

Second Opinions and the Humility Threshold*

David Lagziel[†] and Yevgeny Tsodikovich[‡]

October 30, 2022

Abstract

An uninformed individual sequentially approaches two experts who possess private signals regarding an unknown state. Each expert 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 the experts' expertise levels to the equilibria of this game, showing that: (i) better-informed experts may generate worse recommendations in equilibrium; and (ii) ordering the experts so that the lower-level one provides the second opinion typically improves the outcome. Moreover, we show that limiting the experts' 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, and SING17 – 17th European Meeting on Game Theory. 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 multiple experts, instead of just one, before making important decisions. Such decisions could vary from assessing an academic paper for publication, to career choices, and even to 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 an uninformed individual (henceforth, UI), who faces a decision under uncertainty and sequentially consults two experts. The first expert provides a preliminary assessment, whereas the second is considered to be the *second opinion*. Every expert naturally strives to provide the correct assessment. Whenever he is incorrect, he may wish the other expert to be incorrect as well, so that they share the possible liability. This also holds when he is indeed correct, as he establishes superiority over the other expert. Since all assessments are public, the second expert observes the recommendation of the first. Our goal is to shed light on the strategic interaction between the two privately informed experts, and its implications on the ability of the UI to reach a correct decision.

The described process is, in fact, a sequential two-player Bayesian-game among the two experts. Given that each expert has an individual level of expertise, we ask what are the equilibria of this game and how they evolve as a function of the payoffs and expertise levels. Our analysis shows that there are 4 main equilibria profiles: (i) a *Revealing equilibrium* in which both experts reveal their private information; (ii) a *Herding equilibrium* in which the second expert mimics the first; (iii) a *Guided equilibrium* in which some learning among the experts takes place; and (iv) a *Mixed equilibrium* which builds on the Guided equilibrium (with mixed strategies).

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 to reach the correct decision, to classify each of the mentioned equilibria. Due to the learning involved, the guided equilibrium evidently supports the highest correctness level.

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 experts' expertise levels. Specifically, we show that an infinitesimal *increase* in the second expert'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 an expert 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 *second* expert, rather than the first. 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 UI potentially faces a complex medical procedure. Before making a decision, the UI sequentially consults two physicians who individually possess a private informative signal regarding his condition. The probability that the private signal of physician i matches the correct medical condition is referred to as the expertise level q_i of physician i . We typically assume that: (i) all recommendations are public; (ii) every physician prefers to be correct, while hoping the other physician is incorrect;¹ and (iii) in case the two physicians provide contradicting assessments, the UI follows the second opinion (otherwise, the process of consulting two experts becomes redundant).

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 second expert affects the UI's final decision. Figure 1a illustrates that the relation between these two elements is not necessarily monotone. The figure depicts the probability that the UI's

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

decision is correct, generally referred to as *correctness*, as a function of the second physician’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 second physician is relatively small, he remains attentive to the assessment of the first physician in a way that allows for some learning. Once the second physician crosses this threshold, the game shifts to a Revealing equilibrium in which the second physician completely ignores the assessment of the first, thus decreasing the probability to reach a correct decision.

The impact of the second physician on the UI’s correctness




Figure 1: In Figure 1a, an infinitesimal increase in the second physician’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 second physician to ignore the first, 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 less-informed expert provides the second opinion, and the dashed horizontal (green) line presents the equilibrium, a Revealing one, in case the more-informed expert provides the second opinion. One can see that below and close to the threshold (grey area), it is better to consult the less-informed expert second.

The second effect, referred to as *The Intern Effect*, considers the ordering of physicians according to their expertise levels. It shows that in some cases it is better to request a second opinion from the less-informed physician. This phenomenon is also a result of shifting between two equilibria, where a less-informed second physician learns from the assessment of the first

in equilibrium, whereas reversing the order eliminates this possibility. We however stress that identifying the less-informed expert is not easy, especially when positioned second, because the latter’s ex-post success rate would typically (not always!) appear better.

The driving force behind these effects is the disparity between the experts’ (game-based) expected payoffs and the ex-post success rate of the UI. Yet, it should be clear that experts are not oblivious to their ex-post reputation. The fact that *both* experts 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 complements the results of Ely and Välimäki (2003), where experts’ long-run reputation considerations over-shed any short-term payoffs and eliminate all the surplus.

1.1 Related research

The current research encompasses elements from several different sub-fields, thus its exact position in the literature is somewhat ambiguous. The primary natural candidate is the field of 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 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 as the same time. So in our set-up, the highest/lowest payoff is achieved

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

when a player chooses the ex-post correct/incorrect action alone.

Another critical difference relies 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 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.

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 underline 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 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.

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).

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

An uninformed individual (UI) wishes to identify a binary unknown state. Since he has neither informative signals concerning the state, nor the ability to process such signals, he approaches two heterogeneous experts (i.e., players), who possess private informative signals. For example, the UI could be a firm that approaches two advisers concerning a decision to enter a new market; an individual who consults two salespersons regarding an expensive purchase; or even a research student discussing a potential research question with faculty members. The UI approaches them sequentially, so that the second expert observes the recommendation of the first, and can be thought of as “a second opinion”. Each expert 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 (expert) falls short of this goal. This generates a sequential Bayesian game where the experts’ 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 $i = 1, 2$ (henceforth, Expert i) 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 Expert i , namely a measure of Expert i ’s ability to provide the correct recommendation. The action set of every Expert i is denoted by $A = \{0, 1\}$, where every action $a_i \in A$ denotes a recommendation regarding the realized state.

The game evolves as follows. First, nature chooses a state θ according to a common, publicly known, prior p . Then, every Expert i receives a private signal s_i based on the previously defined

information structure. Expert 1 is the first to act by posting a public recommendation $a_1 \in A$. After observing a_1 , Expert 2 provides his public recommendation $a_2 \in A$.




Figure 2: The evolution of the game.

We assume that every expert prefers giving the accurate recommendation – one that matches the state. To establish superiority, every expert also prefers that the recommendation of the other expert proves to be false. The latter condition ensures that the expert 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 Expert i is characterized by four possible payoffs: $U_1 > U_2 > U_3 > U_4$, where U_1 is the payoff in case Expert i is the only one to provide a correct recommendation, 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 Expert i when he is the only one to provide the wrong recommendation. 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.

		Expert 2	
		$a_2 = \theta$	$a_2 = 1 - \theta$
Expert 1	$a_1 = \theta$	(1, 1)	(α , $-\alpha$)
	$a_1 = 1 - \theta$	($-\alpha$, α)	(-1, -1)

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

Denote the strategy of Expert 1 by $\sigma_1 : S \rightarrow \Delta(A)$, and the strategy of Expert 2 by $\sigma_2 : S \times A \rightarrow \Delta(A)$. Since the game is sequential, the second expert has a structural advantage

of observing the recommendation of the first.

To ensure that the experts' recommendations are indeed informative, independently of the state, we assume that $\min\{q_1, q_2\} \geq p$, which is equivalent to $\Pr(\theta = x | s_i = x) > \frac{1}{2}$ for every (i, x) . Otherwise, in a single expert scenario where $q_i < p$, the state $\theta = 0$ is more likely than $\theta = 1$, regardless of the expert's information and assessment.

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_1, q_2, 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 3 illustrates these classes in the (q_1, q_2) -plane):⁵

- A profile (σ_1, σ_2) is a *Revealing equilibrium* if the actions of both experts match their private signal, i.e., if $\sigma_i = s_i$ for every i, s_i . This typically occurs when the second expert is much more informed than the first, so that he ignores previous recommendations.
- A profile (σ_1, σ_2) is a *Herding equilibrium* if the recommendation of Expert 2 matches that of Expert 1 in every realization of the game, while Expert 1 provides his private signal $\sigma_1 = s_1$. This typically occurs when the first expert is much more informed than the second, so that the second ignores his private signal.
- A profile (σ_1, σ_2) is a *Guided equilibrium* if both experts provide their private signal, $\sigma_i = s_i$, with the exception of $\sigma_2(a_1 = 0, s_2 = 1) = 0$. That is, the second expert provides a recommendation of $a_2 = 0$ whenever Expert 1 recommends $a_1 = 0$. This typically occurs when the two experts have close expertise levels, so that the second expert learns from the first, and weighs-in both signals into his recommendation.
- A profile (σ_1, σ_2) is a *Mixed equilibrium* if both experts use a mixed strategy over their recommendations. In particular, Expert 1 reports his signal if it is $s_1 = 0$, and uses a mixed actions otherwise, $\sigma_1(s_1 = 1) \in \Delta(A)$. Expert 2 reports his signal, $\sigma_2 = s_2$, with

⁵Other than on the boundaries between different areas in this parameter space, where different equilibria coincide.

the exception of $\sigma_2(a_1 = 0, s_2 = 1) \in \Delta(A)$. 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.

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 experts on the UI’s final decision. For that purpose we use the notion of “correctness”, which measures the probability of being correct, in each of the mentioned equilibria.⁶ Formally, given a profile (σ_1, σ_2) , define the *correctness* of Expert i by $C_i = \Pr(\sigma_i = \theta | \sigma_{-i})$. In addition, we are interested in the *correctness of the process*, which is the probability that the UI learns the true state of the world. We assume that the UI adopts the second opinion, namely the recommendation of the second expert, so $C(q_1, q_2) = C_2$. Before we continue, let us clarify this assumption.

The assumption that the UI follows the second opinion is based on the idea that, in many cases, a UI who approaches experts is a layman in terms of the decision process. Specifically, the UI is unfamiliar with the underline statistics of the problem, which include the experts’ expertise levels and the joint distributions of the event and signals. Thus, the relevant options for the UI are either to follow the first expert, or to follow the second (note that a simple randomization among the two experts would not produce a superior outcome). However, the first expert cannot learn from Expert 2’s recommendation which he does not see. So following the recommendation of the first expert would only yield a correctness of at most q_1 , making the second expert redundant. In other words, our focus on the learning process, vis-à-vis the recommendation of Expert 2 in equilibrium, becomes rather straightforward.

Remark 1. *We emphasize that the UI’s decision (to follow the second expert) does not change the analysis of the aforementioned game. The payoffs of the experts depend on their recommendations and on the true state of the world, and neither on the actions nor on the information of the UI. Hence, our analysis remains valid even if the UI follows any other decision rule, or can only observe the recommendation of just one expert.*

⁶See Arieli et al. (2018) for more details.

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 Expert 2's expertise level, so that an infinitesimal increase of q_2 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 Expert 2 ignores the action (and the conveyed information) of Expert 1 if and only if the second expert's expertise level is above the threshold. Thus, only a less-informed expert would take both signals into account to provide an even more accurate recommendation than the first, more-informed, expert.

In addition, in Proposition 3 we establish the importance of aligning the experts in an optimal manner, so that the second opinion is taken from the *less*-informed expert (the one with a lower expertise level), rather than the better-informed one. We refer to this result as *the intern effect* because it allows a more beneficial learning process, in equilibrium, compared to the alternative ordering of experts.

3.1 The case of symmetric payoffs

Our preliminary analysis deals with the symmetric case where the experts' individual payoffs offset when they provide different assessments, 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_1, q_2, 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 Expert 1 is significantly more accurate than Expert 2, whereas a significantly superior Expert 2 yields a Revealing equilibrium. On the other hand, the more sophisticated equilibria (the Guided and Mixed ones) emerge when neither of the two experts 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 $\tilde{q}_1 \geq \tilde{q}_2 + \tilde{p}$.*
2. *There exists a Revealing equilibrium if and only if $\tilde{q}_2 \geq \tilde{q}_1 + \tilde{p}$.*
3. *There exists a Guided equilibrium if and only if $\tilde{q}_i \leq \tilde{q}_{-i} + \tilde{p}$ for every $i = 1, 2$, and*

$$\tilde{q}_1 \geq \tilde{p} + \ln \left(\frac{2 \exp(-\tilde{q}_2) + \alpha + 1}{(\alpha + 1) \exp(-\tilde{q}_2) + 2} \right). \quad (1)$$

4. *There exists a Mixed equilibrium if and only if $\tilde{q}_i \leq \tilde{q}_{-i} + \tilde{p}$ for every $i = 1, 2$, 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 3 provides a visualization to the results of Theorem 1. It shows the disjoint equilibrium regions in the (q_1, q_2) -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 expert 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 of the experts have clear dominance over the other in terms of expertise level. This leads Expert 2 to rely on Expert 1’s truthful action in borderline situations, namely when Expert 2’s signal is

$s_2 = 1$ and this contradicts Expert 1's action of $a_1 = 0$ 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 3, is based on Expert 1's relative expertise level. When Expert 1's expertise level q_1 is rather close to the prior p , the high uncertainty leads both experts to hedge between their two available actions, i.e., a Mixed equilibrium.

Equilibria regions in the (q_1, q_2) -plane



Figure 3: 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).

Let us provide some intuition to the Guided equilibrium outcome. This equilibrium is reached whenever the expertise levels of both experts are rather close to one another, as evident from the condition $\tilde{q}_i \leq \tilde{q}_{-i} + \tilde{p}$, for every $i = 1, 2$. Thus, when the second expert receives a signal that matches the action of the first, it is optimal for him to recommend accordingly. Otherwise, he resorts to the prior which is biased towards a $a_2 = 0$ 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_1, q_2 , 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 experts always recommend $a_i = 0$, irrespective of their signals, is an equilibrium, if and only if $q_1 = q_2 = p$. We generally disregard such examples because they do not provide meaningful insights into the strategic interaction between the two experts. On the other hand, it is easy to verify that non-informative strategies, i.e., providing the same recommendation independently of the signals, do not consist an equilibrium, unless $q_1 = q_2 = p$. This originates from the combination of informative signals and symmetric payoffs, so that being correct strictly dominates being incorrect, assuming that the other expert is uninformative.

Building on the results of Theorem 1, we can now provide some insights into the correctness of this decision process, from the UI's perspective, under the different equilibria. The following proposition shows that a Guided equilibrium, specifically, provides a greater accuracy compared to what each expert can achieve by himself, whereas the other three are limited to the experts' 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 $C(q_1, q_2) = \max\{q_1, q_2\}$, and given a Mixed equilibrium the correctness is $C(q_1, q_2) = q_2$. However, in case there exists only a Guided equilibrium, the correctness of the process is $C(q_1, q_2) > \max\{q_1, q_2\}$.*

The two straightforward cases in the analysis of Proposition 1 are the Herding and Revealing equilibria, that yield, by definition, a correctness of q_1 and q_2 , respectively. To compare, a Guided equilibrium generates a strictly higher correctness compared to $\max\{q_1, q_2\}$ because it supports a learning process among the two experts. Figure 4 illustrates this result as a function of Expert 2's expertise level, for a fixed q_1 . It is important to note that these results, and the notion of correctness in general, are a function of the experts' strategies profile, i.e., the equilibrium type. So, from a mechanism-design perspective, the UI 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, similarly to the

Guided equilibrium, its correctness remains limited to q_2 . This originates from the indifference that both experts 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 first expert generates through his randomization.

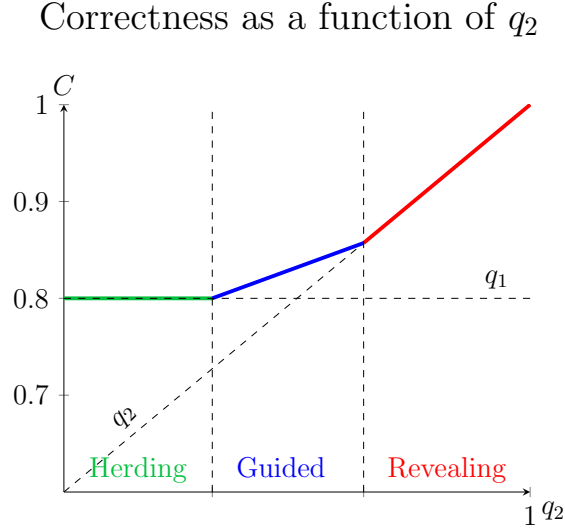


Figure 4: The correctness of the process as a function of the second expert’s level, and a fixed level of $q_1 = 0.8$ for the first expert, and $(p, \alpha) = (0.6, 2)$. A Mixed equilibrium does not exist under these parameters (see Fig. 3). The dashed vertical lines divide the axis to the different 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 experts.

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 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 $\gamma = \frac{U_1 - (-1)}{1 - U_4}$.

		Expert 2	
		$a_2 = \theta$	$a_2 = 1 - \theta$
Expert 1	$a_1 = \theta$	$(1, 1)$	(U_1, U_4)
	$a_1 = 1 - \theta$	(U_4, U_1)	$(-1, -1)$

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

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 of 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 $\tilde{\gamma} \leq \tilde{q}_1 + \tilde{q}_2 - \tilde{p}$. The reason is that when $\tilde{\gamma}$ is too large, it becomes the sole driving force behind the equilibria and subsequent results, irrespective of the experts' private and public information. Specifically, suppose that both signals are $s_1 = s_2 = 1$, and that Expert 1 truthfully reports $a_1 = 1$. This is the highest possible posterior on the event $\{\theta = 1\}$. Still, if the following inequality holds

$$(1 - p)q_1q_2 \cdot 1 + p(1 - q_1)(1 - q_2) \cdot (-1) \leq (1 - p)q_1q_2 \cdot U_4 + p(1 - q_1)(1 - q_2) \cdot U_1,$$

then Expert 2's benefits from "gambling" on the low-probability event $\{\theta = 0 | s_1 = s_2 = 1\}$, simply because it opposes the action of Expert 1. The last inequality is indeed $\tilde{q}_1 + \tilde{q}_2 - \tilde{p} \leq \tilde{\gamma}$.

We start with a simple generalization of previous results. The following theorem extends Theorem 1 to asymmetric payoffs, that is, when $\tilde{\gamma} \neq 1$. Note that for the case of a Guided equilibrium, we impose an additional condition that Expert 1 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_1, q_2 , and γ .*

- *There exists a Herding equilibrium if and only if $\tilde{\gamma} < \tilde{q}_1 - \tilde{q}_2 - \tilde{p}$.*
- *If Expert 1 reveals his private signal in equilibrium and $\tilde{q}_1 - \tilde{q}_2 - \tilde{p} < \tilde{\gamma} < \tilde{q}_1 - \tilde{q}_2 + \tilde{p}$, then the unique equilibrium is a Guided one.*

- *There exists a Revealing equilibrium if and only if $\tilde{q}_1 - \tilde{q}_2 + \tilde{p} < \tilde{\gamma} < \tilde{q}_1 + \tilde{q}_2 - \tilde{p}$.*

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 5 illustrates this transition through the different intervals, with respect to $\tilde{\gamma}$, for each of the mentioned equilibria.




Figure 5: 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 Expert 2 does not necessarily improve the overall decision/correctness of the UI. In fact, one can find two disjoint intervals such that *every* expertise level of Expert 2 in the lower interval generates a 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_1, γ, p) such that $\tilde{q}_1 > \tilde{\gamma} > \tilde{p}$ and assume that Expert 1 reveals his private signal in equilibrium. Then, there exist $\bar{q}_2 > \underline{q}_2$ such that $C(q_1, \underline{q}_2) > C(q_1, \bar{q}_2)$. Moreover, there exists q_2^* such that for every $q_2 \in (\underline{q}_2, q_2^*)$ and every $q_2' \in (q_2^*, \bar{q}_2)$, the correctness is higher when the second expert has lower quality signal, i.e. $C(q_1, q_2) > C(q_1, q_2')$.*

Two clarifications are in order. First, the preliminary assumption that $\tilde{q}_1 > \tilde{\gamma} > \tilde{p}$ enables us to shift from a Guided equilibrium to a Revealing one, by varying the expertise level q_2 of Expert 2. Otherwise, e.g., in case $\tilde{\gamma}$ is significantly larger than $\tilde{q}_1 + \tilde{p}$, we are left only with a Revealing equilibrium, independently of q_2 . Second, note that the potential loss from crossing the humility threshold q_2^* is not necessarily a mild one. Figure 6 illustrates the potential

magnitude of this non-monotone effect, where an infinitesimal increase in q_2 triggers a drop from 0.81 to 0.75 in the overall correctness value.

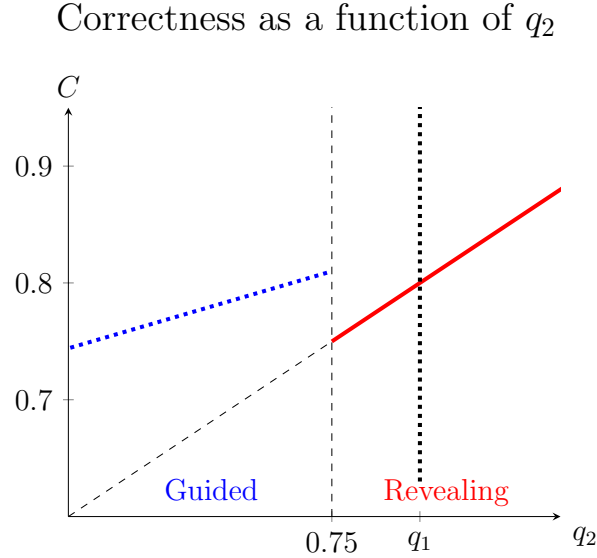


Figure 6: The correctness of the process as a function of Expert 2’s expertise level, given $p = 0.6$, $q_1 = 0.8$, and $\gamma = 2$. The dotted (blue) line describes the correctness under a Guided equilibrium, and the solid (red) line describes the correctness under a Revealing equilibrium. Note that the Humility threshold is below q_1 .

The next proposition relates to the possibilities of the UI in choosing the order of experts. Though the UI is, potentially, a layman in terms of the given problem, there is still the possibility of switching the order of experts in order to reach a better expected outcome. Proposition 3 shows that the UI may prefer taking the second opinion from the less-informed expert, namely the one with the lower expertise level. We refer to this result as *the intern effect* since one can think of the second expert as an intern who takes into account previous recommendations. The alternative ordering, where the better-informed expert is second, does not allow for this learning process to take place, since the better-informed expert 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 better-informed agent.

Proposition 3 (The intern effect). *Assume $\tilde{\gamma} > \tilde{p}$ and fix $q^H > q^L$ such that $\tilde{q}^H - \tilde{q}^L < \tilde{p}$. For every pair of experts such that $(q_1, q_2) = (q^H, q^L)$ yields a Guided equilibrium with cor-*

rectness $C(q^H, q^L)$, then reversing their order to $(q_1, q_2) = (q^L, q^H)$ 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 second opinion is given by the expert with the lower quality signal.

To get some intuition for this result, consider the signals that both experts receive. Once the high-level expert is second, it is sub-optimal for him to deviate from his relatively high-accuracy signal, thus generating a Revealing equilibrium. In practice, the ordering $(q_1, q_2) = (q^L, q^H)$ yields a Revealing equilibrium, which implies that the second recommendation is based entirely on the second signal. Once the ordering reverses, the second expert takes into account all available information, and reverts to the recommendation of the first expert when it is also supported by the prior.

A simple and important way to extend Proposition 3 is by examining a considerably better-informed expert, i.e., a significantly large q^H relative to q^L . In case q^H is indeed sufficiently large, then approaching this expert first 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 order of experts would yield a Revealing equilibrium with a correctness of q^H . So, it becomes weakly better to approach the highly informed expert first, even if a Guided equilibrium is not always achievable.

From a practical perspective, the classification of experts by their expertise levels is anything but trivial. This problem originates from the fact that Expert 2 sees the recommendation of Expert 1 in each of the mentioned equilibria in Proposition 3, and this enables him to maintain a higher correctness level relative to Expert 1. 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 Expert 2 remains well below the correctness of Expert 1. This occurs in equilibrium and although the second expert observes the recommendation of the first, and a mimicking strategy is indeed feasible, but sub-optimal. We derive this conclusion from the following observation.⁷

⁷The proof follows directly from Theorem 1* and Proposition 1, thus omitted.

Observation 1. Fix p, q_1, q_2 , and assume that Expert 1 reveals his private signal in equilibrium. Then, the correctness $C(q_1, q_2)$ in equilibrium is not a monotone function of γ :

- under a Herding equilibrium, i.e., $\tilde{\gamma} < \tilde{q}_1 - \tilde{q}_2 - \tilde{p}$, the correctness is q_1 ;
- under a Guided equilibrium, i.e., $\tilde{q}_1 - \tilde{q}_2 - \tilde{p} < \tilde{\gamma} < \tilde{q}_1 - \tilde{q}_2 + \tilde{p}$, the correctness is $C(q_1, q_2) > \min\{q_1, q_2\}$; and
- under a Revealing equilibrium, i.e., $\tilde{q}_1 - \tilde{q}_2 + \tilde{p} < \tilde{\gamma} < \tilde{q}_1 + \tilde{q}_2 - \tilde{p}$, the correctness is $C(q_1, q_2) = q_2$.

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_2 , by definition. Under symmetric payoffs (i.e., $\tilde{\gamma} = 0$), a Revealing equilibrium implies that $q_2 > q_1$, whereas under asymmetric payoffs (i.e., $\tilde{\gamma} \neq 0$), a Revealing equilibrium exists even if $q_2 < q_1$, 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 value among the two experts, whereas an optimal design of γ can facilitate either a Guided, or a Herding equilibrium, to achieve the maximal possible correctness level, given the two available experts. 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 among experts, than all other equilibria. Our analysis depicts a mapping from the experts' expertise levels to the equilibria of this game, showing that: (i) better-informed experts may generate worse recommendations in equilibrium; and (ii) ordering the experts so that the lower-level one provides the second opinion can typically improve the outcome. This contraries the common belief that the more skilled expert should be the one to provide the second opinion.

Moreover, if the UI can influence the experts' liability in case of an error, then he can ensure an equilibrium in which his correctness is maximized. This optimal liability is generally robust and remains the same, irrespective of either small changes in the experts' quality, or the prior probability of the state of the world. Thus, although the UI is typically uninformed of either the experts' assessment processes, or their quality, he can still facilitate cooperation between the two and optimize his probability of learning the true state of the world.

References

- Ali, S. N. and N. Kartik (2012, November). Herding with collective preferences. *Economic Theory* 51(3), 601–626.
- Arieli, I. (2017, July). Payoff externalities and social learning. *Games and Economic Behavior* 104, 392–410. Publisher: Academic Press.
- Arieli, I., M. Koren, and R. Smorodinsky (2018, June). 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, August). 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, May). 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.
- Clippel, G., K. Eliaz, D. Fershtman, and K. Rozen (2021). On selecting the right agent. *Theoretical Economics* 16(2), 381–402.
- Dasgupta, A. (2000). Social learning with payoff complementarities. *Working paper* (November 1999).
- Debo, L. G., C. Parlour, and U. Rajan (2011, December). 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, September). Congested observational learning. *Games and Economic Behavior* 87, 519–538. Publisher: Academic Press.
- 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.

- 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, May). 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, March). Pathological outcomes of observational learning. *Econometrica* 68(2), 371–398. Publisher: John Wiley & Sons, Ltd.
- Van Zandt, T. and X. Vives (2007, May). 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, March). Herding in Queues with Waiting Costs: Rationality and Regret. *Manufacturing and Service Operations Management* 13(3), 329–346. Publisher: INFORMS.

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_1, q_2, p , and γ . Assume that Expert 1's action matches his signal, as in all the equilibria stated in the Theorem, other than the Mixed equilibrium. Our analysis is divided into the four different states depending on a_1 and s_2 .

Consider first the case where $a_1 = 0$ and $s_2 = 0$. This state occurs with probability q_1q_2 when $\theta = 0$, and with probability $(1 - q_1)(1 - q_2)$ when $\theta = 1$. Expert 2's best response in this state is $a_2 = 0$ if and only if

$$pq_1q_2 \cdot 1 + (1 - p)(1 - q_1)(1 - q_2) \cdot (-1) \geq pq_1q_2 \cdot U_4 + (1 - p)(1 - q_1)(1 - q_2) \cdot U_1,$$

(the best response is $a_2 = 1$ when the inequality is reversed). The last inequality can be rearranged into

$$\frac{p}{1 - p} \cdot \frac{q_1}{1 - q_1} \cdot \frac{q_2}{1 - q_2} \cdot \frac{1 - U_4}{1 + U_1} \geq 1,$$

or, equivalently,

$$\tilde{p} + \tilde{q}_1 + \tilde{q}_2 - \tilde{\gamma} \geq 0. \tag{2}$$

Next, consider the case where $a_1 = 0$ and $s_2 = 1$. This state occurs with probability $q_1(1 - q_2)$ when $\theta = 0$, and with probability $(1 - q_1)q_2$ when $\theta = 1$. Expert 2's best response in this state is $a_2 = 0$ if and only if

$$pq_1(1 - q_2) \cdot 1 + (1 - p)(1 - q_1)q_2 \cdot (-1) \geq pq_1(1 - q_2) \cdot U_4 + (1 - p)(1 - q_1)q_2 \cdot U_1,$$

which can be rearranged into

$$\tilde{p} + \tilde{q}_1 - \tilde{q}_2 - \tilde{\gamma} \geq 0. \tag{3}$$

Third, consider the case where $a_1 = 1$ and $s_2 = 0$. This state occurs with probability $(1 - q_1)q_2$ when $\theta = 0$, and with probability $q_1(1 - q_2)$ when $\theta = 1$. Expert 2's best response in this state is $a_2 = 0$ if and only if

$$p(1 - q_1)q_2 \cdot U_1 + (1 - p)q_1(1 - q_2) \cdot U_4 \geq p(1 - q_1)q_2 \cdot (-1) + (1 - p)q_1(1 - q_2) \cdot 1,$$

which can be rearranged into

$$\tilde{p} - \tilde{q}_1 + \tilde{q}_2 + \tilde{\gamma} \geq 0. \quad (4)$$

Lastly, consider the case where $a_1 = 1$ and $s_2 = 1$. This state occurs with probability $(1 - q_1)(1 - q_2)$ when $\theta = 0$, and with probability q_1q_2 when $\theta = 1$. Expert 2's best response in this state is $a_2 = 0$ if and only if

$$p(1 - q_1)(1 - q_2) \cdot U_1 + (1 - p)q_1q_2 \cdot U_4 \geq p(1 - q_1)(1 - q_2) \cdot (-1) + (1 - p)q_1q_2 \cdot 1,$$

which can be rearranged into

$$\tilde{p} - \tilde{q}_1 - \tilde{q}_2 + \tilde{\gamma} \geq 0. \quad (5)$$

In a **Revealing equilibrium**, $a_2(\cdot, s_2) = s_2$, so Ineq. (2) and Ineq. (4) should hold while Ineq. (3) and Ineq. (5) are reversed:

$$\tilde{p} + \tilde{q}_1 + \tilde{q}_2 - \tilde{\gamma} \geq 0, \quad (6a)$$

$$\tilde{p} + \tilde{q}_1 - \tilde{q}_2 - \tilde{\gamma} \leq 0, \quad (6b)$$

$$\tilde{p} - \tilde{q}_1 + \tilde{q}_2 + \tilde{\gamma} \geq 0, \quad (6c)$$

$$\tilde{p} - \tilde{q}_1 - \tilde{q}_2 + \tilde{\gamma} \leq 0. \quad (6d)$$

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 $\min\{\tilde{q}_1, \tilde{q}_2\} \geq \tilde{p}$, so the aforementioned inequalities reduce to Ineq. (6b) with $\tilde{\gamma} = 0$, i.e., $\tilde{q}_1 + \tilde{p} \leq \tilde{q}_2$, as stated in the theorem.

Hence, if Ineq. (6b) with $\tilde{\gamma} = 0$ holds, whenever Expert 1 reports his signal, the best response of Expert 2 is to report his signal as well. Clearly, the converse is true. Suppose Expert 2 reports his signal, so the action of Expert 1 have no effect on the action of Expert 2. 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**, $a_2(a_1, \cdot) = a_1$, so Ineq. (2) and Ineq. (3) should hold while

Ineq. (4) and Ineq. (5) are reversed:

$$\tilde{p} + \tilde{q}_1 + \tilde{q}_2 - \tilde{\gamma} \geq 0, \quad (7a)$$

$$\tilde{p} + \tilde{q}_1 - \tilde{q}_2 - \tilde{\gamma} \geq 0, \quad (7b)$$

$$\tilde{p} - \tilde{q}_1 + \tilde{q}_2 + \tilde{\gamma} \leq 0, \quad (7c)$$

$$\tilde{p} - \tilde{q}_1 - \tilde{q}_2 + \tilde{\gamma} \leq 0. \quad (7d)$$

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., $\tilde{q}_1 \geq \tilde{q}_2 + \tilde{p}$, as stated in the theorem.

Hence, if Ineq. (7c) with $\tilde{\gamma} = 0$ holds, whenever Expert 1 reports his signal, the best response of Expert 2 is to ignore his signal and report a_1 . Clearly, the converse is true. Suppose Expert 2 repeats the action of Expert 1. Expert 1 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**, $a_2(a_1, s_2) = s_2$ except for $a_2(0, 1) = 0$ so Ineq. (2), Ineq. (3) and Ineq. (4) should hold while Ineq. (5) is reversed:

$$\tilde{p} + \tilde{q}_1 + \tilde{q}_2 - \tilde{\gamma} \geq 0, \quad (8a)$$

$$\tilde{p} + \tilde{q}_1 - \tilde{q}_2 - \tilde{\gamma} \geq 0, \quad (8b)$$

$$\tilde{p} - \tilde{q}_1 + \tilde{q}_2 + \tilde{\gamma} \geq 0, \quad (8c)$$

$$\tilde{p} - \tilde{q}_1 - \tilde{q}_2 + \tilde{\gamma} \leq 0. \quad (8d)$$

Note that Ineq. (8d) implies that Ineq. (8a) holds, and Ineqs. (8c) and (8b) are equivalent to $\tilde{\gamma} - \tilde{p} \leq \tilde{q}_1 - \tilde{q}_2 \leq \tilde{\gamma} + \tilde{p}$. Since $q_2 \geq p$, one can show that Ineq. (8b) yields Ineq. (8d) as follows,

$$\tilde{q}_1 + \tilde{q}_2 \geq \tilde{q}_1 + \tilde{p} \geq \tilde{\gamma} + \tilde{q}_2 \geq \tilde{\gamma} + \tilde{p}.$$

So, one only needs to sustain the inequalities $\tilde{\gamma} - \tilde{p} \leq \tilde{q}_1 - \tilde{q}_2 \leq \tilde{\gamma} + \tilde{p}$. Again, take $\tilde{\gamma} = 0$, and Expert 2 follows the Guided equilibrium given that $-\tilde{p} \leq \tilde{q}_1 - \tilde{q}_2 \leq \tilde{p}$, as stated in the theorem.

Suppose Expert 1 observes $s_1 = 0$. By reporting $a_1 = 0$, he guides the second expert to report $a_2 = 0$ as well. By reporting $a_1 = 1$, the second expert would act according to his own

signal. Reporting $a_1 = 0$ is better if and only if

$$pq_1 \cdot 1 + (1-p)(1-q_1) \cdot (-1) \geq pq_1q_2 \cdot (-\alpha) + pq_1(1-q_2) \cdot (-1) + (1-p)(1-q_1)q_2 \cdot 1 + (1-p)(1-q_1)(1-q_2) \cdot \alpha,$$

which is equivalent to

$$\frac{p}{1-p} \cdot \frac{q_1}{1-q_1} \geq \frac{1+q_2+\alpha-\alpha q_2}{1+q_2\alpha+1-q_2}.$$

The last inequality holds since the LHS is bounded from below by 1 (recall that $p, q_1 \geq 0.5$), while the RHS is bounded from above by 1 (since $q_2 \geq 0.5$ and $\alpha \geq 1$).

Now suppose Expert 1 observes $s_1 = 1$. Reporting $a_1 = 1$ is better if and only if

$$p(1-q_1)q_2 \cdot (-\alpha) + p(1-q_1)(1-q_2) \cdot (-1) + (1-p)q_1(1-q_2) \cdot \alpha + (1-p)q_1q_2 \cdot 1 \geq p(1-q_1) \cdot 1 + (1-p)q_1 \cdot (-1),$$

which is equivalent to

$$\frac{q_1}{1-q_1} \cdot \frac{1-p}{p} \geq \frac{2+q_2(\alpha-1)}{1+\alpha+q_2(1-\alpha)},$$

or,

$$\tilde{q}_1 - \tilde{p} \geq \ln \left(\frac{2 \exp(-\tilde{q}_2) + \alpha + 1}{(\alpha + 1) \exp(-\tilde{q}_2) + 2} \right),$$

as stated in the theorem.

In a **Mixed equilibrium**, the signal $s_1 = 0$ yields $a_1 = 0$, while $s_1 = 1$ leads to a mixed action $(r, 1-r)$ of Expert 1. Similarly, Expert 2 plays according to his own signal, $a_2 = s_2$, unless $(a_1, s_2) = (0, 1)$, a case in which he plays a Mixed action $(\rho, 1-\rho)$.

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

- If $(a_1, s_2) = (0, 0)$, then $a_2 = 0$ if and only if

$$pq_2[q_1+r(1-q_1)] - (1-p)(1-q_2)[q_1r+(1-q_1)] \geq (1-p)(1-q_2)\alpha[q_1r+(1-q_1)] - pq_2\alpha[q_1+r(1-q_1)],$$

which reduces to $\frac{p}{1-p} \cdot \frac{q_2}{1-q_2} \geq \frac{q_1r+1-q_1}{q_1+r(1-q_1)}$, and the last inequality evidently holds.

- If $(a_1, s_2) = (1, 0)$, then $a_2 = 0$ if and only if

$$\alpha p(1-q_1)q_2(1-r) - \alpha(1-p)q_1(1-q_2)(1-r) \geq (1-p)q_1(1-q_2)(1-r) - p(1-q_1)q_2(1-r),$$

which reduces to $\tilde{q}_1 - \tilde{q}_2 \leq \tilde{p}$, and stated in the theorem.

- If $(a_1, s_2) = (1, 1)$, then $a_2 = 1$ if and only if

$$(1-p)q_1q_2(1-r) - p(1-q_1)(1-q_2)(1-r) \geq \alpha p(1-q_1)(1-q_2)(1-r) - \alpha(1-p)q_1q_2(1-r),$$

which reduces to $\tilde{q}_1 + \tilde{q}_2 \geq \tilde{p}$.

- If $(a_1, s_2) = (0, 1)$, then Expert 2 plays a Mixed action if and only if

$$p(1-q_2)[q_1+(1-q_1)r] - (1-p)q_2[q_1r+(1-q_1)] = \alpha(1-p)q_2[q_1r+(1-q_1)] - \alpha p(1-q_2)[q_1+r(1-q_1)],$$

which reduces to $r = \frac{(1-p)(1-q_1)q_2 - pq_1(1-q_2)}{p(1-q_1)(1-q_2) - (1-p)q_1q_2}$. Note that $r \in [0, 1]$ as long as $\tilde{p} \leq \tilde{q}_1 + \tilde{q}_2$ and $-\tilde{p} \leq \tilde{q}_1 - \tilde{q}_2$, as stated in the theorem.

Moving on to the actions of Expert 1, if $s_1 = 0$ then $a_1 = 0$ if and only if

$$pq_1[1 + \alpha q_2 + (1 - q_2)(\rho + (1 - \rho)\alpha)] \geq (1 - p)(1 - q_1)[1 + \alpha(1 - q_2) + q_2(\rho + (1 - \rho)\alpha)],$$

which reduces to $\frac{p}{1-p} \cdot \frac{q_1}{1-q_1} \geq \frac{1+\alpha+q_2\rho(1-\alpha)}{1+\alpha\rho(1-q_2)(1-\alpha)}$, and the last inequality holds for any $\alpha \geq 1$ and $\rho \in [0, 1]$. In addition, if $s_1 = 1$, then Expert 1 is indifferent between $a_1 = 0$ and $a_1 = 1$ if and only if

$$p(1 - q_1)[1 + \alpha + (1 - q_2)\rho(1 - \alpha)] = (1 - p)q_1[1 + \alpha q_2\rho(1 - \alpha)],$$

which yields $\rho = \frac{1+\alpha}{\alpha-1} \cdot \frac{(1-p)q_1 - p(1-q_1)}{(1-p)q_1q_2 - p(1-q_1)(1-q_2)}$. Note that $\rho \in [0, 1]$ as long as

$$\tilde{q}_1 - \tilde{p} \leq \ln \left(\frac{2 \exp(-\tilde{q}_2) + \alpha + 1}{(\alpha + 1) \exp(-\tilde{q}_2) + 2} \right),$$

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

Proof of Proposition 1. By definition, the correctness under a Herding equilibrium is q_1 , which equals $\max\{q_1, q_2\}$ since $\tilde{q}_1 \geq \tilde{q}_2 + \tilde{p}$. Similarly, the correctness under a Revealing equilibrium is $q_2 = \max\{q_1, q_2\}$. So, let us consider a Guided equilibrium, in which $a_2 = s_2$ in case $a_1 = 1$, and $a_1 = 0$ leads to $a_2 = 0$. In other words, in a Guided equilibrium, Expert 2 deviates from his private signal if and only if $(a_1, s_2) = (0, 1)$. Thus,

$$\begin{aligned} C(q_1, q_2) &= \Pr(\sigma_2 = \theta | \sigma_1) \\ &= q_2 + pq_1(1 - q_2) - (1 - p)(1 - q_1)q_2 \\ &= pq_2 + pq_1 + q_1q_2 - 2pq_1q_2, \end{aligned}$$

and the last expression is symmetric w.r.t. substituting q_1 and q_2 . So, w.l.o.g., assume that $q_2 = \max\{q_1, q_2\}$. Theorem 1 states that a Guided equilibrium implies that $\tilde{p} + \tilde{q}_1 \geq \tilde{q}_2$, or equivalently, $\frac{p}{1-p} \cdot \frac{q_1}{1-q_1} \geq \frac{q_2}{1-q_2}$, thus

$$C(q_1, q_2) = q_2 + pq_1(1 - q_2) - (1 - p)(1 - q_1)q_2 \geq q_2,$$

and the inequality is strict in case $\frac{p}{1-p} \cdot \frac{q_1}{1-q_1} > \frac{q_2}{1-q_2}$, as needed.

Given a Mixed equilibrium, Expert 2 deviates from his private signal if and only if $(a_1, s_2) = (0, 1)$. Thus, the correctness is given by

$$\begin{aligned} C(q_1, q_2) &= q_2 + p(1 - q_2)[q_1 + (1 - q_1)r][\rho - (1 - \rho)] + (1 - p)q_2[q_1r + (1 - q_1)][(1 - \rho) - \rho] \\ &= q_2 + (2\rho - 1)\{p(1 - q_2)[q_1 + (1 - q_1)r] - (1 - p)q_2[q_1r + (1 - q_1)]\} = q_2, \end{aligned}$$

where the last inequality follows from Expert 2's indifference condition

$$\begin{aligned} (1 + \alpha)p(1 - q_2)[q_1 + (1 - q_1)r] &= (1 + \alpha)(1 - p)q_2[q_1r + (1 - q_1)] \\ p(1 - q_2)[q_1 + (1 - q_1)r] &= (1 - p)q_2[q_1r + (1 - q_1)] \end{aligned}$$

given in Theorem 1. □

Proof of Theorem 1*. Fix p, q_1, q_2 , and let us follow the cases as given in the theorem.

Assume that $\tilde{\gamma} < \tilde{q}_1 - \tilde{q}_2 - \tilde{p}$. According to the proof of Theorem 1, a Herding equilibrium exists if and only if $\tilde{\gamma} + \tilde{p} < \tilde{q}_1 - \tilde{q}_2$, which holds by assumption. Now, assuming that Expert 2 replicates the action of Expert 1, a straightforward computation would show that the optimal action of Expert 1 is to follow his signal as well, independently of $\tilde{\gamma}$. Thus, a Herding equilibrium exists if and only if $\tilde{\gamma} < \tilde{q}_1 - \tilde{q}_2 - \tilde{p}$.

Moving on to the case where $\tilde{q}_1 - \tilde{q}_2 - \tilde{p} < \tilde{\gamma} < \tilde{q}_1 - \tilde{q}_2 + \tilde{p}$. 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, Expert 1 follows his own signal, the equilibrium is also unique.

Lastly, assume that $\tilde{q}_1 - \tilde{q}_2 + \tilde{p} < \tilde{\gamma} < \tilde{q}_1 + \tilde{q}_2 - \tilde{p}$. 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 Expert 1 is $a_1 = s_1$, given that $a_2 = s_2$, and independently of $\tilde{\gamma}$. This establishes the existence of a Revealing equilibrium and concludes our proof. \square

Proof of Proposition 2. Fix (q_1, γ, p) so that $\tilde{q}_1 > \tilde{\gamma} > \tilde{p}$. We start by proving that there exist $q_2^* < q_1$ and $\epsilon_1 > 0$ such that for every $q_2 \in (q_2^* - \epsilon_1, q_2^*)$ there exists a unique equilibrium, namely a Guided one, and for every $q_2 \in (q_2^*, q_2^* + \epsilon_1)$ there exists a unique equilibrium – a Revealing one.

Take q_2^* such that $\tilde{\gamma} - \tilde{p} = \tilde{q}_1 - \tilde{q}_2^*$. The condition $\tilde{q}_1 > \tilde{\gamma} > \tilde{p}$ implies that $p < q_2^* < q_1$, and for every sufficiently close q_2 to q_2^* (from below), we get $\tilde{\gamma} < \tilde{q}_1 - \tilde{q}_2 + \tilde{p}$, as needed for a Guided equilibrium according to Theorem 1*. In addition, since $q_2^* < q_1$, then for every sufficiently close q_2 to q_2^* (from above), but still below q_1 , we get that $\tilde{q}_1 - \tilde{q}_2 < \tilde{\gamma} - \tilde{p}$, as needed for a Revealing equilibrium according to Theorem 1*. Thus, one can fix $\epsilon_1 \in (0, \min\{q_1 - q_2^*, q_2^* - p\})$, such that for every $q_2 \in (q_2^* - \epsilon_1, q_2^*)$ there exists a Guided equilibrium, and for every $q_2 \in (q_2^*, q_2^* + \epsilon_1)$ there exists a Revealing equilibrium. Note that the assumption that Expert 1 follows his own signal in equilibrium, along with the fact that the best response of Expert 2 is unique in each of these cases, suggest that the mentioned equilibria are indeed unique.

The correctness of Expert 2 under a unique Revealing equilibrium is simply q_2 . Since $q_1 > q_2^*$ and given $\epsilon_1 \in (0, \min\{q_1 - q_2^*, q_2^* - p\})$, it follows that $q_2 = \min\{q_1, q_2\}$ is also 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 $q_2 \geq q_1$ in a Revealing equilibrium. On the other hand, the correctness under a Guided equilibrium is

$$C(q_1, q_2) = q_2 + pq_1(1 - q_2) - (1 - p)(1 - q_1)q_2.$$

Since $\tilde{p} + \tilde{q}_1 - \tilde{q}_2^* = \tilde{\gamma} > \tilde{p}$, we deduce that $pq_1(1 - q_2^*) > \frac{p}{1-p}(1 - p)(1 - q_1)q_2^*$, which leads to

$$pq_1(1 - q_2^*) - (1 - p)(1 - q_1)q_2^* > \left(\frac{p}{1-p} - 1 \right) (1 - p)(1 - q_1)q_2^*.$$

Thus,

$$\begin{aligned}
q_2^* + pq_1(1 - q_2^*) - (1 - p)(1 - q_1)q_2^* &> q_2^*[1 + (2p - 1)(1 - q_1)] \\
&= q_2^*[q_1 + 2p(1 - q_1)] \\
&> q_2^* + \epsilon_2
\end{aligned}$$

for every $0 < \epsilon_2 < q_2^*[q_1 + 2p(1 - q_1) - 1]$. By continuity, one can take a positive and sufficiently small $\epsilon < \min\{\epsilon_1, \epsilon_2\}$ such that $C(q_1, q_2) = q_2 + pq_1(1 - q_2) - (1 - p)(1 - q_1)q_2 > q_2^* + \epsilon$, for every $q_2 \in (q_2^* - \epsilon, q_2^*)$.

Denote $\bar{q}_2 = q_2^* + \epsilon$ and $\underline{q}_2 = q_2^* - \epsilon$. Thus, for every $q_2 \in (\underline{q}_2, q_2^*)$, the correctness is strictly higher than for every $q_2' \in (q_2^*, \bar{q}_2)$. That is, $C(q_1, q_2) > C(q_1, q_2')$, 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 $(q_1, q_2) = (q^H, q^L)$ 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 $(q_1, q_2) = (q^L, q^H)$. 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_1, q_1) . 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