Partially Constructed Sequential Equilibrium (Note)

Joel Watson*

27 November 2023

Abstract

This note identifies sufficient conditions under which a partial construction of sequential equilibrium—characterizing beliefs and behavior for only a subset of the information sets—is sufficient for the existence of a sequential equilibrium that coincides with the partial construction in a suitable way.

1 Introduction

Various applications of noncooperative game theory call for the solution concept of sequential equilibrium, requiring "full consistency" of beliefs across information sets, in contrast to versions of perfect Bayesian equilibrium that utilize weaker consistency notions. Constructing a sequential equilibrium for a complex game can be challenging; the analyst must describe a sequence of fully mixed behavior strategies that determines the beliefs at all information sets and that converges to the equilibrium strategy profile. This paper identifies sufficient conditions under which a *partial construction* of sequential equilibrium—characterizing beliefs and behavior for only a subset of the information sets—is sufficient for the existence of a sequential equilibrium that coincides with the partial construction in a suitable way.

For a simple warm-up example, consider the extensive-form game shown on the left side of Figure 1, denotes as "First game." Suppose we want to focus on information sets labelled (in blue) h_1, h_2, and h_3, and ignore behavior and beliefs at information sets h'_2 and h'_3. Let $J \equiv \{h_1, h_2, h_3\}$ and write $-J = \{h'_2, h'_3\}$. By limiting attention to J, we can describe (i) the strategy profile restricted to J and (ii) the players' beliefs at the information sets in J about the behavior at these same information sets.

Regarding the strategy, suppose that at h_1 player 1 chooses action a, at h_2 player 2 selects action d, and at h_3 player 3 mixes between r and t with equal probability. Let us express this partial strategy profile as adr and adt each with probability $1/2$, and note that it describes behavior at only the information sets in J. The resulting path of play reaches the terminal node following the action sequence (a, d), where the payoff vector is $(3, 2, 1)$.

As for the beliefs, suppose that at h_1 player 1 believes that players 2 and 3 would behave as the partial strategy prescribes for h_2 and h_3. Likewise, at h_2 player 2 believes that player 1

*UC San Diego; http://econ.ucsd.edu/∼jwatson/. Thanks go to Kelvin Leong for helpful research assistance and to seminar participants at UC San Diego and Washington University in St. Louis for their feedback.

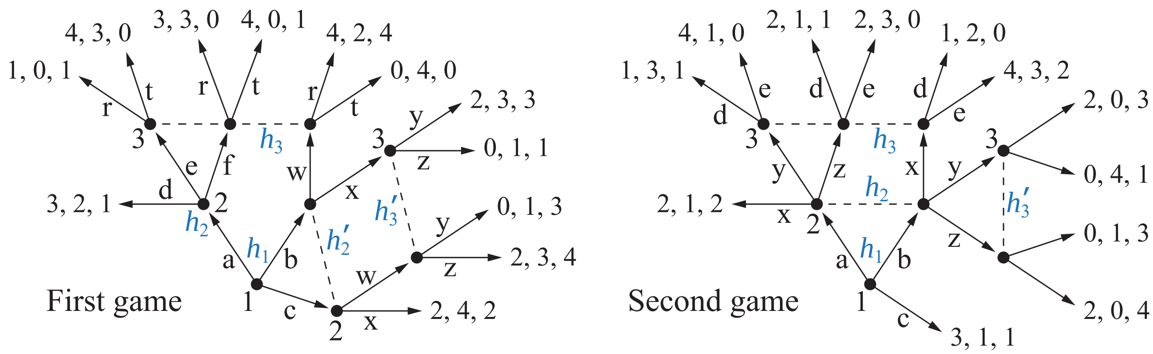

Figure 1: Two examples.

has for sure chosen action a and that player 3 would behave as prescribed for h_3 (choosing r and t with equal probability). Finally, at h_3 player 3 puts probability $1/2$ on the combination of actions ae chosen at h_1 and h_2, and puts probability $1/2$ on the combination af. Note that these are partial beliefs; they are only for the information sets in J and they address only the behavior at the information sets in J.

Though we have not discussed information sets h_2' and h_3', the partial strategies are clearly sequentially rational at h_1, h_2, and h_3, given the partial beliefs we have described. At h_1, action a leads to the payoff 3 for player 1, whereas the expected payoffs of actions b and c cannot exceed 2, regardless of the play at h_2' and h_3'. For instance, if player 1 deviates to take action b and player 2 happens to choose w, then player 1's expected payoff is 2 because player 1 believes that player 3 randomizes between r and t with equal probability at h_3. Likewise, player 2's action d is best at h_2 given her belief that player 3 would randomize in this way. Finally, player 3 is indifferent between r and t at h_3 because her partial belief implies that the path to h_3 was (a, e) with probability $1/2$ and (a, f) with probability $1/2$.

Additionally, the partial beliefs are fully consistent in the following sense: There is a sequence of fully mixed partial behavior strategies that converges to the partial strategy specified above (adr and adt each with probability $1/2$) and has the property that, for each information set in J, the sequence of distributions conditional on reaching this information set converges to the partial belief specified above. Here "conditional on reaching this information set" means conditioning on the subset of partial strategy profiles that would reach the information set when combined with at least one strategy specification for $-J$. For instance, partial strategy bdt reaches h_3 because the path of play would reach h_3 if player 2 were to select action w at $h_2' \in -J$. The set of partial strategy profiles for players 1 and 2 that reach h_3 is $\{\mathrm{ae}, \mathrm{af}, \mathrm{bd}, \mathrm{be}, \mathrm{bf}\}$.

A sequence that works is defined thus, for integer index k: At h_1, the probabilities of actions a, b, and c are, in order, $1 - 2(1/k)^2$, $(1/k)^2$, and $(1/k)^2$. At h_2, the probabilities of actions d, e, and f are, in order, $1 - 2(1/k)$, $1/k$, and $1/k$. And at h_3, the probabilities of actions r and t are each $1/2$. As $k \to \infty$, the conditional probabilities converge to the partial beliefs described above. For instance, at h_3 the condition distribution over $\{\mathrm{a}, \mathrm{b}, \mathrm{c}\} \times \{\mathrm{d}, \mathrm{e}, \mathrm{f}\}$ puts the following probability on both ae and af:

$$\frac{(1 - 2(1/k)^2)(1/k)}{2(1 - 2(1/k)^2)(1/k) + (1/k)^2}$$

As k approaches ∞, this number converges to $1/2$.

Hence, in this example, it seems that we have the makings of a sequential equilibrium. Indeed, there is a way to enrich the partial strategies, partial beliefs, and sequence of fully mixed partial strategies to define a strategy profile for the entire game, beliefs at all information sets, and a sequence of fully mixed behavior strategies that satisfy the conditions of sequential equilibrium. Further, on the margin corresponding to J, the resulting strategy profile and beliefs coincide with the partial strategy and beliefs that we specified above. In other words, our partial sequential-equilibrium construction can be extended into a completely specified sequential equilibrium.

For instance, supplementing the probabilities described in the previous paragraph, let the probability of x at h_2' be $1 - (1/k)$, and let the probability of y at h_3' be $1 - (1/k)$. It is easy to verify that the resulting sequence of fully mixed behavior strategies supports a sequential equilibrium with the strategy profile in which player 1 plays strategy a, player 2 plays dx, and player 3 plays ry and ty with equal probability.

The Theorem of this paper establishes conditions under which, for a given game and set of information sets J, a partial equilibrium construction is guaranteed to extend to a fully specified sequential equilibrium, meaning that we can ignore behavior and beliefs at the information sets in $-J$. Two conditions must be satisfied. The first condition is obvious: Sequential rationality at the information sets in J must be robust to any behavior at the information sets in $-J$. The second condition is subtle and relates to the whether beliefs about actions taken at information sets in J could be disrupted by the specification of behavior at the other information sets. I call the second condition the *rectangular margin-support condition*, for reasons that will become clear. The basic idea is that the information-set structure must allow for independent updating on dimensions J and $-J$ with respect to a particular subset of the strategy space that relates to the support of beliefs in the partial construction.

In the first example, the second condition relates to the following fact: Given the partial-belief sequence specified for J, player 3's belief at h_3 must put probability 0 on path (b, w), regardless of the probabilities put on the actions at h_2' and h_3'. That is, the probabilities on the actions at h_2' and h_3' turn out to be irrelevant for determining player 3's belief about the nodes in h_3, because the probability of path (b, w) is bounded above by $(1/k)^2$, whereas the probabilities of paths (a, e) and (a, f) are both $(1 - 2(1/k)^2)(1/k)$.

For a deeper illustration of the second condition of the Theorem, consider the game on the right side of Figure 1 (Second game). Let $J = \{h_1, h_3\}$, meaning that we wish to ignore information sets h_2 and h_3', and suppose that our partial equilibrium construction is as follows: The sequence of mixing probabilities for h_1 and h_3 is that action c is played with probability $1 - (1/k) - (1/k)^2$, action a with probability $(1/k)$, action b with probability $(1/k)^2$, action d with probability $1 - (1/k)$, and action e with probability $1/k$. In line with these probabilities, we want player 3 at h_3 to believe that action a was chosen for sure at h_1. Then the partial strategy in which player 1 chooses c and player 3 chooses d at h_3' is sequentially rational.

It would seem that we could not guarantee that player 3's belief at h_3 about player 1's action will hold up in a full construction of a sequential equilibrium. This is because, unlike in the first example, the belief at h_3 critically depends on the behavior at $h_2 \in -J$ specified in conjunction with the behavior already described for J. For instance, if actions y and z are both assigned probability $(1/k)^2$ then, as k approaches ∞, player 3's belief at h_3 must put probability 0 on action a having been chosen at h_1. That is, he thinks the path to h_3 was (b, x), which is contrary to what we wanted. This problem arises, in fact, for *any* probabilities that we could pick for actions a and b.

It turns out, however, that for the second example one can still guarantee the existence of a sequential equilibrium of the entire game that preserves the desired belief at h_3, that is, that player 1 chose action a earlier. Thus, we can safely ignore h_2 in the partial construction, so long as sequential rationality at the information sets in J holds regardless of the actions taken at information sets h_2 and h'_3, which is the case.

In this example and for the general result, the construction is accomplished by imposing a lower bound on the probabilities of actions at the information sets in $-J$ that goes to zero more slowly than does the bound on the actions at the information sets in J. We can then use a standard existence result to find Nash equilibria in the space of bounded mixed strategies on domain $-J$ for each k. The key novel element of the proof is then to characterize the limit of the conditional probabilities derived from the combination of mixed partial behavior strategies on domains J and $-J$ under the rectangular margin-support condition.[1]

To conduct the general analysis, it is essential to express beliefs as probability distributions over the space of strategy profiles, which I call *appraisals* (Watson 2017). This is in contrast to expressing beliefs as probability distributions over nodes at information sets (*assessments*, Kreps and Wilson 1982) in conjunction with continuation strategies.[2] The additional structure of appraisals is needed to separate belief components, such the marginal of actions at any particular information set and correlation between actions taken at different information sets. Expressing beliefs in this way is common for defining rationalizability, as in Pearce's (1984) original definition and Battigalli's (1997) restatement, and has been utilized for equilibrium analysis, such as Battigalli's (1996) exploration of independence underlying equilibrium notions and Govindan and Wilson's (2009) study of forward induction. I have refined the approach by incorporating the player's own strategy and paring down the notation to the essentials.

The theory presented here was developed not for simple examples such as those shown in Figure 1 (here to illustrate the theoretical components), but to help with more complex applications. In fact, the Theorem is used to prove the main result in Watson (2023), which examines a class of network contracting games, in which private messages are exchanged between every pair of n players in each of $2n - 1$ stages of time, so there is a huge number of information sets and asymmetric information throughout. Hopefully the theory will give at least one other theorist a useful tool to simplify the process of constructing a sequential equilibrium in an applied setting.

The next section lays out notation and reviews definitions pertaining to the extensive-form representation, beliefs, full consistency, and sequential equilibrium. Section 3 describes what is required for a partial equilibrium construction. Section 4 contains the Theorem and Section 5 presents a proof.

[1]For an illustration of where the rectangular margin-support condition fails, let $J = \{h_1, h_3\}$ in the second example and suppose that we seek a partial equilibrium construction in which at h_3 player 3's belief puts twice the probability on action a as on action b selected at h_1. Clearly there is no way of ensuring this belief without nailing down player 2's behavior at h_2. A more interesting example appears later in this paper.

[2]Appraisals imply assessments, and the two approaches are equivalent for equilibrium definitions.

2 Definitions

Extensive-form representation

Consider finite extensive-form games of perfect recall, with n players and nature taking the role of "player 0." Assume that the probability distribution of nature's action choices is specified in the representation of the game and is common knowledge between the strategic players. For a given such extensive-form game, let us use the following notation and terminology.

Let $N \equiv \{0, 1, \ldots, n\}$, and let $N_+ \equiv N \setminus \{0\}$ denote the set of strategic players. Let H be the set of information sets, which I call *situations*. This set is partitioned into disjoint sets $H_0, H_1, \ldots, H_n$, where H_i denotes the set of situations for player i. Let $H_+ \equiv \cup_{i \in N_+} H_i$ be the set of situations for the strategic players.

Denote by S the space of pure strategy profiles, including nature's strategy. A strategy profile can be expressed as a mapping $s\colon H \to A$ satisfying $s(h) \in A(h)$ for every $h \in H$, where A is the action space and $A(h)$ is the set of feasible actions at situation $h \in H$. The strategy profile can equivalently be represented as a tuple $s = (s_h)_{h \in H}$, where for every $h \in H$, s_h is the action chosen in situation h.

Let ΔS denote the space of probability distributions over S, that is, mixtures of strategy profiles. For any subset $T \subset S$, let us take "ΔT" to mean the subset of ΔS with support in T. Let $u\colon S \to \mathbb{R}^n$ be the payoff function, extended to the space of mixed strategies by the usual expected payoff calculation, assuming von Neumann-Morgenstern utility.

Strategy components and distributions

For any subset of situations $L \subset H$, let s_L denote the restriction of s to the subdomain L. In other words, s_L gives the profile of actions that strategy s specifies for the situations in L. For any $L \subset H$, define $-L \equiv H \setminus L$. Note that we can then write $s = (s_L, s_{-L})$. For any $X \subset S$, define $X_L \equiv \{s_L \mid s \in X\}$.

In the case of $L = \{h\}$ for a single $h \in H$, we simplify notation by dropping the brackets; so, for instance, we write X_h and s_h instead of $X_{\{h\}}$ and $s_{\{h\}}$. Note that S_h is the set of actions available at situation h. Also, for a given player i, the subscript "i" refers to the situations H_i. For example, s_i means the same thing as s_{H_i}. Likewise, "$-i$" refers to H_{-i}. Thus, subscripts "i" and "$-i$" have their usual meaning of identifying the strategies of player i and the other players, including nature.

Definition 1: *Given $L \subset H$ and $X \subset S$, say that* ***X is an L-product set*** *if $X = X_L \times X_{-L}$.*

Note that, trivially, S is an h-product set for every $h \in H$. The next definition identifies whether a mixture of strategy profiles treats a specific set of situations $L \subset H$ independently of the rest, meaning that it can be expressed as the product of the marginal distribution on L and the marginal distribution on $-L$.

Definition 2: *Given $L \subset H$ and any distribution $\sigma \in \Delta S$, say that* ***σ exhibits L-independence*** *if, for every L-product set $X \subset S$, we have $\sigma(X) = \sigma_L(X_L) \cdot \sigma_{-L}(X_{-L})$. Say that* ***σ is a behavior strategy profile*** *if it exhibits h-independence for every $h \in H$, and nature's component σ_0 is as given in the representation.*

Assumption 1: Nature's strategy σ_0 is a behavior strategy with full support.

Next let us recall the definitions of conditional probability on strictly positive-probability events and marginal distribution. We can constrain attention to finite state spaces.

Definition 3: *Given a finite set W, a probability distribution $\pi \in \Delta W$, and an event $Z \subset W$ for which $\pi(Z) > 0$, the* ***distribution conditional on Z****, denoted by $p(\cdot|Z)$, is defined by $p(X|Z) \equiv p(X \cap Z)/p(Z)$ for every $X \subset W$.*

Definition 4: *Given a distribution $p \in \Delta S$ and a set $L \subset H$, the* ***marginal distribution on dimension L****, denoted by p_L, is defined by $p_L(X_L) \equiv p(X_L \times S_{-L})$ for every $X_L \subset S_L$.*

For expressing beliefs in equilibrium definitions, it is convenient to represent information sets as subsets of strategy profiles. For each $h \in H$ and $s \in S$, let us say that s *reaches* h if the path of strategy profile s includes a node in h. Denote by $S(h)$ the set of strategy profiles that reach h. Note that, for any $L \subset H$, $S(h)_L$ is the set of action profiles for the situations in L that are consistent with h being reached.[3]

Because of perfect recall, the situations for an individual player have a particular product structure and precedence relation. For every pair of situations $h, h' \in H_i$ for a given player i, it is the case that $S(h)$ is a product set relative to h'. Further, for $h, h' \in H_i$ with $h \neq h'$, either h *is a successor of* h', in which case $S(h) \subset S(h')$; or h *is a predecessor of* h', in which case $S(h') \subset S(h)$; or neither, in which case $S(h) \cap S(h') = \emptyset$. If h' is a successor of h then every path through h' also passes through h.

Beliefs

It is convenient to express the players' beliefs at their various situations as probability distributions over the strategy space, called *appraisals* (Watson 2019). To model belief updating, we think of the players as having artificial situations that refer to "before the game begins." To be precise, for each player $i \in N_+$, we define the *initial situation* $\underline{h}_i$ to have $S(\underline{h}_i) = S$, and we label these differently across players so as not to duplicate the situations in the game, H. For each player i, let $\underline{H}_i \equiv H_i \cup \{\underline{h}_i\}$ denote the extended set of situations. Let $\underline{H}_+ \equiv \cup_{i \in N_+} \underline{H}_i$.

Definition 5: *For any strategic player i and $h \in \underline{H}_i$, call a distribution $p^h \in \Delta S$* ***an appraisal at h*** *if $p^h \in \Delta S(h)$ and if p^h exhibits independence relative to every $h' \in H_i$. An* ***appraisal system*** *is a collection of appraisals, one for each situation of the strategic players, written $P = \left(p^h\right)_{h \in \underline{H}_+}$.*

Note that an appraisal contains two things: The marginal on S_i gives player i's own strategy and the marginal on S_{-i} gives player i's belief about the strategy profile of the other

[3]Expressing extensive-form information sets as subsets of strategy profiles is not unusual, and it is something I am trying to promote. Mailath, Samuelson, and Swinkels (1993) formulate solution concepts on the basis of "normal form situations," where there is no reference to an extensive form, and Shimoji and Watson (1998) take such "restrictions" as given (whether or not they are derived from extensive-form situations). Note that, here, I am taking the conventional approach of examining standard extensive-form information sets but simply represent them as subsets of the strategy space.

players. The condition $p^h \in \Delta S(h)$ means that the appraisal at h puts probability one on reaching h. The independence condition means that, at any situation $h \in \underline{H}_+$, player i views his own strategy as independent of the other players' strategy profile, and player i's strategy is represented as a behavior strategy.

Full consistency

Definition 6: *Call a behavior strategy profile σ **fully mixed** if $\sigma(s) > 0$ for every $s \in S$.*

A fully mixed behavior strategy reaches every situation with strictly positive probability. Thus, for every $h \in H$ we have $\sigma(S(h)) > 0$ and so the conditional probability formula applies. Then for every $X \subset S(h)$, the probability of X conditional on h is $p(X|S(h))$. Full consistency requires that the players' beliefs at every situation are derived in this way, using a fully mixed strategy profile that is arbitrarily close to the strategy profile that the players are supposed to be using in the game.

Definition 7: *Take as given any finite set W, a subset $Z \subset W$, a sequence of probability distributions $\{\pi^k\} \subset \Delta W$, and a distribution $\pi \in \Delta W$. Say that **$\{\pi^k\}$ induces π conditional on Z** if $\pi^k(Z) > 0$ for every k and $\{\pi^k(\cdot|Z)\}$ converges to π.*

Definition 8: *Say that **an appraisal system P is fully consistent** if there is a behavior strategy σ and a sequence of fully mixed behavior strategies $\{\sigma^k\}$ that converges to σ and induces p^h conditional on $S(h)$, for every $h \in \underline{H}_+$.*

Note that since $S(\underline{h}_i) = S$, this implies that $p^{\underline{h}_i} = \sigma$ for every $i \in N_+$. The next definition means that, at the beginning of the game, the players' appraisals coincide on a given strategy profile.

Definition 9: *Consider a behavior strategy profile σ. Say that an **appraisal system $P = \left(p^h\right)_{h \in \underline{H}_+}$ conforms to σ** if, for every player $i \in N_+$, $p^{\underline{h}_i} = \sigma$.*

Sequential best response

Next consider the test of whether an appraisal at h specifies rational behavior for the player on the move, meaning that the actions given positive probability at situation h maximize the player's expected payoff.

Definition 10: *For a given situation $h \in H_+$ and two appraisals p^h and $\hat{p}^h$, say that **$\hat{p}^h$ is an h-deviation from p^h** if $p^h_{-h} = \hat{p}^h_{-h}$ (they are identical on the other situations).*

Definition 11: *For a strategic player i and a situation $h \in H_i$, say that an appraisal **p^h is rational at h** if $u_i(p^h) \geq u_i(\hat{p}^h)$ for every h-deviation $\hat{p}^h$. Say that an **appraisal system $P = \left(p^h\right)_{h \in \underline{H}_+}$ is sequentially rational** if p^h is rational at h, for every $h \in H_+$.*

Here sequential rationality is defined in terms of "one-shot deviations," meaning that we evaluate player i's rationality at a given situation $h \in H_i$ by looking just at alternative choices at h rather than alternatives that would also adjust player i's behavior at other situations that may be reached in the continuation of the game.[4] The familiar *one-deviation principle*—equivalence between single-deviation optimality and strategy-deviation optimality—holds assuming that the appraisal system is "minimally consistent," which is implied by the full consistency assumption of sequential equilibrium, defined next.

Sequential Equilibrium

With the essential ingredients in place, here is the definition of sequential equilibrium (Kreps and Wilson 1982), expressed using the terminology of appraisals:

Definition 12: *Taking nature's strategy σ_0 as given, say that* ***an appraisal system P is a sequential equilibrium*** *if there is a behavior strategy profile σ such that P conforms to σ, P is fully consistent, and P is sequentially rational. In this case we say that* ***behavior strategy profile σ is a sequential-equilibrium strategy profile****.*

3 Partial Equilibrium Construction

Let us consider a partial equilibrium construction that focuses on a given set of situations $J \subset H$ assumed to contain nature's situations H_0. For each strategic player i, let $\underline{J}_i \equiv (J \cap H_i) \cup \{\underline{h}_i\}$, and let $\underline{J}_+ \equiv \cup_{i \in N_+} \underline{J}_i$. Let ϕ denote a mixture over ΔS_J, in other words, a strategy for only the situations in J. Call ϕ a *J-behavior strategy* if it exhibits h-independence for every $h \in J$, and nature's component σ_0 is as given in the representation. Independence is defined as before, applied here on the reduced domain of situations in J.

I next describe behavior, partial beliefs, and rationality at the situations in J. These definitions are straightforward variants of those in the previous section. Let us begin with a partial appraisal system, which gives the beliefs at each situation in $\underline{J}_+$ about the behavior at the situations in J (and not at other situations).

Definition 13: *For $h \in \underline{J}_+$, a distribution $q^h \in \Delta S_J$ is called a* ***J-appraisal at h*** *if $q^h \in \Delta S(h)_J$ and if q^h exhibits independence relative to every $h' \in H_i \cap J$. A* ***J-partial appraisal system*** *is given by $Q = \left(q^h\right)_{h \in \underline{J}_+}$.*

Definition 14: *Call a J-behavior strategy* ***fully mixed*** *if $\phi(s_J) > 0$ for every $s_J \in S_J$.*

Definition 15: *Say that* **J-partial appraisal system Q is** ***J-fully consistent*** *if there is a J-behavior strategy ϕ and a sequence of fully mixed J-behavior strategies $\{\phi^k\}$ that converges to ϕ and induces q^h conditional on $S(h)_J$, for every $h \in \underline{J}_+$.*

[4]Note that since all strategy profiles in the support of p^h and $\hat{p}^h$ reach h, the expected payoffs shown in the rationality definition are conditional on reaching situation h.

Because the partial equilibrium construction will not address the behavior at situations outside of J, the sequential-rationality condition for the analysis of situations in J requires robustness to any behavior at the other situations.

Definition 16: *For a strategic player i and a situation $h \in J_i$, say that **J-appraisal q^h is J-rational at h** if $u_i(p^h) \geq u_i(\hat{p}^h)$ for every appraisal p^h at h satisfying $p^h_J = q^h$ and for every h-deviation $\hat{p}^h$. Say that **J-appraisal system Q is J-sequentially rational** if q^h is rational at h, for every $h \in J_+$.*

Definition 17: *For any J-behavior strategy ϕ, say that **J-appraisal system Q conforms to ϕ** if, for every player $i \in N_+$, $q^{\underline{h}_i} = \phi$.*

The components above are next used to define partial sequential equilibrium.

Definition 18: *Taking nature's strategy σ_0 as given, say that **a J-appraisal system Q is a J-partial sequential equilibrium** if there is a J-behavior strategy profile ϕ such that Q conforms to ϕ, Q is fully consistent, and Q is J-sequentially rational.*

4 Theorem

The question to be addressed is whether any given J-partial sequential equilibrium can be extended to construct a sequential equilibrium that coincides on J. Coincides means the following two properties: First, the actions specified for the situations in J are the same as in the J-partial sequential equilibrium. Second, for each situation in J, the marginal belief on dimension J is the same as the belief in the J-partial sequential equilibrium. That is, given J and a J-sequential equilibrium Q, we would need to find a sequential equilibrium P such that $p^h_J = q^h$ for every $h \in \underline{J}_+$.

The requirement of J-sequential rationality does not get in the way, because it is defined to be robust to the behavior at the situations in $-J$. But there is no guarantee that a J-partial sequential equilibrium can be extended, because generally the appraisals at situations in J must depend on the distribution of actions taken at situations in $-J$, for whatever sequence of fully mixed behavior strategies is needed for full consistency in the entire game. In other words, for a given situation $h \in J$, p^h depends on the probabilities that $\{\sigma^k\}$ puts on actions at situations in $H \setminus J$.

The main creative step here is to find a condition on the J-partial appraisal system Q guaranteeing that appraisals to be constructed for situations in J will effectively treat the behavior in J as *independent* of the behavior in $H \setminus J$, regardless of what is specified on the latter dimension. Consider the following condition on the situations for a given set $J \subset H$ and any given J-partial appraisal system Q.

Rectangular margin-support condition: $\{s \in S(h) \mid s_J \in \operatorname{supp} q^h\}$ is a J-product set, for every $h \in J_+$.

Note that $\{s \in S(h) \mid s_J \in \operatorname{supp} q^h\}$ contains the strategy profiles that reach h and whose J components are in the support of q^h, the given belief at h regarding behavior in the partial

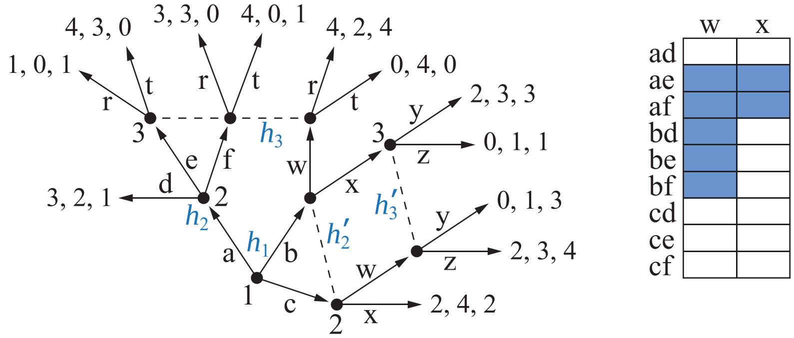

Figure 2: First example.

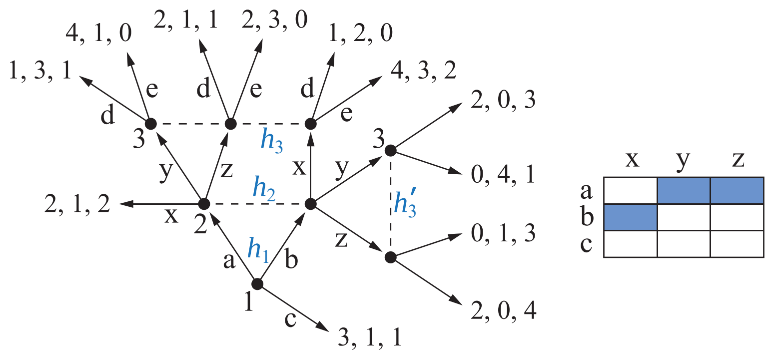

Figure 3: Second example.

equilibrium construction. A sufficient condition for $\{s \in S(h) \mid s_J \in \operatorname{supp} q^h\}$ being a J-product set is that $(\operatorname{supp} q^h) \times S(h)_{-J} \subset S(h)$.

Figures 2 and 3 provide a graphical depiction of the rectangular margin-support condition in the two examples discussed in the Introduction. In Figure 2, the table at the right represents the possible combinations of action profiles taken at situations h_1, h_2, and h_2'. These are arranged in terms of dimensions J and $-J$ that we set earlier, with the profiles of actions at h_1 and h_2 (the ones in J) as the rows and the actions at h_2' (the situation in $-J$) as columns. The profiles shaded in blue are the combinations that reach h_3. That is, at situation h_3, player 3 has observed that one of the shaded profiles has been played. Figure 3 has a similar table, giving the information that player 3 has at h_3 about the actions taken at h_1 and h_2.

Note that the shaded region in Figure 2 is not a product set; it is not the product of $\{\mathrm{ae}, \mathrm{af}, \mathrm{bd}, \mathrm{be}, \mathrm{bf}\}$ and $\{\mathrm{w}, \mathrm{z}\}$. Even with a fully mixed and independent initial belief (as would be the case in the perturbation along the sequence for a sequential equilibrium construction), player 3's belief conditional on the set of profiles that reach h_3 would not necessarily exhibit independence, and the constructed behavior regarding $-J$ would affect the updated belief about J. But if the partial construction makes the conditional probability about J converge to something with support in $\{\mathrm{ae}, \mathrm{af}\}$, as was the case discussed in the Introduction, then this distribution is preserved as the marginal distribution in the full construction.

Likewise, the shaded region in Figure 3 is not a product set. However, if the partial construction makes the conditional probability about J converge to probability 1 on $\{\mathrm{a}\}$, as was true in the partial construction given in the Introduction, then this can be preserved as the marginal distribution in the full construction, so long as we apply a suitable bound on action probabilities in the full construction.

It turns out that such a construction is possible in general, assuming that the subset of action profiles that reach the given situation, when restricted to the support of q^h on the J dimension, is a product set. This is the rectangular margin-support condition. The main result follows.

Theorem: *Take an given a finite extensive-form game, any set of situations $J \subset H$ that contains H_0, and a J-appraisal system Q. If the rectangular margin-support condition is satisfied and Q is a J-partial sequential equilibrium, then there exists a sequential equilibrium P with the property that $p^h_J = q^h$ for every $h \in \underline{J}_+$.*

For an example in which the rectangular margin-support condition fails and the Theorem does not apply, consider the game shown in Figure 4. Let the situations of interest be $J = \{h_1, h_2, h_2'\}$. Define a sequence of fully mixed J-behavior strategies $\{\phi^k\}$ so that the probabilities of actions a and b at h_1 are both $1/k$, the probabilities of actions d and e at h_2 are both $1/k$, and the probabilities of actions w and z at h_2' are both $1/2$. It is clear that this sequence supports the J-fully consistent appraisal system Q given as follows: q^{h_1} puts probability $1/2$ on both partial strategy profiles cfw and cfz; and both q^{h_2} and $q^{h_2'}$ put probability $1/4$ on each of the partial strategy profiles afw, afz, bfw, and bfz. It is easy to confirm that Q is J-sequentially rational. Note that for the rationality test at h_2', the only appraisal $p^{h_2'}$ satisfying $p^{h_2'}_J = q^{h_2'}$ is that which has player 3 choosing x and y with equal probability at h_3.

We have shown that Q is a J-partial sequential equilibrium. However, the rectangular margin-support condition fails at h_2', as can be seen with a quick look at the table in Figure 4, where the shaded cells are the combinations of actions at h_1 and h_3 that reach h_2'. Because $q^{h_2'}$ puts probability $1/4$ on each of the partial strategy profiles afw, afz, bfw, and bfz, $\{s \in S(h_2') \mid s_J \in \operatorname{supp} q^{h_2'}\}$ is *not* a J-product set. There is no sequential equilibrium P with the property that $p^h_J = q^h$ for every $h \in \underline{J}_+$. In fact, the only sequential equilibrium strategy profiles are aexz and bfxz.

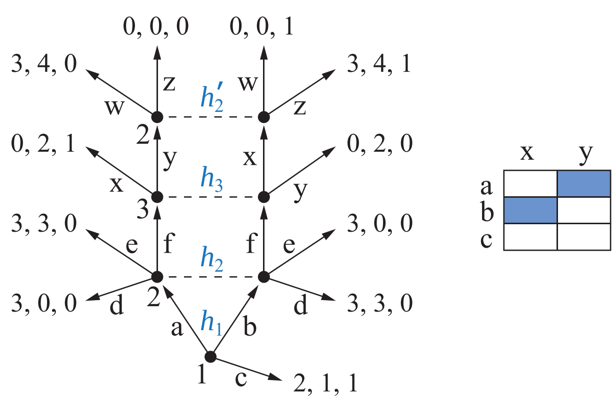

Figure 4: Third example.

5 Proof

This section contains a proof of the Theorem, starting with two component lemmas. To avoid confusion, note that throughout the analysis below, $-J \equiv H \backslash J$ (no initial situations included).

Lemma 1: *Take as given a sequence of fully mixed behavior strategies $\{\sigma^k\}$, a set $Z \subset S$, a distribution $\pi \in \Delta S$, and a set $Y \subset Z$. If $\{\sigma^k\}$ induces π conditional on Z, and if $\operatorname{supp} \pi \subset Y$, then $\{\sigma^k\}$ induces π conditional on Y.*

Proof. The claim is derived by the following equalities, which hold for any $X \subset S$:

$$\pi(X) = \lim_{k\to\infty} \sigma^k(X|Z) = \lim_{k\to\infty} \frac{\sigma^k(X \cap Z)}{\sigma^k(Z)} = \lim_{k\to\infty} \frac{\sigma^k(X \cap Z)}{\sigma^k(Y)} \cdot \frac{\sigma^k(Y)}{\sigma^k(Z)}$$

$$= \lim_{k\to\infty} \frac{\sigma^k(X \cap Z)}{\sigma^k(Y)} \cdot \lim_{k\to\infty} \frac{\sigma^k(Y)}{\sigma^k(Z)} = \lim_{k\to\infty} \frac{\sigma^k(X \cap Z)}{\sigma^k(Y)} \cdot \lim_{k\to\infty} \sigma^k(Y|Z) = \lim_{k\to\infty} \frac{\sigma^k(X \cap Z)}{\sigma^k(Y)} \cdot \pi(Y)$$

$$= \lim_{k\to\infty} \frac{\sigma^k(X \cap Z)}{\sigma^k(Y)} = \lim_{k\to\infty} \frac{\sigma^k(X \cap Y) + \sigma^k(X \cap Z \setminus Y)}{\sigma^k(Y)}$$

$$= \lim_{k\to\infty} \frac{\sigma^k(X \cap Y)}{\sigma^k(Y)} + \lim_{k\to\infty} \frac{\sigma^k(X \cap Z \setminus Y)}{\sigma^k(Y)} = \lim_{k\to\infty} \frac{\sigma^k(X \cap Y)}{\sigma^k(Y)} = \lim_{k\to\infty} \sigma^k(X|Y).$$

The first line is from the definition of $\{\sigma^k\}$ inducing π conditional on Z, the definition of conditional probability, that $Y \subset Z$, and because $\sigma^k(Y) > 0$ due to σ^k being fully mixed and $Y \neq \emptyset$ (which is implied by $\operatorname{supp} \pi \subset Y$). The equalities on the second line hold because the latter limit exists due to $\pi(Y)$ being well defined. The first equality of the third line follows from $\pi(Y) = 1$, since $\operatorname{supp} \pi \subset Y$. The equalities on the fourth line are from the definitions of conditional probability and "induce," and use the fact that $\sigma^k(Y)$ converges to 1 (due to $\operatorname{supp} \pi \subset Y$). So we have $\lim_{k\to\infty} \sigma^k(X|Y) = \pi(X)$ for every $X \subset S$, and therefore $\{\sigma^k\}$ induces π conditional on Y. □

Lemma 2: *Take as given a sequence of fully mixed behavior strategies $\{\sigma^k\}$, a set $J \subset H$, a J-product set $Y \subset S$, and a distribution $\pi \in \Delta S$. Let $\lambda^k \equiv \sigma^k_J$ denote the marginal of σ^k on dimension J and assume that $\lambda^k(\cdot|Y_J)$ converges to some distribution $\psi \in \Delta S_J$. If $\{\sigma^k\}$ induces π conditional on Y, then π exhibits J-independence and $\pi_J = \psi$.*

Proof. For any J-product set $X \subset S$, the following equalities hold:

$$\pi(X) = \lim_{k\to\infty} \sigma^k(X|Y) = \lim_{k\to\infty} \frac{\sigma^k(X \cap Y)}{\sigma^k(Y)} = \lim_{k\to\infty} \frac{\sigma^k_J(X_J \cap Y_J) \cdot \sigma^k_{-J}(X_{-J} \cap Y_{-J})}{\sigma^k_J(Y_J) \cdot \sigma^k_{-J}(Y_{-J})}.$$

The third equality is due to σ^k being a behavior strategy and X and Y being J-product sets, implying that $(X \cap Y)_J = X_J \cap Y_J$ and $(X \cap Y)_{-J} = X_{-J} \cap Y_{-J}$. Because $\{\sigma^k(\cdot|Y)\}$ and $\{\lambda^k(\cdot|Y_J)\}$ converge, we have

$$\lim_{k\to\infty} \frac{\sigma^k_J(X_J \cap Y_J) \cdot \sigma^k_{-J}(X_{-J} \cap Y_{-J})}{\sigma^k_J(Y_J) \cdot \sigma^k_{-J}(Y_{-J})} = \lim_{k\to\infty} \frac{\sigma^k_J(X_J \cap Y_J)}{\sigma^k_J(Y_J)} \cdot \lim_{k\to\infty} \frac{\sigma^k_{-J}(X_{-J} \cap Y_{-J})}{\sigma^k_{-J}(Y_{-J})}$$

$$= \pi_J(X_J \cap Y_J) \cdot \pi_{-J}(X_{-J} \cap Y_{-J}) = \pi_J(X_J) \cdot \pi_{-J}(X_{-J}).$$

The last equality is due to $\pi(Y) = 1$, which follows from $\sigma^k(Y|Y) = 1$ and the definition of $\{\sigma^k\}$ inducing π conditional on Y. That $p_J = \psi$ is obvious by evaluating $\pi(X_J \times S_{-L})$. $\square$

To prove the Theorem, let $J \subset H$ be a set of situations that contains H_0 and let Q be a J-partial sequential equilibrium such that the rectangular margin-support condition is satisfied. Let ϕ be a J-behavior strategy and let $\{\phi^k\}$ be a sequence of fully mixed J-behavior strategies that converges to ϕ and induces q^h conditional on $S(h)_J$, for every $h \in \underline{J}_+$. That Q conforms to ϕ is an implication. For every $h \in \underline{J}_+$, define

$$Y^h \equiv \{s \in S(h) \mid s_J \in \operatorname{supp} q^h\}.$$

We know that Y^h is a J-product set from the rectangular margin-support condition.

Note that for every $h \in \underline{J}_+$, $\phi^k(Y^h_J | S(h)_J)$ converges to 1, since $\{\phi^k\}$ induces q^h conditional on $S(h)_J$ and $Y^h_J = \operatorname{supp} q^h$. This implies that $\phi^k(S(h)_J \backslash Y^h_J)/\phi^k(Y^h_J)$ converges to 0. We can presume that $\{\phi^k\}$ is defined so that $\phi^k(S(h)_J \backslash Y^h_J)/\phi^k(Y^h_J) \in (0, 1/|A|^2)$ for every positive integer k and every $h \in \underline{J}_+$, because this condition must hold for k large enough anyway. (Recall that A is the action space of the game, and so $|A|$ is the number of different actions.) Further, define sequence $\{\delta^k\}$ by:

$$\delta^k \equiv \max_{h \in \underline{J}_+} \left(\frac{\phi^k(S(h)_J \backslash Y^h_J)}{\phi^k(Y^h_J)} \right)^{\frac{1}{2}}.$$

We have that $\delta^k \in (0, 1/|A|)$

Consider any positive integer k. Note that we can define an artificial game in the Bayesian normal form as follows. Nature makes the choices at all situations in J, and nature's strategy is ϕ^k. This means that at each situation $h \in J$, nature takes the action according to distribution ϕ^k_h, which is independent of the choices made at the other situations due to ϕ^k being a behavior strategy. In the artificial game, strategic player $i \in N_+$ makes the choices at situations in $H_i \backslash J$, and let us assume that this player is restricted to put probability of at least δ^k on each available action. Standard arguments and use of Kakutani's fixed-point theorem imply that the artificial game has a Nash equilibrium in fully mixed behavior strategies, θ^k. Let $\sigma^k \in \Delta S$ be defined by $\sigma^k(s) \equiv \phi^k(s_J) \cdot \theta^k(s_{-J})$. Note that σ^k is a fully mixed behavior strategy.

Because the game is finite and the space of feasible behavior strategies in artificial game k converges to the space of behavior strategies in the given game, we can find a subsequence $\{\sigma^{k^m}\}$ and a behavior strategy σ such that $\{\sigma^{k^m}\}$ converges to σ and $\{\sigma^{k^m}(\cdot|S(h))\}$ converges for every $h \in H$. Letting $p^h \equiv \lim_{m\to\infty} \sigma^{k^m}(\cdot|S(h))$ for every $h \in H$, we have thus constructed an appraisal system P that is fully consistent and conforms to σ.

Because σ^{k^m} is fully mixed, every situation is reached in the artificial game, implying that θ^{k^m} is rational at every $h \in -J$. By continuity of payoffs in the action probabilities and because σ^{k^m} converges to σ, it must be that p^h is rational at every $h \in -J$. Regarding situations in J, we need to establish rationality and that the appraisals at the J margin agree with Q.

Take any $h \in J$. The rectangular margin-support condition implies that $(S(h) \backslash Y^h) \cap (Y^h_J \times S_{-J}) = \emptyset$, which implies that

$$\sigma^k(S(h) \backslash Y^h) \leq \sigma^k((S(h)_J \backslash Y^h_J) \times S_{-J}) \leq \phi^k(S(h)_J \backslash Y^h_J).$$

Likewise, we have $\sigma^k(Y^h) \geq \phi^k(Y_J^h) \cdot \delta^k$, because for every $s \in Y^h$, it is the case that $s_J \in Y_J^h$ and $\sigma^k(s) = \phi^k(s_J) \cdot \theta^k(s_{-J}) \geq \phi^k(s_J) \cdot \delta^k$. Combining these inequalities and recalling that $\delta^k \geq (\phi^k(S(h)_J \backslash Y_J^h)/\phi^k(Y_J^h))^{1/2}$, we obtain

$$\frac{\sigma^k(S(h)\backslash Y^h)}{\sigma^k(Y^h)} \leq \frac{\phi^k(S(h)_J \backslash Y_J^h)}{\phi^k(Y_J^h) \cdot \delta^k}$$

$$\leq \frac{\phi^k(S(h)_J \backslash Y_J^h)}{\phi^k(Y_J^h)} \cdot \left(\frac{\phi^k(Y_J^h)}{\phi^k(S(h)_J \backslash Y_J^h)}\right)^{\frac{1}{2}} = \left(\frac{\phi^k(S(h)_J \backslash Y_J^h)}{\phi^k(Y_J^h)}\right)^{\frac{1}{2}}.$$

The last expression on the right side converges to 0 as $k \to \infty$, because $\phi^k(S(h)_J \backslash Y_J^h)/\phi^k(Y_J^h)$ converges to 0. Therefore $\sigma^k(S(h)\backslash Y^h)/\sigma^k(Y^h)$ converges to 0. Recall that $\{\sigma^{k^m}\}$ induces p^h conditional on $S(h)$, so that $p^h(Y^h) = \lim_{m \to \infty} \sigma^{k^m}(Y^h | S(h))$. Using the definition of conditional probability, this implies that $\operatorname{supp} p^h \subset Y^h$.

Writing $Z = S(h)$ and $\pi = p^h$, Lemma 1 then implies that $\{\sigma^{k^m}\}$ induces p^h conditional on Y. Further, noting that ϕ^k is the marginal of σ^k on dimension J, and writing $\psi = q^h$, we can apply Lemma 2 to establish that p^h exhibits J-independence and $p_J^h = q^h$. This combined with the fact that Q is J-sequentially rational implies that P is rational at every $h \in J$.

In summary, we have shown that P is fully consistent, P sequentially rational, and P conforms to σ. Therefore P is a sequential equilibrium. Further, we have shown that $p_J^h = q^h$ for every $h \in J$, completing the proof.

References

Battigalli, Pierpaolo (1996): "Strategic Independence and Perfect Bayesian Equilibria," *Journal of Economic Theory* 70: 201–234.

Battigalli, Pierpaolo (1997): "On Rationalizability in Extensive Form Games," *Journal of Economic Theory* 74: 40–61.

Govindan, Srihari and Robert Wilson (2009): "On Forward Induction," *Econometrica* 77 (1, January): 1–28.

Kreps, D. and R. Wilson (1982): "Sequential Equilibrium," *Econometrica* 50 (4): 863–894.

Pearce, David G. (1984), "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica* 52 (July): 1029–1050.

Watson, J. (2017): "A General, Practicable Definition of Perfect Bayesian Equilibrium," unpublished draft, UC San Diego.

Watson, J. (2023): "Contractual Chains," unpublished manuscript, UC San Diego.