

ADVISORS WITH HIDDEN MOTIVES*

Paula Onuchic
University of Oxford

July, 2022

Abstract

I study a model of advisors with hidden motives: a seller discloses information about an object's value to a potential buyer, who doesn't know the object's value or how profitable the object's sale is to the seller (the seller's motives). I characterize optimal disclosure rules, used by the seller to *steer* sales from lower- to higher-profitability objects. I investigate the effects of mandated transparency of the object's profitability, and ask whether an advisee is always better off when they understand their advisor's motives. The answer to this question may be negative: I show that, by removing the seller's *steering incentive*, transparency can dissuade the seller from disclosing information about the object's value, and also from acquiring that information in the first place. Finally, I consider variations of the communication protocol and show that the effect of transparency on the advisor's informativeness is linked to their ability to commit to disclosure policies.

1 Introduction

People frequently take advice from advisors with hidden motives: broker-dealers and other financial advisors often counsel investors, but also receive sales commissions from financial product providers; doctors inform patients of the effectiveness of different drugs and procedures, but are often sponsored by pharmaceutical companies; online platforms gather reviews from product users, but may disclose them coarsely to benefit some producers; magazines and newspapers selectively publish pieces of reporting that align with their editorial bias. In all the mentioned settings,

*I am grateful for guidance from Debraj Ray, Erik Madsen and Ariel Rubinstein, and for long discussions with Joshua Weiss and Samuel Kapon. I also thank Margaret Meyer, David Pearce, Ludvig Sinander, Dilip Abreu, Ricardo Alonso, Arjada Bardhi, Heski Bar-Isaac, Dhruva Bhaskar, Sylvain Chassang, Navin Kartik, Laurent Mathevet, Luis Rayo, Mauricio Ribeiro, Deszo Szalay, Nikhil Velodi and Basil Williams for their helpful comments.

information receivers understand that information providers may be biased, but are not fully aware of the extent of the conflict. For example, in the context of financial brokers, clients understand that brokers receive sales commissions from some product providers, but may not know the size of the commissions on each product.

Such conflicts of interest between advisors and advisees receive much regulatory attention. In all these contexts, the most common recommendation, and often instituted regulation, is *transparency*: an advisor should let an advisee know their real interests behind their recommendations.¹ The intuition is that, if a person knows the intention behind the information they receive, then they can take the advice with “a grain of salt.” The advisor, anticipating the advisees “healthy skepticism,” should then have the right incentives to provide informative advice.

In this paper, I scrutinize that intuition, and investigate whether an advisee is always better off when their advisor’s motives are transparent. To do so, I propose a model where a seller discloses information about an object to a buyer, who can choose to acquire the object. The seller is an advisor with *hidden motives*, because the buyer does not know how interested they are in the object’s sale. In this context, I show that a policy that informs the buyer about the seller’s interests does not necessarily improve the informativeness of the seller’s advice to the buyer. This can happen through two channels: transparency may decrease the seller’s willingness to disclose information about the object’s value; and it may also hinder the seller’s incentives to acquire information about the object’s value in the first place.

I introduce the model in more detail in section 2. The seller wishes to sell an object to a buyer who does not know the object’s *value* or how *profitable* its sale is to the seller. The seller knows the object’s profitability and sees some signal about its value. Prior to observing the object’s profitability or the realization of the signal about its value, the seller commits to a rule to disclose signal realizations to the buyer. This rule assigns to each realization of the signal and each profitability a probability that the realization is disclosed to the buyer. The buyer is Bayesian and updates their belief about the object’s value based on any information the seller reveals and on the seller’s disclosure policy itself. The probability that they purchase the object is given by the buyer’s “demand function,” which is increasing in the posterior expected value of the object.

When choosing what information to share with the buyer, the seller wants to disclose evidence that the object is valuable, so as to incentivize the buyer to purchase it. They also know that they can hide bad signal realizations, but if they do

¹Think of the SEC mandating that financial advisors disclose commissions they receive from financial product providers, or Instagram asking digital influencers to mark their posts as “paid content” when they are sponsored by brands. In section 6, I mention some empirical literature that considers the effects of transparency in the context of financial advice. Ershov and Mitchell (2022) measure the effects of disclosure regulations on Instagram.

so, the buyer becomes skeptical and reads non-disclosure itself as a negative sign, which lowers the probability of sale. In Theorem 1, I show that optimal disclosure rules follow a simple *threshold structure*. Each signal realization about the object's value is classified as either "good news," if above some endogenously determined threshold, or "bad news." For each good news realization, there is a profitability threshold such that the realization is disclosed if the object's profitability is above the threshold, and concealed otherwise. Conversely, each piece of bad news is concealed if the object is sufficiently profitable, and revealed otherwise. Using such a disclosure policy, the seller effectively *steers* purchases from low- to high-profitability objects.

Theorem 1 also characterizes the seller's optimal profitability thresholds, and Corollaries 1 and 2 show that their shape depends on the curvature of the buyer's demand function. Next, I consider introducing a transparency policy, which institutes that the buyer observes the object's profitability to the seller. I show that, under transparency, the seller cannot use strategic disclosure in order to steer sales across objects with different profitability levels; and so the policy alters their incentives to disclose information about the object's value to the buyer. Interestingly, it is not necessarily true that this policy *increases* the seller's incentives to inform the buyer. Indeed, Proposition 1, and its Corollary 3, show that if the buyer's demand function is strictly concave, then under mandated transparency the seller is Blackwell less informative than when their motives are hidden. (The opposite conclusion holds when the buyer's demand function is strictly convex.)

The intuition behind this result is that, when the seller's motives are hidden, some selective information disclosure allows them to steer sales from low- to high-profitability objects. Importantly, in order to effectively steer sales, the seller *must disclose some information*. Under the transparency policy, precisely because the buyer knows the seller's motives, the "steering incentive" for information disclosure vanishes, sometimes leading the seller to disclose strictly less information.

Having considered the effects of transparency on the seller's incentives to disclose information, in section 4 I augment the model with an ex-ante stage in which the seller invests in acquiring information about the object's value in. I find that the "steering incentive" is also an important driver of the seller's decision to invest in information acquisition. Proposition 3 shows that increasing transparency of the seller's motives unequivocally decreases the precision of the signal they acquire about the object's value. In Proposition 4 and Corollary 4, I show that, through this information acquisition channel, increasing transparency may indeed reduce the seller's informativeness.

In section 5, I consider the robustness of the results mentioned above with respect to changes in the communication protocol. The benchmark model makes two main assumptions regarding the communication protocol. First, the seller can com-

mit to a rule to disclose signal realizations, prior to the object’s profitability or the “evidence” about its value being drawn. Second, the seller can use only disclosure policies, rather than committing to more general messaging strategies. I consider two alternative communication protocols, where each of these assumptions is dispensed with, and show that, also in these scenarios, transparency sometimes decreases the seller’s informativeness.

Considering the disclosure protocol with no commitment, I find that due to the sender’s hidden motives, partially uninformative equilibria may not unravel, as in traditional voluntary disclosure games. The unravelling logic fails because, upon observing non-disclosure, the buyer may not know whether the seller has “good news” but negative profitability or “bad news” but positive profitability. The existence of these partially informative equilibria implies that, in some cases, the optimal disclosure rule (with commitment) is also an equilibrium of the disclosure game without commitment.

Interestingly, in the no-commitment disclosure game, imposing transparency of the seller’s motives unambiguously improves their informativeness. Indeed the common intuition applies: if the buyer knows the intention behind the information they receive, then they take the seller’s advice with “the right degree of skepticism,” and the seller has the right incentives to provide informative advice. However, under this protocol, I note that transparency may still depress a seller’s incentives to acquire precise information about the object; and thus transparency may still negatively affect the seller’s informativeness.

Finally in section 6, I discuss the main assumptions of the model in light of an application where the seller is a financial advisor and the buyer is an investor, and relate my contribution to previous work in that applied context.

1.1 Related Literature

This paper contributes to the literature studying the disclosure of hard evidence, started with Grossman (1981) and Milgrom (1981).² In particular, in my model the sender uses a disclosing technology as in Dye (1985) and Jung and Kwon (1988), whereby evidence realization is either entirely disclosed to the receiver or concealed. My model mainly departs from that literature by (i) considering a disclosure problem where the sender has hidden motives, and (ii) considering a protocol where the sender commits to a disclosure rule prior to the realization of the evidence. To my knowledge, this is the first paper to consider disclosure by a sender with hidden motives, and to assess the effects of mandated transparency in that context. Regarding the commitment assumption, sections 5 and 6 discuss at length

²For a survey, see Milgrom (2008)

the implications of commitment on the communication outcome.

In assuming that the sender can commit ex-ante to a signaling technology, my paper relates to the large literature on information design, mainly stemming from Kamenica and Gentzkow (2011) and Rayo and Segal (2010).³ Indeed, my model is a version of Rayo and Segal’s (2010) model where communication happens through disclosure of hard evidence. Rayo and Segal (2010) do not describe their model as representing advisors *with hidden motives*, and do not consider the effect of mandated transparency on the informativeness of the sender, which is the focus of my analysis.⁴ Moreover, I contribute to the understanding of the benchmark model in Rayo and Segal (2010), by characterizing optimal disclosure for general receiver “demand functions,” while most of Rayo and Segal’s (2010) characterization results apply to the linear specification.⁵ Finally, my model also departs from this benchmark by considering costly information acquisition.

Within information design, my paper contributes to a growing literature studying *constrained information design*. These are problems where the sender is subject to additional constraints, beyond Bayesian plausibility, when choosing a signal structure. For example, Mensch (2021) and Onuchic & Ray (2022) consider monotonicity constraints.⁶ This paper considers a sender who is constrained to disclosure strategies. The characterization of optimal signals is a largely open problem in the information design literature and, because this characterization is elusive, we also know very little about comparative statics in information design models.⁷ In this paper, I provide an example of a comparative static exercise which is simplified by the assumption that the sender is constrained to a reasonable class of signals.

Kartik, Lee and Suen (2017) and Che and Kartik (2009) study different environments with endogenous information acquisition and voluntary disclosure, and also argue that disagreement between sender and receiver can incentivize information acquisition.⁸ Kartik, Lee and Suen (2017) show that an advisee may prefer to

³For a survey, see Kamenica (2019)

⁴The term “transparent motives” was coined by Lipnowski and Ravid (2020) to refer to a problem where the sender’s preferences are state-independent. In my model, under mandated transparency, the seller acts *as if* they have transparent motives, in that sense.

⁵With nonlinear demand functions, the optimal disclosure rule in my model sometimes “pools ordered prospects” – in Rayo and Segal’s (2010) language, ordered prospects are two objects whose values and profitabilities are ordered in the same way. For an example, see the optimal disclosure rule depicted in the right panel of Figure 3. One of Rayo and Segal’s (2010) main characterization results is that, in the linear benchmark, optimal signals never pool ordered prospects.

⁶See Doval and Skreta (2021) and references therein for other examples.

⁷In a recent paper, Curello and Sinander (2022) provide some general comparative statics results, by circumventing the characterization problem.

⁸Shishkin (2022) and DeMarzo, Kremer, and Skrzypacz (2019) also study information acquisition by the sender in a Dye (1985) framework. Szalay (2005) and Ball and Gao (2021) study information acquisition by a biased agent in a delegation context.

solicit advice from just one biased expert even when others (of equal or opposite bias) are available, because in the presence of more advisors, each individual expert free rides on the information acquired by the other experts. In Che and Kartik (2009), though sender and receiver share the same preferences, they hold different priors over the distribution of states. Che and Kartik (2009) show that the sender may invest more in acquiring information when the disagreement between their priors is larger.

2 Environment

A seller wishes to sell an object to a buyer who does not know their value for the object, or how profitable the object’s sale is to the seller. The object’s *value* to the buyer, denoted $x \in \mathcal{X} = [x_{min}, x_{max}]$, and its *profitability* to the seller, denoted $y \in \mathcal{Y} = [y_{min}, y_{max}]$, are drawn from some joint distribution commonly known by seller and buyer. We assume the joint distribution and its marginals have no mass points. Denote by F_Y the cdf of the marginal profitability distribution, and assume it has a strictly positive mean.

Before the buyer decides whether to purchase the object, the seller can convey to them some information about the object’s value (see below for details on the communication protocol). After observing any conveyed information, the buyer forms a Bayesian posterior belief about the object’s value, with some expected value \hat{x} . Given the object’s expected value $\hat{x} \in \mathcal{X}$, it is purchased by the buyer with probability $p(\hat{x})$, where $p : \mathcal{X} \rightarrow [0, 1]$ is some strictly increasing and differentiable “demand function.”⁹ If the object is bought, the seller receives a payoff equal to the object’s profitability, y . Otherwise, the seller’s payoff is 0.

Information Disclosure Protocol. The seller sees the object’s profitability, y , as well as a signal of the value of the object to the buyer. Before observing the productivity and the signal realization, the seller *commits* to a rule to *disclose* signal realizations to the buyer.

⁹The “demand function” p is a reduced-form representation of the buyer’s behavior, as a function of their belief about the object’s value. This approach assumes that (i) the buyer’s purchase decision depends only on their belief about the object’s expected value, and (ii) the probability of sale is increasing in the buyer’s belief.

The demand p can be microfounded by many different models. For example, suppose the buyer has some outside option of privately known value and chooses between purchasing the object or keeping their outside option. Further, suppose that, from the seller’s perspective, the outside option is distributed according to some continuous cdf p . The buyer purchases the object whenever their posterior mean \hat{x} is larger than the outside option. From the seller’s perspective, that happens with probability $p(\hat{x})$, which is thus a strictly increasing “demand function.”

The signal, $\pi : \mathcal{X} \times \mathcal{Y} \rightarrow \Delta\mathcal{M}$, is a map between the object’s value and profitability and a distribution of messages. Each signal realization $m \in \mathcal{M}$, if disclosed, induces on the buyer a Bayesian posterior belief about the object’s value, whose expectation is the buyer’s *posterior mean* value $\hat{x}(m) = \mathbb{E}(x|m, \pi)$. We assume that the signal structure π is such that, given a signal realization m , the buyer’s posterior mean is *independent* of the object’s profitability: for all $y \in \mathcal{Y}$, $\mathbb{E}(x|m, \pi, y) = \mathbb{E}(x|m, \pi) = \hat{x}(m)$.

Given this independence, we refer to a signal realization and to the posterior mean it induces interchangeably. In other words, if the buyer observes a message m , they interpret it as “the object’s expected value is $\hat{x}(m)$,” independently of their belief about the object’s profitability. Further, under this independence assumption, we can equate the signal π with its induced distribution of posterior means. This distribution, call its cdf $F_{X|y}$, may still vary with the object’s profitability y (despite the independence assumption), because the underlying value of the object and its profitability may be correlated.

Each of the realizations can be disclosed by the seller to the buyer, and this choice may depend on the profitability of the object. A *disclosure rule* is a measurable map from profitabilities and signal realizations into disclosure probabilities: $d : \mathcal{Y} \times \mathcal{X} \rightarrow [0, 1]$. When a signal realization \hat{x} is disclosed, the receiver’s posterior mean after observing it is exactly \hat{x} . If otherwise a signal realization is not disclosed, the receiver’s posterior mean is computed using Bayes’ Rule, accounting for the disclosure rule. Formally, if $\int_{\mathcal{Y}} \int_{\mathcal{X}} [1 - d(y, x)] dF_{X|y}(x) dF_Y(y) \neq 0$,¹⁰

$$x^{ND} = \frac{\int_{\mathcal{Y}} \int_{\mathcal{X}} x (1 - d(y, x)) dF_{X|y}(x) dF_Y(y)}{\int_{\mathcal{Y}} \int_{\mathcal{X}} (1 - d(y, x)) dF_{X|y}(x) dF_Y(y)}, \quad (1)$$

which is the expected value conditional on non-disclosure. Define y^{ND} analogously to be the average object profitability given that a realization is not disclosed.

Hidden Motives and Informativeness. In the environment described above, the buyer’s decision to purchase the object depends only on the object’s *value* – the buyer does not inherently care about the *profitability* of the sale to the seller. However, the buyer understands that the the object’s sale may be more or less profitable, and that the seller may be more or less motivated to share information about the object’s value depending on the object’s profitability.

Because the buyer does not observe the seller’s profitability, we say that the seller is an advisor *with hidden motives*. This paper tackles the question of whether the buyer would be better off if they were able to observe the advisor’s motives. To that end, we at times consider an alternative benchmark *with transparency*, where

¹⁰Otherwise, non-disclosure is “off-path,” and we fix x^{ND} at some value in \mathcal{X} .

the buyer observes the signal realizations disclosed by the seller, as well as the object's profitability. We compare the advisor with hidden motives to the transparent benchmark with respect to how informative they are about the object's value.

Given a disclosure rule d , let $F^B(\cdot, d)$ be the distribution of posterior means observed by the buyer.¹¹ We say a disclosure rule d is *more informative than* disclosure rule d' if $F^B(\cdot, d)$ is more informative than $F^B(\cdot, d')$ in the Blackwell order. Equivalently, if $F^B(\cdot, d)$ is a mean preserving spread of $F^B(\cdot, d')$.

3 Optimal Disclosure

Suppose the seller commits to a disclosure rule d . The probability that the object is sold, conditional on its profitability being y is

$$P(y, d) = \int_{\mathcal{X}} [d(y, x)p(x) + (1 - d(y, x))p(x^{ND})] dF_{X|y}(x). \quad (2)$$

To understand (2), first remember that F_y is the distribution of signal realizations given that the object has profitability y . Suppose a signal realization x is disclosed, which happens with probability $d(y, x)$. Then the object is sold with probability $p(x)$, which is reflected in the first term inside the integral of (2). As for the second term, with probability $1 - d(y, x)$ the realization x is not disclosed. In that case, the object is sold with probability $p(x^{ND})$, where x^{ND} is as given in (1).

The seller's expected payoff from committing to disclosure rule d is then

$$\Pi(d) = \mathbb{E}[yP(y, d)] = \mathbb{E}(y)\mathbb{E}[P(y, d)] + \text{Cov}[y, P(y, d)]. \quad (3)$$

In (3), I split the seller's payoff into two terms, expressing that the seller's objective can be seen as twofold. According to the first term, the seller wishes to maximize the overall expected probability of sale, which is multiplied by the average profitability. Per the second term, they seek to maximize the covariance between the object's profitability and its probability of sale.

¹¹The distribution of posterior mean values observed by the buyer is the only object relevant for the buyer's surplus. For a given d , $F^B(\cdot, d)$ can be calculated as follows: If $x < x^{ND}$,

$$F^B(x, d) = \int_{\mathcal{Y}} \int_{[x_{min}, x]} d(y, \hat{x}) dF_{X|y}(\hat{x}) dF_Y(y).$$

If $x \geq x^{ND}$,

$$F^B(x, d) = \int_{\mathcal{Y}} \int_{[x_{min}, x]} d(y, \hat{x}) dF_{X|y}(\hat{x}) dF_Y(y) + \int_{\mathcal{Y}} \int_{\mathcal{X}} (1 - d(y, \hat{x})) dF_{X|y}(\hat{x}) dF_Y(y).$$

Theorem 1 provides a *threshold characterization* of disclosure rules that maximize the seller's value. There is a threshold value \bar{x} such that each signal realization is classified as either *good news*, if x is larger than \bar{x} , or *bad news*, if $x < \bar{x}$. Additionally, for good news realization, there is a profitability threshold such that the realization is disclosed if and only if the object's profitability is above that threshold. Conversely, each bad news realization is disclosed if and only if the object's profitability is below some threshold.

Theorem 1. *An optimal disclosure rule exists and every optimal rule d^* has a threshold structure: There is a threshold value $\bar{x} \in \mathcal{X}$ and a threshold profitability $\bar{y} : \mathcal{X} \rightarrow \mathcal{Y}$ such that d^* almost everywhere satisfies $d^*(y, x) \in \{0, 1\}$ and*

$$d^*(y, x) = 1 \Leftrightarrow (x - \bar{x})(y - \bar{y}(x)) \geq 0.$$

The threshold value satisfies $\bar{x} = x^{ND}$, and, for $x \neq x^{ND}$, $\bar{y}(x)$ satisfies

$$\bar{y}(x) = y^{ND} \left[\frac{p'(x^{ND})(x^{ND} - x)}{p(x^{ND}) - p(x)} \right]. \quad (4)$$

I will now argue that any disclosure rule that does not satisfy the described threshold structure described can be improved upon, but the complete proof of Theorem 1 is in the Appendix. Start with a disclosure rule d which does not satisfy the threshold structure, and let x^{ND} be its implied non-disclosure posterior mean. Now define an alternative rule, \hat{d} , that discloses each realization x with the same probability as d , but has a threshold structure. That is, if $x \leq x^{ND}$, let

$$\hat{d}(y, x) = \begin{cases} 1, & \text{if } y \leq \hat{y}(x) \\ 0, & \text{if } y > \hat{y}(x) \end{cases} \quad (5)$$

and if $x > x^{ND}$, let

$$\hat{d}(y, x) = \begin{cases} 0, & \text{if } y < \hat{y}(x) \\ 1, & \text{if } y \geq \hat{y}(x) \end{cases} \quad (6)$$

where the thresholds \hat{y} are calibrated such that, for each realization x ,

$$\int_{\mathcal{Y}} \hat{d}(y, x) dF_{Y|x}(y) = \int_{\mathcal{Y}} d(y, x) dF_{Y|x}(y), \quad (7)$$

where $F_{Y|x}$ is the profitability distribution conditional on a signal realization x . Condition (7) implies that d and \hat{d} induce the same x^{ND} , and so \hat{d} satisfies all the requirements of the threshold structure in Theorem 1. Figure 1 shows a disclosure rule d that does not have a threshold structure and an improvement \hat{d} .

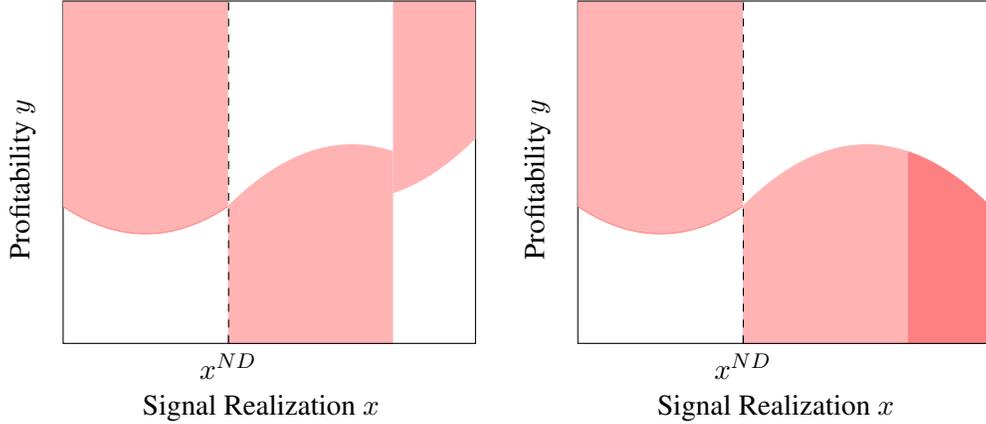


Figure 1: Each panel depicts a disclosure rule, with the colored areas representing zones of no disclosure and the white areas representing zones of disclosure. The disclosure rule depicted on the right panel has a threshold structure, and is an improvement over the disclosure rule in the left panel, which does not have a threshold structure.

By moving from d to \hat{d} , the seller shifts the disclosure probability of bad news to low profitability objects and of good news to high profitability objects, while maintaining the distribution of posterior mean values that is induced on the buyer. It is easy to see (and I argue formally in Appendix A), that: 1. d and \hat{d} produce the same overall probability of sale, because the distribution of posterior mean values is unchanged; and 2. \hat{d} induces a strictly larger covariance between sales and profitability than d , because the change increases the probability that very profitable objects are sold, and decreases that probability for less profitable objects. These facts imply that \hat{d} yields a strictly larger expected seller payoff than d , proving that the optimal rule must have a threshold structure.

Computationally, Theorem 1 achieves a significant simplification of the seller’s problem. All disclosure rules with the described threshold structure are almost everywhere defined by two numbers: x^{ND} and y^{ND} . As such, the seller’s problem can be approached as a two-dimensional maximization problem. We see from equation (4) that the distribution of profitability and signal realizations – and the correlation between the object’s value and its profitability – affects the shape of the optimal thresholds only through x^{ND} and y^{ND} . Further, some features of the optimal profitability threshold are fully defined by p , the buyer’s “demand function” – these features are further explored in Propositions (1) and (2) below.

Linear Demand. Corollary 1 applies Theorem 1 when p is an affine function.

Corollary 1 (to Theorem 1). *If the demand function p is affine, then all optimal*

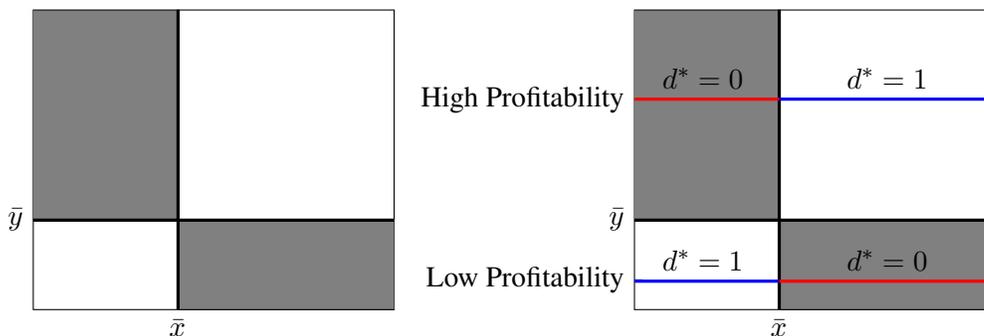


Figure 2: Optimal disclosure rule when p is **affine**. The gray areas represent signal realizations that are optimally concealed by the seller, and white areas are optimally disclosed.

disclosure rules have a threshold structure, with $\bar{x} = x^{ND} \in \text{int}(\mathcal{X})$ and $\bar{y}(x) = y^{ND} \in \text{int}(\mathcal{Y})$ for all $x \in \mathcal{X}$.¹²

When the demand function p faced by the seller is affine, as an implication of the martingale property of beliefs, all disclosure functions yield the same overall probability of sale. That is, for any two disclosure functions d and d' , $\mathbb{E}[P(y, d)] = \mathbb{E}[P(y, d')]$. This means that, in choosing a disclosure policy, there is no scope for the seller to increase or decrease the overall probability that the buyer purchases the object. Rather, the seller solely distributes this constant sale probability between high and low profitability objects.

In order to optimally *steer sales* from low- to high-profitability objects, the seller uses a disclosure rule that assigns objects to two classes, based on their profitability: high profitability, with $y > y^{ND}$, and low profitability, with $y < y^{ND}$. One such disclosure rule is depicted in Figure 2. Given this rule, the buyer's posterior mean for low profitability objects is *at best* equal to x^{ND} , which is the posterior mean induced when the signal realizations above x^{ND} are not disclosed. Conversely, the buyer's induced posterior mean for high profitability is *at least* equal to x^{ND} , which is induced when signal realizations below x^{ND} are not disclosed.

Nonlinear Demand. If the demand function p is not affine, then the informativeness of the seller's disclosure rule impacts the overall probability of sale.

Observation 1. Suppose disclosure rule d is strictly more informative than disclosure rule d' . Then (i) if p is strictly convex, d yields a higher probability of sale than d' ; and (ii) if p is strictly concave, d yields a lower probability of sale than d' .

¹²The fact that the thresholds must be interior does not follow directly from Theorem 1. We provide a proof of that claim in the Appendix.

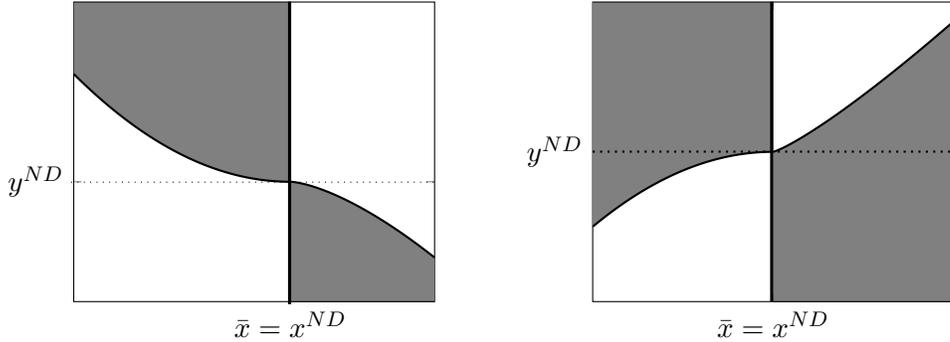


Figure 3: Optimal disclosure rule when p is **strictly convex** (left panel) and **strictly concave** (right panel). The gray areas represent signal realizations that are optimally concealed by the seller, and the white areas are optimally disclosed.

With a nonlinear demand function, the seller sales from low- to high-profitability objects at the expense of the total probability that the object is sold. If the seller faces a convex demand function, he wishes to disclose more information in order to maximize the sale probability. Conversely, if facing a concave demand, overall sale probability increases when information is concealed from the buyer.

Corollary 2 (to Theorem 1). *If p is strictly convex (concave), any optimal disclosure rule has a threshold structure with a strictly decreasing (increasing) profitability threshold function $\bar{y}(x)$, satisfying $\bar{y}(x^{ND}) = y^{ND}$.*

An optimal disclosure rule when p is strictly convex is represented in the left panel of Figure 3. A striking difference when comparing to the linear demand case is that some “very bad news” may be optimally disclosed even when the object is highly profitable; and, conversely, some “very good news” are disclosed to the receiver even when the object has low profitability. This feature is derived from the tension between the desire to maximize the probability of sale (which pushes towards disclosure) and to covary the probability of sales with profitability (which requires some concealment). In the opposite case, where p is strictly concave, due to an analogous tension, the sender optimally conceals some “very bad news” about low-profitability objects and some “very good news” about high-profitability objects. This is depicted in the right panel of Figure 3.¹³

¹³Corollary 2 shows that the “shape” of optimal disclosure rules depends on the curvature of the demand function p . In the current model, the demand function is taken as a primitive, and its curvature does not necessarily reflect some economically interpretable property. However, in a related model, Hwang, Kim and Boleslavsky (2019) study an oligopolistic market where agents compare the expected value of objects offered by different firms, and purchase the highest-value object. In

3.1 Is Transparency Good?

Let’s consider the following policy intervention, which makes the motives of the seller transparent to the buyer: When the buyer observes the seller’s “advice” – either the disclosed signal realization or non-disclosure – they also perfectly observe how profitable the object’s sale is to the seller. This should be thought of as a transparency mandate imposed on the seller by some regulator, forcing them to always reveal their interest in the object’s sale.

Naturally, anticipating that the buyer will see the profitability of the object, the seller optimally uses disclosure strategy which differs from the payoff maximizing disclosure rule without transparency. Indeed, with transparency, the seller’s problem becomes continuum of *separate* disclosure problems, one for each profitability.

Lemma 1. *With mandated transparency, the seller chooses, for each y , a disclosure policy $d_y : \mathcal{X} \rightarrow [0, 1]$ to maximize $yP(y, d_y)$, where*

$$P(y, d_y) = \int_{\mathcal{X}} [d_y(x)p(x) + (1 - d_y(x))p(x_w^{ND})] dF_{X|y}(x)$$

$$\text{and } x_y^{ND} = \frac{\int_{\mathcal{X}} [1 - d_y(x)]x dF_{X|y}(x)}{\int_{\mathcal{X}} [1 - d_y(x)] dF_{X|y}(x)}.$$

Under transparency, the seller cannot use strategic disclosure in order to steer sales across objects with different profitability levels. Rather, they separately maximize the probability that the object with each profitability level gets sold (or minimize it, if the profitability is negative). Proposition 1 below describes the seller’s optimal disclosure strategy under transparency.

Proposition 1. *Consider the benchmark with mandated transparency. If $\text{sign}(y)p$ is strictly concave (convex), then the seller optimally conceals (discloses) all signal realizations when the object’s profitability is y .*

Whether transparency is good or bad in terms of the informativeness of the seller’s advice to the buyer depends not only on the curvature of the demand function, but also on the sign of the object’s profitability. Suppose p is strictly concave. Then, under transparency, the seller conceals all signal realizations when y is positive, and discloses all realizations when y is negative. Thus, conditional on y being

that model, the “demand function” for a particular firm’s object is given by the distribution of the maximum value among the objects offered by that firm’s competitors. Hwang, Kim and Boleslavsky (2019) show that, as the number of competitors increases, that demand function becomes “more convex.” From that perspective, the convexity of the demand function is related to the competitiveness of the market.

positive (negative), the seller is less (more) informative under transparency. And the converse is true if p is strictly concave.

The corollary below states a version of this argument. It shows that, if the object's profitability is always positive, then transparency enhances the seller's informativeness if the demand function p is strictly convex; and makes the seller *less informative* if the demand function is strictly concave.

Corollary 3 (to Proposition 1). *Suppose $y_{min} > 0$. If p is strictly concave (convex) then under mandated transparency, the seller is Blackwell less informative (Blackwell more informative) than with hidden motives.*

If the demand function is affine, then under mandated transparency, the seller is indifferent between all disclosure rules. In particular, full transparency is one optimal policy. This result is, again, an implication of the martingale property of beliefs. Interestingly, if there is even a small probability that the seller's motives will not be transparent to the buyer, then the seller optimally discloses *as if* his motives are fully hidden (with the disclosure rule described in Corollary 1). This failure of upper hemicontinuity in the set of seller's optimal disclosure policies is registered in Proposition 2 below.

Proposition 2. *If the demand function p is affine, and the buyer observes the seller's motives with probability $\tau \in [0, 1]$, then*

- (a) *If $\tau = 1$, the seller is indifferent between all disclosure rules.*
- (b) *If $\tau \neq 1$, the seller's optimal disclosure rule is invariant to τ and satisfies the conditions in Corollary 1.*

4 Hidden Motives and Information Acquisition

In the baseline model, the signal about the object's value that is observed by the seller is taken as given. I now study how the seller's hidden motives affect their incentive to acquire information about the object's value in the first place, knowing that they can selectively disclose the acquired information to the buyer. To study this problem, we focus on the linear demand case.

[Assumption] For the remainder of the section, we assume that p is affine.

4.1 Signal Acquisition Technology

Prior to the object's value and profitability realizing, the seller commits to two actions: they *acquire a costly signal* of the object's value and (as before) choose a

rule to disclose that signal. The signal acquisition is *overt*, so that the buyer knows the signal acquired by the seller, even if the signal realization itself is not disclosed.

Signals are indexed by their precision $\theta \in \mathbb{R}_+$, and are acquired at a cost $c(\theta)$, where c is strictly increasing. The set of signals available to the seller is such that

- (i) For every θ and every $y \in \mathcal{Y}$, $F_{X|y}$ has full support over some interval $\mathcal{X}(\theta)$.¹⁴
- (ii) More precise signals are more informative than less precise ones: $\theta' > \theta$ implies that, for every $y \in \mathcal{Y}$, $F_{X|y}(\cdot, \theta')$ is a mean-preserving spread of $F_{X|y}(\cdot, \theta)$.
- (iii) The signal with precision $\theta = 0$ is perfectly uninformative: for every $y \in \mathcal{Y}$, $F_{X|y}(\cdot, 0)$ is a degenerate distribution.

4.2 Transparency and Information Acquisition

We want to consider how mandated transparency of the seller's motives affect their decision to acquire information about the object's value. Consider the policy that discloses the object's profitability to the receiver with probability $\tau \in [0, 1]$.

Proposition 3. *If $\tau = 1$, the seller optimally acquires the least precise signal, that is, $\theta = 0$. Moreover, take $\tau > \tau'$ and let θ and θ' be signal precisions optimally acquired under transparency policies τ and τ' , respectively. Then $\theta \leq \theta'$.*

Proposition 3 shows that transparency negatively affects the seller's incentives to acquire information about the object's value. Indeed, the seller always acquires a *less precise signal* when faced with a policy that mandates *more transparency*. Moreover, under full transparency, the seller has no incentives to acquire any information, and chooses the least precise (and least costly) signal available.

Remember that, when the seller's motives are hidden, they can strategically disclose information about the object's value to the buyer in order to steer their purchases from low- to high-profitability objects. But in order to selectively inform the buyer, the seller must have information about the object's value in the first place. Indeed, the more precise a signal the seller acquires, the more profitably they can steer the buyer's action. Now consider the policy that mandates transparency: when the seller's motives are (with some probability) made transparent to the buyer, their ability to steer sales between objects with different profitabilities is hindered. In turn, their incentives to acquire information about the object's value are also dampened, and they choose to invest in a less precise signal.

While Proposition 3 unequivocally establishes that transparency hinders information acquisition, it does not imply that the buyer is unequivocally better informed when the seller's motives are less transparent. Because the seller does not

¹⁴Along with assumption (ii), this implies that $\theta' > \theta \Rightarrow \mathcal{X}(\theta) \subseteq \mathcal{X}(\theta')$.

fully disclose the information they acquire, the comparison between the seller's informativeness under two different transparency policies is not always unambiguous (in the Blackwell sense). We evaluate the informativeness of a seller who chooses a pair of precision and disclosure rule (θ, d) in terms of the distribution of posterior means (about the object's value) that is induced on the buyer, denoted $F^B(\cdot, \theta, d)$.

Proposition 4 establishes some comparisons of the seller's informativeness under different transparency regimes. First, it shows that transparency does not make the seller more Blackwell informative. (Note that, because the informativeness order is not complete, a not more informative signal is not necessarily less informative.) Next, Proposition 4 states that, under full transparency, the seller is Blackwell less informative than under partial or no transparency.

Finally, I establish some conditions on the structure of signals available to the seller such that more transparency necessarily implies that the seller becomes Blackwell less informative. I say a signal of precision θ is *regular* if there is a unique pair $(\hat{x}, \hat{y}) \in \text{int}(\mathcal{X} \times \mathcal{Y})$ such that

$$\mathbb{E}_\theta [x | (x - \hat{x})(y - \hat{y}) < 0] = \hat{x} \text{ and } \mathbb{E}_\theta [y | (x - \hat{x})(y - \hat{y}) < 0] = \hat{y}.$$

Regularity guarantees that the seller has a unique optimal disclosure rule, given acquired precision θ . Next, I say a signal of precision θ is *symmetric* if, for each y , $F_{X|y}(\cdot, \theta)$ is a symmetric distribution of signal realizations; and, for each signal realization x , the profitability distribution conditional on x realizing, $F_{Y|x}(\cdot, \theta)$, is symmetric. Symmetry holds, for example, if the object's profitability and the distribution of signals about its value are independently distributed, and the distribution of signals is symmetric.

Proposition 4. *Let $\tau > \tau'$, and let (θ, d) and (θ', d') be signal precisions and disclosure rules optimally chosen by the seller under transparency policies τ and τ' , respectively. Then:*

1. $F^B(\cdot, \theta, d)$ is not strictly Blackwell more informative than $F^B(\cdot, \theta', d')$.
2. If $\tau = 1$, then $F^B(\cdot, \theta', d')$ is Blackwell more informative than $F^B(\cdot, \theta, d)$.
3. If signals of precision θ and θ' are regular and symmetric, then $F^B(\cdot, \theta', d')$ is Blackwell more informative than $F^B(\cdot, \theta, d)$.

We can also evaluate the seller's informativeness according to some completion of the Blackwell order, call it \succsim .¹⁵ Say that the set of signals available to the seller

¹⁵A completion of the Blackwell order \succsim is a complete order over $\Delta\mathcal{X}$ such that $G, G' \in \Delta\mathcal{X}$ and G is Blackwell more informative than G' implies that $G \succsim G'$.

is such that *precision is beneficial* if endowing the seller with a more precise signal weakly improves the seller’s optimally disclosed signal, according to the order \succsim .

Formally, fix some transparency level $\tau \in [0, 1)$, and suppose that, for every θ , the signal of precision θ is regular – so that for each θ , there is a unique optimal disclosure rule d_θ . Also note (per Proposition 2) that such optimal disclosure rule is independent of τ , so long as $\tau \neq 1$. We say that *precision is beneficial* if $\theta \geq \theta'$ implies that $F^B(\cdot, \theta, d_\theta) \succsim F^B(\cdot, \theta', d_{\theta'})$ – or, equivalently, if $F^B(\cdot, \theta, d_\theta)$ is \succsim -more-informative than $F^B(\cdot, \theta', d_{\theta'})$.

One possible such completion of the Blackwell order is the buyer’s surplus. In that case, precision is beneficial if a buyer would prefer to be advised by a seller who has more precise information. Corollary 4 below shows that, when precision is beneficial, transparency is always detrimental if evaluated in terms of \succsim -informativeness.

Corollary 4 (to Proposition 3). *Assume that precision is beneficial. Let $\tau > \tau'$, and let (θ, d) and (θ', d') be signal precisions and disclosure rules optimally chosen by the seller under transparency policies τ and τ' , respectively. Then $F^B(\cdot, \theta', d')$ is \succsim -more-informative than $F^B(\cdot, \theta, d)$.*

4.3 An Example

Suppose the object’s value and profitability are independent and both are distributed according to $U[0, 1]$. Further, the seller can acquire signals of precision $\theta \in [0, 1]$ about the object’s value, subject to a strictly increasing cost $c(\theta)$.

Given the acquired signal, a piece of hard evidence about the object’s value realizes, denoted $e \in [0, 1]$. With probability θ , e is equal to the true value of the object, and with probability $(1 - \theta)$, e is a random number drawn from $U[0, 1]$, uncorrelated with the object’s true value. The mean posterior implied by a piece of evidence $e \in [0, 1]$ is

$$x = \theta e + (1 - \theta) \frac{1}{2}.$$

And so the distribution of mean posteriors implied by θ is $F_{X|y}(\cdot, \theta) = U\left[\frac{1-\theta}{2}, \frac{1+\theta}{2}\right]$, for every y . It is easy to see that every signal θ is symmetric, and that $\theta' > \theta$ implies that, for every $y \in \mathcal{Y}$, $F_{X|y}(\cdot, \theta')$ is a mean-preserving spread of $F_{X|y}(\cdot, \theta)$. I also show in the Appendix that every signal θ is regular.

Because every signal θ is regular and symmetric, any optimal disclosure rule has a threshold structure with thresholds $\bar{x} = \mathbb{E}(x) = 1/2$ and $\bar{y} = \mathbb{E}(y) = 1/2$.¹⁶

¹⁶In the proof of Statement 3 in Proposition 4, I show that regularity and symmetry implies that $\bar{x} = \mathbb{E}(x)$ and $\bar{y} = \mathbb{E}(y)$ in any optimal disclosure rule.

In short, if the object has “above average” profitability, the seller discloses evidence about that object’s value if and only if the evidence induces an “above average” posterior. Conversely, if the object has “below average” profitability, the seller discloses evidence about that object’s value if and only if the evidence induces a “below average” posterior.

Accounting for the acquired precision θ and the optimal disclosure rule, the distribution of posterior means induced on the buyer is

$$F^B(x, \theta, d) = \begin{cases} 0, & \text{if } x < \frac{1-\theta}{2}, \\ \frac{1}{2\theta} \left[x - \frac{1-\theta}{2} \right], & \text{if } x \in \left[\frac{1-\theta}{2}, \frac{1}{2} \right), \\ \frac{1}{2} + \frac{1}{2\theta} \left[x - \frac{1-\theta}{2} \right], & \text{if } x \geq \frac{1}{2}. \end{cases}$$

In words, with probability $1/2$ the seller does not disclose the signal realization, and the buyer’s posterior mean is equal to $x^{ND} = 1/2$. With half probability the seller discloses the signal realization – either below average realizations for less profitable objects or above average realizations for the more profitable objects – in which case the distribution of posteriors induced on the buyer is $U \left[\frac{1-\theta}{2}, \frac{1+\theta}{2} \right]$.

Because $\theta' > \theta$ implies that $U \left[\frac{1-\theta'}{2}, \frac{1+\theta'}{2} \right]$ is a strict mean-preserving spread of $U \left[\frac{1-\theta}{2}, \frac{1+\theta}{2} \right]$, we can conclude that, when transparency increases, the seller acquires strictly less precise signals and provides strictly less information to the buyer. (This example illustrates result 3 in Proposition 4.)

5 Alternative Communication Protocols

The model makes two main assumptions regarding the communication protocol. The first one is that the seller can commit to a rule to disclose signal realizations, prior to the object’s profitability or the “evidence” about its value being drawn. This commitment assumption is typical to the information design literature, recently surveyed in Kamenica (2019). Second, we assume that the seller can use only disclosure policies, à la Dye (1985), rather than committing to more general signal structures.

In this section, we consider two variations of the model, dropping each of these two assumptions. Under each alternative communication protocol, we investigate whether mandating transparency about the seller’s motives improves the informativeness of their advice to the buyer.

5.1 Disclosure with No Commitment

Consider the following disclosure protocol with no commitment: First, the seller observes the object’s profitability and a signal realization. After that, the seller

chooses whether to disclose the signal realization (a piece of *evidence* about the object’s value) to the buyer. The buyer observes the disclosed evidence – or the fact that no evidence was disclosed – and forms a posterior mean about the object’s value, taking into account the seller’s equilibrium strategy. As before, the buyer does not observe the object’s profitability.¹⁷ The equilibrium notion is Perfect Bayesian Equilibrium.¹⁸

Proposition 5. *For any demand function p , an equilibrium exists, and any equilibrium disclosure strategy d^* has a threshold structure: $d^*(y, x) \in \{0, 1\}$ and¹⁹*

$$d^*(y, x) = 1 \Leftrightarrow (x - \bar{x})y \geq 0,$$

for some $\bar{x} \in \mathcal{X}$, satisfying $\bar{x} = \mathbb{E}[x | (x - \bar{x})y < 0]$ if $\{(x - \bar{x})y < 0\} \neq \emptyset$.

In evidence disclosure models where the advisor’s motives are transparent – for example, in the seminal works of Grossman (1981), Milgrom (1981) and Milgrom and Roberts (1986) – equilibria feature full information revelation. The main insight is that any candidate equilibrium with some evidence concealment would “unravel” due to the receiver’s (the buyer’s) rational skepticism.

Indeed, the threshold equilibria described in Proposition 5 can be consistent with the unravelling logic. Suppose the object’s profitability is always positive ($y_{min} \geq 0$), and conjecture a threshold equilibrium with some evidence concealment – that is, suppose $\bar{x} \in (x_{min}, x_{max})$. Then, because y is always greater than 0, it must be that $\bar{x} > \mathbb{E}[x | (x - \bar{x})y < 0]$, and thus the equilibrium condition in Proposition 5 is not satisfied. Such a conjectured equilibrium would unravel: the buyer’s posterior upon observing non-disclosure would be $\mathbb{E}[x | (x - \bar{x})y < 0]$, which is strictly smaller than \bar{x} . Consequently, when the seller draws a signal realization just under \bar{x} , they strictly prefer to reveal it to the buyer, which is inconsistent with the initially conjectured equilibrium. An analogous unravelling argument applies if the object’s profitability is always negative ($y_{max} \leq 0$).

However, when profitability can be both positive or negative, so that $0 \in (y_{min}, y_{max})$, then any equilibrium involves partial disclosure. There is some interior threshold \bar{x} such that the seller discloses (only and) all “good news” about the object ($x \geq \bar{x}$) when profitability is positive, and (only and) all “bad news” ($x \leq \bar{x}$) about the object when profitability is negative. Partially uninformative

¹⁷We assume that the seller cannot choose to disclose the object’s profitability to the buyer. This assumption is not restrictive, because as we shall see, in equilibrium the seller would never want to disclose the object’s profitability.

¹⁸Because of the evidence structure and our restriction to the evidence disclosure protocol, there is no “informed principal” problem when defining an equilibrium.

¹⁹Assuming that, when indifferent, the seller discloses the signal realization.

equilibria do not unravel, because, upon observing non-disclosure, the buyer does not know whether the seller has “good news” but negative profitability or “bad news” but positive profitability.

In such an equilibrium, the seller’s strategy is akin to *sanitization strategies* as in Dye (1985): conditional on any particular profitability level y , the seller discloses only sufficiently favorable information to the buyer. (With the caveat that, to a seller with negative profitability, only “bad news” are sufficiently favorable.) Also note that equilibrium disclosure strategies are independent of the shape of the demand function p – in contrast to the model with commitment, where optimal disclosure rules differ based on the curvature of p . The following Corollary to Proposition 5 tells us that, when p is affine, it is possible for the optimal disclosure rule with commitment to also be an equilibrium disclosure strategy in the game without commitment.²⁰

Corollary 5 (to Proposition 5). *Suppose p is affine, and suppose the seller has an optimal disclosure rule (with commitment) with profitability threshold $\bar{y} = 0$. Then there is an optimal disclosure rule d^* , with $\bar{y} = 0$, which is also an equilibrium disclosure strategy in the evidence disclosure game with no commitment.*

Now let’s turn to the question of whether mandated transparency incentivizes the seller to disclose information about the object’s value to the buyer. When the seller’s evidence structure is taken as given (that is, when we do not consider the seller’s problem of acquiring information in the first place), then mandating transparency is unequivocally beneficial for informativeness. Indeed, under transparency, the seller is fully informative.

Proposition 6. *Under transparency, full disclosure is the unique equilibrium.*

5.1.1 Overt Information Acquisition

Proposition 6 shows that, under the no commitment protocol, transparency never hinders the seller’s incentive to provide information. However, even without commitment in the disclosure stage, transparency hurts the seller’s incentives to acquire information. Suppose (as before) that we augment the disclosure model with the initial stage where the seller (overtly) chooses to acquire a costly signal about the object’s value. This stage happens before the object’s profitability is drawn. When

²⁰Some of the literature on voluntary disclosure, for example Glazer and Rubinstein (2004, 2006), Sher (2011), and more recently Titova (2022), argues that, under some restrictions on the relation between the sender’s and receiver’s payoffs, the commitment assumption is superfluous. Corollary 5 shows that commitment may also be redundant in a voluntary disclosure model where the sender has hidden motives.

p is affine, Proposition 3 holds, even with the no commitment communication protocol – the proof is exactly parallel to the proof of Proposition 3. Further, results (analogous to) Proposition 4 and Corollary 4 also apply.

5.2 Unconstrained Signaling Technology

The usual assumption in information design models is that a sender (the seller, in our case) commits a signal a map from states of the world (the pieces of evidence about the object’s value) into distributions of messages to inform a receiver (the buyer) about the state. The sender’s choice of such a map is unrestricted. Contrastingly, in this paper, I assumed that the sender is restricted to a class of “signaling strategies”: the class of disclosure rules, whereby the sender’s message either fully conveys the information in a piece of evidence, or is “silent.” Such silence is a message in itself, which conveys to the receiver that the state of the world is “one of the evidence-profitability pairs that would lead the sender to stay silent.”

In the *constrained information design* problem I study, Theorem 1 provides quite a complete characterization of the seller’s optimal messaging strategy. In comparison, a characterization of the sender’s optimal signal in the equivalent unconstrained design problem is elusive – Rayo and Segal (2010) provide a partial characterization of the optimal signal when the demand function p is affine.

Regardless, some of the results in this paper also hold in the “unconstrained design version” of the problem. For example, Proposition 1, and its Corollary 3 would still be true, so that transparency may be detrimental to the seller’s incentives to relay information about the object’s value to the buyer. Further, Proposition 3 would also be true in that unconstrained benchmark, highlighting that also in that case transparency hinders the sellers motivation to acquire information about the object’s value. However, precisely because we cannot characterize optimal signals in the unconstrained problem, it is hard to pinpoint if and when the extra information would be relayed to the receiver. Consequently, the effect of transparency on the seller’s overall informativeness to the receiver are unclear.

6 Discussion of an Application

The model describes a situation where a seller informs a potential buyer about the value of an object for sale, where the buyer does not know how profitable the sale of such object would be to the seller. Below, I interpret features of the model in the context of a brokerage company who provides financial advice to investors, but also receives commissions for the sale of financial products. In the context of financial brokers, clients understand that brokers receive sales commissions from

some product sales, but may not know the size of the commissions on each product (and hence their motives are hidden).

There are many other contexts in which people consume information disclosed by entities with hidden motives: doctors inform patients of the effectiveness of different drugs and procedures, but are often sponsored by pharmaceutical companies; online platforms gather reviews from product users, but may disclose them coarsely to benefit some producers; magazines and newspapers selectively publish pieces of reporting that align with their editorial bias. In all the mentioned settings, information receivers understand the information providers may be biased, but are not fully aware of the extent of the conflict.

6.1 Financial Advisors with Hidden Motives

When an investor opens an account at a brokerage company, they usually get some access to product reports put together by the company's research team. In addition to publicly available data about companies and industries, a report on a particular product includes forecasts and evaluations produced by the team. Upon seeing the provided research, the investor compares the perceived value of the product to an outside option, which could depend on their current appetite for investment, desire to reallocate their current portfolio, or even independent information she may have sourced about the financial product at hand.

Hidden Motives. The incentives of the advisor/broker may not be fully understood by the investor. For instance, some of the products available in the brokerage system are proprietary products that are issued or managed by the broker themselves. Upon selling one of these assets, the broker receives extra compensation. Another source of conflict is that third parties commonly pay the broker for marketing and selling their products, which may make the sale of some products more desirable than others. These considerations can fuel the broker's desire to produce advice that steers investors to the more profitable products. Indeed, it is well documented that brokers conceal bad news about companies in which they have financial interests and that financial advisors' recommendations favor products with high commissions.²¹ In line with our modeling assumption, in this context, clients understand that some products are more profitable to the brokers than others, but the specific profitability of each product's sale to the broker is *hidden*.

Disclosure. In the model, the seller produces information about a product and

²¹See, for example, Anagol, Cole and Sarkar (2017) and Eckardt and Rathke-Doppner (2010) about insurance brokers; Chalmers and Reuter (2020) on retirement plans; Inderst and Ottaviani (2012.2) on general financial advice. For a survey on quality disclosure and certification, see Dranove and Jin (2010).

chooses whether to share the outcome of the research with the buyer or not. This feature of the model matches the existence of regulatory “Chinese walls” within financial institutions. To ensure that the analyses produced by a broker are not biased by the companies’ profit-seeking interests, various regulations institute walls between research departments and client-facing teams. These walls ensure that the research production is not tainted by the profitability considerations of the brokerage company. As in the model, the reports produced can be understood as unbiased evidence about a product’s quality, while the disclosure of this evidence remains under the discretion of the sales team.

Regulation. Because investors often rely on professional advice when making investment decisions, financial advisory is a highly regulated field. (And has become increasingly so since the Great Recession and the institution of the Dodd-Frank Act.) In the United States, the regulation varies across states, and also across categories within the field – for example, financial advisors and insurance agents are often subject to more stringent regulation than broker-dealers, who often also provide advice. Specific details notwithstanding, much of the regulation focuses on making the interests of financial advisors transparent to investors who consult them. For example, they might be required to disclose other sources of compensation received beyond service fees, professional affiliations with another broker-dealer or securities issuer, and other potential or existing conflicts of interest. Other types of regulation also forbid advisors from receiving some types of compensation. For example, the UK and the Netherlands have altogether imposed bans on commission payments for some types of financial advisors.

Commitment. In information design models, a common interpretation of the commitment assumption is that it stands in for the sender’s desire to build and sustain *credibility* with the receiver in a game of repeated interaction. Think of a circumstance where a financial advisor repeatedly interacts with an investor, informing them about assets’ values over a long time. And suppose that this investor follows the advisor’s recommendation, but also sees the outcome of their investment decisions (so they can tell, to some extent, whether the advice was sound). Over time, the investor can “punish” unsatisfactory advice from the seller with future incredulity – maybe by not following their recommendation, maybe by seeking another advisor. Through that mechanism, in the repeated game, the advisor’s optimal disclosure strategy would be such that the buyer’s interpretation of a “no disclosure” message coincides with the interpretation yielded in the one-shot game under the commitment assumption. This point is formally made by Best and Quigley (2020), and a similar argument linking commitment and reputation is provided by Mathevet, Pearce and Stacchetti (2022).

A recent paper, Liu and Lin (2022) propose a new notion of credibility for information design problems: A disclosure policy is credible if the sender cannot profit from tampering with her messages while keeping the message distribution unchanged. It's useful to note that all optimal disclosure rules characterized in Theorem 1 are credible, according to this definition. In other words, if the sender can commit to a distribution of posterior means to induce on the receiver, then the optimal way to correlate these posterior means with the object's profitability is attained with threshold-structure disclosure rules as described in the Theorem. Consequently, the seller would not "tamper" with messages, conditional on maintaining the distribution of induced posteriors.

Related Literature. In a series of papers in 2012, Inderst and Ottaviani (2012.1, 2012.2, 2012.3) propose models of brokers and financial advisors compensated through commissions. In their model, competing sellers offer commissions to the advisor, knowing that he will steer business to the seller that offers highest compensation. The price of the asset is set equal to buyers' expected value for it, so that the consumer retains no surplus. An implication is that information is not valuable to the consumer. The authors show that biased commissioned agents may improve welfare by providing buyers with less information and steering business to high-commission firms who are also more cost efficient. My model takes an alternative approach, taking the sender's distribution of profitability as given and focusing on the value of the information provided to buyers. In my environment, information is always beneficial to the consumer and I show that hidden motives can improve surplus precisely by increasing the amount of information provided to the consumer.

There is also a large literature that studies the provision of information by Credit Rating Agencies that are financed by fees paid by issuers of financial products. Some key papers in this literature are Bolton, Freixas, and Shapiro (2012), Opp, Opp and Harris (2013), Bar-Isaac and Shapiro (2012) and Skreta and Veldkamp (2009). In this literature, the CRA receives payments equally from all issuers of financial products. In my model, the main concern is that the advisor benefits some products over others because they have different profitabilities.

References

- Anagol, S., S. Cole, and S. Sarkar. (2017) "Understanding the Advice of Commissions Motivated Agents: Evidence from the Indian Life Insurance Market." *Review of Economics and Statistics*, **99**, 1-15.
- Ball, I. and X. Gao. (2021) "Benefitting from Bias." *working paper*.

- Bar-Isaac, H., and J. Shapiro. (2013) “Ratings Quality over the Business Cycle.” *Journal of Financial Economics*, **108**, 62-78.
- Best, J., and D. Quigley. (2020) “Persuasion for the Long Run.” *working paper*.
- Bolton, P., X. Freixas, and J. Shapiro. (2012) “The Credit Ratings Game.” *The Journal of Finance*, **67**, 85-111. NBR 6023
- Chalmers, J. and J. Reuter. (2020) “Is Conflicted Investment Advice Better than no Advice?.” *Journal of Financial Economics*, forthcoming.
- Che, Y.K., and N. Kartik (2009) “Opinions as Incentives.” *Journal of Political Economy*, **117**, 815-860.
- Curello, G. and L. Sinander (2022) “The Comparative Statics of Persuasion,” *working paper*.
- DeMarzo, P. M., I. Kremer and A. Skrzypacz (2019) “Test Design and Minimum Standards.” *American Economic Review*, **109**: 2173-2207.
- Doval, L., and V. Skreta. (2021) “Constrained Information Design,” *working paper*.
- Dranove, D. and G. Jin. (2010) “Quality Disclosure and Certification: Theory and Practice.” *Journal of Economic Literature*, **48**: 935-63.
- Dye, R. A. (1985) “Disclosure of Nonproprietary Information.” *Journal of Accounting Research*, **23**: 123-145.
- Eckardt, M., and S. Rathke-Doppner. (2010) “The Quality of Insurance Intermediary Services – Empirical Evidence for Germany.” *Journal of Risk and Insurance*, **77**, 667-701.
- Ershov, D. and M. Mitchell. (2022) “The Effects of Advertising Disclosure Regulations on Social Media: Evidence From Instagram.” *working paper*.
- Glazer, J., and A. Rubinstein. (2004) “On the Optimal Rules of Persuasion.” *Econometrica*, **72**, 1715-1736.
- Glazer, J., and A. Rubinstein. (2006) “A Study in the Pragmatics of Persuasion: A Game Theoretical Approach.” *Theoretical Economics* **1**, 395-410.
- Grossman, S. (1981) “The Informational Role of Warranties and Private Disclosure about Product Quality.” *Journal of Law & Economics*, **24**: 461-484.
- Hwang, I., K. Kim, and R. Boleslavsky (2019). “Competitive Advertising and Pric-

ing,” *working paper*.

Inderst, R., and M. Ottaviani. (2012) “Competition through Commissions and Kickbacks.” *American Economic Review*, **102**, 780-809.

Inderst, R., and M. Ottaviani. (2012) “Financial Advice”. *Journal of Economic Literature*, **50**, 494-512.

Inderst, R., and M. Ottaviani. (2012) “How (not) to Pay for Advice: A Framework for Consumer Financial Protection.” *Journal of Financial Economics*, **105**, 393-411.

Jung, W.-O. and Y. K. Kwon (1988) “Disclosure When the Market Is Unsure of Information Endowment of Managers.” *Journal of Accounting Research*, **26**: 146-153.

Kamenica, E. (2019) “Bayesian Persuasion and Information Design.” *Annual Review of Economics*, **11**.

Kamenica, E., and M. Gentzkow. (2011) “Bayesian Persuasion.” *American Economic Review*, **101**, 2590-2615.

Kartik, N., F. X. Lee, and W. Suen. (2017) “Investment in Concealable Information by Biased Experts.” *The RAND Journal of Economics*, **48**: 24-43.

Lipnowski, E. and D. Ravid. (2020) “Cheap Talk with Transparent Motives.” *Econometrica*, **88**: 1631-1660.

Lin, X. and C. Liu. (2022) “Credible Persuasion,” *working paper*.

Mathevet, L., D. Pearce, and E. Stacchetti. (2022) “Reputation for a Degree of Honesty.” *working paper*.

Mensch, J. (2021) “Monotone Persuasion.” *Games and Economic Behavior*, **130**: 521-542.

Milgrom, P. (1981) “Good News and Bad News: Representation Theorems and Applications.” *The Bell Journal of Economics*, **12**: 380-391.

Milgrom, P. (2008) “What the Seller Won’t Tell You: Persuasion and Disclosure in Markets.” *Journal of Economic Perspectives*, **22**: 115-131.

Milgrom, P. and J. Roberts. (1986) “Relying on the Information of Interested Parties.” *RAND Journal of Economics*: 18-32.

- Onuchic, P. (f) D. Ray. (2022) “Conveying Value via Categories,” *working paper*.
- Opp, C., M. Opp, and M. Harris. (2013) “Rating Agencies in the Face of Regulation.” *Journal of Financial Economics*, **108**, 46-61.
- Rayo, L., and I. Segal (2010) “Optimal Information Disclosure.” *Journal of Political Economy*, **118**, 949–987.
- Shishkin, D. (2022) “Evidence Acquisition and Voluntary Disclosure.” *working paper*.
- Sher, I. (2011) “Credibility and Determinism in a Game of Persuasion.” *Games and Economic Behavior*, **71**, 409-419.
- Skreta, V., and L. Veldkamp. (2009) “Ratings Shopping and Asset Complexity: A Theory of Ratings Inflation.” *Journal of Monetary Economics*, **56**, 678-695.
- Szalay, D. (2005) “The Economics of Clear Advice and Extreme Options.” *The Review of Economic Studies*, **72**: 1173-1198.
- Titova, M. (2022) “Persuasion with Verifiable Information.” *working paper*.

Appendix

Proof of Theorem 1

I first complete the proof in the main text, showing that any disclosure rule that does not satisfy the described threshold structure described can be improved upon. To that end, consider d and \hat{d} as in (5) and (3). We prove the following two claims used in the main text.

Claim 1. d and \hat{d} produce the same overall probability of sale.

Proof.

$$\begin{aligned}
\mathbb{E}[P(y, \hat{d})] - \mathbb{E}[P(y, d)] &= \\
&= \int_{\mathcal{Y}} \int_{\mathcal{X}} [p(x) - p(x^{ND})] [\hat{d}(y, x) - d(y, x)] dF_{X|y}(x) dF_Y(y) \\
&= \int_{\mathcal{X}} [p(x) - p(x^{ND})] \int_{\mathcal{Y}} [\hat{d}(y, x) - d(y, x)] dF_{Y|x}(y) dF_X(x) = 0
\end{aligned}$$

where $F_{Y|x}$ is the profitability distribution conditional on a signal realization x and F_X is the marginal distribution of signal realizations. The first equality uses the

definition of $P(y, d)$ and the third is due to d and \hat{d} disclosing each realization with the same probability, as in (7). ■

Claim 2. \hat{d} induces a larger covariance between sales and profitability than d .

Proof.

$$\begin{aligned}
& \text{Cov} \left[y, P(y, \hat{d}) \right] - \text{Cov} \left[y, P(y, d) \right] = \mathbb{E} \left[\left(P(y, \hat{d}) - P(y, d) \right) (y - \mathbb{E}(y)) \right] \\
& = \int_{\mathcal{Y}} \int_{\mathcal{X}} [p(x) - p(x^{ND})] \left[\hat{d}(y, x) - d(y, x) \right] [w - \mathbb{E}(y)] dF_{X|y}(x) dF_Y(y) \\
& = \int_{\mathcal{X}} [p(x) - p(x^{ND})] \int_{\mathcal{Y}} \left[\hat{d}(y, x) - d(y, x) \right] [y - \mathbb{E}(y)] dF_{Y|x}(y) dF_X(x)
\end{aligned} \tag{8}$$

By the definition of \hat{d} , for $x < x^{ND}$, $\hat{d}(y, x) - d(y, x) \geq 0$ when $y < \hat{y}(x)$ and $\hat{d}(y, x) - d(y, x) \leq 0$ when $y > \hat{y}(x)$ – but, as given by (7), the expected difference $\hat{d}(y, x) - d(y, x)$ is 0. This, along with the fact that $y - \mathbb{E}(y)$ is increasing in y , implies that

$$\int_{\mathcal{Y}} \left[\hat{d}(y, x) - d(y, x) \right] [y - \mathbb{E}(y)] dF_{Y|x}(y) \leq 0,$$

when $x \leq x^{ND}$. Analogously, we can show that

$$\int_{\mathcal{Y}} \left[\hat{d}(w, x) - d(w, x) \right] [w - \mathbb{E}(y)] dF_{Y|x}(y) \geq 0,$$

when $x > x^{ND}$. Moreover, these inequalities are strict for a positive measure of signal realizations. These observations, along with the fact that p is strictly increasing, deliver that the expression in (8) is strictly positive. ■

Now we prove the following Lemma:

Lemma 2. *In any optimal disclosure rule, the profitability threshold satisfies (4).*

Proof. Using (3), for $y \in \mathcal{Y}$ and $x \in \mathcal{X}$, we can take a derivative of the sender's value with respect to $d(y, x)$, to get

$$\begin{aligned}
\frac{\partial \Pi}{\partial d(y, x)} &= y (p(x) - p(x^{ND})) dF_{X|y}(x) dF_Y(y) \\
&\quad + \left(\int_{\mathcal{Y}} \int_{\mathcal{X}} \tilde{y} [1 - d(\tilde{y}, \tilde{x})] dF_{X|\tilde{y}}(\tilde{x}) dF_Y(\tilde{y}) \right) p'(x^{ND}) \frac{\partial x^{ND}}{\partial d(y, x)}
\end{aligned}$$

Now from (1), we get

$$\frac{\partial x^{ND}}{\partial d(y, x)} = \frac{\int_{\mathcal{Y}} \int_{\mathcal{X}} (\tilde{x} - x)(1 - d(\tilde{y}, \tilde{x})) dF_{X|\tilde{y}}(\tilde{x}) dF_Y(\tilde{y})}{\left(\int_{\mathcal{Y}} \int_{\mathcal{X}} (1 - d(\tilde{y}, \tilde{x})) dF_{X|\tilde{y}}(\tilde{x}) dF_Y(\tilde{y}) \right)^2} dF_{X|y}(x) dF_Y(y)$$

Substituting this into the previous equation, we have

$$\frac{\partial \Pi}{\partial d(y, x)} = [y(p(x) - p(x^{ND})) - y^{ND} p'(x^{ND})(x - x^{ND})] dF_{X|y}(x) dF_Y(y), \quad (9)$$

where y^{ND} is the average object profitability given non-disclosure. It is easy to check that, if $x < x^{ND}$,

$$\frac{\partial \Pi}{\partial d(y, x)} \begin{cases} > 0, & \text{if } y < y^{ND} \left[\frac{p'(x^{ND})(x^{ND} - x)}{p(x^{ND}) - p(x)} \right] \\ < 0, & \text{if } y > y^{ND} \left[\frac{p'(x^{ND})(x^{ND} - x)}{p(x^{ND}) - p(x)} \right] \end{cases}$$

Conversely, if $x > x^{ND}$,

$$\frac{\partial \Pi}{\partial d(y, x)} \begin{cases} > 0, & \text{if } y > y^{ND} \left[\frac{p'(x^{ND})(x^{ND} - x)}{p(x^{ND}) - p(x)} \right] \\ < 0, & \text{if } y < y^{ND} \left[\frac{p'(x^{ND})(x^{ND} - x)}{p(x^{ND}) - p(x)} \right] \end{cases}$$

Now suppose that, with positive probability, either $d(w, x) \neq 0$ when (9) is negative or $d(w, x) \neq 1$ when (9) is positive. Then d cannot be optimal. So any optimal disclosure rule must have a threshold structure where the profitability threshold satisfies

$$\bar{y}(x) = y^{ND} \left[\frac{p'(x^{ND})(x^{ND} - x)}{p(x^{ND}) - p(x)} \right]$$

for all $x \neq x^{ND}$. Equivalently, the profitability threshold satisfies (4). \blacksquare

To complete the proof of Theorem 1, we show that an optimal disclosure rule exists.

Lemma 3. *An optimal disclosure rule exists.*

Proof. Let \mathcal{D} be the set of all disclosure rules $d : \mathcal{Y} \times \mathcal{X} \rightarrow [0, 1]$. Also define the following subset of \mathcal{D} , the set of threshold disclosure rules: \mathcal{T} is the set of disclosure rules such that $d(y, x) \in \{0, 1\}$ and $d(y, x) = 1 \Leftrightarrow (x - \bar{x})(y - \hat{y}(x)) > 0$ for some $\bar{x} \in \mathcal{X}$ and $\hat{y}(x)$ satisfying, for some $\bar{y} \in \mathcal{Y}$

$$\bar{y}(x) = \bar{y} \left[\frac{p'(\bar{x})(\bar{x} - x)}{p(\bar{x}) - p(x)} \right].$$

Further, define the following “discretized version” of \mathcal{D} : for each $N > 0$, $d \in \mathcal{D}^N$ if $d \in \mathcal{D}$ and $d(y, x) = d(y', x')$ whenever $y_{\min} + \frac{n^Y}{N}(y_{\max} - y_{\min}) \leq y, y' < y_{\min} + \frac{n^Y+1}{N}(y_{\max} - y_{\min})$ and $x_{\min} + \frac{n^X}{N}(x_{\max} - x_{\min}) \leq x, x' < x_{\min} + \frac{n^X+1}{N}(x_{\max} - x_{\min})$ for some $n^Y, n^X \in \{0, \dots, N-1\}$. Note that the disclosure rules in \mathcal{D}^N are fully defined by N^2 numbers in $[0, 1]$.

And define also a discretized version of \mathcal{T} : for every $N > 0$, $d \in \mathcal{T}^N$ if $d \in \mathcal{D}^N$ and there is some $\bar{x} \in \mathcal{X}$ and $\bar{y} \in \mathcal{Y}$ such that (i) $d(x, y) = 1$ if $(x' - \bar{x})(y' - \hat{y}(x)) > 0$ for every x' and y' with $y_{\min} + \frac{n^Y}{N}(y_{\max} - y_{\min}) \leq y, y' < y_{\min} + \frac{n^Y+1}{N}(y_{\max} - y_{\min})$ and $x_{\min} + \frac{n^X}{N}(x_{\max} - x_{\min}) \leq x, x' < x_{\min} + \frac{n^X+1}{N}(x_{\max} - x_{\min})$ for some $n^Y, n^X \in \{0, \dots, N-1\}$; and (ii) $d(x, y) = 0$ if $(x' - \bar{x})(y' - \hat{y}(x)) < 0$ for every x' and y' with $y_{\min} + \frac{n^Y}{N}(y_{\max} - y_{\min}) \leq y, y' < y_{\min} + \frac{n^Y+1}{N}(y_{\max} - y_{\min})$ and $x_{\min} + \frac{n^X}{N}(x_{\max} - x_{\min}) \leq x, x' < x_{\min} + \frac{n^X+1}{N}(x_{\max} - x_{\min})$ for some $n^Y, n^X \in \{0, \dots, N-1\}$.

Now we use the following facts about these sets:

Fact 1. For every N , $\sup_{d \in \mathcal{D}^N} \Pi(d) = \sup_{d \in \mathcal{T}^N} \Pi(d)$.

Proof of Fact 1. As previously noted, the disclosure rules in \mathcal{D}^N are fully defined by N^2 numbers in $[0, 1]$. This, along with the continuity of the objective, implies that there exists some $\hat{d} \in \mathcal{D}^N$ such that $\Pi(\hat{d}) = \sup_{d \in \mathcal{D}^N} \Pi(d)$. But, with an argument analogous to the proof of Lemma 2, we can show that if $d \neq \mathcal{T}^N$, then $\Pi(d) < \sup_{d \in \mathcal{D}^N} \Pi(d)$. And so it must be that $\hat{d} \in \mathcal{T}^N$, completing the proof. ■

Fact 2. There exists some d^* such that $\Pi(d^*) = \sup_{d \in \mathcal{T}} \Pi(d)$.

Proof of Fact 2. Disclosure rules in \mathcal{T} are fully defined by two numbers, $\bar{x} \in \mathcal{X}$ and $\bar{y} \in \mathcal{Y}$, whenever $(x - \bar{x})(y - \hat{y}(x)) \neq 0$. But, because the joint distribution of profitabilities and signal realizations have no mass points, the disclosure rule at the equality points do not affect the seller’s expected payoff. And so, effectively, when choosing a disclosure rule in \mathcal{T} , the seller chooses only two numbers, from a compact set. This, along with continuity of the objective, implies the existence of some d^* such that $\Pi(d^*) = \sup_{d \in \mathcal{T}} \Pi(d)$. ■

Now we have $\lim_{N \rightarrow \infty} \sup_{d \in \mathcal{D}^N} \Pi(d) = \sup_{\mathcal{D}} \Pi(d)$ and $\lim_{N \rightarrow \infty} \sup_{d \in \mathcal{T}^N} \Pi(d) = \sup_{\mathcal{T}} \Pi(d)$. But, by Fact 1, $\sup_{d \in \mathcal{D}^N} \Pi(d) = \sup_{d \in \mathcal{T}^N} \Pi(d)$ for every N . And so $\sup_{d \in \mathcal{D}} \Pi(d) = \sup_{d \in \mathcal{T}} \Pi(d)$. But then, by Fact 2, there exists some d^* such that $\Pi(d^*) = \sup_{d \in \mathcal{T}} \Pi(d) = \sup_{d \in \mathcal{D}} \Pi(d)$, which concludes the existence proof. ■

Proof of Corollary 1

If a disclosure rule conceals signal realizations with strictly positive probability, then it must be that $x^{ND} \in \text{int}(\mathcal{X})$ and $y^{ND} \in \text{int}(\mathcal{Y})$. Next, I argue that the fully revealing disclosure rule is not optimal.

Remember that the demand function is affine, so $p(x) = a+bx$. Then it must be that for any disclosure rule d , $\mathbb{E}[P(y, d)] = a\mathbb{E}(y) + b\mathbb{E}(x)$. (This is a straightforward consequence of the martingale property of beliefs, or “Bayesian plausibility.”) Therefore the overall probability of sale is independent of the disclosure rule.

Now consider a full revelation disclosure rule d^1 and some other threshold disclosure rule d , defined by interior thresholds \bar{x} and \bar{y} . It must be that, for $y < \bar{y}$, $P(y, d) < P(y, d^1)$ and, for $y > \bar{y}$, $P(y, d) > P(y, d^1)$, and so the covariance of sales and profitability is larger under d than under d^1 ; and thus d^1 is not an optimal disclosure rule. (This last argument relies on the fact that the distribution of profitabilities and signal realizations has full support.) ■

Proof of Lemma 1

Suppose the seller chooses disclosure rule d . Upon observing non-disclosure and profitability y (under mandated transparency), the buyer’s mean posterior is

$$x_y^{ND} = \frac{\int_{\mathcal{X}} [1 - d(y, x)] x dF_{X|y}(x)}{\int_{\mathcal{X}} [1 - d(y, x)] dF_{X|y}(x)}.$$

And so the probability of sale of object y is

$$P(y, d) = \int_{\mathcal{X}} [d(y, x)p(x) + (1 - d(y, x))p(x_y^{ND})] dF_{X|y}(x).$$

Note that the probability of sale of object y is thus independent of the disclosure rule used for objects with profitability $y' \neq y$. And so the seller’s problem is separable across profitability levels. Therefore maximizing $\Pi(d)$ over $d : \mathcal{Y} \times \mathcal{X} \rightarrow [0, 1]$ is equivalent to maximizing, for each $y \in \mathcal{Y}$, $yP(y, d_y)$, over $d_y : \mathcal{X} \rightarrow [0, 1]$. ■

Proof of Proposition 1

For any given $d_y : \mathcal{X} \rightarrow [0, 1]$,

$$yP(y, d_y) = \int_{\mathcal{X}} yp(x) dF_y^B(x, d_y),$$

where F_y^B is the distribution of posterior means induced on the buyer when profitability is y and the disclosure rule is d_y . Suppose $\text{sign}(y)p$ is everywhere strictly

concave, and take two disclosure rules d_y and d'_y , where $F_y^B(\cdot, d_y)$ is a mean preserving spread of $F_y^B(\cdot, d'_y)$. This immediately implies that, $\mathbb{E}[P(d_y)] \leq \mathbb{E}[P(d'_y)]$. Now let $d^0 : \mathcal{X} \rightarrow [0, 1]$ be the fully concealing disclosure rule – that is, for all $x \in \mathcal{X}$, $d^0(x) = 0$. And note that, for any d_y , $F_y^B(\cdot, d_y)$ is a mean preserving spread of $F_y^B(\cdot, d^0)$, and so d^0 maximizes $\mathbb{E}[yP(y, d_y)]$. By Lemma 1, it is then a solution to the seller's problem.

An analogous argument shows that, if $\text{sign}(y)p$ is everywhere strictly convex, then the fully disclosing rule d^1 maximizes $\mathbb{E}[yP(y, d_y)]$, and d^1 is thus a solution to the seller's problem. ■

Proof of Proposition 2

Suppose $p(x) = a + bx$, with $a, b \in \mathbb{R}$. Then, under full transparency, for any disclosure rule d ,

$$yP(y, d) = \int_{\mathcal{X}} y[a + bx]dF_y^B(x, d_y),$$

where $d_y(\cdot) = d(y, \cdot)$ and F_y^B is the distribution of posterior means induced on the buyer when profitability is y and the disclosure rule is d_y . And so

$$yP(y, d) = ya + yb \int_{\mathcal{X}} xdF_y^B(x, d_y) = ya + yb\mathbb{E}[x|y].$$

The second equality is due to the fact that, under any disclosure rule, F_y^B has average equal to $\mathbb{E}[x|y]$, the underlying expected value of the object given that the profitability is y . (F_y^B is a mean preserving spread of the degenerate distribution at the object's expected value, given the profitability that is observed by the buyer due to the mandated transparency.) Therefore, under full transparency, we have

$$\Pi^1(d) = \mathbb{E}[yP(y, d)] = a\mathbb{E}(y) + b\mathbb{E}[y\mathbb{E}(x|y)], \quad (10)$$

which is independent of the disclosure rule, proving part (a) of the proposition. (Denote by Π^1 the seller's value under full transparency; and by Π^0 the value when profitability is not revealed to the buyer.)

Now if $\tau \neq 1$, then the seller's value is

$$\Pi(d) = \tau\Pi^1(d) + (1 - \tau)\Pi^0(d) = \tau \{a\mathbb{E}(y) + b\mathbb{E}[y\mathbb{E}(x|y)]\} + (1 - \tau)\Pi^0(d).$$

Because the first term is constant, it follows that the seller's optimal disclosure rule maximizes $\Pi^0(d)$, proving part (b) of the proposition. ■

Proof of Proposition 3

Let $\Pi_\tau(d, \theta)$ be the expected payoff to the seller who has the signal of precision θ , uses disclosure rule d , and has their motives disclosed to the buyer with probability τ . And let $\bar{\Pi}_\tau(\theta) = \max_d \Pi_\tau(d, \theta)$.

Lemma 4. $\bar{\Pi}_0(\theta)$ is weakly increasing in θ .

Proof of Lemma 4. Remember that $p(x) = a + bx$. Then for each θ , and each y ,

$$yP(y, d, \theta) = ay + by \left[\int_{\mathcal{X}} xd(y, x)dF_{X|y}(x, \theta) + \int_{\mathcal{X}} x^{ND}(1 - d(y, x))dF_{X|y}(x, \theta) \right],$$

$$\text{where } x^{ND} = \frac{\int_{\mathcal{Y}} \int_{\mathcal{X}} x(1 - d(y, x))dF_{X|y}(x, \theta)dY(y)}{\int_{\mathcal{Y}} \int_{\mathcal{X}} (1 - d(y, x))dF_{X|y}(x, \theta)dY(y)}.$$

Now take two precision levels $\theta' > \theta$, and remember that, for each y , $F_{X|y}(\cdot, \theta')$ is a mean-preserving spread of $F_{X|y}(\cdot, \theta)$. And so, for any disclosure rule d , there exists some disclosure rule d' such that

$$\int_{\mathcal{X}} xd'(y, x)dF_{X|y}(x, \theta') = \int_{\mathcal{X}} xd(y, x)dF_{X|y}(x, \theta),$$

$$\text{and } \int_{\mathcal{X}} d'(y, x)dF_{X|y}(x, \theta') = \int_{\mathcal{X}} d(y, x)dF_{X|y}(x, \theta),$$

for every $y \in \mathcal{Y}$. And thus $yP(y, d', \theta') = yP(y, d, \theta)$, and $\Pi_0(d', \theta') = \Pi_0(d, \theta)$. Because this is true for any disclosure rule d , it must be that $\bar{\Pi}_0(\theta') \geq \bar{\Pi}_0(\theta)$. ■

Lemma 5. $\bar{\Pi}_1(\theta)$ is independent of θ .

Proof. From equation (10), we know that under full transparency,

$$\bar{\Pi}_1(\theta) = a\mathbb{E}(y) + b\mathbb{E}[y\mathbb{E}(x|y, \theta)] = a\mathbb{E}(y) + b\mathbb{E}[y\mathbb{E}(x|y)]. \quad (11)$$

The last equality is due to the fact that the signal's precision does not affect the underlying expected value of the object, for a given y . ■

By Lemmas 4 and 5, and the fact that $\Pi(d, \theta) = \tau\Pi^1(d, \theta) + (1 - \tau)\Pi^0(d, \theta)$ for every disclosure rule d , we have

$$\bar{\Pi}_\tau(\theta) = \tau\bar{\Pi}_1(\theta) + (1 - \tau)\bar{\Pi}_0(\theta),$$

where the first term is constant and the second term is weakly increasing in θ .

Because $\bar{\Pi}_1(\theta)$ is constant and c is strictly increasing in θ , it must be that if $\tau = 1$, the seller acquires $\theta = 0$ (proving the first statement in the proposition).

Now take two transparency levels $\tau' > \tau$ and suppose θ and θ' are signal precisions optimally acquired by the seller under each transparency level. If $\theta = \theta'$, then the second statement in the proposition is true. Suppose instead that $\theta \neq \theta'$. Then, by the optimality of θ under τ ,

$$\begin{aligned} \tau \bar{\Pi}_1(\theta) + (1 - \tau) \bar{\Pi}_0(\theta) - c(\theta) &\geq \tau \bar{\Pi}_1(\theta') + (1 - \tau) \bar{\Pi}_0(\theta') - c(\theta'). \\ \Rightarrow (1 - \tau) \bar{\Pi}_0(\theta) - c(\theta) &\geq (1 - \tau) \bar{\Pi}_0(\theta') - c(\theta'). \end{aligned} \quad (12)$$

And by the optimality of θ' under τ' ,

$$\begin{aligned} \tau' \bar{\Pi}_1(\theta') + (1 - \tau') \bar{\Pi}_0(\theta') - c(\theta') &\geq \tau' \bar{\Pi}_1(\theta) + (1 - \tau') \bar{\Pi}_0(\theta) - c(\theta). \\ \Rightarrow (1 - \tau') \bar{\Pi}_0(\theta') - c(\theta') &\geq (1 - \tau') \bar{\Pi}_0(\theta) - c(\theta). \end{aligned} \quad (13)$$

Adding up the two inequalities, we have

$$(\tau' - \tau) \bar{\Pi}_0(\theta) \geq (\tau' - \tau) \bar{\Pi}_0(\theta') \Rightarrow \bar{\Pi}_0(\theta) \geq \bar{\Pi}_0(\theta').$$

Note that $\bar{\Pi}_1(\theta) = \bar{\Pi}_1(\theta')$ contradicts either (12) or (13), because c is strictly increasing in θ . And so it must be that $\bar{\Pi}_0(\theta) > \bar{\Pi}_0(\theta')$, and so $\theta > \theta'$, completing the proof. \blacksquare

Proof of Proposition 4

Statement 2 in the proposition follows trivially from the fact that, under $\tau = 1$, the seller acquires $\theta = 0$, and the signal $\theta = 0$ is perfectly uninformative.

Statement 1. $F^B(\cdot, \theta, d)$ is not strictly Blackwell more informative than $F^B(\cdot, \theta', d')$.

Proof of Statement 1. If $\tau = 1$, then statement 1 is a consequence of statement 2. Assume instead that $\tau \neq 1$. Because d is an optimal disclosure rule under τ , it is a threshold disclosure rule (by Theorem 1), and so $F^B(\cdot, \theta, d)$ is defined as follows. There is some $\bar{x} \in \text{int}(\mathcal{X})$ and $\bar{y} \in \text{int}(\mathcal{Y})$ such that:

◦ If $x < \bar{x}$,

$$F^B(x, \theta, d) = \int_{[y_{\min}, \bar{y})} F_{X|y}(x, \theta) dF_Y(y).$$

◦ If $x \geq \bar{x}$,

$$\begin{aligned} F^B(x, \theta, d) &= \int_{\mathcal{Y}} F_{X|y}(x, \theta) dF_Y(y) + \int_{[y_{\min}, \bar{y})} (1 - F_{X|y}(x, \theta)) dF_Y(y) \\ &= F_Y(\bar{y}) + \int_{[\bar{y}, y_{\max}]} F_{X|y}(x, \theta) dF_Y(y). \end{aligned}$$

The distribution $F^B(\cdot, \theta', d')$ can be expressed analogously, substituting θ with θ' and \bar{x} and \bar{y} with some potentially different $\bar{x}' \in \text{int}(\mathcal{X})$ and $\bar{y}' \in \text{int}(\mathcal{Y})$. Consider three cases:²²

Case 1. $\bar{x} = \bar{x}'$ and $\bar{y} = \bar{y}'$. In this case, the fact that for each y , $F_{X|y}(\cdot, \theta')$ is a mean-preserving spread of $F_{X|y}(\cdot, \theta)$ implies that $F^B(x, \theta', d')$ is a mean preserving spread of $F^B(x, \theta, d)$. (And so the former is Blackwell more informative.)

Case 2. $\bar{y} \neq \bar{y}'$. Suppose first that $\bar{y}' > \bar{y}$. Then there is some $x < \min\{\bar{x}, \bar{x}'\}$ such that

$$\begin{aligned} \int_{x_{\min}}^x F^B(\hat{x}, \theta', d') d\hat{x} &= \int_{x_{\min}}^x \int_{[y_{\min}, \bar{y}']} F_{X|y}(\hat{x}, \theta') dF_Y(y) d\hat{x} \\ &= \int_{y_{\min}}^{\bar{y}'} \int_{x_{\min}}^x F_{X|y}(\hat{x}, \theta') d\hat{x} dF_Y(y) > \int_{y_{\min}}^{\bar{y}} \int_{x_{\min}}^x F_{X|y}(\hat{x}, \theta') d\hat{x} dF_Y(y) \\ &\geq \int_{y_{\min}}^{\bar{y}} \int_{x_{\min}}^x F_{X|y}(\hat{x}, \theta) d\hat{x} dF_Y(y) = \int_{x_{\min}}^x F^B(\hat{x}, \theta, d) d\hat{x}. \end{aligned} \quad (14)$$

The first inequality used the fact that $\bar{y}' > \bar{y}$ and that for each y , $F_{X|y}(\cdot, \theta')$ has full support over $\mathcal{X}(\theta')$ (and thus over $\mathcal{X}(\theta) \subseteq \mathcal{X}(\theta')$). The second inequality used the fact that for each y , $F_{X|y}(\cdot, \theta')$ is a mean-preserving spread of $F_{X|y}(\cdot, \theta)$. But (14) implies that $F^B(\hat{x}, \theta, d)$ is not a mean preserving spread of $F^B(\hat{x}, \theta', d')$.

An analogous argument can be used to show that, if $\bar{y}' < \bar{y}$, there is some $x > \max\{\bar{x}, \bar{x}'\}$ such that $\int_{[x, x_{\max}]} (1 - F^B(\hat{x}, \theta', d')) d\hat{x} > \int_{[x, x_{\max}]} (1 - F^B(\hat{x}, \theta, d)) d\hat{x}$, and so $F^B(\hat{x}, \theta, d)$ is not a mean preserving spread of $F^B(\hat{x}, \theta', d')$.²³

Case 3. $\bar{y} \neq \bar{y}'$ and $\bar{x} \neq \bar{x}'$. Suppose first that $\bar{x}' < \bar{x}$. Then there exists some $x \in (\bar{x}', \bar{x})$ such that

$$\begin{aligned} \int_{x_{\min}}^x F^B(\hat{x}, \theta', d') d\hat{x} &= \int_{x_{\min}}^x \left[Y(\bar{y}) + \int_{\bar{y}}^{y_{\max}} F_{X|y}(\hat{x}, \theta) dF_Y(y) \right] d\hat{x} \\ &> \int_{x_{\min}}^x \int_{y_{\min}}^{\bar{y}} F_{X|y}(\hat{x}, \theta') dF_Y(y) d\hat{x} = \int_{y_{\min}}^{\bar{y}} \int_{x_{\min}}^x F_{X|y}(\hat{x}, \theta') d\hat{x} dF_Y(y) \\ &\geq \int_{y_{\min}}^{\bar{y}} \int_{x_{\min}}^x F_{X|y}(\hat{x}, \theta) d\hat{x} dF_Y(y) = \int_{x_{\min}}^x F^B(\hat{x}, \theta, d) d\hat{x}. \end{aligned}$$

The first inequality used the fact that $\bar{y}' > \bar{y}$ and that for each y , $F_{X|y}(\cdot, \theta')$ has full support over $\mathcal{X}(\theta')$ (and thus over $\mathcal{X}(\theta) \subseteq \mathcal{X}(\theta')$). The second inequality used the

²²The following arguments repeatedly use the definition of mean preserving spreads in Shaked and Shanthikumar (2007), equation (3.A.8).

²³This argument would use the equivalent definition of mean preserving spreads, in equation (3.A.7) of Shaked and Shanthikumar (2007).

fact that for each y , $F_{X|y}(\cdot, \theta')$ is a mean-preserving spread of $F_{X|y}(\cdot, \theta)$. Again, this implies that $F^B(\hat{x}, \theta, d)$ is not a mean preserving spread of $F^B(\hat{x}, \theta', d')$. An analogous argument can be used to show that if $\bar{x}' > \bar{x}$, there is some $x \in (\bar{x}, \bar{x}')$ such that $\int_{[x, x_{max})} (1 - F^B(\hat{x}, \theta', d')) d\hat{x} > \int_{[x, x_{max})} (1 - F^B(\hat{x}, \theta, d)) d\hat{x}$ and so $F^B(\hat{x}, \theta, d)$ is not a mean preserving spread of $F^B(\hat{x}, \theta', d')$. ■

Statement 3. If signals of precision θ and θ' are regular and symmetric, then $F^B(\cdot, \theta', d')$ is Blackwell more informative than $F^B(\cdot, \theta, d)$.

Proof of Statement 3. If a signal of precision θ is regular, then if

$$\mathbb{E}_\theta [x|(x - \hat{x})(y - \hat{y}) < 0] = \hat{x} \text{ and } \mathbb{E}_\theta [y|(x - \hat{x})(y - \hat{y}) < 0] = \hat{y}, \quad (15)$$

for some $\hat{x} \in \text{int}(\mathcal{X})$ and $\hat{y} \in \text{int}(\mathcal{Y})$, then the threshold disclosure rule with thresholds \hat{x} and \hat{y} is the unique optimal disclosure rule. (Corollary 1 implies that any optimal disclosure rule satisfies (15) and regularity implies that there is a unique such pair (\hat{x}, \hat{y}) .)

Now suppose the signal of precision θ is symmetric and note that

$$\begin{aligned} \mathbb{E}_\theta [x|(x - \mathbb{E}(x))(y - \mathbb{E}(y)) < 0] &= \quad (16) \\ &= \int_{y_{min}}^{\mathbb{E}(y)} \int_{\mathbb{E}(x)}^{x_{max}} x dF_{X|y}(x, \theta) dF_Y(y) + \int_{\mathbb{E}(y)}^{y_{max}} \int_{x_{min}}^{\mathbb{E}(x)} x dF_{X|y}(x, \theta) dF_Y(y) \\ &= \int_{y_{min}}^{\mathbb{E}(y)} \int_{\mathbb{E}(x)}^{x_{max}} x dF_{X|y}(x, \theta) dF_Y(y) + \int_{y_{min}}^{\mathbb{E}(y)} \int_{x_{min}}^{\mathbb{E}(x)} x dF_{X|y}(x, \theta) dF_Y(y) \\ &= \int_{y_{min}}^{\mathbb{E}(y)} \int_{x_{min}}^{x_{max}} x dF_{X|y}(x, \theta) dF_Y(y) = \mathbb{E}(x), \end{aligned}$$

where $\mathbb{E}(x)$ and $\mathbb{E}(y)$ are the underlying unconditional expected value and profitability of the object, respectively. The second and last equalities in (16) use the fact that the signal is symmetric. Analogously, we can show that

$$\mathbb{E}_\theta [y|(x - \mathbb{E}(x))(y - \mathbb{E}(y)) < 0] = \mathbb{E}(y).$$

And so the optimal disclosure rule given acquired signal θ must use thresholds $\bar{x} = \mathbb{E}(x)$ and $\bar{y} = \mathbb{E}(y)$. But the same is true for θ' , because signal θ' is also regular and symmetric. And so, by Case 1 in the proof of Statement 1, it must be that $F^B(x, \theta', d')$ is a mean preserving spread of $F^B(x, \theta, d)$. (And so the former is Blackwell more informative.) ■

Proof of Claim in Example 4.3

In Example 4.3, I made the following claim:

Claim 3. *The signal structure described by $Y = U[0, 1]$ and $F_y(\cdot; \theta) = U\left[\frac{1-\theta}{2}, \frac{1+\theta}{2}\right]$ for every y is regular.*

Proof. For cleaner algebra, I write the proof for the case with $\theta = 1$, so that $F_y(\cdot; \theta) = U[0, 1]$. The proof for $\theta \in [0, 1)$ follows the same steps.

Remember that a signal of precision θ is *regular* if there is a unique pair $(\hat{x}, \hat{y}) \in \text{int}(\mathcal{X} \times \mathcal{Y})$ such that

$$\mathbb{E}_\theta [x|(x - \hat{x})(y - \hat{y}) < 0] = \hat{x} \text{ and } \mathbb{E}_\theta [y|(x - \hat{x})(y - \hat{y}) < 0] = \hat{y}. \quad (17)$$

Given the distributions in the example, (17) becomes

$$\frac{\hat{y}\hat{x}\frac{\hat{x}}{2} + (1 - \hat{y})(1 - \hat{x})\frac{1+\hat{x}}{2}}{\hat{x}\hat{y} + (1 - \hat{y})(1 - \hat{x})} = \hat{x} \text{ and } \frac{\hat{y}\hat{x}\frac{\hat{y}}{2} + (1 - \hat{y})(1 - \hat{x})\frac{1+\hat{y}}{2}}{\hat{x}\hat{y} + (1 - \hat{y})(1 - \hat{x})} = \hat{y}.$$

$$\Rightarrow \hat{y}\hat{x}\hat{x} = (1 - \hat{y})(1 - \hat{x})(1 - \hat{x}) \text{ and } \hat{y}\hat{y}\hat{x} = (1 - \hat{y})(1 - \hat{y})(1 - \hat{x}).$$

Remember that we are looking for an interior solution, so $\hat{x}, \hat{y} \in (0, 1)$. Dividing the first equation by the second, we have

$$\frac{\hat{x}}{\hat{y}} = \frac{1 - \hat{x}}{1 - \hat{y}},$$

which has a unique solution at $\hat{x} = \hat{y} = 1/2$. ■

Proof of Proposition 5

Assume $y_{min} < 0 < y_{max}$. (The case where all profitabilities are positive or all profitabilities are negative is discussed in the main text.)

Suppose the buyer's mean posterior about the object's value after observing non-disclosure is $\hat{x} \in \mathcal{X}$. Then a seller with profitability $y > 0$, upon observing a signal realization x , will choose to disclose it if and only if $x \geq \hat{x}$ (assuming that the seller discloses when indifferent). Conversely, a seller with profitability $y < 0$ will disclose signal realization x if and only if $x \leq \hat{x}$. A seller with profitability $y = 0$ is always indifferent between disclosing and not disclosing, and we resolve this indifference with disclosure.

And so Bayesian consistency requires that, in this equilibrium,

$$\hat{x} = \mathbb{E}[x|(x - \hat{x})y < 0],$$

where we note that the set $\{(x - \bar{x})y < 0\}$ has positive probability, because the joint distribution of profitabilities and signal realizations has full support.

Regarding existence, I remark that there exists some \hat{x} satisfying

$$\hat{x} = \mathbb{E}[x|(x - \hat{x})y < 0].$$

To see that, note that $\mathbb{E}[x|(x - \hat{x})y < 0] \in (x_{min}, x_{max})$ for both $\hat{x} = x_{min}$ and $\hat{x} = x_{max}$ (both because the joint distribution of profitabilities and signal realizations has full support). And so, by continuity of $\mathbb{E}[x|(x - \hat{x})y < 0]$, there must be some solution to $\hat{x} = \mathbb{E}[x|(x - \hat{x})y < 0]$. Given such \hat{x} , it is easy to see that $d^*(y, x) \in \{0, 1\}$ and $d^*(y, x) = 1 \Leftrightarrow (x - \hat{x})y \geq 0$ defines an equilibrium disclosure strategy. ■

Proof of Proposition 6

Under mandated transparency, the buyer forms mean posteriors about the object separately for each profitability level y . And so it follows from standard unravelling arguments that there can be no concealment in equilibrium. ■