Equivalence of Cheap Talk and Bayesian Persuasion in a Finite Continuous Model

Elliot Lipnowski*†
Columbia University

June 1, 2020

Abstract

This note confirms a conjecture posed by Françoise Forges concerning sender-receiver games of cheap talk with finite parameters. I show that, if the sender's value function is continuous, then she can attain in equilibrium the same payoff as under communication with commitment.

*e.lipnowski@columbia.edu
†I thank Françoise Forges and Doron Ravid for helpful comments.

Model

A sender (S) and receiver (R) interact in a game of strategic information transmission (i.e., cheap talk). A state $\theta \in \Theta$ is realized according to some distribution $\mu_0 \in \Delta\Theta$ and directly observed by the sender, then the sender transmits a message $m \in M$ to the receiver, after which the receiver chooses an action $a \in A$. Each player $i \in \{S, R\}$ enjoys payoff $u_i(a, \theta)$ for some given $u_i : A \times \Theta \to \mathbb{R}$, and wishes to maximize her expectation of this payoff. Throughout, I assume that Θ, A, and M, are all finite nonempty sets with $|M| \geq \min\{|A|, |\Theta|\}$, and μ_0 is of full support.

Define sender and receiver (behavioral) strategy spaces as $\Sigma_S := (\Delta M)^\Theta$ and $\Sigma_R := (\Delta A)^M$, respectively. For each strategy profile $(\sigma_S, \sigma_R) \in \Sigma_S \times \Sigma_R$, define the associated payoff for each player $i \in \{S, R\}$ to be

$$U_i(\sigma_S, \sigma_R) := \sum_{\theta\in\Theta,\ m\in M,\ a\in A} \mu_0(\theta)\ \sigma_S(m|\theta)\ \sigma_R(a|m)\ u_i(a, \theta).$$

I focus attention on Nash equilibria, which I henceforth call "equilibria" for brevity.

Recall the following objects, familiar from the literature on communication games.

$$\begin{aligned}
A_R^* : \Delta\Theta &\rightrightarrows A \\
\mu &\mapsto \operatorname{argmax}_{a\in A} \sum_{\theta\in\Theta} \mu(\theta)\ u_R(a, \theta) \\
A_S^* : \Delta\Theta &\rightrightarrows A \\
\mu &\mapsto \operatorname{argmax}_{a\in A_R^*(\mu)} \sum_{\theta\in\Theta} \mu(\theta)\ u_S(a, \theta) \\
v : \Delta\Theta &\to \mathbb{R} \\
\mu &\mapsto \max_{a\in A_R^*(\mu)} \sum_{\theta\in\Theta} \mu(\theta)\ u_S(a, \theta) = \sum_{\theta\in\Theta} \mu(\theta)\ u_S\left(A_S^*(\mu), \theta\right).
\end{aligned}$$

In words, A_R^* is the receiver's best response correspondence, mapping every belief $\mu \in \Delta\Theta$ to the nonempty set of actions that the receiver finds optimal given a belief μ concerning the state. The correspondence A_S^* describes sender-preferred selection among the receiver's best responses, mapping every belief $\mu \in \Delta\Theta$ to the nonempty set of actions in $A_R^*(\mu)$ that the sender would prefer if, hypothetically, the sender too had belief μ concerning the state. Finally, v is the sender's value function, mapping every belief $\mu \in \Delta\Theta$ to the highest attainable expected sender payoff conditional on state distribution μ and the receiver responding optimally to the same.

Let $\hat{v} : \Delta\Theta \to \mathbb{R}$ denote the concave envelope of v, that is, the pointwise smallest concave function $\Delta\Theta \to \mathbb{R}$ that is everywhere weakly above v. It is well known (e.g., see Kamenica, 2019; Forges, 2020, for related surveys) that $\hat{v}$ summarizes the highest attainable sender value if the sender could commit to a strategy ex ante. In particular, prior work demonstrates that $\hat{v}(\mu_0)$ is the highest sender payoff $U_S(\sigma_S, \sigma_R)$ among any $(\sigma_S, \sigma_R) \in \Sigma_S \times \Sigma_R$ such that σ_R is a best response for the receiver to σ_S. As said optimality property for the receiver must also hold in any equilibrium, the following well-known fact immediately follows.

Fact 1. *No equilibrium generates a sender payoff strictly greater than $\hat{v}(\mu_0)$.*

Proposition 1 below establishes that this upper bound can be attained by some equilibrium whenever v is continuous.

Analysis

Toward proving Proposition 1, I establish the relevant incentive property that continuity of the value function enables.

Lemma 1. *If v is continuous, then A_S^* is upper hemicontinuous.*

Proof. The graph of A_R^* is closed by Berge's theorem, and the set

$$\left\{(\mu, a) \in \Delta\Theta \times A : \sum_{\theta\in\Theta} \mu(\theta)\, u_S(a, \theta) = v(\mu)\right\}$$

is closed because v is continuous. The graph of A_S^*, being the intersection of the above two sets, is closed too. That is, A_S^* is upper hemicontinuous. □

Lemma 2. *Suppose $\mu \in \Delta\Theta$; A_S^* is upper hemicontinuous at μ; $\varphi : \Delta\Theta \to \mathbb{R}$ is affine; $\varphi \geq v$; and $\varphi(\mu) = v(\mu)$. Then, some $\alpha \in \Delta A_S^*(\mu)$ exists such that $\sum_{a\in A} \alpha_\mu(a)\, u_S(a, \theta) \leq \varphi(\delta_\theta)$ for every $\theta \in \Theta$.*

Proof. Upper hemicontinuity of A_S^* at μ delivers a neighborhood $N \subseteq \Delta\Theta$ of μ such that $A_S^*(\tilde{\mu}) \subseteq A_S^*(\mu)$ for every $\tilde{\mu} \in N$. Defining the convex function (as a maximum of affine functions),

$$\begin{aligned} \psi : \Delta\Theta &\to \mathbb{R} \\ \tilde{\mu} &\mapsto \max_{a\in A_S^*(\mu)} \left\{\varphi(\tilde{\mu}) - \sum_{\theta\in\Theta} \tilde{\mu}(\theta)\, u_S(a, \theta)\right\}, \end{aligned}$$

then, every $\tilde{\mu} \in N$ has

$$\psi(\tilde{\mu}) \geq \max_{a \in A_S^*(\tilde{\mu})} \left\{ \varphi(\tilde{\mu}) - \sum_{\theta \in \Theta} \tilde{\mu}(\theta)\, u_S(a, \theta) \right\} = \varphi(\tilde{\mu}) - v(\tilde{\mu}) \geq 0.$$

The function ψ is convex, zero at μ, and nonnegative on a neighborhood of μ; it is therefore globally nonnegative. Finally, the minimax theorem implies

$$\begin{aligned} \max_{\alpha \in \Delta A_S^*(\mu)} \min_{\theta \in \Theta} \left\{ \varphi(\delta_\theta) - \sum_{a \in A_S^*(\mu)} \alpha(a)\, u_S(a, \theta) \right\} &= \min_{\tilde{\mu} \in \Delta\Theta} \max_{a \in A_S^*(\mu)} \left\{ \varphi(\tilde{\mu}) - \sum_{\theta \in \Theta} \tilde{\mu}(\theta)\, u_S(a, \theta) \right\} \\ &= \min_{\tilde{\mu} \in \Delta\Theta} \psi(\tilde{\mu}) \geq 0, \end{aligned}$$

so that some $\alpha \in A_S^*(\mu)$ is as desired. □

Proposition 1. *If v is continuous, then an equilibrium exists that yields sender payoff $\hat{v}(\mu_0)$.*

Proof. By Theorem 23.4 of Rockafellar (1970), the concave function $\hat{v}$ is superdifferentiable at μ_0. That is, some affine $\varphi : \Delta\Theta \to \mathbb{R}$ exists such that $\varphi \geq \hat{v}$ and $\varphi(\mu_0) = \hat{v}(\mu_0)$. Let $D := \{\mu \in \Delta\Theta : \ v(\mu) = \varphi(\mu)\}$. Then, by Lemmata 1 and 2, each $\mu \in D$ admits some $\alpha_\mu \in \Delta A_S^*(\mu)$ such that $\sum_{a \in A} \alpha_\mu(a)\, u_S(a, \theta) \leq \varphi(\delta_\theta)$ for every $\theta \in \Theta$.

I now construct an equilibrium that delivers payoff $\hat{v}(\mu_0)$ to the sender. To that end, fix some $(\sigma_S^*, \sigma_R^*) \in \Sigma_S \times \Sigma_R$ that maximizes U_S subject to the receiver responding optimally to the sender's strategy; hence $U_S(\sigma_S^*, \sigma_R^*) = \hat{v}(\mu_0)$. Without loss of generality, assume $M_0 := \{m \in M : \ \sigma_S^*(m|\theta) = 0 \ \forall \theta \in \Theta\}$ is empty.[1] Below, I show that one can alter the receiver's strategy to construct an equilibrium that generates the same sender payoff.[2]

For each $m \in M$, define the total probability $\tau_m := \sum_{\tilde{\theta} \in \Theta} \mu_0(\tilde{\theta})\, \sigma_S^*(m|\tilde{\theta}) > 0$ of message m, and the posterior belief $\mu_m \in \Delta\Theta$ associated with Bayesian updating from message m, that is, $\mu_m(\theta) := \frac{\mu_0(\theta)\ \sigma_S^*(m|\theta)}{\tau_m}$ for each $\theta \in \Theta$. A given strategy $\sigma_R \in \Sigma_R$ is a best response to σ_S^* if and only if $\sigma_R\left(A_R^*(\mu_m)|m\right) = 1$ for every $m \in M$. Among such strategies σ_R, the sender has a payoff of $U_S(\sigma_S^*, \sigma_R) = \hat{v}(\mu_0)$ if and only if $\sigma_R\left(A_S^*(\mu_m)|m\right) = 1$ for every $m \in M$.

[1] Indeed, if not, one could fix some $m^* \in M \setminus M_0$, replace $\sigma_S^*(\cdot|\theta)$ with $\sigma_S^*(\cdot|\theta) - \frac{\sigma_S^*(m^*|\theta)}{2}\delta_{m^*} + \frac{\sigma_S^*(m^*|\theta)}{2|M_0|}\sum_{m \in M_0}\delta_m$ for every $\theta \in \Theta$, and replace $\sigma_R^*(\cdot|m)$ with $\sigma_R^*(\cdot|m^*)$ for every $m \in M_0$. It is easy to see that the modified strategy profile is incentive compatible for the receiver because the original one is.

[2] As every non-terminal history is reached with strictly positive probability, the resulting Nash equilibrium is, in fact, a sequential equilibrium.

Observe that, because

$$\varphi(\mu_0) = \hat{v}(\mu_0) = \sum_{m \in M} \tau_m v(\mu_m) \le \sum_{m \in M} \tau_m \varphi(\mu_m) = \varphi\left(\sum_{m \in M} \tau_m \mu_m\right) = \varphi(\mu_0),$$

it must be that the inequality holds with equality. Therefore, $\{\mu_m\}_{m \in M} \subseteq D$. Define, then, the receiver strategy $\sigma_R^{**} \in \Sigma_R$ by letting $\sigma_R^{**}(\cdot|m) := \alpha_{\mu_m}$ for every message $m \in M$. By construction, σ_R^{**} is a best response to σ_S^* for the receiver and $U_S(\sigma_S^*, \sigma_R^{**}) = \hat{v}(\mu_0)$. It remains to verify sender optimality. To that end, observe that the highest payoff the sender can achieve from any deviation is

$$\begin{aligned}
\max_{\sigma_S \in \Sigma_S} U_S(\sigma_S, \sigma_R^{**}) &= \sum_{\theta \in \Theta} \mu_0(\theta) \max_{m \in M} \sum_{a \in A} \alpha_{\mu_m}(a)\, u_S(a, \theta) \\
&\le \sum_{\theta \in \Theta} \mu_0(\theta) \max_{m \in M} \varphi(\delta_\theta) \\
&= \varphi\left(\sum_{\theta \in \Theta} \mu_0(\theta)\, \delta_\theta\right) \\
&= \varphi(\mu_0) = \hat{v}(\mu_0) = U_S(\sigma_S^*, \sigma_R^{**}).
\end{aligned}$$

Therefore, $(\sigma_S^*, \sigma_R^{**})$ is an equilibrium as required. □

The following example, presented informally, shows that one must allow both players to mix for Proposition 1's conclusion to hold.[3]

Example 1. *Suppose* $A = \{0, 2, 4\}$, $\Theta = \{0, 4\}$, $M = \{0, 3\}$, μ_0 *is uniform,* $u_R(a, \theta) = -(a - \theta)^2$, *and*

$$u_S(a, \theta) = \begin{cases} a & : \ a \neq 2 \\ 2(\theta - 1) & : \ a = 2. \end{cases}$$

Observe first that this specification satisfies the continuity condition. Identifying a belief $\mu \in \Delta\Theta$ *with its associated expectation in* $[0, 4]$ *of the state, the receiver optimally chooses action* 0 *when* $\mu \in [0, 1]$*, optimally chooses action* 2 *when* $\mu \in [1, 3]$*, and optimally chooses action* 4 *when* $\mu \in [3, 4]$*. The sender's value function is then given by* $v(\mu) = \min\{\max\{0, 2(\mu - 1)\}, 4\}$*, which is continuous.*

The unique equilibrium (up to switching the two messages) attaining sender payoff $\hat{v}(\mu_0)$ *has the sender transmitting message* 3 *when the state is* 4*, the sender transmitting message* 0 *with probability* $\frac{2}{3}$ *and message* 3 *with probability* $\frac{1}{3}$ *when the state is* 0*, the receiver choosing*

[3] Example 1 from Forges (2020), with two states and four actions, illustrates the same.

action 0 *when the message is* 0*, and the receiver choosing action* 2 *with probability* $\frac{2}{3}$ *and action* 4 *with probability* $\frac{1}{3}$ *when the message is* 3*. In particular, neither players employs a pure strategy. Uniqueness of the sender strategy obtains because the two induced posterior beliefs (*0 *from message* 0 *and* 3 *from message* 3*) are the only two posterior beliefs at which the value function coincides with its concave envelope's tangent line at the prior. The receiver, who has one optimal action following message* 0 *and two optimal actions following message* 3*, must mix with the described probabilities to make the sender indifferent between the two messages when the state is* 0*.*

I conclude with a stronger sufficient condition for the main result to apply, also proposed by Françoise Forges. The condition says that, whenever the receiver has multiple best responses to a given belief, the sender would (at the same belief) be indifferent between these receiver best responses.

Corollary 1. *If* $A_S^* = A_R^*$*, then an equilibrium exists that yields sender payoff* $\hat{v}(\mu_0)$*.*

The corollary can be derived from Proposition 1 by showing that v is continuous whenever $A_S^* = A_R^*$. Alternatively, modifying the proof of Proposition 1, one can invoke Lemma 2 without appealing to Lemma 1 in this case, by appealing to Berge's theorem.

References

Forges, Françoise. 2020. "Games with incomplete information: from repetition to cheap talk and persuasion." *Annals of Economics and Statistics* (137):3–30.

Kamenica, Emir. 2019. "Bayesian Persuasion and Information Design." *Annual Review of Economics* 11 (1).

Rockafellar, R Tyrrell. 1970. *Convex Analysis*. 28. Princeton University Press.