Perfect Bayesian Persuasion*

Elliot Lipnowski†
Columbia University

Doron Ravid
University of Chicago

Denis Shishkin
UC San Diego

June 10, 2021

Abstract

A sender publicly commits to an experiment to inform a receiver's decision. We study attainable sender payoffs, accounting for her incentives at the experiment choice stage, and not presupposing a receiver tie-breaking rule when indifferent. We characterize when the sender has a unique equilibrium payoff, which therefore coincides with her optimal value in Kamenica and Gentzkow (2011). A sufficient condition is that every action which is a receiver best response to some belief over a set of states is a unique best response to some other such belief—a generic property in the finite case.

*Lipnowski and Ravid acknowledge support from the National Science Foundation (grant SES-1730168).
†`e.lipnowski@columbia.edu`, `dravid@uchicago.edu`, `dshishkin@ucsd.edu`

1. Introduction

In this paper, we concern ourselves with the model of communication with commitment from Kamenica and Gentzkow (2011), hereafter KG . A receiver (R, he) must choose an action $a \in A$, but a sender (S, she) controls R's available information about a payoff-relevant state $\theta \in \Theta$, which is distributed according to some prior probability distribution $\mu_0 \in \Delta\Theta$. Specifically, S first announces a Blackwell experiment concerning the state, that is, a measurable function $\psi : \Theta \to \Delta M$ for a given space M of messages. Then, after observing both ψ and a realized message $m \in M$, R chooses an action. Each player $i \in \{S, R\}$ seeks to maximize the expectation of an objective $u_i(a, \theta)$. Finally, we maintain the technical assumptions that both A and Θ are nonempty compact metrizable spaces; that M is a Polish space of at least the cardinality of A;[1] and that the objectives $u_S, u_R : A \times \Theta \to \mathbb{R}$ are continuous.

The main result of KG is a characterization of the S-optimal equilibrium payoff. We defer a formal definition of equilibrium to the following section, but a brief overview of KG's analysis is in order. They adopt a belief-based approach, casting S's optimization problem as one of directly choosing $p \in \Delta\Delta\Theta$, the ex-ante distribution of R's posterior belief μ concerning the state. Because R is Bayesian (and S's experiment choice is made in ignorance of the state), it must be that p belongs to $\mathcal{I}(\mu_0)$, the set of belief distributions that average to the prior; KG term this condition Bayes plausibility. But what payoff does S derive from a given Bayes-plausible belief distribution? By R rationality, R will choose from his best responses $A_R^*(\mu) \subseteq A$ whenever her posterior belief is μ. S's expected payoff from such an action $a \in A_R^*(\mu)$ is then equal to $\int u_S(a, \cdot)\, \mathrm{d}\mu$. Given KG's focus on S-optimal equilibrium, they can assume without loss that R always breaks any indifferences in S's favor. Thus, KG can summarize S's payoff from inducing belief μ as $v(\mu) := \max_{a \in A_R^*(\mu)} \int u_S(a, \cdot)\, \mathrm{d}\mu$. We call v the *value function*. Hence, S's best equilibrium value is given by $\hat{v}(\mu_0) = \max_{p \in \mathcal{I}(\mu_0)} \int v \, \mathrm{d}p$.

But what happens if R may choose best responses that differ from those S would prefer? If, in the worst case, R always chooses S's least favorite of his best responses, then from inducing R belief μ, S can only expect a payoff of $w(\mu) := \min_{a \in A_R^*(\mu)} \int u_S(a, \cdot)\, \mathrm{d}\mu$. Accordingly, S would have a profitable deviation if her payoff were ever strictly below $\hat{w}(\mu_0) = \sup_{p \in \mathcal{I}(\mu_0)} \int w \, \mathrm{d}p$. In what follows, we verify that this payoff lower bound is the only additional constraint imposed by S's experiment-choice incentives. Using this result,

[1] One could, similar to KG, allow S to choose the message space for the experiment as well, and require it to be a measurable subset of the universal space M (or a finite subset if $\min\{|A|, |\Theta|\}$ were finite). Our assumption that M is large enough (which would, for example, always be satisfied by $M = [0, 1]$) ensures that exogenous restrictions on how much information can be conveyed do not constrain attainable S payoffs. If Θ were finite, an alternative restriction that $|M| \geq |\Theta|$ would deliver identical results via Carathéodory's theorem.

we go on to fully characterize when S has a unique equilibrium payoff, and to provide meaningful sufficient conditions for the same.

2. The Equilibrium Payoff Set

We now formally define an equilibrium concept for the persuasion game. Note two features of the definition. First, while R must respond optimally to his belief, we make no direct assumption on which of his best responses he chooses in the case that he is indifferent.[2] Second we explicitly include an optimality condition for S at the experiment-choice stage (which would have no bite under S-favorable tie-breaking, and so is rarely included in the literature).

Definition 1. *Let Ψ denote the set of all measurable functions $\tilde{\psi} : \Theta \to \Delta M$ (a.k.a. experiments), which we view as a measurable space endowed with the discrete σ-algebra. A sender strategy is an experiment $\psi \in \Psi$; a receiver strategy is a measurable function $\rho : \Psi \times M \to \Delta A$; and a receiver belief map is a measurable function $\beta : \Psi \times M \to \Delta\Theta$. An* ***equilibrium*** *is a triple of such maps $\langle \psi, \rho, \beta \rangle$ such that*

1. *The sender's choice satisfies*

$$\psi \in \text{argmax}_{\tilde{\psi} \in \Psi} \int_\Theta \int_M \int_A u_S(a, \theta)\, d\rho(a|\tilde{\psi}, m)\, d\tilde{\psi}(m|\theta)\, d\mu_0(\theta);$$

2. *Every $\tilde{\psi} \in \Psi$ and $m \in M$ have*

$$\rho\left(\text{argmax}_{a \in A} \int_\Theta u_R(a, \theta)\, d\beta(\theta|\tilde{\psi}, m) \,\middle|\, \tilde{\psi}, m\right) = 1;$$

3. *Every $\tilde{\psi} \in \Psi$, Borel $\hat{M} \subseteq M$, and Borel $\hat{\Theta} \subseteq \Theta$ have*

$$\int_\Theta \int_{\hat{M}} \beta(\hat{\Theta}|\tilde{\psi}, m)\, d\tilde{\psi}(m|\theta)\, d\mu_0(\theta) = \int_{\hat{\Theta}} \tilde{\psi}(\hat{M}|\theta)\, d\mu_0(\theta).$$

In such a case, we say the induced ***equilibrium sender payoff*** *is*

$$\int_\Theta \int_M \int_A u_S(a, \theta)\, d\rho\,(a|\psi, m)\; d\psi(m|\theta)\, d\mu_0(\theta).$$

[2] Still, it will often be the case that *mutual* best response requires that R break indifferences in S's favor on the path of play, just as a recipient of a zero offer must accept the offer in subgame-perfect equilibrium of an ultimatum game.

The interpretation is as follows. First, S publicly chooses an experiment $\tilde{\psi} \in \Psi$.[3] The experiment then produces a message $m \in M$ that R observes. Then, R updates his beliefs according to the message and the chosen experiment, and chooses an action $a \in A$. We require that S only choose experiments that maximize her expected payoffs, that R (having seen the realized experiment and message) only choose actions that maximize his expected payoffs with respect to his belief about the payoff state, and that R's beliefs conform to Bayesian updating.[4]

In what follows, we document the set of attainable equilibrium S payoffs, with a particular focus on understanding when it is unique.

Remark 1. *Although our focus is on equilibrium S payoffs rather than behavior, our results have natural implications for behavior as well. In particular, when S's equilibrium payoff is unique, our results imply that R breaks indifferences in S's favor with probability 1 on path in every equilibrium. Hence, in this case, the results of KG (and many subsequent papers surveyed in Kamenica, 2019) are robust to allowing arbitrary tie-breaking for R and to accounting for S's experiment-choice incentives.*

2.1. Characterizing Equilibrium Payoffs

We begin by stating a characterization of the equilibrium S payoff set as a function of the parameters of our game. This set is a compact interval, with highest value equal to KG's commitment solution, and lowest value equal to the supremum value S can guarantee when R breaks her indifferences adversarially. In the special case in which the state space is finite, this result is exactly Proposition 1 from Wu (2020). Although no substantive new arguments are required for the general case, we include a proof for the sake of completeness.

Proposition 1. *The set of equilibrium S payoffs is* $[\hat{w}(\mu_0), \hat{v}(\mu_0)]$.

Necessity is essentially immediate, and the proof of sufficiency is constructive. By degrading information from an S-optimal (under favorable tie-breaking) experiment and allowing for R to mix among optimal choices in the degraded experiment, one can find an experiment for S to choose and R best response to target any payoff in the given interval.

[3] One could easily extend the model to allow S to mix over experiment choice. Doing so would entail added notational burden but would have no effect on the resulting S payoff set because the experiment choice is public, not informed by private information, and not simultaneous to any other decisions.

[4] Moreover, we assume that S cannot signal what she does not know. Indeed, our Bayesian condition implies that every $\tilde{\psi} \in \Psi$ and Borel $\hat{\Theta} \subseteq \Theta$ have $\int_\Theta \int_M \beta(\hat{\Theta}|\tilde{\psi}, m)\, \mathrm{d}\tilde{\psi}(m|\theta)\, \mathrm{d}\mu_0(\theta) = \mu_0(\hat{\Theta})$, so that the chosen experiment alone does not cause belief updating by R about the payoff state.

Then, having R break indifference adversarially to S following off-path experiment choices ensures that this experiment choice is indeed optimal for S.

It is apparent that Proposition 1 depends only on the value correspondence $V = [w, v]$, and moreover (as is clear from our proof) the only substantive property required of the environment is that the attainable S payoffs from R responding optimally to a given belief be convex.[5] In addition to making Proposition 1 more tractable to apply, this feature also expands its applicability beyond the basic model we have considered. For example, the proposition can be applied to settings in which a receiver is subject to independent private payoff shocks. Additionally, the proposition applies to public persuasion of a set of agents who play a game, so long as the set of induced payoffs for the sender is convex for every public belief. The latter condition holds, for instance, if the receivers have a unique equilibrium of their induced game, or if the receivers observe a rich public randomization device after the experiment choice but before their gameplay.

3. Equilibrium Payoff Uniqueness

In this section we ask, when does S have a unique equilibrium payoff? Whenever she does, the traditional analysis that focuses on S-optimal equilibrium (and so assumes S-favorable tie-breaking by R) is essentially without loss.

As a starting observation, because uniqueness follows directly from Proposition 1 whenever $v = w$, a sufficient condition for S to have a unique equilibrium payoff is immediate.

Corollary 1. *S has a unique equilibrium payoff if, at any belief, S is indifferent between between all of R's best responses.*

Although restrictive, the above condition nevertheless captures many cases of interest. For example, if the action space is a convex subset of some linear space with R's payoff being strictly concave in his action (e.g., Crawford and Sobel, 1982; Chakraborty and Harbaugh, 2010), then he has a unique best response to every belief, and so the corollary applies.

The following result gives an alternative characterization of when the sender attains her Bayesian persuasion value in all equilibria at all priors. It says such uniqueness always holds if and only if information can always serve as a stand-in for favorable tie-breaking.

[5] Our analysis also uses the fact that V is nonempty-compact-valued and upper hemicontinuous, and that the set of optimal R choices is a weakly measurable correspondence of his belief. These features are satisfied in all past applications of which we are aware.

Proposition 2. *S has a unique equilibrium payoff for every prior (holding other parameters fixed) if and only if* $\hat{w} \geq v$.

The proof shows, given an arbitrary experiment choice, how it can be augmented with additional information to replicate favorable tie-breaking.

3.1. "Direct" Sufficient Conditions for Uniqueness

The above exact conditions for payoff uniqueness were expressed in terms of the derived objects v, w, $\hat{v}$, and $\hat{w}$. While these functions are primitive to the environment, it is desirable to find more interpretable sufficient conditions on the players' preferences. We now develop sufficient conditions "directly" on R's preferences that ensure S has a unique equilibrium payoff. The following condition, satisfied in many applications of interest, is our key such condition.

Definition 2. *Say the* ***perturbed unique best response (PUBR)*** *property holds if, for any* $\mu \in \Delta\Theta$ *and* $a \in A^*_R(\mu)$*, some* $\mu' \in \Delta\Theta$ *exists such that* $A^*(\mu') = \{a\}$ *and* $supp(\mu') \subseteq supp(\mu)$*.*

The above property says that any action that is optimal for R at some belief is in fact uniquely optimal at some alternative belief, where the alternative belief can be assumed to rule out any open neighborhood of states that the original belief rules out.[6] Note that the PUBR property depends only on R's preferences (together with the spaces of states and actions), not on the prior or on S's preferences.

The next result shows the PUBR property is sufficient to guarantee equilibrium selection. In light of Proposition 2, it is enough to show that information can serve as a stand-in for selection. The key observation is that, for any given belief, the PUBR property helps locate nearby beliefs at which *every* best response gives S an expected payoff nearly as high as favorable tie-breaking would at the belief itself. Using such beliefs, we can construct slightly informative experiments that similarly give S a high payoff under all R best responses.

Proposition 3. *Under the perturbed unique best response property, S has a unique equilibrium payoff.*[7]

[6] Note, this condition is strictly stronger than the requirement that R has no duplicate actions. The latter condition is insufficient for ensuring uniqueness, as witnessed by $A = \Theta = \{0, 1\}$, $\mu_0 = \frac{1}{2}$, $u_S(a, \theta) = -a$, and $u_R(a, \theta) = a\theta$.

[7] A converse trivially holds in the case that A is finite. If the PUBR property fails for u_R, as witnessed by some $a \in A$ and $\mu \in \Delta\Theta$, then equilibrium uniqueness fails for prior μ and S preferences $u_S(\tilde{a}, \tilde{\theta}) := \mathbf{1}_{\tilde{a}=a}$. Indeed, such preferences generate $v(\mu) = 1$ and $\hat{w} = 0$, so that S can get any payoff in $[0, 1]$ in equilibrium.

The following result states that, generically, unique equilibrium S payoffs obtain in finite environments.[8] The proof shows that the PUBR property is generic, which suffices by the previous result. Intuitively, a failure of the PUBR property says that, for some fixed action and fixed set of states, R's highest possible expected payoff benefit from using said action rather than another is exactly zero—a knife-edge condition.

Proposition 4. *If A and Θ are finite, then an open dense $\mathcal{U}_R \subseteq \mathbb{R}^{A\times\Theta}$ of full Lebesgue measure exists such that S has a unique equilibrium payoff (for any prior and any S preferences) as long as $u_R \in \mathcal{U}_R$.*

References

Aliprantis, Charalambos D and Kim Border. 2006. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer Science & Business Media.

Chakraborty, Archishman and Rick Harbaugh. 2010. "Persuasion by cheap talk." *American Economic Review* 100 (5):2361–2382.

Crawford, Vincent P and Joel Sobel. 1982. "Strategic information transmission." *Econometrica* 50 (6):1431–1451.

de Clippel, Geoffroy. 2008. "An axiomatization of the inner core using appropriate reduced games." *Journal of Mathematical Economics* 44 (3-4):316–323.

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

Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review* 101 (6):2590–2615.

Li, Fei and Peter Norman. 2020. "Sequential persuasion." *Available at https://econtheory.org/ojs/index.php/te/article/viewForthcomingFile/3474/27882/1* .

Lipnowski, Elliot and Laurent Mathevet. 2018. "Disclosure to a psychological audience." *American Economic Journal: Microeconomics* 10 (4):67–93.

Wu, Wenhao. 2020. "Sequential Bayesian Persuasion." .

[8] A related result follows directly from Proposition 3 of Li and Norman (2020) in the generic finite case. That result implies that all equilibria with S-favorable tie-breaking result in the same state-contingent action distribution. Hence, *assuming* R breaks indifferences in S's favor generically yields a behavioral uniqueness property, not just payoff uniqueness.

A. Proofs

Before proceeding to formal proofs, we review for convenience several key notations.

$$
\begin{aligned}
A_R^* : \Delta\Theta &\rightrightarrows A \\
\mu &\mapsto \operatorname{argmax}_{a\in A} \int u_R(a,\cdot)\,\mathrm{d}\mu \\
V : \Delta\Theta &\to \mathbb{R} \\
\mu &\mapsto \operatorname{co}\left\{\int u_S(a,\cdot)\,\mathrm{d}\mu :\ a \in A_R^*(\mu)\right\} \\
v : \Delta\Theta &\to \mathbb{R} \\
\mu &\mapsto \max V(\mu) \\
w : \Delta\Theta &\to \mathbb{R} \\
\mu &\mapsto \min V(\mu) \\
\mathcal{I} : \Delta\Theta &\rightrightarrows \Delta\Delta\Theta \\
\mu &\mapsto \left\{p \in \Delta\Delta\Theta :\ \int \tilde{\mu}\,\mathrm{d}p(\tilde{\mu}) = \mu\right\} \\
\hat{v} : \Delta\Theta &\to \mathbb{R} \\
\mu &\mapsto \max_{p\in\mathcal{I}(\mu)} \int v\,\mathrm{d}p \\
\hat{w} : \Delta\Theta &\to \mathbb{R} \\
\mu &\mapsto \sup_{p\in\mathcal{I}(\mu)} \int w\,\mathrm{d}p.
\end{aligned}
$$

Proof of Proposition 1. To begin, we recall some well-known facts about experiments and Bayesian updating—which collectively tell us that S choosing from Ψ and choosing from $\mathcal{I}(\mu_0)$ are equivalent formalisms. First, any experiment $\tilde{\psi} \in \Psi$ admits some compatible belief map, that is, some measurable $\tilde{\beta} = \tilde{\beta}_{\tilde{\psi}} : M \to \Delta\Theta$ such that every Borel $\hat{M} \subseteq M$ and Borel $\hat{\Theta} \subseteq \Theta$ have $\int_\Theta \int_{\hat{M}} \tilde{\beta}(\hat{\Theta}|m)\,\mathrm{d}\tilde{\psi}(m|\theta)\,\mathrm{d}\mu_0(\theta) = \int_{\hat{\Theta}} \tilde{\psi}(\hat{M}|\theta)\,\mathrm{d}\mu_0(\theta)$. Second, given $\tilde{\psi} \in \Psi$ if we define the belief distribution $p_{\tilde{\psi},\tilde{\beta}} \in \Delta\Delta\Theta$ via $p_{\tilde{\psi},\tilde{\beta}}(D) := \int_\Theta \tilde{\psi}\left(\tilde{\beta}^{-1}(D) \,\middle|\, \theta\right)\,\mathrm{d}\mu_0(\theta)$ for each Borel $D \subseteq \Delta\Theta$, then $p_{\tilde{\psi},\tilde{\beta}} = p_{\tilde{\psi},\tilde{\beta}'}$ for any two such compatible $\tilde{\beta}$ and $\tilde{\beta}'$. We therefore refer to the associated belief distribution simply as $p_{\tilde{\psi}}$. Third, every $\tilde{\psi} \in \Psi$ has $p_{\tilde{\psi}} \in \mathcal{I}(\mu_0)$. Fourth, every $p \in \mathcal{I}(\mu_0)$ with $|\mathrm{supp}(p)| \leq |M|$ admits some $\tilde{\psi}_p \in \Psi$ such that $p_{\tilde{\psi}_p} = p$.

Now we proceed to show $s \in [\hat{w}(\mu_0), \hat{v}(\mu_0)]$ is necessary and sufficient for s to be an equilibrium S payoff.

First, to see the condition is necessary, fix an arbitrary equilibrium $\langle \psi, \rho, \beta \rangle$, and let $s \in \mathbb{R}$ be the induced S payoff; we will show $s \in [\hat{w}(\mu_0), \hat{v}(\mu_0)]$. For any $\tilde{\psi} \in \Psi$ and $m \in M$ the R optimality condition implies $\rho\left(A_R^*\left(\beta(\cdot|\tilde{\psi}, m)\right) \,\middle|\, \tilde{\psi}, m\right) = 1$, so that $\int_A u_S(a,\theta)\,\mathrm{d}\rho\left(a\middle|\tilde{\psi}, m\right) \in$

$V\left(\beta(\cdot|\tilde{\psi},m)\right)$. Therefore, any $\tilde{\psi}\in\Psi$ has

$$\begin{aligned}
\int_\Theta\int_M\int_A u_S(a,\theta)\,\mathrm{d}\rho(a|\tilde{\psi},m)\,\mathrm{d}\tilde{\psi}(m|\theta)\,\mathrm{d}\mu_0(\theta) &= \int_\Theta\int_M V\left(\beta(\cdot|\tilde{\psi},m)\right)\,\mathrm{d}\tilde{\psi}(m|\theta)\,\mathrm{d}\mu_0(\theta)\\
&= \int_{\Delta\Theta} V\,\mathrm{d}p_{\tilde{\psi}}\\
&= \left[\int_{\Delta\Theta} w\,\mathrm{d}p_{\tilde{\psi}},\ \int_{\Delta\Theta} v\,\mathrm{d}p_{\tilde{\psi}}\right]\\
&\subseteq \left[\int_{\Delta\Theta} w\,\mathrm{d}p_{\tilde{\psi}},\ \hat{v}(\mu_0)\right].
\end{aligned}$$

Hence, $s\le\hat{v}(\mu_0)$. Moreover, any $p\in\mathcal{I}(\mu_0)$, taking $\tilde{\psi}=\psi_p$ implies (by S rationality)

$$\begin{aligned}
&\int_\Theta\int_M\int_A u_S(a,\theta)\,\mathrm{d}\rho(a|\psi,m)\,\mathrm{d}\psi(m|\theta)\,\mathrm{d}\mu_0(\theta)\\
\ge\ &\int_\Theta\int_M\int_A u_S(a,\theta)\,\mathrm{d}\tilde{\rho}(a|\tilde{\psi},m)\,\mathrm{d}\tilde{\psi}(m|\theta)\,\mathrm{d}\mu_0(\theta)\\
\ge\ &\int_{\Delta\Theta} w\,\mathrm{d}p.
\end{aligned}$$

Applying this observation to every $p\in\mathcal{I}(\mu_0)$ implies $s\ge\hat{w}(\mu_0)$.

Conversely, take any $s\in[\hat{w}(\mu_0),\hat{v}(\mu_0)]$. Letting $p_1\in\mathcal{I}(\mu_0)$ with $\int_{\Delta\Theta} v\,\mathrm{d}p_v=\hat{v}(\mu_0)$ and $|\mathrm{supp}(p_v)|\le|A|$—which exists by the revelation principle—define $p_\lambda:=\int\delta_{\lambda\mu+(1-\lambda)\mu_0}\,\mathrm{d}p_1(\mu)\in\mathcal{I}(\mu_0)$ for each $\lambda\in[0,1]$; observe $|\mathrm{supp}(p_\lambda)|\le|\mathrm{supp}(p_1)|\le|M|$. As $\lambda\mapsto p_\lambda$ is continuous, it follows that the $\lambda\mapsto\int V\,\mathrm{d}p_\lambda$ is nonempty-compact-convex-valued and upper hemicontinuous because V is. Moreover,

$$\int V\,\mathrm{d}p_0=V(\mu_0)\ni w(\mu_0)\le s\le\hat{v}(\mu_0)\in\int V\,\mathrm{d}p_1.$$

The intermediate value theorem for correspondences (e.g., Lemma 2 from de Clippel, 2008) therefore delivers some $\lambda\in[0,1]$ such that $s\in\int V\,\mathrm{d}p_\lambda$. Some measurable $\zeta:\Delta\Theta\to[0,1]$ then exists such that $s=\int[(1-\zeta)w+\zeta v]\,\mathrm{d}p_\lambda$. By the measurable maximum theorem (Theorem 18.19 from Aliprantis and Border, 2006), a pair of measurable functions $\alpha_w,\alpha_v:\Delta\Theta\to\Delta A$ exist such that, each $\mu\in\Delta\Theta$ has $\int_{A\times\Theta}u_S\,\mathrm{d}\left[\alpha_w(\cdot|\mu)\otimes\mu\right]=w(\mu)$, $\int_{A\times\Theta}u_S\,\mathrm{d}\left[\alpha_v(\cdot|\mu)\otimes\mu\right]=v(\mu)$, and $\alpha_w\left(A_R^*(\mu)\middle|\mu\right)=\alpha_v\left(A_R^*(\mu)\middle|\mu\right)=1$. With these objects in hand, we can define our candidate $\psi:\Theta\to\Delta M$, $\rho:\Psi\times M\to\Delta A$, and $\beta:\Psi\times M\to\Delta\Theta$ via

$$\begin{aligned}
\psi(\cdot|\theta) &:= \tilde{\psi}_{p_\lambda}\\
\beta(\cdot|\tilde{\psi},m) &:= \tilde{\beta}_{\tilde{\psi}}(\cdot|m)\\
\rho(\cdot|\tilde{\psi},m) &:= \begin{cases}(1-\zeta)\alpha_w\left(\cdot\,\middle|\,\beta(\cdot|\tilde{\psi},m)\right)+\zeta\alpha_v\left(\cdot\,\middle|\,\beta(\cdot|\tilde{\psi},m)\right) & :\ \tilde{\psi}=\tilde{\psi}_{p_\lambda}\\ \alpha_w\left(\cdot\,\middle|\,\beta(\cdot|\tilde{\psi},m)\right) & :\ \tilde{\psi}\ne\tilde{\psi}_{p_\lambda}.\end{cases}
\end{aligned}$$

It is immediate from the construction that all three maps are measurable and that R rationality and the Bayesian property are both satisfied. Moreover, direct computation shows that choosing experiment $\tilde{\psi} \in \Psi$ gives S a continuation payoff of

$$\int_\Theta \int_M \int_A u_S(a,\theta)\, \mathrm{d}\rho(a|\tilde{\psi},m)\, \mathrm{d}\tilde{\psi}(m|\theta)\, \mathrm{d}\mu_0(\theta) = \begin{cases} s & : \tilde{\psi} = \tilde{\psi}_{p_\lambda} \\ \int_{\Delta\Theta} w\, \mathrm{d}p_{\tilde{\psi}} & : \tilde{\psi} \neq \tilde{\psi}_{p_\lambda} \end{cases}$$

Therefore, S gets payoff s if the triple is an equilibrium. Finally, the triple is indeed an equilibrium: S rationality is confirmed because $s \geq \hat{w}(\mu_0) \geq \int w\, \mathrm{d}p_{\tilde{\psi}}$ for every alternative $\tilde{\psi} \in \Psi$. □

Proof of Proposition 2. By Proposition 1, it suffices to show that $\hat{w} = \hat{v}$ if and only if $\hat{w} \geq v$. As it is immediate that $\hat{v} \geq \hat{w}$ and $\hat{v} \geq v$, we need only establish that $\hat{w} \geq \hat{v}$ if $\hat{w} \geq v$.[9] Supposing $\hat{w} \geq v$, let us show, for arbitrary prior $\mu_0 \in \Delta\Theta$ and $\epsilon > 0$, that $\hat{w}(\mu_0) > \hat{v}(\mu_0) - \epsilon$.

Defining the set $\mathcal{O} := \left\{p \in \Delta\Delta\Theta : \int w\, \mathrm{d}p - v\left(\int \mu\, \mathrm{d}p(\mu)\right) > -\epsilon\right\}$, let us show that the correspondence $\mathcal{O} \cap \mathcal{I} : \Delta\Theta \rightrightarrows \Delta\Delta\Theta$ admits a measurable selector. First, because the barycentre map is continuous, the correspondence $\mathcal{I}$ is upper hemicontinuous and compact-valued, hence weakly measurable. Next, because the barycentre map is continuous and w and v are lower and upper semicontinuous, respectively, the set $\mathcal{O}$ is open. But then, for every open $\mathcal{Q} \subseteq \Delta\Delta\Theta$, the set $\{\mu \in \Delta\Theta : \mathcal{I}(\mu) \cap \mathcal{O} \cap \mathcal{Q}\}$ is measurable, implying $\mathcal{O} \cap \mathcal{I}$ is weakly measurable. Moreover, that $\hat{w} \geq v$ implies $\mathcal{O} \cap \mathcal{I}$ is nonempty-valued as well. Therefore, the Kuratowski and Ryll-Nardzewski measurable selection theorem delivers a measurable selector φ of $\mathcal{O} \cap \mathcal{I}$.

We can now establish that $\hat{w}(\mu_0) > \hat{v}(\mu_0) - \epsilon$. For any $p \in \mathcal{I}(\mu_0)$, that $q_p := \int \varphi\, \mathrm{d}p \in \mathcal{I}(\mu_0)$ too implies

$$\hat{w}(\mu_0) \geq \int w\, \mathrm{d}q_p = \int \int w\, \mathrm{d}\varphi(\cdot|\mu)\, \mathrm{d}p(\mu) > \int [v(\mu) - \epsilon]\, \mathrm{d}p(\mu) = \int v\, \mathrm{d}p - \epsilon.$$

Maximizing over $p \in \mathcal{I}(\mu_0)$ yields $\hat{w}(\mu_0) > \hat{v}(\mu_0) - \epsilon$, establishing the claim. □

Proof of Proposition 3. Suppose the PUBR property holds, and take an arbitrary $\bar{\mu} \in \Delta\Theta$ and $\epsilon > 0$. We will show that $\hat{w}(\bar{\mu}) \geq s := v(\bar{\mu}) - \epsilon$, which will establish the result by Proposition 2. To that end, let $a \in A_R^*(\bar{\mu})$ be such that $\int u_S(a,\cdot)\, \mathrm{d}\bar{\mu} = v(\bar{\mu})$. Letting $\hat{\Theta} := \mathrm{supp}(\bar{\mu})$, the PUBR property then delivers some $\mu' \in \Delta\hat{\Theta}$ with $A^*(\mu') = \{a\}$. From linearity of expected utility in beliefs, it follows that $A_R^*(\mu) = \{a\}$ for any proper convex combination μ of μ' and $\bar{\mu}$. We may therefore assume without loss, replacing μ' with such a convex combination close enough to $\bar{\mu}$, that $\int u_S(a,\cdot)\, \mathrm{d}\mu' > s$.

Define now the sets of beliefs,

$$D := \mathrm{co}\{\mu \in \Delta\hat{\Theta} : \ w(\mu) > s\} \text{ and } \tilde{D} := \left\{\tilde{\mu} \in \Delta\hat{\Theta} : \ [\mathrm{co}\{\bar{\mu},\tilde{\mu}\}] \setminus \{\bar{\mu}\} \subseteq D\right\}.$$

[9] This result is immediate for the case that Θ is finite. In this case, $\hat{v}$ is the pointwise smallest concave function above v, so that $\hat{w} \geq \hat{v}$ if $\hat{w} \geq v$. In the general case, $\hat{v}$ is the pointwise smallest concave *upper semicontinuous* function above v, so that $\hat{w} \geq \hat{v}$ would directly follow if we knew $\hat{w}$ were upper semicontinuous. Instead of proving such upper semicontinuity, we show directly that $\hat{w} \geq \hat{v}$.

Let us establish that $\tilde{D}$ contains a nonempty set that is relatively open in $\Delta\hat{\Theta}$. Our choice of μ' ensures that $\mu' \in \tilde{D}$. We will now show that, in the relative topology on $\Delta\hat{\Theta}$, the belief μ' is in fact interior in $\tilde{D}$. To that end, observe first that D is open in $\Delta\hat{\Theta}$, because w is lower semicontinuous and the convex hull of an open set is open. As $\mu' \in \tilde{D} \subseteq D$, some open neighborhood $N \subseteq \mathrm{ca}(\Theta)$ of the zero measure therefore exists such that $(\mu' + N) \cap \Delta\hat{\Theta} \subseteq D$. To see that $(\mu' + \frac{1}{2}N) \cap \Delta\hat{\Theta} \subseteq \tilde{D}$, consider an arbitrary $\eta \in N$ such that $\mu' + \frac{1}{2}\eta \in \Delta\hat{\Theta}$ and an arbitrary $\lambda \in (0, 1]$. Observe that

$$(1-\lambda)\bar{\mu} + \lambda(\mu' + \tfrac{1}{2}\eta) = (1 - \tfrac{\lambda}{2})\left(\frac{1-\lambda}{1-\frac{\lambda}{2}}\,\bar{\mu} + \frac{\frac{\lambda}{2}}{1-\frac{\lambda}{2}}\,\mu'\right) + \tfrac{\lambda}{2}\,(\mu' + \eta),$$

which is in D because $\mu' \in \tilde{D}$ and D is convex.

Above, we showed that $\tilde{D}$ has nonempty interior, in the relative topology on $\Delta\hat{\Theta}$. Because the set $\{\mu \in \Delta\hat{\Theta} : \gamma\mu \le \bar{\mu}$ for some $\gamma \in (0, 1)\}$ is dense in $\Delta\hat{\Theta}$ by Lemma 2 from Lipnowski and Mathevet (2018), it follows that some $\tilde{\mu} \in \tilde{D}$ and $\gamma \in (0, 1)$ exist with $\gamma\tilde{\mu} \le \bar{\mu}$. Finally, because $\hat{w}$ is concave, $\tilde{\mu} \in \tilde{D}$, and

$$\bar{\mu} = \tfrac{\gamma}{\gamma+\lambda(1-\gamma)}\left[(1-\lambda)\bar{\mu} + \lambda\tilde{\mu}\right] + \tfrac{\lambda(1-\gamma)}{\gamma+\lambda(1-\gamma)}\left[\frac{\bar{\mu} - \gamma\tilde{\mu}}{1-\gamma}\right],$$

it follows that every $\lambda \in (0, 1]$ has

$$\begin{aligned} \hat{w}(\bar{\mu}) \quad &\ge \quad \tfrac{\gamma}{\gamma+\lambda(1-\gamma)}\hat{w}\left((1-\lambda)\bar{\mu} + \lambda\tilde{\mu}\right) + \tfrac{\lambda(1-\gamma)}{\gamma+\lambda(1-\gamma)}\hat{w}\left(\tfrac{\bar{\mu}-\gamma\tilde{\mu}}{1-\gamma}\right) \\ &\ge \quad \tfrac{\gamma}{\gamma+\lambda(1-\gamma)}s + \tfrac{\lambda(1-\gamma)}{\gamma+\lambda(1-\gamma)}\min u_S(A \times \Theta) \\ &\to \quad s \text{ as } \lambda \to 0. \end{aligned}$$

Therefore $\hat{w}(\bar{\mu}) \ge s$, as desired. □

Lemma 1. *Given finite A and Θ, the set $\mathcal{U}_R \subseteq \mathbb{R}^{A\times\Theta}$ of R objectives satisfying the PUBR property is open, dense, and of full Lebesgue measure.*

Proof. Toward showing these properties, let us note an algebraic characterization of $\mathcal{U}_R$. Define the finite index set $\mathbb{I} := \{(a, \hat{\Theta}) : \ a \in A,\ \emptyset \ne \hat{\Theta} \subseteq \Theta\}$ and, for each $i = (a, \hat{\Theta}) \in \mathbb{I}$, define

$$\begin{aligned} \varphi_i : \mathbb{R}^{A\times\Theta} \quad &\to \quad \mathbb{R} \\ u_R \quad &\mapsto \quad \max_{\mu\in\Delta\hat{\Theta}} \min_{a'\in A\setminus\{a\}} \int \left[u_R(a,\cdot) - u_R(a',\cdot)\right] \mathrm{d}\mu. \end{aligned}$$

Then, clearly, $\mathcal{U}_R = \{u_R \in \mathbb{R}^{A\times\Theta} : \ \varphi_i(u_R)$ is nonzero for every $i \in \mathbb{I}\} = \bigcap_{i\in\mathbb{I}} \varphi_i^{-1}(\mathbb{R} \setminus \{0\})$.

We can therefore show $\mathcal{U}_R \subseteq \mathbb{R}^{A\times\Theta}$ is open, dense, and of full Lebesgue measure by establishing that $\varphi_i^{-1}(\mathbb{R} \setminus \{0\})$ enjoys these properties for every $i = (a, \hat{\Theta}) \in \mathbb{I}$. First, note it is open because (by Berge's theorem) φ_i is continuous. To show it is of full measure (hence also dense), define $\bar{z} := [\mathbf{1}_{\tilde{a}=a}]_{\tilde{a}\in A,\ \tilde{\theta}\in\Theta} \in \mathbb{R}^{A\times\Theta}$, and observe that $\varphi_i(u_R + \lambda\bar{z}) = \varphi_i(u_R) + \lambda$ for any $u_R \in \mathbb{R}^{A\times\Theta}$ and any $\lambda \in \mathbb{R}$. Now, fixing some $\bar{\theta} \in \Theta$, observe that we can decompose the vector space of all R objectives as the direct sum $\mathbb{R}^{A\times\Theta} = Y \oplus Z$, where

$$Y := \{u_R \in \mathbb{R}^{A\times\Theta} : \; u_R(a, \bar{\theta}) = 0\} \text{ and } Z := \{\lambda \bar{z} : \; \lambda \in \mathbb{R}\}.$$

As $\varphi_i(y + \lambda\bar{z}) = \varphi(y) + \lambda$ for any $y \in Y$ and $\lambda \in \mathbb{R}$, and a singleton is Lebesgue-null in $\mathbb{R} \cong Z$, it follows from the law of iterated expectations that $\varphi_i^{-1}(0)$ is Lebesgue-null as well. The proposition follows. $\square$

Proof of Proposition 4. The result follows directly from Proposition 3 and Lemma 1. $\square$