

Second Opinions and the Humility Threshold

David Lagziel[†] and Yevgeny Tsodikovich

Discussion Paper No. 23-05

August 2023

Monaster Center for
Economic Research
Ben-Gurion University of the Negev
P.O. Box 653
Beer Sheva, Israel

Fax: 972-8-6472941
Tel: 972-8-6472286

Second Opinions and the Humility Threshold ^{*}

David Lagziel[†] and Yevgeny Tsodikovich[‡]

July 2022

Abstract

We study a possible conflict of interests between a manager and a consultant, where each possesses a private signal regarding an unknown state. Each of them strives to provide an accurate recommendation regarding the state, while preferring to be the only one to do so. Our analysis depicts a mapping from their expertise levels to the equilibria of this game, showing that: (i) better-informed players may generate worse recommendations in equilibrium; and (ii) ordering the players so that the lower-level one provides the second opinion typically improves the outcome. Moreover, we show that limiting the players' liability facilitates cooperation and increases the probability of reaching a correct decision.

Journal classification numbers: C72, D82, D83

Keywords: Sequential Bayesian Games; Guided Equilibrium; Learning.

^{*}For their valuable comments, the authors thank Nicolas Vieille, Eran Shmaya, Amnon Schreiber, Eilon Solan, Xavier Venel, Ehud Lehrer, as well as the participants of the 7th Workshop on Stochastic Methods in Game Theory in Erice, SING17 – 17th European Meeting on Game Theory, and the 12th annual conference of the Israeli Chapter of the Game Theory Society. Lagziel acknowledges the support of the Israel Science Foundation, Grant #513/19. Tsodikovich acknowledges the support of the Israel Science Foundation Grants #2566/20, #1626/18 & #448/22.

[†]Ben-Gurion University of the Negev, Beer-Sheva 8410501, Israel. E-mail: Davidlag@bgu.ac.il

[‡]Department of Economics, Bar Ilan University, Israel. e-mail: yevgets@gmail.com

1 Introduction

“Two are better than one, because they have a good return for their labor. For if one falls down, his companion can lift him up.” Ecclesiastes 4:9-10

These simple words from Ecclesiastes capture the natural idea that two individuals can provide backup to one another: in case one fails, the other can still deliver. This is also the reason why in situations of growing complexity and uncertainty, people feel the need to consult experts, before making important decisions. Such decisions could vary from assessing an academic paper for publication, to career choices, and even complex medical procedures. In this paper, we aim to answer in what sense, and under which conditions, two are indeed better than one.

We consider a game between a consultant and a manager. Each receives a private informative signal regarding an unknown state. The consultant provides a preliminary assessment, whereas the manager takes the final decision, based on his private information and the consultant’s recommendation. In this setting, every player naturally strives to provide the correct assessment. Whenever a player is incorrect, he may wish the other player to be incorrect as well, so that they share the possible liability. This also holds when a player is indeed correct, as he establishes superiority over the other. For example, an accurate assessment by the consultant, along with an incorrect one from the manager, underscores the importance of the consultant’s services. Our goal is to shed light on the strategic interaction between the manager and consultant, and its implications on the ability to reach a correct decision.

The described process is, in fact, a sequential two-player Bayesian-game between the consultant and the manager. Our analysis shows that there are 4 main equilibria profiles: (i) a *Revealing equilibrium* in which both players reveal their private information; (ii) a *Herding equilibrium* in which the manager mimics the consultant; (iii) a *Guided equilibrium* in which some learning among the players takes place; and (iv) a *Mixed equilibrium* which builds on the Guided equilibrium (with mixed strategies). A broader discussion of these equilibria is given in Section 2.

To study the nature of these equilibria, we extend our analysis in two ways. First, we map the parameter-space (of expertise levels and payoffs) to the different equilibria, which

results in four disjoint sets, one for each equilibrium. Second, we use the notion of *correctness*, defined through the probability that the manager reaches a correct decision, to classify each of the mentioned equilibria. Due to the learning involved, the guided equilibrium evidently supports the highest correctness level, which is evidently of high importance for the firm and its shareholders.

We then combine these two research paths to present two intriguing effects. The first effect shows that the correctness of the process is not a monotonic function of the players' expertise levels. Specifically, we show that an infinitesimal *increase* in the manager's expertise level may yield a stark drop in the correctness of the process. We refer to the level at which this phenomenon occurs as *The Humility Threshold*.

The Humility Threshold extends the well-known herding phenomenon in two ways. Herding typically occurs when a player is less informed than his predecessors. In our setup, however, the drop in the correctness of the process follows from the increase in the expertise level of the *manager*, rather than the consultant. Moreover, the Humility Threshold is also driven by payoff externalities (i.e., other-regarding preferences), rather than information alone.

To better understand this effect, one needs to consider the transition between the previously mentioned equilibria. Consider, for example, a hypothetical situation in which a firm faces a complex situation. Before the manager takes a decision, based on his private signal, he approaches a consultant who possesses a private informative signal. The probability that the private signal of each player matches the correct state is referred to as the expertise level q_i of player i . We typically assume that: (i) all recommendations are public; (ii) every player prefers to be correct, while hoping the other to be incorrect;¹ and (iii) in any case, the manager is the one who takes the final decision.

The equilibria of such games are naturally determined according to the specific payoffs and exogenous probabilities. Yet, *ceteris paribus*, we can still ask how does the expertise level of the manager affects the final decision. Figure 1a illustrates that the relation between these two elements is not necessarily monotone. The figure depicts the probability that the manager's

¹Our analysis also accounts for the limit case in which the manager and consultant are indifferent to the other's correctness.

final decision is correct, generally referred to as *correctness*, as a function of the manager’s expertise level, q_2 . The blunt drop occurs when q_2 crosses the Humility Threshold because the game transitions from a Guided equilibrium to a Revealing one. As long as the expertise level of the manager is relatively small, he remains attentive to the assessment of the consultant in a way that allows for some learning. Once the manager crosses this threshold, the game shifts to a Revealing equilibrium in which the manager completely ignores the assessment of the consultant, thus decreasing the probability to reach a correct decision.

This emphasizes a possible conflict of interest between the manager and the shareholders. While the shareholders want to maximize the probability that the correct decision is being taken, the manager might have an incentive to “gamble” due to large payoff externalities, whenever he is correct and the consultant gave a false recommendation. Notice that this conflict of interest is limited, because the manager would always prefer taking the correct decision for the shareholders (i.e., the one that matches the state), had he known the realized state, and independently of the consultant’s assessment. Moreover, in Theorem 1* , we show that by correctly incentivizing the manager, this conflict of interest can be removed.

The second effect, referred to as *The Intern Effect*, focuses on the relative expertise levels of the manager and the consultant. Given the ability to choose, we ask whether a better-informed manager is more beneficial, in terms of correctness, compared to a better-informed consultant. Our analysis shows that in some cases it is better to have a less-informed manager and a better-informed consultant. This phenomenon also follows from the transition between two equilibria, where a less-informed manager learns from the consultant’s recommendation in equilibrium, whereas the alternative eliminates this possibility.

The driving force behind these effects is the disparity between the players’ (game-based) expected payoffs and the ex-post success rate of the manager’s decision. Yet, it should be clear that both players are not oblivious to their ex-post reputation. The fact that *both* players provide an assessment for every incident, rather than just one, supports a payoff structure that incorporates their reputation concerns. In this sense, the stated effects complement the results of Ely and Välimäki (2003), where experts’ long-run reputation considerations over-shed any short-term payoffs and eliminate all the surplus.

The correctness of the manager’s decision

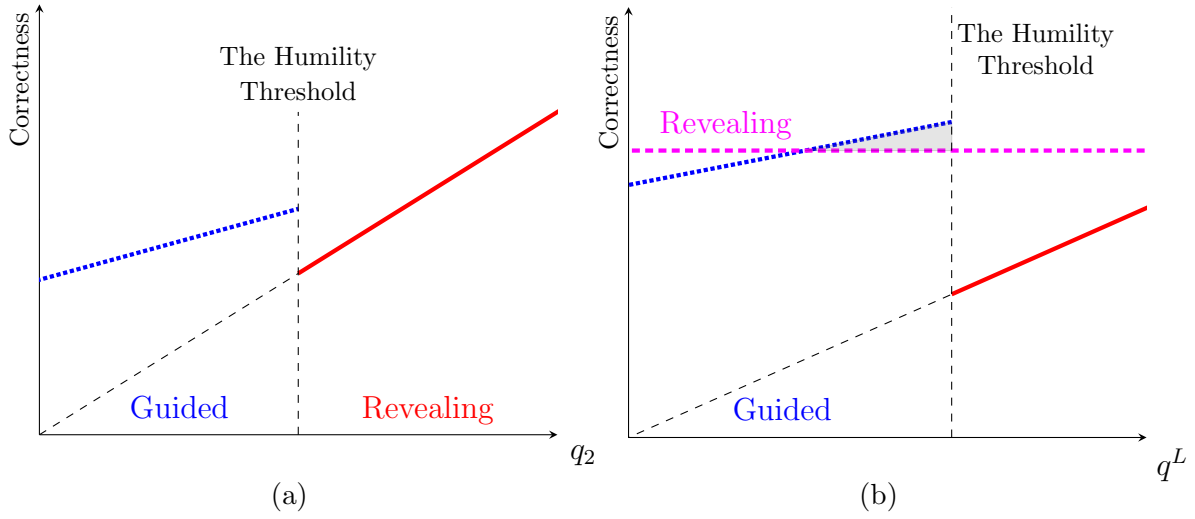


Figure 1: In Figure 1a, an infinitesimal increase in the manager’s expertise level shifts the game from a Guided equilibrium, described by the dotted (blue) line, to a Revealing equilibrium, depicted by the solid (red) line. This leads the manager to ignore the consultant, an action that eventually results in a worse expected outcome. Figure 1b presents two possible equilibria to the left of the Humility threshold. The dotted (blue) line presents the equilibrium, a guided one, given that the manager is less-informed, and the dashed horizontal (pink) line presents the equilibrium, a Revealing one, in case the manager is more-informed.

1.1 Related research

This paper encompasses elements from several different sub-fields, the primary being social learning with externalities, which expanded remarkably since the studies of Banerjee (1992), Bikhchandani et al. (1992), and Smith and Sorensen (2000).² Social learning with externalities typically relates to either positive or negative congestion costs, under which players’ payoffs either increase or decrease depending on the number of players who choose the same action (see, e.g., Veeraraghavan and Debo (2011), Debo et al. (2011), and Eyster et al. (2014), among others). On the one hand, our study matches this line of research through the possibility of learning under negative externalities among experts, as well as the basic information structure and actions. For example, Eyster et al. (2014) show that backward-looking negative externalities (i.e., players’ actions are less profitable the more they are played by others) prevent action fixation and improve social learning. This resembles our notion of a Guided equilibrium, where

²For a recent extensive review of this topic, see Bikhchandani et al. (2021).

the second player learns from the first, through payoff (and information) considerations. On the other hand, we study the *equilibria* of a Bayesian-game, which includes both positive and negative externalities at the same time. So in our set-up, the highest/lowest payoff is achieved when a player chooses the ex-post correct/incorrect action alone.

Another critical difference relays on the fact that our model is a two-stage two-player Bayesian-game, thus eliminating the possibility of either asymptotic learning as in Arieli (2017) and Mossel et al. (2020), or repeated interactions as in De Clippel et al. (2021). A study more related to ours is Dasgupta (2000), which also focuses on finite and sequential Bayesian-games.³ It identifies transitions between equilibria outcomes through a trigger strategy, that depends on the private belief of the stage player. Though the payoffs and information structure (as well as the monotone likelihood ratio property) are rather different from our work, the shifts between equilibria through a threshold strategy resemblances our equilibria mapping and transitions. This is in contrast to the study of Ali and Kartik (2012) that identifies a unique Herding equilibrium in case the players' preferences are aligned, in the sense that each player prefers other players to be correct too. One could also find some connection between our study and the two-stage model of Feldman et al. (2019), which focuses on social learning regarding the product design of a monopolist firm (similarly to Crapis et al., 2017), with the obvious distinction regarding the underlying game.

Our study also relates to the research agenda of Bayesian games and Bayesian comparative statics.⁴ The more relevant and recent studies in this field are Jensen (2018) and Mekonnen and Vizcaíno (2022), who study how the distribution of individual decisions and equilibria outcomes vary with changes in the underlying economic parameters. They follow the studies of Athey (2001) and Van Zandt and Vives (2007) who prove that Bayesian-Nash equilibrium profiles are point-wise monotone functions of the players' beliefs. Mekonnen and Vizcaíno (2022) study how a higher accuracy level of one player influences the equilibrium actions of all others. This

³Another related and more recent study by Li and Norman (2021) takes a more general perspective on the topic of sequential Bayesian games, and provides a tractable characterization of equilibria outcomes.

⁴This line of research originates from the field of statistical decision-making, who study how variations in the information structure influence a single decision maker, going back to the studies of Blackwell (1951, 1953), Lehmann (1988), Quah and Strulovici (2009), and more recently Lagziel and Lehrer (2019, 2022).

again resembles our analysis, and specifically, the two previously discussed effects (the Humility threshold and The Intern Effect). The main differences between these studies and the current work are: (i) their focus on one-stage Bayesian games (rather than sequential); (ii) their use of a continuum of actions and different payoff functions (quasi-concave differences/supermodular); and (iii) the relevant information structures.

Structure of the paper. In Section 2 we describe the model and the key definitions. In Section 3 we present the main results, divided into two subsections: in Subsection 3.1 we deal with the case of symmetric payoffs, and in Subsection 3.2 we extend the discussion to asymmetric payoffs, where the intern effect and the humility threshold exist. Concluding remarks are given in Section 4. To facilitate readability, proofs are relegated to the appendix.

2 The model

A manager wishes to identify a binary unknown state. To do so, he approaches a consultant for a recommendation, and based on it and on the manager's private information, he takes a decision. Each of the two strives to establish superiority over the other, in the sense that they individually prefer to provide an accurate assessment of the realized state, whilst hoping the other falls short of this goal. This generates a sequential Bayesian game where their payoffs depend on the realized profile of recommendations and on the realized state.

Formally, consider the following two-player, incomplete-information sequential game G . There are two states denoted by $\theta \in \Theta = \{0, 1\}$, and a prior probability $p = \Pr(\theta = 0) > \frac{1}{2}$. Given θ , every player $i \in \{c, m\}$ receives an independent, noisy and informative signal $s_i \in S = \{0, 1\}$, such that $\Pr(s_i = \theta | \theta) = q_i$. One can think of q_i as *the expertise level* of player i , namely a measure of player i 's ability to identify correctly the state. The action set of every player i is denoted by $A = \{0, 1\}$.

The game evolves as follows. First, nature chooses a state θ according to a common, publicly known, prior p . Then, every player i receives a private signal s_i based on the previously defined information structure. The consultant is the first to act by posting a public recommendation

$a_c \in A$. After observing a_c , the manager makes a decision $a_m \in A$.

We assume that every player prefers choosing the action that matches the state. To establish superiority, every player also prefers that the action of the other player proves to be incorrect. The latter condition ensures that the player is either the only correct one or that both provided false assessments so that they jointly share the blame. To model these preferences, the utility of every agent is characterized by four possible payoffs: $U_1 > U_2 > U_3 > U_4$, where U_1 is the payoff in case player i is the only one to choose the correct action, U_2 is the payoff when both experts are correct, U_3 is the payoff in case both experts are incorrect, and U_4 is the payoff of player i when he is the only one to choose the wrong action. Since games are strategically equivalent under an affine transformation of payoffs, it is without loss of generality that we normalize $U_2 = 1$ and $U_3 = -1$. At this point, our analysis is divided into two parts: in Section 3.1 we assume symmetry between the two extreme cases, so $U_1 = -U_4 \equiv \alpha > 1$, whereas in Section 3.2 we study the general case where $U_1 \neq -U_4$. The realized symmetric payoffs are summarized in Table 1.

		Manager	
		$a_m = \theta$	$a_m = 1 - \theta$
Consultant	$a_c = \theta$	(1, 1)	$(\alpha, -\alpha)$
	$a_c = 1 - \theta$	$(-\alpha, \alpha)$	(-1, -1)

Table 1: The payoffs matrix given a realized state of $\theta \in \Theta$.

Denote the strategy of the consultant by $\sigma_c : S \rightarrow \Delta(A)$, and the strategy of the manager by $\sigma_m : S \times A \rightarrow \Delta(A)$. Since the game is sequential, the manager has a structural advantage of observing the recommendation of the consultant.

To ensure that the players' signals are indeed informative, independently of the state, we assume that $\min\{q_c, q_m\} \geq p$, which is equivalent to $\Pr(\theta = x | s_i = x) > \frac{1}{2}$ for every (i, x) . Otherwise, in a single-player scenario where $q_i < p$, the state $\theta = 0$ is more likely than $\theta = 1$, regardless of the player's information.⁵

⁵One can find some resemblance between our model and the example provided in Section 2 of Smith et al.

We analyze this game according to the standard Bayesian-Nash equilibrium. Our first goal is to identify an equilibrium for every composition of the parameters (p, q_c, q_m, U_1, U_4) . To achieve this goal, we define the following four equilibrium structures, and our analysis would show that for every choice of the mentioned parameters, exactly one of these equilibria exists (Figure 2 illustrates these classes in the (q_c, q_m) -plane):⁶

- A profile (σ_c, σ_m) is a *Revealing equilibrium* if the actions of both players match their private signal, i.e., if $\sigma_i = s_i$ for every i, s_i . This typically occurs when the manager is much more informed than the consultant, so he ignores his recommendation.
- A profile (σ_c, σ_m) is a *Herding equilibrium* if the decision of the manager matches the recommendation of the consultant in every realization of the game, while the consultant provides his private signal $\sigma_c = s_c$. This typically occurs when the consultant is much more informed than the manager so that the manager ignores his own private signal.
- A profile (σ_c, σ_m) is a *Guided equilibrium* if both players provide their private signal, $\sigma_i = s_i$, with the exception of $\sigma_m(a_c = 0, s_m = 1) = 0$. That is, the manager takes the decision $a_m = 0$ whenever the consultant recommends $a_c = 0$. This typically occurs when the two have close expertise levels, so that the manager learns from the consultant, and weighs-in both signals into his decision.
- A profile (σ_c, σ_m) is a *Mixed equilibrium* if both players use a mixed strategy over their recommendations. In particular, the consultant reports his signal if it is $s_c = 0$, and uses a mixed action otherwise, $\sigma_c(s_c = 1) \in \Delta(A)$. The manager follows his signal, $\sigma_m = s_m$, with the exception of $\sigma_m(a_c = 0, s_m = 1) \in \Delta(A)$. This typically occurs in similar situations to the ones described under the Guided equilibrium, but when expertise levels are rather low, relative to the prior.⁷

(2021), with the obvious distinction that we account for a broader game.

⁶Other than on the boundaries between different areas in this parameter space, where different equilibria coincide.

⁷In part, this resembles the mixed-action equilibrium given in Lipnowski and Ravid (2020), in the sense that both the sender and the receiver exercise mixed actions, to support the indifference of the opposing side, in equilibrium.

Using the previous classification and analysis of the different equilibria, we proceed to our main goal – to study the impact of the strategic interaction between the consultant and the manager on the latter’s final decision. For that purpose, we use the notion of “correctness”, which measures the probability of reaching a correct final decision, in each of the mentioned equilibria.⁸ Formally, given a profile (σ_c, σ_m) , define the *correctness of the process* to be the probability that in equilibrium the manager matches the true state of the world, so $C(q_1, q_2) = \Pr(\sigma_m = \theta | \sigma_c)$. Note that in general, C also depends on the prior p and on the payoffs U_1 and U_4 . In what follows, we will assume p, U_1 , and U_4 are fixed and study how the correctness changes with the players’ expertise levels.

Finally, we introduce a logit notation to represent probabilities, so

$$\tilde{p} := \ln \left(\frac{p}{1-p} \right), \quad \tilde{q}_i := \ln \left(\frac{q_i}{1-q_i} \right).$$

This notation allows us to simplify some of the equations and conditions resulting from Bayesian updating and present them as linear functions.

3 Main results

Our main results are divided into two parts: in Section 3.1 we study the case of symmetric payoffs, and in Section 3.2 we focus on asymmetric ones. The preliminary focus on the symmetric case facilitates the exposition of the key features of the model and resulting equilibria. Using these features, we present the more advanced and insightful aspects of the second-opinion problem in Section 3.2.

Specifically, in Proposition 2 we prove that the correctness is not a monotone function of the manager’s expertise level, so that an infinitesimal increase of q_m may trigger a stark drop in the correctness of the process. More precisely, we show that there is a threshold expertise level, namely *the humility threshold*, so that the manager ignores the action (and the conveyed information) of the consultant if and only if the manager’s expertise level is above the threshold. Thus, only a less-informed manager would take both signals into account to take an even more accurate action than the more-informed consultant.

⁸See Arieli et al. (2018) for more details.

In addition, in Proposition 3 we compare two scenarios: one in which the manager is better-informed than the consultant, and the reverse possibility. Our analysis shows that, in terms of correctness, it is better to have a *less*-informed manager and a *better*-informed consultant. We refer to this result as *the intern effect* because it accommodates a superior learning process in equilibrium, compared to the alternative ordering of expertise levels/players.

3.1 The case of symmetric payoffs

Our preliminary analysis deals with the symmetric case where the players' individual payoffs offset when they choose different actions, namely, $U_1 = -U_4 \equiv \alpha > 1$. The analysis is divided into two parts: in Theorem 1 we depict sufficient and necessary conditions for the previously mentioned equilibria to arise, and in Proposition 1 we study the correctness of the process under each of these equilibria.

Starting with Theorem 1, we divide the *entire* parameter-space (p, q_c, q_m, α) into four disjoint parts, where each supports a different equilibrium type. Generally speaking, it shows that a Herding equilibrium arises if and only if the consultant is significantly more accurate than manager, whereas a significantly superior manager yields a Revealing equilibrium. On the other hand, the more sophisticated equilibria (the Guided and Mixed ones) emerge when neither of the two is significantly more accurate than the other.

Theorem 1. *Consider the previously defined game G .*

1. *There exists a Herding equilibrium if and only if $\tilde{q}_c \geq \tilde{q}_m + \tilde{p}$.*
2. *There exists a Revealing equilibrium if and only if $\tilde{q}_m \geq \tilde{q}_c + \tilde{p}$.*
3. *There exists a Guided equilibrium if and only if $\tilde{q}_i \leq \tilde{q}_{-i} + \tilde{p}$ for every $i \in \{c, m\}$, and*

$$\tilde{q}_c \geq \tilde{p} + \ln \left(\frac{2 \exp(-\tilde{q}_m) + \alpha + 1}{(\alpha + 1) \exp(-\tilde{q}_m) + 2} \right). \quad (1)$$

4. *There exists a Mixed equilibrium if and only if $\tilde{q}_i \leq \tilde{q}_{-i} + \tilde{p}$ for every $i \in \{c, m\}$, and Ineq. (1) does not hold.*

Though all proofs are deferred to the appendix, let us point out that, in order to avoid repetitions, the proof of Theorem 1 builds on the general set-up of the game. Hence, it partially supports subsequent results in Section 3.2.

Figure 2 provides a visualization of the results of Theorem 1. It shows the disjoint equilibrium regions in the (q_c, q_m) -plane, for fixed α and p . One can see that the two regions of the Herding and Revealing equilibria are rather straightforward – a significantly higher expertise level of one over the other. On the other hand, the regions of the Guided and Mixed equilibria are more intriguing. First, the Guided and Mixed equilibria arise when neither has a clear dominance over the other in terms of expertise level. This leads the manager to rely on the consultant’s truthful action in borderline situations, namely when the manager’s signal is $s_m = 1$ and this contradicts the consultant’s recommendation of $a_c = 0$ and the prior (which is biased towards $\theta = 0$). Second, the distinction between a Guided equilibrium and a Mixed one, as given in Equation (1) and presented in Figure 2, is based on the consultant’s relative expertise level, q_c . When it is rather close to the prior p , the high uncertainty leads both players to hedge between their two available actions, i.e., a Mixed equilibrium.

Let us provide some intuition to the Guided equilibrium outcome. This equilibrium is reached whenever the expertise levels of both players are rather close to one another, as evident from the condition $\tilde{q}_i \leq \tilde{q}_{-i} + \tilde{p}$, for every $i \in \{c, m\}$. Thus, when the manager receives a signal that matches the recommendation of the consultant, it is optimal for him to recommend it accordingly. Otherwise, he resorts to the prior which is biased towards a $a_m = 0$ recommendation. Note that the stated proximity condition is independent of α , and specifically independent of the fact that $\alpha > 1$ (i.e., it depends solely on q_c, q_m , and p).

Though Theorem 1 provides a complete mapping from any set of parameters to an equilibrium, one may wonder whether other equilibria exist. The answer to this question is yes, but these equilibria are quite null. For example, the profile in which both players always recommend $a_i = 0$, irrespective of their signals, is an equilibrium, if and only if $q_c = q_m = p$. We generally disregard such examples because they do not provide meaningful insights into the strategic interaction between the two. On the other hand, it is easy to verify that non-informative strate-

Equilibria regions in the (q_c, q_m) -plane

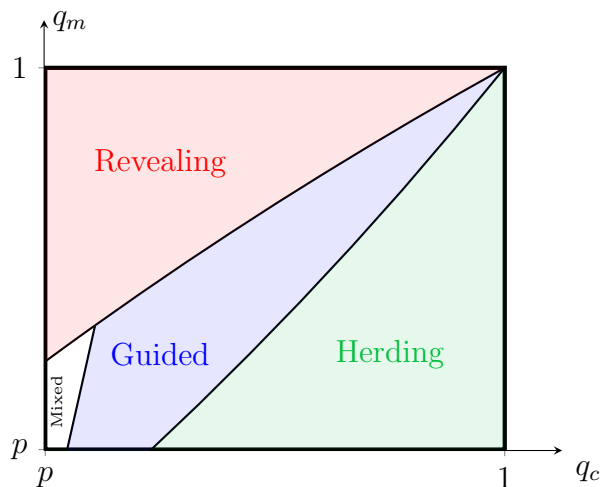


Figure 2: The different equilibria regions in the (q_1, q_2) -plane. Given $p = 0.6$ and $\alpha = 2$, each region corresponds to exactly one of the four possible equilibria: Revealing (red), Guided (blue), Herding (green), Mixed (white).

gies, i.e., providing the same recommendation independently of the signals, do not consist an equilibrium, unless $q_c = q_m = p$. This originates from the combination of informative signals and symmetric payoffs so that being correct strictly dominates being incorrect, assuming that the other player is uninformative.

Building on the results of Theorem 1, we can now provide some insights into the correctness of this decision process under the different equilibria. The following proposition shows that a Guided equilibrium, specifically, provides a greater accuracy compared to what each player can achieve by himself, whereas the other three are limited to the players' individual levels. This improvement originates from the fact that a Guided equilibrium supports a form of *strategic learning*, in equilibrium, which strictly increases their aggregate accuracy level.

Proposition 1. *Given either a Herding or a Revealing equilibrium, the correctness of the process is $C(q_c, q_m) = \max\{q_c, q_m\}$, and given a Mixed equilibrium the correctness is $C(q_c, q_m) = q_m$. However, in case there exists only a Guided equilibrium, the correctness of the process is $C(q_c, q_m) > \max\{q_c, q_m\}$.*

The two straightforward cases in the analysis of Proposition 1 are the Herding and Revealing equilibria, that yield, by definition, a correctness of q_c and q_m , respectively. To compare, a Guided equilibrium generates strictly higher correctness compared to $\max\{q_c, q_m\}$ because it supports a learning process between the manager and the consultant. Figure 3 illustrates this result as a function of the manager's expertise level, for a fixed q_c . It is important to note that these results, and the notion of correctness in general, are a function of the strategies profile, i.e., the equilibrium type. So, from a mechanism-design perspective, the shareholders/firm can only benefit from supporting a Guided equilibrium, when feasible, relative to the other equilibria types.

Interestingly, though the Mixed equilibrium supports a form of learning, similar to the Guided equilibrium, its correctness remains limited to q_m . This originates from the indifference that both players generate, in equilibrium, by hedging between their two possible actions. So although some learning can take place, it is concealed by the inherent disinformation that the consultant generates through his randomization.




Figure 3: The correctness of the process as a function of the manager's expertise level, and a fixed level of $q_c = 0.8$ for the consultant, and $(p, \alpha) = (0.6, 2)$. A Mixed equilibrium does not exist under these parameters (see Fig. 2). The dashed vertical lines divide the axis into the different equilibria regions (from left to right): Herding, Guided, and Revealing. Notably, the Guided regime supports a correctness level that individually supersedes the expertise levels of both players.

3.2 The case of asymmetric payoffs

Using the basic features of the second-opinion problem and its possible solutions, as given in Section 3.1, we turn to study the general set-up in which the experts' payoffs do not necessarily offset when they provide different assessments, i.e., when $U_1 \neq -U_4$. Still, without loss of generality, the normalization of $U_2 = 1$ and $U_3 = -1$ remains as before, and the new payoffs are summarized in Table 2. Hence, $U_1 + 1$ is the potential gain from non-conformity (being

		Manager	
		$a_m = \theta$	$a_m = 1 - \theta$
Consultant	$a_c = \theta$	$(1, 1)$	(U_1, U_4)
	$a_c = 1 - \theta$	(U_4, U_1)	$(-1, -1)$

Table 2: The payoff matrix given a realized state of $\theta \in \Theta$ with asymmetric payoffs.

correct alone versus being incorrect with the other expert), and $1 - U_4$ is the potential gain from conformity (being correct with the other expert versus being incorrect alone). We define the *non-conformity gain ratio* as the ratio between these two numbers, and denote it by $\gamma = \frac{U_1 - (-1)}{1 - U_4}$. To be inline with the logit representation of probabilities, we also denote $\tilde{\gamma} = \ln \gamma$.

In general, we use the non-conformity gain ratio γ to study how the equilibria of the game evolve, as a function of the payoffs. Though one can perform an analysis for every $\tilde{\gamma}$ (and for every U_1 and U_4), we limit the discussion to the range $\tilde{\gamma} \leq \tilde{q}_c + \tilde{q}_m - \tilde{p}$. The reason is that when $\tilde{\gamma}$ is too large, it becomes the sole driving force behind the equilibria and subsequent results, irrespective of private and public information. Specifically, suppose that both signals are $s_c = s_m = 1$, and that the consultant truthfully reports $a_c = 1$. This is the highest possible posterior on the event $\{\theta = 1\}$. Still, if the following inequality holds

$$(1 - p)q_cq_m \cdot 1 + p(1 - q_c)(1 - q_m) \cdot (-1) \leq (1 - p)q_cq_m \cdot U_4 + p(1 - q_c)(1 - q_m) \cdot U_1,$$

then the manager benefits from “gambling” on the low-probability event $\{\theta = 0 | s_c = s_m = 1\}$, simply because it opposes the action of the consultant. The last inequality is indeed $\tilde{q}_c + \tilde{q}_m - \tilde{p} \leq \tilde{\gamma}$.

We start with a simple generalization of previous results. The following theorem extends Theorem 1 to asymmetric payoffs, that is, when $\tilde{\gamma} \neq 0$. Note that for the case of a Guided equilibrium, we impose an additional condition that the consultant reveals his private signal in equilibrium. This allows us to avoid the differentiation between a Guided and a Mixed equilibrium, which does not provide further meaningful insights to the analysis in this section.

Theorem 1*. *Consider the previously defined game G with fixed parameters p, q_c, q_m , and γ .*

- *There exists a Herding equilibrium if and only if $\tilde{\gamma} < \tilde{q}_c - \tilde{q}_m - \tilde{p}$.*
- *If the consultant reveals his private signal in equilibrium and $\tilde{q}_c - \tilde{q}_m - \tilde{p} < \tilde{\gamma} < \tilde{q}_c - \tilde{q}_m + \tilde{p}$, then the unique equilibrium is a Guided one.*
- *There exists a Revealing equilibrium if and only if $\tilde{q}_c - \tilde{q}_m + \tilde{p} < \tilde{\gamma} < \tilde{q}_c + \tilde{q}_m - \tilde{p}$.*

The ability to vary the payoffs of the game, through the non-conformity gain ratio γ , shows how the transition from a Herding equilibrium to a Revealing one, passes through a Guided equilibrium. Figure 4 illustrates this transition through the different intervals, with respect to $\tilde{\gamma}$, for each of the mentioned equilibria.




Figure 4: The transitions between equilibria regimes as a function of γ .

The following Proposition 2, which builds on Theorem 1*, presents one of the main insights of this paper. It shows that a better-informed manager does not necessarily improve the overall correctness of the process. In fact, one can find two disjoint intervals such that *every* expertise level of the manager in the lower interval generates strictly higher correctness than *every* expertise level in the higher interval. The limit value that separates these two intervals is referred to as *the humility threshold* – the highest level that allows for a strategic learning process in equilibrium.

Proposition 2 (The humility threshold). Fix (q_c, γ, p) such that $\tilde{q}_c > \tilde{\gamma} > \tilde{p}$ and assume that the consultant reveals his private signal in equilibrium. Then, there exist $\bar{q}_m > \underline{q}_m$ such that $C(q_c, \underline{q}_m) > C(q_c, \bar{q}_m)$. Moreover, there exists q_m^* such that for every $q_m \in (\underline{q}_m, q_m^*)$ and every $q'_m \in (q_m^*, \bar{q}_m)$, the correctness is higher when the manager has lower quality signal, i.e. $C(q_c, q_m) > C(q_c, q'_m)$.

Three clarifications are in order. First, the preliminary assumption that $\tilde{q}_c > \tilde{\gamma} > \tilde{p}$ enables us to shift from a Guided equilibrium to a Revealing one, by varying the expertise level q_m of the manager. Otherwise, e.g., in case $\tilde{\gamma}$ is significantly larger than $\tilde{q}_c + \tilde{p}$, we are left only with a Revealing equilibrium, independently of q_m . Second, this transition (between equilibria) hinges on the fact $\tilde{\gamma} > \tilde{p} > 0$, which implies that $U_1 < -U_4$. In other words, the positive externalities (among players) are small compared to the negative ones, so the manager is more inclined to follow the consultant further along the guided equilibrium, instead of uniquely following his own private signal. Third, note that the potential loss from crossing the humility threshold q_m^* is not necessarily a mild one. Figure 5 illustrates the potential magnitude of this non-monotone effect, where an infinitesimal increase in q_m triggers a drop from 0.81 to 0.75 in the overall correctness value.

The next proposition deals with the question of who should be better-informed: the manager or the consultant? That is, we consider two opposite scenarios in which either the manager or the consultant are the better-informed player. Proposition 3 shows that, in term of correctness, it might be better to have a manager with a lower expertise level rather than such a consultant. We refer to this result as *the intern effect* since one can think of the manager as a player who typically takes into account previous recommendations. The alternative set-up, in which the manager is relatively better-informed, limits the learning process among players since a better-informed manager does not feel the need to consider previous assessments. Notably, this entire process occurs in equilibrium, so the disregard for less-informed opinions is not a behavioral artifact such as arrogance or recklessness, but the *optimal strategic reaction* of a well-informed rational player.

Proposition 3 (The intern effect). Assume $\tilde{\gamma} > \tilde{p}$ and fix $q^H > q^L$ such that $\tilde{q}^H - \tilde{q}^L < \tilde{p}$. In

Correctness as a function of q_m

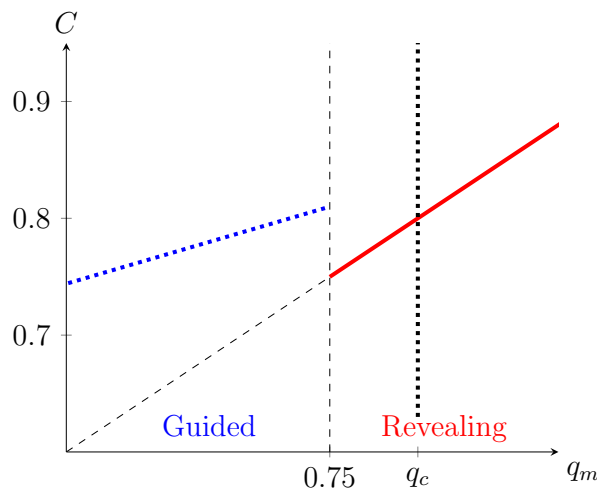


Figure 5: The correctness of the process as a function of the manager’s expertise level, given $p = 0.6$, $q_c = 0.8$, and $\gamma = 2$. The dotted (blue) line describes the correctness under a Guided equilibrium, and the solid (red) line describes the correctness under a Revealing equilibrium. Note that the Humility threshold is below q_c .

case $(q_c, q_m) = (q^H, q^L)$ yields a Guided equilibrium with correctness $C(q^H, q^L)$, then reversing the expertise levels to $(q_c, q_m) = (q^L, q^H)$ would generate a Revealing equilibrium with a lower correctness of $C(q^L, q^H) < C(q^H, q^L)$. Hence, the higher correctness is obtained when the manager is given the lower quality signal.

To get some intuition for this result, consider the signals that both players receive. Once the manager is better-informed, it is sub-optimal for him to deviate from his relatively high-accuracy signal towards the consultant’s less-accurate assessment. In practice, the ordering $(q_c, q_m) = (q^L, q^H)$ yields a Revealing equilibrium, which implies that the manager’s decision is based entirely on his own (superior) signal. In case the relative accuracy reverses, the rational less-informed manager takes into account all available information, and reverts to the consultant’s recommendation if it is also supported by the prior.

A simple and important way to extend Proposition 3 is by examining a considerably better-informed player, i.e., a significantly large q^H relative to q^L . In case q^H is indeed sufficiently large, then approaching a consultant with an expertise level of q^H would yield either a Herding equilibrium with a correctness of q^H , or a Guided one, with a correctness that exceeds q^H .

To compare, switching the relative expertise levels would yield a Revealing equilibrium with a correctness of q^H . So, it becomes weakly better to approach a highly informed consultant, even if a Guided equilibrium is not always achievable.

From a practical perspective, revealing the players' expertise levels is anything but trivial. This problem originates from the fact that the manager sees the recommendation of the consultant in each of the mentioned equilibria in Proposition 3, and this enables him to maintain a higher correctness level, as an individual player, relative to the consultant. As an example, one can think of the structural advantage that top executives maintain by their ability to typically conclude meetings, and how this procedure supports the image of well-thought-out individuals. Yet, in some cases, the correctness of the manager, as an individual player, remains well below the correctness of the consultant. This occurs in equilibrium and although the manager observes the recommendation of the consultant, and a mimicking strategy is indeed feasible, but sub-optimal. We derive this conclusion from the following observation.⁹

Observation 1. *Fix p, q_c, q_m , and assume that the consultant reveals his private signal in equilibrium. Then, the correctness $C(q_c, q_m)$ in equilibrium is not a monotone function of γ :*

- *under a Herding equilibrium, i.e., $\tilde{\gamma} < \tilde{q}_c - \tilde{q}_m - \tilde{p}$, the correctness is q_c ;*
- *under a Guided equilibrium, i.e., $\tilde{q}_c - \tilde{q}_m - \tilde{p} < \tilde{\gamma} < \tilde{q}_c - \tilde{q}_m + \tilde{p}$, the correctness is $C(q_c, q_m) > \min\{q_c, q_m\}$; and*
- *under a Revealing equilibrium, i.e., $\tilde{q}_c - \tilde{q}_m + \tilde{p} < \tilde{\gamma} < \tilde{q}_c + \tilde{q}_m - \tilde{p}$, the correctness is $C(q_c, q_m) = q_m$.*

Before we conclude, let us point to one striking difference between Observation 1 and Proposition 1 regarding the case of a Revealing equilibrium. The correctness under a Revealing equilibrium is q_m , by definition. Under symmetric payoffs (i.e., $\tilde{\gamma} = 0$), a Revealing equilibrium implies that $q_m > q_c$, whereas under asymmetric payoffs (i.e., $\tilde{\gamma} \neq 0$), a Revealing equilibrium exists even if $q_m < q_c$, depending on \tilde{p} and $\tilde{\gamma}$. Hence, from a mechanism-design perspective, a sub-optimal design of incentives (through γ) can cause the correctness to drop to the minimal

⁹The proof follows directly from Theorem 1* and Proposition 1, thus omitted.

possible value, whereas an optimal design of γ can facilitate either a Guided, or a Herding equilibrium, to achieve the maximal possible correctness level. As evident from Observation 1, the ability to fix γ is quite robust since the equilibria are supported on a wide range of γ values.

4 In conclusion

This paper provides an analysis of a sequential, two-player, Bayesian game, in which one equilibrium supports a higher level of strategic learning, than all other equilibria. Our analysis depicts a mapping from the expertise levels to the equilibria of this game, showing that: (i) better-informed players may generate worse recommendations in equilibrium; and (ii) ordering the players so that the lower-level one provides the second opinion can typically improve the outcome. This contraries the common belief that the high-skill player should be the one to provide the second opinion and that a better-informed manager is better.

We note that since the shareholders can (indirectly) influence the manager's liability in case of an error, they can ensure an equilibrium in which the correctness is maximized, which is optimal for them. This optimal liability is generally robust and remains the same, irrespective of either small changes in the manager's and consultant's quality, or the prior probability of the state of the world. Thus, although the shareholders are typically uninformed of the details of the decision processes in the company, they can still facilitate cooperation between the manager and the consultant and optimize the probability of choosing the correct decision.

A Proofs

Proof of Theorem 1. To avoid repetitions, we start our analysis with the more general payoffs given in Table 2. Recall $\tilde{\gamma} = \ln(\gamma) = \ln\left[\frac{U_1+1}{1-U_4}\right]$, whereas the payoffs in Table 1 yield $\tilde{\gamma} = 0$.

Fix q_c, q_m, p , and γ . Assume that the consultant's recommendation matches his signal, as in all the equilibria stated in the Theorem, other than the Mixed equilibrium. Our analysis is divided into four different states depending on a_c and s_m .

Consider first the case where $a_c = 0$ and $s_m = 0$. This state occurs with probability $q_c q_m$ when $\theta = 0$, and with probability $(1 - q_c)(1 - q_m)$ when $\theta = 1$. The manager's best response in this state is $a_m = 0$ if and only if

$$pq_c q_m \cdot 1 + (1 - p)(1 - q_c)(1 - q_m) \cdot (-1) \geq pq_c q_m \cdot U_4 + (1 - p)(1 - q_c)(1 - q_m) \cdot U_1,$$

(the best response is $a_m = 1$ when the inequality is reversed). The last inequality can be rearranged into

$$\frac{p}{1 - p} \cdot \frac{q_c}{1 - q_c} \cdot \frac{q_m}{1 - q_m} \cdot \frac{1 - U_4}{1 + U_1} \geq 1,$$

or, equivalently,

$$\tilde{p} + \tilde{q}_c + \tilde{q}_m - \tilde{\gamma} \geq 0. \quad (2)$$

Next, consider the case where $a_c = 0$ and $s_m = 1$. This state occurs with probability $q_c(1 - q_m)$ when $\theta = 0$, and with probability $(1 - q_c)q_m$ when $\theta = 1$. The manager's best response in this state is $a_m = 0$ if and only if

$$pq_c(1 - q_m) \cdot 1 + (1 - p)(1 - q_c)q_m \cdot (-1) \geq pq_c(1 - q_m) \cdot U_4 + (1 - p)(1 - q_c)q_m \cdot U_1,$$

which can be rearranged into

$$\tilde{p} + \tilde{q}_c - \tilde{q}_m - \tilde{\gamma} \geq 0. \quad (3)$$

Third, consider the case where $a_c = 1$ and $s_m = 0$. This state occurs with probability $(1 - q_c)q_m$ when $\theta = 0$, and with probability $q_c(1 - q_m)$ when $\theta = 1$. The manager's best response in this state is $a_m = 0$ if and only if

$$p(1 - q_c)q_m \cdot U_1 + (1 - p)q_c(1 - q_m) \cdot U_4 \geq p(1 - q_c)q_m \cdot (-1) + (1 - p)q_c(1 - q_m) \cdot 1,$$

which can be rearranged into

$$\tilde{p} - \tilde{q}_c + \tilde{q}_m + \tilde{\gamma} \geq 0. \quad (4)$$

Lastly, consider the case where $a_c = 1$ and $s_m = 1$. This state occurs with probability $(1 - q_c)(1 - q_m)$ when $\theta = 0$, and with probability $q_c q_m$ when $\theta = 1$. The manager's best response in this state is $a_m = 0$ if and only if

$$p(1 - q_c)(1 - q_m) \cdot U_1 + (1 - p)q_c q_m \cdot U_4 \geq p(1 - q_c)(1 - q_m) \cdot (-1) + (1 - p)q_c q_m \cdot 1,$$

which can be rearranged into

$$\tilde{p} - \tilde{q}_c - \tilde{q}_m + \tilde{\gamma} \geq 0. \quad (5)$$

In a **Revealing equilibrium**, $a_m(\cdot, s_m) = s_m$, so Ineq. (2) and Ineq. (4) should hold while Ineq. (3) and Ineq. (5) are reversed:

$$\tilde{p} + \tilde{q}_c + \tilde{q}_m - \tilde{\gamma} \geq 0, \quad (6a)$$

$$\tilde{p} + \tilde{q}_c - \tilde{q}_m - \tilde{\gamma} \leq 0, \quad (6b)$$

$$\tilde{p} - \tilde{q}_c + \tilde{q}_m + \tilde{\gamma} \geq 0, \quad (6c)$$

$$\tilde{p} - \tilde{q}_c - \tilde{q}_m + \tilde{\gamma} \leq 0. \quad (6d)$$

Note that Ineq. (6d) implies that Ineq. (6a) holds, and Ineq. (6b) yields Ineq. (6c). To simplify the analysis, we now revert to the conditions of Theorem 1 by taking $\tilde{\gamma} = 0$. Moreover, recall that $\min\{\tilde{q}_c, \tilde{q}_m\} \geq \tilde{p}$, so the aforementioned inequalities reduce to Ineq. (6b) with $\tilde{\gamma} = 0$, i.e., $\tilde{q}_c + \tilde{p} \leq \tilde{q}_m$, as stated in the theorem.

Hence, if Ineq. (6b) with $\tilde{\gamma} = 0$ holds, whenever the consultant reports his signal, the best response of the manager is to report his signal as well. Clearly, the converse is true. Suppose that the manager reports his signal, so the action of the consultant has no effect on the action of the manager. The expected payoff of being correct is positive while the expected payoff of being incorrect is negative, so it is better to report the signal.

In a **Herding equilibrium**, $a_m(a_c, \cdot) = a_c$, so Ineq. (2) and Ineq. (3) should hold while Ineq. (4) and Ineq. (5) are reversed:

$$\tilde{p} + \tilde{q}_c + \tilde{q}_m - \tilde{\gamma} \geq 0, \quad (7a)$$

$$\tilde{p} + \tilde{q}_c - \tilde{q}_m - \tilde{\gamma} \geq 0, \quad (7b)$$

$$\tilde{p} - \tilde{q}_c + \tilde{q}_m + \tilde{\gamma} \leq 0, \quad (7c)$$

$$\tilde{p} - \tilde{q}_c - \tilde{q}_m + \tilde{\gamma} \leq 0. \quad (7d)$$

Note that Ineq. (7d) implies that Ineq. (7a) holds, and Ineq. (7c) yields Ineq. (7b) and Ineq. (7d). We now revert to the conditions of Theorem 1 by taking $\tilde{\gamma} = 0$, so the aforementioned inequalities reduce to Ineq. (7c) with $\tilde{\gamma} = 0$, i.e., $\tilde{q}_c \geq \tilde{q}_m + \tilde{p}$, as stated in the theorem.

Hence, if Ineq. (7c) with $\tilde{\gamma} = 0$ holds, whenever the consultant reports his signal, the best response of the manager is to ignore his signal and report a_c . Clearly, the converse is true. Suppose that the manager repeats the action of the consultant. The consultant receives 1 if he is correct and -1 if he is wrong, so he should report his signal which has a higher probability to match the true state of the world.

In a **Guided equilibrium**, $a_m(a_c, s_m) = s_m$ except for $a_m(0, 1) = 0$ so Ineq. (2), Ineq. (3) and Ineq. (4) should hold while Ineq. (5) is reversed:

$$\tilde{p} + \tilde{q}_c + \tilde{q}_m - \tilde{\gamma} \geq 0, \quad (8a)$$

$$\tilde{p} + \tilde{q}_c - \tilde{q}_m - \tilde{\gamma} \geq 0, \quad (8b)$$

$$\tilde{p} - \tilde{q}_c + \tilde{q}_m + \tilde{\gamma} \geq 0, \quad (8c)$$

$$\tilde{p} - \tilde{q}_c - \tilde{q}_m + \tilde{\gamma} \leq 0. \quad (8d)$$

Note that Ineq. (8d) implies that Ineq. (8a) holds, and Ineqs. (8c) and (8b) are equivalent to $\tilde{\gamma} - \tilde{p} \leq \tilde{q}_c - \tilde{q}_m \leq \tilde{\gamma} + \tilde{p}$. Since $q_m \geq p$, one can show that Ineq. (8b) yields Ineq. (8d) as follows,

$$\tilde{q}_c + \tilde{q}_m \geq \tilde{q}_c + \tilde{p} \geq \tilde{\gamma} + \tilde{q}_m \geq \tilde{\gamma} + \tilde{p}.$$

So, one only needs to sustain the inequalities $\tilde{\gamma} - \tilde{p} \leq \tilde{q}_c - \tilde{q}_m \leq \tilde{\gamma} + \tilde{p}$. Again, take $\tilde{\gamma} = 0$, and the manager follows the Guided equilibrium given that $-\tilde{p} \leq \tilde{q}_c - \tilde{q}_m \leq \tilde{p}$, as stated in the theorem.

Suppose that the consultant observes $s_c = 0$. By reporting $a_c = 0$, he guides the manager to choose $a_m = 0$. By reporting $a_c = 1$, the manager would act according to his own signal. Reporting $a_c = 0$ is better if and only if

$$pq_c \cdot 1 + (1-p)(1-q_c) \cdot (-1) \geq pq_c q_m \cdot (-\alpha) + pq_c(1-q_m) \cdot (-1) + (1-p)(1-q_c)q_m \cdot 1 + (1-p)(1-q_c)(1-q_m) \cdot \alpha,$$

which is equivalent to

$$\frac{p}{1-p} \cdot \frac{q_c}{1-q_c} \geq \frac{1+q_m+\alpha-\alpha q_m}{1+q_m\alpha+1-q_m}.$$

The last inequality holds since the LHS is bounded from below by 1 (recall that $p, q_c \geq 0.5$), while the RHS is bounded from above by 1 (since $q_m \geq 0.5$ and $\alpha \geq 1$).

Now suppose that the consultant observes $s_c = 1$. Reporting $a_c = 1$ is better if and only if

$$p(1-q_c)q_m \cdot (-\alpha) + p(1-q_c)(1-q_m) \cdot (-1) + (1-p)q_c(1-q_m) \cdot \alpha + (1-p)q_cq_m \cdot 1 \geq p(1-q_c) \cdot 1 + (1-p)q_c \cdot (-1),$$

which is equivalent to

$$\frac{q_c}{1-q_c} \cdot \frac{1-p}{p} \geq \frac{2+q_m(\alpha-1)}{1+\alpha+q_m(1-\alpha)},$$

or,

$$\tilde{q}_c - \tilde{p} \geq \ln \left(\frac{2 \exp(-\tilde{q}_m) + \alpha + 1}{(\alpha + 1) \exp(-\tilde{q}_m) + 2} \right),$$

as stated in the theorem.

In a **Mixed equilibrium**, the signal $s_c = 0$ yields $a_c = 0$, while $s_c = 1$ leads to a mixed action $(r, 1-r)$ of the consultant. Similarly, the manager plays according to his own signal, $a_m = s_m$, unless $(a_c, s_m) = (0, 1)$, a case in which he plays a mixed action $(\rho, 1-\rho)$.

We begin with an analysis of the manager's actions. To shorten the exposition, we depict the conditions that follow Ineqs. (2)-(5), adjusted to the Mixed-equilibrium strategies.

- If $(a_c, s_m) = (0, 0)$, then $a_m = 0$ if and only if

$$pq_m[q_c+r(1-q_c)] - (1-p)(1-q_m)[q_cr+(1-q_c)] \geq (1-p)(1-q_m)\alpha[q_cr+(1-q_c)] - pq_m\alpha[q_c+r(1-q_c)],$$

which reduces to $\frac{p}{1-p} \cdot \frac{q_m}{1-q_m} \geq \frac{q_cr+1-q_c}{q_c+r(1-q_c)}$, and the last inequality evidently holds.

- If $(a_c, s_m) = (1, 0)$, then $a_m = 0$ if and only if

$$\alpha p(1-q_c)q_m(1-r) - \alpha(1-p)q_c(1-q_m)(1-r) \geq (1-p)q_c(1-q_m)(1-r) - p(1-q_c)q_m(1-r),$$

which reduces to $\tilde{q}_c - \tilde{q}_m \leq \tilde{p}$, and stated in the theorem.

- If $(a_c, s_m) = (1, 1)$, then $a_m = 1$ if and only if

$$(1-p)q_cq_m(1-r) - p(1-q_c)(1-q_m)(1-r) \geq \alpha p(1-q_c)(1-q_m)(1-r) - \alpha(1-p)q_cq_m(1-r),$$

which reduces to $\tilde{q}_c + \tilde{q}_m \geq \tilde{p}$.

- If $(a_c, s_m) = (0, 1)$, then the manager plays a mixed action if and only if

$$p(1-q_m)[q_c+(1-q_c)r]-(1-p)q_m[q_cr+(1-q_c)] = \alpha(1-p)q_m[q_cr+(1-q_c)]-\alpha p(1-q_m)[q_c+r(1-q_m)],$$

which reduces to $r = \frac{(1-p)(1-q_c)q_m-pq_c(1-q_m)}{p(1-q_c)(1-q_m)-(1-p)q_cq_m}$. Note that $r \in [0, 1]$ as long as $\tilde{p} \leq \tilde{q}_c + \tilde{q}_m$ and $-\tilde{p} \leq \tilde{q}_c - \tilde{q}_m$, as stated in the theorem.

Moving on to the actions of the consultant, if $s_c = 0$ then $a_c = 0$ if and only if

$$pq_c[1 + \alpha q_m + (1 - q_m)(\rho + (1 - \rho)\alpha)] \geq (1 - p)(1 - q_c)[1 + \alpha(1 - q_m) + q_m(\rho + (1 - \rho)\alpha)],$$

which reduces to $\frac{p}{1-p} \cdot \frac{q_c}{1-q_c} \geq \frac{1+\alpha+q_m\rho(1-\alpha)}{1+\alpha\rho(1-q_m)(1-\alpha)}$, and the last inequality holds for any $\alpha \geq 1$ and $\rho \in [0, 1]$. In addition, if $s_c = 1$, then the consultant is indifferent between $a_c = 0$ and $a_c = 1$ if and only if

$$p(1 - q_c)[1 + \alpha + (1 - q_m)\rho(1 - \alpha)] = (1 - p)q_c[1 + \alpha + q_m\rho(1 - \alpha)],$$

which yields $\rho = \frac{1+\alpha}{\alpha-1} \cdot \frac{(1-p)q_c-p(1-q_c)}{(1-p)q_cq_m-p(1-q_c)(1-q_m)}$. Note that $\rho \in [0, 1]$ as long as

$$\tilde{q}_c - \tilde{p} \leq \ln \left(\frac{2 \exp(-\tilde{q}_m) + \alpha + 1}{(\alpha + 1) \exp(-\tilde{q}_m) + 2} \right),$$

as stated in the theorem. This concludes our proof. \square

Proof of Proposition 1. By definition, the correctness under a Herding equilibrium is q_c , which equals $\max\{q_c, q_m\}$ since $\tilde{q}_c \geq \tilde{q}_m + \tilde{p}$. Similarly, the correctness under a Revealing equilibrium is $q_m = \max\{q_c, q_m\}$. So, let us consider a Guided equilibrium, in which $a_m = s_m$ in case $a_c = 1$, and $a_c = 0$ leads to $a_m = 0$. In other words, in a Guided equilibrium, the manager deviates from his private signal if and only if $(a_c, s_m) = (0, 1)$. Thus,

$$\begin{aligned} C(q_c, q_m) &= \Pr(\sigma_c = \theta | \sigma_c) \\ &= q_m + pq_c(1 - q_m) - (1 - p)(1 - q_c)q_m \\ &= pq_m + pq_c + q_cq_m - 2pq_cq_m, \end{aligned}$$

and the last expression is symmetric w.r.t. substituting q_c and q_m . So, w.l.o.g., assume that $q_m = \max\{q_c, q_m\}$. Theorem 1 states that a Guided equilibrium implies that $\tilde{p} + \tilde{q}_c \geq \tilde{q}_m$, or equivalently, $\frac{p}{1-p} \cdot \frac{q_c}{1-q_c} \geq \frac{q_m}{1-q_m}$, thus

$$C(q_c, q_m) = q_m + pq_c(1 - q_m) - (1 - p)(1 - q_c)q_m \geq q_m,$$

and the inequality is strict in case $\frac{p}{1-p} \cdot \frac{q_c}{1-q_c} > \frac{q_m}{1-q_m}$, as needed.

Given a Mixed equilibrium, the manager deviates from his private signal if and only if $(a_c, s_m) = (0, 1)$. Thus, the correctness is given by

$$\begin{aligned} C(q_c, q_m) &= q_m + p(1 - q_m)[q_c + (1 - q_c)r][\rho - (1 - \rho)] + (1 - p)q_m[q_cr + (1 - q_c)][(1 - \rho) - \rho] \\ &= q_m + (2\rho - 1)\{p(1 - q_m)[q_c + (1 - q_c)r] - (1 - p)q_m[q_cr + (1 - q_c)]\} = q_m, \end{aligned}$$

where the last inequality follows from the manager's indifference condition

$$\begin{aligned} (1 + \alpha)p(1 - q_m)[q_c + (1 - q_c)r] &= (1 + \alpha)(1 - p)q_m[q_cr + (1 - q_c)] \\ p(1 - q_m)[q_c + (1 - q_c)r] &= (1 - p)q_m[q_cr + (1 - q_c)] \end{aligned}$$

given in Theorem 1. □

Proof of Theorem 1*. Fix p, q_c, q_m , and let us follow the cases as given in the theorem.

Assume that $\tilde{\gamma} < \tilde{q}_c - \tilde{q}_m - \tilde{p}$. According to the proof of Theorem 1, a Herding equilibrium exists if and only if $\tilde{\gamma} + \tilde{p} < \tilde{q}_c - \tilde{q}_m$, which holds by assumption. Now, assuming that the manager replicates the action of the consultant. A straightforward computation would show that the optimal action of the consultant is to follow his signal as well, independently of $\tilde{\gamma}$. Thus, a Herding equilibrium exists if and only if $\tilde{\gamma} < \tilde{q}_c - \tilde{q}_m - \tilde{p}$.

Moving on to the case where $\tilde{q}_c - \tilde{q}_m - \tilde{p} < \tilde{\gamma} < \tilde{q}_c - \tilde{q}_m + \tilde{p}$. These inequalities are equivalent to Ineqs. (8b) and (8c), and according to the proof of Theorem 1, they imply that a Guided equilibrium exists. Since, by assumption, the consultant follows his own signal, the equilibrium is also unique.

Lastly, assume that $\tilde{q}_c - \tilde{q}_m + \tilde{p} < \tilde{\gamma} < \tilde{q}_c + \tilde{q}_m - \tilde{p}$. Using the proof of Theorem 1, a Revealing equilibrium exists if and only if Ineqs. (6b) and (6d) holds. Indeed, both inequalities hold

under the given parametric assumptions over $\tilde{\gamma}$. Moreover, similarly to the Herding analysis, the optimal action of the consultant is $a_c = s_c$, given that $a_m = s_m$, and independently of $\tilde{\gamma}$. This establishes the existence of a Revealing equilibrium and concludes our proof. \square

Proof of Proposition 2. Fix (q_c, γ, p) so that $\tilde{q}_c > \tilde{\gamma} > \tilde{p}$. We start by proving that there exist $q_m^* < q_c$ and $\epsilon_1 > 0$ such that for every $q_m \in (q_m^* - \epsilon_1, q_m^*)$ there exists a unique equilibrium, namely a Guided one, and for every $q_m \in (q_m^*, q_m^* + \epsilon_1)$ there exists a unique equilibrium – a Revealing one.

Take q_m^* such that $\tilde{\gamma} - \tilde{p} = \tilde{q}_c - \tilde{q}_m^*$. The condition $\tilde{q}_c > \tilde{\gamma} > \tilde{p}$ implies that $p < q_m^* < q_c$, and for every sufficiently close q_m to q_m^* (from below), we get $\tilde{\gamma} < \tilde{q}_c - \tilde{q}_m + \tilde{p}$, as needed for a Guided equilibrium according to Theorem 1*. In addition, since $q_m^* < q_c$, then for every sufficiently close q_m to q_m^* (from above), but still below q_c , we get that $\tilde{q}_c - \tilde{q}_m < \tilde{\gamma} - \tilde{p}$, as needed for a Revealing equilibrium according to Theorem 1*. Thus, one can fix $\epsilon_1 \in (0, \min\{q_c - q_m^*, q_m^* - p\})$, such that for every $q_m \in (q_m^* - \epsilon_1, q_m^*)$ there exists a Guided equilibrium, and for every $q_m \in (q_m^*, q_m^* + \epsilon_1)$ there exists a Revealing equilibrium. Note that the assumption that the consultant follows his own signal in equilibrium, along with the fact that the best response of the manager is unique in each of these cases, suggest that the mentioned equilibria are indeed unique.

Since $q_c > q_m^*$ and given $\epsilon_1 \in (0, \min\{q_c - q_m^*, q_m^* - p\})$, it follows that $q_m = \min\{q_c, q_m^*\}$ is the correctness of the process under the aforementioned Revealing equilibrium. This observation deviates from Proposition 1, as the latter specifically relates to cases where $\tilde{\gamma} = 0$, which entails that $q_m \geq q_c$ in a Revealing equilibrium. On the other hand, the correctness under a Guided equilibrium is

$$C(q_c, q_m) = q_m + pq_c(1 - q_m) - (1 - p)(1 - q_c)q_m.$$

Since $\tilde{p} + \tilde{q}_c - \tilde{q}_m^* = \tilde{\gamma} > \tilde{p}$, we deduce that $pq_c(1 - q_m^*) > \frac{p}{1-p}(1 - p)(1 - q_c)q_m^*$, which leads to

$$pq_c(1 - q_m^*) - (1 - p)(1 - q_c)q_m^* > \left(\frac{p}{1-p} - 1 \right) (1 - p)(1 - q_c)q_m^*.$$

Thus,

$$\begin{aligned}
q_m^* + pq_c(1 - q_m^*) - (1 - p)(1 - q_c)q_m^* &> q_m^*[1 + (2p - 1)(1 - q_c)] \\
&= q_m^*[q_c + 2p(1 - q_c)] \\
&> q_m^* + \epsilon_2
\end{aligned}$$

for every $0 < \epsilon_2 < q_m^*[q_c + 2p(1 - q_c) - 1]$. By continuity, one can take a positive and sufficiently small $\epsilon < \min\{\epsilon_1, \epsilon_2\}$ such that $C(q_c, q_m) = q_m + pq_c(1 - q_m) - (1 - p)(1 - q_c)q_m > q_m^* + \epsilon$, for every $q_m \in (q_m^* - \epsilon, q_m^*)$.

Denote $\bar{q}_m = q_m^* + \epsilon$ and $\underline{q}_m = q_m^* - \epsilon$. Thus, for every $q_m \in (\underline{q}_m, q_m^*)$, the correctness is strictly higher than for every $q'_m \in (q_m^*, \bar{q}_m)$. That is, $C(q_c, q_m) > C(q_c, q'_m)$, as stated in the theorem. \square

Proof of Proposition 3. Fix (γ, p, q^H, q^L) such that $q^H > q^L$, where $\tilde{q}^H - \tilde{q}^L < \tilde{p}$, and $(q_c, q_m) = (q^H, q^L)$ yields a Guided equilibrium with correctness

$$\begin{aligned}
C(q^H, q^L) &= q^L + pq^H(1 - q^L) - (1 - p)(1 - q^H)q^L \\
&= q^H + pq^L(1 - q^H) - (1 - p)(1 - q^L)q^H > q^H,
\end{aligned}$$

where the first equality follows from the symmetry of the expression w.r.t. q^H and q^L , and the inequality follows from the assumption that $\tilde{p} + \tilde{q}^L > \tilde{q}^H$.

Now consider $(q_c, q_m) = (q^L, q^H)$. Following the proof Theorem 1 (to establish that there exists a Revealing equilibrium), one needs to show that Ineq. (6b) and Ineq. (6d) hold. First, note that Ineq. (8d) and Ineq. (6d) are identical, independently of the ordering of (q_c, q_m) . So one only needs to show that Ineq. (6b) holds, i.e., $\tilde{q}^L - \tilde{q}^H \leq \tilde{\gamma} - \tilde{p}$. This inequality follows from the fact that $\tilde{\gamma} > \tilde{p}$ and $q^H > q^L$. Thus, we established that a Revealing equilibrium exists, and its correctness is $C(q^L, q^H) = q^H < C(q^H, q^L)$, as needed. \square

References

Ali, S. N. and N. Kartik (2012, November). Herding with collective preferences. *Economic Theory* 51(3), 601–626.

- Arieli, I. (2017). Payoff externalities and social learning. *Games and Economic Behavior* 104, 392–410. Publisher: Academic Press.
- Arieli, I., M. Koren, and R. Smorodinsky (2018). The one-shot crowdfunding game. *ACM EC 2018 - Proceedings of the 2018 ACM Conference on Economics and Computation* 18, 213–214. arXiv: 1805.11872 Publisher: Association for Computing Machinery, Inc ISBN: 9781450358293.
- Athey, S. (2001). Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information. *Econometrica* 69(4), 861–889. Publisher: [Wiley, Econometric Society].
- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics* 107(3), 797–817. Publisher: Oxford Academic.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* 100(5), 992–1026. Publisher: University of Chicago Press.
- Bikhchandani, S., D. A. Hirshleifer, O. Tamuz, and I. Welch (2021). Information Cascades and Social Learning. *SSRN Electronic Journal*.
- Blackwell, D. (1951). Comparison of Experiments. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 93–102.
- Blackwell, D. (1953). Equivalent Comparisons of Experiments. *Annals of Mathematical Statistics* 24(2), 265–272. Publisher: Institute of Mathematical Statistics.
- Crapis, D., B. Ifrach, C. Maglaras, and M. Scarsini (2017, November). Monopoly Pricing in the Presence of Social Learning. *Management Science* 63(11), 3586–3608. Publisher: INFORMS.
- Dasgupta, A. (2000). Social learning with payoff complementarities. *Working paper* (November 1999).

- De Clippel, G., K. Eliaz, D. Fershtman, and K. Rozen (2021). On selecting the right agent. *Theoretical Economics* 16(2), 381–402.
- Debo, L. G., C. Parlour, and U. Rajan (2011). Signaling Quality via Queues. *Management Science* 58(5), 876–891. Publisher: INFORMS.
- Ely, J. C. and J. Välimäki (2003, August). Bad Reputation. *Quarterly Journal of Economics* 118(3), 785–814.
- Eyster, E., A. Galeotti, N. Kartik, and M. Rabin (2014). Congested observational learning. *Games and Economic Behavior* 87, 519–538. Publisher: Academic Press.
- Feldman, P., Y. Papanastasiou, and E. Segev (2019, April). Social Learning and the Design of New Experience Goods. *Management Science* 65(4), 1502–1519. Publisher: INFORMS.
- Jensen, M. K. (2018, January). Distributional Comparative Statics. *Review of Economic Studies* 85(1), 581–610. Publisher: Oxford Academic.
- Lagziel, D. and E. Lehrer (2019). A Bias of Screening. *American Economic Review: Insights* 1(3), 343–356.
- Lagziel, D. and E. Lehrer (2022, November). Screening Dominance: A Comparison of Noisy Signals. *American Economic Journal: Microeconomics* 14(4), 1–24.
- Lehmann, E. L. (1988, June). Comparing Location Experiments. *Annals of Statistics* 16(2), 521–533. Publisher: Institute of Mathematical Statistics.
- Li, F. and P. Norman (2021). Sequential persuasion. *Theoretical Economics* 16(2), 639–675.
- Lipnowski, E. and D. Ravid (2020). Cheap Talk With Transparent Motives. *Econometrica* 88(4), 1631–1660. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15674](https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA15674).
- Mekonnen, T. and R. L. Vizcaíno (2022, January). Bayesian comparative statics. *Theoretical Economics* 17(1), 219–251. Publisher: John Wiley & Sons, Ltd.

- Mossel, E., M. Mueller-Frank, A. Sly, and O. Tamuz (2020). Social Learning Equilibria. *Econometrica* 88(3), 1235–1267. arXiv: 1207.5895 Publisher: John Wiley & Sons, Ltd.
- Quah, J. K.-H. and B. Strulovici (2009, November). Comparative Statics, Informativeness, and the Interval Dominance Order. *Econometrica* 77(6), 1949–1992.
- Smith, L. and P. Sorensen (2000). Pathological outcomes of observational learning. *Econometrica* 68(2), 371–398. Publisher: John Wiley & Sons, Ltd.
- Smith, L., P. N. Sørensen, and J. Tian (2021, October). Informational Herding, Optimal Experimentation, and Contrarianism. *The Review of Economic Studies* 88(5), 2527–2554.
- Van Zandt, T. and X. Vives (2007). Monotone equilibria in Bayesian games of strategic complementarities. *Journal of Economic Theory* 134(1), 339–360. Publisher: Academic Press.
- Veeraraghavan, S. K. and L. G. Debo (2011). Herding in Queues with Waiting Costs: Rationality and Regret. *Manufacturing and Service Operations Management* 13(3), 329–346. Publisher: INFORMS.