, using the same principle. A can offer a hybrid signal structure to reduce the bypass price: if v is above $1/2$ then A informs B of that fact, and otherwise tells B the precise value of v. Then, B's incentive to buy information is the same as in the case of the binary signal above for prices above $1/2$, but is greater for lower ones since the hybrid structure is more informative. This reduces the bypass price without affecting S's payoff from inducing signal purchase, opening room for raising t. In fact the optimum offer for A, subject to trade taking place if and only if $v \geq 1/2$, is this hybrid signal,[14] with $t \simeq 0.052$.

However, A can improve upon this hybrid signal with a different binary signal, for which the threshold is below $1/2$. Lowering the threshold boosts S's profit from inducing information purchase since the probability of trade is higher, discouraging bypass; but it also boosts S's profit from bypassing if the threshold is low enough, by reducing the value of information for B. The two effects balance out at the optimum, which is characterized by the optimal threshold $\widehat{\theta} \simeq 0.297$, corresponding fee $\hat{t} \simeq 0.07$, seller price $\hat{p} \simeq 0.55$ and bypass price $\simeq 0.386$. Notice that the bypass price exceeds the threshold, which implies that the hybrid signal structure does not improve on the threshold form, because no further information on $v < \widehat{\theta}$ can benefit B when the price exceeds $\widehat{\theta}$. We show in Section 4 that this threshold structure and fee are optimal for A among all possible offers of a single contract. As we show in Section

[14] Note that A strictly prefers the hybrid signal to the apparently equivalent binary signal derived from the hybrid signal by pooling all messages that lead to identical purchase decisions by B, a signal which would simply reveal whether $v \geq 1/2$ or not. We have established above that the maximal fee A can get for this binary signal is $1/24$. If A were to offer it for fee $t \simeq 0.052 > 1/24$, S would induce B to reject A's offer and buy the good outright, by setting the bypass price.

5, A can reduce the bypass price further by offering a menu of contracts. The form of the optimal menu is pinned down by the logic of maximum deterrence against the seller setting a bypass price. In the optimal menu there is a continuum of thresholds which generate a unit-elastic demand function for the seller.

4. Single-Contract Game Γ_1

In the single-contract game Γ_1, A's offer of a contract (ψ, t) defines a game between S and B. An equilibrium of Γ_1 is specified by a continuation equilibrium for each contract $(\psi, t) \in \mathcal{C}$, and a contract which is optimal for A given the specified continuations. Since B's ultimate decision is a binary one (either buy the good or not) it might be thought that, to identify A's optimal contracts, it is sufficient to consider only binary signals. However, this is not so. The offered signal, by shaping B's decision given S's price, determines the demand function which S faces, thereby influencing S's optimal pricing. This point is illustrated by the example in Section 3: A's maximal fee obtainable with a hybrid signal ($\simeq 0.052$) is unobtainable with a binary signal that induces trade for the same values of v (i.e., $v \geq 1/2$; see footnote 14). This shows that because the buyer's "type" (the seller's price in our context) is endogenous[15] it does not follow from the Revelation Principle that there is an optimal binary signal; hence we must consider all possible signals. Nevertheless, as Proposition 1 shows, there exists a simple, pure-strategy equilibrium in which A offers a binary threshold contract and any other equilibrium, pure or mixed, is outcome-equivalent to this one.

Proposition 1 *There exists an equilibrium of Γ_1 in which, on path, A offers a binary threshold contract $(T_{\widehat{\theta}}, \hat{t})$, S sets a price $\hat{p} > 0$, and B accepts the contract and then buys the good if and only if $v \geq \widehat{\theta}$. Any other equilibrium is outcome-equivalent to this equilibrium; that is, on path, A receives fee $\hat{t}$, S's price is $\hat{p}$ and, except possibly for a zero-measure set of v, the good is traded if and only if $v \geq \widehat{\theta}$. The value of $(\widehat{\theta}, \hat{t}, \hat{p})$ depends on c and is given as follows.*
(a) If $c \geq \mu$, $\widehat{\theta} = \hat{p} = c$ and $\hat{t} = \int_c^1 (v - c)dF$. Hence, the outcome is efficient and A extracts the full surplus.

[15] This argument does not apply in the menu game Γ_M; see Lemma 3 in Section 5.

(*b*) *If* $c < \mu$, $\widehat{\theta}$ *is the unique solution to*

$$(\theta - c)(1 + F(\theta))^2 \;=\; \mu - c; \tag{2}$$

$c < \widehat{\theta} < \mu$ *(hence the outcome is inefficient) and* $\widehat{\theta}$ *strictly increases in* c*; A's fee* $\hat{t} = t(\widehat{\theta}) > 0$ *where* $t(\cdot)$ *is given by*

$$t(\theta) \;=\; \int_{\theta}^{1} v dF - \frac{\mu}{1 + F(\theta)} + \frac{cF(\theta)^2}{1 + F(\theta)} \;>\; 0\,; \tag{3}$$

S's price

$$\hat{p} \;=\; \frac{\mu - c[F(\widehat{\theta})]^2}{1 - [F(\widehat{\theta})]^2} \;>\; \mu\,;$$

B's expected payoff is 0 *and S's expected payoff is*

$$\frac{\mu - c}{1 + F(\widehat{\theta})} \;=\; (\widehat{\theta} - c)(1 + F(\widehat{\theta})).$$

If $c \geq \mu$ (case (a)), so that there can be no surplus if B does not buy information, A can offer a threshold signal T_c that informs B whether v exceeds c or not, for a fee $\hat{t}$ equal to B's expected surplus from buying the good at price c if and only if $v \geq c$. In equilibrium B then buys neither information nor the good if $p > c$, S sets $p = c$, the outcome is efficient and A captures all the surplus, so offering T_c is optimal for A. There are other signals which can deliver $\hat{t}$ for A, e.g. a hybrid signal as discussed in Section 3, but all equilibria are outcome-equivalent.

If $c < \mu$ the analysis is more involved. The proof first characterizes equilibrium continuation strategies following the announcement of an arbitrary contract (ψ, t). If, given S's price p, B rejects (ψ, t), his expected payoff is $u_o(p) \equiv \max\{\mu - p, 0\}$ as he buys the good if $p \leq \mu$ and not otherwise. If he accepts, he optimally buys the good if and only if the signal s (posterior valuation) exceeds p, so his expected payoff is

$$u_I(p|(\psi, t)) \;\equiv\; \int_{p}^{1} (s - p) dH^{\psi} - t. \tag{4}$$

This is decreasing and convex in p, with left-derivative at p equal to the negative of the probability of trade of the good, $-(1 - H^{\psi}_{-}(p))$, where $H^{\psi}_{-}(p) \equiv \lim_{\rho \uparrow p} H^{\psi}(p)$. For $t > 0$ not too large, B accepts (ψ, t) if p is in an intermediate interval $(\underline{p}(\psi, t), \bar{p}(\psi, t)]$,

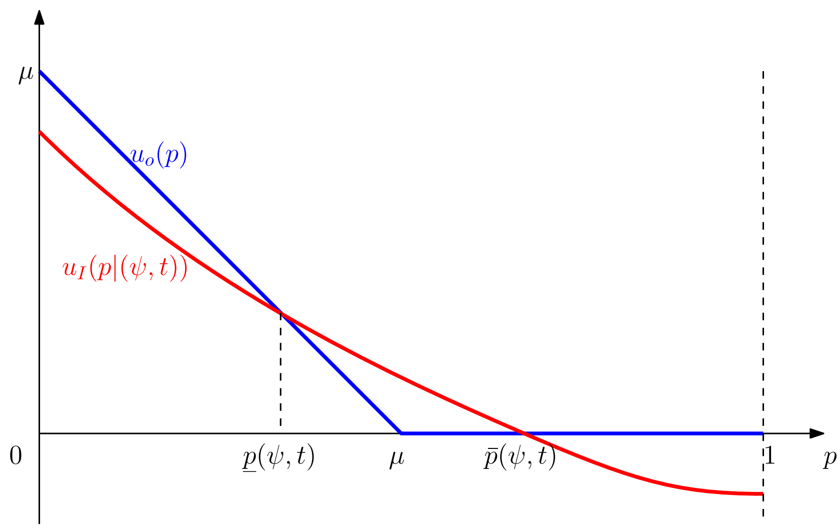

Figure 1 $u_I(p|(\psi,t))$ illustrates B's expected payoff from accepting (ψ,t), and $u_o(p)$ that from rejecting it.

where $\underline{p}(\psi,t)$ and $\bar{p}(\psi,t)$ are the two prices at which $u_I(p|(\psi,t)) = u_o(p)$ as illustrated in Figure 1,[16] buys outright for lower prices, and buys neither information nor the good for higher ones. S's profit therefore is either

$$\pi_I(p|(\psi,t)) \;=\; (p-c)(1-H_-^{\psi}(p)) \quad \text{for} \quad p \in (\underline{p}(\psi,t),\bar{p}(\psi,t)], \tag{5}$$

or else $\pi_o(p)$, which equals $p-c$ if $p \leq \underline{p}(\psi,t)$ and equals 0 if $p > \bar{p}(\psi,t)$. S's price is either a maximizer of (5), or the "bypass price" $\underline{p}(\psi,t)$, whichever gives higher profit.

We say that, for signal $\psi \in \Psi$, a fee $t > 0$ is *achievable with* ψ if $\pi_I(p|(\psi,t)) \geq \pi_o(\underline{p}(\psi,t))$, where p maximizes (5). If the inequality is strict, i.e., if S strictly prefers to induce information purchase, A could earn a slightly higher fee, since the interval $(\underline{p}(\psi,t),\bar{p}(\psi,t)]$ shrinks continuously as t rises (hence u_I falls) and $\pi_I(p|(\psi,t))$ is left-continuous in p. This implies the

Equal Profit Principle For any $\psi \in \Psi$, if some $t > 0$ is achievable with ψ then the maximal fee achievable with ψ exists; $t > 0$ is the maximal such fee if and only if $\pi_I(p|(\psi,t)) = \pi_o(\underline{p}(\psi,t)) > 0$, where p maximizes (5).

Suppose that, for a given θ, $t > 0$ is the maximal fee achievable with threshold

[16] As the proof shows, it is innocuous, for the purpose of characterizing equilibrium, to assume that at $\underline{p}(\psi,t)$ he buys outright and at $\bar{p}(\psi,t)$ he accepts the contract.

signal T_θ. Given (T_θ, t), S's optimal price is $\bar{p}(T_\theta, t)$ since sale probability is constant at $1 - F(\theta)$ for $p \in (\underline{p}(T_\theta, t), \bar{p}(T_\theta, t)]$. By the Equal Profit Principle S is indifferent between this price and the bypass price $\underline{p}(T_\theta, t)$, which pins down t as $t(\theta)$ in (3). The optimal threshold, therefore, is $\widehat{\theta}$ that maximizes $t(\theta)$, thus satisfies the first order condition (2), which has a unique solution $\widehat{\theta} \in (c, \mu)$. This characterizes the unique $(T_{\widehat{\theta}}, t(\widehat{\theta}))$ which is optimal among threshold contracts and its associated price $\bar{p}(T_{\widehat{\theta}}, t(\widehat{\theta}))$, which is $\hat{p}$ in Proposition 1. Also, in any equilibrium, A's payoff is at least $\hat{t} = t(\widehat{\theta})$ because, for any $\epsilon > 0$, A could offer $(T_{\widehat{\theta}}, \hat{t} - \epsilon)$. This has a lower bypass price and a higher $\bar{p}$, so S must price to induce information purchase, giving $\hat{t} - \epsilon$ to A.

The argument that any equilibrium contract is outcome-equivalent to this threshold one is essentially as follows. Suppose the maximal achievable fee with ψ is $t \geq \hat{t}$, which implies $c < \underline{p}(\psi, t)$. A can obtain the same payoff, t, with a specific threshold contract. Let q_* be the probability with which B would buy the good if he had the signal ψ and the price was $\underline{p}(\psi, t)$. Let T_{θ_*} be the threshold signal that induces the same trade probability q_*, i.e., $1 - F(\theta_*) = q_*$. Comparing (T_{θ_*}, t) with (ψ, t), (i) the bypass price is lower, i.e., $\underline{p}(T_{\theta_*}, t) \leq \underline{p}(\psi, t)$, and (ii) S's payoff (from setting $\bar{p}(T_{\theta_*}, t)$) is at least as high as she obtains from (ψ, t). Thus, since S weakly prefers to induce information purchase than to bypass when the contract is (ψ, t), the same is true when the contract is (T_{θ_*}, t) so A will indeed obtain t from (T_{θ_*}, t).

(i) follows from a simple observation that the most efficient way to trade the good with a certain probability, say q_*, is to do so if and only if v is above the appropriate threshold, in this case θ_*. So B's surplus is higher when trading with T_{θ_*} than with ψ at $\underline{p}(\psi, t)$, that is, $u_I(\underline{p}(\psi, t) | (T_{\theta_*}, t)) \geq u_I(\underline{p}(\psi, t) | (\psi, t)) = u_0(\underline{p}(\psi, t))$, implying (i).

(ii) is true because the total surplus achieved under (T_{θ_*}, t) is weakly higher than under (ψ, t) (it would be higher, conditional on information acquisition, at price $\underline{p}(\psi, t)$ and, as price increases from $\underline{p}(\psi, t)$, the sale probability, hence also the surplus, is constant in the former case but, since $c < \underline{p}(\psi, t)$, weakly falls in the latter) while B's payoff is zero (hence weakly lower) and A's payoff, t, is the same.

Hence $(T_{\theta_*}, t) = (T_{\hat{\theta}}, \hat{t})$ and so the best contract for A is no better than the best threshold contract. In fact, any equilibrium contract $(\psi, \hat{t})$ must be outcome-equivalent to $(T_{\hat{\theta}}, \hat{t})$. This is because S's payoff from $(T_{\hat{\theta}}, \hat{t})$ cannot be strictly higher than from $(\psi, \hat{t})$, otherwise, by (i), the Equal Profit Principle would be violated. Hence, total surplus, conditional on information acquisition, must be the same under ψ as under $T_{\hat{\theta}}$ both at $\underline{p}(\psi, \hat{t})$ and as price rises from $\underline{p}(\psi, \hat{t})$. This implies outcome-equivalence.

4.4. Effect of the Adviser on Welfare.

Does the presence of A increase or decrease total surplus, compared with a situation in which B is uninformed? Secondly, how does it affect the payoffs of B and S?

If $c \geq \mu$ then, without A, the outcome would be inefficient: if $c > \mu$ then there would be no trade and if $c = \mu$, trade would happen at price c, even if $v < c$. The advisor strictly increases total surplus, to its maximum, but is of no benefit to B or S since they both get zero whether A is present or not.

If $c < \mu$ then, again, B does not benefit since he gets zero in either case. S is strictly worse off when A is present. Without A, trade takes place at price μ and S obtains payoff $\mu - c$. With A present, S's expected payoff, by Proposition 1, is $(\mu - c)/(1 + F(\widehat{\theta})) < \mu - c$.

Whether A increases total surplus depends on the value of c. Total surplus with A present is $\int_{\widehat{\theta}}^{1}(v - c)dF$ where $\widehat{\theta}$ is a function of c. Therefore, surplus increases if this exceeds $\mu - c$, i.e., if $\int_0^{\widehat{\theta}}(v - c)dF < 0$, and decreases if the inequality is reversed. Note that $\int_0^{\widehat{\theta}}(v - c)dF > 0$ for $c = 0$ and $\int_0^{\widehat{\theta}}(v - c)dF < 0$ for c close to μ (since $c < \widehat{\theta} < \mu$). Substituting $\int_0^{\theta}(v - c)dF = 0$, i.e., $c = \int_0^{\theta} v dF/F(\theta)$, in (2) gives

$$\theta + (2 + F(\theta)) \int_0^{\theta} (\theta - v)dF \;=\; \mu.$$

The LHS strictly increases in θ, so, by continuity, there is a unique $\widehat{\theta}$, hence a unique c, at which $\int_0^{\widehat{\theta}}(v - c)dF = 0$. A increases total welfare if cost is above this level and reduces welfare if cost is below it. Summarizing,

Corollary. *There exists $\widehat{c} \in (0, \mu)$ such that, compared with the outcome without the advisor, in the equilibrium of Γ_1: (i) social surplus is strictly higher if $c > \widehat{c}$ and strictly lower if $c < \widehat{c}$; (ii) the seller's payoff is always lower; and (iii) the buyer's surplus is always zero with or without the advisor.*

When $c < \mu$ the seller in fact is made strictly worse off and so has an interest in lobbying to prevent the advisor operating; when the seller is relatively inefficient (c close to μ) such a restriction of information trade would be surplus-destroying.

5. Equilibrium in the Menu Game Γ_M

In this section, we analyze the game in which A may offer a menu of contracts.

If $c \geq \mu$ then the outcome of any equilibrium of the menu game Γ_M is the same as the unique equilibrium outcome of Γ_1 since A can extract all the surplus by offering a menu containing only the optimal single contract. Hence we consider the case $c < \mu$ below. We show that the main equilibrium properties of Γ_1 continue to hold; namely, on the equilibrium path trade takes place according to a threshold rule, the equal-profit principle prevails, S sets the maximal price which induces information purchase, and B's net surplus is zero. It turns out, however, that, unlike the optimal single contract, the optimal menu induces a unit-elastic demand function: S is indifferent between all prices which induce information purchase. The additional contracts in the menu serve to lower the price at which S can bypass A.

To facilitate exposition, we adopt an innocuous convention that a menu M always includes the *null contract* $(T_0, 0)$, which offers no information for a zero fee. Then B always selects one contract from M; selecting $(T_0, 0)$ is equivalent to rejecting all contracts. If an optimal menu for A is to exist, a continuation equilibrium must exist following any menu A may offer. At least one will exist if the menu contains a single contract, but if it contains a continuum of contracts it might be that there is no optimal contract for B, hence no continuation equilibrium. For this reason, we assume that in Γ_M A may offer a menu from the set of menus that have a continuation equilibrium, denoted by Υ. We study perfect Bayesian equilibria of Γ_M.

We examine A's maximal feasible payoff, denoted by t^*, where a payoff is feasible if it is obtained in some continuation equilibrium following (A's offer of) some menu in Υ. For expositional ease, we characterize t^* presuming it exists, deriving key equilibrium properties in the process, and show that t^* indeed exists in the Appendix. Then we show that A's payoff is t^* in every equilibrium of Γ_M. We start with a useful observation, that A obtains t^* in a pure-strategy continuation equilibrium following announcement of some $M \in \Upsilon$.

Lemma 1 *If A obtains t^e in a continuation equilibrium following $M \in \Upsilon$, there exists $M' \in \Upsilon$ such that there is a pure strategy continuation equilibrium following M' in which A obtains at least t^e.*

The logic of the proof is that if, in some mixed-strategy equilibrium following a menu, the best on-path price for A is p^e, the best on-path payoff for A following p^e is $\tilde{t} \geq t^e$, $\tilde{q}$ is the corresponding trade probability and $1 - F(\widetilde{\theta}) = \tilde{q}$, then A could add a threshold contract $(T_{\widetilde{\theta}}, \tilde{t})$ to the equilibrium menu and there would be a pure strategy continuation equilibrium in which A gets $\tilde{t}$.

To characterize t^*, therefore, we focus on pure-strategy continuation equilibria following some menu M in Υ. In each such equilibrium, for each $p \in [0,1]$, B would select an optimal contract, say $(\psi,t) \in M$, and derive a utility $u_I(p|(\psi,t))$ as defined in (4) in Section 4. Hence, B's optimized utility for $p \in [0,1]$ in the continuation game is

$$U_I(p|M) \;\equiv\; \max_{(\psi,t)\in M} u_I(p|(\psi,t)).$$

U_I is the upper envelope of B's payoff functions, derived from the contracts in M.[17] Since each $u_I(p|(\psi,t))$ is convex in p, so is U_I.

We represent B's strategy as a contract schedule $(\boldsymbol{\psi}(\cdot), t(\cdot)) : [0,1] \to \Psi \times \mathbb{R}$ where $(\boldsymbol{\psi}(p), t(p)) \in M$ is the optimal contract that B selects for $p \in [0,1]$. If S sets p and B selects $(\boldsymbol{\psi}(p), t(p))$, trade takes place when the posterior induced by $\boldsymbol{\psi}(p)$ is p or higher. We denote the probability of trade in this case by $q(p)$ and S's profit by $\pi(p|M) = (p-c)q(p)$. S sets an optimal price, say p^e, that gives her highest expected profit, determining A's payoff as $t(p^e)$. Then $q(p) = -u'_{I-}(\rho|(\boldsymbol{\psi}(p), t(p)))|_{\rho=p}$ (see (**??**)); it decreases in p because $U_I(p|M)$ is convex in p.

We may ignore contracts that B would never select for any p, as they do not affect the continuation equilibrium. Hence, a pure-strategy continuation equilibrium following some menu in Υ is represented by a strategy profile $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ that satisfies

(a) $U_I(p|M) = u_I(p|(\boldsymbol{\psi}(p), t(p)))$ for all $p \in [0,1]$ where $M = \{(\boldsymbol{\psi}(p), t(p))\}_{p\in[0,1]}$, and

(b) $p^e \in \arg\max_{p\in[0,1]}(p-c)q(p)$ where $q(p) = -u'_{I-}(\rho|(\boldsymbol{\psi}(p), t(p)))|_{\rho=p}$.

We call such a strategy profile a "PCE" (for pure-strategy continuation equilibrium).

We say that a PCE $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ is *optimal* if it delivers t^* for A, i.e., if $t(p^e) = t^*$, and a menu M is *optimal* if it corresponds to an optimal PCE.

The following Lemma shows that all fees are non-negative in an optimal PCE, since otherwise A could do better by making a small equal increase in all fees for

[17] Note that $u_I(p|(T_0, 0)) = \max\{\mu - p, 0\} = u_0(p)$.

non-null signals in the menu.

Lemma 2 *If $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ is an optimal PCE, $t(p) \geq 0$ for all $p \in [0,1]$.*

If all contracts have nonnegative fees, it is optimal for B to select a null contract when $p = 0$ or $p = 1$ for the same reason as in Γ_1. Therefore, by convexity of $U_I(p|M)$, for a PCE $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ with $t(\cdot) \geq 0$ there exist $\underline{p}(M) \in [0, \mu)$ and $\bar{p}(M) \in (\mu, 1]$, where M is as defined in (a) above, such that B buys the good outright if $p \leq \underline{p}(M)$; buys neither information nor the good if $p > \bar{p}(M)$; and otherwise selects contract $(\boldsymbol{\psi}(p), t(p))$, buying the good with probability $q(p)$.

Unlike in the single-contract game of Section 4, we can appeal to an argument related to the Revelation Principle and focus, when characterizing equilibrium outcomes, on optimal PCE's in which all the signals offered are binary 'buy' recommendations[18] and, furthermore, are threshold signals.

Lemma 3 *Suppose $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ is an optimal PCE. Then there is an optimal PCE $[\tilde{\boldsymbol{\psi}}(\cdot), \tilde{t}(\cdot), p^e]$ such that (a) for each $p \in [0,1]$, $\tilde{\boldsymbol{\psi}}(p)$ is a single-threshold signal with the same probability of trade, $q(p)$, as given by $\boldsymbol{\psi}(p)$; and (b) $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ and $[\tilde{\boldsymbol{\psi}}(\cdot), \tilde{t}(\cdot), p^e]$ are outcome-equivalent.*

(b) implies that in an optimal PCE trade occurs, on path, if and only if v exceeds some threshold (ignoring zero-probability events). That is, the contract-price pair $[\boldsymbol{\psi}(p^e), t(p^e), p^e]$ is single-threshold-equivalent, in the following sense.

Definition. A triple $(\psi, t, p) \in \mathcal{C} \times \mathbb{R}_+$ is single-threshold-equivalent if, for some threshold $\theta \in (0,1)$, ψ generates a posterior expectation $s \geq p$ if and only if $v \geq \theta$, i.e., $\psi([0,\theta) \times [p,1]) = \psi([\theta,1] \times [0,p)) = 0$.

In other words, if the price is p then ψ provides the same useful information to B as T_θ would.

In the proof of part (a) of Lemma 3, the signal $\tilde{\boldsymbol{\psi}}(p)$, in effect, sends message *buy* if $v \geq \theta(p)$ and otherwise sends message *don't buy*, where $1 - F(\theta(p)) = q(p)$. $\tilde{t}(p)$ is the fee such that B's expected payoff, given p, from accepting $(\tilde{\boldsymbol{\psi}}(p), \tilde{t}(p))$ is the same as that from accepting $(\boldsymbol{\psi}(p), t(p))$. Note that $\tilde{t}(p) \geq t(p)$, with strict inequality if $\psi(p)$ is not single-threshold-equivalent, because B's expected gross payoff is maximized,

[18] This requires the recommendation to be conditioned on the buyer's 'type', i.e. S's price, which is not possible in Γ_1, since in that game A can only offer a single contract.

subject to a trade probability of $q(p)$, when trade takes place if and only if $v \geq \theta(p)$. Given p, choosing $(\tilde{\boldsymbol{\psi}}(p), \tilde{t}(p))$ is optimal for B when the menu is $[\tilde{\boldsymbol{\psi}}(\cdot), \tilde{t}(\cdot)]$ (for any deviation there is a payoff-equivalent deviation when the menu is $[\boldsymbol{\psi}(\cdot), t(\cdot)]$). p^e is an optimal price for S if B follows this strategy since the demand function is then unchanged from the original equilibrium, so $[\tilde{\boldsymbol{\psi}}(\cdot), \tilde{t}(\cdot), p^e]$ is an optimal PCE. Part (b) of Lemma 3 follows because if $(\boldsymbol{\psi}(p^e), t(p^e), p^e)$ is not single-threshold-equivalent then $\tilde{t}(p^e) > t(p^e)$ so $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ could not have been optimal.[19]

Propositions 2 and 3 below show that the requirement of optimality pins down the structure of the menu of threshold contracts to a specific form, driven by the need to deter S from bypassing. In any optimal PCE (if one exists) the thresholds and corresponding fees are designed in such a way that the seller obtains identical profits from all prices in $[\underline{p}(M), \bar{p}(M)]$, i.e., she is indifferent between all prices which induce information purchase, and selects the highest of these prices.

Proposition 2 *Suppose $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ is an optimal PCE. Then*
(a) $\bar{p}(M) = p^e$;
(b) $\pi(p|M) = \pi(p^e|M) = \underline{p}(M) - c$ *for all* $p \in [\underline{p}(M), \bar{p}(M)]$;
(c) $q(p^e) \leq 1 - F(c)$.

As in the equilibrium of Γ_1, S charges the highest price at which B buys information (by (a)), so that B's net surplus is zero; and, by (b), S is indifferent between doing so and bypassing A (by charging $\underline{p}(M)$). However, by (b), now S is also indifferent between all prices in between. (b) implies that the limit of $q(p)$, as $p \to \underline{p}(M)$ from above, is 1 and that, as p increases from $\underline{p}(M)$, $q(p)$ decreases continuously, in a unit-elastic manner,[20] to $q(p^e)$ at p^e, and then drops to zero. Given the threat of bypass, this turns out to be the way for A to maximize extractable consumer surplus.

For the proof of part (a), suppose that $p^e < \bar{p}(M)$. A could modify the menu of threshold contracts by removing all thresholds higher than the one chosen at price p^e. Then S could increase price above p^e while maintaining the sale probability at $q(p^e)$, hence strictly increasing her profit. Since she then would strictly prefer to choose

[19] Although the on-path signal must be single-threshold-equivalent, the off-path signals need not be: for $p' \neq p^e$, any contract $(\hat{\boldsymbol{\psi}}(p'), \hat{t}(p'))$ may replace $(\tilde{\boldsymbol{\psi}}(p'), \tilde{t}(p'))$ in the menu without affecting the equilibrium, as long as the two contracts give the same trade probability, $q(p')$, and the same buyer utility.

[20] If $c = 0$. If $c > 0$ the demand function $q(p)$ is unit-elastic with respect to mark-up.

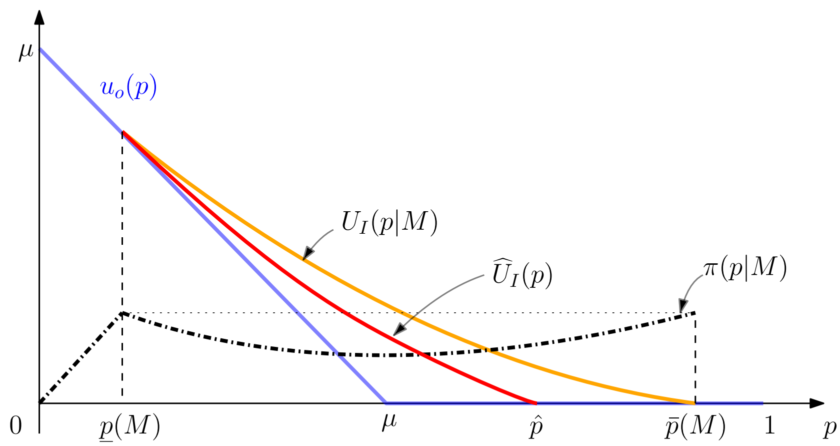

Figure 2 The menu can be modified to increase A's payoff if Proposition 2(b) fails.

such a price than to bypass A, A could further modify the menu by making a small equal increase in all the fees for non-null contracts and S would still price so as to induce information purchase, contradicting optimality of M.

Next we sketch the argument in the proof of part (b) of Proposition 2, for the case $c = 0$. Suppose $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ is optimal and satisfies (a). By optimality of p^e, $pq(p) \leq p^e q(p^e)$, i.e., since the slope of U_I at p equals the slope of the optimal u_I at p, $pq(p) = \pi(p|M) = -pU_I'(p|M) \leq -p^e U_I'(p^e|M) = \pi(p^e|M) = p^e q(p^e)$ for all $p \in [\underline{p}(M), \bar{p}(M))$. Suppose that ($b$) is not satisfied: $pq(p) < p^e q(p^e)$ for some $p \in (\underline{p}(M), \bar{p}(M)]$. We construct a utility schedule $\widehat{U}_I$ such that $\widehat{U}_I(\underline{p}(M)) = U_I(\underline{p}(M)|M)$ and, for each $p \in [\underline{p}(M), \bar{p}(M))$, the absolute value of its slope (i.e. the corresponding probability $\hat{q}(p)$ of trade of the good) satisfies $p\hat{q}(p) = p^e q(p^e)$, so that each such p would give payoff $p^e q(p^e)$ to S. $\widehat{U}_I$ must be steeper than U_I since (b) does not hold, hence hits 0 at some $\hat{p} < \bar{p}(M) = p^e$, as illustrated in Figure 2.

To construct a PCE in which B's optimal payoff function is $\widehat{U}_I(\cdot)$, define $\hat{\theta}(p)$ by $p(1 - F(\hat{\theta}(p))) = p^e q(p^e) = \pi(p^e|M)$ and let $\hat{t}(p)$ satisfy $u_I(p|(T_{\hat{\theta}(p)}, \hat{t}(p))) = \widehat{U}_I(p)$ for $p \in [\underline{p}(M), \hat{p}]$. Then, $[T_{\hat{\theta}(\cdot)}, \hat{t}(\cdot), \hat{p}]$ is a PCE.[21] This is because (i) B's optimal contract for price $p \in [\underline{p}(M), \hat{p}]$ is $(T_{\hat{\theta}(p)}, \hat{t}(p))$ since $\widehat{U}_I(\cdot)$ is the upper envelope of $\{u_I(\cdot|(T_{\hat{\theta}(p')}, \hat{t}(p')))\}_{p' \in [\underline{p}(M), \hat{p}]}$ and is convex, and (ii) this gives S a profit of $p\hat{q}(p) = -p\widehat{U}_I'(p) = p^e q(p^e)$; hence, since S is indifferent between all such prices, and the highest bypass price is $\underline{p}(M)$ as before, $\hat{p}$ is an optimal price. Note, from $\hat{p}(1 -$

[21]We show in the Appendix that $\hat{t}(\cdot) \geq 0$.

$F(\hat{\theta}(\hat{p}))) = p^e q(p^e)$ and $\hat{p} < p^e$, that $1 - F(\hat{\theta}(\hat{p})) > q(p^e)$. Therefore total surplus is strictly higher in this PCE than in $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$. However, B and S obtain the same payoffs as in $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$, namely zero and $\pi(p^e|M)$ respectively, so A is strictly better off, a contradiction. This sketches the argument that S is indifferent between all prices in $(\underline{p}(M), \bar{p}(M)]$. We show in the Appendix that $\underline{p}(M)$ also gives the same profit (the equal-profit principle).

The properties in Proposition 2 and Lemmas 2 and 3 define an optimization problem the solution of which gives the optimal PCE in single-threshold signals, which we represent by $[\theta(\cdot), t(\cdot), p^e]$ where $\theta(p)$ is the threshold chosen by B when the price is p, i.e., $\boldsymbol{\psi}(p) = T_{\theta(p)}$. Proposition 2($b$) shows that $U_I(p|M)$ satisfies the differential equation $-U_I'(p) = (p^e - c)q(p^e)/(p - c)$, and Proposition 2($a$) gives boundary condition $U_I(p^e|M) = 0$. The solution is

$$U_I(p|M) = (p^e - c)q^e \ln\left(\frac{p^e - c}{p - c}\right) \quad \text{for all} \quad p \in [(p^e - c)q^e + c, p^e] \tag{6}$$

where $q^e = q(p^e)$. Since $U_I(\underline{p}(M)|M) = \mu - \underline{p}(M)$ and, by Proposition 2(b), $\underline{p}(M) = (p^e - c)q^e + c$, (p^e, q^e) satisfies

$$p^e = \frac{\mu - c}{q^e(1 - \ln(q^e))} + c. \tag{7}$$

Since $q^e(1 - \ln(q^e))$ converges to 0 as q^e tends to 0, there exists a constant $\underline{q}(\mu, c) > 0$ such that $p^e \in [0, 1]$ implies $q^e \geq \underline{q}(\mu, c)$. Denote by (P) the maximization problem:

$$\max_{q^e \in [\underline{q}(\mu,c),1]} \quad \int_{\theta(p^e)}^1 v dF(v) - p^e q^e \;=\; \int_{F^{-1}(1-q^e)}^1 v dF(v) - \frac{\mu - c}{1 - \ln(q^e)} - cq^e \tag{8}$$

$$s.t. \int_{\theta(p)}^1 v dF(v) - pq(p) - (p^e - c)q^e \ln\left(\frac{p^e - c}{p - c}\right) \geq 0 \quad \forall p \in [(p^e - c)q^e + c, p^e] \tag{9}$$

where p^e is as in (7), $q(p) = (p^e - c)q^e/(p - c)$ and $1 - F(\theta(p)) = q(p)$.

The feasible set for this problem is non-empty as $q^e = 1$ satisfies (9), and is compact. Hence, a solution exists by the Weierstrass Theorem. We show that this solution constitutes an optimal PCE as characterized below.

Proposition 3 *$[\theta(\cdot), t(\cdot), p^*]$ is an optimal PCE if and only if the following three conditions hold:*

(i) $p^* = (\mu - c)/(q^*(1 - \ln(q^*))) + c$ where q^* solves (P);

(ii) for all $p \in [(p^* - c)q^* + c, p^*]$*,* $1 - F(\theta(p)) = q(p) = (p^* - c)q^*/(p - c)$ *and* $t(p)$ *is given by the LHS of (9) with* $(p^e, q^e) = (p^*, q^*)$*; and*

(iii) for all $p \in [0, (p^* - c)q^* + c) \bigcup (p^*, 1]$*,* $(\theta(p), t(p)) = (0, 0)$*.*

Moreover, in any equilibrium of Γ_M*, A's payoff is* t^**.*

The expression for $q(p)$ in *(ii)* comes from Proposition 2(b). The expression for $t(p)$, the LHS of (9), obtains because $U_I(p|M)$ in (6) is B's expected payoff when he buys contract $(T_{\theta(p)}, t(p))$ and then buys the good if and only if $v \geq \theta(p)$. Therefore Proposition 2(a)-(b) and Lemma 2 imply that if $[\theta(\cdot), t(\cdot), p^e]$ is an optimal PCE then it must satisfy the constraints of (P), with $t(\cdot)$ equal to the LHS of (9), and A's equilibrium payoff is given by the maximand of (P). Hence the maximized value of (P) is an upper bound for t^*. Furthermore, we show in the Appendix that if q^e satisfies the constraints of (P) then the associated triple $[\theta(\cdot), t(\cdot), p^e]$, with $t(\cdot)$ given by the LHS of (9), is a PCE, with A's equilibrium payoff given by (8). So the maximized value of (P) is also a lower bound for t^*, hence equal to t^*, i.e., $t^* = t(p^*)$.

Finally, any equilibrium of Γ_M must give t^* to A. This is because A could offer a perturbation of the optimal menu, in which $(T_{\theta(p^*)}, t(p^*) - \epsilon)$ replaces $(T_{\theta(p^*)}, t(p^*))$; there is a unique continuation equilibrium and A obtains $t(p^*) - \epsilon$. Since this applies for any small $\epsilon > 0$, A's payoff in any equilibrium must be $t(p^*) = t^*$.

In any equilibrium of Γ_M, therefore, any on-path PCE is either the optimal one characterized in Proposition 3, or outcome-equivalent to that although the off-path signals may not be single-threshold-equivalent. However, $q(p)$ and $U_I(p)$ must be as given by Proposition 3 and (6).

For many distributions F the non-negativity constraint (9) in (P) can be ignored. To see this, note that the fee schedule $t(p)$ for the solution, the LHS of (9), is 0 at $p = (p^* - c)q^* + c$ (because $\theta(p) = 0$ at the bypass price, by Proposition 2(b), and (p^*, q^*) satisfies (7)) and has the first derivative

$$t'(p) = -\theta'(p)\theta(p)F'(\theta(p)) - q(p) - pq'(p) + q^* \frac{p^* - c}{p - c} = \big(p - \theta(p)\big)q^* \frac{p^* - c}{(p - c)^2}, \quad (10)$$

where the second equality follows because $(p-c)q(p)$ is constant, with value $(p^*-c)q^*$. Therefore, if $\theta(p) \leq p$ for all $p \in ((p^* - c)q^* + c, p^*)$, $t'(p) \geq 0$ and so (9) is satisfied.

Proposition 4 below shows that this condition holds[22] if F is uniformly distributed. Moreover, the solution to (P) is unique in this case. Consequently, Γ_M has a unique equilibrium outcome and, as in the equilibrium of Γ_1, the good is traded if and only if the value v is above a threshold θ^* which is in the interval (c, μ) and increases in c.

Proposition 4 *If F is uniformly distributed on $[0,1]$ then (P) has a unique solution and it is obtained by solving (8), ignoring (9). Thus, there is a unique equilibrium outcome of Γ_M. In equilibrium trade occurs if and only if $v \geq \theta^*$, where $\theta^* \in (c, \mu)$ increases in c.*

In fact, we show in our working paper (Evans and Park, 2025) that Proposition 4 can be generalized to a much wider class of distributions, including any distribution F that is not overly concentrated, in the sense that the following holds:

$$F'(v) \leq 2 \quad \text{for all} \quad v \in [0,1]. \tag{11}$$

To illustrate Proposition 4: if F is uniform on $[0,1]$ and $c = 0$ the problem (P) reduces to

$$\max_{q^e \in [\underline{q}(0.5,0),1]} \; q^e\left(1 - \frac{q^e}{2}\right) - p^e q^e \quad \text{where} \quad p^e = \frac{0.5}{q^e(1 - \ln(q^e))}.$$

The solution to this problem is approximately $q^e = 0.63$, $p^e = 0.543$. For this example, as noted in Section 3, the maximum fee which A can charge is 0.07 when he is restricted to offering a single contract, by (3) and (2). He does substantially better in Γ_M since, by (8), the equilibrium fee is approximately 0.089.

Effect of the Adviser on Welfare. The effect of the advisor on welfare, compared to the case in which B is uninformed, is broadly similar to his effect when only a single contract can be offered. If $c \geq \mu$ the effect is identical since a menu is redundant in that case. The following Proposition summarizes the welfare properties for the case in which $c < \mu$. The equilibrium of the menu game, like that of Γ_1, is inefficient. S is strictly worse off than if A were not present (B, of course, obtains no benefit from A since his payoff remains at zero). As in the case of Γ_1, A's effect on total social surplus depends on c: he decreases it if c is low and increases it if c is high.

[22]Note also that if (p^e, q^e) satisfies (7) and $(p^e - c)q^e$ exceeds the monopoly profit $\pi^m \equiv \max_p (p - c)(1 - F(p))$ then (p^e, q^e) satisfies (9). This follows since then $(p - c)\big(1 - F(\theta(p))\big) = (p^e - c)q^e \geq \pi^m \geq (p - c)(1 - F(p))$, so $\theta(p) \leq p$.

Proposition 5 (a) *If $c < \mu$, every equilibrium of Γ_M is inefficient and, compared to a situation in which A is not present and S makes a take-it-or-leave-it price offer to B,*

(i) *S is worse off in Γ_M, and*

(ii) *there exist two thresholds, $\widehat{c}$ and $\widetilde{c}$, where $0 < \widehat{c} \leq \widetilde{c} < \mu$, such that social surplus in Γ_M is strictly lower if $c < \widehat{c}$ and strictly higher if $c > \widetilde{c}$.*

(b) *If Γ_M has a unique equilibrium, obtained by solving (8), ignoring (9), then equilibrium social surplus is lower in Γ_M than in Γ_1.*

In Evans and Park (2025) we show that if F satisfies (11) then $\widehat{c} = \widetilde{c}$. Furthermore, the working paper's generalization of Proposition 4, mentioned above, and Proposition 5(b) jointly imply that the unique equilibrium of Γ_M is less efficient than that of Γ_1 if F satisfies (11). Together with the Corollary to Proposition 1, this implies that if (11) holds, and cost is relatively high, there may be a public policy benefit to allowing a monopoly information firm to operate, but restricting its ability to offer a menu.

6. Comparison with the Full-Information Case

Here we discuss two variants of our model, in each of which the monopoly outcome prevails in the sense that S charges the monopoly price as if B knew the realization of his value v, and A offers a contract that effectively equips B with full information. We then compare the equilibrium welfare with the welfare achieved in our main model. The first variant differs from the main model in that the order of moves of A and S is reversed: first S publicly sets her price p and then A offers B a contract $(\psi, t) \in \mathcal{C}$. (There is no need to offer a menu if p is set first and observed by A. Note also that the case in which S sets p first and p is unobserved by A is accommodated by Γ_M.) In the second variant there are multiple informed third-party advisors who competitively offer contracts to B.

Seller Moves First The analysis of this game is straightforward. For an arbitrary $p \in \mathbb{R}_+$, suppose S has set price p. Then B's reservation payoff (i.e., without A) is $\max\{\mu - p, 0\}$, while his surplus would be maximal at $\int_p^1 (v - p)dF$ when he buys the good if and only if $v \geq p$. Therefore, it is optimal for A to offer the single-threshold signal T_p (or its equivalent, as it has maximal information value for B) for a fee equal to its value of information, $t = \int_p^1 (v - p)dF - \max\{\mu - p, 0\}$. B

will pay the fee and then buy the good if and only if $v \geq p$, generating a profit of $(p - c)(1 - F(p))$ for S. Anticipating this, the seller will charge the monopoly price $p^m(c) \in \arg\max_p (p - c)(1 - F(p))$, as if B knows v. Hence B's expected payoff is $\max\{\mu - p^m(c), 0\}$. The presence of A benefits B in the case in which $p^m(c) < \mu$ since, with no advisor, S would simply charge μ and B's payoff would be zero.

Competitive Advisors Suppose there are multiple competitive advisors, all fully-informed, who can offer any contract. Then it is easy to see that there is an equilibrium in which competition drives the equilibrium fee down to zero for a signal which tells B whether or not v exceeds p for any price p that S may have set, e.g., full information.[23] Anticipating this, S sets the monopoly price $p^m(c)$. The only difference between this case and the previous one is that the buyer retains the consumer surplus (in excess of the buyer's reservation payoff $\max\{\mu - p^m(c), 0\}$), whereas in the previous case the monopoly advisor extracts it as fee.

How does the welfare (in the sense of total surplus) achieved in the full-information monopoly outcome compare with that of our main model Γ_M? We established above that the equilibrium outcome of Γ_M is fully efficient (i.e., the good is traded if and only if $v \geq c$) if $c \geq \mu$ and that it converges to the fully efficient outcome as $c \to \mu$ from below. On the other hand, the monopoly price is strictly greater than c for all $c \in [0, 1)$, hence the level of inefficiency at the monopoly outcome is bounded away from zero for all $c \in [0, \mu]$. Consequently, there is a threshold $\bar{c} < \mu$, which depends on the distribution F, such that welfare is strictly higher in the equilibrium of Γ_M than in the monopoly outcome if $c > \bar{c}$.

In addition, welfare is higher in the equilibrium of Γ_M than in the monopoly outcome if F is uniform because then the equilibrium threshold θ^* in Γ_M satisfies $\theta^* \in (c, \mu)$ by Proposition 4, whereas the monopoly price $p^m(c)$ exceeds μ. In fact, we show in Evans and Park (2025) that the same conclusion holds in a broader class of cases, namely, when F satisfies (11) and either (i) $p^m(c) > \mu$ or (ii) $c > \mu/2$.

Hence, from a welfare perspective, it is often better to have a monopoly advisor (with commitment power) than competitive advisors. However, as noted above, the

[23] This equilibrium is robust in the sense that it exists for a variety of specific extensive forms, e.g., whether advisors move simultaneously or sequentially, whether they may revise offers based on the rival's, whether the market is contestable or not. However, we do not know if this is the unique equilibrium outcome.

buyer is better off with competitive advisors since he can then retain the consumer surplus. He is also strictly better off if S, rather than the monopoly advisor, moves first if F and c are such that $p^m(c) < \mu$.

Proposition 6 (a) *There is a threshold $\bar{c} < \mu$ such that equilibrium total surplus in Γ_M is greater than that in the full-information monopoly outcome if $c > \bar{c}$.*
(b) *The equilibrium total surplus in Γ_M is greater than that in the full-information monopoly outcome if F is uniform on* $[0, 1]$.

7. Buyer-Optimal Signals and Hard Information

In this Section we discuss the relation of our results to those of some recent papers in two related literatures; in the first of these a buyer of a good selects an optimal signal structure before the seller of the good sets the price, and in the second an information firm sells hard information to a seller of an asset or a good.

7.1. *Buyer-Optimal Signals*

In Roesler and Szentes (2017), henceforth RS, B may freely choose any signal of his valuation and then S sets her price after observing B's signal, though not its realization. The buyer-optimal signal is designed in such a way that S faces a unit-elastic demand over an interval of prices, and sets the lowest price.

RS define an *outcome* (referred to below as an *RS-outcome*) as a pair (G, p) where G is a feasible distribution of B's posterior expectation of v (i.e., F is a mean-preserving spread of G) and p is optimal for S given G. If $c = 0$ the buyer-optimal outcome is efficient[24] and gives rise to a unit-elastic demand: the least-informative buyer-optimal signal takes the form $(G^z_{\pi^*}, \pi^*)$, where the cdf $G^z_{\pi^*}(s) = 1 - (\pi^*/s)$ for $s \in [\pi^*, z)$ and $G^z_{\pi^*}(s) = 1$ for $s \geq z$. S is indifferent between all prices in $[\pi^*, z]$, the support of B's posterior, chooses π^*, and trades with probability 1.

Why, in our model, does the adviser not want to offer this buyer-optimal signal? One way to understand the stark difference between our result in Section 4, that

[24]For $c > 0$, buyer-optimal outcomes of RS are not generally efficient; they show that the good is traded whenever valuation exceeds c (Proposition 2 of the online Supplementary Appendix) so any inefficiency is due to too much trade. In contrast, inefficiency in our optimal outcome is due to too little trade (i.e., $c < \hat{\theta}$) when $c < \mu$. The welfare comparison between the two outcomes can go either way. In Example 1 of the Online Appendix of RS, for instance, welfare is higher in their outcome when $c = 0$ but in our outcome when $c = 1/2$.

optimality requires single-threshold equivalence, and that of RS is that the RS signal is designed to make it optimal for the seller to charge a low price. Our advisor, however, does not want to induce too low a price from S because that would enhance the value of buying the good outright for B, reducing B's willingness to pay for the signal offered. In the case where $c = 0$, B would in fact have no incentive to pay any positive price for the RS signal since he would know in advance that its realization would be above S's price π^*. Proposition 1 shows that a single-threshold signal achieves the dual aims of inducing an appropriately high price from S and also a high gross consumer surplus for B, to be extracted via the fee.

RS's result that S's demand function is unit-elastic is at first sight reminiscent of the results of our analysis of the menu game. However, there are several important differences. For example, in RS the seller sets the lowest S-optimal price whereas in our analysis she sets the highest. The most important contrast, however, is that the optimum in the menu game cannot be achieved by a single signal − only by a menu with multiple contracts. This follows from the fact that the optimum single signal is single-threshold-equivalent, by Proposition 1, whereas no single-threshold signal can produce the unit-elastic demand function which, by Proposition 2, is optimal in Γ_M — among prices which induce information purchase, S strictly prefers the highest if the signal is single-threshold.

The following Lemma sheds further light on the relation between RS-outcomes and the equilibrium of Γ_M. We assume that $c = 0$.

Lemma 4 Let (p^*, q^*) be the price and trade probability in an optimal PCE of Γ_M. (i) $(G^{p^*}_{p^*q^*}, p)$ is an RS-outcome for every $p \in [p^*q^*, p^*]$; and (ii) there exists $\tilde{p} \in [p^*q^*, p^*]$ such that B's surplus from the RS-outcome $(G^{p^*}_{p^*q^*}, \tilde{p})$ is $t(p^*)$.

By (i), there exists an RS-outcome which gives the same demand function, price p^* and producer surplus p^*q^* as the equilibrium of Γ_M. However, B's surplus is nil in the RS-outcome (since B buys only if his posterior value is p^*) while in Γ_M it is $t(p^*) > 0$, which is transferred to A. By (ii) there exists an RS-outcome in which B's surplus is $t(p^*)$. Why can A not achieve $t(p^*)$ with this signal? The answer is that $\tilde{p} < \mu$, so that B would prefer to buy the good outright than to pay $t(p^*)$ for the signal.

In Ravid, Roesler and Szentes (2022) B privately acquires a signal at an exogenous cost (increasing in informativeness of the signal) before S sets a price. Since S sets her price without observing B's signal, the game is strategically simultaneous-move. In the benchmark case in which signal acquisition is free, equilibria are Pareto-ranked (under a mild condition on F): in the best equilibrium B learns v and S sets monopoly price, and in the worst one B learns the buyer-optimal signal of RS discussed above but S sets the highest price at which she would obtain the same profit if B had the full information F instead. When signal acquisition is costly, B acquires a signal that generates unit-elastic demand on an interval of prices which S randomizes over; as the cost vanishes this equilibrium converges to the worst equilibrium of the zero-cost benchmark described above. They stress the significant welfare loss when information acquisition is costly, even if the cost is minuscule, as opposed to when it is free. Note that, since the limit outcome of Ravid, Roesler and Szentes (2022) is worse than the monopoly outcome, it follows from Proposition 6 that Γ_M results in a more efficient outcome than the former in the environments defined in that Proposition.

7.2. Hard Information

In Ali, Haghpanah, Lin and Siegel (2022) a monopoly information intermediary (e.g., credit rating agency) sells hard information to a seller in a competitive asset market, by offering a signal on the seller's asset value and a pair of fees, a test fee and a disclosure fee. If the seller pays for a test she receives the signal realization and then decides whether to pay to disclose it to the market. The signal that maximizes the intermediary's payoff in the least favorable continuation equilibrium generates an exponential distribution of its realization, which is therefore very different from our threshold signals, but somewhat reminiscent of the logarithmic form of the buyer's indirect utility function in our menu version. However, the strategic logic at play is very different. The form of their signal is driven by the fact that the intermediary is selling to the seller, who decides whether to pay the disclosure fee after learning the realized signal.[25] In our model, even if the information sold to the buyer is hard, the advisor's payoff is independent of its content.

Asseyer and Weksler (2024) study a monopoly certifier in a common-value setting (where sellers' opportunity costs are correlated with buyers' valuations) with a posi-

[25] If the intermediary can charge only test fees, the problem is simpler and a binary signal is optimal but so are all signals that separate below-mean values from the rest, e.g., the fully revealing one.

tive certification cost. They characterize certifier-optimal designs: any feasible signal maximizing sellers' willingness to pay is optimal, and only a fraction of sellers obtain certification, resulting in partial market disclosure.

Celik and Strausz (2025) examine the interaction between hard and soft information. A monopolist certifier sells hard information about quality to a partially informed seller while screening the seller's private information. They show that it is optimal to fully screen soft information but not fully disclose hard information.

Appendix

Proof of Proposition 1 Define $u_I(.|(\psi,t)), u_o(.), \pi_I(.|(\psi,t))$ and $\pi_o(.)$ as in the text following the statement of Proposition 1. We formalize $\underline{p}(\psi,t)$ and $\bar{p}(\psi,t)$, introduced informally in the main text, as the two price levels such that

$$0 < \underline{p}(\psi,t) < \mu < \bar{p}(\psi,t) \qquad \text{and} \qquad u_I(p|(\psi,t)) = u_o(p).$$

Since $u_o(\cdot)$ is piecewise-linear while $u_I(\cdot|(\psi,t))$ is convex, both price levels are uniquely defined for $t > 0$ so long as $u_I(\mu|(\psi,t)) > 0$.

First, we show that, for any contract $(\psi,t) \in \Psi \times \mathbb{R}$, an equilibrium exists in the continuation game of Γ_1 which follows A's offer of (ψ,t). (Henceforth, *continuation equilibrium* always refers to an equilibrium in the continuation game beginning immediately after A announces a contract). It suffices to show that S's optimal price exists given any contract $(\psi,t) \in \Psi \times \mathbb{R}$ and B's strategy specified below (which is optimal for B). Since S's profit is negative if $p < c$, we focus on $p \in [c,1]$ to identify S's optimal price.

<u>B's strategy:</u> [1] If $t < 0$, so that $u_I(p|(\psi,t)) > u_o(p)$ for all $p \in [0,1]$, accept (ψ,t) for any $p \in [0,1]$; [2] If $u_I(p|(\psi,t)) \leq u_o(p)$ for all $p \in [0,1]$, reject (ψ,t) for any $p \in [0,1]$; [3] Otherwise (i.e., when $\underline{p}(\psi,t)$ and $\bar{p}(\psi,t)$ are well-defined as above), accept (ψ,t) if and only if $p \in (\underline{p}(\psi,t), \bar{p}(\psi,t)]$. After accepting (rejecting, resp.) the contract, B buys the good if and only if the realized signal is $s \geq p$ ($\mu \geq p$, resp.).

In case [2] above, $\max\{c,\mu\}$ is an optimal price for S. In case [3], S faces a demand function $D(p) = 1 - H^{\psi}_{-}(p)$ for $p \in (\underline{p}(\psi,t), \bar{p}(\psi,t)]$ where $H^{\psi}_{-}(p) \equiv \lim_{\rho\uparrow p} H^{\psi}(\rho)$, $D(p) = 1$ if $p \leq \underline{p}(\psi,t)$ and $D(p) = 0$ if $p > \bar{p}(\psi,t)$. In case [1], S faces the same

$D(p)$ where $\underline{p}(\psi, t) = 0$ and $\bar{p}(\psi, t) = 1$. In each case, S's optimal price exists by the next result, proving that a continuation equilibrium exists following any (ψ, t).

<u>Claim 1</u> *Given $D(p)$ above, S's maximum profit exists, denoted by π^*.*

Proof. If $c \geq \bar{p}(\psi, t)$ S has a maximum profit of zero. Since $1 - H^{\psi}_{-}(p)$ is decreasing and left-continuous, so is $D(p)$. Hence, S's profit function $\pi(p) \equiv D(p)(p - c) \geq 0$ is left-continuous and falls at discontinuity points on $[c, \bar{p}(\psi, t)]$, i.e., $\pi(\hat{p}) > \lim_{p \downarrow \hat{p}} \pi(p)$ if π is discontinuous at $\hat{p}$. Let $\pi^* \equiv \sup_{p \in [c, \bar{p}(\psi, t)]} \pi(p)$ and take a convergent sequence $p_n \to p^* \in [c, \bar{p}(\psi, t)]$ such that $\pi(p_n) \to \pi^*$. If π is continuous at p^* then $\pi(p^*) = \pi^*$. If π is discontinuous at p^*, p_n must converge to p^* from below since π falls at p^*, which implies $\pi(p^*) = \pi^*$ because π is left-continuous. This proves the Claim. $\square$

(a) Assume $c \geq \mu$. Suppose A offers $(T_c, \hat{t})$ where $\hat{t} = \int_c^1 (v - c) dF$, i.e., B's expected surplus from buying the good at price c if and only if $v \geq c$. If this contract is offered B will not buy information (and so nor will he buy the good) if $p > c$ because then his surplus from trade would fall short of $\hat{t}$. Thus, it is optimal for S to set price c and, given price c, for B to buy information and then trade if and only if $v \geq c$. In the continuation equilibrium in which S and B behave in this way, the outcome is efficient and A captures all the surplus. For all other contracts the continuation equilibrium is arbitrary. Clearly it is optimal for A to announce $(T_c, \hat{t})$, establishing a strategy profile as described in part (a) as an equilibrium.

To see that every equilibrium of Γ_1 is outcome-equivalent, suppose that A offers $(T_c, \hat{t} - \epsilon)$. If S sets a price just above c B will accept this contract, so S must price so that B accepts with probability 1. Therefore, in any continuation equilibrium, A obtains at least $\hat{t} - \epsilon$ for any small $\epsilon > 0$, so gets $\hat{t}$ in every equilibrium. This implies that trade occurs if and only if $v \geq c$ and B and S each obtain zero, so $p = c$.

(b) The rest of the proof pertains to the case in which $c < \mu$. We make the innocuous assumption that if B is indifferent between buying the good and not, he buys it.

If A offers (F, ϵ), i.e., full information for a small fee $\epsilon > 0$, B accepts (F, ϵ) for every price $p \in (\underline{p}(F, \epsilon), \bar{p}(F, \epsilon))$. $(\underline{p}(F, \epsilon), \bar{p}(F, \epsilon)) \to (0, 1)$ as $\epsilon \to 0$, so, for sufficiently small $\epsilon > 0$, S optimally sets the monopoly price, $\arg\max_p (p - c)(1 - F(p))$, and A obtains payoff ϵ. To find equilibrium contracts, therefore, it suffices to consider continuation equilibria following (the announcement of) some contract (ψ, t) with $t > 0$, in which (ψ, t) is accepted with strictly positive probability.

Consider such a contract (ψ, t) with such a continuation equilibrium, denoted by $\hat{\sigma}$. Let $\hat{\pi}$ be S's equilibrium payoff in $\hat{\sigma}$. $\pi_I(\mu|(\psi, t)) > 0$, since $c < \mu$ and $H_{-}^{\psi}(\mu) < 1$, so $\hat{\pi} > 0$ since if S sets price μ his profit is strictly positive whether B accepts (ψ, t) or not. Claim 2 below implies that, for the purpose of identifying A's highest continuation equilibrium payoff, we can assume that B accepts (ψ, t) for sure at $\bar{p}(\psi, t)$ and buys the good outright at $\underline{p}(\psi, t)$.

<u>Claim 2</u> *There exists a continuation equilibrium σ following announcement of (ψ, t) in which (i) B accepts the contract if and only if $p \in (\underline{p}(\psi, t), \bar{p}(\psi, t)]$; (ii) S sets p such that $\pi_I(p|(\psi, t)) \geq \pi_o(\underline{p}(\psi, t))$ and p maximizes (5); and (iii) A's payoff is t.*

Proof. $u_I(\mu|(\psi, t)) \geq 0$. Suppose $u_I(\mu|(\psi, t)) > 0$. There exists a price $\hat{p} \in [\underline{p}(\psi, t), \bar{p}(\psi, t)]$, which gives payoff $\hat{\pi}$ to S, such that B accepts (ψ, t) with strictly positive probability when price is $\hat{p}$, hence, since $t > 0$, buys the good with probability below 1. If $\hat{p} = \underline{p}(\psi, t)$, since $\hat{\pi} > 0$ implies $\hat{p} = \underline{p}(\psi, t) > c$, S can do strictly better by shading her price slightly and selling for sure. Hence $\hat{p} \in (\underline{p}(\psi, t), \bar{p}(\psi, t)]$. It cannot be that $u_I(\mu|(\psi, t)) = 0$ since then S's equilibrium price in $\hat{\sigma}$ would be μ but, as above, since B would buy the good with probability strictly below 1, S could profitably deviate by shading price.

Define σ as the profile in which B's strategy is as in (i) of Claim 2 and S sets price $\hat{p}$. The Claim follows if $\hat{p}$ maximizes (5) and $\pi_I(\hat{p}|(\psi, t) \geq \pi_o(\underline{p}(\psi, t))$ since $\hat{p}$ is then optimal for S given B's strategy. Since $\hat{p}$ achieves $\hat{\pi}$ in $\hat{\sigma}$ and B accepts (ψ, t) for sure at $p \in (\underline{p}(\psi, t), \bar{p}(\psi, t))$, $\hat{\pi} \geq \lim_{p \to \bar{p}(\psi, t)} \pi_I(p|(\psi, t)) = \pi_I(\bar{p}(\psi, t)|(\psi, t))$, by left-continuity of π_I, so $\hat{p}$ maximizes (5). Also $\pi_I(\hat{p}|(\psi, t)) \geq \pi_o(\underline{p}(\psi, t))$ since in $\hat{\sigma}$ S could choose p arbitrarily close to, and below, $\underline{p}(\psi, t)$ and sell for sure. $\square$

Claim 2 implies that $t > 0$ is the maximal payoff for A over all continuation equilibria following contracts which involve ψ if and only if t is the maximal fee achievable with ψ, as defined in the main text.

Now, we prove the Equal Profit Principle defined in the main text. Consider any $\psi \in \Psi$ such that some $t > 0$ is achievable with ψ, i.e., $\pi_I(p|(\psi, t)) \geq \pi_o(\underline{p}(\psi, t))$ where p maximizes (5). This implies $\pi_I(p|(\psi, t)) \geq \pi_I(\mu|(\psi, t)) > 0$. By definition, t is the maximal fee achievable with ψ if and only if $t + \epsilon$ is not achievable with ψ for any

$\epsilon > 0$. Characterizing when this is the case establishes the Equal Profit Principle.

To check if $t+\epsilon$ is achievable with ψ, note that if A offers contract $(\psi, t+\epsilon)$ then $u_I(\cdot|(\psi, t+\epsilon)) = u_I(\cdot|(\psi,t)) - \epsilon$ so, by continuity, the interval of prices for which B would accept $(\psi, t+\epsilon)$ shrinks continuously, as ϵ increases from 0:

[A] $\underline{p}(\psi, t+\epsilon)$ increases and $\bar{p}(\psi, t+\epsilon)$ decreases in ϵ, both continuously and both strictly, until they reach μ at some $\tilde{t} = t + \tilde{\epsilon} > t$.

First, consider the case that $\pi_I(p|(\psi,t)) = \pi_o(\underline{p}(\psi,t))$ where p maximizes (5). By [A], $\pi_o(\underline{p}(\psi, t+\epsilon)) > \pi_o(\underline{p}(\psi,t))$ and $\pi_I(p'|(\psi,t+\epsilon)) = \pi_I(p'|(\psi,t)) \leq \pi_I(p|(\psi,t))$ for every $p' \in (\underline{p}(\psi,t+\epsilon), \bar{p}(\psi,t+\epsilon)]$ if $\epsilon \in (0,\tilde{\epsilon})$, while $(\underline{p}(\psi,t+\epsilon), \bar{p}(\psi,t+\epsilon)] = \emptyset$ if $\epsilon \geq \tilde{\epsilon}$. Hence, $t+\epsilon$ is not achievable with ψ for any $\epsilon > 0$, establishing that $t > 0$ is the maximal achievable fee with ψ if $\pi_I(p|(\psi,t)) = \pi_o(\underline{p}(\psi,t)) > 0$.

Next, consider the case that $\pi_I(p|(\psi,t)) > \pi_o(\underline{p}(\psi,t))$ where p maximizes (5). For $\epsilon \geq 0$, let $S(\epsilon) \equiv \max_{p' \in (\underline{p}(\psi,t+\epsilon),\bar{p}(\psi,t+\epsilon)]} \pi_I(p'|(\psi,t+\epsilon))$ and let $\Delta(\epsilon) \equiv S(\epsilon) - \pi_0(\underline{p}(\psi,t+\epsilon))$. $S(\epsilon)$ is well-defined if $\pi_I(p'|(\psi,t+\epsilon)) \geq \pi_0(\underline{p}(\psi,t+\epsilon))$ for some $p' \in (\underline{p}(\psi,t+\epsilon), \bar{p}(\psi,t+\epsilon)]$ since π_I is left-continuous and falls at discontinuity points and $\pi_I(\underline{p}(\psi,t+\epsilon)|(\psi,t+\epsilon)) \leq \pi_0(\underline{p}(\psi,t+\epsilon))$. In the considered case, $\pi_I(p|(\psi,t)) - \pi_o(\underline{p}(\psi,t)) = \Delta(0) > 0$. By [A], as ϵ increases from 0, (i) $\pi_o(\underline{p}(\psi,t+\epsilon))$ increases continuously, and (ii) $S(\epsilon)$ exists and weakly decreases in ϵ so long as $\Delta(\epsilon) > 0$.

If $\Delta(\epsilon) > 0$, for any $p'' \in P(\epsilon) \equiv \arg\max_{p' \in (\underline{p}(\psi,t+\epsilon),\bar{p}(\psi,t+\epsilon)]} \pi_I(p'|(\psi,t+\epsilon))$ we have $\underline{p}(\psi,t+\epsilon) < p''$. As ϵ increases slightly, therefore, $S(\epsilon)$ stays constant if $P(\epsilon) \neq \{\bar{p}(\psi,t+\epsilon)\}$, and, since $\pi_I(\cdot|(\psi,t+\epsilon))$ is left-continuous at $\bar{p}(\psi,t+\epsilon)$, $S(\epsilon)$ changes continuously if $P(\epsilon) = \{\bar{p}(\psi,t+\epsilon)\}$. So $\Delta(\epsilon)$ decreases continuously in ϵ and hits 0 at some $\bar{\epsilon} < \tilde{\epsilon}$, i.e., $\Delta(\bar{\epsilon}) = 0$, because $\lim_{\epsilon\to\tilde{\epsilon}} \pi_o(\underline{p}(\psi,t+\epsilon)) = \mu - c > \pi_I(\mu|(\psi,t)) = \lim_{\epsilon\to\tilde{\epsilon}} \sup_{p' \in (\underline{p}(\psi,t+\epsilon),\bar{p}(\psi,t+\epsilon)]} \pi_I(p'|(\psi,t+\epsilon))$; the inequality applies since B accepts (ψ,t) given price μ, hence must buy the good with probability strictly below 1.

This shows that $t+\epsilon$ is achievable with ψ for all $\epsilon \in (0,\bar{\epsilon}]$, but not achievable for any $\epsilon > \bar{\epsilon}$. Hence, $t+\bar{\epsilon}$ is the maximal achievable fee with ψ and satisfies $S(\bar{\epsilon}) = \pi_0(\underline{p}(\psi,t+\bar{\epsilon})) > 0$, proving the Equal Profit Principle.

We are now ready to prove part (b) of Proposition 1. Consider a single-threshold contract (T_θ, t), where $\theta \in (0,1)$, which has a subsequent continuation equilibrium in which S sets p which maximizes $\pi_I(.|(T_\theta,t))$ on $(\underline{p}(T_\theta,t), \bar{p}(T_\theta,t)]$, and A's payoff is $t > 0$. Trade probability is constant on $(\underline{p}(\psi,t), \bar{p}(\psi,t)]$ so S's optimal price in this interval is $\bar{p}(\psi, t)$. $u_I(p|(T_\theta, t)) = \int_\theta^1 v dF - p(1-F(\theta)) - t$. Since $\underline{p}(T_\theta, t)$ and $\bar{p}(T_\theta, t)$ satisfy $u_I(p|(T_\theta, t)) = \mu - p$ and $u_I(p|(T_\theta, t)) = 0$, respectively,

$$\underline{p}(T_\theta, t) \;=\; \frac{\int_0^\theta v dF + t}{F(\theta)} \quad \text{and} \quad \bar{p}(T_\theta, t) \;=\; \frac{\int_\theta^1 v dF - t}{1 - F(\theta)}.$$

By the Equal Profit Principle A's maximal payoff achievable with T_θ, denoted by $t(\theta)$, satisfies $\pi_o(\underline{p}(T_\theta, t(\theta))) = \pi_I(\bar{p}(T_\theta, t(\theta))|(T_\theta, t(\theta)))$, i.e.,

$$\underline{p}(T_\theta, t(\theta)) - c \;=\; \big(\bar{p}(T_\theta, t(\theta)) - c\big)(1 - F(\theta)),$$

so, after rearrangement, we get (3). An optimal threshold $\widehat{\theta}$ maximizes $t(\theta)$. Since

$$t'(\theta) \;=\; \Big[-\theta + \frac{\mu + 2cF(\theta) + cF(\theta)^2}{(1+F(\theta))^2} \Big] F'(\theta),$$

$t'(\theta) = 0$ if and only if (2) holds. The LHS of (2) is negative for $\theta < c$ and strictly increases from 0 when $\theta = c$ to $4(1-c)$ when $\theta = 1$, so (2) has a unique solution $\widehat{\theta}$ and $\widehat{\theta} \in (c, \mu)$. Since $t(0) = 0, t(1) = (c-\mu)/2 < 0$ and $t'(0) = \mu F'(0) > 0$, $t(\theta)$ is a maximum and strictly positive at $\widehat{\theta}$. This establishes that $\widehat{\theta}$ exists and is the unique solution to (2). Hence $t(\widehat{\theta})$ is the maximum payoff achievable by A in any continuation equilibrium following announcement of a single-threshold contract, and it is achieved in the equilibrium following $(T_{\widehat{\theta}}, t(\widehat{\theta}))$, in which S sets price $\bar{p}(T_{\widehat{\theta}}, t(\widehat{\theta}))$.

$\widehat{\theta}$ increases in c because the partial derivatives of $(\theta - c)(1 + F(\theta))^2 + c$, with respect to θ and c, are of opposite sign for $\theta > c$. S's optimal price is

$$\bar{p}(T_{\widehat{\theta}}, t(\widehat{\theta})) \;=\; \frac{\int_{\widehat{\theta}}^1 v dF - t(\widehat{\theta})}{1 - F(\widehat{\theta})} \;=\; \frac{\mu - c[F(\widehat{\theta})]^2}{1 - [F(\widehat{\theta})]^2} \;>\; \mu$$

where the second equality is from (3) and the inequality from $c < \mu$; S's expected payoff is $(\bar{p}(T_{\widehat{\theta}}, t(\widehat{\theta})) - c)(1 - F(\widehat{\theta})) = \frac{\mu - c}{1 + F(\widehat{\theta})} = (\widehat{\theta} - c)(1 + F(\widehat{\theta}))$ from (2).

The argument in the text shows that in any equilibrium of Γ_1, if one exists, A's payoff is at least $t(\widehat{\theta})$. Next we show that $t(\widehat{\theta})$ is the maximal achievable fee with any contract in Ψ, which will establish existence of an equilibrium of Γ_1 as specified in Proposition 1.

Suppose, to the contrary, that $t > t(\widehat{\theta})$ is achievable with some $\psi \in \Psi$. By the Equal Profit Principle, we may assume that t is the maximal achievable fee with ψ and $\pi_o(\underline{p}(\psi, t)) = \pi_I(p|(\psi, t)) > 0$ where p maximizes (5). We show that then t is achievable with some single-threshold contract, a contradiction.

Define q_* and θ_* as in the main text, i.e., $q_* = 1 - H^\psi(\underline{p}(\psi,t)) = 1 - F(\theta_*)$. Conditional on trade probability q_*, total surplus is maximized when trade takes place if and only if $v \geq \theta_*$, so

$$\int_{\theta_*}^{1} (v-c)dF(v) \geq \int_{[\underline{p}(\psi,t),1]} (s-c)dH^\psi(s) \geq \int_{[p,1]} (s-c)dH^\psi(s) \tag{12}$$

where the second inequality is due to $c \leq \underline{p}(\psi,t)$.

If the contract is (T_{θ_*}, t), S's price is $\underline{p}(\psi,t)$ and B accepts the contract, then B will buy the good if and only if $v \geq \theta_*$, because, conditional on $v < \theta_*$, his expected valuation is less than his valuation, under signal ψ, conditional on realizations for which he does not buy, i.e., $\int_0^{\theta_*} vdF(v) \leq \int_{[0,\underline{p}(\psi,t))} sdH^\psi(s)$, by (12) and $H^\psi(\underline{p}(\psi,t)) = F(\theta_*)$. Therefore, $u_I(\underline{p}(\psi,t)|(T_{\theta_*},t))$ equals

$$\begin{aligned} \int_{\theta_*}^{1} (v-c)dF(v) - t - (\underline{p}(\psi,t)-c)q_* &\geq \int_{[\underline{p}(\psi,t),1]} (s-c)dH^\psi(s) - t - (\underline{p}(\psi,t)-c)q_* \\ &= u_I(\underline{p}(\psi,t)|(\psi,t)) = u_0(\underline{p}(\psi,t)), \end{aligned}$$

which implies $\underline{p}(T_{\theta_*},t) \leq \underline{p}(\psi,t)$. Given (T_{θ_*},t), S's optimal price in $(\underline{p}(T_{\theta_*},t), \bar{p}(T_{\theta_*},t)]$ is $p^* = \bar{p}(T_{\theta_*},t)$, and $u_I(p^*|(T_{\theta_*},t)) = 0$. $\pi_I(p^*|(T_{\theta_*},t))$ is equal to

$$\int_{\theta_*}^{1} (v-c)dF(v) - t \geq \int_{[p,1]} (s-c)dH^\psi(s) - t - u_I(p|(\psi,t)) = \pi_I(p|(\psi,t)). \tag{13}$$

Hence $\pi_I(p^*|(T_{\theta_*},t)) \geq \pi_I(p|(\psi,t)) \geq \underline{p}(\psi,t) - c \geq \underline{p}(T_{\theta_*},t) - c = \pi_o(\underline{p}(T_{\theta_*},t))$, i.e., t is achievable with T_{θ_*}.

This shows that there is no continuation equilibrium in which A's payoff exceeds $t(\widehat{\theta})$, the best available from a single-threshold contract. Therefore there is an equilibrium of Γ_1 in which A offers $(T_{\widehat{\theta}}, t(\widehat{\theta}))$ and, after every other offer, the continuation equilibrium is arbitrary. In addition, A's payoff is $t(\widehat{\theta})$ in every equilibrium.

It remains to prove outcome-equivalence of all equilibria. Consider another equilibrium in which A announces a contract (ψ,t) such that $t = t(\widehat{\theta})$ and, by the Equal Profit Principle, S sets p that maximizes (5) in the continuation equilibrium. Then, t is achievable with the threshold contract (T_{θ_*},t) by the same argument as above. Unless both inequalities in (12) are tight $\pi_I(p^*|(T_{\theta_*},t)) > \pi_I(p|(\psi,t))$, in which case $\pi_I(p^*|(T_{\theta_*},t)) > \pi_o(\underline{p}(T_{\theta_*},t))$, since $\underline{p}(T_{\theta_*},t) \leq \underline{p}(\psi,t)$, and thus, by the Equal Profit Principle, there exists an equilibrium in which A obtains more than $t = t(\widehat{\theta})$. This contradiction shows that the inequalities in (12) hold as equalities. Therefore, since $\theta_* = \widehat{\theta}$, under signal ψ trade occurs if and only if $v \geq \widehat{\theta}$ both at price $\underline{p}(\psi,t)$ and price p. Thus, the surplus achieved in each equilibrium is the same. $\pi_I(p^*|(T_{\theta_*},t)) = \pi_I(p|(\psi,t))$, i.e. S's payoff is the same in the two equilibria so B's payoff is also the same (i.e. zero), which implies $p = \bar{p}(T_{\widehat{\theta}}, t(\widehat{\theta}))$. So the two equilibria are outcome-equivalent. QED.

Proof of Lemma 1 Consider a menu $M \in \Upsilon$ such that A obtains t^e in the continuation equilibrium, denoted by σ. Following M, there is an on-path price p^e after which A's expected payoff is $\tilde{t} \geq t^e$. Let $\tilde{q}$ denote the probability of trade after p^e and let $1 - F(\widetilde{\theta}) = \tilde{q}$. Suppose that A announces the menu $M' \equiv M \bigcup \{(T_{\widetilde{\theta}}, \tilde{t})\}$ and in the continuation strategy profile, denoted σ': (i) B selects $(T_{\widetilde{\theta}}, \tilde{t})$ after p^e (and buys the good iff $v \geq \tilde{\theta}$) while, after $p \neq p^e$, B selects $(T_{\widetilde{\theta}}, \tilde{t})$ if it is optimal but, if not, selects an optimal contract such that $\pi(p|\sigma') \leq \pi(p|\sigma)$, where $\pi(p|\sigma)$ is S's profit in the strategy profile σ at price p, and $\pi(p|\sigma')$ is similarly defined; (ii) S sets $\tilde{p} = \max\{p | B \text{ selects } (T_{\widetilde{\theta}}, \tilde{t}) \text{ after } p\}$, which exists because $u_I(p|(\psi,t))$ is continuous in p for every $(\psi,t) \in M'$. B's strategy is optimal since $T_{\widetilde{\theta}}$ maximizes his payoff conditional on trading with probability $\tilde{q}$, so his expected payoff from $(T_{\widetilde{\theta}}, \tilde{t})$ is at least as high as from any element of M after p^e. S's strategy is optimal because $\pi(\tilde{p}|\sigma') = (\tilde{p} - c)\tilde{q} \geq \pi(p^e|\sigma') = (p^e - c)\tilde{q} = \pi(p^e|\sigma) \geq \pi(p|\sigma')$ for all $p \neq \tilde{p}, p^e$. Hence, A's payoff is $\tilde{t}$. This shows that there is a pure strategy continuation equilibrium σ' following M' in which A's payoff is at least t^e. QED.

Proof of Lemma 2 For $M = \{(\boldsymbol{\psi}(p), t(p))\}_{p \in [0,1]}$, let $t' = \inf\{t | (\psi,t) \in M\}$ and suppose that $t' < 0$. $U_I(p|M) \geq \max\{\mu - p, 0\} - t' > \max\{\mu - p, 0\}$ for all $p \in [0,1]$ since B can buy a contract with approximate fee t' and ignore the information. Construct a new menu M_1 defined by $(\psi,t) \in M \setminus \{(T_0,0)\}$ if and only if $(\psi, t+\epsilon) \in M_1 \setminus \{(T_0,0)\}$, where $\epsilon > 0$ is sufficiently small that $U_I(p|M_1) > \max\{\mu - p, 0\}$ for all $p \in [0,1]$. There exists an equilibrium for M_1 in which, for each $p \in [0,1]$, B's choice (and, hence, S's profit) is the same as in the original equilibrium, since B's payoff from each non-null choice is reduced by ϵ while the null contract is dominated, and S chooses p^e as before. A's payoff is higher by ϵ in this equilibrium, contradicting the optimality of M. QED

Proof of Lemma 3 (a) For each price p, define $\tilde{\boldsymbol{\psi}}(p)$ as in the text and consider menu $[\tilde{\boldsymbol{\psi}}(.), \tilde{t}(.)]$ where

$$\tilde{t}(p) = \int_{\theta(p)}^{1} v dF - \int v[pr_{\boldsymbol{\psi}(p)}(s \geq p|v)]dF + t(p) \tag{14}$$

and $pr_{\psi}(\cdot|\cdot)$ is conditional probability for signal ψ. If price is p, then, for any $p' \in [0,1]$, B's payoff from accepting $(\tilde{\boldsymbol{\psi}}(p'), \tilde{t}(p'))$ and obeying is $\int_{\theta(p')}^{1} v dF - q(p')p - \tilde{t}(p')$, i.e., using (14), $\int v[pr_{\boldsymbol{\psi}(p')}(s \geq p'|v)]dF - q(p')p - t(p')$, which equals the payoff from accepting $(\boldsymbol{\psi}(p'), t(p'))$ and then buying the good if and only if $s \geq p'$. Thus, since $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$ is a PCE, B cannot benefit by deviating from $(\tilde{\boldsymbol{\psi}}(p), \tilde{t}(p))$, subject to obeying the recommendation. His best payoff obtainable by ignoring information is $\max\{\mu - p, 0\} - \min_{p'} \tilde{t}(p') \leq \max\{\mu - p, 0\} - \min_{p'} t(p')$, a deviation payoff available in the original PCE. Hence it is an optimal strategy for B to choose $(\tilde{\boldsymbol{\psi}}(p), \tilde{t}(p))$ if price is p. p^e is optimal for S since sale probability for each p is $q(p)$, as in the original PCE. Hence $[\tilde{\boldsymbol{\psi}}(\cdot), \tilde{t}(\cdot), p^e]$ is an optimal PCE. (b) follows since $\int_{\theta(p^e)}^{1} v dF > \int v[pr_{\psi(p^e)}(s \geq p^e|v)]dF$ if $(\boldsymbol{\psi}(p^e), t(p^e), p^e)$ is not single-threshold-equivalent. QED

Proof of Proposition 2 Consider an optimal PCE $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$, where $t(p) \geq 0$ for all $p \in [0,1]$ by Lemma 2 and $(\boldsymbol{\psi}(p^e), t(p^e), p^e)$ is single-threshold-equivalent by Lemma 3. We show below that if any of the conditions (a)–(c) fails, there is another PCE in which A's payoff is strictly higher.

Lemma 3 shows that, for any optimal $[\boldsymbol{\psi}(\cdot), t(\cdot), p^e]$, there is a menu of threshold contracts with an outcome-equivalent PCE which has the same trade function $q(\cdot)$, utility schedule $U_I(\cdot|M)$, hence the same interval $[\underline{p}(M), \bar{p}(M)]$ and profit levels $\pi(p|M)$. To prove (a)–(c), therefore, it is without loss to consider PCE's consisting of only single-threshold signals, which we denote by $[\theta(\cdot), t(\cdot), p^e]$, where $T_{\theta(p)}$ is the contract chosen by B for price p. We will also, where the meaning is clear, write θ for T_θ. Let $q^e = q(p^e)$ and let $M = \{(\theta(p), t(p))\}_{p \in [0,1]}$. Note that S obtains a strictly positive profit from any price $p \in (c, \mu)$, because the posterior expectation, s, exceeds

p with strictly positive probability for any signal. Hence $(p^e - c)q^e > 0$, so that $p^e > c$. Moreover, $\underline{p}(M) - c \leq (p^e - c)q^e$ by optimality of p^e for S.

Proof of part (a). The claim is equivalent to $U_I(p^e|M) = 0$. Suppose, to the contrary, that $U_I(p^e|M) > 0$. Let M_2 be a menu obtained from M by dropping all thresholds $\theta > \theta(p^e)$. $U_I(\cdot|M_2)$ coincides with $U_I(\cdot|M)$ for $p \leq p^e$ and its graph is linear with slope $-q^e$ for $p \in [p^e, \tilde{p}]$, where $U_I(\tilde{p}|M_2) = 0$. Given menu M_2 there is an equilibrium continuation in which B selects $\theta(p)$ for all $p \in [\underline{p}(M), p^e]$ and selects $\theta(p^e)$ for all $p \in [p^e, \tilde{p}]$. S's optimal price is $\tilde{p}$ since the probability of trade is constant at q^e on $[p^e, \tilde{p}]$, and the highest profit S can obtain by setting a price $p' < p^e$ is at most $(p^e - c)q^e < (\tilde{p} - c)q^e$. In particular, $\underline{p}(M) - c < (\tilde{p} - c)q^e$. A's payoff is $t(p^e)$. Now consider a menu $M_2(\epsilon)$ which is the same as M_2 except that all fees for non-null contracts have been increased by the same small $\epsilon > 0$. Then $\underline{p}(M_2(\epsilon))$ is slightly above $\underline{p}(M)$ and $\bar{p}(M_2(\epsilon))$ is slightly below $\tilde{p}$. By continuity $\underline{p}(M_2(\epsilon)) - c < (\bar{p}(M_2(\epsilon)) - c)q^e$ for small $\epsilon > 0$. There is a PCE following $M_2(\epsilon)$ in which (i) B chooses $\theta(p)$ for each $p \in (\underline{p}(M_2(\epsilon)), p^e)$ and chooses $\theta(p^e)$ for $p \in [p^e, \bar{p}(M_2(\epsilon))]$, (ii) S charges $\bar{p}(M_2(\epsilon))$ and sells with probability q^e, and (iii) A's payoff is $t(p^e) + \epsilon > t(p^e)$. Thus there is another PCE where A's payoff is higher, proving (a).

Proof of part (c). The claim is equivalent to $\theta(p^e) \geq c$. Suppose instead that $\theta(p^e) < c$. Consider menu M_3 obtained by adding $(\theta', t(p^e))$ to M, where $\theta' = \theta(p^e) + \epsilon \in (\theta(p^e), c)$ and ϵ is small. Then $u_I(\cdot|(\theta', t(p^e)))$ is slightly flatter than $u_I(\cdot|(\theta(p^e), t(p^e)))$ and slightly greater at all $p \in [\underline{p}(M), p^e] \cap [c, p^e]$ since $\theta(p^e) < \theta' < c < p^e$, so $\bar{p}(M_3)$ is slightly higher than $\bar{p}(M)$ while $\underline{p}(M_3) \leq \underline{p}(M)$. If S sets $\bar{p}(M_3)$ then B optimally selects $(\theta', t(p^e))$ and S's profit strictly exceeds $(p^e - c)q^e$ because B's payoff is zero but total surplus is higher than under M. It follows that if the fee for θ' were increased by small $\eta > 0$ S would still price so that B selects $(\theta', t(p^e) + \eta)$, i.e., there is another PCE where A's payoff is higher, proving (c).

Proof of part (b). Denote $(p^e - c)q^e + c$ by $\underline{p}$. $(p - c)q(p) \leq (p^e - c)q^e$ for any $p \in [\underline{p}, p^e]$ by optimality of p^e. Furthermore $U_I(\underline{p}|M) \geq \mu - \underline{p}$ since $\underline{p}(M) - c \leq (p^e - c)q^e$ implies $\underline{p}(M) \leq \underline{p}$.

Define a differentiable function $g : [\underline{p}, p^e] \to \mathbb{R}$ by

$$g'(p) = \frac{-(p^e - c)q^e}{(p - c)} \quad \forall\, p \in [\underline{p}, p^e] \quad \text{and} \quad g(\underline{p}) = \mu - \underline{p}. \tag{15}$$

Let $\bar{p}(g)$ satisfy $g(\bar{p}(g)) = 0$. Then g is a lower bound for U_I on $[\underline{p}, \bar{p}(g)]$. To see this, suppose that $U_I(p'|M) < g(p')$ for $p' \in [\underline{p}, \bar{p}(g)]$. Since $U_I(\underline{p}|M) \geq g(\underline{p})$, it must be that, for some $p'' \in [\underline{p}, p']$ at which U_I is differentiable, $-U_I'(p''|M) > -g'(p'')$ and so $q(p'') > -g'(p'')$. So $(p'' - c)q(p'') > (p^e - c)q^e$ by (15), contradicting optimality of p^e.

Suppose there is a menu $\widehat{M}$, consisting of threshold contracts with nonnegative fees, such that $U_I(p|\widehat{M}) = g(p)$ for all $p \in [\underline{p}, \bar{p}(g)]$. Then, given $\widehat{M}$, S would be indifferent between all prices in $[\underline{p}, \bar{p}(g)]$ and her optimal profit would be the same as that given M, namely $(p^e - c)q^e$. Therefore there would be a PCE following $\widehat{M}$, denoted by $[\hat{\theta}(\cdot), \hat{t}(\cdot), \hat{p}^e = \bar{p}(g)]$, in which S charges $\hat{p}^e$ and B buys with probability $\hat{q}^e = -g'(\hat{p}^e)$. S's profit would be $(\hat{p}^e - c)\hat{q}^e = (p^e - c)q^e$ and B's payoff would be zero, since $g(\hat{p}^e) = 0$. These are the same payoffs as in the PCE following M.

Claim 3. If there exists a menu $\widehat{M}$ and PCE as above and $U_I(p|M) \neq g(p)$ at some $p \in [\underline{p}, \hat{p}^e]$, another PCE exists that gives A strictly more than $t(p^e)$.

Proof: Suppose that $\widehat{M}$ exists and $U_I(p|M) \neq g(p)$ for some $p \in [\underline{p}, \hat{p}^e]$. Then $\hat{p}^e < p^e$ since $U_I'(p|M) \geq g'(p)$ for all $p \in [\underline{p}, \hat{p}^e]$ such that U_I is differentiable (because $-U_I'(p|M) = q(p)$ and $(p - c)q(p) \leq (p^e - c)q^e$). By (a) and (c), $U_I(p^e|M) = 0$ and $\theta(p^e) \geq c$. From (15), $-g'(p^e) = q^e$ and consequently, $\hat{q}^e > q^e$ so that $\hat{\theta}(\hat{p}^e) < \theta(p^e)$.

If $\hat{\theta}(\hat{p}^e) \in [c, \theta(p^e))$, then A's fee in the PCE following $\widehat{M}$ strictly exceeds that following M, because total surplus is higher while the payoffs of S and B are the same, at $(p^e - c)q^e$ and zero, respectively. Hence suppose $\hat{\theta}(\hat{p}^e) < c$.

Let $\tilde{t} = \int_c^1 (v - c)dF - (p^e - c)q^e \geq \int_{\theta(p^e)}^1 (v - c)dF - (p^e - c)q^e = t(p^e)$. If $p = \bar{p}(T_c, \tilde{t})$ and B accepts $(T_c, \tilde{t})$, payoffs for S and B are $(p^e - c)q^e$ and 0, as in the PCE following M. Note that $u_I(\hat{p}^e|(T_c, \tilde{t})) > 0$ because $(\hat{p}^e - c)(1 - F(c)) < (p^e - c)q^e$.

(1) Suppose $u_I(\underline{p}|(T_c, \tilde{t})) > \mu - \underline{p}$, so that $\underline{p}(T_c, \tilde{t}) < \underline{p}$. For small enough $\epsilon > 0$ and menu $\{(T_c, \tilde{t} + \epsilon), (T_0, 0)\}$, there is a PCE in which S selects $\bar{p}(T_c, \tilde{t} + \epsilon)$, B accepts $(T_c, \tilde{t} + \epsilon)$, and A gets $\tilde{t} + \epsilon > t(p^e)$.

(2) If $u_I(\underline{p}|(T_c, \tilde{t}) \leq \mu - \underline{p}$, let $\tilde{t}' = \hat{t}(\hat{p}^e) - u_I(\hat{p}^e|(T_c, \tilde{t}))$ so $u_I(\hat{p}^e|(\hat{\theta}(\hat{p}^e), \tilde{t}')) = u_I(\hat{p}^e|(T_c, \tilde{t}))$. $\hat{\theta}(\hat{p}^e) < c$, so $u_I(p|(\hat{\theta}(\hat{p}^e), \tilde{t}'))$ crosses $u_I(p|(T_c, \tilde{t}))$ at $p = \hat{p}^e$ from above.

(i) Suppose $u_I(\underline{p}|(\widehat{\theta}(\hat{p}^e), \tilde{t}')) > \mu - \underline{p}$. Let $\tilde{t}'' > \tilde{t}'$ be such that (α) $u_I(\underline{p}|(\widehat{\theta}(\hat{p}^e), \tilde{t}'')) = \mu - \underline{p} + \epsilon$ and $\epsilon > 0$ is small enough that (β) $\epsilon < (\underline{p} - c)(F(\widehat{\theta}(\hat{p}^e)) - F(\widehat{\theta}(\underline{p})))$. Note that $\tilde{t}'' \geq 0$. This is because, at price $\underline{p}$: first, total surplus is higher with $(\hat{\theta}(\hat{p}^e), \tilde{t}'')$ than with $(\hat{\theta}(\underline{p}), \hat{t}(\underline{p}))$ since $\underline{p} < \hat{p}^e$, so $\hat{\theta}(\underline{p}) < \hat{\theta}(\hat{p}^e)$; second, by (α), B's payoff is higher by ϵ; and, third, by (β), S's payoff is lower by more than ϵ. Hence $\tilde{t}'' > \hat{t}(\underline{p})$, and $\hat{t}(\underline{p}) \geq 0$ by assumption. Let $u_I(p'|(\hat{\theta}(\hat{p}^e), \tilde{t}'')) = u_I(p'|(T_c, \tilde{t}))$, so $p' < \hat{p}^e$. Given $M' = \{(\hat{\theta}(\hat{p}^e), \tilde{t}''), (T_c, \tilde{t}), (T_0, 0)\}$ S's potential optimal prices are $\bar{p}(T_c, \tilde{t})$, (giving payoff $(p^e - c)q^e$), p' (giving $(p' - c)\hat{q}^e < (\hat{p}^e - c)\hat{q}^e = (p^e - c)q^e$) and $\underline{p}(M')$ (giving $\underline{p}(M') - c < \underline{p} - c = (p^e - c)q^e$, where the inequality follows from (α)). Hence there is a PCE following M' in which A obtains $\tilde{t}$.

(ii) Suppose $u_I(\underline{p}|(\hat{\theta}(\hat{p}^e), \tilde{t})) \leq \mu - \underline{p}$. For each $p \in (\underline{p}, \hat{p}^e)$ let $\tilde{t}(p) = \hat{t}(p) - u_I(\hat{p}^e|(T_c, \tilde{t}))$, so that $u_I(\cdot|(\hat{\theta}(p), \tilde{t}(p))) = u_I(\cdot|\hat{\theta}(p), \hat{t}(p)) + u_I(\hat{p}^e|(T_c, \tilde{t}))$. Find $\tilde{p} \in (\underline{p}, \hat{p}^e)$ such that $u_I(\underline{p}|(\hat{\theta}(\tilde{p}), \tilde{t}(\tilde{p}))) = \mu - \underline{p} + \epsilon$ for small enough $\epsilon > 0$. If A offers $M' = \{(\hat{\theta}(p), \tilde{t}(p))|p \in [\tilde{p}, \hat{p}^e]\} \cup \{(T_c, \tilde{t}), (T_0, 0)\}$, it is optimal for S to price at $\bar{p}(T_c, \tilde{t})$ and B to accept $(T_c, \tilde{t})$ (since S's payoff is constant at $(p^e - c)q^e$ for prices in $[\tilde{p}, \hat{p}^e]$ and lower for prices in $[\underline{p}(M'), \tilde{p}) \bigcup (p^e, \bar{p}(T_c, \tilde{t}))$). Also, $\tilde{t}(\tilde{p}) > \hat{t}(\underline{p}) \geq 0$ for the same reason as in (i), and $\tilde{t}(p)$ increases in $p \in (\tilde{p}, \hat{p}^e)$ because total surplus from trading with $(\hat{\theta}(p), \tilde{t}(p))$ increases while B's payoff decreases and S's profit is constant.

Finally, since $\underline{p}(M') < \underline{p}$ in both (i) and (ii), A may modify $(T_c, \tilde{t})$ to $(T_c, \tilde{t} + \epsilon')$ for small enough $\epsilon' > 0$ and still induce S to price at $\bar{p}(T_c, \tilde{t} + \epsilon')$ and B to accept $(T_c, \tilde{t} + \epsilon')$, giving A a payoff of $\tilde{t} + \epsilon' > t(p^e)$. This proves Claim 3. $\square$

Therefore, if $\widehat{M}$ exists, it will follow that, for an optimal menu M, $U_I(.|M)$ must coincide with g on $[\underline{p}, \bar{p}(g) = p^e]$, so that S is indifferent between all prices in $[\underline{p}, p^e]$ and $U_I'(\underline{p}|M) = -1$. This will prove (b). Finally, then, we construct the menu $\widehat{M}$, given a PCE $[\theta(\cdot), t(\cdot), p^e]$ where $t(p) \geq 0$ for all $p \in [0, 1]$ and $q^e = 1 - F(\theta(p^e))$.

For $p \in [\underline{p}, \bar{p}(g)]$, let $\hat{q}(p) = -g'(p)$, define $\hat{\theta}(p)$ by $1 - F(\hat{\theta}(p)) = \hat{q}(p)$, and let

$$\hat{t}(p) = \int_{\hat{\theta}(p)}^{1} v dF(v) - p\hat{q}(p) - g(p) = \int_{\hat{\theta}(p)}^{1} (v - p) dF(v) - g(p). \tag{16}$$

The menu $\widehat{M}$ is then given by $\{(\hat{\theta}(p), \hat{t}(p))|p \in [\underline{p}, \bar{p}(g)]\} \cup (T_0, 0)$. Suppose $\hat{t}(p) \geq 0$ for all $p \in [\underline{p}, \bar{p}(g)]$. Then, for all such p, the graph of $u_I(\cdot|(\hat{\theta}(p), \hat{t}(p)))$ is linear wherever it lies above the graph of $u_0(\cdot)$. (The graph of $u_I(\cdot|(\hat{\theta}(p), \hat{t}(p)))$ is piecewise linear with three pieces; the value at p_1 on the left-hand piece is $\mu - p_1 - \hat{t}(p) \leq \mu - p_1 = \max\{\mu - p_1, 0\} = u_o(p_1)$ and the value at p_2 on the right-hand piece is $-\hat{t}(p) \leq 0 = u_o(p_2)$.) By construction, the graph of $u_I(\cdot|(\hat{\theta}(p), \hat{t}(p)))$ is tangent to the convex function $g(\cdot)$ at $p \in [\underline{p}, \bar{p}(g)]$. This implies that g is the upper envelope of the locally linear functions $u_I(\cdot|(\hat{\theta}(p),\hat{t}(p)))$ on $[\underline{p},\bar{p}(g)]$. Therefore it remains only to show that $\hat{t}(p)\geq 0$ for all $p\in[\underline{p},\bar{p}(g)]$. From (16) and $(p-c)\hat{q}(p)=(p^e-c)q^e$, we get

$$\hat{t}'(p) \;=\; -\hat{\theta}(p)F'(\hat{\theta}(p))\hat{\theta}'(p)-c\hat{q}'(p)+\hat{q}(p) \;=\; [\hat{\theta}(p)-c]\hat{q}'(p)+\hat{q}(p) \;=\; \hat{q}(p)\frac{p-\hat{\theta}(p)}{p-c}.$$

where the second equality follows because $\hat{q}(p)=1-F(\hat{\theta}(p))$ implies $-F'(\hat{\theta}(p))\hat{\theta}'(p)=\hat{q}'(p)$, and the third because $(p-c)\hat{q}(p)$ is constant so that $\hat{q}(p)+(p-c)\hat{q}'(p)=0$.

$\hat{q}(\underline{p})=-g'(\underline{p})=1$, so $\hat{\theta}(\underline{p})=0$ and, by (16), $\hat{t}(\underline{p})=0$. As p increases, the fee increases as long as $p\geq\hat{\theta}(p)$. Suppose that $\hat{t}(p)<0$ for some $p\in[\underline{p},\bar{p}(g)]$. Then there exists $\tilde{p}\in[\underline{p},\bar{p}(g)]$ such that $\hat{t}(\tilde{p})<0$ and $\tilde{p}<\hat{\theta}(\tilde{p})$.

Since $U_I'(p|M)\geq g'(p)$ for all $p\in[\underline{p},\bar{p}(g)]$, $q(\tilde{p})\leq\hat{q}(\tilde{p})$ and so $\theta(\tilde{p})\geq\hat{\theta}(\tilde{p})$. From $t(\tilde{p})=\int_{\theta(\tilde{p})}^1(v-\tilde{p})dF(v)-U_I(\tilde{p}|M)$ and (16), we get

$$\hat{t}(\tilde{p})-t(\tilde{p}) \;=\; \int_{\hat{\theta}(\tilde{p})}^{\theta(\tilde{p})}(v-\tilde{p})dF(v)+[U_I(\tilde{p}|M)-g(\tilde{p})] \;\geq\; 0$$

where the inequality follows because $U_I(\tilde{p}|M)\geq g(\tilde{p})$ and $\tilde{p}<\hat{\theta}(\tilde{p})$. Therefore, since $t(\tilde{p})\geq 0$ by Lemma 2, $\hat{t}(\tilde{p})\geq 0$. This shows that $\hat{t}(p)\geq 0$ for all $p\in[\underline{p}=(p^e-c)q^e+c,\bar{p}(g)]$. QED

Proof of Proposition 3 Let V^* denote the maximized value of (P). Lemma 3 shows that

[I] *for any PCE there is a threshold-contract PCE* $[\theta(\cdot),t(\cdot),p^e]$, *with* $t(\cdot)\geq 0$ *and satisfying the conditions in Proposition 2(a)–(c), which A weakly prefers.*

We now show that $(p^e,q^e=1-F(\theta(p^e)))$ satisfies the constraints of (P). To see this, recall that $U_I(\cdot|M)$ (M is the menu associated with PCE $[\theta(\cdot),t(\cdot),p^e]$) is (6), the solution to a differential equation implied by conditions of Proposition 2(a)–(c). Since $U_I((p^e-c)q^e+c|M)=\mu-(p^e-c)q^e-c$ given that $\underline{p}(M)=(p^e-c)q^e+c$ by Proposition 2(b), (p^e,q^e) satisfies (7). By $U_I(p|M)=\int_{\theta(p)}^1 vdF(v)-pq(p)-t(p)$ and (6), $t(p)$ equals the LHS of (9). Hence, (9) is satisfied because $t(p)\geq 0$ by [I].

Therefore, V^* is an upper bound of A's payoff in any PCE. Hence, if

[II] *a solution to (P) exists and constitutes a PCE* $[\theta(\cdot),t(\cdot),p^e]$,

then t^* exists and $V^*=t^*$, so a PCE $[\theta(\cdot),t(\cdot),p^e]$ is optimal iff (i)–(iii) apply.

We now prove [II]. We have shown that a solution to (P) exists, denoted by (p^e,q^e)

satisfying (7) and (9), where, for $p \in [\underline{p} \equiv (p^e - c)q^e + c, p^e]$, $q(p) = (p^e - c)q^e/(p - c)$ and $1 - F(\theta(p)) = q(p)$. For such p let $t(p)$ equal the LHS of (9) and let $g(p) = (p^e - c)q^e \ln((p^e - c)/(p - c))$. For p outside this interval, let $\theta(p) = 0 = t(p)$. Then $[\theta(\cdot), t(\cdot), p^e]$ is a PCE. To see this, note that $g(\cdot)$ is strictly convex and decreasing, $g(p^e) = 0$ and, by (7), $g(\underline{p}) = \mu - \underline{p}$; also, for $p \in [\underline{p}, p^e]$, $u_I(p|(\theta(p), t(p))) = g(p)$ and $u_I'(p|(\theta(p), t(p))) = -q(p) = g'(p)$. It follows that the graph of $u_I(\cdot|(\theta(p), t(p)))$ is a supporting tangent to g at p. Hence $(\theta(p), t(p))$ is an optimal choice for B from the menu $\{\theta(p), t(p)\}_{p \in [0,1]}$ when the price is p (see the argument following (16) above). Therefore, given this menu, it is an equilibrium for B to choose $(\theta(p), t(p))$ for $p \in [0, 1]$ and, since $(p - c)q(p)$ is then constant on $[\underline{p}, p^e]$ and lower outside this interval, for S to choose price p^e. This proves [II].

To show that A's payoff is t^* in every equilibrium of Γ_M, let $M(\epsilon)$ be the same as the optimal menu $\{(\theta(p), t(p))\}_{p \in [0,1]}$ obtained from the solution q^e to (P) as described above, except that $(\theta(p^e), t(p^e) - \epsilon)$ replaces $(\theta(p^e), t(p^e))$, where $\epsilon > 0$ is small. Given $M(\epsilon)$, if S sets p slightly above p^e B's unique optimal choice is $(\theta(p^e), t(p^e) - \epsilon)$, giving S a profit strictly higher than $(p^e - c)q^e$. Hence, given $M(\epsilon)$, there is a unique equilibrium continuation and it gives payoff $t(p^e) - \epsilon$ to A. Since ϵ is arbitrary, A's payoff is $t(p^e) = t^*$ in every equilibrium. QED

Proof of Proposition 4 Assume $F(\theta) = \theta$. First, we show that the solution to (8), ignoring (9), is unique. The objective function (8), given (7), reduces to

$$\max_{q^e \in [\underline{q}(0.5,c),1]} \; O(q^e) \;\equiv\; q^e\left(1 - c - \frac{q^e}{2}\right) - \frac{0.5 - c}{1 - \ln(q^e)}.$$

We show first that (i) for some $q' \in (0, 0.5]$, $O(q)$ is convex for $q < q'$ and concave for $q > q'$; and (ii) $O(q)$ has a unique maximizer q^* and $q^* \in (0.5, 1 - c)$. To see this, note that $O'(q) = 1 - c - q - \frac{0.5-c}{q(1-\ln(q))^2}$, $O''(q) = -1 - \frac{(0.5-c)(1+\ln(q))}{q^2(1-\ln(q))^3}$ and $O'''(q) = \frac{(2c-1)(1+\ln(q)+\ln(q)^2)}{q^3(1-\ln(q))^4}$. Since $1 + \ln(q) + \ln(q)^2 = (1 + \frac{\ln(q)}{2})^2 + \frac{3\ln(q)^2}{4} > 0$ and $c < \mu = 0.5$, $O''(q)$ decreases in q. $O''(q) > 0$ for q close to zero and, because $1 + \ln(0.5) > 0$, $O''(q) < 0$ if $q \geq 0.5$. Thus, $O(q)$ is convex for low q, then concave. $O'(0.5) = (0.5-c)(1-\frac{1}{0.5(1-\ln(0.5))^2}) > 0$. $1-c > 0.5$, so $O'(1-c) = -\frac{0.5-c}{(1-c)(1-\ln(1-c))^2} < 0$. $O(q) \leq 0$ as $q \to 0$ and $O(0.5) > 0$. Therefore $O(q)$ is maximized at some unique $q^* \in (0.5, 1 - c)$.

It is routinely verified that p^e in (7) decreases in q^e, with values $p^e(0.5) \approx 0.59 - 0.18c$ at $q^e = 0.5$ and 0.5 at $q^e = 1$. This means that $p^* = p^e(q^*) \in (0.5, p^e(0.5))$, thus $q^* \in [\underline{q}(0.5, c), 1]$, establishing that q^* is the unique solution to (8).

Next, we show $\theta(p) \leq p$ for all $p \in (c, p^*)$, which verifies that (9) is satisfied as explained in the main text, thus that q^* is the unique solution to (P). Recall

$$\theta(p) = 1 - \frac{(p^* - c)q^*}{p - c} = 1 - \frac{0.5 - c}{(1 - \ln(q^*))(p - c)}$$

Clearly, $\theta(p^*) = 1 - q^* < p^*$ because $p^*, q^* > 0.5$. Note that $\theta(p)$ increases because

$$\theta'(p) = \frac{0.5 - c}{(1 - \ln(q^*))(p - c)^2} \;>\; \theta_l'(c) \equiv \frac{0.5 - c}{(1 - \ln(0.5))(p^e(0.5) - c)^2} \approx \frac{0.423}{0.5 - c}$$

for $p \in (c, p^*)$, where the inequality is due to $0.5 < q^*$ and $p < p^* < p^e(0.5)$; and $\theta(p)$ is concave because $\theta''(p) < 0$. Therefore, $p - \theta(p)$ decreases in $p < p^*$ if $\theta_l'(c) > 1$, hence $\theta(p) \leq p$ holds for all $p \in (c, p^*)$. This is the case if $c \geq 0.1$ because $\theta_l'(0.1) > 1$.

If $c < 0.1$, note from (7) that

$$(p^e - c)q^e = \frac{0.5 - c}{(1 - \ln(q^*))} \;>\; \frac{0.5 - c}{(1 - \ln(0.5))} \;>\; \frac{(1 - c)^2}{4} = \max_{p \in (0,1)} (p - c)(1 - p).$$

Since $(p^e - c)q^e = (p - c)(1 - \theta(p))$, this implies that $(p - c)(1 - \theta(p)) > (p - c)(1 - p)$, hence $\theta(p) < p$ as desired, establishing that q^* is the unique solution to (P).

Lastly, since $q^* \in (0.5, 1 - c)$, $\theta(p^*) = 1 - q^* \in (c, 0.5)$. Moreover, q^* decreases in c because $O'(q)$ decreases in c, which in turn implies that $\theta^* = \theta(p^*)$ increases in c. This completes the proof. QED

Proof of Proposition 5 (a) To prove inefficiency, note first that $\theta(p^e) \geq c$ by Proposition 2(c). The derivative of total surplus with respect to $\theta(p^e)$ is $-(\theta(p^e) - c)F'(\theta(p^e))$, which equals zero if $\theta(p^e) = c$, whereas S, by (7), gets payoff $(\mu - c)/(1 - \ln(q^e))$, which increases in q^e, hence reduces in $\theta(p^e)$. Therefore A's payoff increases locally as the threshold rises from c, so $\theta(p^e) > c$, which is inefficient.

(i): S is strictly worse off since her payoff is $(p^e - c)q^e < \mu - c$ by (7).

(ii): The statement for low c follows because the equilibrium of Γ_M is inefficient. The outcome without A is efficient if $c = 0$, hence, by continuity, social surplus is lower for low $c > 0$. Consider the limit outcome as $c \to \mu$. The limit equilibrium outcome of Γ_1 is efficient by Proposition 1 (because $c < \widehat{\theta} < \mu$) and A's fee converges to the full social surplus because S's expected payoff, $(\widehat{\theta}-c)(1+F(\widehat{\theta}))$, converges to zero. Therefore, in Γ_M too A's equilibrium fee must converge to the social surplus; hence, in the limit, the equilibrium outcome is efficient. The limit outcome in the absence of A is, however, bounded away from efficiency: the limit amount of inefficiency is $\lim_{c\uparrow\mu}\int_0^c(c-v)dF=\int_0^\mu(\mu-v)dF>0$. Hence, for c near μ, surplus is higher in Γ_M.

(b) Consider the relaxed version of (P) ignoring (9), referred to as (P_r). The value of the objective function (8) is 0 at the upper boundary $q^e=1$; and is negative at the lower boundary $q^e=\underline{q}(c,\mu)$ because $\int_{\theta(p^e)}^1 vdF(v)$ is the unconditional expected value of $v>\theta(p^e)$, which is lower than p^eq^e when $p^e=1$ since $q^e=1-F(\theta(p^e))$.

The first order condition for an interior solution to (P_r) is

$$F^{-1}(1-q^e)-\left(\frac{\mu-c}{q^e(1-\ln(q^e))^2}+c\right) \;=\; 0. \tag{17}$$

Since the LHS is $-\mu<0$ at the upper boundary $q^e=1$ (i.e., the objective function decreases at $q^e=1$) and the objective function value is negative at the lower boundary $q^e=\underline{q}(c,\mu)$ as shown above, (P_r) has interior solution, q^*, that satisfies (17). Since $q^*=1-F(\theta(p^*))$, (17) gives

$$\theta(p^*) \;=\; \frac{\mu-c}{(1-F(\theta(p^*)))(1-\ln(1-F(\theta(p^*))))^2}+c.$$

To show that total welfare is lower in Γ_M than in Γ_1, we show that $\theta(p^*)>\widehat{\theta}$. Recall

$$\widehat{\theta}=\frac{\mu-c}{(1+F(\widehat{\theta}))^2}+c$$

from (2), where the RHS is monotone decreasing in θ. For $F(\theta)\in(0,1)$

$$\frac{\mu-c}{(1-F(\theta))(1-\ln(1-F(\theta)))^2}>\frac{\mu-c}{(1+F(\theta))^2}$$

because $(1+F(\theta))^2-(1-F(\theta))(1-\ln(1-F(\theta)))^2=0$ at $F(\theta)=0$ and its derivative with respect to $F(\theta)$ is $1+2F(\theta)+\ln(1-F(\theta))^2>0$. Since $F(\theta)$ is increasing in θ, it follows that $\widehat{\theta}<\theta(p^*)$. QED

Proof of Lemma 4 (i) By Proposition 3, (p^*,q^*) satisfies (7) and (9), which is equivalent to $t(p)\geq 0$ for all $p\in[p^*q^*,p^*]$. $(G_{p^*q^*}^{p^*},p)$ is an RS-outcome for all $p\in[p^*q^*,p^*]$ if F is a mean-preserving spread of $G_{p^*q^*}^{p^*}$.$G_{p^*q^*}^{p^*}$ has an atom of q^* at p^*, so its mean is

$$\int_{p^*q^*}^{p^*} \frac{p^*q^*}{v} dv + p^*q^* = p^*q^* \ln\Big(\frac{1}{q^*}\Big) + p^*q^*$$

which equals μ by (7). Therefore F is a mean-preserving spread of $G^{p^*}_{p^*q^*}$ if

$$\int_{p^*q^*}^{p} \Big(1 - \frac{p^*q^*}{v}\Big) dv \leq \int_0^p F(v)dv \quad \text{for all} \quad p \in [p^*q^*, p^*], \quad \text{i.e., if}$$

$$\phi(p) \equiv \int_0^p F(v)dv - p + p^*q^* - p^*q^* \ln\Big(\frac{p^*q^*}{p}\Big) \geq 0 \quad \text{for all} \quad p \in [p^*q^*, p^*]. \tag{18}$$

By (9),

$$t(p) = \int_{\theta(p)}^{1} v dF(v) - p^*q^* - p^*q^* \ln\Big(\frac{p^*}{p}\Big) \geq 0 \quad \text{for all} \quad p \in [p^*q^*, p^*],$$

which, after integrating by parts and using (7) and $F(\theta(p)) = 1 - (p^*q^*/p)$, gives

$$t(p) = \int_0^{\theta(p)} F(v)dv - \theta(p) + \frac{\theta(p)p^*q^*}{p} - p^*q^* \ln\Big(\frac{p^*q^*}{p}\Big) \geq 0 \ \text{ for all } \ p \in [p^*q^*, p^*]. \tag{19}$$

$\phi'(p) = F(p) - F(\theta(p))$ and, by (10), $t'(p) = q(p)(p - \theta(p))/p$, so that $t'(\cdot)$ and $\phi'(\cdot)$ always have the same sign; moreover, at any turning-point, i.e., for any p such that $p = \theta(p)$, t and ϕ have the same value. (18) then follows from (19) since $\phi(p^*q^*) \geq 0$ and $t(\cdot)$ is non-decreasing at p^*, otherwise A would get a higher fee from a slightly lower seller price, contradicting optimality of the menu.

(ii) As p falls from p^* to p^*q^* in $(G^{p^*}_{p^*q^*}, p)$, B's surplus increases continuously from zero to $\mu - p^*q^*$. Since $\mu - p^*q^* \geq t(p^*)$ because μ is the maximal social surplus and the total equilibrium surplus of Γ_M is $t(p^*) + p^*q^*$, there exists $\tilde{p} \in [p^*q^*, p^*]$ such that B's surplus from the RS-outcome $(G^{p^*}_{p^*q^*}, \tilde{p})$ is $t(p^*)$. QED

References

Admati, A. and P. Pfleiderer (1986), "A Monopolistic Market for Information", *J. Econ. Theory*, 39, 400-38.

Albano, G. and A. Lizzeri (2001), "Strategic Certification and Provision of Quality", *Int. Econ. Rev.*, 42(1), 267-83.

Ali, S.N. N. Haghpanah, X. Lin and R. Siegel (2022), "How to Sell Hard Information", *Quart. J. Econ.*, 137(1), 619-78.

Asseyer, A. and R. Weksler (2024), "Certification Design with Common Values", *Econometrica*, 92(3), 651-86.

Bergemann, D. and A. Bonatti (2019), “Markets for Information: An Introduction”, *Annual Revi. Econ.*, 11, 85-107.
Bergemann, D., A. Bonatti and A. Smolin (2018), “The Design and Price of Information”, *Amer. Econ. Rev.*, 108(1), 1-48.
Bergemann, D. and S. Morris (2019), “Information Design: A Unified Perspective”, *J. Econ. Perspectives*, 57(1), 44-95.
Bergemann, D. and M. Pesendorfer (2007), “Information Structures in Optimal Auctions”, *J. Econ. Theory*, 137(1), 580-609.
Calzolari, G. and A. Pavan (2006), “On the Optimality of Privacy in Sequential contracting”, *J. Econ. Theory*, 130, 168-204.
Celik, G. and R. Strausz (2025), “Informative Certification: Screening vs. Acquisition”, Discussion paper, Collaborative Research Center Transregio 190.
Eso, P. and B. Szentes (2007), “Optimal Information Disclosure in Auctions and the Handicap Auction”, *Rev. Econ. Stud.*, 74(3), 705-31.
Evans, R. and I.-U. Park (2025), “Selling Information for Bilateral Trade”, SSRN: https://ssrn.com/abstract=5362376
Gershkov, A., B. Moldovanu, P. Strack and M. Zhang (2021), “A Theory of Auctions with Endogenous Valuations”, *J. Pol. Econ.*, 129(4), 1011-51.
Guo, Y., H. Li and X. Shi (2024), “Optimal Discriminatory Disclosure”, unpublished.
Hörner, J. and A. Skrzypacz (2016), “Selling Information”, *J. Pol. Econ.*, 124(6), 1515-62.
Inderst, R. and M. Ottaviani (2012), “Competition through Commissions and Kickbacks”, *Amer. Econ. Rev.*, 102(2), 780-809.
Kamenica, E. (2019), “Bayesian Persuasion and Information Design”, *Annual Rev. Econ.*, 11, 249-72.
Kamenica, E. and M. Gentzkow (2011), “Bayesian Persuasion”, *Amer. Econ. Rev.*, 101(6), 2590-615.
Lee, C. (2021), “Optimal Recommender System Design”, working paper, U. Penn.
Lewis, T. and D. Sappington (1994), “Supplying Information to Facilitate Price Discrimination”, *Int. Econ. Rev.*, 35(2), 309-27.
Li, H. and X. Shi (2017), “Discriminatory Information Disclosure”, *Amer. Econ. Rev.*, 107(11), 3363-85.
Lizzeri, A. (1999), “Information Revelation and Certification Intermediaries”, *Rand J. Econ.*, 30(2), 214-31.
Luca, M. T. Wu, S. Couvidat and D. Frank (2015), “Does Google Content Degrade Google Search? Experimental Evidence”, Harvard Bus. Sch. WP 16-035.
Mensch, J. (2022), “Screening Inattentive Buyers”, *Amer. Econ. Rev.*, 112, 1949-84.
Myerson, R. (1986), “Multistage Games with Communication,” *Econometrica*, 54, 323-58.
Ravid, D. A.-K. Roesler and B. Szentes (2022), “Learning Before Trading: On the Inefficiency of Ignoring Free Information”, *J. Pol. Econ.*, 130(2), 346-87.
Roesler, A-K. and B. Szentes (2017), “Buyer-Optimal Learning and Monopoly Pricing”, *Amer. Econ. Rev.*, 107(7), 2072-80.
Strassen, V. (1965), “The Existence of Probability Measures with Given Marginals”, *Annals Math. Stat.*, 36(2), 423-39.

Watson, J. (2024), “Perfect Bayesian Equilibrium: Consistency Conditions for Practical Definitions”, working paper, UCSD.

Wei, D. and B. Green (2024), “(Reverse) Price Discrimination with Information Design”, *Ameri. Econ. J.: Microeconomics*, 16(2), 267-95.

Yang, K, H. (2019), “Equivalence in Business Models for Informational Intermediaries”, working paper, U. Chicago.

Yang, K. H. (2022), “Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences”, *Amer. Econ. Rev.*, 112(4), 1364-93.