Bayesian Elicitation

Supplementary Appendix (For Online Publication)

Mark Whitmeyer*

February 6, 2024

Contents

A Sufficient Conditions for a Positive Value of Information

The text of Whitmeyer (2019) contains Proposition 5.5, which states that the value of information is positive in two-action signaling games in which the sender's messages are costless. Here, we provide additional sufficient conditions that may be of interest. In particular, we establish that the positive value of information result holds if there are just two states (regardless of whether the game is cheap talk).

Theorem A.1. *In signaling games without opacity design, the value of information is always positive for the receiver, provided*

1. *There are two states (or fewer); or*
2. *The receiver has two actions (or fewer); and*

*Department of Economics, Arizona State University. Email: mark.whitmeyer@gmail.com.

(a) There are three states (or fewer); or

(b) There are two messages (or fewer); or

(c) The game is cheap talk.

Note that throughout the proof of this theorem the receiver may not design the opacity in the game that follows. In addition, we observe that the result holds trivially if there exists a (fully) separating equilibrium and so in the proof(s) below we restrict attention to the cases where such separating equilibria do not exist.

First, we establish that it is WLOG to restrict attention to equilibria in which there is no state in which the sender mixes over messages following which the receiver strictly prefers different actions.

Lemma A.2. *In games, there exists a receiver-optimal equilibrium in the sender, in any state, mixes over messages that induce beliefs at which the receiver strictly prefers different actions.*

Proof. Let each action be strictly optimal in at least one state (or else the result is trivial). We may partition the set of states $\Theta = \Theta_1 \sqcup \Theta_2$, where Θ_1 is the set of states in which the receiver strictly prefers to choose action a_1, and Θ_2 is the set of states in which the receiver strictly prefers to choose action a_2. It is WLOG to suppose that there are no states in which the receiver is indifferent between her two actions.

Equivalently, $v_i := u_R(a_1, \theta_i) > u_R(a_2, \theta_i) =: w_i$ for all $\theta_i \in \Theta_1$, and $w_i > v_i$ for all $\theta_i \in \Theta_2$. Denote

$$\Theta_1 = \{\theta_1, \dots, \theta_t\}, \quad \text{and} \quad \Theta_2 = \{\theta_{t+1}, \dots, \theta_n\}$$

Consider an equilibrium in which the sender mixes in at least one state, θ_k. WLOG, let $\theta_k \in \Theta_1$, i.e., in that state the receiver strictly prefers to choose action a_1. Next, suppose that θ_k mixes over a subset of the set of messages, M_k, where

$$M_k = \{m_1, \dots, m_l\}$$

with a generic element $m_k \in M_k$. Moreover, M_k is partitioned by three sets, M_k^0, M_k^1 and M_k^2, where M_k^0 is the set of messages after which the receiver is indifferent between her two actions, M_k^1 is the set of messages after which the receiver strictly prefers a_1, and

M_k^2 is the set of messages after which the receiver strictly prefers a_2. By assumption, neither M_k^1 nor M_k^2 is the empty set, in which case we can pick two messages, $m_1 \in M_k^1$ and $m_2 \in M_k^2$. The receiver's expected payoff at this equilibrium can be written as

$$V(\mu) = \mu(\theta_k)(\sigma_k(m_1)v_k + \sigma_k(m_2)w_k) + \gamma$$

where γ is the remainder of the receiver's payoff that–crucially for the sake of this proof–does not depend on $\sigma_k(m_1)$ or $\sigma_k(m_2)$. However, the receiver's payoff strictly increases if instead θ_k modified his mixed strategy so that $\hat{\sigma}_k(m_1) = \sigma_k(m_1) + \sigma_k(m_2)$ and $\hat{\sigma}_k(m_2) = 0$ since $v_k > w_k$ (recall that we stipulated that $\theta_k \in \Theta_1$). Moreover, it is easy to see that this is also an equilibrium: θ_k is indifferent over any pure strategy in the support of his mixed strategy; and following messages m_1 and m_2 under the new mixture, the receiver still finds it optimal to choose a_1 and a_2, respectively (and the receiver's beliefs and payoffs following any other message are unchanged).

Since $m_1 \in M_k^1$ and $m_2 \in M_k^2$ were two arbitrary messages, and θ_k was an arbitrary state, the result follows. ■

Second, we discover that if there is a receiver-optimal equilibrium at belief μ_0 in which at most two messages are used, then any information benefits the receiver. Formally,

Lemma A.3. *In any signaling game, if there is a receiver-optimal equilibrium at belief μ_0 in which at most two messages are used, then any initial experiment benefits the receiver.*

Proof. Again, let each action be uniquely optimal in at least one state, and let two messages be used in the receiver-optimal equilibrium at belief μ_0 (if only one message is used, the receiver obtains the pooling payoff at μ_0, and hence any initial experiment must be to her profit).

In addition, we may, WLOG, impose that at μ_0 there is an equilibrium such that, following each message, m_1 and m_2, different actions, a_1 and a_2, respectively, are strictly optimal. Otherwise, this would just yield the pooling payoff and the result would be trivial. By Lemma A.2, this imposition ensures that the sender is choosing a pure strategy in each state.

Next, partition set Θ_1 into two sets, G_1^e and G_1^d, which correspond to the states in which the sender chooses m_1 and m_2, respectively. Likewise, partition set Θ_2 into two sets G_2^e and G_2^d which correspond to the states in which the sender chooses m_2 and m_1, respectively.

These sets are, explicitly,

$$G_1^e := \{\theta_1, \dots, \theta_m\}, \qquad G_1^d := \{\theta_{m+1}, \dots, \theta_t\}$$
$$G_2^e := \{\theta_{t+1}, \dots, \theta_r\}, \qquad G_2^d := \{\theta_{r+1}, \dots, \theta_n\}$$

Moreover, note that it is possible that some are the empty set; although of course if G_2^d is nonempty then G_1^e cannot be empty, and similarly for G_1^d and G_2^e.

Next, since we have imposed that an equilibrium of the above form exists at μ_0, such an equilibrium must exist at any belief μ such that the following condition holds:

Condition A.4.

$$\sum_{i=1}^{m} \mu(\theta_i) v_i + \sum_{j=r+1}^{n} \mu\left(\theta_j\right) v_j \geq \sum_{i=1}^{m} \mu(\theta_i) w_i + \sum_{j=r+1}^{n} \mu\left(\theta_j\right) w_j$$

and

$$\sum_{i=t+1}^{r} \mu(\theta_i) w_i + \sum_{j=m+1}^{t} \mu\left(\theta_j\right) w_j \geq \sum_{i=t+1}^{r} \mu(\theta_i) v_i + \sum_{j=m+1}^{t} \mu\left(\theta_j\right) v_j$$

Accordingly, at μ_0, the receiver's payoff is

$$V^T(\mu_0) = \sum_{i=1}^{m} \mu_0(\theta_i) v_i + \sum_{j=r+1}^{n} \mu_0\left(\theta_j\right) v_j + \sum_{i=t+1}^{r} \mu_0(\theta_i) w_i + \sum_{j=m+1}^{t} \mu_0\left(\theta_j\right) w_j$$

WLOG, we may suppose that there are just three signal realizations: one, A, after which Condition A.4 holds; one, B, after which there is a pooling equilibrium for which a_1 is optimal; and one, C, after which there is a pooling equilibrium for which a_2 is optimal. We may make this assumption since the receiver is only aided by multiple signal realizations in each pooling region. It is possible that there may be other, better equilibria for the receiver in such regions, but the proof requires only use of the pooling payoff which bounds the receiver's payoff from below. Then the receiver's expected payoff from the initial signal is bounded below by

$$\begin{aligned} & p\left[\sum_{i=1}^{m}\mu_A(\theta_i)v_i+\sum_{j=r+1}^{n}\mu_A\left(\theta_j\right)v_j+\sum_{i=t+1}^{r}\mu_A(\theta_i)w_i+\sum_{j=m+1}^{t}\mu_A\left(\theta_j\right)w_j\right] \\ & +q\left[\sum_{i=1}^{m}\mu_B(\theta_i)v_i+\sum_{j=r+1}^{n}\mu_B\left(\theta_j\right)v_j+\sum_{i=t+1}^{r}\mu_B(\theta_i)v_i+\sum_{j=m+1}^{t}\mu_B\left(\theta_j\right)v_j\right] \\ & +s\left[\sum_{i=1}^{m}\mu_C(\theta_i)w_i+\sum_{j=r+1}^{n}\mu_C\left(\theta_j\right)w_j+\sum_{i=t+1}^{r}\mu_C(\theta_i)w_i+\sum_{j=m+1}^{t}\mu_C\left(\theta_j\right)w_j\right] \end{aligned} \tag{A1}$$

where

$$p := \Pr(A), \quad q := \Pr(B), \quad s := \Pr(C), \quad \text{and} \quad p+q+s=1,$$

and $p\mu_A(\theta_i)+q\mu_B(\theta_i)+s\mu_C(\theta_i)=\mu_0(\theta_i)$ for all i. Expression A1 can be simplified to

$$V^T(\mu_0)+q\Upsilon+s\Gamma$$

where

$$\Upsilon := \left[\sum_{i=t+1}^{r}\mu_B(\theta_i)(v_i-w_i)+\sum_{j=m+1}^{t}\mu_B\left(\theta_j\right)\left(v_j-w_j\right)\right]$$

and

$$\Gamma := \left[\sum_{i=1}^{m}\mu_C(\theta_i)(w_i-v_i)+\sum_{j=r+1}^{n}\mu_C\left(\theta_j\right)\left(w_j-v_j\right)\right]$$

Since a_1 is optimal in the pooling equilibrium following B and a_2 is optimal in the pooling equilibrium following C we must have

$$\sum_{i=1}^{m}\mu_B(\theta_i)(v_i-w_i)+\sum_{j=r+1}^{n}\mu_B\left(\theta_j\right)\left(v_j-w_j\right) \geq \sum_{i=t+1}^{r}\mu_B(\theta_i)(w_i-v_i)+\sum_{j=m+1}^{t}\mu_B\left(\theta_j\right)\left(w_j-v_j\right) \tag{A2}$$

and

$$\sum_{i=1}^{m}\mu_C(\theta_i)(v_i-w_i)+\sum_{j=r+1}^{n}\mu_C\left(\theta_j\right)\left(v_j-w_j\right)\leq\sum_{i=t+1}^{r}\mu_C(\theta_i)(w_i-v_i)+\sum_{j=m+1}^{t}\mu_C\left(\theta_j\right)\left(w_j-v_j\right)$$

Moreover, since Condition A.4 does not hold for belief μ_B, we must have either

$$\sum_{i=1}^{m}\mu_B(\theta_i)v_i+\sum_{j=r+1}^{n}\mu_B\left(\theta_j\right)v_j<\sum_{i=1}^{m}\mu_B(\theta_i)w_i+\sum_{j=r+1}^{n}\mu_B\left(\theta_j\right)w_j \tag{A3}$$

and/or

$$\sum_{i=t+1}^{r}\mu_B(\theta_i)w_i+\sum_{j=m+1}^{t}\mu_B\left(\theta_j\right)w_j<\sum_{i=t+1}^{r}\mu_B(\theta_i)v_i+\sum_{j=m+1}^{t}\mu_B\left(\theta_j\right)v_j \tag{A4}$$

Suppose that Inequality $A3$ holds. Then we may substitute it into Inequality $A2$ and cancel:

$$\sum_{i=t+1}^{r}\mu_B(\theta_i)(v_i-w_i)+\sum_{j=m+1}^{t}\mu_B\left(\theta_j\right)\left(v_j-w_j\right)\geq 0$$

Hence, Υ is positive. On the other hand, if Inequality $A4$ holds then we may substitute it directly into Υ, which again must be positive.

A symmetric procedure works at belief μ_C to establish that Γ also must be positive. Since Γ and Υ are both positive, Expression $A1$, the receiver's payoff from learning, must be at least weakly greater than $V^T(\mu_0)$. We have exhausted every case, and so conclude that any initial experiment benefits the receiver. ■

The costless nature of messages in cheap talk games, in conjunction with Lemma A.2, allows us to conclude that there must be a receiver-optimal equilibrium at belief μ_0 in which at most two messages are used. Refer to the text for more on this portion of the result. Next, we establish that if there are just two states, then there is always a receiver-optimal equilibrium in which at most two messages are used.

Lemma A.5. *Let the number of states, n, equal two. Then, any equilibrium that yields the receiver a payoff of v in which $l > 2$ messages are used, there exists an equilibrium in which at most two messages are used that yields the receiver a payoff that is weakly higher than v.*

Proof. Suppose that $l > 2$ messages are used. If there exist messages in the support of each state's strategy at which the receiver is certain of the state, then there must exist a separating equilibrium; hence, the result is trivial.

Suppose that there exist no separating equilibria. Then, it is WLOG to focus on just two classes of l-message equilibria: i. The sender chooses a mixed strategy with full support in both states, or ii. In one state (say θ_H), the sender chooses a mixed strategy with full support, and in the other chooses a mixed strategy with support on all but one message.[1]

In both cases, there will be l resulting equilibrium beliefs $\mu_1' < \mu_2' < \cdots < \mu_l'$, where in case ii. $\mu_l' = 1$. However since the sender is mixing in each state, he must be indifferent over each message in the support of their mixed strategy (in each state). Hence, there must also be an equilibrium in both cases in which only two messages are used, which induce beliefs μ_1' and μ_l'. Indeed such an equilibrium can be constructed by taking each on-path message m_i with the associated induced belief μ_i' with $i \neq 1, l$ and moving weight from each player's mixed strategy on m_i to message m_l at the ratio

$$\frac{\Delta(\sigma_H(m_i))}{\Delta(\sigma_L(m_i))} = \frac{\Delta(\sigma_H(m_l))}{\Delta(\sigma_L(m_l))} = \frac{(1-\mu)\mu_l'}{(1-\mu_l')\mu}$$

Such a process decreases μ_i' and by construction maintains μ_l'. This can be done until $\mu_i' = \mu_1'$ for each i.

The Blackwell experiment that corresponds to this new, binary, distribution of posteriors is more informative than in the original situation, where l messages were used. Hence, the receiver's payoff must be weakly higher in the two message equilibrium. ■

Lemmata A.3 and A.5 imply that in any signaling game with two states and two actions, *ex ante* learning always benefits the receiver. Perhaps surprisingly, the value of information is also always positive for the receiver in three state, two action signaling games. *Viz*,

[1] Indeed, any other l message equilibria in a game that does not admit a separating equilibria must have multiple messages that are chosen in one state, and given the existence of such an equilibrium, there must exist an equilibrium in which just one of those "single state" messages is used.

Lemma A.6. *In signaling games with three states and two actions, the receiver's payoff is convex in the prior.*

Proof. Denote the set of (three) states by $\Theta = \{\theta_L, \theta_M, \theta_H\}$. Recall our convention $v_i := u_R(a_1, \theta_i)$ and $w_i := u_R(a_2, \theta_i)$, for all $i = L, M, H$. WLOG, we may assume[2] that action a_1 is strictly optimal in states θ_L and θ_H, and action a_2 is strictly optimal in state θ_M: $v_L > w_L$, $v_H > w_H$, and $w_M > v_M$.

Next, observe that we can picture any belief in the (x, y)-coordinate plane, where the x-axis corresponds to μ_M, and the y-axis corresponds to μ_H. Define region R_1 as the region in which a_1 is optimal and R_2 as the region in which a_2 is optimal. Each region, R_1 and R_2, is compact and convex, and the two regions share a boundary that is a line segment. Define R to be the simplex of beliefs, $R = R_1 \cup R_2$.

From Lemma A.2, at the prior μ_0, it suffices to consider just three possible arrangements of the posteriors that are induced at equilibrium:

Case 1: All posteriors lie in one region.

Case 2: All posteriors that follow messages chosen by θ_M fall in region R_2, where there is at least one posterior that does not lie on the boundary $R_1 \cap R_2$; and all posteriors that follow messages chosen by θ_L and θ_H fall in region R_1, where there is at least one posterior that does not lie on the boundary $R_1 \cap R_2$.

Case 3: All posteriors that follow messages chosen by θ_H (θ_L) fall in region R_1, where there is at least one posterior that does not lie on the boundary $R_1 \cap R_2$; and all posteriors that follow messages chosen by θ_L (θ_H) and θ_M fall in region R_2, where there is at least one posterior that does not lie on the boundary $R_1 \cap R_2$. By symmetry, we need focus only on the case where the posteriors that follow θ_H's messages fall in region R_1.

Note that throughout this proof, by Lemma A.2, each belief that is not on the line segment $R_1 \cap R_2$ must lie on the boundary of the triangle (2-simplex) of beliefs.

In the first case, the receiver clearly benefits from any initial experiment. The payoff under the prior is the pooling payoff, and so *ex ante* learning can only aid the receiver. The second case is trickier: there are two sub-cases that we need to examine:

[2] Otherwise, this is just the two state case.

Case 2a: The mixed strategy in each state has support on at least one message that induces a belief that is *not* on the boundary $R_1 \cap R_2$.

Case 2b: In one state, say θ_L, the sender mixes only over messages that induce beliefs that are on the boundary $R_1 \cap R_2$.

In case 2a, it is easy to see that there must also be a receiver-optimal equilibrium in which θ_H and θ_L each choose one message (possibly the same message) that induces a belief in $R \setminus R_2$, and θ_M chooses one message that induces a belief in $R \setminus R_1$. This is clearly an equilibrium, since the sender in each state already has support of his mixed strategy on his respective message; is optimal for the receiver, since this yields the receiver the maximum possible payoff (the separating payoff); and, moreover, does not depend on the prior. Hence, any initial experiment benefits the receiver.

In case 2b, there must exist a receiver-optimal equilibrium in which θ_H sends just one message, m_H, which induces a belief in $R \setminus R_2$; θ_L sends just one message, m_L, which induces a belief on the boundary $R_1 \cap R_2$; and θ_M mixes between m_L and m_M, the latter which induces a belief in $R \setminus R_1$. We will return to this distribution of posteriors shortly.

Case 3 also must be divided into two cases:

Case 3a: θ_L mixes only over messages that induce beliefs that are on the boundary $R_1 \cap R_2$.

Case 3b: θ_L mixes over at least one message that induces a belief in $R \setminus R_1$.

Case 3a is identical to case 2b. In case 3b, there must exist a receiver-optimal equilibrium in which θ_H sends just one message, m_H, that induces a belief in $R \setminus R_2$; and θ_L and θ_M pool on one message m_p, that induces a belief in $R \setminus R_1$. Here, only two messages are used and so by Lemma A.3 any initial experiment benefits the receiver.

Consequently, it remains to consider the scenario that case 2b reduces to: for prior μ_0 just three messages are used as follows: θ_H separates and chooses message m_H, θ_L chooses message m_L and θ_M mixes between two messages, m_L and m_M, in such a way that the receiver is indifferent over her actions following m_L (note, that there is an equivalent scenario that is obtained by interchanging θ_H and θ_L).

The prior μ_0 must be such that

$$\mu_L^0 \leq \frac{w_M - v_M}{v_L - w_L}\mu_M^0$$

and the receiver's payoff is

$$v := V(\mu_0) = \mu_H^0 v_H + \mu_M^0 w_M + \mu_L^0 w_L$$

Call this equilibrium $S^\dagger$. It is clear that WLOG we may focus on an initial experiment that is binary, and which yields just two beliefs μ_1 and μ_2, where μ_1 is a belief such $S^\dagger$ is feasible, and μ_2 is a belief such that $S^\dagger$ is infeasible. To see that this is WLOG, note that if there are multiple initial experiment realizations after which $S^\dagger$ is feasible, the receiver achieves at least the payoff as in the case when there is just one such initial experiment realization. Likewise, if there are multiple initial experiment realizations after which $S^\dagger$ is infeasible, since we need only assume the pooling payoff in this case, it is again clear that the receiver achieves at least the payoff as in the case where there is just one such initial experiment realization.

Thus, the initial experiment, ζ, yields $\mu_1 = \left(\mu_L^1, \mu_M^1, \mu_H^1\right)$ with probability p and $\mu_2 = \left(\mu_L^2, \mu_M^2, \mu_H^2\right)$ with probability $(1-p)$, where $p\mu_1 + (1-p)\mu_2 = \mu_0$. For belief μ_2, the receiver's payoff is bounded below by the pooling equilibrium payoff, and so we assume that this is indeed the payoff. Note that since μ_2 is not a belief for which $S^\dagger$ is feasible, we must have

$$\mu_L^2 > \frac{w_M - v_M}{v_L - w_L}\mu_M^2 \tag{A5}$$

Claim A.7. *For belief μ_2, action a_1 is optimal.*

Proof. Suppose for the sake of contradiction that a_1 is not optimal. That is,

$$\mu_L^2 w_L + \mu_M^2 w_M + \mu_H^2 w_H > \mu_L^2 v_L + \mu_M^2 v_M + \mu_H^2 v_H$$

But then

$$\mu_L^2 w_L + \mu_M^2 w_M + \mu_H^2 w_H > \mu_L^2 w_L + \mu_M^2 w_M + \mu_H^2 v_H$$
$$\mu_H^2 w_H > \mu_H^2 v_H$$

where the second inequality follows from Inequality A5. This is a contradiction. ■

Thus, a_1 is optimal and so the receiver's expected payoff is

$$V = p\left[\mu_H^1 v_H + \mu_M^1 w_M + \mu_L^1 w_L\right] + (1-p)\left[\mu_L^2 v_L + \mu_M^2 v_M + \mu_H^2 v_H\right]$$

which reduces to

$$V = v + (1-p)\left[\mu_L^2 (v_L - w_L) - \mu_M^2 (w_M - v_M)\right]$$

which is greater than v by Inequality A5. ■

We finish by showing that without opacity design, the value of information in games with two states and n actions is always positive. Before proceeding to the lemma, we encounter two new pieces of jargon. First, we call the experiment induced by the receiver-optimal equilibrium at belief μ_0 the **Null-Optimal Experiment**, $\eta\colon \Theta \to \Delta(M)$. Second, the realization of the initial experiment, y, begets a posterior belief μ_y, and we call the experiment induced by the receiver-optimal equilibrium at this belief the **y-Equilibrium Experiment**, $\gamma_y\colon \Theta \to \Delta(M)$.

Lemma A.8. *In any two state, n action, simple signaling game, the receiver's payoff is convex in the prior.*

Proof. Recall that we need consider only signaling games that do not admit separating equilibria since otherwise the result is trivial. Moreover, Lemma A.5 allows us to suppose that only two messages are used in any equilibrium. There are three cases to consider.

Case 1: For prior μ_0, the receiver-optimal equilibrium is one in which the sender pools. Observe that in this case, the null-optimal experiment is a completely uninformative experiment, so it is obvious that *ex ante* information benefits the receiver.

Case 2: For prior μ_0, the receiver-optimal equilibrium is one in which the sender mixes in one state and chooses a pure strategy in the other. Observe that in this case, the null-optimal experiment begets two posteriors: one that is in the interior on $[0,1]$ and the other that is either 0 or 1.

WLOG (the other cases follow analogously) suppose that θ_H mixes and chooses message m_1 with probability σ and θ_L chooses message m_1. Following an observation of message m_2, the receiver's belief is 1 and following message m_1 it is $\mu_j < \mu_0$.

Consider any initial experiment, ζ, with $k \geq 2$ realizations. Observe that for any realization that yields a belief $\mu_i \geq \mu_j$, an equilibrium in which θ_H mixes and θ_L does not must also exist, and hence the receiver's equilibrium payoffs for each of these beliefs must be bounded below by the payoff for that equilibrium. Suppose that in each case that this equilibrium is optimal (and hence that the receiver's payoffs are at their lower bounds).

For each y such that $\mu_y \geq \mu_j$, the y-equilibrium experiment is one that sends the posteriors to μ_j and 1. Consequently, it is WLOG to suppose that ζ just has a single experiment realization that yields a belief above μ_j. Moreover, any realization of experiment ζ that yields a posterior $\mu_y < \mu_j$ must beget an equilibrium payoff bounded below by the pooling payoff. Hence, we suppose that for each such realization y, the optimal equilibrium is the pooling equilibrium. Moreover, the resulting payoff from this distribution over pooling payoffs itself is bounded below by the payoff were the initial experiment to have merely a single signal y that begets a belief below μ_j. Hence we suppose that is the case.

To summarize, ζ has just two realizations, y_1 and y_2, corresponding to beliefs $\mu_1 < \mu_j$ and $\mu_2 > \mu_j$, respectively. γ_1 has just one realization, corresponding to belief μ_1. γ_2 has two realizations, corresponding to beliefs μ_j and 1. Hence, ξ–the experiment that corresponds to the information ultimately acquired by the receiver following the initial learning and the resulting equilibrium play in the signaling game–has three realizations, corresponding to beliefs μ_1, μ_j and 1. The null-optimal experiment η has two realizations, corresponding to beliefs μ_j and 1.

The resulting distribution over posteriors induced by η has support on 1 and μ_j. Likewise, the resulting distribution over posteriors induced by ξ has support on 1, μ_j and μ_1. Since $\mu_1 < \mu_j$, ξ is more (Blackwell) informative than η and so the receiver prefers ξ–the receiver prefers learning.

Case 3: For prior μ_0, the receiver-optimal equilibrium is one in which the sender mixes in both states. Observe that in this case, the null-optimal experiment begets two posteriors, both of which are in the interior of $[0,1]$.

Let the sender in the high state choose a mixed strategy σ_H and let the sender in the low state choose a mixed strategy σ_L. This begets two posteriors, $\mu_j > \mu_0 > \mu_l$.

The remainder proceeds in the same way as in the second case, any experiment realization that yields a belief in the interval $[\mu_l, \mu_j]$ leads to an equilibrium payoff bounded below by the optimal equilibrium payoff at belief μ_0, and any experiment realization that yields a belief outside that interval leads to an equilibrium payoff bounded below by the pooling payoff.

Ultimately, the null-optimal experiment, η, is less (Blackwell) informative than ξ (defined as in the first two cases), so learning must always be beneficial.

We have gone through each case, and the result is shown. ■

References

Mark Whitmeyer. Bayesian elicitation. *Mimeo*, 2019.