The epistemic spirit of divinity

Pierpaolo Battigalli,* Emiliano Catonini†

January 22, 2021

Abstract: We study a signaling game where players have common belief in an outcome distribution and the receiver believes that the sender's first-order beliefs are independent of her payoff-type. We capture these epistemic hypotheses through a rationalizability procedure with second-order belief restrictions. Our solution concept is related to, but weaker than Divine Equilibrium (Banks and Sobel, 1987). First, we do not obtain sequential equilibrium, but just Perfect Bayesian Equilibrium with heterogeneous off-path beliefs (Fudenberg and He, 2016). Second, when we rationalize a deviation, we take into account that some types could have preferred another deviation, and we show this is natural and relevant via an economic example.

1 Introduction

We investigate the interaction between sender (*she*) and receiver (*he*) in signaling games under the following hypotheses. There is common belief that sender and receivers choose their actions in a way that is consistent with a given distribution over terminal nodes. When the receiver observes an "off-path" message, he interprets it as follows. First, he maintains that the sender believed in his expected on-path behavior. Second, if possible, he thinks that the deviation is justified as a best response to some belief about his off-path behavior. While doing so, the receiver assumes that this belief has not been influenced by the sender's type. Finally, there is common belief that the receiver rationalizes off-path messages in this way.

These strategic reasoning hypotheses can be justified in many contexts. The existence of an expected outcome distribution can arise from the observation of past play, or from a pre-play non-binding agreement among the players. How the receiver reasons after off-path messages can be interpreted in two ways. The first is that there is an ex-ante stage of the game where the sender forms beliefs about the receiver before learning her type. However, this interpretation cannot apply to an economic context where the sender already knows her type before the game starts. The second way is a kind of principle of insufficient reason for the receiver: If there is no clear direction in which the sender's type would influence her

*Università Bocconi, pierpaolo.battigalli@unibocconi.it

†National Research University Higher School of Economics, emiliano.catonini@gmail.com

beliefs (for instance the sender is exposed to the same information regardless of her type), it seems wise to reason about what induced the sender to deviate independently of her type.

To capture in a transparent way the behavioral implications of our strategic reasoning hypotheses, we employ a variation of Strong-Δ-Rationalizability (Battigalli 2003, Battigalli and Siniscalchi 2003), a notion of rationalizability with belief restrictions for games in extensive form with clear epistemic foundations (Battigalli and Siniscalchi 2007, Battigalli and Prestipino 2013). While the baseline notion of Strong-Δ-Rationalizability only features restrictions to first-order beliefs, we formalize and restrict the receiver's second-order beliefs to capture the effects of our independence hypothesis.

In signaling games, the rationalization of deviations in light of an expected on-path behavior has been captured by the Intuitive Criterion (Cho and Kreps, 1987). On top of this, Divine Equilibrium (Banks and Sobel, 1987) is inspired by the idea that the sender's beliefs are independent of her type. We aim to capture the spirit of Divine Equilibrium starting from primitive assumptions on players' beliefs, and investigate its implications for players' incentives to stick to the expected on-path behavior. We obtain the following: Every outcome distribution of a Divine Equilibrium is compatible with our hypotheses, but not the other way round. There are two reasons for this. The first reason is that, differently from Divine Equilibrium, we do not require the sender to assign positive probability only to reactions of the receiver that best respond to the same belief. This expands both the set of beliefs the sender can have, making it easier to find a belief that keeps all types of the sender on path, and the possible combinations of types that find a deviation profitable under the same belief, expanding the set of possible reactions of the receiver. As a consequence, while Divine Equilibrium refines sequential equilibrium (Kreps and Wilson, 1982), the outcome distributions that are compatible with our hypotheses are induced by a Perfect Bayesian Equilibrium with possibly heterogeneous off-path beliefs (Fudenberg and He, 2016). Requiring the sender to be certain of the receiver's belief after an unexpected message is in line with the spirit of sequential equilibrium However, the hypothesis that the sender has purposedly deviated *despite* her belief in the equilibrium distribution over terminal nodes implies that the sender does not believe in the equilibrium reaction to the deviation. This destroys the motivation for certainty of off-path beliefs. The second reason why we cannot rule out every equilibrium that is not divine is that in Divine Equilibrium beliefs are refined after each off-path message one by one. This induces the receiver to raise the relative probability of type θ with respect to type θ' whenever θ finds that particular deviation profitable for a larger set of responses. But θ may find another deviation even more profitable, and choose it under a belief that induces θ' to stick to the first deviation. As we will show in Section 2 through an economic example, this difference is not just technical, it has relevant and intuitive implications in simple and meaningful games.

Our analysis is closely related to that of Sobel, Stole and Zapater (1990). Sobel et. al. consider the complete-information game where the sender forms beliefs at the ex-ante stage, before observing the chance move that determines her "type". Then, given a sequential equilibrium of the signaling game, they substitute the equilibrium messages with one message m^* that directly gives to the sender the equilibrium expected payoff. Finally, they apply extensive-form rationalizability a la Pearce (1984) (with minor differences) to the modified game. With this, they obtain a “Fixed-Equilibrium Rationalizable Outcome” (FERO) of the original game if m^* survives the elimination procedure. Then, they provide an example of a FERO that is not divine, because the reactions of the receiver to two different deviations can only be justified by the belief, after both deviations, that the same type of the sender would have chosen the other deviation. (Our main example will further highlight that an equilibrium can be fail to be divine even when all deviations can be rationalized under the same theory, instead of under mutually contradictory theories.) So, they refine FERO with the notion of “fixed-equilibrium signal-by-signal rationalizable outcome”, and obtain equivalence with “co-divine equilibria”, where the mixed best responses of the receiver are expanded to mixes of pure best responses. (While doing so, they keep focusing on sequential equilibrium, although, for the reason above, co-divinity would not rule out all non-sequential equilibria.) In the appendix, we consider the complete-information scenario and show that FERO (extended to non-sequential equilibria) is equivalent to Strong-Δ-Rationalizability with the sender's first-order beliefs restricted by the expected outcome distribution and by the independence assumption,[1] which in turn is equivalent to our approach with second-order belief restrictions for the incomplete-information scenario.

Battigalli and Siniscalchi (2003) restrict beliefs on a particular outcome distribution without the independence assumption. With this, they show the equivalence between non-emptiness of Strong-Δ-Rationalizability and passing the Iterated Intuitive Criterion, and that every outcome distribution compatible with strategic reasoning is induced by a self-confirming equilibrium (Fudenberg and Levine, 1993; Battigalli 1987). Cho (1987) strengthens the Intuitive Criterion by requiring that *one* distribution over reactions of the receiver that are compatible with the criterion induces *all* types of the sender to stay on path, coherently with the spirit of Nash equilibrium. We directly obtain Nash equilibrium because of the independence hypothesis: The receiver must be able to believe that all the types of the sender stay on-path for the same first-order belief.

The paper is structured as follows. Section 2 shows similarities and differences between our approach and divinity through an example. Section 3 carries out the formal analysis. Section 4 formalizes the relation between our solution concept and Divine Equilibrium, and the solution of the example of Section 2.

2 Job market example

A potential employee can be of two types, good (θ) and bad (θ'). She can stop studying after graduating from the BSc (m_1), or she can continue to a MSc (m_2) or to a PhD program (m_3). The employer can hire the employee in three different positions, a_1, a_2, a_3, with increasing level of responsibilities. The employer prefers to hire a good employee in

[1] Plus an hypothesis of structural consistency (Kreps and Wilson, 1982) of the *receiver*'s beliefs. In our baseline framework, we will impose an analogous condition as the quickest way to capture the receiver's hypothesis of independence on the sender's first-order beliefs after an unexpected message. In the framework with restrictions to the first-order beliefs of the sender, structural consistency of the receiver's beliefs does not have such clean interpretation. See the appendix for more details.

the position she is best qualified for: a BSc graduate in position a_1, a MSc graduate in position a_2, and a PhD graduate in position a_3. In particular, there is a productivity boost from education to a good employee for positions a_2 and a_3, but overqualification does not carry any additional benefit (i.e., a PhD is equivalent to a MSc for position a_2). There is no productivity boost from education to a bad employee, so the employer always prefers to hire her in position a_1. (Any additional benefit of hiring a good type rather than a bad type independently of education is immaterial for the analysis.) Education has an increasing cost, which is higher for the bad type, but worth paying for both types if and only if the employee will not end up overqualified for the position. The following table summarizes players' payoffs — the first entry in each box is the payoff of the employee.

m_1	a_1	a_2	a_3
θ	$0,3$	$4,2$	$9,0$
θ'	$0,3$	$4,2$	$9,0$

m_2	a_1	a_2	a_3
θ	$-2,3$	$2,5$	$7,3$
θ'	$-3,3$	$1,2$	$6,0$

m_3	a_1	a_2	a_3
θ	$-5,3$	$-1,5$	$4,6$
θ'	$-8,3$	$-4,2$	$1,0$

Suppose that the two types are a priori equally likely. Consider the pooling equilibrium where the two types choose m_1 and the employer chooses a_1 after m_1 and m_2, and a_2 after m_3. This equilibrium is not divine. The reason is the following. Under the belief that m_1 will lead to position a_1, all the beliefs that induce θ' to prefer m_2 to m_1 (also after eliminating the dominated response a_3) also induce θ to strictly prefer m_2 to m_1. Hence, according to divinity, the employer cannot raise the probability of θ' after m_2, and for every belief where θ is not less likely than θ', the optimal response is a_2.

However, the pooling equilibrium is consistent with our epistemic hypotheses. Suppose the employer interprets a deviation to m_2 or m_3 as follows: the employee expects to get the position she is qualified for, that is, position a_i after m_i for each $i = 1, 2, 3$. Given this belief, θ strictly prefers m_3, while θ' is indifferent between a_2 and a_3.[2] Then, whenever the employer believes that θ' picks m_2 with positive probability, he must assign probability 1 to θ' after m_2, and this justifies a_1. Moreover, if the employer believes that θ' breaks her tie at random, she must assign probability 1/3 to θ' after m_3 and this justifies a_2.[3]

Although unneeded in this example, in general the receiver may have to rationalize different unexpected messages with different theories about the sender. We will always allow the receiver to come up with a different theory after each unexpected message. This is in line with the view that the receiver only formulates alternative theories once the initial theory fails to explain the sender's observed behavior. However, it would make sense also to assume that the receiver has an initial array of theories that cover all possible messages, from which he will pick in a lexicographic order to rationalize the observed message. This further restriction would refine the set of equilibria, and the example of Sobel et al. (1990, pag 321) is a case in point: the equilibrium responses to two different deviations can only be justified under two different theories about the sender where the same type chooses the

[2] The tie between the payoffs of θ' after (m_2, a_2) and (m_3, a_3) only serves the purpose of focusing on deterministic beliefs: θ' could be indifferent between m_2 and m_3 anyway under a probabilistic belief, or the employer could give positive probability to two different beliefs of the sender that induce θ' to choose, respectively, m_2 and m_3.

[3] a_2 is also the unique best response if θ' is believed to pick m_3 with higher probability than m_2, or if the receiver assigns higher probability to a belief of the sender that induces θ' to choose m_3 with respect to another belief that induces θ' to choose m_2 (see the previous footnote).

other deviation. On the other hand, the equilibrium above, albeit not divine, would survive this refinement, and so would every divine equilibrium: the equilibrium response to each unexpected message m must be optimal under a theory where the sender only chooses m or the equilibrium messages, therefore any lexicographic ordering of these theories would induce the receiver to use each theory after the corresponding message.

3 Main analysis

We consider the following signaling game. There is a payoff-relevant parameter θ, and it is common knowledge that θ is drawn from a finite set Θ. The sender ($i = 1$) knows the true value of θ (henceforth, the sender's "type"), and chooses a message m from the finite set M (M does not depend on θ). The receiver ($i = 2$), who does not know the true value of θ, observes m and then chooses an action a from the finite set A (A does not depend on m). The payoffs of sender and receiver are given by

$$u_i : \Theta \times M \times A \to \mathbb{R}, \quad i = 1, 2.$$

Suppose there exists a common prior $p \in \Delta(\Theta)$ such that $p(\theta) > 0$ for all $\theta \in \Theta$. For each type θ of the sender, fix a probability measure $\nu^\theta \in \Delta(M)$ over messages. Let M^* be the set of messages m such that $\nu^\theta(m) > 0$ for some $\theta \in \Theta$. For each message $m \in M^*$, fix a probability measure $\nu^m \in \Delta(A)$ over the actions of the receiver that (to avoid uninteresting cases) are optimal once the prior is updated given $(\nu^\theta)_{\theta\in\Theta}$ and m.

We want to capture the behavioral implications of the following epistemic hypotheses on players' strategic reasoning. We assume that players are rational, i.e., they maximize subjective expected utility and update their beliefs via the chain rule of probability. We let $(\nu^\theta)_{\theta\in\Theta}$ and $(\nu^m)_{m\in M^*}$ represent the on-path behavior that sender and receiver initially expect from each other. The initial belief of the receiver is also consistent with the prior over types, and with the hypothesis that the sender's first-order belief is independent of her type. Then, after an off-path message $m \in M \backslash M^*$, the receiver has to revise his belief about the sender. We assume he does so by first formulating a theory about the sender where (i) her first-order belief is independent of her type, (ii) she expects the receiver to play according to $\nu^{m'}$ after each $m' \in M^*$, but (iii) her belief about the receiver's reactions to the unexpected messages induces at least one type to choose m. With this, he can update his prior belief about the sender's type after observing m. Finally, we assume *common strong belief* (Battigalli and Siniscalchi 2002) of the all the previous hypotheses. Strong belief in an event means belief with probability 1 at every information set that is consistent with the event. (Here, for the sender, we will only need to assign a belief at the initial history, therefore belief and strong belief coincide.) Common *correct* strong belief of our basic conditions on rationality and beliefs, where "correct" means that the basic conditions do hold, is constructed recursively as follows.

1. **Sender:** each type of the sender is rational and has belief ν^m about the receiver's action after each $m \in M^*$.

Receiver: the receiver is rational, has an initial belief about the sender that is consistent with the prior p and with each ν^{θ}, and always believes that the sender's first-order belief is independent of her type.

2. **Sender:** each type of the sender satisfies 1.S and believes that the receiver satisfies 1.R.

   Receiver: the receiver satisfies 1 and strongly believes that each type of the sender satisfies 1.S.

3. **Sender:** each type of the sender satisfies 2.S and believes that the receiver satisfies 2.R.

   Receiver: the receiver satisfies 2 and strongly believes that each type of the sender satisfies 2.S.

∞. For every n, each type of the sender satisfies n.S and the receiver satisfies n.R.

Common correct strong belief of the basic conditions may also be impossible: in our case, this happens when the belief that all types choose an on-path message under the same first-order belief is at odds with some step of strategic reasoning.

There is an important qualification to make about the receiver's strong belief in n.S and his hypothesis of independence. Both must apply to the theory about the sender the receiver formulates to rationalize an unexpected message m. So, if m is compatible with n.S, this theory must assign probability 1 to the type-message-belief triples that are compatible with n.S, and at the same time capture the independence hypothesis. The receiver's belief conditional on m will feature correlations between types and their first-order beliefs, as some types find optimal to send m and others don't given the same belief. However, in our formal description of the receiver's beliefs we will introduce the belief conditional on m and not the theory about the sender that rationalizes m, in line with standard practice. This will require to impose the requirements on the theory directly in the rationalizability procedure, as we will articulate later.

Our goal is now to construct a rationalizability procedure that either highlights the impossibility of common correct strong belief of our basic epistemic hypotheses, or captures the behavioral implications thereof. We will construct a version of Strong-Δ-Rationalizability (Battigalli 2003, Battigalli and Siniscalchi 2003) that accommodates not only the restrictions to first-order beliefs given by the prior and by the expected on-path behaviors, but also the restriction to the second-order beliefs of the receiver given by his independence hypothesis. The baseline definition of Strong-Δ-Rationalizability has been given the epistemic justification above by Battigalli and Siniscalchi (2007) for the case of first-order belief restrictions;[4] the extension to second-order belief restrictions is straightforward.

We start by formalizing the belief restrictions. Let

$$\Delta_1 = \left\{ \mu_1 \in \Delta(A^M) \,|\forall m \in M^*, \forall a \in A, \mu_1 \left(\left\{ \sigma_2 \in A^M \,| \sigma_2(m) = a \right\} \right) = \nu^m(a) \right\}$$

[4]Battigalli and Prestipino (2013) provide an alternate epistemic justification where the first-order belief restrictions are transparent, i.e. there is common belief at every node of the game that the restrictions hold.

denote the set of first-order beliefs of the sender that induce belief ν^m after each message $m \in M^*$. Note that, while the sender can have correlated beliefs about the actions of the receiver after different messages, these correlations are immaterial for her choice problem.

Let P_2 denote the set of all finite-support probability measures ν over $\Theta \times M \times \Delta(A^M)$ such that:

B1 for each $\theta \in \Theta$, $\nu(\{\theta\} \times M \times \Delta(A^M)) = p(\theta)$;

B2 for each $(\theta, \mu_1) \in \Theta \times \Delta(A^M)$,

$$\nu(\{\theta\} \times M \times \{\mu_1\}) = \nu(\{\theta\} \times M \times \Delta(A^M)) \cdot \nu(\Theta \times M \times \{\mu_1\}).$$

Thus, P_2 is the set of second-order beliefs of the receiver that conform to the prior and feature no correlation between the sender's type and first-order belief. Let P_2^* the set of all $\nu \in P_2$ such that

B3 for each $(\theta, m) \in \Theta \times M$,

$$\nu\left(\{(\theta, m)\} \times \Delta(A^M)\right) = p(\theta) \cdot \nu^\theta(m).$$

Thus, P_2^* is the subset of P_2 where the receiver believes that every sender's type θ chooses messages according to ν^θ.

Since the sender has finite sets of types and messages, the focus on finite-support probability measures over her first-order beliefs is without loss of generality for the receiver's belief over justifiable messages of the sender.

The receiver has an initial theory about the sender, which he updates or revises after each message m. In particular, the receiver has a *Conditional Probability System* of beliefs over $\Theta \times M \times \Delta(A^M)$: an array $(\mu_2(\cdot|h^0), (\mu_2(\cdot|m))_{m \in M})$ where $\mu_2(\cdot|h^0)$ is the belief at the initial history h^0, and each $\mu_2(\cdot|m)$ is concentrated on $\Theta \times \{m\} \times \Delta(A^M)$ and derived from $\mu_2(\cdot|h^0)$ with the chain rule whenever possible. We require $\mu_2(\cdot|h^0) \in P_2^*$. Then for each $m \in M^*$, $\mu_2(\cdot|m)$ is derived from $\mu_2(\cdot|h^0)$. As anticipated, for each $m \in M \backslash M^*$, we will also require $\mu_2(\cdot|m)$ to be derived from some new theory $\nu \in P_2$. As a stand-alone requirement, this is immaterial for $\mu_2(\cdot|m)$: every belief over $\Theta \times \{m\} \times \Delta(A^M)$ can be derived by conditioning from some $\nu \in P_2$.[5] Yet, requiring $\nu \in P_2$ is not immaterial for the updated belief $\mu_2(\cdot|m)$ when ν is also required to assign probability 1 to a subset of $\Theta \times M \times \Delta(A^M)$ — Example 1 will clarify this point. Since the theory ν has no place in the array of beliefs of a CPS, in order to keep the current simple description of beliefs we will require $\nu \in P_2$ directly in the rationalizability procedure.

Formally, for each $\nu \in P_2$ and $m \in M$ with $\nu(\Theta \times \{m\} \times \Delta(A^M)) > 0$, let $\nu|m$ denote the probability measure derived from ν by conditioning on m, i.e.,

$$\forall\, (\theta, \mu_1) \in \Theta \times \Delta(A^M), \quad (\nu|m)\,(\theta, m, \mu_1) = \frac{\nu(\theta, m, \mu_1)}{\nu(\Theta \times \{m\} \times \Delta(A^M))}.$$

[5] The explicit construction of ν starting from the conditional belief on m is available upon request.

So, we are going to focus on the following set of receiver's CPS's:

$$\Delta_2 = \left\{ (\mu_2(\cdot|h^0), (\mu_2(\cdot|m))_{m\in M}) \,\middle|\, \begin{array}{c} \exists \nu^* \in P_2^*, \mu_2(\cdot|h^0) = \nu^* \\ \forall m \in M^*, \mu_2(\cdot|m) = \nu^*|m \end{array} \right\};$$

Conditional on each $m \in M^*$, every $\nu \in P_2^*$ induces the same belief η^m over the sender's types, the one derived with Bayes rule from the prior and the expected behaviors $(\nu^\theta)_{\theta\in\Theta}$. So we have $\mathrm{marg}_\Theta \mu_2(\cdot|m) = \eta^m$ for each $m \in M^*$.

We now introduce our version of Strong-Δ-Rationalizability, which for future reference we call "Path-rationalizability with second-order independence". For each $(\theta, m, \mu_1) \in \Theta \times M \times \Delta(A^M)$, abusing notation, let $u_1(\theta, m, \mu_1)$ denote the expected payoff of type θ given message m and the probability measure over the receiver's actions after m induced by μ_1. Similarly, for each $(\eta, m, a) \in \Delta(\Theta) \times M \times A$,

$$u_2(\eta, m, a) = \sum_{\theta\in\Theta} \eta(\theta) u_2(\theta, m, a).$$

Definition 1 *Consider the following reduction procedure.*

Step 0 *For each $\theta \in \Theta$, let $\Sigma^0_{1,\theta} = M \times \Delta_1$. Let $\Sigma^0_2 = A^M \times \Delta_2$.*

Step $n > 0$ *For each $\theta \in \Theta$, for each $(m, \mu_1) \in \Sigma^{n-1}_{1,\theta}$, let $(m, \mu_1) \in \Sigma^n_{1,\theta}$ if:*

S1 $\mu_1\left(\mathrm{Proj}_{A^M} \Sigma^{n-1}_2\right) = 1$;

S2 for every $m' \in M$,

$$u_1(\theta, m, \mu_1) \geq u_1(\theta, m', \mu_1).$$

For each $(\sigma_2, \mu_2) \in \Sigma^{n-1}_2$, let $(\sigma_2, \mu_2) \in \Sigma^n_2$ if:

R1 $\mu_2\left(\cup_{\theta\in\Theta}\left(\{\theta\} \times \Sigma^{n-1}_{1,\theta}\right) |h^0\right) = 1$;

R2 For each $m \in M \backslash M^*$, if $(m, \overline{\mu}_1) \in \Sigma^{n-1}_{1,\overline{\theta}}$ for some $(\overline{\theta}, \overline{\mu}_1) \in \Theta \times \Delta_1$, there is $\nu \in P_2$ such that $\nu\left(\cup_{\theta\in\Theta}\left(\{\theta\} \times \Sigma^{n-1}_{1,\theta}\right)\right) = 1$ and $\mu_2(\cdot|m) = \nu|m$;

R3 For every $m \in M$ and $a \in A$,

$$u_2(\mathrm{marg}_\Theta \mu_2(\cdot|m), m, \sigma_2(m)) \geq u_2(\mathrm{marg}_\Theta \mu_2(\cdot|m), m, a).$$

Finally, let $\Sigma^\infty_{1,\theta} = \cap_{n\geq 0} \Sigma^n_1$ for each $\theta \in \Theta$, and let $\Sigma^\infty_2 = \cap_{n\geq 0} \Sigma^n_2$. The elements in Σ^∞_1 and Σ^∞_2 are called path-rationalizable with second-order independence.

Path-rationalizability with second-order independence works as follows. At every step n, by $\mu_1 \in \Delta_1$, the sender believes that receiver will play as expected, i.e., according to ν^m, after each on-path message $m \in M^*$. After every other message $m \in M\backslash M^*$, by S1, she believes that the receiver will play actions that are consistent with his previous step of reasoning. Then, by S2, each type of the sender chooses a message that is optimal given her belief about the receiver. The receiver reasons as follows. At the beginning of the game, by

$\mu_2 \in \Delta_2$, he believes that all types of the sender will choose messages as expected, that is, according to $(\nu^\theta)_{\theta\in\Theta}$ (by B3), under the same belief about the receiver (by B2). However, by R1, he also believes that each type of the sender will choose messages that are consistent with her previous step of reasoning. The two requirements may be mutually inconsistent: there may not be any $\mu_2 \in \Delta_2$ that satisfies R1. This is the only way Path-rationalizability with second-order independence can yield the empty set, highlighting the importance of having an initial belief of the receiver; we will expand on this later. After receiving a message that is consistent with the initial belief, i.e., after $m \in M^*$, the receiver simply updates the initial belief (see the construction of Δ_2); then, by R3, he chooses an optimal action. (Recall that the actions in the support of ν^m are optimal by assumption.) After an unexpected message $m \in M \backslash M^*$, the receiver follows a best rationalization principle. If the message is inconsistent with step $n-1$, he cannot refine his beliefs with respect to the previous step of reasoning — to capture this, it is enough that every $\sigma_2 \in \mathrm{Proj}_{A^M}\Sigma_2^n$ is taken from $\mathrm{Proj}_{A^M}\Sigma_2^{n-1}$. If the message is consistent with step $n-1$ for at least one type of the sender, the receiver revises his belief according to R2. First, he comes up with a new theory about the sender where each type has the prior probability (by B1), all types have the same belief (by B2), the beliefs they may have are consistent with (the previous step of) strategic reasoning, and at least one of these beliefs $\overline{\mu}_1$ induces at least one type $\overline{\theta}$ to play m. Then, the receiver updates this new theory based on the observation of m. In this case, the independency restriction on second-order beliefs cannot induce any inconsistency between Δ_2 and R2 (Δ_2 does not restrict the receiver's off-path beliefs), but it does refine the set of actions that, by R3, the receiver could choose optimally after m. The following example of Path-rationalizability with second-order independence illustrates this fact. The first entry in each box is the payoff of the sender.

Example 1

m_1	a_1	a_2	a_3
θ	$1,1$	$0,0$	$0,0$
θ'	$1,1$	$0,0$	$0,0$

m_2	a_1	a_2	a_3
θ	$3,0$	$0,3$	$0,2$
θ'	$2,3$	$0,0$	$0,2$

Example 1

Let $p(\theta) = p(\theta') = 1/2$, $\nu^\theta(m_1) = \nu^{\theta'}(m_1) = 1$, $\nu^{m_1}(a_1) = 1$. So, Δ_1 is the set of beliefs that give probability 1 to strategies $\sigma \in A^M$ with $\sigma(m_1) = a_1$. For the receiver, every initial belief $\nu \in P_2^*$ is the product of the following marginals (by B2):[6] *the prior (by B1), a Dirac on m_1 (by B3), and a (finite-support) probability measure over $\Delta(A^M)$. So, we have*

$$\Delta_2 = \left\{\mu_2 | \forall \mu_1 \in \Delta(A^M), \mu_2((\theta, m_1, \mu_1)|h^0) = \mu_2((\theta', m_1, \mu_1)|h^0)\right\}.$$

Given the belief in a_1 after m_1, the sender has the incentive to deviate to m_2 only if she assigns sufficiently high probability to a_1 after m_2: at least $1/3$ for type θ and $1/2$ for θ'. So we have

$$\begin{aligned}
\Sigma^1_{1,\theta} &= \{m_1\} \times \{\mu_1 \in \Delta_1 | \mu_1(a_1.a_1) \leq 1/3\} \cup \{m_2\} \times \{\mu_1 \in \Delta_1 | \mu_1(a_1.a_1) \geq 1/3\}, \\
\Sigma^1_{1,\theta'} &= \{m_1\} \times \{\mu_1 \in \Delta_1 | \mu_1(a_1.a_1) \leq 1/2\} \cup \{m_2\} \times \{\mu_1 \in \Delta_1 | \mu_1(a_1.a_1) \geq 1/2\}.
\end{aligned}$$

[6] B2 does not imply a product structure in general, it does here because the marginal over messages is a Dirac.

For the receiver, the first step eliminates the strategies that prescribe the dominated actions a_2 and a_3 after m_1. This has no impact on the sender's beliefs at the second step. For the receiver, the initial beliefs $\nu \in P_2^*$ that satisfy R1 at the second step are the product of the prior, the Dirac on m_1, and a finite-support probability measure over $\{\mu_1 \in \Delta_1 | \mu_1(a_1.a_1) \leq 1/3\}$. The off-path beliefs after m_2 that satisfy R2 do not justify a_1: every $\nu \in P_2$ with $\nu(\{\theta\} \times \Sigma^1_{1,\theta}) = \nu(\{\theta'\} \times \Sigma^1_{1,\theta'}) = 1/2$ must assign non-lower probability to (θ, m_2) than to (θ', m_2), because by B2 types must be associated with the same beliefs, and whenever a belief justifies m_2 for θ', m_2 is the only optimal message for θ. So, calling Δ_2^2 the set of $\mu_2 \in \Delta_2$ that satisfy R1 and R2, we have

$$\begin{aligned} \Sigma_2^2 &= \{a_1.a_2\} \times \left\{\mu_2 \in \Delta_2^2 \,\middle|\, \mu_2(\{\theta\} \times \Sigma^1_{1,\theta} | m_2) \geq 2/3 \right\} \cup \\ &\quad \{a_1.a_3\} \times \left\{\mu_2 \in \Delta_2^2 \,\middle|\, \mu_2(\{\theta\} \times \Sigma^1_{1,\theta} | m_2) \in [1/2, 2/3] \right\}. \end{aligned}$$

At the third step, both types of the sender eliminate m_2, because every belief over Σ_2^2 justifies only m_1. Therefore, we have

$$\Sigma^3_{1,\theta} = \Sigma^3_{1,\theta'} = \{m_1\} \times \left\{\mu_1 \in \Delta_1 | \mu_1\left(\mathrm{Proj}_{A^M} \Sigma_2^{n-1}\right) = 1\right\}.$$

At the fourth step, the receiver refines his initial beliefs by giving probability 1 to the beliefs of the sender compatible with step 3, but cannot refine the beliefs after m_2 because it is incompatible with step 3 for both types of the sender. Therefore, the path-rationalizable strategies of the receiver are $\{a_1.a_2, a_1.a_3\}$, and for each type of the sender the only path-rationalizable message is m_1. Note that not only ν^θ, $\nu^{\theta'}$, and ν^{m_1} are compatible with strategic reasoning, but they also assign probability 1 to the only path-rationalizable behavior of the opponent. This is far from true in general; in the next section we will provide an example where the expected behavior is compatible with strategic reasoning, but also different behaviors are. $\triangle$

Given the assumption that each action in the support of ν^m is optimal under the belief η^m induced by $(\nu^\theta)_{\theta\in\Theta}$, as long as Σ_2^{n-1} is non-empty, it contains strategies that make S1 compatible with Δ_1, thus each $\Sigma^n_{1,\theta}$ is non-empty as well. Instead, we obtain an empty Σ_2^n when $m \notin \mathrm{Proj}_M \Sigma^{n-1}_{1,\theta}$ for some $(\theta, m) \in \Theta \times M^*$ with $\nu^\theta(m) > 0$. In this case, R1 cannot be satisfied by any $\mu_2 \in \Delta_2$, because incompatible with B3. The interpretation is that playing according to ν^θ is not compatible with strategic reasoning for type $\theta \in \Theta$ of the sender. But Σ_2^n can be empty even if $m \in \mathrm{Proj}_M \Sigma^{n-1}_{1,\theta}$ for every $(\theta, m) \in \Theta \times M^*$ with $\nu^\theta(m) > 0$. This happens when different types of the sender play according to ν^θ only for different beliefs. Then, R1 and B3 are not compatible with B2.[7] The interpretation is the following: believing that the sender will play according to $(\nu^\theta)_{\theta\in\Theta}$ even when her beliefs are not influenced by her type is not compatible with strategic reasoning for the receiver. Here is an example of this second kind of inconsistency.

[7] R2 is instead always compatible with B2, because $\mathrm{Proj}_{\Delta(A^M)} \Sigma^{n-1}_{1,\theta}$ is the same for all $\theta \in \Theta$, thus every type can be associated with $\overline{\mu}_1$.

Example 2

m_1	a_1	a_2
θ	$1,0$	$0,0$
θ'	$1,0$	$0,0$

m_2	a_1	a_2
θ	$3,0$	$0,0$
θ'	$0,0$	$3,0$

Example 2

Let $p(\theta) = p(\theta') = 1/2$*,* $\nu^{\theta}(m_1) = \nu^{\theta'}(m_1) = 1$*,* $\nu^{m_1}(a_1) = 1$*. So,* Δ_1 *is the set of beliefs that give probability* 1 *to strategies* $\sigma \in A^M$ *with* $\sigma(m_1) = a_1$*, and*

$$\Delta_2 = \left\{\mu_2 | \forall \mu_1 \in \Delta(A^M), \mu_2((\theta, m_1, \mu_1)|h^0) = \mu_2((\theta', m_1, \mu_1)|h^0)\right\}.$$

Given the belief in a_1 *after* m_1*, types* θ *and* θ' *have the incentive to deviate to* m_2 *if they assign at least probability* $1/3$ *to, respectively,* a_1 *and* a_2*. So we have*

$$\begin{aligned}
\Sigma^1_{1,\theta} &= \{m_1\} \times \{\mu_1 \in \Delta_1 | \mu_1(a_1.a_1) \le 1/3\} \cup \{m_2\} \times \{\mu_1 \in \Delta_1 | \mu_1(a_1.a_1) \ge 1/3\}, \\
\Sigma^1_{1,\theta'} &= \{m_1\} \times \{\mu_1 \in \Delta_1 | \mu_1(a_1.a_2) \le 1/3\} \cup \{m_2\} \times \{\mu_1 \in \Delta_1 | \mu_1(a_1.a_2) \ge 1/3\}.
\end{aligned}$$

Therefore, there is no $\mu_2 \in \Delta_2$ *that satisfies R1 at step* 2*.* $\triangle$

Observe now the following. For every $\mu_2 \in \Delta_2$, B2 requires all types of the sender to be associated with the same belief in the support of $\mu_2\left(\cdot|h^0\right)$; at the same time, B3 requires the types to be associated with their on-path messages; therefore, if μ_2 believes in the rationality of the sender, each of those beliefs must justify an on-path message for each type $\theta \in \Theta$. Note that, if every message m with $\nu^{\theta}(m) > 0$ is justified by some $\mu_1 \in \Delta_1$, μ_1 justifies them all: every $\mu_1 \in \Delta_1$ induces the same beliefs $(\nu^m)_{m \in M^*}$ after the messages in M^*, so if θ had a strict ranking of the on-path messages under some μ_1, this ranking would be the same under all $\mu_1 \in \Delta_1$. Therefore, compatibility of R1 with B2 and B3 at any step $n > 1$ guarantees that some belief of the sender justifies all on-path messages for all types.

Lemma 1 *Fix* $n > 1$*. We have* $\Sigma^n_2 \neq \emptyset$ *if and only if there exists* $\overline{\mu}_1 \in \Delta_1$ *with* $\overline{\mu}_1\left(\text{Proj}_{A^M}\Sigma^{n-2}_2\right) = 1$ *such that, for each* $(\theta, m) \in \Theta \times M^*$ *with* $\nu^{\theta}(m) > 0$*,*

$$m \in \arg\max_{m' \in M} u_1(\theta, m', \overline{\mu}_1). \tag{1}$$

Proof. Sufficiency. Construct $\nu^* \in P^*_2$ as $\nu^*\left(\theta, m, \overline{\mu}_1\right) = p(\theta) \cdot \nu^{\theta}(m)$ for each $(\theta, m) \in \Theta \times M$, thus $\nu^*\left(\Theta \times M \times \left(\Delta(A^M) \backslash \{\overline{\mu}_1\}\right)\right) = 0$. it is easy to check that B1, B2 and B3 are satisfied. Fix $m \in \text{Proj}_M\left(\cup_{\theta \in \Theta}\Sigma^{n-1}_{1,\theta}\right)$. Fix $\left(\overline{\theta}, \mu_1\right) \in \Theta \times \Delta_1$ such that $(m, \mu_1) \in \Sigma^{n-1}_{1,\overline{\theta}}$. Let $m(\overline{\theta}) = m$. For each $\theta \neq \overline{\theta}$, fix $m(\theta) \in M$ such that $(m(\theta), \mu_1) \in \Sigma^{n-1}_{1,\theta}$. (Note that $\text{Proj}_{\Delta(A^M)}\Sigma^{n-1}_{1,\theta}$ is independent of θ.) Construct $\nu(m) \in P_2$ as $\nu(m)\left(\theta, m(\theta), \mu_1\right) = p(\theta)$ for each $\theta \in \Theta$;[8] it is easy to check that B1 and B2 are satisfied. Construct $\mu_2 \in \Delta_2$ as $\mu_2(\cdot|h^0) = \nu^*$, $\mu_2(\cdot|m) = \nu^*|m$ for each $m \in M^*$, and $\mu_2(\cdot|m) = \nu(m)|m$ for each $m \in M \backslash M^*$. For each $(\theta, m, \mu_1) \in \Theta \times M \times \Delta(A^M)$, if $\nu^*\left(\theta, m, \mu_1\right) > 0$, then $\nu^{\theta}(m) > 0$

[8] To complete the definition of $\nu(m)$, $\nu(m)\left(\{\theta\} \times (M \backslash \{m(\theta)\}) \times \left(\Delta(A^M) \backslash \{\mu_1\}\right)\right) = 0$ for each $\theta \in \Theta$.

and $\mu_1 = \overline{\mu}_1$. Then, by (1), $(m, \mu_1) \in \Sigma_{1,\theta}^{n-1}$. Hence, μ_2 satisfies R1. Moreover, μ_2 satisfies R2 by construction of each $\nu(m)$. Hence $\Sigma_2^n \neq \emptyset$.

Necessity. By $\Sigma_2^n \neq \emptyset$, there exists $\mu_2 \in \Delta_2$ that satisfies R1. Fix $(\overline{\theta}, \overline{m}, \overline{\mu}_1) \in \Theta \times M \times \Delta(A^M)$ such that $\mu_2\left((\overline{\theta}, \overline{m}, \overline{\mu}_1)|h^0\right) > 0$. Thus, $(\overline{m}, \overline{\mu}_1) \in \Sigma_{1,\overline{\theta}}^{n-1}$. Hence, by $n > 1$ and S2,

$$\overline{m} \in \arg\max_{m' \in M} u_1(\overline{\theta}, m', \overline{\mu}_1).$$

Fix $m \neq \overline{m}$ with $\nu^{\overline{\theta}}(m) > 0$. Recall that $\mu_2\left(\cdot|h^0\right) \in P_2^*$. Thus, by B3,

$$\mu_2\left(\left\{(\overline{\theta}, m)\right\} \times \Delta(A^M)|h^0\right) > 0.$$

So, by R1, there is $\mu_1 \in \Delta_1$ such that $(m, \mu_1) \in \Sigma_{1,\overline{\theta}}^{n-1}$. Hence, by $n > 1$ and S2,

$$m \in \arg\max_{m' \in M} u_1(\overline{\theta}, m', \mu_1).$$

Since μ_1 and $\overline{\mu}_1$ induce the same belief ν^m after each message in M^*, we have

$$\begin{aligned} u_1(\overline{\theta}, \overline{m}, \overline{\mu}_1) &= u_1(\overline{\theta}, \overline{m}, \mu_1), \\ u_1(\overline{\theta}, m, \overline{\mu}_1) &= u_1(\overline{\theta}, m, \mu_1). \end{aligned}$$

Therefore, optimality of $\overline{m}$ under $\overline{\mu}_1$ and of m under μ_1 imply optimality of both under both beliefs. We conclude that, for each $m \in M^*$ with $\nu^{\overline{\theta}}(m) > 0$,

$$m \in \arg\max_{m' \in M} u_1(\overline{\theta}, m', \overline{\mu}_1).$$

Fix $\theta \neq \overline{\theta}$. Recall that $\mu_2\left(\cdot|h^0\right) \in P_2^*$. Thus, by B2,

$$\mu_2(\{\theta\} \times M \times \{\overline{\mu}_1\}|h^0) = \mu_2(\{\theta\} \times M \times \Delta(A^M)|h^0) \cdot \mu_2(\Theta \times M \times \{\overline{\mu}_1\}|h^0) > 0.$$

By B3, for each m with $\nu^\theta(m) = 0$, $\mu_2((\theta, m, \overline{\mu}_1)|h^0) = 0$. So, there is $\overline{m}$ with $\nu^{\overline{\theta}}(\overline{m}) > 0$ such that $\mu_2((\theta, \overline{m}, \overline{\mu}_1)|h^0) > 0$. Hence, repeating the argument used for $\overline{\theta}$, we conclude that for each m with $\nu^\theta(m) > 0$,

$$m \in \arg\max_{m' \in M} u_1(\theta, m', \overline{\mu}_1).$$

Having shown this for each $\theta \in \Theta$, (1) obtains. ■

Lemma 1 guarantees that, if $(\nu^\theta)_{\theta\in\Theta}$ and $(\nu^m)_{m\in M^*}$ are compatible with strategic reasoning, there is a belief $\overline{\mu}_1 \in \Delta_1$ over the strategically sophisticated strategies of the receiver so that every type of the sender has the incentive to stay on path. Such strategies of the receiver need not be optimal under the same off-path beliefs. Then, $(\nu^\theta)_{\theta\in\Theta}$ and $\overline{\mu}_1$, which induces ν^m after each $m \in M^*$, define the behavioral strategies of a Perfect Bayesian Equilibrium with heterogenous off-path beliefs (PBH; Fudenberg and He, 2016).

Proposition 1 *Suppose that $\Sigma_2^3 \neq \emptyset$. Then, there exists a PBH $(\beta_1, \beta_2) \in (\Delta(M))^\Theta \times (\Delta(A))^M$ such that $\beta_1(\theta) = \nu^\theta$ for each $\theta \in \Theta$ and $\beta_2(m) = \nu^m$ for each $m \in M^*$.*

Proof. By Lemma 1, there exists $\overline{\mu}_1 \in \Delta_1$ with $\overline{\mu}_1 \left(\mathrm{Proj}_{A^M}\Sigma_2^1\right) = 1$ such that, for each $(\theta, m) \in \Theta \times M^*$ with $\nu^\theta(m) > 0$,

$$m \in \arg\max_{m' \in M} u_1(\theta, m', \overline{\mu}_1).$$

Consider the profile of behavioral strategies $(\beta_1, \beta_2) \in (\Delta(M))^\Theta \times (\Delta(A))^M$ defined by:

1. $\beta_1(\theta) = \nu^\theta$ for each $\theta \in \Theta$;
2. $\beta_2(m) = \nu^m$ for each $m \in M^*$;
3. for each $m \in M \backslash M^*$ and $a \in A$,
$$\beta_2(m)[a] = \overline{\mu}_1 \left(\left\{\sigma_2 \in A^M | \sigma_2(m) = a\right\}\right).$$

For each $(\theta, m) \in \Theta \times M^*$ with $\beta_1(\theta)[m] > 0$, i.e., $\nu^\theta(m) > 0$, m is optimal against β_2 because β_2 and $\overline{\mu}_1$ induce the same belief after each $m' \in M$. Moreover, β_1 induces belief η^m after every $m \in M^*$, and by assumption the actions in the support of $\beta_2(m) = \nu^m$ are optimal given η^m. Finally, for each $m \in M \backslash M^*$ and $a \in A$ with $\beta_2(m)[a] > 0$, we have $\overline{\mu}_1 \left(\left\{\sigma_2 \in A^M | \sigma_2(m) = a\right\}\right) > 0$, thus there is $(\sigma_2, \mu_2) \in \Sigma_2^1$ with $\sigma_2(m) = a$ such that, by R3,

$$\forall a' \in A, \quad u_2\left(\mathrm{marg}_\Theta \mu_2(\cdot|m), m, \sigma_2(m)\right) \geq u_2\left(\mathrm{marg}_\Theta \mu_2(\cdot|m), m, a'\right).$$

Hence, a is optimal given $\mathrm{marg}_\Theta \mu_2(\cdot|m)$. ■

We obtain a PBH and not just a self-confirming equilibrium (Fudenberg and Levine 1993, Battigalli 1987) because of our restriction on the receiver's second-order beliefs.[9] Battigalli and Siniscalchi (2003) have shown that when the first-order beliefs are restricted by a given outcome distribution (as the prior p, $(\nu^\theta)_{\theta \in \Theta}$, and $(\nu^m)_{m \in M^*}$ do), non-emptiness of Strong-Δ-Rationalizability guarantees that the distribution is induced by a self-confirming equilibrium. Moreover, they show that in a signaling game non-emptiness of Strong-Δ-Rationalizability is equivalent to the Iterated Intuitive Criterion (Cho and Kreps 1987). Since our restrictions to first-order beliefs are of the same kind, also Path-rationalizability with second-order independence, when non-empty, guarantees that the corresponding PBH satisfies the Iterated Intuitive Criterion. The examples above illustrate how the independence restriction on the receiver's second-order beliefs further refines his first-order beliefs: in Example 1, at step 2, the intuitive criterion would allow the receiver to assign any probability to θ' after m_2 and thus to play a_1, because in his mind θ and θ' could have different beliefs where θ' has the incentive to play m_2 and θ has the incentive to play m_1; in Example 2, again, ν^θ and $\nu^{\theta'}$ would be justified by different beliefs about the receiver's off-path behavior,[10] so we would not get the empty set.

[9] On top of this, we obtain a PBH and not just a Bayes-Nash equilibrium because the first step of reasoning guarantees that the receiver best replies to some off-path beliefs.

[10] As observed, for a given type $\theta \in \Theta$ all the actions in the support of ν^θ would be anyway justified by the same belief, because of indifference among all of them given $(\nu^m)_{m \in M^*}$. Fudenberg and Kamada (2015) call this property "unitary beliefs".

4 Comparison with divinity

Banks and Sobel (1987) call a sequential equilibrium $(\beta_1^*, \beta_2^*) \in (\Delta(M))^\Theta \times (\Delta(A))^M$ "divine" when β_2^* survives an iterated procedure of refinement of off-path beliefs inspired by the idea that the beliefs of the sender do not differ across types. However, there are some relevant differences with our analysis. To appreciate similarities and differences between our analysis and Divine Equilibrium, it is sufficient to focus on the first two steps of reasoning. To facilitate this comparison, we report here the first three steps of divine equilibrium and we call them "divinity criterion". Banks and Sobel focus directly on sequential equilibrium, which in signaling games coincides with Perfect Bayesian Equilibrium (with common off-path beliefs). However, we start from a Bayes-Nash equilibrium (β_1^*, β_2^*) to show that the restrictions to off-path beliefs of PBE emerge endogenously from their conditions. Let $M^* := \cup_{\theta\in\Theta}\mathrm{Supp}\beta_1^*(\cdot|\theta)$, and for any $m \in M\backslash M^*$ and any map $\sigma \in [0,1]^\Theta$ that assigns to each $\theta \in \Theta$ a probability of playing m, let η_σ^m denote the probability distribution over Θ derived from the prior and σ by Bayes rule. Let $\mathrm{Conv}(Y)$ denote the convex hull of a set Y.

Definition 2 *Fix a Bayes-Nash equilibrium (β_1^*, β_2^*). For each $m \in M\backslash M^*$, let*

$$\begin{aligned}
\Sigma_1^d(m) &:= \left\{\sigma \in [0,1]^\Theta \,\middle|\, \exists \alpha \in \Delta(A), \forall \theta \in \Theta, \sigma(\theta) \in \arg\max_{\pi\in[0,1]} \pi u_1(\theta, m, \alpha) + (1-\pi) u_1(\theta, \beta_1^*, \beta_2^*)\right\};\\
\Gamma(m) &:= \left\{p' \in \Delta(\Theta) \,\middle|\, \exists \sigma \in \Sigma_1^d(m)\backslash\left\{\overrightarrow{0}\right\}, p' = \eta_\sigma^m\right\};\\
\Sigma_2^d(m) &:= \left\{\alpha \in \Delta(A) \,\middle|\, \exists p' \in \mathrm{Conv}(\Gamma(m)), \mathrm{Supp}\alpha \subseteq \arg\max_{a'\in A} u_2(p', m, a')\right\}
\end{aligned}$$

We say that (β_1^*, β_2^*) satisfies the ***divinity criterion*** *if for each $m \in M\backslash M^*$*

$$\begin{aligned}
\Gamma(m) &\neq \emptyset \Rightarrow \beta_2^*(\cdot|m) \in \Sigma_2^d(m),\\
\Gamma(m) &= \emptyset \Rightarrow \exists p' \in \Delta(\Theta), \mathrm{Supp}\beta_2^*(\cdot|m) \subseteq \arg\max_{a'\in A} u_2(p', m, a').
\end{aligned}$$

Note that all the actions in the support of each $\beta_2^*(\cdot|m)$ must best reply to the same belief over the sender's types. Then, we have the following.

Remark 1 *If (β_1^*, β_2^*) satisfies the divinity criterion, then it is a Perfect Bayesian Equilibrium.*

The divinity criterion guarantees that the equilibrium distributions over messages and actions are compatible with the first two steps of reasoning under our hypotheses.

Theorem 1 *Fix a Bayes-Nash Equilibrium (β_1^*, β_2^*) that satisfies the divinity criterion and let $\nu^\theta := \beta_1^*(\cdot|\theta)$ for each $\theta \in \Theta$, $\nu^m := \beta_2^*(\cdot|m)$ for each $m \in M^*$. Then*

$$\times_{m\in M}\mathrm{Supp}\beta_2^*(\cdot|m) \subseteq \mathrm{Proj}_{A^M}\Sigma_2^2.$$

Proof. Fix $m \in M \backslash M^*$. If $\Gamma(m) = \emptyset$, there is $p' \in \Delta(\Theta)$ such that

$$\text{Supp}\beta_2^*(\cdot|m) \subseteq \arg\max_{a' \in A} u_2(p', m, a'),$$

so fix $\nu_{2,m} \in P_2$ such that $\nu_{2,m}|m = p'$ (it exists because the prior has full support).

Now suppose that $\Gamma(m) \neq \emptyset$. Then, there is $p' \in \text{Conv}(\Gamma(m))$ such that $\text{Supp}\beta_2^*(\cdot|m) \subseteq \arg\max_{a' \in A} u_2(p', m, a')$. Thus, $p' = \gamma^1 p^1 + ... + \gamma^n p^n$ where $p^j \in \Gamma(m)$ for each $j = 1, ..., n$ and $(\gamma^j)_{j=1,...,n}$ is a convex combination. For each $j = 1, ..., n$, there is $\sigma^j \in \Sigma_1^d(m)$ such that $p^j = \eta_{\sigma^j}^m$. Hence, there is $\alpha \in \Delta(A)$ such that

$$\sigma^j(\theta) \in \arg\max_{\pi \in [0,1]} \pi u_1(\theta, m, \alpha) + (1-\pi) u_1(\theta, \beta_1^*, \beta_2^*) \tag{2}$$

for each $\theta \in \Theta$. Construct $\mu_1^j \in \Delta(A^M)$ that induces α after m and $\beta_2^*(\cdot|m')$ after each $m' \neq m$. Since (β_1^*, β_2^*) is an equilibrium, for each $\theta \in \Theta$ we have

$$u_1(\theta, m', \mu_1^j) \geq u_1(\theta, m'', \mu_1^j)$$

for each $m' \in M^*$ and $m'' \in M \backslash M^*$ with $m'' \neq m$. But then, for each $\theta \in \Theta$, by (2) we get

$$\begin{aligned} m &\in \arg\max_{m'} u_1(\theta, m', \mu_1^j) \quad \text{if } \sigma^j(\theta) > 0, && (3) \\ \text{Supp}\nu^\theta &\subseteq \arg\max_{m'} u_1(\theta, m', \mu_1^j) \quad \text{if } \sigma^j(\theta) < 1. && (4) \end{aligned}$$

Construct $\nu_{2,m}^j \in P_2$ as, for each $\theta \in \Theta$,

$$\begin{aligned} \nu_{2,m}^j(\{\theta\} \times \text{Supp}\nu^\theta \times \{\mu_1^j\}) &= p(\theta) \cdot (1 - \sigma^j(\theta)), \\ \nu_{2,m}^j((\theta, m, \mu_1^j)) &= p(\theta) \cdot \sigma^j(\theta). \end{aligned}$$

For each $\theta \in \Theta$ with $\nu_{2,m}^j((\theta, m, \mu_1^j)) > 0$, we have $\sigma^j(\theta) > 0$, and thus by (3) $(m, \mu_1^j) \in \Sigma_{1,\theta}^1$. For each $\theta \in \Theta$ with $\nu_{2,m}^j(\{\theta\} \times \text{Supp}\nu^\theta \times \{\mu_1^j\}) > 0$, we have $\sigma^j(\theta) < 1$, and thus by (4) $(m', \mu_1^j) \in \Sigma_{1,\theta}^1$ for each $m' \in \text{Supp}\nu^\theta$. So, $\nu_{2,m}^j(\cup_{\theta \in \Theta} \left(\{\theta\} \times \Sigma_{1,\theta}^{n-1}\right)) = 1$. Now, for each $j = 1, ..., n$ let

$$\begin{aligned} \widetilde{\delta}^j &= \frac{\gamma^j}{\sum\limits_{\theta \in \Theta} p(\theta) \sigma^j(\theta)}, \\ \delta^j &= \frac{\widetilde{\delta}^j}{\sum\limits_{k=1,...n} \widetilde{\delta}^k}. \end{aligned}$$

For future reference, observe that

$$\frac{\delta^j}{\sum\limits_{k=1,...n} \delta^k \sum\limits_{\theta \in \Theta} p(\theta) \sigma^k(\theta)} = \frac{\gamma^j}{\sum\limits_{\theta \in \Theta} p(\theta) \sigma^j(\theta) \cdot \sum\limits_{k=1,...n} \widetilde{\delta}^k} \cdot \frac{\sum\limits_{k=1,...n} \widetilde{\delta}^k}{\sum\limits_{k=1,...n} \gamma^k} = \frac{\gamma^j}{\sum\limits_{\theta \in \Theta} p(\theta) \sigma^j(\theta)}. \tag{5}$$

Let $\nu_{2,m} = \delta^1 \nu^1_{2,m} + ... + \delta^n \nu^n_{2,m}$. Clearly, $\nu_{2,m} \in P_2$ and $\nu_{2,m}(\cup_{\theta \in \Theta} \left(\{\theta\} \times \Sigma^{n-1}_{1,\theta} \right)) = 1$. Moreover, for each $\overline{\theta} \in \Theta$, we have

$$
\begin{aligned}
(\nu_{2,m}|m) \left(\left\{ \left(\overline{\theta}, m \right) \right\} \times \Delta(A^M) \right) &= \frac{\nu_{2,m}(\left\{ \left(\overline{\theta}, m \right) \right\} \times \Delta(A^M))}{\nu_{2,m}(\Theta \times \{m\} \times \Delta(A^M))} = \\
&= \frac{\delta^1 \nu^1_{2,m}(\left\{ \left(\overline{\theta}, m \right) \right\} \times \Delta(A^M)) + ... + \delta^n \nu^n_{2,m}(\left\{ \left(\overline{\theta}, m \right) \right\} \times \Delta(A^M))}{\delta^1 \nu^1_{2,m}(\Theta \times \{m\} \times \Delta(A^M)) + ... + \delta^n \nu^n_{2,m}(\Theta \times \{m\} \times \Delta(A^M))} = \\
&= \frac{\delta^1 p(\overline{\theta}) \sigma^1(\overline{\theta})}{\sum\limits_{k=1,...n} \delta^k \sum\limits_{\theta \in \Theta} p(\theta) \sigma^k(\theta)} + ... + \frac{\delta^n p(\overline{\theta}) \sigma^n(\overline{\theta})}{\sum\limits_{k=1,...n} \delta^k \sum\limits_{\theta \in \Theta} p(\theta) \sigma^k(\theta)} = \\
&= \gamma^1 \frac{p(\overline{\theta}) \sigma^1(\overline{\theta})}{\sum\limits_{\theta \in \Theta} p(\theta) \sigma^1(\theta)} + ... + \gamma^n \frac{p(\overline{\theta}) \sigma^n(\overline{\theta})}{\sum\limits_{\theta \in \Theta} p(\theta) \sigma^n(\theta)} = \\
&= \gamma^1 p^1(\overline{\theta}) + ... + \gamma^n p^n(\overline{\theta}) = p'(\overline{\theta}),
\end{aligned}
$$

where the fourth equality follows from (5). Construct $\mu_2 \in \Delta_2$ such that $\mu_2(\cdot|h^0) = \nu^*$ for some $\nu^* \in P^*_2$, and $\mu_2(\cdot|m) = \nu_{2,m}|m$ for each $m \in M \backslash M^*$. Fix $\sigma_2 \in A^M$ such that $\sigma_2(m) \in \mathrm{Supp}\beta^*_2(\cdot|m)$ for each $m \in M$. By construction, (σ_2, μ_2) satisfies R1, R2, and R3 at step $n = 2$, so $(\sigma_2, \mu_2) \in \Sigma^2_2$. ■

The other way round is not true: even if $\times_{m \in M} \mathrm{Supp}\beta^*_2(\cdot|m) \subseteq \mathrm{Proj}_{A^M} \Sigma^\infty_2$, (β^*_1, β^*_2) might not satisfy the divinity criterion. To see this, we now formalize the solution to the example of Section 2.

Example 3

m_1	a_1	a_2	a_3
θ	$0,3$	$4,2$	$9,0$
θ'	$0,3$	$4,2$	$9,0$

m_2	a_1	a_2	a_3
θ	$-2,3$	$2,5$	$7,3$
θ'	$-3,3$	$1,2$	$6,0$

m_3	a_1	a_2	a_3
θ	$-5,3$	$-1,5$	$4,6$
θ'	$-8,3$	$-4,2$	$1,0$

The prior is $p(\theta) = p(\theta') = 1/2$. *Consider the equilibrium* (β^*_1, β^*_2) *with* $\beta^*_1(m_1|\theta) = \beta^*_1(m_1|\theta') = 1$, $\beta^*_2(a_1|m_1) = \beta^*_2(a_1|m_2) = \beta^*_2(a_2|m_3) = 1$. *We have*

$$\Sigma^d_1(m_2) = \Sigma^d_1(m_3) = \{[0,1] \times \{0\}\} \cup \{\{1\} \times [0,1]\},$$

where the two components are justified by beliefs of the sender that make, respectively, θ *indifferent between* m_1 *and* m_i*, thus* θ' *strictly prefer* m_1*, and* θ' *indifferent between* m_1 *and* m_i*, thus* θ *strictly prefer* m_i. *With this, we get*

$$\Gamma(m_2) = \Gamma(m_3) = \left\{ p' \in \Delta(\Theta) \left| p'(\theta) \geq 1/2 \right. \right\}.$$

But then (abusing notation),

$$\Sigma^d_2(m_2) = \{a_2\}, \quad \Sigma^d_2(m_3) = \Delta(a_2, a_3).$$

so $\beta_2^*(\cdot|m_2) \notin \Sigma_2^d(m_2)$*:* (β_1^*, β_2^*) *does not satisfy the divinity criterion (and no equilibrium with* $\beta_1^*(m_1|\theta) = \beta_1^*(m_1|\theta') = 1$ *would).*

Now we turn to Path-rationalizability with second-order independence. We have

$$\begin{aligned} \{(m_1, \delta_{a_1.a_1.a_2}), (m_3, \delta_{a_1.a_2.a_3})\} &\subset \Sigma^1_{1,\theta}, \\ \{(m_1, \delta_{a_1.a_1.a_2})\} \cup (\{m_2, m_3\} \times \{\delta_{a_1.a_2.a_3}\}) &\subset \Sigma^1_{1,\theta'}. \end{aligned}$$

For the receiver, note that $\times_{m\in M}\mathrm{Supp}\beta_2^*(\cdot|m) = a_1.a_1.a_2$*. We check whether* $a_1.a_1.a_2 \in \mathrm{Proj}_{A^M}\Sigma_2^2$*. Consider the following belief* $\mu_2 \in \Delta_2$*. Let* $\mu_2(\cdot|h^0)$ *assign probability* $1/2$ *to each element of* $\{\theta, \theta'\} \times \{m_1\} \times \{\delta_{a_1.a_1.a_2}\}$*; thus,* μ_2 *satisfies B1, B2, B3, and R1. Let* $\mu_2(\cdot|m_2)$ *and* $\mu_2(\cdot|m_3)$ *be both derived by updating from* ν *with*

$$\begin{aligned} \nu((\theta, m^3, \delta_{a_1.a_2.a_3})) &= 1/2, \\ \nu((\theta', m^2, \delta_{a_1.a_2.a_3})) &= \nu((\theta', m^3, \delta_{a_1.a_2.a_3})) = 1/4; \end{aligned}$$

thus, μ_2 *satisfies R2. We get*

$$\begin{aligned} \mathrm{marg}_\Theta(\nu|m_2)(\theta') &= 1, \\ \mathrm{marg}_\Theta(\nu|m_3)(\theta') &= 1/3; \end{aligned}$$

thus $(a_1.a_1.a_2, \mu_2)$ *satisfies R3. Hence,* $(a_1.a_1.a_2, \mu_2) \in \Sigma_2^2$*.*

Since also $a_1.a_2.a_3$ *survives the second step, an easy inductive argument shows that the equilibrium survives all steps of path rationalizability.* $\triangle$

To conclude, note that, after m_3, a_1 and a_3 best reply only to disjoint sets of beliefs about the sender's type. This is immaterial for Path-rationalizability with second-order independence: the sender can still assign positive probability both to a_1 and a_3 at all steps. For Divine Equilibrium, this is not allowed, because a mix of a_1 and a_3 is not a best response to any belief. Sobel et al. (1990, pag 316) provide an example where this fact matters.

5 Appendix

In this appendix, we analyze our signaling game as a complete information game with asymmetric observation of an initial chance move. Then, we show the equivalence of the analysis with the analysis of the incomplete information game of the main body, and with the notion of fixed-equilibrium rationalizable outcome (henceforth, FERO) introduced by Sobel et al. (1990).

The timing of the game is as follows.

1. The pseudo-player chance chooses the value of θ from Θ.
2. The sender ($i = 1$) observes θ and chooses a message m from M.
3. The receiver ($i = 2$) observes m but not θ and chooses an action a from A.

4. The payoffs of the sender and the receiver are realized as a function

$$u_i : \Theta \times M \times A \to \mathbb{R}, \quad i = 1, 2,$$

of the terminal history (θ, m, a).

Chance chooses θ according to the commonly known probability measure p. Let M^Θ and A^M denote the sets of strategies of the sender and of the receiver, respectively. It will be useful to introduce the following notation:

$$\begin{aligned} M^\Theta(\theta, m) &: = \left\{ \sigma_1 \in M^\Theta \left| \sigma_1(\theta) = m \right. \right\}, \\ M^\Theta(m) &: = \cup_{\theta \in \Theta} M^\Theta(\theta, m), \\ S(\theta, m) &: = \{\theta\} \times M^\Theta(\theta, m), \\ S(m) &: = \cup_{\theta \in \Theta} S(\theta, m); \end{aligned}$$

that is, $M^\Theta(\theta, m)$ is the set of strategies that assign m to θ, $M^\Theta(m)$ is the set of strategies that assign m to some θ, $S(\theta, m)$ is the set of type-strategy pairs where the type is θ and the strategy assigns m to θ, and $S(m)$ is the set of all type-strategy pairs where the strategy assigns m to the type.

At the beginning of the game, the sender and the receiver have a belief over the strategies of the other player. The sender believes there is no correlation between the move of chance and the strategy of the receiver, and this fact is commonly believed. Therefore, we can simply model the belief of the sender as a probability measure $\mu_1 \in \Delta(A^M)$. On top of this, at the beginning of the game, the sender and the receiver commonly believe that the opponent will play according to the probability measures $(\nu^\theta)_{\theta \in \Theta}$ and $(\nu^m)_{m \in M}$ that we introduced in the main body.

Altogether, the sender and the receiver will have beliefs in this restricted set:

$$\Delta_1 = \left\{ \mu_1 \in \Delta(A^M) \left| \forall (m, a) \in M \times A, \mu_1 \left(\left\{ \sigma_2 \in A^M \left| \sigma_2(m) = a \right. \right\} \right) = \nu^m(a) \right. \right\},$$

This set coincides with the one defined in the main body; the additional hypothesis that the sender's belief is independent of the observed chance move will simply be captured by the fact that she will form one belief $\mu_1 \in \Delta_1$ before observing the chance move, without updating it afterwards.

The receiver will have a CPS over $\Theta \times M^\Theta$, that is, an array of beliefs $\mu_2 = (\mu_2(\cdot|h^0), (\mu_2(\cdot|m))_{m \in M})$ such that, for each $m \in M$, $\mu_2(\cdot|m)$ assigns probability 1 to $S(m)$, and is derived from $\mu_2(\cdot|h^0)$ by conditioning whenever possible. We will focus on the set Δ_2 of CPS's such that,[11] for each $(\theta, m) \in \Theta \times M$,

$$\mu_2(\{\theta\} \times \left\{ \sigma_1 \in M^\Theta | \sigma_1(\theta) = m \right\} | h^0) = p(\theta) \cdot \nu^\theta(m). \tag{6}$$

[11] This definition of Δ_2 differs from the one defined in the main body in that we do not need to restrict the receiver's second-order beliefs and thus we do not model them. Note that any correlation in $\mu_2(\cdot|h^0)$ between the chance move and the strategy of the sender is immaterial, as all that counts is the message that a strategy prescribes after the associated chance move.

Hence, for each $m \in M^*$, $\mu_2(\cdot|m)$ is derived from $\mu_2(\cdot|h^0)$. After observing each $m \in M \backslash M^*$, the receiver must abandon his initial theory $(\nu^\theta)_{\theta\in\Theta}$ of the sender's behavior, but we do not want the receiver to abandon his prior belief about the chance move. So, we want $\mu_2(\cdot|m)$ to be derived from a theory $\nu \in \Delta(\Theta \times M^\Theta)$ that does not contradict the prior: $\nu(\{\theta\} \times M^\Theta) = \theta$ for each $\theta \in \Theta$. As a stand-alone requirement, it is immaterial for $\mu_2(\cdot|m)$: every belief over $S(m)$ can be derived by conditioning from a belief ν over $\Theta \times M^\Theta$ that follows the prior. Yet, this requirement is not immaterial for the updated belief $\mu_2(\cdot|m)$ when ν is also required to assign probability 1 to a subset of $\Theta \times M^\Theta$. Since the theory ν has no place in the array of beliefs of a CPS, we will introduce this requirement directly in the rationalizability procedure.

With this, we can define Strong-Δ-Rationalizability for the problem at hand, which for future reference we call "Path-rationalizability with first-order independence". For every $(\theta, m, \mu_1) \in \Theta \times M \times \Delta_1^\mu$, let $u_1(\theta, m, \mu_1)$ denote the sender's expected payoff after θ and m given the marginal distribution over actions of the receiver induced by μ_1.

Definition 3 *Consider the following reduction procedure.*

Step 0 *Let $\Sigma_1^0 = M^\Theta$ and $\Sigma_2^0 = A^M$.*

Step n *For each $\sigma_1 \in \Sigma_1^{n-1}$, let $\sigma_1 \in \Sigma_1^n$ if there exists $\mu_1 \in \Delta_1$ such that:*

S1' $\mu_1(\Sigma_2^{n-1}) = 1$;

S2' for each $\theta \in \Theta$ and $m \in M$,

$$u_1(\theta, \sigma_1(\theta), \mu_1) \geq u_1(\theta, m, \mu_1).$$

For each $\sigma_2 \in \Sigma_2^{n-1}$, let $\sigma_2 \in \Sigma_2^n$ if there exists $\mu_2 \in \Delta_2$ such that:

R1' $\mu_2(\Theta \times \Sigma_1^{n-1}|h^0) = 1$;

R2' for each $m \in M \backslash M^*$,[12] if $\sigma_1 \in \Sigma_1^{n-1} \cap M^\Theta(m) \neq \emptyset$, there is $\eta \in \Delta(M^\Theta)$ such that $\eta(\Sigma_1^{n-1}) = 1$ and $\mu_2(\cdot|m) = (p \times \eta)|S(m)$;

R3' for every $m \in M$ and $a \in A$,

$$u_2(\mathrm{marg}_\Theta \mu_2(\cdot|m), m, \sigma_2(m)) \geq u_2(\mathrm{marg}_\Theta \mu_2(\cdot|m), m, a).$$

Finally, let $\Sigma_1^\infty = \cap_{n\geq 0}\Sigma_1^n$ and $\Sigma_2^\infty = \cap_{n\geq 0}\Sigma_2^n$. The strategies in Σ_1^∞ and Σ_2^∞ are called path-rationalizable with first-order independence.

We anticipated that, for each $m \in M \backslash M^*$, we wanted $\mu_2(\cdot|m)$ to be derived by conditioning from a theory $\nu \in \Delta(\Theta \times M^\Theta)$ that follows the prior. Here R2' is also requiring that ν be a product measure between the prior and a probability measure over the sender's strategies. Since we impose $\nu(\Theta \times \Sigma_1^{n-1}) = 1$, one would expect that the independence between the chance move and the behavior of the sender follows from the fact that, for each $n > 1$, Σ_1^{n-1} has been derived under the independence assumption between the chance move

[12] Recall that M^* is the set of all $m \in M$ such that $\nu^\theta(m) > 0$ for some $\theta \in \Theta$.

and the sender's first-order belief. This is not true, because if ν associates different strategies in Σ_1^1 with different chance moves, it is as if the sender is deemed changing her view and behavior after seeing the chance move. We need therefore $\mu_2(\cdot|m)$ to be derived from a product measure $p \times \eta$ — an assumption of *structural consistency* (Kreps and Wilson, 1982) of the receiver's beliefs. Thus, imposing independence of the sender's first-order beliefs is not enough for our purpose; we also need a restriction that has to do with independence for the receiver's first-order beliefs.

Now we present the solution concept used by Sobel et al. (1990). We will slightly modify their definition in terms of language and by letting both players do each step of reasoning instead of alternating between them. For simplicity, we also let the sender have the same available messages after every chance move. Fix an equilibrium (β_1^*, β_2^*). Modify the game by substituting the on-path messages

$$M^* := \{m \in M | \exists \theta \in \Theta, \beta_1^*(m|\theta) > 0\}$$

with a unique message m^* that terminates the game. Let $\widetilde{M} = M \backslash M^*$, $\widehat{M} = \widetilde{M} \cup \{m^*\}$ and $u_1(\theta, m^*, \cdot) = u_1(\theta, m, \beta_2^*)$ for any $m \in M^*$ with $\beta_1^*(m|\theta) > 0$.

Definition 4 *Consider the following reduction procedure.*

Step 0 *Let* $\widehat{\Sigma}_1^0 = \widehat{M}^\Theta$ *and* $\widehat{\Sigma}_2^0 = A^{\widetilde{M}}$.

Step n *For each* $\hat{\sigma}_1 \in \widehat{\Sigma}_1^{n-1}$, *let* $\hat{\sigma}_1 \in \widehat{\Sigma}_1^n$ *if there exists* $\mu_1 \in \Delta(A^{\widetilde{M}})$ *such that:*

SF1 $\mu_1(\widehat{\Sigma}_2^{n-1}) = 1$;

SF2 for each $\theta \in \Theta$ *and* $m \in \widehat{M}$,

$$u_1(\theta, \hat{\sigma}_1(\theta), \mu_1) \geq u_1(\theta, m, \mu_1).$$

For each $\hat{\sigma}_2 \in \widehat{\Sigma}_2^{n-1}$, *let* $\hat{\sigma}_2 \in \widehat{\Sigma}_2^n$ *if, for each* $m \in \widetilde{M}$ *with* $\widehat{\Sigma}_1^n \cap \widehat{M}^\Theta(m) \neq \emptyset$, *there exists* $\eta \in \Delta(\widehat{M}^\Theta)$ *such that:*

RF1 $\eta\left(\widehat{\Sigma}_1^{n-1}\right) = 1$;

RF2 $\eta\left(\widehat{M}^\Theta(m)\right) > 0$;

RF3 for every $a \in A$,

$$u_2(p', m, \hat{\sigma}_2(m)) \geq u_2(p', m, a),$$

where $p' \in \Delta(\Theta)$ *is derived from* p *and* η *with Bayes rule.*

Say that (β_1^*, β_2^*) *determines a fixed-equilibrium rationalizable outcome if there is* $\widehat{\sigma}_1 \in \cap_{n>0}\widehat{\Sigma}_1^n$ *with* $\widehat{\sigma}_1(\theta) = m^*$ *for each* $\theta \in \Theta$.

The belief over chance moves derived with Bayes rule from p and η given message m coincides with the marginal over chance moves of $p\times\eta$ after conditioning on $\widehat{S}(m)$: for each $\theta\in\Theta$, we have

$$\begin{aligned}
p'(\theta) &= \frac{\eta\left(\widehat{M}^{\Theta}(\theta,m)\right)p(\theta)}{\sum_{\theta'\in\Theta}\eta\left(\widehat{M}^{\Theta}(\theta',m)\right)p(\theta')} = \frac{(p\times\eta)\left(\{\theta\}\times\widehat{M}^{\Theta}(\theta,m)\right)}{\sum_{\theta'\in\Theta}(p\times\eta)\left(\{\theta'\}\times\widehat{M}^{\Theta}(\theta',m)\right)} = \\
&= \frac{(p\times\eta)\left(\widehat{S}(\theta,m)\right)}{\sum_{\theta'\in\Theta}(p\times\eta)\left(\widehat{S}(\theta',m)\right)} = \frac{(p\times\eta)\left(\widehat{S}(\theta,m)\right)}{(p\times\eta)\left(\widehat{S}(m)\right)} = \\
&= \frac{(p\times\eta)\left(\left(\{\theta\}\times\widehat{M}^{\Theta}\right)\cap\widehat{S}(m)\right)}{(p\times\eta)\left(\widehat{S}(m)\right)} = (p\times\eta)\left(\{\theta\}\times\widehat{M}^{\Theta}\,\middle|\,\widehat{S}(m)\right).
\end{aligned}$$

Now we can formalize the equivalence between the two procedures

Proposition 2 *Fix an equilibrium (β_1^*,β_2^*) and let $\left(\nu^{\theta}\right)_{\theta\in\Theta}=(\beta_1^*(\cdot|\theta))_{\theta\in\Theta}$, $(\nu^m)_{m\in M}=(\beta_2^*(\cdot|m))_{m\in M}$. Thus, (β_1^*,β_2^*) determines a fixed-equilibrium rationalizable outcome if and only if $\Sigma_1^{\infty}\neq\emptyset$.*

Proof. Let φ be the map that associates each $m\in\widetilde{M}$ with itself and each $m\in M^*$ with m^*. Let ς be the map that associates each $\sigma_1=(\sigma_1(\theta))_{\theta\in\Theta}$ with $\varsigma(\sigma_1)=(\varphi(\sigma_1(\theta)))_{\theta\in\Theta}\in\widehat{M}^{\Theta}$. Let ϱ be the map that associates each $\sigma_2=(\sigma_2(m))_{m\in M}$ with $\varrho(\sigma_2)=(\sigma_2(m))_{m\in\widetilde{M}}\in A^{\widetilde{M}}$. Call $\widehat{\sigma}_1^*$ the strategy $\widehat{\sigma}_1\in\widehat{M}^{\Theta}$ such that $\widehat{\sigma}_1(\theta)=m^*$ for each $\theta\in\Theta$. Fix $n\geq 0$ and assume by way of induction that $\widehat{\Sigma}_1^n=\varsigma(\Sigma_1^n)$. The induction hypothesis is true for $n=0$ because $\varsigma(M^{\Theta})=\widehat{M}^{\Theta}$ and $\varrho(A^M)=A^{\widetilde{M}}$. We are going to show that the induction hypothesis is true for $n+1$ as long as $\widehat{\sigma}_1^*\in\widehat{\Sigma}_1^n$, otherwise $\Sigma_1^n=\emptyset$. So, if $\widehat{\sigma}_1^*\in\cap_{n>0}\widehat{\Sigma}_1^n$, by finiteness of the game there is $\sigma_1\in\varsigma^{-1}(\widehat{\sigma}_1^*)$ such that $\sigma_1\in\cap_{n>0}\Sigma_1^n$, establishing $\Sigma_1^{\infty}\neq\emptyset$; else, $\Sigma_1^{\infty}=\emptyset$, completing the proof.

Suppose first that $\widehat{\sigma}_1^*\notin\widehat{\Sigma}_1^n$. Then, by the induction hypothesis, for every $\sigma_1\in\Sigma_1^n$, there exists $\theta\in\Theta$ such that $\sigma_1(\theta)\notin M^*$. But then, there is no $\mu_2\in\Delta_2$ with $\mu_2(\Sigma_1^n|h^0)=1$, thus $\Sigma_2^{n+1}=\emptyset$.

Suppose now that $\widehat{\sigma}_1^*\in\widehat{\Sigma}_1^n$. Then, by the induction hypothesis, there exists $\sigma_1^*\in\Sigma_1^n$ such that $\sigma_1(\theta)\in M^*$.for each $\theta\in\Theta$. Fix $\mu_1^*\in\Delta_1$ with $\mu_1^*(\Sigma_2^{n-1}|h^0)=1$ such that, for each $\theta\in\Theta$ and $m'\in M$,

$$u_1(\theta,\sigma_1^*(\theta),\mu_1^*)\geq u_1(\theta,m',\mu_1^*).$$

For each $\theta\in\Theta$ and $m\in M^*$ with $\nu^{\theta}(m)>0$, since (β_1^*,β_2^*) is an equilibrium and $\mu_1^*\in\Delta_1$, we have

$$u_1(\theta,m,\mu_1^*)\geq u_1(\theta,\sigma_1^*(\theta),\mu_1^*).$$

Therefore, μ_1^* justifies every $\sigma_1\in M^{\Theta}$ such that $\nu^{\theta}(\sigma_1(\theta))>0$ for each $\theta\in\Theta$. Hence, there exists $\mu_2^*\in\Delta_2$ such that $\mu_2^*(\Sigma_1^n|h^0)=1$. We will use μ_2^* later.

Fix $\sigma_2\in\Sigma_2^n$. Thus, for each $m\in\widetilde{M}$, there is $\eta^m\in\Delta(M^{\Theta})$ such that $\eta^m\left(\Sigma_1^{n-1}\right)=1$ and for every $a\in A$,

$$u_2(\mathrm{marg}_{\Theta}((p\times\eta^m)\,|S(m)),m,\sigma_2(m))\geq u_2(\mathrm{marg}_{\Theta}((p\times\eta^m)\,|S(m)),m,a), \tag{7}$$

Define $\widehat{\eta}^m \in \Delta(\widehat{M}^\Theta)$ as follows: for each $\widehat{\sigma}_1 \in \widehat{M}^\Theta$, let $\widehat{\eta}^m(\widehat{\sigma}_1) = \eta^m(\varsigma^{-1}(\widehat{\sigma}_1))$. Thus, by the induction hypothesis, $\widehat{\eta}^m(\Sigma_1^{n-1}) = 1$, satisfying R1*. For each $\theta \in \Theta$, $\widehat{\sigma}_1(\theta) = m$ if and only if, for each $\sigma_1 \in \varsigma^{-1}(\widehat{\sigma}_1)$, $\sigma_1(\theta) = m$; therefore, $\widehat{\eta}^m(\widehat{M}^\Theta(m)) > 0$, satisfying R2', and

$$\mathrm{marg}_\Theta((p \times \eta^m)\,|S(m)) = \mathrm{marg}_\Theta((p \times \widehat{\eta}^m)\,|\widehat{S}(m)), \tag{8}$$

Since $\mathrm{marg}_\Theta((p \times \widehat{\eta}^m)\,|\widehat{S}(m))$ is the posterior p' derived from p and $\widehat{\eta}^m$ by Bayes rule, by (9), (9), and $\varrho(\sigma_2)(m) = \sigma_2(m)$, $\varrho(\sigma_2)$ and $\widehat{\eta}^m$ satisfy R3'. Thus, $\varrho(\sigma_2) \in \widehat{\Sigma}_2^n$.

Now fix $\sigma_1 \in \Sigma_1^{n+1}$ and $\mu_1 \in \Delta_1$ with $\mu_1(\Sigma_2^n) = 1$ that justifies σ_1. Define $\widehat{\mu}_1 \in \Delta(A^{\widetilde{M}})$ with $\widehat{\mu}_1(\widehat{\Sigma}_2^n) = 1$ as $\widehat{\mu}_1(\widehat{\sigma}_2) = \widehat{\mu}_1(\varrho^{-1}(\widehat{\sigma}_2))$ for each $\widehat{\sigma}_2 \in A^{\widetilde{M}}$. Then, the sender expects the same payoff under $\widehat{\mu}_1$ and μ_1 after each $m \in \widetilde{M}$, and the equilibrium payoff after m^* / each $m \in M^*$. Hence, $\widehat{\mu}_1$ justifies every $\widehat{\sigma}_1 \in \varsigma(\sigma_1)$. This establishes $\widehat{\Sigma}_1^{n+1} \supseteq \varsigma(\Sigma_1^{n+1})$.

Fix $\widehat{\sigma}_2 \in \widehat{\Sigma}_2^n$. Thus, for each $m \in \widetilde{M}$, there is $\widehat{\eta}^m \in \Delta(\widehat{M}^\Theta)$ such that $\widehat{\eta}^m\left(\widehat{\Sigma}_1^{n-1}\right) = 1$ and for every $a \in A$,

$$u_2(p', m, \hat{\sigma}_2(m)) \geq u_2(p', m, a), \tag{9}$$

where $p' \in \Delta(\Theta)$ is derived from p and $\widehat{\eta}^m$ with Bayes rule (thus $p' = \mathrm{marg}_\Theta((p \times \widehat{\eta}^m)\,|\widehat{S}(m))$). Define $\eta^m \in \Delta(M^\Theta)$ as follows: for each $\widehat{\sigma}_1 \in \widehat{\Sigma}_1^{n-1}$, let $\eta^m(\sigma_1) = \widehat{\eta}^m(\widehat{\sigma}_1)$ for some $\sigma_1 \in \Sigma_1^{n-1} \cap \varsigma^{-1}(\widehat{\sigma}_1)$, which is non-empty by the induction hypothesis. Thus, $\widehat{\eta}^m(\Sigma_1^{n-1}) = 1$. For each $\theta \in \Theta$, $\sigma_1(\theta) = m$ if and only if $\widehat{\sigma}_1(\theta) = m$, therefore

$$\mathrm{marg}_\Theta((p \times \eta^m)\,|S(m)) = \mathrm{marg}_\Theta((p \times \widehat{\eta}^m)\,|\widehat{S}(m)) = p'. \tag{10}$$

Fix $\sigma_2 \in \varrho^{-1}(\widehat{\sigma}_2)$ such that $\nu^m(\sigma_2(m)) > 0$ for each $m \in M^*$. Construct $\mu_2 \in \Delta_2$ that satisfies R1' and R2' at step n as $\mu_2(\cdot|h^0) = \mu_2^*(\cdot|h^0)$ and, for each $m \in \widetilde{M}$, $\mu_2(\cdot|m) = (p \times \eta^m)\,|S(m)$. For each $m \in \widetilde{M}$, by (9), (9), and $\sigma_2(m) = \widehat{\sigma}_2(m)$, $\sigma_2(m)$ and $\mu_2(\cdot|m)$ satisfy R3'. For each $m \in M^*$, $\sigma_2(m)$ and $\mu_2(\cdot|m)$ satisfy R3' as well, because $\sigma_2(m)$ is an equilibrium action and $\mu_2(\cdot|m)$ is the equilibrium belief. Thus, $\sigma_2 \in \Sigma_2^n$.

Now fix $\widehat{\sigma}_1 \in \widehat{\Sigma}_1^{n+1}$ and $\widehat{\mu}_1 \in \Delta(A^{\widetilde{M}})$ with $\widehat{\mu}_1(\widehat{\Sigma}_2^n) = 1$ that justifies $\widehat{\sigma}_1$. Transform $\widehat{\mu}_1$ into $\mu_1 \in \Delta_1$ with $\mu_1(\Sigma_2^n) = 1$ by redistributing each probability value $\widehat{\mu}_1(\widehat{\sigma}_2)$ over the strategies $\sigma_2 \in \varrho^{-1}(\widehat{\sigma}_2)$ such that $\nu^m(\sigma_2(m)) > 0$ for each $m \in M^*$. Then, the sender expects the same payoff under $\widehat{\mu}_1$ and μ_1 after each $m \in \widetilde{M}$, and the equilibrium payoff after m^* / each $m \in M^*$. Hence, μ_1 justifies every $\sigma_1 \in \varsigma^{-1}(\widehat{\sigma}_1)$ such that, for each $\theta \in \Theta$, if $\widehat{\sigma}_1(\theta) = m^*$, then $\nu^\theta(\sigma_1(\theta)) > 0$. This establishes $\widehat{\Sigma}_1^{n+1} \subseteq \varsigma(\Sigma_1^{n+1})$. ■

Finally, we show the equivalence between Path-rationalizability with first-order independence and Path-rationalizability with second-order independence. To distinguish them, we put a bar on the sets and delta-restrictions of Path-rationalizability with second-order independence.

Proposition 3 *A strategy of the receiver is path-rationalizable with first-order independence if and only if is path-rationalizable with second-order independence. For every strategy of the sender that is path-rationalizable with first-order independence, every prescribed message is path-rationalizable with second-order independence for the type; for every type and for every message that is path-rationalizable with second-order independence, there is a strategy that is path-rationalizable with first-order independence and prescribes the message to the type.*

Proof.

Assume by way of induction that (i) $\Sigma_2^n = \mathrm{Proj}_{A^M}\overline{\Sigma}_2^n$, (ii) for every $\sigma_1 \in \Sigma_1^n$, for every $\theta \in \Theta$, $\sigma_1(\theta) \in \mathrm{Proj}_M\overline{\Sigma}_{1,\theta}^n$, and (iii) for every $\overline{\mu}_1 \in \cup_{\theta\in\Theta}\mathrm{Proj}_{\Delta(A^M)}\overline{\Sigma}_{1,\theta}^n$, for every $\sigma_1 \in M^\Theta$ such that $(\sigma_1(\theta), \overline{\mu}_1) \in \overline{\Sigma}_{1,\theta}^n$ for each $\theta \in \Theta$, we have $\sigma_1 \in \Sigma_1^n$. All is trivally true for $n = 0$.

Recall that $\Delta_1 = \overline{\Delta}_1$, and by the induction hypothesis (i), R1 and R1' coincide. Then, (ii) and (iii) for $n+1$ follow.

To prove (i) for $n+1$, observe preliminarly that by the set of messages prescribed by some strategy $\sigma_1 \in \Sigma_1^n$ to some type coincides with the set of messages m such that $m \in \mathrm{Proj}_M\overline{\Sigma}_{1,\theta}^n$ for some $\theta \in \Theta$: one inclusion follows directly from induction hypothesis (ii), the opposite inclusion follows from induction hypothesis (iii) and the observation that if $\overline{\mu}_1 \in \cup_{\theta\in\Theta}\mathrm{Proj}_{\Delta(A^M)}\overline{\Sigma}_{1,\theta}^n$, then $\overline{\mu}_1 \in \mathrm{Proj}_{\Delta(A^M)}\overline{\Sigma}_{1,\theta}^n$ for every $\theta \in \Theta$. Let $M^\#$ denote the set of such messages.

Now fix $(\overline{\sigma}_2, \overline{\mu}_2) \in \overline{\Sigma}_2^{n+1}$. Construct a CPS μ_2 as follows.

For each $(\theta, m) \in \Theta \times M$ such that $\overline{\mu}_2(\{(\theta, m)\} \times \Delta(A^M)|h^0) > 0$, since $\overline{\mu}_2$ satisfies R1, $m \in \mathrm{Proj}_M\overline{\Sigma}_{1,\theta}^n$. By the induction hypothesis, there is $\sigma_1 \in \Sigma_1^n$ such that $\sigma_1(\theta) = m$. Let $\mu_2(\cdot|h^0)$ assign to (θ, σ_1) probability $\overline{\mu}_2(\{(\theta, m)\} \times \Delta(A^M)|h^0)$.[13] Thus, $\mu_2(\Theta \times \Sigma_1^n|h^0) = 1$, satisfying R1'. Since $\overline{\mu}_2(\cdot|h^0)$ satisfies B1 and B3, for each $(\theta, m) \in \Theta \times M$ with $\nu^\theta(m) > 0$ we have $\overline{\mu}_2(\{(\theta, m)\} \times \Delta(A^M)|h^0) = p(\theta) \cdot \nu^\theta(m)$. Thus, by construction,

$$\mu_2(\{\theta\} \times \left\{\sigma_1 \in M^\Theta | \sigma_1(\theta) = m\right\} |h^0) = p(\theta) \cdot \nu^\theta(m), \tag{11}$$

thus $\mathrm{marg}_\Theta\mu_2(\cdot|h^0) = \mathrm{marg}_\Theta\overline{\mu}_2(\cdot|h^0)$.

For each $m \in M^*$, derive $\mu_2(\cdot|m)$ from $\mu_2(\cdot|h^0)$ by conditioning.

For each $m \in M^\#\backslash M^*$, since $\overline{\mu}_2$ satisfies R2, there is $\nu \in \Delta(\Theta \times M \times \Delta(A^M))$ that satisfies B1 and B2 such that $\nu(\cup_{\theta\in\Theta}\left(\{\theta\} \times \overline{\Sigma}_{1,\theta}^n\right)) = 1$ and $\overline{\mu}_2(\cdot|m) = \nu|m$. Let B be the finite set of all μ_1 such that $\nu(\Theta \times M \times \{\mu_1\}) > 0$. Construct $\nu' \in \Delta(\Theta \times M^\Theta \times \Delta(A^M))$ as follows: for each $(\theta, \sigma_1, \mu_1) \in \Theta \times M^\Theta \times \Delta(A^M)$, if $\mu_1 \in B$ (thus by B2 $\nu((\{\theta'\} \times M \times \{\mu_1\})) > 0$ for every $\theta' \neq \theta$), let

$$\nu'(\theta, \sigma_1, \mu_1) = \nu(\theta, M, \mu_1) \cdot \frac{\prod_{\theta'\in\Theta}\nu(\theta', \sigma_1(\theta'), \mu_1)}{\sum\limits_{\sigma_1'\in M^\Theta}\prod_{\theta'\in\Theta}\nu(\theta', \sigma_1'(\theta'), \mu_1)}$$

(thus the denominator is positive), otherwise let $\nu'(\theta, \sigma_1, \mu_1) = 0$. For each $(\theta, m) \in \Theta \times M$,

[13] Note that, for every type θ, every associated message m gives rise to a different strategy σ_1, so $\mu_2(\cdot|h^0)$ is well-defined.

we have

$$\begin{aligned}
\nu'(\{\theta\} \times M^{\Theta}(\theta,m) \times \Delta(A^M)) &= \sum_{\mu_1 \in B} \sum_{\sigma_1 \in M^{\Theta}(\theta,m)} \nu'((\theta,\sigma_1,\mu_1)) = \\
&= \sum_{\mu_1 \in B} \sum_{\sigma_1 \in M^{\Theta}(\theta,m)} \nu(\theta, M, \mu_1) \cdot \frac{\prod_{\theta' \in \Theta} \nu(\theta', \sigma_1(\theta'), \mu_1)}{\sum\limits_{\sigma_1' \in M^{\Theta}} \prod_{\theta' \in \Theta} \nu(\theta', \sigma_1'(\theta'), \mu_1)} = \\
&= \sum_{\mu_1 \in B} \nu((\theta, m, \mu_1)) \cdot \frac{\sum\limits_{\sigma_1 \in M^{\Theta}(\theta,m)} \nu(\theta, M, \mu_1) \cdot \prod_{\theta' \neq \theta} \nu(\theta', \sigma_1(\theta'), \mu_1)}{\sum\limits_{\sigma_1' \in M^{\Theta}} \prod_{\theta' \in \Theta} \nu(\theta', \sigma_1'(\theta'), \mu_1)} = \\
&= \sum_{\mu_1 \in B} \nu((\theta, m, \mu_1)) \cdot \frac{\sum\limits_{\sigma_1 \in M^{\Theta}(\theta,m)} \sum\limits_{m' \in M} \nu(\theta, m', \mu_1) \prod_{\theta' \neq \theta} \nu(\theta', \sigma_1(\theta'), \mu_1)}{\sum\limits_{\sigma_1' \in M^{\Theta}} \prod_{\theta' \in \Theta} \nu(\theta', \sigma_1'(\theta'), \mu_1)} = \\
&= \sum_{\mu_1 \in B} \nu((\theta, m, \mu_1)) = \\
&= \nu(\{(\theta, m)\} \times \Delta(A^M)).
\end{aligned}$$

Hence, $\mathrm{marg}_{\Theta}(\nu'|S(m)) = \mathrm{marg}_{\Theta}(\nu|m) = \mathrm{marg}_{\Theta}\overline{\mu}_2(\cdot|m)$ (and

$$\nu'(\Theta \times M^{\Theta} \times \Delta(A^M)) = \sum_{(\theta,m) \in \Theta \times M} \nu'(\{\theta\} \times M^{\Theta}(\theta,m) \times \Delta(A^M)) = \sum_{(\theta,m) \in \Theta \times M} \nu(\{(\theta,m)\} \times \Delta(A^M)) = 1,$$

i.e., ν' is well-defined). Furthermore, for each $(\theta, \sigma_1) \in \Theta \times M^{\Theta}$,

$$\begin{aligned}
\nu'(\{(\theta,\sigma_1)\} \times \Delta(A^M)) &= \sum_{\mu_1 \in B} \nu(\theta, M, \mu_1) \frac{\prod_{\theta' \in \Theta} \nu(\theta', \sigma_1(\theta'), \mu_1)}{\sum\limits_{\sigma_1' \in M^{\Theta}} \prod_{\theta' \in \Theta} \nu(\theta', \sigma_1'(\theta'), \mu_1)} = \\
&= p(\theta) \cdot \sum_{\mu_1 \in B} \nu(\Theta, M, \mu_1) \frac{\prod_{\theta' \in \Theta} \nu(\theta', \sigma_1(\theta'), \mu_1)}{\sum\limits_{\sigma_1' \in M^{\Theta}} \prod_{\theta' \in \Theta} \nu(\theta', \sigma_1'(\theta'), \mu_1)}.
\end{aligned}$$

where the second equality follow from the fact that ν satisfies B1 and B2. Hence, $\mathrm{marg}_{\Theta \times M^{\Theta}} \nu'$ is a product measure $p \times \mathrm{marg}_{M^{\Theta}} \nu'$. Finally, $\nu'((\theta, \sigma_1, \overline{\mu}_1)) > 0$ if and only if $\nu((\theta', \sigma_1(\theta'), \overline{\mu}_1)) > 0$ for every $\theta' \in \Theta$, hence, by $\nu(\cup_{\theta \in \Theta} (\{\theta\} \times \overline{\Sigma}^n_{1,\theta})) = 1$ and the induction hypothesis (iii), $\nu'(\Theta \times \Sigma^n_1 \times \Delta(A^M)) = 1$. Thus, $\mu_2(\cdot|m) = \mathrm{marg}_{\Theta \times M^{\Theta}} \nu' | S(m)$ satisfies R2'.

For each $m \notin M^{\#}$,[14] let $\mu_2(\cdot|m)$ be any $\nu \in \Delta(\Theta \times M^{\Theta})$ such that $\mathrm{marg}_{\Theta}\mu_2(\cdot|m) = \mathrm{marg}_{\Theta}\overline{\mu}_2(\cdot|m)$.

It is easy to see that μ_2 is a CPS. Moreover, μ_2 satisfies (6) by (11), hence $\mu_2 \in \Delta_2$.

Since $(\overline{\sigma}_2, \overline{\mu}_2)$ satisfies R3 and $\mathrm{marg}_{\Theta}\mu_2(\cdot|m) = \mathrm{marg}_{\Theta}\overline{\mu}_2(\cdot|m)$ for each $m \in M$, $(\overline{\sigma}_2, \mu_2)$ satisfies R3', thus $\overline{\sigma}_2 \in \overline{\Sigma}^{n+1}_2$.

Now fix $\sigma_2 \in \Sigma^{n+1}_2$. Fix $\mu_2 \in \Delta_2$ that satisfies R1', R2', and R3' with σ_2. Construct a CPS $\overline{\mu}_2$ over $\Theta \times M \times \Delta(A^M)$ as follows. Let ς be the map that associates each $(\theta, \sigma_1) \in$

[14] That $M^* \subseteq M^{\#}$ follows from the existence of $(\overline{\sigma}_2, \overline{\mu}_2) \in \overline{\Sigma}^{n+1}_2$.

$\Theta \times M^\Theta$ with $(\theta, \sigma_1(\theta)) \in \Theta \times M$. Let τ be the map that associates each $\sigma_1 \in \Sigma_1^n$ with any $\mu_1 \in \Delta_1$ such that, if $n > 0$, σ_1 and μ_1 satisfy S1' and S2' at step n. Let $\nu \in \Delta(\Theta \times M)$ be the pushforward of $\mu_2(\cdot|h^0)$ through ς. Since $\mu_2(\cdot|h^0)$ satisfies R1', by the induction hypothesis (ii) we have $\nu(\cup_{\theta\in\Theta} \{\theta\} \times \mathrm{Proj}_M \overline{\Sigma}_{1,\theta}^n) = 1$, and we can fix $\mu_1 \in \Delta_1$ and $\sigma_1 \in M^\Theta$ with $\mu_2(\Theta \times \{\sigma_1\})|h^0) > 0$ that, if $n > 1$, satisfy S1' and S2' at step n. Hence, $\mu_1 \in \overline{\Delta}_1$, and if $n > 0$, $\mu_1(\Sigma_2^{n-1}) = 1$, so by the induction hypothesis (i), $\mu_1(\overline{\Sigma}_2^{n-1}) = 1$. Moreover, since $\mu_2 \in \Delta_2$, $\nu^\theta(\sigma_1(\theta)) > 0$ for every $\theta \in \Theta$, and thus, if $n > 1$, μ_1 justifies every $m \in M^*$ with $\nu^\theta(m) > 0$ (recall that all messages played in equilibrium by a type must give the same expected payoff under every equilibrium conjecture). Let $\overline{\mu}_2(\cdot|h^0) = \nu \times \delta_{\mu_1}$. Thus, $\overline{\mu}_2(\cdot|h^0)$ satisfies B2 and it inherits B1 and B3 from (6). Moreover, it satisfies R1 because for each $(\theta, m) \in \Theta \times M$ with $\nu((\theta, m)) > 0$, $(m, \mu_1) \in \overline{\Sigma}_{1,\theta}^n$. For each $m \in M^*$, derive $\mu_2(\cdot|m)$ by conditioning. Thus, $\mathrm{marg}_\Theta \overline{\mu}_2(\cdot|m) = \mathrm{marg}_\Theta \mu_2(\cdot|m)$. For each $m \in M^\# \backslash M^*$, by R2', $\mu_2(\cdot|m) = (p \times \eta)\,|S(m)$ for some $\eta \in \Delta(\Sigma_1^n)$. Let ψ be the map that associates each $(\theta, \sigma_1) \in \Theta \times M^\Theta$ with $(\varsigma(\theta, \sigma_1), \tau(\sigma_1)) \in \Theta \times M \times \Delta(A^M)$. Let ν be the pushfoward of $p \times \eta$ through ψ. It is easy to see that ν satisfies B1 and B2. Moreover, by the induction hypothesis (ii), $\nu(\cup_{\theta\in\Theta} \{\theta\} \times \mathrm{Proj}_M \overline{\Sigma}_{1,\theta}^n) = 1$. So, $\overline{\mu}_2(\cdot|m) = \nu|m$ satisfies R2, and $\mathrm{marg}_\Theta \overline{\mu}_2(\cdot|m) = \mathrm{marg}_\Theta \mu_2(\cdot|m)$. For each $m \in M \backslash M^\#$, let $\overline{\mu}_2(\cdot|m) = \mu_2(\cdot|m)$. It is easy to see that $\overline{\mu}_2 \in \Delta_2$. Since $\overline{\mu}_2(\cdot|m) = \mathrm{marg}_\Theta \mu_2(\cdot|m)$ for every $m \in M$, since σ_2 satisfies R3' with μ_2, it satisfies R3' with $\overline{\mu}_2$. Hence, $\sigma_2 \in \overline{\Sigma}_2^{n+1}$. ■

References

[1] Banks J.S. and J. Sobel (1987): "Equilibrium selection in signaling games," *Econometrica,* 55(3), 647-661.

[2] Battigalli P. (2003): "Rationalizability in Infinite, Dynamic Games of Incomplete Information," *Research in Economics,* 57, 1-38.

[3] Battigalli P. and M. Siniscalchi (2003): "Rationalization and Incomplete Information," *The B.E. Journal of Theoretical Economics*, 3(1), 1-46.

[4] Battigalli P. and M. Siniscalchi (2007): "Interactive epistemology in games with payoff uncertainty," *Research in Economics*, 61, 165-184.

[5] Battigalli P. and A. Prestipino (2013): "Transparent Restrictions on Beliefs and Forward Induction Reasoning in Games with Asymmetric Information", *The B.E. Journal of Theoretical Economics*, 13(1), 79-130.

[6] Cho, I.K. (1987): "A refinement of sequential equilibrium," *Econometrica*, 55(6), 1367-1389.

[7] Cho, I.K. and D. Kreps (1987): "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 102(2), 179-221.

[8] Fudenberg, D. and K. He (2018): "Learning and type compatibility in signaling games," *Econometrica*, 86, 1215-1255.

[9] Fudenberg, D. and Y. Kamada (2015): "Rationalizable Partition-Confirmed Equilibrium,". *Theoretical Economics*, 10, 775-806.

[10] Kreps, D. and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50, 863-894.

[11] Sobel, J., L. Stole, I. Zapater (1990): ""Fixed-Equilibrium Rationalizability in Signaling Games," *Journal of Economic Theory*, 52, 304-331.