ADVISORS WITH HIDDEN MOTIVES*

Paula Onuchic
London School of Economics

July, 2025

Abstract

An advisor discloses evidence about an object to a potential buyer, who doesn't know the object's value or the profitability of its sale (the advisor's motives). I characterize optimal disclosure rules that balance two goals: maximizing the overall probability of sale, and steering sales from lower- to higher-profitability objects. I consider the implications of a regulation that forces the advisor to always reveal her motives to the buyer. I show that whether such policies induce the advisor to disclose more evidence about the object's value hinges on the curvature of the buyer's demand for the object. This result refines our understanding of effective regulation of advisor-advisee communication with and without commitment.

1 Introduction

People frequently take advice from advisors with hidden motives: broker-dealers and other financial advisors counsel investors, but also receive sales commissions from financial product providers; digital influencers provide product reviews to their followers, but these are often sponsored content; doctors inform patients of the effectiveness of different drugs and procedures, but may be rewarded by pharmaceutical companies; magazines and newspapers selectively publish pieces of reporting that align with their editorial bias. In all the mentioned settings, information receivers understand that information providers may be biased, but do not know the extent of the conflict in each interaction: clients understand that brokers receive sales commissions from some product providers, but may not know the size of the commissions on each product; in the social media context, followers understand that influencers post sponsored content, but may not know which exact publications are paid advertising.

In this paper, I propose a model where a potentially biased sender informs a receiver about an underlying state. The sender communicates with the receiver by committing to a policy to disclose information about the state. The main results are twofold. First, I characterize features of optimal communication,

*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, Nageeb Ali, Ricardo Alonso, Arjada Bardhi, Heski Bar-Isaac, Sylvain Chassang, Navin Kartik, Laurent Mathevet, Luis Rayo, Mauricio Ribeiro, and Dezsö Szalay for their helpful comments.

and relate them to the sender's tradeoffs between informing the receiver about the state and steering them towards the sender's preferred action. In the context of social media marketing, I relate this characterization to influencers' strategic usage of two information tools, "reviews" and "endorsements." Second, I compare the informativeness of advisors with hidden motives to a benchmark in which conflicts are made transparent to their advisees. This transparent benchmark can be interpreted as arising from regulatory interventions or as a "direct marketing" counterpart to influencer-intermediated marketing. I find that transparency need not make the receiver better off; indeed, depending on the primitives, the receiver may obtain better information from a sender whose motives are hidden.

Section 2 introduces the model. The sender is an advisor who communicates about an object's value to a buyer, who then chooses whether to acquire the object. Apart from its value to the buyer, the object is also described by the profitability of its potential sale to the advisor. This profitability, which determines the extent to which the advisor wishes to "push" the object's sale, is unknown to the buyer and therefore corresponds to the sender's hidden motives. Prior to the realization of the object's value and profitability, the advisor commits to a rule to disclose evidence that conveys the object's value to the buyer. This rule, which can depend on the object's profitability, assigns to each evidence realization a probability that it is disclosed to the buyer. In the context of social media influencers, we can interpret a disclosed signal as the provision of a truthful review of the object, and non-disclosure can be interpreted as an endorsement of the object by the influencer, who "brands" it with their approval stamp without providing any direct information about its value. Considering that setting, we can therefore view the sender's problem as that of choosing the optimal endorsing-versus-reviewing strategy.

The buyer is Bayesian and updates their belief about the object's value based on any information the advisor reveals *or fails to reveal* (since strategic non-disclosure is itself a signal of the object's value). The probability that the buyer purchases the object is given by their "demand function," which is increasing in their posterior expectation of the object's value. This demand function is taken as a primitive of the model, but in Section 2.1, I provide a possible micro-foundation in which the demand function as arises from the buyer choosing between the object being offered by the advisor and one or more alternative outside options. The proposed micro-foundation introduces a possible interpretation of the curvature of the demand function as a measure of competition in the market to which the advisor belongs. Specifically, a "more convex" demand function would arise in a more competitive market, where the buyer has access to more potential outside options.

In Theorem 1, I show that optimal disclosure rules feature *aligned disclosure*: the object's value is conveyed to the buyer (only) when the interests of the advisor and buyer are aligned. This happens either when both the object's value and its profitability are high, or when both are low, where high versus low values and profitabilities are defined by optimally chosen thresholds.[1] Conversely, the advisor chooses no disclosure after observing misaligned object realizations — either profitable objects with low value or high value objects with low profitability. By committing to an aligned disclosure rule, the advisor induces an ambiguous meaning to "no disclosure:" the advisee cannot fully disentangle whether "no

[1]The Theorem, along with its Corollaries 1 and 2, describes how the thresholds that define optimal disclosure policies are determined by a key primitive of the communication environment, the curvature of the buyer's demand function.

news" means "bad news about the object's value (and that the object's sale is profitable), or whether "no news" means "the object's sale is not profitable to the advisor" (and good news about the object's value). The advisor profits from this ambiguity, steering purchases from low- to high-profitability objects.

In section 4, I compare the advisor's optimal disclosure policy to that in a benchmark in which their motives are made transparent to the receiver, that is, a setting in which on top of observing any information about the object that is willingly disclosed by the advisor, the buyer also sees the object's profitability to the advisor. The goal in this comparison is to assess whether the buyer is better off when receiving information from a sender with hidden or transparent motives, a question that is relevant from a perspective of regulating advice markets.

Proposition 1 shows that making the advisor's motives transparent can increase or decrease the informativeness of the optimal disclosure policy, depending on features of the buyer's demand function. Under a concave demand, the advisor voluntarily discloses some evidence to the buyer when their motives are hidden, which makes them more informative than in the transparent benchmark in which the optimal policy is to not disclose any information. The opposite is true when the buyer's demand function is convex: the transparent advisor optimally discloses all evidence, and is therefore more informative than in the case of hidden motives. Propositions 2 and 3 further describe the relation between features of the buyer's demand function and the comparison between the transparent and hidden motives settings.

From a theoretical perspective, this result illustrates that, in a model of communication with commitment, the 'alignment between sender and receiver preferences' and the 'opaqueness of the sender's motives' are distinct objects; and it is not necessarily true that regulations that reduce the latter would also make the sender's interests more aligned with those of the receiver. In terms of a regulatory takeaway, these results argue that the effectiveness of a transparency policy can depend on the curvature of the buyer's demand function, highlighting that there is no "one size fits all" policy to optimally regulate advice markets. Rather, it is important to fit the regulation to specific features of each market. Using the interpretation introduced in section 2.1, which relates the curvature of the buyer's demand function to the degree of competition in the relevant market, we learn that mandating transparency is an effective policy in very competitive markets, but may not be so in markets where the buyer does not have many alternatives to the object being offered by the advisor.

To further contrast the theoretical distinction between the (mis-)alignment of sender and receiver preferences and the opaqueness of the sender's motives in a communication model with commitment, in section 5, I consider a variation of the environment, in which the advisor makes disclosure choices without commitment. In this case, unlike in the commitment setting, Propositions 4 and 5 show that a transparency policy is effective in increasing the information provided to the buyer.

This paper contributes to the literature on strategic communication by considering a "disclosure with commitment" communication protocol, combining (simple) evidence disclosure, as in Grossman (1981) and Milgrom (1981), with the timing of committed communication typical of the Bayesian persuasion literature following Kamenica and Gentzkow (2011).[2] A recent paper that similarly combines these two

[2]Both Grossman (1981) and Milgrom (1981) consider disclosure games with a richer evidence structure (in which the

ingredients is Antic and Chakraborty (2024). Our papers differ in that we consider different sender and receiver preferences, as well as distinct evidence structures. More broadly, disclosure with commitment is connected to a 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 (2023) consider monotonicity constraints (see Doval and Skreta (2024) and references therein for other examples).

My paper, like Rayo and Segal's (2010) work, departs from most of the literature on information design in that it considers a problem with a multidimensional state. And it differs from Rayo and Segal (2010) both because I consider a disclosure communication protocol, which allows me to characterize optimal policies when the buyer's demand function is not linear, and because I study the introduction of policies that make the advisor's motives transparent.[3] Most of Rayo and Segal's (2010) characterization results apply to the linear specification. With 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. One of Rayo and Segal's (2010) main results is that, in the linear benchmark, optimal signals never pool ordered prospects.

In terms of regulation of advice markets, Matthews and Postlewaite (1985) show that mandatory disclosure can lead a seller to acquire less information and thereby worsen information sharing. In my paper, the seller does not acquire information, and mandated transparency is about her motives, not the object's value. Other works that consider disclosure with costly information acquisition are Che and Kartik (2009), Kartik, Lee, and Suen (2017), and Libgober (2022).[4]

In section 5, I also consider a model in which the advisor makes disclosure choices without commitment. If the object's sale is always profitable or always unprofitable to the advisor, unraveling ensues as in Grossman (1981) and Milgrom (1981). However, if the object's sale is sometimes profitable and sometimes unprofitable, then equilibrium disclosures are only partial.The existence of partly-uninformative equilibria as described in Proposition 4 is in line with Seidmann and Winter's (1997) observation that in disclosure environments where the sender's preferred action depends on their type, there may exist equilibria in which their type is not always revealed in equilibrium.[5] I use this characterization to argue that, absent commitment, mandated transparency always makes the buyer (weakly) better off.

Finally, this paper also relates to previous work studying cheap talk models in which the sender has hidden motives. For example, Sobel (1985), Morris (2001), and Morgan and Stocken (2003), study environments with cheap talk communication in which the receiver does not know whether the sender's

sender can use any truthful message of the type "the state belongs to set D"). The evidence structure in my model is what Hagenbach and Koessler (2017) refer to as "simple evidence," where the sender simply chooses between a fully informative message (disclosure) or a fully uninformative message (no dislcosure).

[3]In section 5.2, I briefly comment on a variation in which the seller can commit to signaling structures more general than simple disclosure policies. That variation corresponds to the problem considered in Rayo and Segal's (2010).

[4]See also Shishkin (2023) and DeMarzo, Kremer, and Skrzypacz (2019) about information acquisition by the sender in a Dye (1985) framework. And see Szalay (2005) and Ball and Gao (2024) about information acquisition by a biased agent in a delegation context. An earlier version of this paper also shows that the introduction of a transparency policy can hinder the advisor's incentives to acquire information about the objects value; and therefore also affect the informativeness of their advice through that channel.

[5]Mezzetti (2025) makes a similar observation.

preferences are aligned with their own. My paper is most related to Li and Madarasz (2008), which studies a version of Crawford and Sobel's (1982) cheap talk environment, with the additional assumption that the receiver does not know the size or direction of the sender's bias. Like me, Li and Madarasz (2008) ask whether instituting a policy that mandates the transparency of the sender's motives (their bias in that case) is gainful for the receiver, and find that such mandated transparency policy may be ineffective.

2 Environment

A an advisor 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 advisor. (The advisor could be the original seller of the object, for example, or a salesperson who obtains a commission from influencing the buyer to purchase the good.) The object's *value* to the buyer, denoted $x \in \mathcal{X} = [x_{min}, x_{max}]$, and its *profitability* to the advisor, denoted $y \in \mathcal{Y} = [y_{min}, y_{max}]$, with $y_{min} \geqslant 0$, are drawn from a joint distribution commonly known by advisor and buyer.

Before the buyer decides whether to purchase the object, the advisor can disclose to them some evidence conveying 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 $\hat{x} \in \mathcal{X}$. Given their posterior $\hat{x}$ about the object's value, the buyer purchases the object with probability $p(\hat{x})$, where $p : \mathcal{X} \to [0, 1]$ is a strictly increasing and continuously differentiable "demand function." If the object is purchased, the advisor receives a payoff equal to the object's profitability, y. Otherwise, the advisor's payoff is 0.

Information Disclosure Protocol. The advisor has access to a piece of evidence that conveys the object's value.[6] I denote by F be the joint distribution over $\mathcal{X} \times \mathcal{Y}$ of the object's value and its profitability. F_Y is the marginal profitability distribution, and $F_{X|y}$ denotes the distribution of values conditional on a profitability $y \in \mathcal{Y}$. I assume that distributions F_Y and $F_{X|y}$, for each $y \in \mathcal{Y}$, have strictly positive densities f_Y and $f_{X|y}$ for each y.

At an initial stage, the advisor commits to a rule to disclose evidence to the buyer. A disclosure rule is a measurable map from the object's realized value *and its profitability* into a probability that the realization is disclosed, $d : \mathcal{X} \times \mathcal{Y} \to [0, 1]$.[7] Note that the disclosure decision depends both on the evidence realization and on the profitability of the object, so in practice the advisor commits to rules to disclose the object's value, conditional on its profitability.

[6] The analysis extends to the case in which the evidence does not fully reveal the object's value. In that instance, the evidence is a signal $\pi : \mathcal{X} \times \mathcal{Y} \to \Delta\mathcal{M}$, which is a measurable function mapping the object's value and profitability to a distribution of messages in a measurable set of possible messages, $\mathcal{M}$. All the results hold under the additional assumption that knowing the profitability of the object conveys no additional information about the object's value; that is, $\mathbb{E}(x|m, y) = \mathbb{E}(x|m) = \hat{x}(m)$. In other words, if an agent observes message m from signal π, they interpret it as "the object's expected value is $\hat{x}(m)$," independently of their belief about the object's profitability.

[7] Or more explicitly, a disclosure policy is a signal $D : \mathcal{X} \times \mathcal{Y} \to \Delta(\mathcal{X} \cup \{ND\})$ (where ND denotes no disclosure), with the restriction to simple evidence disclosure, so that $supp(D(x, y)) \in \{x, ND\}$.

If a piece of evidence indicating the object's value x is disclosed, the buyer's posterior mean after observing it is, by definition, exactly x. If otherwise the evidence is not disclosed, the buyer's posterior mean is computed using Bayes' Rule, accounting for the disclosure strategy. Formally, if $\int_{\mathcal{Y}} \int_{\mathcal{X}} [1 - d(x, y)] \, dF_{X|y}(x) dF_Y(y) > 0$,[8]

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

which is the expected value conditional on non-disclosure. The average object profitability given that a realization is not disclosed is analogously given by

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

2.1 Micro-Foundations for the "Demand Function" p

Throughout this paper, I regard the buyer as a passive agent, a "receiver" who sees (or does not see) information about the object, forms a belief about the object's expected value, and buys it or not according to some exogenously given "demand function" p. A possible micro-foundation for this demand function, which follows the description in Rayo and Segal (2010), is given below.

Privately-known outside option. A risk-neutral buyer chooses between acquiring the object being sold by the advisor or taking an outside option. An example of an outside option would be of buying another object somewhere else, or refraining from buying altogether. The value of the buyer's outside option x_o is private information, distributed according to F_o. Once the buyer sees all the provided information about the object's value, and forms belief $\hat{x}$, they purchase the object if $\hat{x} > x_o$ and do not purchase it otherwise. From the advisor's perspective, a purchase then happens with probability $F_o(\hat{x})$, the probability that the expected value of the object is greater than that of the outside option. In this case, the "demand function" p coincides with the distribution of outside option values F_o.[9] The assumption that p is increasing and continuously differentiable requires then that the cdf F_o be continuously differentiable.

Competing sellers. A different micro-foundation, in the spirit of Hwang, Kim, and Boleslavsky (2023), describes a buyer who has access to the object being offered by a seller and to various potential outside options, perhaps referring to objects being sold by alternative sellers. The value of the product sold by each of the outside sellers is unknown to the inside seller, but each is distributed according to F_o, and value draws are independent across inside and each of the outside objects. Therefore, if there are n outside sellers, the best outside option available to the buyer, $\max(x_o)$ is distributed according to

[8] Otherwise, non-disclosure is "off-path," and we fix x^{ND} at some value in $\mathcal{X}$.

[9] Alternatively, one may think that the advisor communicates with a population of possible buyers, and each buyer decides between making a purchase from the advisor and taking their own private outside option. In that case, F_o is the distribution of outside options in that population of buyers, and $p = F_o$ describes the amount of sales to be made by the advisor when they induce a certain posterior belief about the object's value on the population of buyers.

$\max(x_o) \sim (F_o)^n$. The buyer purchases the inside object then if the expected value of the inside object, $\hat{x}$, is greater than $\max(x_o)$. This happens with probability $(F_o)^n(\hat{x})$. In this case, the "demand function" p therefore coincides with $(F_o)^n$. Again, the assumption that p is increasing and continuously differentiable requires then that the cdf F_o be continuously differentiable.

Note moreover, that an *increase in competitiveness*, in the sense of an increase in the number of competitors n, maps into an increase in the convexity of the demand function. Formally, if $n' > n$, then demand function $p' = (F_o)^{n'}$ is more convex than demand function $p = (F_o)^n$, because the former function is a strictly increasing and convex transformation of the latter. And moreover, if additionally F_o has a derivative bounded away from 0 and n is large enough, then the demand function $p = (F_o)^n$ is a convex function in the entirety of its support.

The results stated in the paper show that the characterization of optimal disclosure rules, and policy implications in terms of transparency mandates, depend on the curvature of the demand function. Once these results are stated, I refer back to this interpretation of the convexity of the demand function as related to the competitiveness in the market to which the advisor belongs.

3 Optimal Disclosure

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

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

To understand (2), first remember that $F_{X|y}$ is the distribution of evidence realizations given that the object has profitability y. Suppose a realization x is disclosed, which happens with probability $d(x,y)$. 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 $1 - d(x,y)$ 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 advisor's expected payoff from committing to disclosure rule d is then

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

In (3), I split the advisor's payoff into two terms, expressing that the advisor's objective can be seen as twofold. According to the first term, the advisor 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. This covariance term reflects the advisor's desire to steer the buyer from purchasing low-profitability objects to purchasing those with high profitability. These two objectives are sometimes at odds, and the characterization provided below illustrates how optimal disclosure balances the two goals.

The first result, Theorem 1, provides a of disclosure rules that maximize the advisor's value, showing

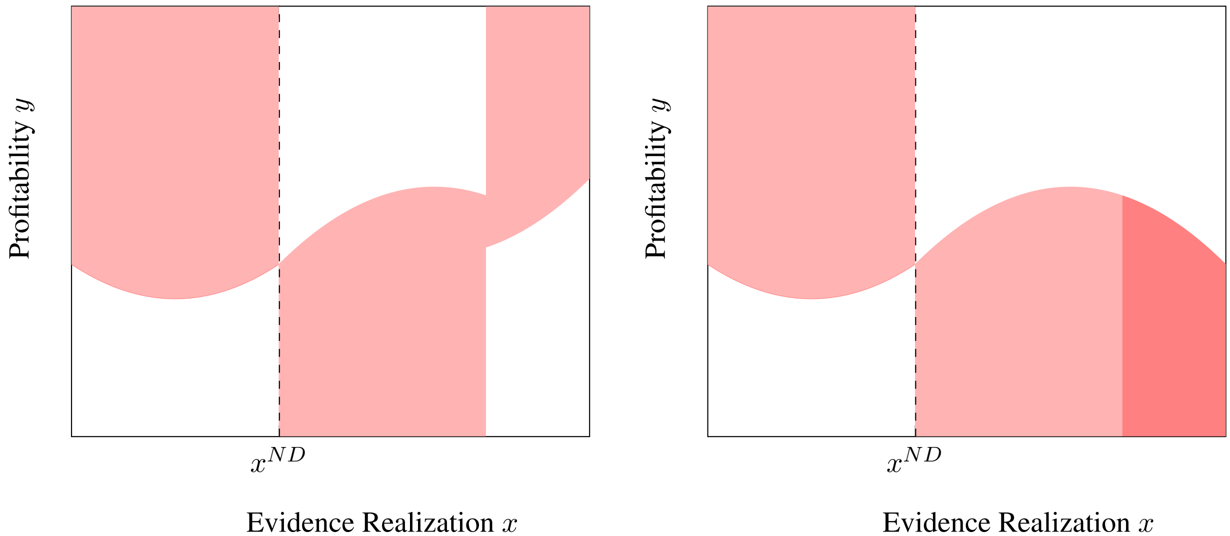

Figure 1: Building disclosure rule $\hat{d}$ from disclosure rule d: in each panel, the colored areas representing zones of no disclosure and the white areas representing zones of disclosure. The left-hand side panel illustrates a possible disclosure rule d that does not have a "threshold structure." The right-hand side panel depicts a disclosure rule $\hat{d}$, which has a threshold structure and is an improvement over d, derived from d according to (6), (7), and (8).

that they feature *aligned disclosure*. There is a threshold value $\bar{x}$ such that evidence is classified as either good news, if x is larger than $\bar{x}$, or bad news, if $x < \bar{x}$. For a good news realization $x > \bar{x}$, there is a profitability threshold $\bar{y}(x)$ such that the evidence is disclosed if and only if the object's profitability is above that threshold. Conversely, each bad news realization $x < \bar{x}$ is disclosed if and only if the object's profitability is below the threshold $\bar{y}(x)$.

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

$$d^*(x, y) = 1 \Leftrightarrow (x - \bar{x})(y - \bar{y}(x)) \geqslant 0. \tag{4}$$

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

$$\bar{y}(x) = y^{ND} \left[\frac{p'(x^{ND})(x^{ND} - x)}{p(x^{ND}) - p(x)} \right]. \tag{5}$$

A complete proof of Theorem 1, including the characterization of the threshold profitability function $\bar{y}(\cdot)$ is provided in the Appendix. For intuition, let's see that any disclosure rule which leads to no disclosure with positive probability and does not satisfy the threshold structure described in (4) can be improved upon by a rule that does satisfy (4). Start with one such disclosure rule d according to which no disclosure happens with positive probability, and which does not satisfy the threshold structure. (For example, the disclosure rule depicted in the left-hand side panel of Figure 1.) Let x^{ND} be its implied non-disclosure posterior mean. Define then an alternative rule, $\hat{d}$, that discloses each realization x with the same probability as d, but has a threshold structure. Taking the case in Figure 1, this means maintaining the "size" of the pink no-disclosure segment for each value x, but adjusting the profitability levels for which the realization is disclosed or not disclosed — d is plotted on the left-hand panel and the "improvement" $\hat{d}$ is in the right-hand panel. For $x \leqslant x^{ND}$, the no disclosure region is not adjusted, it remains corresponding to high-profitability realizations. As for $x > x^{ND}$, under d, some realizations with high profitability are not disclosed; instead, under $\hat{d}$, no disclosure is assigned to realizations with low profitability. Formally, if $x \leqslant x^{ND}$, let

$$\hat{d}(x,y) = \begin{cases} 1, \text{ if } y \leqslant \hat{y}(x) \\ 0, \text{ if } y > \hat{y}(x) \end{cases} \tag{6}$$

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

$$\hat{d}(x,y) = \begin{cases} 0, \text{ if } y < \hat{y}(x) \\ 1, \text{ if } y \geqslant \hat{y}(x), \end{cases} \tag{7}$$

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

$$\int_{\mathcal{Y}} \hat{d}(x,y) dF_{Y|x}(y) = \int_{\mathcal{Y}} d(x,y) dF_{Y|x}(y), \tag{8}$$

where $F_{Y|x}$ is the profitability distribution conditional on value x. Condition (8) implies that d and $\hat{d}$ induce the same x^{ND}, and so $\hat{d}$ satisfies (4), with $\bar{x} = x^{ND}$.

By moving from d to $\hat{d}$, the advisor 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 the Appendix), that: *Claim 1.* d and $\hat{d}$ produce the same overall probability of sale, because the distribution of posterior mean values is unchanged; and *Claim 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 advisor payoff than d, as desired.

The value and profitability thresholds described in Theorem 1 partition the value-profitability space into four "quadrants" — see, for example, Figures 2 and 3, which illustrate the quadrants defined by the optimal disclosure rules. The first and third quadrants represent regions where the advisor and buyer have aligned interests, either because both value and profitability are high or because both value and profitability are low. Evidence realizations in these "alignment" regions are optimally disclosed to the buyer. Conversely, the second and fourth quadrants represent areas of misalignment between advisor and advisee, and therefore these realizations are optimally concealed.

Aligned disclosure rules still encompass the possibility of full disclosure — if $x^{ND} = inf(\mathcal{X})$ and

$y^{ND} = inf(\mathcal{Y})$ or $x^{ND} = sup(\mathcal{X})$ and $y^{ND} = sup(\mathcal{Y})$ — or no disclosure — if $x^{ND} = inf(\mathcal{X})$ and $y^{ND} = sup(\mathcal{Y})$ or $x^{ND} = sup(\mathcal{X})$ and $y^{ND} = inf(\mathcal{Y})$. Proposition 6, stated in Appendix B, proposes three conditions on primitives that guarantee "interior solutions," ensuring that both disclosure and non-disclosure happen with positive probability in the optimal disclosure rule.

Steering and Credibility of Optimal Disclosure. The characterization of optimal disclosure rules highlights the steering logic behind optimal advice in for advisors with hidden motives. By committing to a threshold disclosure rule, the advisor can induce an ambiguous meaning to "no disclosure:" the advisee cannot fully disentangle whether "no news" means "bad news about the object's value (and that the object's sale is profitable), or whether "no news" means "the object's sale is not profitable to the advisor" (and good news about the object's value). By creating such ambiguity, the advisor can profitably steer the advisee's purchases from low- to high-profitability objects.

In this model, steering is made possible because the advisor is able to commit to a disclosure rule, and therefore influence the buyer's interpretation of "no disclosure." But is it reasonable to expect such commitment power from the advisor? A recent paper by Lin and Liu (2023) proposes a 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. Their reasoning is that if a sender "deviates" from an information policy in a way that keeps the marginal probability of sending each message unchanged, then that deviation would be "undetectable"; and if a policy is such that there are "undetectable" deviations that would benefit the sender, then such policy is not credible. According to this definition, the optimal disclosure rules described in Theorem 1 are credible.[10]

3.1 Linear Demand

Sections 3.1 and 3.2 now provide further characterization of optimal disclosure under different assumptions about the buyer's "demand function" p. Corollary 1 applies Theorem 1 when p is affine.

Corollary 1. *If the demand function p is affine, then the thresholds defining an optimal disclosure rule satisfy $\bar{x} = x^{ND} \in int\,(\mathcal{X})$ and $\bar{y}(x) = y^{ND} \in int\,(\mathcal{Y})$ for all $x \in \mathcal{X}$.*

When the demand function p faced by the advisor is affine, all disclosure rules yield the same overall probability of sale. That is, for any two disclosure rules d and d', $\mathbb{E}\left[P(y,d)\right] = \mathbb{E}\left[P(y,d')\right]$. This fact is an implication of the martingale property of posterior beliefs: we know that the expected posterior belief of the buyer about the value of the object has to equal the buyer's prior belief about the object's value. Because the probability of sale is an affine function of such posterior belief, the expected probability of sale has to equal the "ex-ante probability fo sale."

[10]To see, note that any disclosure rule other than the optimal rule as given in Theorem 1, but which induces that same marginal distribution over messages, must invariably swap the disclosure of a realization (x, y) and a different realization (x, y'); that is, two states with the same evidence realization x. But as per the argument in the proof of Theorem 1, if one such "undetectable swap" were available and were indeed profitable, then the starting disclosure rule must not have been optimal in the first place.

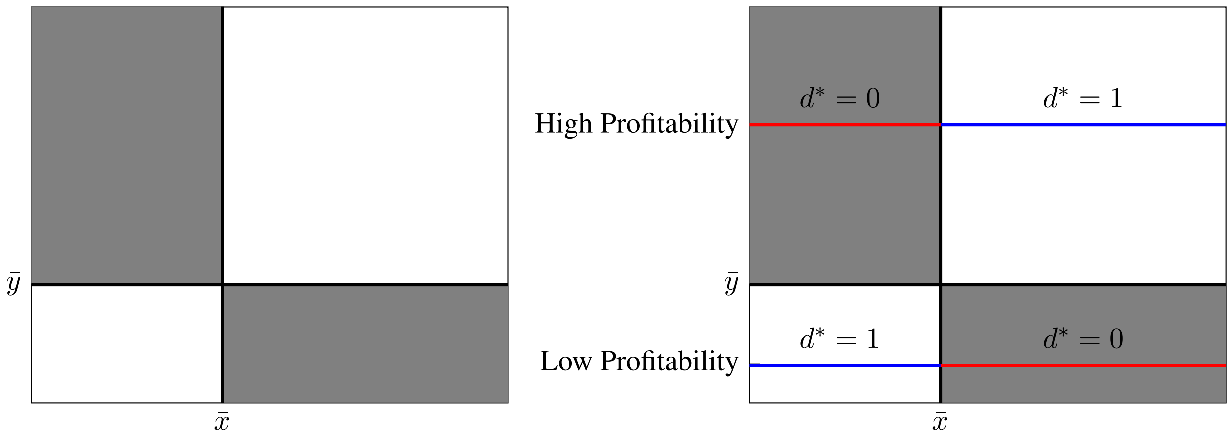

Figure 2: Left-hand panel: optimal disclosure rule when p is **affine**. The gray areas represent evidence that is optimally concealed by the advisor, and white areas are optimally disclosed. The right-hand panel illustrates that, in this case, the object is classified as either a high-profitability object or a low-profitability object. For the former, values are disclosed if and only if they are "good news," whereas for the latter values are disclosed if and only if they represent "bad news."

Consequently, in choosing a disclosure policy, there is no scope for the advisor to increase or decrease the overall probability that the buyer purchases the object. Rather, the advisor 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 advisor uses a constant threshold $\bar{y}(x) = y^{ND}$, thereby effectively assigning the object to one of two classes high profitability, with $y > y^{ND}$, and low profitability, with $y < y^{ND}$. Such an optimal disclosure rule is depicted in Figure 2; its right-hand panel illustrates that, for high-profitability objects, the object's value is disclosed if and only if it is "good news," while the opposite is true for low-profitability objects.

3.2 Nonlinear Demand

If the demand function p is not affine, then the amount of information disclosed by the advisor impacts the overall probability of sale. Observation 1 below states that increasing the amount of information about the object's value that is disclosed to the buyer increases the overall probability that the object is sold if the demand function is convex, and decreases it if the demand function is concave. To formally state this result, we say that a disclosure rule d has *more disclosure* than a disclosure rule d' if $d(x, y) \geqslant d'(x, y)$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$; strictly so if this inequality is strict for a subset of $\mathcal{X} \times \mathcal{Y}$ with positive measure.

Observation 1. *Suppose d has more disclosure than d'. Then if p is strictly convex (concave), d yields a higher (lower) sale probability than d'; and strictly so if d has strictly more disclosure than d'.*

With a nonlinear demand function, the advisor transfers sales from low- to high-profitability objects *at the expense of* the total probability that the object is sold. If the advisor faces a convex demand function, he wishes to disclose more information in order to maximize the sale probability; but maximizing

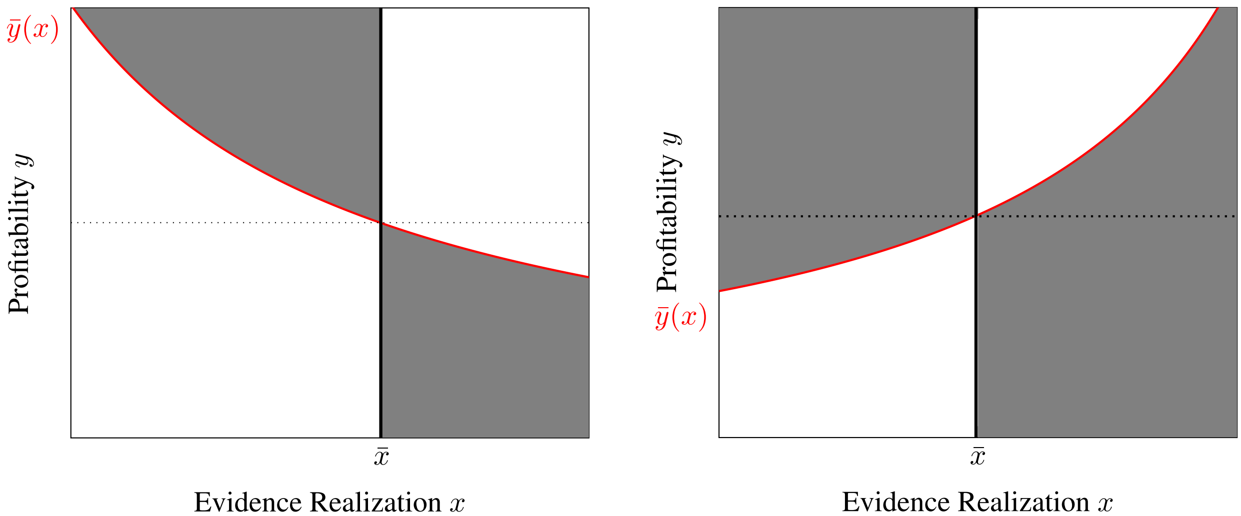

Figure 3: Optimal disclosure rule when the buyer's demand function is **convex** (left panel) and **concave** (right panel). These images correspond to the optimal disclosure rules corresponding to the examples in sections 3.2.1 and 3.2.2, respectively. The gray areas represent evidence realizations that are optimally concealed by the advisor, and the white areas are optimally disclosed.

the covariance between profitability and probability of sale requires the concealment of some realizations. Conversely, if the demand function is concave, overall sale probability increases when information is concealed from the buyer, but to improve the covariance between sales and profitability, the advisor must disclose some value realizations. The following corollary describes features of optimal disclosure in these cases.

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

For an illustration, see Figure 3, which depicts optimal disclosure rules in examples 3.2.1 and 3.2.2 below. In each panel, the x-axis represents the expected value of the object induced by a realization of the evidence, and the y-axis represents the object's profitability. The threshold $\bar{x}$ divides good-news and bad-news evidence realizations; and the function $\bar{y}(x)$ is the optimal profitability threshold dividing disclosure and non-disclosure regions. The left panel of Figure 3 illustrates optimal disclosure when the buyer's demand for the object is convex. In that case, the advisor can maximize the overall probability of sale by disclosing all evidence realizations; but optimally chooses to conceal some realizations in order to steer the buyer from low- to high-profitability objects. As a consequence, extreme evidence realizations (very bad news or very good news) are always disclosed, but some bad news are concealed when the object's sale is very profitable, as are some good news if the object is less profitable. This description is summarized by the optimal profitability threshold being a *decreasing* function of the value of the evidence realization. Conversely, if the buyer's demand for the object is concave, then the optimal profitability threshold is an *increasing* function — as depicted in the right panel of Figure 3. When demand is concave, this shape is a result of balancing the advisor's desire to maximize overall sales by concealing all realizations; and to steer sales towards high-profitability objects by selectively disclosing some evidence.

3.2.1 Example: Optimal Disclosure under Convex Demand Function

An object's value and its profitability are independently distributed, each distributed according to the uniform distribution $U[0,1]$, so that $F_Y = U[0,1]$ and $F_{X|y} = U[0,1]$ for each $y \in [0,1]$. Further, let the demand function be given by $p(x) = x^2$.

From Theorem 1, we know that the thresholds defining the optimal disclosure rule should satisfy:

$$\bar{x} = x^{ND}, \text{ and}$$

$$\bar{y}(x) = y^{ND}\left[\frac{2x^{ND}(x^{ND}-x)}{(x^{ND})^2 - x^2}\right] = y^{ND}\left[\frac{2x^{ND}}{x^{ND}+x}\right].$$

Moreover, given these thresholds, x^{ND} and y^{ND} must indeed correspond to the Bayesian posteriors implied by no disclosure: $x^{ND} = \mathbb{E}(x|\text{no disclosure})$, and $y^{ND} = \mathbb{E}(y|\text{no disclosure})$. Numerically, I find that there is a unique pair $(x^{ND}, y^{ND}) \in (0,1)^2$ that satisfies these conditions. The pair is given by

$$x^{ND} = 0.5824 \text{ and } y^{ND} = 0.5098.$$

Because there is a unique such pair, we know that it must correspond to the thresholds in the optimal disclosure rule. The expected payoff to the advisor under full disclosure equals $1/6$. The expected payoff under no disclosure is $1/8$. And the expected payoff under the optimal disclosure rule just described equals 0.1885.

3.2.2 Example: Optimal Disclosure under Concave Demand Function

Again, the object's value and its profitability are independently distributed, each distributed according to the uniform distribution $U[0,1]$, so that $F_Y = U[0,1]$ and $F_{X|y} = U[0,1]$ for each $y \in [0,1]$. But now consider the following concave demand function: $p(x) = 2x - x^2$.

In this case, the thresholds defining the optimal disclosure rule should satisfy:

$$\bar{x} = x^{ND}, \text{ and}$$

$$\bar{y}(x) = y^{ND}\left[\frac{(2-2x^{ND})(x^{ND}-x)}{2x^{ND} - (x^{ND})^2 - (2x - x^2)}\right] = y^{ND}\left[\frac{2(1-x^{ND})}{2 - x^{ND} - x}\right].$$

Again, it must be that $x^{ND} = \mathbb{E}(x|\text{no disclosure})$, and $y^{ND} = \mathbb{E}(y|\text{no disclosure})$. Numerically, I find that the is a unique pair $(x^{ND}, y^{ND}) \in (0,1)^2$ satisfying these conditions, and therefore corresponding to the thresholds in the optimal disclosure rule, is

$$x^{ND} = 0.5397 \text{ and } y^{ND} = 0.5402.$$

In this example, the expected payoff to the advisor under full disclosure equals $1/3$. The expected payoff under no disclosure is $3/8$. And the expected payoff under the optimal disclosure rule just described equals 0.3907.

4 Transparent Motives Benchmark

Consider a policy intervention that makes the motives of the advisor transparent to the buyer. Such transparency policies are common in advice markets.[11]

In the model, the transparency policy imposes that when the buyer observes the "advice" — either the disclosed evidence or non-disclosure — they also perfectly observe how profitable the object's sale is to the advisor. Naturally, anticipating that the buyer will see the profitability of the object, the advisor optimally uses a disclosure strategy that differs from the payoff maximizing disclosure rule without transparency. Indeed, with transparency, the advisor's problem becomes a continuum of separate disclosure problems, one for each profitability. This fact is stated in Theorem 2, along with a characterization of optimal disclosure under mandated transparency.

Theorem 2. *With mandated transparency, the advisor chooses for each $y \in \mathcal{Y}$ a policy $d(\cdot, y) : \mathcal{X} \to [0,1]$ to maximize $P(y, d(\cdot, y))$, where*

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

$$\text{and } x_y^{ND} = \frac{\int_{\mathcal{X}} [1 - d(x,y)] x dF_{X|y}(x)}{\int_{\mathcal{X}} [1 - d(x,y)] dF_{X|y}(x)}. \tag{9}$$

Under mandated transparency, any optimal disclosure rule d^* almost everywhere satisfies

$$p(x) > p\left(x_y^{ND}\right) + p'\left(x_y^{ND}\right)\left(x - x_y^{ND}\right) \Rightarrow d^*(x,y) = 1, \tag{10}$$

$$\text{and } p(x) < p\left(x_y^{ND}\right) + p'\left(x_y^{ND}\right)\left(x - x_y^{ND}\right) \Rightarrow d^*(x,y) = 0. \tag{11}$$

Under transparency, the advisor cannot use strategic disclosure in order to steer sales across objects with different profitability levels. Rather, Theorem 2 shows that they separately maximize the probability that the object with each profitability level gets sold. The optimal disclosure rule may differ across objects with different profitabilities because $F_{X|y}$ may depend on y, which therefore implies the optimal value of x_y^{ND} — which determines the optimal disclosure rule, as per (10) and (11) — varies with y.

[11]For example, financial advisors are often required to disclose professional affiliations with securities issuers, sources of compensation received beyond service fees, and other potential or existing conflicts of interest. As for social media advertising, the FTC and the UK CMA mandate digital influencers to mark their posts as "paid content" when they are sponsored by brands, and take other steps to clarify their relation with product producers to their content consumers.

4.1 Curvature of p and the Effectiveness of Mandated Transparency

Proposition 1 below is a direct consequence of Theorem 2, and shows that mandating transparency of the advisor's motives may increase or decrease the amount of evidence the advisor voluntarily discloses to the buyer, and that this effect depends on the curvature of the buyer's demand function p.[12]

Proposition 1. *Let d and d' be optimal disclosure policies to the advisor under mandated transparency and hidden motives, respectively.*

1. *If $p(\cdot)$ is strictly* ***convex****, mandated transparency improves evidence disclosure, that is, d has* ***more disclosure*** *than d'.*

2. *If $p(\cdot)$ is strictly* ***concave****, mandated transparency harms evidence disclosure, that is, d has* ***less disclosure*** *than d'.*

If the demand function is convex, then the mandated transparency policy increases the set of evidence pieces that the advisor would voluntarily choose to disclose to the buyer. In fact, given that demand regime, any optimal disclosure policy under mandated transparency involves the voluntary disclosure of *all evidence* – so that the optimal d involves full disclosure, regardless of the object's profitability. To see this, note that for any value of x_y^{ND}, strict convexity of p implies that the condition on the right-hand side of (10) is satisfied strictly for any $x \neq x_y^{ND}$. An optimal disclosure rule must therefore satisfy $d^*(x, y) = 1$ almost everywhere.

Conversely, if the demand function is concave, then mandated transparency has the opposite effect on evidence disclosure: the optimal disclosure policy is for the advisor to not disclose any of the realized evidence. Again, this can be seen directly from (10). Because p is strictly concave, then for any x_y^{ND}, the condition on the right-hand side of (10) fails for all $x \neq x_y^{ND}$, which implies that almost all evidence is optimally concealed from the buyer.

Does mandated transparency help or hurt the buyer? To answer this question, I consider how transparency affects the Blackwell informativeness of the advice. (A virtue of this approach is that it is agnostic about the buyer's objective, so long as the buyer's posterior mean is a sufficient statistic for his welfare.) Denote the transparency policy by $\tau \in \{0, 1\}$, with $\tau = 1$ indicating the mandated transparency environment, and $\tau = 0$ the hidden motives environment. Given a disclosure rule d and transparency policy τ, let $F^B(\cdot, d, \tau)$ be the distribution of posterior means observed by the buyer.[13] We say the pair (d, τ) is *more informative than* the pair (d', τ') if $F^B(\cdot, d, \tau)$ is more informative than $F^B(\cdot, d', \tau')$ in the Blackwell order.

[12] If the demand function is affine, a case that is not contemplated in Proposition 1, then under mandated transparency, the advisor is indifferent between all disclosure rules, including full disclosure and no disclosure. The effect of mandated transparency on disclosure is therefore indeterminate.

[13] For a given d and τ, $F^B(x, d, \tau)$ is equal to

$$\int_{\mathcal{Y}} \int_{[x_{min}, x)} d(\hat{x}, y) dF_{X|y}(\hat{x}) dF_Y(y) + \int_{\mathcal{Y}} \int_{\mathcal{X}} (1 - d(\hat{x}, y)) \mathbb{1}\{x_y^{ND} \leqslant x\} dF_{X|y}(\hat{x}) dF_Y(y),$$

where for every $y \in \mathcal{Y}$, $x_y^{ND} = x^{ND}$ as given in (1) if $\tau = 0$; and x_y^{ND} is as defined in (9) if $\tau = 1$.

Corollary 3.

1. *If $p(\cdot)$ is strictly **convex**, the advisor is more informative under mandated transparency.*

2. *If $p(\cdot)$ is strictly **concave**, and $\mathbb{E}[x|y] = \mathbb{E}[x]$ for every $y \in \mathcal{Y}$, then the advisor is less informative under mandated transparency.*

Proposition 1 and its Corollary 3 show that the effectiveness of mandated transparency as a regulatory policy depends on the curvature of the demand function p. Returning to the micro-foundations provided in section 2.1, one possible interpretation of the curvature of the demand function is as a measure of the competitiveness in the market to which the advisor belongs. Specifically, if there are sufficiently many "competitors" to the considered advisor, then the demand function p is convex, and mandated transparency is an effective policy. Conversely, if competition is not sufficient, then p may be concave, implying that mandated transparency is not a good policy tool. An interpretation is that competition with outside sellers is a sufficient motivator for the seller to provide information to the buyer — this point is indeed made by Hwang, Kim and Boleslavsky (2023). But, in case there is lack of competition, allowing sellers to profitably steer buyers (by keeping their motives hidden) can provide the necessary incentives for the disclosure of information about the object's value.

4.2 Local Effects of Mandated Transparency

The results in section 4.1 concern extreme cases where p is strictly concave or strictly convex. "Local versions" of those results hold, which depend only on the local curvature of the demand function around the expected value x^{ND}. From (9), we know that under mandated transparency, a value x for an object with profitability y is disclosed if

$$p(x) > p\left(x_y^{ND}\right) + p'\left(x_y^{ND}\right)\left(x - x_y^{ND}\right), \tag{12}$$

and x is not disclosed if the opposite inequality holds. Without mandated transparency, we know from (a rewriting of) the characterization in Theorem 1 that a value x for an object with profitability y is disclosed if

$$p(x) > p\left(x^{ND}\right) + \frac{y^{ND}}{y} p'\left(x^{ND}\right)\left(x - x^{ND}\right), \tag{13}$$

where remember that y^{ND} is the expected profitability conditional on no disclosure under the optimal disclosure rule. Again, the value x is not disclosed if the opposite inequality holds.

Suppose optimal disclosure rules with and without mandated transparency are such that, for some profitability level y, we have $x^{ND} = x_y^{ND}$. And further suppose that p is locally strictly convex around x^{ND}. In that case, condition (12) must hold for realizations x sufficiently close to x^{ND}, meaning that such "local realizations" are disclosed to the buyer. In comparison, consider the optimal policy when the advisor has hidden motives. For high-profitability objects ($y > y^{ND}$) and "local bad news realizations" ($x \uparrow x^{ND}$), the opposite inequality to (13) must hold; so that such local bad news are concealed.

Analogously, for low-profitability objects ($y < y^{ND}$), "local good news realizations" ($x \downarrow x^{ND}$) are concealed from the buyer. This argument implies that, if p is locally convex around x^{ND}, then mandated transparency implies a local increase in disclosure. This result is formally stated below in Proposition 2, along with an opposite result for the case when p is locally concave.

Proposition 2. *Suppose d and d_m are optimal disclosure rules with hidden motives and under mandated transparency, respectively. And suppose, for all $y \in \mathcal{Y}$, it holds that $x^{ND}(d) = x_y^{ND}(d_m) =: \tilde{x} \in int(\mathcal{X})$.*

1. *If p is locally strictly convex around $\tilde{x}$, there exist x' and x'', with $x' < \tilde{x} < x''$ such that $d(x, y) \leqslant d_m(x, y)$ for almost all (x, y) with $x \in (x', x'')$; and strictly so for an open subset of such (x, y).*
2. *If p is locally strictly concave around $\tilde{x}$, there exist x' and x'', with $x' < \tilde{x} < x''$ such that $d(x, y) \geqslant d_m(x, y)$ for almost all (x, y) with $x \in (x', x'')$; and strictly so for an open subset of such (x, y).*

Also from conditions (12) and (13), we can see that, compared to the benchmark with hidden motives, mandated transparency implies a weak increase in the disclosure of bad news, and a weak decrease in the disclosure of good news, about objects with high profitability. To see, if the object's profitability is high ($y > y^{ND}$), and a realization is "bad news" ($x < x^{ND}$), then condition (13) holding implies that condition (12) holds as well. This means that if such bad news are disclosed under hidden motives, then they are disclosed as well under mandated transparency. Conversely, for "good news" ($x > x^{ND}$), condition (12) implies condition (13), so that their disclosure under mandated transparency implies that they are also disclosed under hidden motives. Proposition 3 states this result formally, as well as an analogous result for low profitability objects ($y < y^{ND}$).

Proposition 3. *Suppose d and d_m are optimal disclosure rules with hidden motives and under mandated transparency, respectively. And suppose, for all $y \in \mathcal{Y}$, it holds that $x^{ND}(d) = x_y^{ND}(d_m) =: \tilde{x}$; and let $\tilde{y} := y^{ND}(d)$. The following statement holds for almost all $(x, y) \in \mathcal{X} \times \mathcal{Y}$:*

$$d(x, y) \geqslant d_m(x, y) \Leftrightarrow (x - \tilde{x})(y - \tilde{y}) \geqslant 0.$$

Both Propositions 2 and 3 start from the assumption that optimal disclosure rules with and without mandated transparency lead to the same beliefs of no disclosure $x^{ND} = x_y^{ND}$ (for all profitability levels $y \in \mathcal{Y}$). If the value distribution $F_{X|y}$ is independent from y, then we know from Theorem 2 that in the optimal disclosure rule under mandated disclosure, we have that x_y^{ND} is also independent of y. Moreover, as the distribution of profitabilities becomes more centered around its expected value, it must be that the value x^{ND} in the optimal rule under hidden motives approaches such value x_y^{ND}. Indeed, in the limit as the distribution F_Y becomes the degenerate distribution, it must be that these two values coincide.

5 Alternative Communication Protocols

The model makes two main assumptions regarding the communication protocol. First, the advisor can *commit* to a disclosure rule, prior to the object's profitability or the evidence about its value being drawn. Second, the advisor can use only disclosure policies, either revealing or not revealing a realized piece of information about the object's value, rather than committing to more general signal structures. In this section, I consider a variations of the model, which drops the commitment from the communication protocol. Under that new protocol, I investigate whether mandating transparency about the advisor's motives improves the informativeness of their advice to the buyer. Next, I briefly comment on a variation of the model in which the advisor can commit to more general signal structures.

5.1 No Commitment Disclosure Protocol

A direct reading of the results in section 4 is that they delineate conditions under which a policy maker should or should not institute a transparency policy: such decision should be made based on the curvature of the buyer's demand for the object. But beyond this normative implication, the results also refine our theoretical understanding of regulation in advice markets. Specifically, I highlight that in information design models *with commitment*, the 'alignment between sender and receiver preferences' and the 'opaqueness of the sender's motives' are distinct objects; and it is not necessarily true that regulations that reduce the latter would also make the sender's interests more aligned with those of the receiver. And even in contexts where transparency does improve the alignment between the advisor and advisee's interests — such as when the demand function is convex, in the model — it does so because, as a byproduct the transparency of their motives, the sender's effective objective function becomes "more convex," therefore inducing them to optimally disclose more evidence. This is in contrast with what happens in a disclosure environment without commitment — discussed below — in which the transparency of the advisor's motives induces the unravelling of uninformative equilibria.

To see this contrast, I introduce a version of the model where there is no commitment in the communication protocol. In this section, we allow the support of the object's profitability to include negative profitability values, so we may have $y_{min} < 0 < y_{max}$,[14] in which case the buyer is unsure whether the profitability is such that the advisor wishes to maximize the object's probability of sale or to minimize it. Consider the following disclosure protocol with no commitment: First, the advisor observes the object's profitability and value. After that, the advisor chooses whether to disclose the 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 advisor's equilibrium strategy. As before, the buyer does not observe the object's profitability. The equilibrium notion is Perfect Bayesian Equilibrium.

Proposition 4. *For any demand function p, an equilibrium exists, and any equilibrium disclosure strategy*

[14] Note that the characterization of optimal disclosure rules under the commitment protocol, given in Theorem 1, also applies when $y_{min} < 0$.

d^* has a threshold structure: $d^*(x, y) \in \{0, 1\}$ and[15]

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

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

Suppose the object's profitability is always positive ($y_{min} \geqslant 0$), 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 4 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 advisor draws a 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 ($y_{max} \leqslant 0$). 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 advisor discloses (only and) all "good news" about the object ($x \geqslant \bar{x}$) when profitability is positive, and (only and) all "bad news" ($x \leqslant \bar{x}$) about the object when profitability is negative. Partially uninformative equilibria do not unravel, because, upon observing non-disclosure, the buyer does not know whether the advisor has "good news" but negative profitability or "bad news" but positive profitability.

Note also that equilibrium disclosure strategies as described in the proposition are independent of the shape of the demand function p – in contrast to the model with commitment, where the shape of optimal disclosure rules depends on the curvature of p. This happens because the advisor makes their disclosure decision only after seeing the value; and at that point their best response is guided solely from comparing the given realization to the buyer's belief of no disclosure. If the profitability of the object is positive, then, because the demand function is strictly increasing, the best response is to disclose good news (better than the belief of no disclosure), and conceal otherwise. If instead the profitability of the object is negative, the best response is to disclose bad news, and conceal otherwise. This implies that the "threshold profitability" is always 0, which divides positive profitability objects from negative profitability objects.

Now return to the question of whether mandated transparency incentivizes the advisor to disclose information about the object's value to the buyer. Proposition 5 shows that, in contrast with the benchmark with commitment, in this case mandated transparency induces the advisor to reveal all their evidence about the object's value; and this result holds independently of the curvature of the buyer's demand function.

Proposition 5. *Under a mandated transparency policy that reveals the object's profitability to the buyer, full disclosure is the unique equilibrium of the disclosure game with no commitment.*

[15] Assuming that, when indifferent, the advisor discloses the evidence.

5.2 Unconstrained Signaling Technology

The usual assumption in information design models is that a sender (the advisor, in this case) commits a signal a map from states of the world (the object's value and profitability) 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 simple disclosure rules, in which 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 value-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 advisor'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 advisor's incentives to relay information about the object's value to the buyer, depending on the curvature of the buyer's demand function.

References

Antic, N. and A. Chakraborty (2024) "Selected Facts," *working paper*.

Ball, I., and X. Gao. (2024) "Benefiting from Bias: Delegating to Encourage Information Acquisition." *Journal of Economic Theory*.

Che, Y.K., and N. Kartik (2009) "Opinions as Incentives." *Journal of Political Economy*, **117**, 815-860.

Crawford, V. and J. Sobel. (1982) "Strategic Information Transmission." *Econometrica*: 1431-1451.

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. (2024) "Constrained Information Design," *Mathematics of Operations Research* **49**: 78-106..

Grossman, S. (1981) "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law & Economics*, **24**: 461-484.

Hagenbach, J. and F. Koessler. (2017) "Simple versus Rich Language in Disclosure Games." *Review of Economic Design* **21**: 163-175.

Hwang, I., K. Kim, and R. Boleslavsky (2023). "Competitive Advertising and Pricing," *working paper*.

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.

Li, M. and K. Madarász. (2008) “When Mandatory Disclosure Hurts: Expert Advice and Conflicting Interests.” *Journal of Economic Theory* **139**: 47-74.

Libgober, J. (2022) ”False Positives and Transparency.” *American Economic Journal: Microeconomics* **14**: 478-505.

Lin, X. and C. Liu. (2023) “Credible Persuasion.” *Journal of Political Economy*, forthcoming.

Matthews, Steven, and Andrew Postlewaite. (1985) “Quality Testing and Disclosure.” *RAND Journal of Economics*: 328-340.

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

Mezzetti, C. (2025) “Manipulative Disclosure,” *RAND Journal of Economics*, forthcoming.

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

Morris, S. (2001) “Political Correctness.” *Journal of Political Economy* **109**: 231-265.

Onuchic, P., and D. Ray. (2023) “Conveying Value via Categories,” *Theoretical Economics*, **18**:1407-1439.

Rayo, L., and I. Segal (2010) “Optimal Information Disclosure.” *Journal of Political Economy*, **118**, 949–987.

Seidmann, D., and E. Winter. (1997) “Strategic Information Transmission with Verifiable Messages.” *Econometrica*: 163-169.

Shishkin, D. (2023) “Evidence Acquisition and Voluntary Disclosure.” *working paper*.

Skreta, V., and L. Veldkamp. (2009) “Ratings Shopping and Asset Complexity: A Theory of Ratings Inflation.” *Journal of Monetary Economics*, **56**, 678-695.

Sobel, J. (1985) “A Theory of Credibility.” *Review of Economic Studies* **52**: 557-573.

Szalay, D. (2005) “The Economics of Clear Advice and Extreme Options.” *The Review of Economic Studies*, **72**: 1173-1198.

A Proofs

Statements and Proofs of Claims 1 and 2

This section completes the proof in the main text, showing that any disclosure rule that does not satisfy the threshold structure described in (4) can be improved upon by a rule that does satisfy (4). To that end, consider d and $\hat{d}$ as in (6) and (7). I prove the following two claims used in the main text.

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

Proof.

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

where $F_{Y|x}$ is the profitability distribution conditional on value x and F_X is the marginal value distribution. 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 (8). ■

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

Proof.

$$\begin{aligned}
&\text{Cov}\left[y, P(y,\hat{d})\right] - \text{Cov}\left[y, P(y,\hat{d})\right] = \mathbb{E}\left[\left(P(y,\hat{d}) - P(y,d)\right)(y - \mathbb{E}(y))\right] \\
&= \int_{\mathcal{Y}} \int_{\mathcal{X}} \left[p(x) - p(x^{ND})\right] \left[\hat{d}(x,y) - d(x,y)\right] \left[y - \mathbb{E}(y)\right] dF_{X|y}(x) dF_Y(y) \\
&= \int_{\mathcal{X}} \left[p(x) - p(x^{ND})\right] \int_{\mathcal{Y}} \left[\hat{d}(x,y) - d(x,y)\right] \left[y - \mathbb{E}(y)\right] dF_{Y|x}(y) dF_X(x) \qquad (14)
\end{aligned}$$

By the definition of $\hat{d}$, for $x < x^{ND}$, $\hat{d}(x,y) - d(x,y) \geqslant 0$ when $y < \hat{y}(x)$ and $\hat{d}(x,y) - d(x,y) \leqslant 0$ when $y > \hat{y}(x)$. We thus have, for $x \leqslant x^{ND}$,

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

where the equality follows from the fact that the expected difference $\hat{d}(x,y) - d(x,y)$ is 0 (as given by (8)). Analogously, we can show that

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

when $x > x^{ND}$. Moreover, these inequalities are strict for a positive measure of realizations. These observations, along with the fact that p is strictly increasing, deliver that the expression in (14) is strictly positive. ■

Proof of Theorem 1

Step 1. Suppose $\hat{d}$ is a disclosure rule such that no disclosure happens with positive probability — that is, $\int_{\mathcal{Y}}\int_{\mathcal{X}}\left[1-\hat{d}(x,y)\right]dF_{X|y}(x)dF_Y(y)>0$ — and suppose $\hat{d}$ does not satisfy the characterization in Theorem 1. Then d can be strictly improved.

For any disclosure rule d such that no disclosure happens with positive probability, the advisor's value is given by

$$\int_y yP(y,d)dF_Y(y) = \int_y y\int_{\mathcal{X}}\left[d(x,y)p(x)+(1-d(x,y))p(x^{ND})\right]dF_{X|y}(x)dF_Y(y),$$

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

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

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

Now from the definition of x^{ND}, we get

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

Substituting this into the previous equation, we have

$$\frac{\partial\Pi}{\partial d(x,y)} = \left[y\left(p(x)-p(x^{ND})\right)-y^{ND}p'(x^{ND})(x-x^{ND})\right]dF_{X|y}(x)dF_Y(y), \tag{15}$$

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

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

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

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

Now take disclosure $\hat{d}$, which does not satisfy the characterization in Theorem 1. That is, $\hat{d}$ does not have a threshold structure as in (4) with $\bar{x} = x^{ND}$ and profitability threshold satisfying

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

Then it must be that either $\hat{d}(x,y) \neq 0$ when (15) is negative or $\hat{d}(x,y) \neq 1$ when (15) is positive. In each case, $\hat{d}$ can be strictly improved.

Step 2. Suppose instead that $\hat{d}$ is the "full disclosure" rule — that is, $d(x,y) = 1$ almost everywhere. Then there are two possibilities. First, it may be that full disclosure is a rule that satisfies (4), (5), and $\bar{x} = x^{ND}$ — for $x^{ND} = \inf(\mathcal{X})$ and $y^{ND} = \inf(\mathcal{Y})$ or $x^{ND} = \sup(\mathcal{X})$ and $y^{ND} = \sup(\mathcal{Y})$. In that case, full disclosure satisfies the necessary conditions for optimality, as given by the Theorem.

If instead full disclosure is not a rule that satisfies the necessary conditions given in the Theorem, it must be that the disclosure rule implied by (4), (5), and $\bar{x} = x^{ND}$, with $x^{ND} = \inf(\mathcal{X})$ and $y^{ND} = \inf(\mathcal{Y})$, is such that no disclosure happens with positive probability. For this to be the case, there must exist $x \in \mathcal{X}$ such that

$$\frac{p'\left(\inf(\mathcal{X})\right)\left(x - \inf(\mathcal{X})\right)}{p(x) - p\left(\inf(\mathcal{X})\right)} > 1.$$

Similarly, it must be that the disclosure rule implied by (4), (5), and $\bar{x} = x^{ND}$, with $x^{ND} = \sup(\mathcal{X})$ and $y^{ND} = \sup(\mathcal{Y})$, is such that no disclosure happens with positive probability. For this to be the case, there must exist $x' \in \mathcal{X}$ such that

$$\frac{p'\left(\sup(\mathcal{X})\right)\left(x' - \sup(\mathcal{X})\right)}{p(x') - p\left(\sup(\mathcal{X})\right)} < 1.$$

When such x and x' exist, we know from Proposition 6 and Lemma 2 that the optimal disclosure rule must be interior, and so that no disclosure happens with positive probability. Therefore, $\hat{d}$, the "full disclosure" rule, is not optimal. ■

Proof of Theorem 2

Part 1. Statement of the advisor's problem.

Suppose the advisor chooses disclosure rule d. Upon observing non-disclosure and profitability y

(under mandated transparency), the buyer's mean posterior is

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

And so the probability of sale of object y is

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

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 advisor's problem is separable across profitability levels. Therefore maximizing $\Pi(d)$ over $d : \mathcal{X} \times \mathcal{Y} \to [0,1]$ is equivalent to maximizing, for each $y \in \mathcal{Y}$, $yP(y, d(\cdot, y))$, over $d(\cdot, y) : \mathcal{X} \to [0,1]$.

Part 2. Characterization of optimal disclosure rule.

We proceed similarly to the proof of Theorem 1. Fixing some $y \in \mathcal{Y}$ and $x \in \mathcal{X}$, we can take a derivative of the sender's value with respect to $d(x,y)$, to get

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

Now from the definition of x_y^{ND}, we get

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

Substituting this into the previous equation, we have

$$\frac{\partial \Pi}{\partial d(x,y)} = y \left[\left(p(x) - p(x_y^{ND}) \right) - p'(x_y^{ND})(x - x_y^{ND}) \right] dF_{X|y}(x) dF_Y(y). \tag{16}$$

If $\hat{d}$ is a disclosure rule that induces no disclosure with positive probability and does not satisfy (10) and (11) as given in the Theorem, then there is a positive measure of (x,y) such that (16) is strictly positive but $\hat{d}(x,y) < 1$ or such that (16) is strictly negative but $\hat{d}(x,y) > 0$. This means that $\hat{d}$ can be strictly improved, and cannot be a solution to the advisor's problem.

Now we must consider the possibility that the solution is "full disclosure." For each $\hat{y} \in \mathcal{Y}$, and for each $\hat{x} \in \mathcal{X}$, let $d_{\hat{x}}(\cdot, \hat{y})$ be defined by

$$p(x) \geqslant p(\hat{x}) + p'(\hat{x})(x - \hat{x}) \Rightarrow d_{\hat{x}}(x, \hat{y}) = 1,$$

$$\text{and } p(x) < p(\hat{x}) + p'(\hat{x})(x - \hat{x}) \Rightarrow d_{\hat{x}}(x, \hat{y}) = 0.$$

First suppose for every $\hat{x} \in \mathcal{X}$, $d_{\hat{x}}(\cdot, \hat{y})$ implies that an open interval of realizations in $\mathcal{X}$ are not disclosed. Then the map $\Gamma_{\hat{y}}(\hat{x}) = \mathbb{E}(x | \text{ no disclosure region implied by } d_{\hat{x}}(\cdot, \hat{y}))$ is a continuous self map in the compact convex set $\mathcal{X}$; and therefore has a fixed point. As in Lemma 1, we can show that the disclosure rule implied by this fixed point is a strict improvement over full disclosure. If instead there exists some $\hat{x} \in \mathcal{X}$ such that $d_{\hat{x}}(\cdot, \hat{y})$ implies almost all realizations are disclosed, then full disclosure is a disclosure rule that satisfies the necessary conditions given by in the Theorem. ■

Proof of Proposition 1

If p is strictly convex, then full disclosure is the only disclosure rule that satisfies the necessary conditions (10) and (11) in Theorem 2. It is therefore the optimal disclosure rule for the advisor under mandated transparency; which trivially has more disclosure than the optimal disclosure rule for the advisor with hidden motives.

If p is strictly concave, then no disclosure is the only disclosure rule that satisfies the necessary conditions (10) and (11) in Theorem 2. It is therefore the optimal disclosure rule for the advisor under mandated transparency; which trivially has less disclosure than the optimal disclosure rule for the advisor with hidden motives. ■

Proof of Proposition 2

Proof of Statement 1. If p is locally strictly convex around $\tilde{x}$, there exist $x', x'' \in \mathcal{X}$ with $x' < \tilde{x} < x''$ such that

$$p(x) > p(\tilde{x}) + p'(\tilde{x})(x - \tilde{x}).$$

From Theorem 2, we therefore know that $d_m(x, y) = 1$ for almost all (x, y) with $x \in [x', x'']$. Consequently, $d(x, y) \leqslant d_m(x, y)$ for almost all (x, y) with $x \in (x', x'')$. Further note that, for each $y > y^{ND}$ (the expected profitability of no disclosure implied by d), there exists $\hat{x}(y)$ such that if $x \in (\tilde{x}, \hat{x}(y))$,

$$p(x) < p(\tilde{x}) + \frac{y^{ND}}{y} p'(\tilde{x})(x - \tilde{x}).$$

Therefore, by Theorem 1, almost all such realizations (x, y) are optimally concealed. And therefore there is an open set of such (x, y) with $x \in [x', x'']$ such that $d(x, y) < d_m(x, y)$.

Proof of Statement 2. The proof of the second statement is analogous. ■

Proof of Proposition 3

Suppose (x, y) is such that $(x - \tilde{x})(y - \tilde{y}) > 0$. Then if

$$p(x) > p(\tilde{x}) + \frac{\tilde{y}}{y} p'(\tilde{x})(x - \tilde{x}),$$

it must be that

$$p(x) > p(\tilde{x}) + p'(\tilde{x})(x - \tilde{x})$$

also holds. Therefore, if $(x - \tilde{x})(y - \tilde{y}) > 0$, $d_m(x, y) = 1$ implies $d(x, y) = 1$; and so $d(x, y) \geqslant d_m(x, y)$. If instead (x, y) is such that $(x - \tilde{x})(y - \tilde{y}) < 0$, then

$$p(x) > p(\tilde{x}) + p'(\tilde{x})(x - \tilde{x}) \Rightarrow p(x) > p(\tilde{x}) + \frac{\tilde{y}}{y} p'(\tilde{x})(x - \tilde{x}).$$

Consequently, in that case $d(x, y) = 1$ implies $d_m(x, y) = 1$; and so $d(x, y) \leqslant d_m(x, y)$. ■

Proof of Proposition 4

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 an advisor with profitability $y > 0$, upon observing x, will choose to disclose it if and only if $x \geqslant \hat{x}$ (assuming that the advisor discloses when indifferent). Conversely, an advisor with profitability $y < 0$ will disclose x if and only if $x \leqslant \hat{x}$. An advisor 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 values and profitabilities 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 values and profitabilities 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 $d^*(x, y) \in \{0, 1\}$ and $d^*(x, y) = 1 \Leftrightarrow (x - \hat{x})y \geqslant 0$ defines an equilibrium disclosure strategy.
■

Proof of Proposition 5

Under mandated transparency, the buyer forms mean posteriors about the object separately for each profitability level y. It follows from a standard unravelling argument (reproduced below) that there can be no concealment in equilibrium.

Suppose for some $y \in \mathcal{Y}$, $\mathbb{X} \subseteq \mathcal{X}$ is the set of evidence realizations that the advisor does not disclose to the buyer; that is, so that $d(x,y) < 1$ if $x \in \mathbb{X}$. And suppose $\mathbb{X}$ is a set with positive probability according to $F_{X|y}$. Then upon seeing no disclosure, and that the object's profitability is y, the buyer's "no disclosure" belief is

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

From this construction, it must be that there is a positive probability that $x \in \mathbb{X}$ and $x > x_y^{ND}$. And therefore, if one such x realizes, the advisor's best response is to disclose it to the buyer, thereby deviating from the candidate equilibrium strategy $d(x,y) < 1$. We can therefore conclude that, for each y, any $d(\cdot;y)$ such that no disclosure happens with positive probability according to $F_{X|y}$ cannot be an equilibrium disclosure strategy.

And so it follows that the unique equilibrium disclosure strategy is full disclosure. ■

B Additional Results

B.1 Interior vs. Corner Optimal Disclosure Rules

Theorem 1 shows that an optimal disclosure rule is such that, for some $\hat{x} \in \mathcal{X}$ and $\hat{y} \in \mathcal{Y}$,

$$(x-\hat{x})(y-\bar{y}(x)) \geqslant 0 \Rightarrow d_{\hat{x},\hat{y}}(x,y) = 1, \quad (x-\hat{x})(y-\bar{y}(x)) < 0 \Rightarrow d_{\hat{x},\hat{y}}(x,y) = 0,$$

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

And the following system of two equations must hold:

$$\begin{cases} \hat{x} = \mathbb{E}\left[x|\text{no disclosure region implied by } d_{\hat{x},\hat{y}}\right] \\ \hat{y} = \mathbb{E}\left[y|\text{no disclosure region implied by } d_{\hat{x},\hat{y}}\right] \end{cases} \tag{17}$$

But note that, if $\hat{x} = \inf(\mathcal{X})$ and $\hat{y} = \inf(\mathcal{Y})$, or $\hat{x} = \sup(\mathcal{X})$ and $\hat{y} = \sup(\mathcal{Y})$, it is possible for the no disclosure region implied by $d_{\hat{x},\hat{y}}$ to be empty.[16] In that case, the expected values on the right hand side of the equations in (17) are not defined by Bayesian updating. For convenience, I take the stance that the buyer's "off-path" beliefs of no disclosure are such that (17) is vacuously satisfied. This means that, if

[16]For any other "corner cases," in which $\hat{x} \in \{\inf(\mathcal{X}), \sup(\mathcal{X})\}$ or $\hat{y} \in \{\inf(\mathcal{Y}), \sup(\mathcal{Y})\}$, the implied disclosure rule $d_{\hat{x},\hat{y}}$ is necessarily such that no disclosure happens with positive probability. In the proof of Lemma 1, I argue that these corner disclosure rules cannot be solutions to the advisor's problem, as they do not satisfy system (17).

the no disclosure set is empty at one or both corners — $\hat{x} = \inf(\mathcal{X})$ and $\hat{y} = \inf(\mathcal{Y})$, or $\hat{x} = \sup(\mathcal{X})$ and $\hat{y} = \sup(\mathcal{Y})$ — then these "corner solutions" with full disclosure are candidate optimal disclosure rules.

Lemma 1 and Proposition 6 below delineate conditions so that the optimal disclosure rule does not involve full disclosure and does not correspond to one of the two potential corner solutions. Specifically, Lemma 1 shows that if system (17) has an interior solution, then such solution must imply a disclosure rule that yields strictly higher payoff to the advisor than full disclosure.

Next, Proposition 6 proposes three conditions that guarantee that system (17) has an interior solution. The first condition guarantees that, for any $\hat{x} \in \mathcal{X}$ and $\hat{y} \in \mathcal{Y}$, the implied disclosure rule $d_{\hat{x},\hat{y}}$ induces no disclosure with positive probability; in that case, I show that (17) must have an interior solution. This first condition is satisfied, for example, if the demand function p is strictly concave; but it can also hold for non-concave demand functions. The second condition states that (17) also has an interior solution if the demand function p is affine. Finally, condition 3 shows that the same holds if p is "close enough" to being an affine function; specifically, this condition guarantees that the optimal disclosure rule is interior for strictly convex demand functions that are "not too convex." (In section 3.2.1, I work out an example where the demand function p is convex and the optimal disclosure rule is interior.)

Lemma 1. *If system (17) has an interior solution, with $\hat{x} \in int(\mathcal{X})$ and $\hat{y} \in int(\mathcal{Y})$, then the optimal disclosure rule is interior and such that no disclosure happens with positive probability.*

Proposition 6.

1. *If there exist $x, x' \in \mathcal{X}$ such that*

$$\frac{p'\left(\inf(\mathcal{X})\right)\left(x - \inf(\mathcal{X})\right)}{p(x) - p\left(\inf(\mathcal{X})\right)} > 1 \text{ and } \frac{p'\left(\sup(\mathcal{X})\right)\left(x' - \sup(\mathcal{X})\right)}{p(x') - p\left(\sup(\mathcal{X})\right)} < 1,$$

   then the optimal disclosure rule is interior and no disclosure happens with positive probability.

2. *If p is an affine function, then the optimal disclosure rule is interior and no disclosure happens with positive probability.*

3. *If $p = \alpha p_1 + (1-\alpha)p_2$, where p_1 is an affine function, then there exists $\bar{\alpha} \in (0,1)$ such that if $\alpha > \bar{\alpha}$, the optimal disclosure rule is interior and no disclosure happens with positive probability.*

B.2 Proof of Lemma 1

Suppose system (17) has a solution with $\hat{x} \in int(\mathcal{X})$ and $\hat{y} \in int(\mathcal{Y})$, and let $d_{\hat{x},\hat{y}}$ be the corresponding disclosure rule defined by such $\hat{x}$ and $\hat{y}$.

Step 1. Showing that $d_{\hat{x},\hat{y}}$ yields strictly higher value to the advisor than full disclosure.

For each $\alpha \in [0,1]$, define the following alternative disclosure rule:

$$1 - d_\alpha(x,y) = \begin{cases} 1 - d_{\hat{x},\hat{y}}(x,y), \text{ if } d_{\hat{x},\hat{y}}(x,y) = 1 \\ \alpha\left(1 - d_{\hat{x},\hat{y}}(x,y)\right), \text{ otherwise} \end{cases}$$

Note that, because for all $\alpha > 0$, the probability of no disclosure implied by d_α is proportional to that implied by $d_{\hat{x},\hat{y}}$, it must be that these rules imply the same x^{ND} and y^{ND}. Therefore, because $d_{\hat{x},\hat{y}}$ is such that $\hat{x}$ and $\hat{y}$ solve the system (17), we know that as α decreases, there is an increase in disclosure for realizations (x,y) such that $\partial\Pi/\partial d(x,y)$, as given by (15), is strictly negative. Therefore, the value of d_α to the advisor is strictly increasing in α. Moreover, full disclosure corresponds to the case of $\alpha = 0$, which is therefore dominated by $d_{\hat{x},\hat{y}}$.

Step 2. Showing that any "corner disclosure rule" $d_{\tilde{x},\tilde{y}}$ defined by $\tilde{x} \in \{\inf(\mathcal{X}), \sup(\mathcal{X})\}$ or $\tilde{y} \in \{\inf(\mathcal{Y}), \sup(\mathcal{Y})\}$ such that no disclosure happens with positive probability cannot be an optimal disclosure rule.

For any disclosure rule such that no disclosure happens with positive probability, it must be that $\mathbb{E}(x|\text{ no disclosure}) \in int(\mathcal{X})$ and $\mathbb{E}(y|\text{ no disclosure}) \in int(\mathcal{Y})$ — because F has full support over $\mathcal{X} \times \mathcal{Y}$ and no mass points. Therefore, if $\tilde{x} \in \{\inf(\mathcal{X}), \sup(\mathcal{X})\}$ or $\tilde{y} \in \{\inf(\mathcal{Y}), \sup(\mathcal{Y})\}$, the system (17) is not satisfied, and by Theorem 1 disclosure rule $d_{\tilde{x},\tilde{y}}$ can be strictly improved.

■

B.3 Proof of Proposition 6

For each $\hat{x} \in \mathcal{X}$ and $\hat{y} \in \mathcal{Y}$, define $d_{\hat{x},\hat{y}}$ by

$$(x-\hat{x})(y-\bar{y}(x)) \geqslant 0 \Rightarrow d_{\hat{x},\hat{y}}(x,y) = 1, \quad (x-\hat{x})(y-\bar{y}(x)) < 0 \Rightarrow d_{\hat{x},\hat{y}}(x,y) = 0,$$

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

Proof of Statement 1. Suppose there exist $x, x' \in \mathcal{X}$ such that

$$\frac{p'\left(\inf(\mathcal{X})\right)\left(x - \inf(\mathcal{X})\right)}{p(x) - p\left(\inf(\mathcal{X})\right)} > 1 \text{ and } \frac{p'\left(\sup(\mathcal{X})\right)\left(x' - \sup(\mathcal{X})\right)}{p(x') - p\left(\sup(\mathcal{X})\right)} < 1. \tag{18}$$

Note that, because (18) holds and p is continuously differentiable, for all $\hat{x} \in \mathcal{X}$ and $\hat{y} \in \mathcal{Y}$, $d_{\hat{x},\hat{y}}(x,y) = 0$ for some non-empty open subset of $\mathcal{X} \times \mathcal{Y}$. Therefore, $\mathbb{E}\left[x|\text{no disclosure region implied by } d_{\hat{x},\hat{y}}\right]$ and $\mathbb{E}\left[y|\text{no disclosure region implied by } d_{\hat{x},\hat{y}}\right]$ are well defined. Consider the following two continuous functions

$$\Gamma_1(\hat{x},\hat{y}) = \mathbb{E}\left[x|\text{no disclosure region implied by } d_{\hat{x},\hat{y}}\right] - \hat{x}$$

and $\Gamma_2(\hat{x},\hat{y}) = \mathbb{E}\left[y|\text{no disclosure region implied by } d_{\hat{x},\hat{y}}\right] - \hat{y}$.

For a disclosure rule such that no disclosure happens with positive probability, $\mathbb{E}(x|$ no disclosure$) \in int(\mathcal{X})$ and $\mathbb{E}(y|$ no disclosure$) \in int(\mathcal{Y})$ — because F has full support over $\mathcal{X} \times \mathcal{Y}$ and no mass points. Consequently, for every $\hat{y}$, $\Gamma_1(\inf(\mathcal{X}),\hat{y}) > 0$ and $\Gamma_1(\sup(\mathcal{X}),\hat{y}) < 0$. Similarly, for every $\hat{x}$, $\Gamma_2(\hat{x},\inf(\mathcal{Y})) > 0$ and $\Gamma_1(\hat{x},\sup(\mathcal{Y})) < 0$. Therefore, by the Pointcaré-Miranda Theorem, there exist $\hat{x} \in int(\mathcal{X})$ and $\hat{y} \in int(\mathcal{Y})$ such that both $\Gamma_1(\hat{x},\hat{y})$ and $\Gamma_2(\hat{x},\hat{y})$ are simultaneously equal to 0; which are solutions to system (17). ■

Proof of Statement 2. Suppose p is affine. The proof uses the following claim.

Claim 3. *Let $x_\epsilon = \inf(\mathcal{X}) + \epsilon$ and $y_\epsilon = \inf(\mathcal{Y}) + \epsilon$. Then*

$$\lim_{\epsilon\to 0} \mathbb{E}\left[x|\textit{no disclosure region implied by } d_{x_\epsilon,y_\epsilon}\right] > \inf(\mathcal{X}),$$

$$\textit{and } \lim_{\epsilon\to 0} \mathbb{E}\left[y|\textit{no disclosure region implied by } d_{x_\epsilon,y_\epsilon}\right] > \inf(\mathcal{Y}).$$

Proof. Because p is affine, $\bar{y}(x) = y_\epsilon$, so no disclosure happens if and only if either (i) $x < x_\epsilon$ and $y > y_\epsilon$, or (ii) $x > x_\epsilon$ and $y < y_\epsilon$. Therefore, $\mathbb{E}\left[x|\text{no disclosure region implied by } d_{x_\epsilon,y_\epsilon}\right]$ is given by

$$\inf(\mathcal{X}) + \left[\int_{\inf(\mathcal{Y})}^{y_\epsilon}\int_{\epsilon}^{\Delta} x f_{X|y}(\inf(\mathcal{X})+x)f_Y(y)dxdy + \int_{y_\epsilon}^{\sup(\mathcal{Y})}\int_0^{\epsilon} x f_{X|y}(\inf(\mathcal{X})+x)f_Y(y)dxdy\right]$$

$$\times\left[\int_{\inf(\mathcal{Y})}^{y_\epsilon}\int_{\epsilon}^{\Delta} f_{X|y}(\inf(\mathcal{X})+x)f_Y(y)dxdy + \int_{y_\epsilon}^{\sup(\mathcal{Y})}\int_0^{\epsilon} f_{X|y}(\inf(\mathcal{X})+x)f_Y(y)dxdy\right]^{-1},$$

where I use the notation $\Delta = \sup(\mathcal{X}) - \inf(\mathcal{X})$. To assess the limit of this object as $\epsilon \to 0$, I use L'Hospital's rule. The derivative of the numerator is

$$\frac{\partial \text{NUM}}{\partial \epsilon} = \int_{\epsilon}^{\Delta} x f_{X|y_\epsilon}(\inf(\mathcal{X})+x)f_Y(y_\epsilon)dx - \int_0^{\epsilon} x f_{X|y_\epsilon}(\inf(\mathcal{X})+x)f_Y(y_\epsilon)dx$$
$$- \int_{\inf(\mathcal{Y})}^{y_\epsilon} \epsilon f_{X|y}(\inf(\mathcal{X})+\epsilon)f_Y(y)dy + \int_{y_\epsilon}^{\sup(\mathcal{Y})} \epsilon f_{X|y}(\inf(\mathcal{X})+\epsilon)f_Y(y)dy.$$

The derivative of the denominator is

$$\frac{\partial \text{DEN}}{\partial \epsilon} = \int_{\epsilon}^{\Delta} f_{X|y_\epsilon}(\inf(\mathcal{X})+x)f_Y(y_\epsilon)dx - \int_0^{\epsilon} f_{X|y_\epsilon}(\inf(\mathcal{X})+x)f_Y(y_\epsilon)dx$$
$$- \int_{\inf(\mathcal{Y})}^{y_\epsilon} f_{X|y}(\inf(\mathcal{X})+\epsilon)f_Y(y)dy + \int_{y_\epsilon}^{\sup(\mathcal{Y})} f_{X|y}(\inf(\mathcal{X})+\epsilon)f_Y(y)dy.$$

Therefore we have

$$\lim_{\epsilon\to 0}\mathbb{E}\left[x|\text{no disclosure region implied by } d_{x_\epsilon,y_\epsilon}\right] = \inf(\mathcal{X}) + \frac{\lim_{\epsilon\to 0}\frac{\partial\text{NUM}}{\partial\epsilon}}{\lim_{\epsilon\to 0}\frac{\partial\text{DEN}}{\partial\epsilon}} =$$

$$= \inf(\mathcal{X}) + \frac{\int_0^{\Delta} x f_{X|\inf(\mathcal{Y})}(\inf(\mathcal{X})+x) f_Y(\inf(\mathcal{Y}))dx}{\int_0^{\Delta} f_{X|\inf(\mathcal{Y})}(\inf(\mathcal{X})+x) f_Y(\inf(\mathcal{Y}))dx + \int_{\inf(\mathcal{Y})}^{\sup(\mathcal{Y})} f_{X|y}(\inf(\mathcal{X})) f_Y(y)dy} > \inf(\mathcal{X}),$$

where the inequality is due to f_Y and $f_{X|y}$, for each y, being strictly positive densities. The second part of the second observation — that $\lim_{\epsilon\to 0}\mathbb{E}\left[y|\text{no disclosure region implied by } d_{x_\epsilon,y_\epsilon}\right] > \inf(\mathcal{Y})$ — can be shown analogously. ∎

By Claim 3, we know that there exists some $\delta > 0$ such that $\Gamma_1(\inf(\mathcal{X})+\delta, \inf(\mathcal{Y})+\delta) > 0$ and $\Gamma_2(\inf(\mathcal{X})+\delta, \inf(\mathcal{Y})+\delta) > 0$. Moreover, we have that for any $\hat{y} > \inf(\mathcal{Y})+\delta$, it also holds that

$$\mathbb{E}\left[x|\text{no disclosure}\right] > \delta, \text{ and therefore } \Gamma_1(\inf(\mathcal{X})+\delta, \hat{y}) > 0$$

(This is true because, as $\hat{y}$ increases, the "no disclosure" region with $x > \delta$ becomes larger and the "no disclosure region with $x < \delta$ becomes smaller.) Consequently, for all $\hat{y} > \inf(\mathcal{Y})+\delta$, $\Gamma_1(\inf(\mathcal{X})+\delta, \hat{y}) > 0$. Analogously, for all $\hat{x} > \inf(\mathcal{X})+\delta$, $\Gamma_2(\hat{x}, \inf(\mathcal{Y})+\delta) > 0$.

Now consider the following claim (stated without proof, as it is analogous to that of Claim 3).

Claim 4. *Let $x^\epsilon = \sup(\mathcal{X}) - \epsilon$ and $y^\epsilon = \sup(\mathcal{Y}) - \epsilon$; then*

$$\lim_{\epsilon\to 0}\mathbb{E}\left[x|\textit{no disclosure region implied by } d_{x_\epsilon,y_\epsilon}\right] < \sup(\mathcal{X}),$$

$$\textit{and } \lim_{\epsilon\to 0}\mathbb{E}\left[y|\textit{no disclosure region implied by } d_{x_\epsilon,y_\epsilon}\right] < \sup(\mathcal{Y}).$$

By Claim 4, we know that there exists some $\delta' > 0$ such that $\Gamma_1(\sup(\mathcal{X})-\delta', \sup(\mathcal{Y})-\delta') < 0$ and $\Gamma_2(\sup(\mathcal{X})-\delta', \sup(\mathcal{Y})-\delta') < 0$. Moreover, we have that for any $\hat{y} < \sup(\mathcal{Y})-\delta'$, it also holds that $\Gamma_1(\sup(\mathcal{X})-\delta', \hat{y}) < 0$. (As $\hat{y}$ decreases, the "no disclosure" region with $x < \delta'$ becomes larger and the "no disclosure region with $x > \delta'$ becomes smaller.) Consequently, for all $\hat{y} < \sup(\mathcal{Y})-\delta'$, $\Gamma_1(\sup(\mathcal{X})-\delta', \hat{y}) < 0$. Analogously, for all $\hat{x} < \sup(\mathcal{X})-\delta'$, $\Gamma_2(\hat{x}, \sup(\mathcal{Y})-\delta) < 0$.

Putting all together, there exists a convex compact subset $\tilde{\mathcal{X}}\times\tilde{\mathcal{Y}}$ of $\mathcal{X}\times\mathcal{Y}$, defined by $\tilde{\mathcal{X}} = [\inf(\mathcal{X})+\delta, \sup(\mathcal{X})-\delta']$ and $\tilde{\mathcal{Y}} = [\inf(\mathcal{Y})+\delta, \sup(\mathcal{Y})-\delta']$ such that for all $\hat{y}\in\tilde{\mathcal{Y}}$,

$$\Gamma_1(\inf(\tilde{\mathcal{X}}), \hat{y}) < 0, \text{ and } \Gamma_1(\sup(\tilde{\mathcal{X}}), \hat{y}) > 0.$$

Similarly, for all $\hat{x}\in\tilde{\mathcal{X}}$,

$$\Gamma_2(\hat{x}, \inf(\tilde{\mathcal{Y}})) < 0, \text{ and } \Gamma_2(\hat{x}, \sup(\tilde{\mathcal{Y}})) > 0.$$

Therefore, by the Pointcaré-Miranda Theorem, there exist $\hat{x}\in int(\tilde{\mathcal{X}})$ and $\hat{y}\in int(\tilde{\mathcal{Y}})$ such that both

$\Gamma_1(\hat{x}, \hat{y})$ and $\Gamma_2(\hat{x}, \hat{y})$ are simultaneously equal to 0; which are solutions to system (17). ■

Proof of Statement 3. Let $d^{\alpha}_{\hat{x},\hat{y}}$ be the disclosure rule implied by thresholds $\hat{x}$ and $\hat{y}$ when the demand function is given by $p = \alpha p_1 + (1 - \alpha)p_2$. Because p_2 is continuously differentiable, regardless of its curvature, the disclosure rule $d^{\alpha}_{\hat{x},\hat{y}}$ induces no disclosure with positive probability, except perhaps if $(\hat{x}, \hat{y}) = (\inf(\mathcal{X}), \inf(\mathcal{Y}))$ or $(\hat{x}, \hat{y}) = (\sup(\mathcal{X}), \sup(\mathcal{Y}))$.

Moreover, for each $\hat{x} \in int(\mathcal{X})$ and $\hat{y} \in int(\mathcal{Y})$, the values

$$\Gamma^{\alpha}_1(\hat{x}, \hat{y}) = \mathbb{E}\left[x|\text{no disclosure region implied by } d^{\alpha}_{x_\epsilon, y_\epsilon}\right] - \hat{x}$$

$$\text{and } \Gamma^{\alpha}_2(\hat{x}, \hat{y}) = \mathbb{E}\left[y|\text{no disclosure region implied by } d^{\alpha}_{x_\epsilon, y_\epsilon}\right] - \hat{y}$$

are continuous functions of α. We know by Statement 2 that, for $\alpha = 1$, there is an interior point $(\hat{x}, \hat{y})$ such that Γ^{α}_1 and Γ^{α}_2 are simultaneously equal to 0. By continuity, it must be that there exists $\bar{\alpha} \in (0, 1)$ such that if $\alpha > \bar{\alpha}$, then there is an interior point $(\hat{x}', \hat{y}')$ such that Γ^{α}_1 and Γ^{α}_2 are simultaneously equal to 0 as well. ■

Alternative Proof of Statement 2. This is an additional argument showing that full disclosure is not an optimal disclosure rule when p is an affine function. In that case, the optimal disclosure rule must conceal realizations with positive probability, and therefore induce $x^{ND} \in \text{int}(\mathcal{X})$ and $y^{ND} \in \text{int}(\mathcal{Y})$. By Theorem 1, $\hat{x} = x^{ND}$ and $\hat{y} = y^{ND}$ must therefore be a solution to system (17).

The demand function is affine, given by $p(x) = a + bx$. For any disclosure rule d, the first term in the advisor's payoff expressed in (3) must be $\mathbb{E}(y)\mathbb{E}[P(y, d)] = \mathbb{E}(y)\left[a + b\mathbb{E}(x)\right]$. (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 let d^1 be the full disclosure rule; and take any other disclosure rule $d_{\hat{x},\hat{y}}$, 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 $y > \bar{y}$, $P(y, d) > P(y, d^1)$. Consequently, the covariance of sales and profitability is larger under d than under d^1; and thus d^1 is not an optimal disclosure rule. ■