

UC San Diego

UC San Diego Previously Published Works

Title

Partially constructed sequential equilibrium

Permalink

<https://escholarship.org/uc/item/5nh3x3kf>

Journal

Journal of Economic Theory, 231

ISSN

0022-0531

Author

Watson, Joel

Publication Date

2026

DOI

10.1016/j.jet.2025.106111

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-ShareAlike License, available at <https://creativecommons.org/licenses/by-nc-sa/4.0/>

Peer reviewed

Partially Constructed Sequential Equilibrium

Joel Watson*

30 November 2025

Final version before publication,
Journal of Economic Theory 231 (January 2026) 106111

Abstract

This article 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. Illustrative applications are also presented.

1 Introduction

For some applications of noncooperative game theory, analysts construct a sequential equilibrium (Kreps and Wilson 1982), requiring *full consistency* of beliefs across information sets. For other applications, analysts prefer a solution concept with a weaker consistency requirement but may still find it useful to construct a sequential equilibrium (if the game is finite) to establish existence. Constructing a sequential equilibrium for a complex game can be challenging; one must describe a sequence of fully mixed behavior-strategy profiles that converges to the equilibrium strategy profile and that determines the beliefs at all information sets, as the limits of the corresponding sequences of conditional probabilities. This paper identifies 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.

To get a sense of what is required for a partial construction, consider the extensive-form representation shown in Figure 1. Suppose we want to specify behavior and beliefs for only the 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 $-J = \{h'_2, h'_3\}$.

Specifying a partial strategy profile is straightforward; it is the J component of a behavior-strategy profile. For example, let us prescribe action α at h_1 , action g at h_2 , and mixing

*UC San Diego; <http://econ.ucsd.edu/~jwatson/>. Thanks go to Kelvin Leong for helpful research assistance, to seminar participants at UC San Diego and Washington University in St. Louis for their feedback, to Editor Battigalli and referees for their detailed comments, and to Silvia Amalia Meneghesso for helpful comments and careful proofreading. First public version of the paper: 27 November 2023.




Figure 1: First example.

between actions r and t with equal probability at h_3 . Letting $S_J = \{\alpha, \beta, \gamma\} \times \{g, e, f\} \times \{r, t\}$ denote the J component of the space of strategy profiles, we can express this partial behavior-strategy profile as a probability distribution on S_J , namely αgr and αgt each assigned probability $1/2$. For reference below, let us denote this distribution by ϕ .

Specifying partial beliefs means providing, for each information set in J , the active player's belief about behavior at (and only at) the information sets in J .¹ These beliefs can be written as a probability distribution on S_J . For example, suppose that at h_1 player 1 believes that players 2 and 3 would behave as the specified partial strategy profile prescribes for h_2 and h_3 . With player 1 planning to choose α , this belief puts equal probability on αgr and αgt . Suppose that at h_2 player 2 has the same belief. Finally, suppose that at h_3 player 3 thinks player 1 chose action α for sure and that player 2's action is equally likely to have been e and f . With player 3 planning to mix between r and t , player 3's belief therefore puts equal probability on αer , αet , αfr , and αft . Let us denote these partial beliefs, collectively, by Q .

I explain below why it is helpful to express partial behavior-strategy profiles and partial beliefs both as probability distributions on S_J , rather than in the form of an *assessment* (Kreps and Wilson 1982) that combines a behavior-strategy profile with probability distributions over the nodes in each information set. But first observe that the partial strategy profile and partial beliefs specified above satisfy the conditions of sequential equilibrium in a fashion that is robust to whatever could be specified for information sets h_2' and h_3' .

The partial strategy profile ϕ clearly is sequentially rational at h_1 , h_2 , and h_3 , given the partial beliefs we have described. At h_1 , action α leads to the payoff 3 for player 1, whereas the expected payoffs of actions β and γ cannot exceed 2, regardless of the play at h_2' and h_3' . For instance, if player 1 deviates to take action β and player 2 happens to choose w , then player 1's expected payoff is 2 because player 1 believes that, at h_3 , player 3 randomizes between r and t with equal probability. Likewise, player 2's action g 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 makes player 2's action at h_2' irrelevant.

Additionally, the partial beliefs Q are fully consistent in the following sense: There is

¹Here "belief" incorporates the active player's own plans and past actions, for whichever of her own information sets are in J . This is a more encompassing use of the term "belief" than that adopted by Kreps and Wilson (1982) and much of the literature. The value of this approach will be demonstrated.

a sequence of fully mixed, partial behavior-strategy profiles that converges to ϕ and has the following property: For each information set $h \in J$, the specified partial belief at h is the limit of the sequence of distributions conditional on the subset of S_J that reaches h . A partial strategy profile $s_J \in S_J$ is said to reach h if there exists $s_{-J} \in S_{-J} = \{w, x\} \times \{y, z\}$ such that the path induced by strategy profile (s_J, s_{-J}) reaches h .

A sequence that works is defined as follows, for integer index k : At h_1 , the probabilities of actions α , β , and γ are, in order, $1 - 2(1/k)^2$, $(1/k)^2$, and $(1/k)^2$. At h_2 , actions g , e , and f are chosen with probabilities $1 - 2(1/k)$, $1/k$, and $1/k$. And at h_3 , the probabilities of actions r and t are each $1/2$. Note that the set of partial strategy profiles that reach h_3 is $S(h_3)_J \equiv \{\alpha e, \alpha f, \beta g, \beta e, \beta f\} \times \{r, t\}$. For any k , the distribution on S_J conditional on $S(h_3)_J$ puts the following probability on each of $\alpha e r$, $\alpha e t$, $\alpha f r$, and $\alpha f t$:

$$\frac{\left(1 - 2\left(\frac{1}{k}\right)^2\right) \cdot \frac{1}{k} \cdot \frac{1}{2}}{\left(1 - 2\left(\frac{1}{k}\right)^2\right) \cdot 2\left(\frac{1}{k}\right) + \left(\frac{1}{k}\right)^2}.$$

As k approaches ∞ , this number converges to $1/4$, which matches the partial belief we specified for player 3 at h_3 .

Hence, it seems that we have the makings of a sequential equilibrium. Indeed, for this example, there is a way to extend 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-strategy profiles that satisfy the requirements of sequential equilibrium. Further, on the J margin, the resulting strategy profile and beliefs coincide with the partial strategy profile and beliefs that we specified above.² In other words, our partial sequential-equilibrium construction can be extended to a completely specified sequential equilibrium.

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 completely specified sequential equilibrium. Two conditions must be satisfied: First, sequential rationality at the information sets in J must be robust to any behavior at the information sets in $-J$. Second, a property that I call the *rectangular margin-support condition* must hold for the partial beliefs. The second condition is subtle and relates to whether beliefs about actions taken at information sets in J could be influenced by the specification of behavior at information sets in $-J$. For instance, in the example above, one path to information set h_3 goes through information set h'_2 and requires action w . Therefore, in any completely specified sequential equilibrium, player 3's belief at h_3 depends on the probability of w relative to the probabilities of α , β , e , and f .

The rectangular margin-support condition states that, at every $h \in J$, if we restrict attention to the support of the partial belief on S_J , then the set of strategy profiles that reach h is a product set with respect to J and $-J$ (that is, a *conjunctive event*, Watson 2025). This condition is sufficient to find a sequence of fully mixed partial behavior-strategy profiles for

²For instance, supplementing the probabilities described for the sequence of partial behavior-strategy profiles, let the probability of x at h'_2 be $1 - (1/k)$, and let the probability of y at h'_3 be $1 - (1/k)$. The resulting sequence of fully mixed behavior-strategy profiles supports a sequential equilibrium in which player 1 plays strategy α , player 2 plays gx , and player 3 plays ry and ty with equal probability.

component $-J$ that, when combined with the given sequence for component J , maintains the limit of conditional probabilities on the J -margin. Additional elements of the proof relate to the beliefs and sequential rationality at information sets in $-J$.

To conduct the general analysis in a straightforward manner, behavior-strategy profiles and beliefs are expressed as probability distributions on the space of strategy profiles, as illustrated above. The beliefs, which include the players' plans for unreached information sets, are called *appraisals* following Watson (2025). This is in contrast to expressing beliefs in the form of an assessment—that is, as probability distributions over nodes at information sets, in conjunction with a behavior-strategy profile. I next explain why the appraisal structure is preferred.

Returning to the example shown in Figure 1, consider player 3's information set h_3 and let $J \equiv \{h_1, h_2, h_3\}$ as before. For a partial equilibrium construction, we need to describe player 3's belief about the actions taken at the information sets in J , most notably h_1 and h_2 , in isolation of whatever will be player 3's belief about the action taken at h'_2 . Note that this belief, along with player 3's plan of action for h_3 , is precisely about the J component of the strategy profile. Hence, it is most directly expressed as a probability distribution on S_J . We could take steps to convert this probability distribution into an assessment, but doing so would obscure its meaning and lead to unnecessary complications.

In fact, representing player 3's belief in terms of a probability distribution on the nodes of information set h_3 , as in an assessment, would require describing a *set* of such probability distributions. For instance, suppose player 3 puts probability $1/2$ on both α and β , and she puts probability $1/2$ on both e and f . This belief corresponds to the set of probability distributions given by

$$\left\{ \left(\frac{1}{2+2\nu}, \frac{1}{2+2\nu}, \frac{\nu}{1+\nu} \right) \mid \nu \in [0, 1] \right\},$$

where the vectors give the probabilities of the nodes in h_3 , from left to right.

Carrying around such a set of probability distributions is cumbersome in comparison to handling a single probability distribution on S_J . Further, it defeats the purpose of the partial-equilibrium construction by accounting for behavior on S_{-J} , the component of the strategy profile we hoped to ignore. In the example, ν is effectively the probability player 3 puts on player 2 choosing w at h'_2 . To describe the set requires calculating a posterior distribution for every ν , again defeating the purpose of the partial-construction technique.

To reiterate, the object of analysis for the partial construction is S_J . Therefore, the most direct, straightforward method of analysis, and the one that creates methodological value, is to express partial behavior strategies and partial beliefs as probability distributions on S_J . The only reason to convert beliefs into the assessment form would be to state definitions of sequential rationality and full consistency in the manner that Kreps and Wilson (1982) did, using the same notation. But sequential rationality and full consistency are easily stated in terms of appraisals, and I show formally that the appraisal-based definitions are equivalent to their assessment counterparts for analysis of sequential equilibrium.

Further, the appraisal form appears prominently in the prior literature. Beliefs are put in terms of probability distributions on the space of strategy profiles for rationalizability theory, as in Pearce's (1984) original definition and Battigalli's (1997) restatement, and for some equilibrium analysis, such as Battigalli's (1996) exploration of independence underlying equilibrium notions, Mailath, Samuelson, and Swinkels' (1993) formulation of solution concepts




Figure 2: Second example.

on the basis of “normal-form information sets,” and Govindan and Wilson’s (2009) study of forward induction.³ The appraisal form is more general and particularly well-suited for solution concepts weaker than sequential equilibrium and for analysis of forward induction.⁴

For another illustration of the rectangular margin-support condition, consider the game in Figure 2. 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 specifies action γ played with probability $1 - (1/k) - (1/k)^2$, action α with probability $(1/k)$, action β with probability $(1/k)^2$, action r with probability $1 - (1/k)$, and action t with probability $1/k$. In line with these probabilities, we specify that at h_3 player 3 believes that action α was chosen for sure at h_1 . Then the partial strategy profile in which player 1 chooses γ and player 3 chooses r at h_3 is sequentially rational, regardless of what player 3 might believe about player 2’s action at h_2 .

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 complete construction of a sequential equilibrium. This is because, unlike in the first example, the belief at h_3 critically depends on the behavior at $h_2 \in -J$ in conjunction with the behavior already described for J . For instance, if actions y and z are both assigned probability $(1/k)^2$, then in a complete equilibrium construction (letting $k \rightarrow \infty$), player 3’s belief at h_3 must put probability 0 on action α having been chosen at h_1 . This problem arises for *any* probabilities that we could pick for actions α and β , because of how the actions at h_1 and h_2 interact to determine whether h_3 is reached.

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 α earlier. The construction is accomplished by imposing a sequence of lower bounds on the probabilities of actions at the information sets in $-J$ that goes to zero

³I refine the approach a little by incorporating the player’s own strategy in this player’s beliefs and by paring down the notation to the essentials. This structure is useful also for epistemic and psychological game theory (Battigalli and De Vito 2021, Battigalli et al. 2019).

⁴At an information set h , the active player i ’s appraisal (and what is needed to evaluate rationality) comprises both the belief about what has happened in the past to reach h and what will happen in the future, and both entail not only the behavior of the other players but also the behavior of player i herself. An assessment records the belief about the past as a probability distribution over nodes, whereas the belief about the future is assumed to be given by the single equilibrium strategy profile.

more slowly than do the probabilities on the actions at the information sets in J .⁵ 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.

While the simple examples shown in Figures 1 and 2 do well to illustrate the basic idea behind the Theorem, the primary objective is to help with more complex applications. In fact, the Theorem herein has already been used in the proof of the main theorem in Watson (2024), which examines a class of network contracting games, and to prove a result in Watson (2025) on multistage games with observed actions.

An application along the lines of Watson (2024) is presented here, where I show that it is enough to focus on a class of information sets that are essentially on the equilibrium path or entail a single deviation. A second application, on calculating sequential equilibria in general multistage games with observed actions and independent types, is also presented. I provide a novel result establishing that, to characterize the set of sequential equilibria, one can focus on a class of information sets in which there is at most one public deviation per period. I hope that the theory will give other analysts a useful tool to simplify the process of constructing a sequential equilibrium in applied settings.

The next section lays out notation and reviews definitions pertaining to the extensive-form representation, beliefs, full consistency, and sequential equilibrium. Subsection 2.7 explains how to translate between the appraisal-based and assessment-based definitions, and establishes the equivalence for analysis of sequential equilibrium. Section 3 describes what is required for a partial equilibrium construction. Section 4 contains the Theorem. Section 5 describes and provides the analysis for the two applications noted above, and Section 6 presents proofs of the technical results along with additional lemmas.

2 Basic Definitions

Consider finite games with *perfect recall*, represented in the extensive form. Where reference herein is made to components of the representation, I shall use Osborne and Rubinstein's (1994) description of the extensive form, whereby nodes are defined as action sequences. For a given such game, let N denote the set of players, including Chance (also called Nature in some contexts) as the nonstrategic player 0. Let $N_+ \equiv N \setminus \{0\}$ denote the set of strategic players. Herein, " $\subset$ " denotes (weak) subset, not proper subset; sequences are expressed as $(m^k)_k$, where $k = 1, 2, \dots$ is understood; and statements of convergence of the form $m^k \rightarrow m$ are for $k \rightarrow \infty$.

The first few subsections below review how the fundamental components of analysis—information, behavior-strategy profiles, and beliefs—can all be expressed in terms of the space of strategy profiles S . As discussed in the Introduction, this is the most economical formula-

⁵A standard existence result is used to find Nash equilibria in the space of bounded mixed strategies on domain $-J$ for each k , along with a novel step of characterizing the limit of the conditional probabilities derived from the combination of mixed partial behavior-strategy profiles on domains J and $-J$ under the rectangular margin-support condition. 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 α as on action β selected at h_1 . Clearly there is no way of ensuring this belief without nailing down player 2's behavior at h_2 .

tion for the analysis herein and is essential for separating components of the strategy profile in a partial construction of equilibrium. The subsections after give the corresponding definitions of sequential rationality, full consistency, and sequential equilibrium. The final subsection shows that these definitions are equivalent to the conventional definitions of full consistency and sequential equilibrium, expressed in terms of assessments (Kreps and Wilson 1982).

2.1 Situations, strategies, and information sets

In the extensive-form representation, we customarily use the term *information set* for a set of decision nodes between which a player cannot distinguish, and this defines a distinct *situation* in which the player must make a choice. It will be most useful to account for situations without making reference to nodes, and to account for the players' information in terms of subsets of strategy profiles.⁶ Therefore, I use the word *situation* for the distinct places in the game at which players must choose actions, and I use the term *information set* (in the strategic form) for the information about the strategy profile that players have at these points in the game.

Assume situations are distinctly labeled, let H be the set of situations in the game, denote by H_i the set of situations for a given player i , and let $H_+ \equiv \cup_{i \in N_+} H_i$ be the set of situations for all strategic players. For every $h \in H$, let $A(h)$ denote the set of feasible actions available to the player on the move at h , and let $A \equiv \cup_{h \in H} A(h)$.

A strategy profile, including Chance's strategy, is denoted by $s = (a_h)_{h \in H}$, where $a_h \in A(h)$ for every $h \in H$. The space of strategy profiles is $S = \times_{h \in H} A(h)$. For every $i \in N_+$, let $u_i : S \rightarrow \mathbb{R}$ be player i 's payoff function, extended to probability distributions on the strategy profiles as a von Neumann-Morgenstern utility representation.

For any $s = (a_h)_{h \in H} \in S$ and set of situations $L \subset H$, let $s_L = (a_h)_{h \in L}$ denote the profile of the components in L . For any $L \subset H$, define $-L \equiv H \setminus L$. We can then write $s = (s_L, s_{-L})$. For $X \subset S$, define $X_L \equiv \{s_L \mid s \in X\}$ as the projection to S_L . In the case of $L = \{h\}$ for any $h \in H$, we simplify notation by dropping the brackets. For instance, we write X_h instead of $X_{\{h\}}$. Likewise, let subscript " i " stand in for " H_i ", so that s_i means the same thing as s_{H_i} , which is player i 's strategy. Note that $S_h = A(h)$ and, for a given strategy profile s , s_h is the action chosen at h by the player on the move.

Definition: For any $h \in H$, denote by $S(h)$ the set of strategy profiles that reach situation h .⁷ Call $S(h)$ the *(strategic form) information set at h* .

The information set at h represents what the player on the move has observed about past play, which is that the actual strategy profile must be in $S(h)$. 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 reaching h .

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

⁶Expressing information sets as subsets of strategy profiles is standard for defining extensive-form rationalizability (Pearce 1984) and conditional dominance (Shimoji and Watson 1998). It is also utilized for equilibrium theories, as noted in the Introduction. Differentiating between situations and information is useful beyond the present study, such as for modeling limited memory; see, for instance, Battigalli and Generoso (2024).

⁷That is, a strategy profile is contained in $S(h)$ if and only if, in the extensive form, the induced path includes a node in h , meaning there is a node in h that is a truncation of the action sequence induced by s .

Perfect recall implies that the information sets for an individual player have a product structure. To be precise, for every player $i \in N_+$, $h \in H_i$, and $L \subset H_i$, it is the case that $S(h)$ is an L -product set.⁸

2.2 Probability mass functions

For a given finite space W , let ΔW denote the space of probability distributions on W . Then for any subset $T \subset W$, let us take “ ΔT ” to mean the subset of ΔW with support in T . Typically, W will be S or S_L for some $L \subset H$. Recall the definitions of marginal distribution and conditional probability:

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

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

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: Given $L \subset H$ and a distribution $p \in \Delta S$, say that **p exhibits L -independence** if, for every L -product set $X \subset S$, $p(X) = p_L(X_L) \cdot p_{-L}(X_{-L})$.

2.3 Behavior-strategy profiles

Recall that a *behavior-strategy profile* is commonly defined as a probability distribution on the actions at each situation (extensive-form information set), with independence across situations. That is, it is given by $\pi = (\pi_h)_{h \in H} \in \times_{h \in H} \Delta A(h)$. Notice that a behavior-strategy profile can be equivalently written as a distribution $\sigma \in \Delta S$ defined by $\sigma(s) = \prod_{h \in H} \pi_h(s_h)$ for every $s \in S$. This format will be most convenient for our analysis.

Definition: Call $\sigma \in \Delta S$ a **behavior-strategy profile** if there exists $(\pi_h)_{h \in H} \in \times_{h \in H} \Delta A(h)$ such that $\sigma(s) = \prod_{h \in H} \pi_h(s_h)$ for every $s \in S$. Likewise, for any player $i \in N$, call $\sigma_i \in \Delta S_i$ a **behavior strategy** if there exists $(\pi_h)_{h \in H_i} \in \times_{h \in H_i} \Delta A(h)$ such that $\sigma_i(s_i) = \prod_{h \in H_i} \pi_h(s_h)$ for every $s_i \in S_i$.

Let us denote by $\underline{\Delta S}$ the subset of ΔS comprising the behavior-strategy profiles as defined above. One can show (see Watson 2025) that $\sigma \in \underline{\Delta S}$ if and only if σ exhibits h -independence for every $h \in H$. Clearly, $\underline{\Delta S}$ is homeomorphic to $\times_{h \in H} \Delta A(h)$ and there is no loss in

⁸For any player $i \in N_+$ and $h, h' \in H_i$, it must be that either $S(h) = S_{h'} \times S(h)_{-h'}$ or $S(h)_{h'}$ is a singleton and so $S(h) = S(h)_{h'} \times S(h)_{-h'}$. This is because, if the first equation does not hold, there exist strategy profiles s and s' such that $s_{-h'} = s'_{-h'}$, $s \in S(h)$, and $s' \notin S(h)$, which means that at h player i has already made the choice at h' , and so, by perfect recall, it must be that $S(h)_{h'}$ is a singleton.

expressing behavior-strategy profiles as elements of $\underline{\Delta}S$. Hereinafter, I shall express behavior-strategy profiles in this way, keeping in mind that $\underline{\Delta}S$ is a proper subset of ΔS . Let us assume that Chance's strategy σ_0 is a behavior strategy with full support.

2.4 Beliefs and plans

I shall express also as elements of ΔS the players' beliefs and plans at their various situations, using the term *appraisal*.

Definition: For any $i \in N_+$ and $h \in H_i$, call a distribution $p^h \in \Delta S$ **an appraisal at h** if $p^h \in \Delta S(h)$ and if p^h exhibits h' -independence on S for every $h' \in H_i$. An **appraisal system**, denoted by $P = (p^h)_{h \in H_+}$, gives an appraisal for each situation of the strategic players.⁹

To understand the meaning of an appraisal, note that, since a strategy profile describes the players' actions at every situation, a distribution on S can be interpreted as a theory of how the game is being played. For any player $i \in N_+$ and any situation $h \in H_i$, the appraisal p^h is player i 's theory at h . The condition $p^h \in \Delta S(h)$ means the appraisal puts probability one on reaching h , in line with player i 's knowledge. The independence condition means that player i views his own behavior as independent of the other players' behavior, and that his own behavior is represented as a behavior strategy. Thus, one can always write the appraisal p^h as (σ_i, θ_{-i}) , where $\theta_{-i} = p_{-i}^h \in \Delta S_{-i}$ is player i 's belief about the other players' strategies, and $\sigma_i = p_i^h$ is player i 's belief about his own behavior strategy (combining recalled prior actions and the plan for situations that have not yet been reached). I shall usually say “appraisal p^h ” instead of “appraisal at h, p^h ,” as typically h will be given explicitly.

2.5 Sequential rationality

We can test whether an appraisal at h specifies rational behavior for the player on the move, meaning that the actions assigned positive probability at situation h maximize the player's expected payoff.

Definition: 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 $p_{-h}^h = \hat{p}_{-h}^h$.

To understand the definition of h -deviation, recall that p^h , as an appraisal at h , is a distribution on $S(h)$ that satisfies h' -independence for every $h' \in H_i$, where i is the player who is active at h . Since h itself is in H_i , we therefore have $p^h(s) = p_h^h(s_h) \cdot p_{-h}^h(s_{-h})$ for every $s \in S$. That is, p^h is the product of its marginal on component h and its marginal on the complementary component $-h$. An h -deviation $\hat{p}^h$ replaces the marginal on component h , while keeping the same marginal on the complementary component, so that $\hat{p}^h(s) = \hat{p}_h^h(s_h) \cdot p_{-h}^h(s_{-h})$ for every $s \in S$. Therefore, $\hat{p}^h$ expresses a single deviation from p^h , and it occurs at h .

⁹Watson (2025) defines appraisal systems to include the players' beliefs at artificial “initial situations.” This additional structure is useful for stating alternative belief-consistency notions and various forms of perfect Bayesian equilibrium, but it is not needed here.

Definition: 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** $P = (p^h)_{h \in H_+}$ is **sequentially rational** if p^h is rational at h , for every $h \in H_+$.

Here sequential rationality is defined in terms of single deviations, meaning that we evaluate player i 's rationality at a given situation $h \in H_i$ by looking only 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.¹⁰ The familiar *one-deviation property* (equivalence between single-deviation optimality and continuation-strategy-deviation optimality) holds under the full-consistency condition of sequential rationality to be imposed below.¹¹

2.6 Full consistency and sequential equilibrium

Sequential equilibrium (Kreps and Wilson 1982) requires assessments to satisfy *full consistency*. This subsection puts the definition of full consistency in terms of appraisal systems and gives the definition of sequential equilibrium in the same format.

Definition: Call a behavior-strategy profile $\sigma \in \underline{\Delta}S$ **fully mixed** if $\sigma(s) > 0$ for every $s \in S$, and σ_0 is the strategy for Chance given in the representation of the game.

A fully mixed behavior-strategy profile σ reaches every situation $h \in H_+$ with positive probability, implying a well-defined probability distribution $\sigma(\cdot | S(h))$ calculated from the conditional-probability formula.

Definition: Consider an appraisal system P and behavior-strategy profile $\sigma \in \underline{\Delta}S$. Say that (P, σ) is **fully consistent** if there is a sequence of fully mixed behavior-strategy profiles $(\sigma^k)_k$ such that $\sigma^k \rightarrow \sigma$ and $\sigma^k(\cdot | S(h)) \rightarrow p^h$ for every $h \in H_+$.

Definition: Given an appraisal system P and behavior-strategy profile $\sigma \in \underline{\Delta}S$, say that (P, σ) is a **sequential equilibrium** if (P, σ) is fully consistent and P is sequentially rational.

2.7 Equivalence with assessment-based versions

This subsection shows there is no loss in formulating the players' contingent beliefs and plans as an appraisal system rather than as an assessment. The two formulations are equivalent for all intents and purposes. I start by stating Kreps and Wilson's (1982) definitions, expressing behavior-strategy profiles in the space $\underline{\Delta}S$ for convenience in making comparisons.

Recall that a situation $h \in H_+$, as an extensive-form information set, is a subset of decision nodes. Let $D_+ \equiv \cup_{h \in H_+} h$ denote the set of all decision nodes of the strategic players. An assessment is given by (μ, π) , where π is a behavior-strategy profile and μ is a mapping from

¹⁰Note 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 .

¹¹This result, also (and perhaps better) called the "one-deviation principle," holds under a weaker condition that Watson (2025) calls *minimal consistency*. The result is provided by Perea (2000) for belief systems as modelled here, and by Hendon, Jacobsen, and Sloth (1996) in the context of assessments. Watson (2025) proves a mixed-strategy, n -player version of Perea's (2000) Theorem 3.1, in the form needed for the setting herein.

D_+ to $[0, 1]$ with the property that $\sum_{d \in h} \mu(d) = 1$ for each $h \in H_+$. The interpretation of an assessment is that, for any $i \in N_+$, $h \in H_i$, and $d \in h$, at situation h player i assigns probability $\mu(d)$ to node d and assigns zero probability to nodes outside of h . Further, player i believes that, for every situation h' on a path from d , the player on the move at h' will act according to $\pi_{h'}$. Call μ the *belief function*.

I next demonstrate how to convert an assessment into an equivalent appraisal system. For any node d , let $\underline{L}(d)$ denote the set of situations that are on the path to d , not including the situation containing d .

Definition: For any assessment (μ, π) , *implied appraisal system* $P[\mu, \pi] = (p^h[\mu, \pi])_{h \in H_+}$ is constructed as follows: For each $d \in D_+$, define behavior-strategy profile $\omega^d \in \underline{\Delta}S$ by:

- for $h \in \underline{L}(d)$, let ω_h^d put probability 1 on the action taken at h on the path to d ; and
- for $h \in H \setminus \underline{L}(d)$, let $\omega_h^d \equiv \pi_h$.

Then for each $h \in H_+$, define $p^h[\mu, \pi] \in \Delta S$ by $p^h[\mu, \pi](s) = \sum_{d \in h} \mu(d) \omega^d(s)$ for every $s \in S$, and let $P[\mu, \pi] = (p^h[\mu, \pi])_{h \in H_+}$.

In words, for every $h \in H_+$, probability distribution $p^h[\mu, \pi]$ is a weighted average of behavior-strategy profiles $(\omega^d)_{d \in h}$, indexed by the nodes in h . For each node $d \in h$, ω^d puts probability 1 on the actions taken at predecessor nodes needed to reach d , and otherwise plays according to π . The weighted average puts weight $\mu(d)$ on ω^d .

Lemma 1: For any given assessment (μ, π) , $P[\mu, \pi]$ is an appraisal system. Further, from each $h \in H_+$, (μ, π) and $P[\mu, \pi]$ induce the same distribution over terminal nodes, and hence the same expected payoffs.

Our next task is to state the condition of full consistency for assessments. Extending the definition of $S(h)$ and expanding notation, for any decision node d in the extensive form, let $S(d)$ denote the set of strategy profiles that reach d . Note that for any $h \in H_+$ and distinct nodes $d, d' \in h$, $S(d) \cap S(d') = \emptyset$. Therefore, for any distribution $\sigma \in \Delta S$, we have $\sum_{d \in h} \sigma(S(d)) = \sigma(S(h))$. A fully mixed behavior-strategy profile σ reaches every situation $h \in H$ with positive probability, implying a well-defined belief function derived using the conditional-probability formula as follows:

Definition: Given any fully mixed $\sigma \in \underline{\Delta}S$, the *implied belief function* $\rho[\sigma]: D_+ \rightarrow [0, 1]$ is defined thusly: for $h \in H_+$ and $d \in h$, $\rho[\sigma](d) \equiv \sigma(S(d) | S(h))$.

Clearly $\sum_{d \in h} \rho[\sigma](d) = 1$ for each $h \in H_+$. Full consistency requires that the players' beliefs are given by the limit of such conditional-probability updating, in relation to a sequence of fully mixed behavior-strategy profiles that converges to the equilibrium strategy profile.

Definition: Say that an assessment (μ, π) is *fully consistent* if there is a sequence of fully mixed behavior-strategy profiles $(\sigma^k)_k$ such that $\sigma^k \rightarrow \pi$ and $\rho[\sigma^k] \rightarrow \mu$.

The following result does most of the work in establishing the equivalence of appraisals and assessments for describing the players' beliefs and plans.

Lemma 2: *Given any appraisal system P and behavior-strategy profile π , (P, π) is fully consistent if and only if there is an assessment (μ, π) that is fully consistent and satisfies $P = P[\mu, \pi]$.*

In addition, we must check that the conditions of sequential rationality are equivalent. To evaluate sequential rationality of an assessment, we take each player $i \in N_+$ and situation $h \in H_i$, calculate player i 's expected payoff conditional on h under the given behavior-strategy profile π and belief μ , and compare it to the expected payoff player i would get by deviating at h (a single deviation). From the induced-appraisal construction and Lemma 1, this is equivalent to calculating $u_i(p^h)$ for $p^h = p^h[\mu, \pi]$, and calculating $u_i(\hat{p}^h)$ for every h -deviation $\hat{p}^h$. To see why, observe that deviating to mixed action α_h at situation h gives player i an expected payoff of $u_i(p^h[\mu, \hat{\pi}])$, where $\hat{\pi}$ is the behavior-strategy profile defined by $\hat{\pi}_h = \alpha_h$ and $\hat{\pi}_{h'} = \pi_{h'}$ for every $h' \neq h$. Further, it is clear from its definition that the induced appraisal $p^h[\mu, \hat{\pi}]$ is an h -deviation in relation to $p^h[\mu, \pi]$. Thus, evaluating sequential rationality for an assessment is equivalent to evaluating sequential rationality of its implied appraisal system.

3 Partial Equilibrium Construction

Let us consider a partial equilibrium construction that focuses on a given set of situations J , assumed to contain H_0 and at least one situation of a strategic player. That is, $J \in \mathcal{H}$, where

$$\mathcal{H} \equiv \{L \subset H \mid H_0 \subset L \text{ and } L \cap H_+ \neq \emptyset\}.$$

For each player i , let $J_i \equiv J \cap H_i$, and let $J_+ \equiv J \cap H_+$.

Definition: *Given $J \in \mathcal{H}$, call $\phi \in \Delta S_J$ a **J-behavior-strategy profile** if there exists $(\pi_h)_{h \in J} \in \times_{h \in J} \Delta A(h)$ such that $\phi(s_J) = \prod_{h \in J} \pi_h(s_h)$ for every $s_J \in S_J$.*

Independence is defined as before, applied here on the components of S_J . Let $\underline{\Delta} S_J$ denote the subset of ΔS_J comprising the J -behavior-strategy profiles as defined above.

I next describe 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 J_+ about the behavior at (and only at) the situations in J .

Definition: *Given $J \in \mathcal{H}$, $i \in N_+$, and any $h \in J_i$, 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 $h' \in J_i$. A **J-appraisal system** gives an appraisal for every $h \in J_+$ and is denoted by $Q = (q^h)_{h \in J_+}$.*

Definition: *Given $J \in \mathcal{H}$, call a J -behavior-strategy profile ϕ **fully mixed** if $\phi(s_J) > 0$ for every $s_J \in S_J$, and if σ_0 is the strategy for Chance given in the representation of the game.*

Definition: *Take as given any $J \in \mathcal{H}$, J -appraisal system Q , and J -behavior-strategy profile $\phi \in \underline{\Delta} S_J$. Say that **(J, Q, ϕ) is fully consistent** if there is a sequence of fully mixed J -behavior-strategy profiles $(\phi^k)_k$ such that $\phi^k \rightarrow \phi$ and $\phi^k(\cdot \mid S(h)_J) \rightarrow q^h$ for every $h \in J_+$.*

Because the partial equilibrium construction will not deal with 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: Take as given any $J \in \mathcal{H}$ and J -appraisal system Q . For any player $i \in N_+$ and $h \in J_i$, say that **J -appraisal q^h is r -rational at h** if $u_i(p^h) \geq u_i(\hat{p}^h)$ for every appraisal p^h at h satisfying $p_J^h = q^h$, and for every h -deviation $\hat{p}^h$. Say that **Q is sequentially r -rational** if q^h is r -rational at h , for every $h \in J_+$.

The extra “ r ” in r -rational stands for robust to behavior at the situations in $-J$. To understand this condition, note that p^h in the definition is not required to exhibit J -independence. It could exhibit correlation between sub-components of $-h$, such as between actions $s_{h'}$ and $s_{h''}$ where $h' \in J \setminus \{h\}$ and $h'' \in -J$. Still, as an appraisal at h , it must exhibit h' -independence for every $h' \in H_i$ including for $h' = h$, where i is the active player at h . The same is true for q^h . For the r -rationality test, we look at every such appraisal p^h whose marginal on component J is equal to q^h .

Definition: Take as given any $J \in \mathcal{H}$, J -appraisal system Q , and J -behavior-strategy profile $\phi \in \underline{\Delta}S_J$. Say that **(J, Q, ϕ) is a partial sequential equilibrium** if (J, Q, ϕ) is fully consistent and Q is sequentially r -rational.

4 Theorem

Our concern is whether partial sequential equilibria can be extended to construct sequential equilibria that coincide on J . For any appraisal system P and set $J \subset H$, denote by $P_J = (p_J^h)_{h \in J_+}$ the marginals on strategy component J of the appraisals at situations in J_+ .

Definition: Given $J \subset H$, $\sigma \in \underline{\Delta}S$, $\phi \in \underline{\Delta}S_J$, appraisal system P , and J -appraisal system Q , say that **(J, Q, ϕ) extends to (P, σ)** if $\sigma_J = \phi$ and $P_J = Q$.

Incentive conditions do not get in the way of extending a partial sequential equilibrium to a full sequential equilibrium, due to the robustness requirement of the partial equilibrium definition. The problem is with beliefs, since the appraisals at situations in J , particularly about behavior in J , generally depend on the distribution of actions taken at situations in $-J$, for whatever sequence $(\sigma^k)_k$ of fully mixed behavior-strategy profiles is needed for full consistency in the entire game. In other words, for a given $h \in J$, the marginal p_J^h depends on σ_{-J}^k as $k \rightarrow \infty$. Further, for $h \in -J$, p_{-J}^h generally depends on σ_J^k .

The main creative step is to find a condition guaranteeing that appraisals to be constructed for situations in J will effectively treat the behavior in J as *independent* of the behavior in $-J$, regardless of what is specified on the latter component. Consider the following condition on a given J -appraisal system Q .

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

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




Figure 3: First example.

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

Figures 3 and 4 provide a graphical depiction of the rectangular margin-support condition in the two examples discussed in the Introduction, where $J \equiv \{h_1, h_2, h_3\}$ in the first example and $J \equiv \{h_1, h_3\}$ in the second example. In Figure 3, the table at the right represents the possible combinations of action profiles taken at situations h_1, h_2 , and h'_3 . These are arranged in terms of components J and $-J$, with the profiles of actions at h_1 and h_2 (the ones in J) as the rows and the actions at h'_3 (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 4 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 3 is not a product set. It is not the product of $\{\alpha e, \alpha f, \beta g, \beta e, \beta f\}$ and $\{w, x\}$. Even with a fully mixed behavior-strategy profile (as we have in 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 (randomized) behavior on the $-J$ component would affect the updated belief about the J component. But if the partial construction makes the conditional probability about the actions at h_1 and h_2 converge to something with support in $\{\alpha e, \alpha f\}$, as was the case discussed in




Figure 4: Second example.

the Introduction, then the rectangular margin-support condition is satisfied. Likewise, the shaded region in Figure 4 is not a product set. However, if the partial construction makes the conditional probability about the action at h_1 converge to probability 1 on $\{\alpha\}$, as in the Introduction, then the rectangular margin-support condition is satisfied.

For both of these examples, a sequential equilibrium can be assembled so that the J -margin of player 3's appraisal at h_3 is the same as in the specified partial sequential equilibrium. A key aspect of the full equilibrium construction is to impose a sequence of lower bounds on the probabilities of actions taken at the situations in $-J$, approaching zero more slowly (by an order of magnitude) than does the sequence utilized for the partial-equilibrium construction on J . Under the rectangular margin-support condition, the resulting appraisals in the full construction have the desired relation to those in the partial construction.

The analysis applies in general, leading to the main result, which is stated next.

Theorem: *For any partial sequential equilibrium (J, Q, ϕ) of a finite extensive-form game with perfect recall, if (J, Q, ϕ) satisfies $J \in \mathcal{H}$ and the rectangular margin-support condition, then (J, Q, ϕ) extends to a sequential equilibrium.*

For an example in which the rectangular margin-support condition fails and the Theorem does not apply, consider the game shown in Figure 5. Let the situations of interest be $J = \{h_1, h_2, h'_2\}$. Define a sequence of fully mixed J -behavior-strategy profiles $\{\phi^k\}$ so that the probabilities of actions α and β at h_1 are both $1/k$, the probabilities of actions g and e at h_2 are both $1/k$, and the probabilities of actions w and z at h'_2 are both $1/2$. Note that $S(h'_2)_J = \{\alpha f, \beta f\} \times \{w, z\}$. 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 γfw and γfz ; and both q^{h_2} and $q^{h'_2}$ put probability $1/4$ on each of the partial strategy profiles αfw , αfz , βfw , and βfz . It is easy to confirm that Q is J -sequentially r-rational. Note, by the way, that for the rationality test at h'_2 , the only appraisal $p^{h'_2}$ satisfying $p^{h'_2}_J = q^{h'_2}$ is the one for which the marginal $p^{h'_2}_{-h'_2}$ puts probability $1/2$ on both αfy and βfx .

We have shown that Q is a 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 Fig-




Figure 5: Third example.

ure 5, 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 αfw , αfz , βfw , and βfz , $\{s \in S(h'_2) \mid s_J \in \text{supp } q^{h'_2}\}$ is *not* a J -product set. In fact, there is no sequential equilibrium in which player 1 plays strategy γ with positive probability. To see why, note first that player 3 must choose x with probability 1 in any sequential equilibrium. If there were an equilibrium in which γ were played with positive probability, then player 2 would have to choose action f with strictly positive probability at h_2 . On one hand, if the equilibrium appraisal at h'_2 satisfies $p_{-h'_2}^{h'_2}(\alpha fy) \geq p_{-h'_2}^{h'_2}(\beta fx)$, then for full consistency the appraisal at h_2 must satisfy $p_{h_1}^{h_2}(\alpha) = 1$. But then e is player 2's only rational choice at h_2 , a contradiction. On the other hand, if $p_{-h'_2}^{h'_2}(\alpha fy) < p_{-h'_2}^{h'_2}(\beta fx)$, then player 2 must choose z for sure at h'_2 , which makes β player 1's uniquely rational choice, a contradiction.

5 Applications

To demonstrate the usefulness of the Theorem, in this section I first describe how the result may simplify analysis of applications in a standard class of games with publicly observed actions. Then I provide an economic example in a different setting, of decentralized contracting, in which actions are not publicly observed. In both settings, the Theorem allows for a partial equilibrium construction on a substantially reduced set of situations. For instance, in the decentralize contracting setting, with n players and κ possible productive outcomes, there are in excess of $2^{2n}(3n - 2)n\kappa + 2n$ situations, whereas in the partial equilibrium construction, J contains at most $3n^2$ situations.

5.1 Multistage games with observed actions and independent types

Consider the class of finite multistage games with observed actions and independent types. To keep the notation simple, let us specify different chance-agents that select the types of the different strategic players. There are n strategic players, numbered $1, 2, \dots, n$, and correspondingly n chance-agents called $01, 02, \dots, 0n$. The game begins at stage 0, when, simultaneously, each chance-agent $0i$, at a situation denoted by h_{0i} , chooses the type of strategic player i , denoted by $\theta_i \in \Theta_i$. Types are drawn independently according to an exogenously given probability distribution with full support, and each strategic player i observes her own type θ_i but nothing else. Then in each of the stages $1, 2, \dots, T$, the strategic players simultaneously choose actions, where the sets of feasible actions may depend on previous actions but not on the types. At the end of each stage, the players observe the profile of actions chosen. Payoffs depend on the type vector and sequence of action profiles.¹²

An implication of the Theorem is that, to construct any sequential equilibrium for this class of games, it is enough to specify beliefs and behavior for the set of personal histories

¹²As usual, to represent these multistage games in the extensive form, simultaneous choices would formally appear as sequential moves, with information sets capturing the informational content of simultaneity. For convenience, I express play paths as sequences of action profiles rather than as sequences of individual actions. Incidentally, the analysis below can be generalized to allow the set of feasible actions for a given player i in any stage to depend on player i 's type.

that entail at most one public deviation per stage. A public deviation is a case of a player selecting an action that, according to the equilibrium strategies and beliefs, should not have been chosen with strictly positive probability. In other words, we can ignore situations that come after simultaneous public deviations. Details follow.

Let $a^t = (a_1^t, a_2^t, \dots, a_n^t)$ denote the profile of actions chosen in stage t . A public history is denoted by $b = (a^1, a^2, \dots, a^{T'})$ for any positive integer $T' \leq T$, or by $b = \emptyset$ signifying the null history at the first stage. For each $i \in N_+$, player i 's personal history is given by $h = (b, i, \theta_i)$. Let B denote the set of public histories. Note that, maintaining the notation developed earlier, H denotes the set of all situations, $H_i = \{(b, i, \theta_i) \mid b \in B, \theta_i \in \Theta_i\}$ is the set of situations for each player $i \in N_+$, and H_0 denotes the set of situations for the chance agents. Define

$$\hat{\mathcal{H}} \equiv \{J \subset H \mid H_0 \subset J \text{ and } J \cap H_i \neq \emptyset \text{ for every } i \in N_+\}.$$

For any $J \in \hat{\mathcal{H}}$, define $L_i(J) \equiv J_i \cup \{h_{0i}\}$ for each $i \in N_+$.

For any $b \in B$, let us extend notation to let $S(b)$ denote the subset of strategy profiles (including types chosen by the chance agents) that would make public history b occur. For any player $i \in N_+$ and type θ_i , for which we can write $h = (b, i, \theta_i)$, we also have the previously defined $S(h)$, which is the set of strategy profiles that reach player i 's situation h . For any given $J \in \hat{\mathcal{H}}$, by perfect recall and observed actions, we have

$$S(b, i, \theta_i)_J = S(b, i, \theta_i)_{L_i(J)} \times \left(\times_{j \in N_+ \setminus \{i\}} S(b)_{L_j(J)} \right),$$

for every $i \in N_+$ and $\theta_i \in \Theta_i$.

Take as given a set $J \in \hat{\mathcal{H}}$, and consider a J -appraisal system Q and J -behavior-strategy profile ϕ such that (J, Q, ϕ) is fully consistent. An implication of full consistency and the product structure described in the previous paragraph is that there exists $R = (r_j^b)_{b \in B, j \in N_+}$ such that $r_j^b \in \Delta S_{L_j(J)}$ for every $b \in B$ and $j \in N_+$, and the following holds: For every $(b, i, \theta_i) \in H_+$ and $j \in N_+ \setminus \{i\}$, it is the case that $q^{(b, i, \theta_i)}_{L_j(J)}$ exhibits $L_j(J)$ -independence and $q^{(b, i, \theta_i)}_{L_j(J)} = r_j^b$.

In words, for public history b , there is a “public belief” about the strategy and type of each strategic player. At player i 's corresponding personal history, player i 's appraisal specifies exactly the public belief regarding each of the other players, and it exhibits independence across the players. This claim is easily proved by noting that, for a full-support prior distribution exhibiting independence across two components of the strategy-profile space, the posterior probability conditional on a conjunctive event is the product of the updated marginals (Watson 2025), and then noting that the product characterization extends to the limit.

Definition: For any $b, b' \in B$, say that **b and b' are consecutive public histories** (in this order) if b' is a sequence equal to b appended with one additional action profile.

Definition: Take as given any $J \in \hat{\mathcal{H}}$, J -appraisal system Q , and J -behavior-strategy profile ϕ such that (J, Q, ϕ) is fully consistent. For $b, b' \in B$, say that **(b, b') has the AMOD property** if b and b' are consecutive public histories and $r_j^b(S(b')_{L_j(J)}) = 0$ for at most one $j \in N_+$.

Definition: Say that a partial sequential equilibrium (J, Q, ϕ) is **AMOD-minimal** if J is the minimal subset of $\hat{\mathcal{H}}$ satisfying $(\emptyset, i, \theta_i) \in J$ for every $i \in N_+$ and $\theta_i \in \Theta_i$, and $[(b, b') \text{ has the AMOD property and } (b, i, \theta_i) \in J] \text{ implies } (b', i, \theta_i) \in J$.

Thus, (J, Q, ϕ) is AMOD-minimal if J comprises (only) the chance-agent situations, the personal histories in stage 1, and the personal histories in other stages that entail no simultaneous public deviations. AMOD stands for “at most one (public) deviation.”

Proposition 1: *For any finite multistage game with perfect recall, observed actions, and independent types, every AMOD-minimal partial sequential equilibrium extends to a sequential equilibrium, and every sequential equilibrium is an extension of an AMOD-minimal partial sequential equilibrium.*

A proof of this proposition may be found in Section 6. The first claim is verified by noting that $S(h) = S(h)_J \times S_{-J}$ for every $h \in J_+$, which follows from the fact that no personal history in $-J$ can be reached on the path to any personal history in J . This implies the rectangular margin-support condition, and then the Theorem establishes the existence of a sequential equilibrium to which the given partial sequential equilibrium extends.

A more cumbersome series of steps is required to prove the second claim, but the simple intuition is this: One can derive a set J from a given sequential equilibrium (P, σ) and then use it to define J -behavior-strategy profile $\phi = \sigma_J$ and J -appraisal system $Q = P_J$. One can show that (J, Q, ϕ) is fully consistent using the J -marginals of the sequence of fully mixed behavior-strategy profiles that establishes full consistency of (P, σ) . Because the inequalities for sequential rationality at $h \in J$ depend not on what players would do after any simultaneous public deviations in the continuation of the game (such is a zero-probability event according to the active player’s appraisal at h and for any h -deviation), r -rationality at h is implied for q^h . The final step is to show that (J, Q, ϕ) is AMOD-minimal.

Putting Proposition 1 to work, a sequential equilibrium can be constructed using the following iterative procedure. In round 1 of the procedure, the analyst describes the set B^1 of action sequences that are supposed to occur with positive probability on the equilibrium path. Further, letting J^1 denote the set of personal histories for these action sequences, the analyst specifies the mixed actions to be chosen at each personal history in J^1 . Round 1 concludes by letting $\bar{B}^1 = B^1$. In round 2, the analyst describes the action sequences, B^2 , that entail a single public deviation in stage 1 of the game, in relation to sequences in $\bar{B}^1$, and thereafter occur with positive probability in equilibrium following the public deviation. Letting J^2 denote the set of personal histories in stages 2 through T for the action sequences in B^2 , the analyst specifies the mixed actions to be chosen at each personal history in J^2 . Let $\bar{B}^2 = \bar{B}^1 \cup B^2$.

This process proceeds inductively in rounds $k = 2, 3, \dots, T$. In round k , the analyst describes the action sequences, B^k , that entail a single public deviation in stage $k - 1$ from sequences in $\bar{B}^{k-1}$ and thereafter occur with positive probability in equilibrium. Letting J^k denote the set of personal histories in stages k through T for the action sequences in B^k , the analyst specifies the mixed actions to be chosen at each personal history in J^k , and we let $\bar{B}^k = \bar{B}^{k-1} \cup B^k$. From these steps, $J \equiv J^T$ and J -behavior-strategy profile ϕ have been identified. The analysis then specifies a sequence of fully mixed J -behavior-strategy

profiles to create the J -appraisal system Q , and confirms the conditions for a partial sequential equilibrium.

5.2 Example of efficient decentralized contracting with externalities

I next sketch an example of decentralized contracting along the lines of Watson (2024). A group of n players will play a simultaneous-move underlying game with externally enforced transfers. Before the underlying game is played, some pairs of players have the opportunity to form bilateral contracts by sending messages to each other. A given network specifies which pairs of players can communicate and contract. All messages are private. Each player observes only the messages she sent and received from individual other players. Therefore, at the end of the contracting phase, each player will know only the contracts she formed; she will have observed nothing about the contracts formed by pairs of other players. The underlying game is played in this setting of informational asymmetry.

The grand game consists of interaction in the contracting protocol, constrained so that communication/contracting is possible only on network links, followed by interaction in the underlying game and external enforcement of contracted transfers. Externalities are generally pervasive, as players care about underlying-game actions taken by others with whom they cannot contract. The outcome of the underlying game is assumed to be globally verifiable, so that a contract between a given pair of players may specify transfers between them as a function of the underlying-game actions chosen by others. The key question is whether externalities will be internalized by the contracts that the players form in an equilibrium of the grand game.

Watson (2024) shows that efficient contracting is possible in general, and that it requires a sequential contracting protocol allowing flexibility for players to form and later cancel contractual arrangements. In the equilibrium construction, pairs of players establish *assurance contracts* in which one or both players provide a guarantee (subject to a monetary penalty) of the underlying-game actions that other players will select. Also, for any player i who is linked to both players j and k , player i 's decision of whether to form, or later cancel, a contract with player j is typically conditioned on player i 's experience with player k . Neither player j nor player k observes player i 's contracting choices with the other, but the equilibrium contracts link these relationships in some manner.

Watson's (2024) general "possibility" result establishes the existence of a single contracting protocol that performs well for any given set of underlying games and connected networks. Performing well means that for any underlying game and network, there is an efficient sequential equilibrium of the grand game. The analysis makes significant use of the Theorem herein but, even so, is extraordinarily complicated owing to breadth of the general result and contracting protocol that is needed for the proof, which has a huge number of situations requiring classification.

I seek here to demonstrate the essence of the logic and to highlight the utility of the Theorem in this application. To reduce complexity, I present a version of the model that makes several simplifying assumptions regarding the enforcement technology and contracting protocol. I allow only a narrow class of *assurance contracts* and external enforcement of simple penalties, and I fix the underlying game and network. There will be no moves of Chance/nature in

this model, yet there is a great deal of asymmetric information (and nontrivial considerations of off-equilibrium-path beliefs) because the players all observe different things as the game progresses. There are no proper subgames.

Players labelled $1, 2, \dots, n$ will interact over $n + 1$ periods of time labelled $0, 1, 2, \dots, n$. Since there is no Chance/nature player, we can write $N = \{1, 2, \dots, n\}$. Periods 0 through $n - 1$ are collectively called the *contracting phase*; interaction in this phase determines the bilateral contracts that the players form. Period n is called the *production phase*, at which time the players simultaneously choose actions in the underlying game. An undirected and irreflexive network, denoted by $K \subset N \times N$, gives the set of links between players. A given pair of players (i, j) can contract if and only if $(i, j) \in K$.

Let E_i denote the set of feasible actions (such as effort levels in a team-production application) for player $i \in N$, and let $v_i: E \rightarrow \mathbb{R}$ denote player i 's payoff function in the underlying game, where $E = \times_{j \in N} E_j$ is the set of productive action profiles. Assume that the underlying game has a Nash equilibrium in pure actions, $\underline{a}$, and has an efficient action profile a^* satisfying $a_i^* \neq \underline{a}_i$ for every player $i \in N$. To make things interesting economically, assume also that $n > 2$ and a^* is not a Nash equilibrium of the underlying game.

When referring to a contract for a pair $(i, j) \in K$, let us use the convention of ordering this pair so that $i < j$, to avoid confusion and double counting. An assurance contract is a vector $(\hat{e}, Z, Z')$, where $\hat{e} \in E$ is the target action profile to be chosen in the underlying game, $Z \subset N$ is the set of players whom player i guarantees will select their part of the target action profile, and $Z' \subset N$ is the set of players whom player j guarantees will select their part of the target action profile.

External enforcement is parameterized by a penalty increment ζ that we assume satisfies $\zeta > \max\{v_i(e) - v_i(e') \mid i \in N; e, e' \in E\}$. Enforcement of the contract $(\hat{e}, Z, Z')$ for the pair (i, j) means that, after the underlying game is played, player i is compelled to pay a penalty of $n\zeta \cdot \#\{k \in Z \mid e_k \neq \hat{e}_k\}$ to player j . That is, player i pays $n\zeta$ for every player $k \in Z$ who deviates from $\hat{e}_k$. Likewise, player j is required to pay a penalty of $n\zeta \cdot \#\{k \in Z' \mid e_k \neq \hat{e}_k\}$ to player i . Note that there are no penalties if $\hat{e}$ is the outcome of the underlying game. The set of feasible contracts is denoted by $C \equiv \{(\hat{e}, Z, Z') \mid \hat{e} \in E; Z, Z' \subset N\}$. Let us denote by $\underline{c}$ the absence of a contract.

The contracting phase works as follows. In period 0, simultaneously players submit contract offers and learn whether contracts were formed with their linked partners. To be precise, each player i 's action in period 0 is a vector $(c^{ij})_{j \neq i}$, where $c^{ij} \in C \cup \{\underline{c}\}$ is player i 's contract offer to player j . We interpret the offer $\underline{c}$ as silence. Communication is constrained by the network, in that player i is restricted to choose $c^{ij} = \underline{c}$ for every player j for which $(i, j) \notin K$. If a pair of players (i, j) make the same contract offer to each other, that is if $c^{ij} = c^{ji}$, then these players are informed that this contract has been formed; otherwise, they are informed only that no contract has been formed.

In periods 1 through $n - 1$, communication is not allowed between pairs of players that did not form a contract in period 0. On the other hand, if a pair of players (i, j) formed a contract in period 0, then in each period $t \in \{1, 2, \dots, n - 1\}$ these players simultaneously choose whether to cancel the contract, unless one of both of them had cancelled earlier, and they observe each other's choices. If in period t one player cancels, or if both players cancel simultaneously, then they are released from their contract, but each player who elected to cancel must pay the




Figure 6: Example of a network.

other a cancellation penalty of $t\zeta$. After a cancellation, no further communication is allowed between these players i and j .

In the contracting phase, each player i observes only the contract offers she made to others, whether contracts were formed between herself and her linked partners, and whether she and her linked partners cancelled any of these contracts. Player i observes nothing about contract offers, or whether contracts were formed or cancelled, by the other pairs of players (j, k) for which $j \neq i$ and $k \neq i$. For example, with six players in the network K shown on the left side of Figure 6, player 4 does not observe the contractual interaction of the pairs $(1, 3)$, $(2, 3)$, $(3, 5)$, and $(5, 6)$. With only this limited information, the players choose actions in the underlying game. The grand-game payoff for each player i is this player's underlying-game payoff $u_i(e)$, plus any penalties received, minus any penalties paid.

The issue to be addressed is whether there is an equilibrium of the grand game in which the efficient action profile a^* is played in the production phase with probability one—in other words, whether the players have the incentive to create contracts that induce them to choose a^* . Clearly this would *not* be possible if the contracting phase consisted of just one round, such as simultaneous offers as specified above for period 0 but with no options to cancel. The general reasoning is this: First, to give players the incentive to choose a^* , assurance contracts must be formed in equilibrium, for otherwise players would deviate from a^* since it is not a Nash equilibrium of the underlying game. Second, under the presumption that pairs of players on one side of the network will form assurance contracts and play their part of profile a^* , a player on the other side would typically gain by refusing to contract (sending offer $\underline{c}$) with her partners and then free-riding on the others while choosing a best response in the underlying game.

Proposition 2: *In the decentralized contracting model described above, for any underlying game and connected network, the grand game has an efficient sequential equilibrium.*

I next provide an outline of the proof of Proposition 2, highlighting the role of the Theorem. To keep things simple, I describe the logical steps using the example of $n = 6$ and the network K shown on the left side of Figure 6. It will be clear how to apply the logic more broadly.

Set of situations for the partial equilibrium

Let us find any minimally connected network $\underline{K} \subset K$. For the example, assume it is the one that removes links $(3, 4)$ and $(4, 3)$, so that

$$\underline{K} = \{(1, 3), (3, 1), (2, 3), (3, 2), (3, 5), (5, 3), (4, 5), (5, 4), (5, 6), (6, 5)\},$$

as shown on the right side of Figure 6. Next, we describe the *prescribed path* of play for the equilibrium we intend to construct. The prescribed path is this:

- Period 0 behavior:

Players 1 and 3 offer each other $c^{13} = (a^*, \{1\}, \{2, 3, 4, 5, 6\})$.
 Players 2 and 3 offer each other $c^{23} = (a^*, \{2\}, \{1, 3, 4, 5, 6\})$.
 Players 3 and 5 offer each other $c^{35} = (a^*, \{1, 2, 3\}, \{4, 5, 6\})$.
 Players 4 and 5 offer each other $c^{45} = (a^*, \{4\}, \{1, 2, 3, 5, 6\})$.
 Players 5 and 6 offer each other $c^{56} = (a^*, \{1, 2, 3, 4, 5\}, \{6\})$.
 All other offers are $\underline{c}$.

- Behavior in periods 1 through 5: No contracts are cancelled.
- Period 6 behavior: Each player $i \in N$ chooses a_i^* .

Note that on the prescribed path, pairs (1, 3), (2, 3), (3, 5), (4, 5), and (5, 6) would form the specified contracts and not cancel them. Let $\bar{H}$ be the set of situations reached on the prescribed path.

Next let us define, by induction, a set of situations $\underline{H}$ for periods 1-5 that would be reached following a specific kind of single deviation in period 0 in which exactly one of the players offers $\underline{c}$ to everyone else, instead of making the offers specified for the prescribed path, and where a player who detects that such a deviation may have occurred immediately cancels all of her remaining contracts.

Start with such a deviation by player 1, where this player offers $\underline{c}$ to player 3. We thus include in $\underline{H}$ the situation h_3^1 for player 3 in period 1 in which player 3 sees that no contract was formed with player 1 (so player 3 knows that player 1 deviated) but otherwise player 3 has observed no deviations from the prescribed path. We specify that at h_3^1 player 3 shall cancel her contracts with players 2 and 5. Continuing, we include in $\underline{H}$ the situation h_5^2 for player 5 in period 2 in which player 5 sees that player 3 has cancelled the contract with her (so player 5 may think that this was triggered by player 1 deviating in period 0) but otherwise player 5 has observed nothing that differs from the prescribed path. We specify that, at h_5^2 , player 5 shall cancel her contracts with players 4 and 6. There are no situations for players 2, 4, and 6 to include in $\underline{H}$ yet, because once they learn of a contract cancellation by player 3 or player 5, they have no outstanding contracts.

We repeat this step for players 2-6, in each case starting with a single deviation in period 0 in which one of these players offers $\underline{c}$ to everyone else, and proceeding to situations of the other players that would be reached as they cancel all of their contracts. Thus, $\underline{H}$ is a small set of situations that correspond to one player refusing to form any contracts, and every other player cancelling all of their contracts once they deduce that the initial deviation occurred, either by observing that a contract wasn't formed or that a contract was cancelled by another player. For every situation in $\underline{H}$, the active player is prescribed to cancel all remaining contracts.

Let $\tilde{H}$ comprise the situations in period 6 (the production phase) either that follow situations in $\underline{H}$ (this, for example, applies to players 2-6 in the case that player 1 deviated in period 0) or for a player who deviated by offering $\underline{c}$ to everyone else in period 0. At each situation in $\tilde{H}$ for a given player i , specify that player i choose action $\underline{a}_i$ in the underlying game.

Finally, define $J \equiv \overline{H} \cup \underline{H} \cup \tilde{H}$. Note that the above steps have specified the component s_J of the strategy profile for the grand game. Let $\phi \in \underline{\Delta} S_J$ put probability one on this vector s_J .

Importantly, J is a relatively small subset of H , including only the situations on the prescribed path and on exactly six paths identified with individual players deviating at period 0. Set J does not contain situations that would be reached if there were simultaneous deviations by linked players in period 0 that resulted in them forming a different contract than is prescribed. Also, J does not contain some of the situations that can follow the period-0 deviations featured above, such as, in our example, the one for player 3 where this player's personal history is exactly as expected on the prescribed path, except that in period 2 player 5 had cancelled the contract with her. This situation is not in J because it would not be reached after a period-0 deviation and subsequent play as described above. If, say, player 6 deviated in period 0, then player 5 is prescribed to cancel with player 3 in period 1, not wait until period 2 to do so.

Further, at all situations in J for periods other than 0 and 1, the players do not know whether situations outside of J have already been reached. In fact, for every $h \in J$ that occurs in a period after period 1, it is the case that $S(h)$ is *not* a J -product set.

For instance, consider three situations h_5^0 , h_5^1 , and h_3^2 given as follows and loosely illustrated in Figure 7 (the extensive-form diagram being too complicated to draw). Let h_5^0 be player 5's single situation in period 0. Let h_5^1 be the situation for player 5 in period 1 where player 5 had offered $\underline{c}$ to player 6 in period 0 but made the prescribed offers to the other players (including offering c^{35} to player 3 and c^{45} to player 4), and where the other players sent their prescribed offers to player 5. Finally, let h_3^2 be player 3's situation in period 2 on the prescribed path. Note that $h_5^0 \in J$, $h_5^1 \in -J$, and $h_3^2 \in J$.

Let $\bar{a}^0$ denote the prescribed action for player 5 at h_5^0 , which is to offer c^{56} to player 6, c^{35} to player 3, and c^{45} to player 4. Let $\tilde{a}^0$ denote the action of player 5 at h_5^0 that offers c^{35} to player 3, c^{45} to player 4, and $\underline{c}$ to player 6. Let $\bar{a}^1$ denote the action of player 5 at h_5^1 that cancels with players 3 and 4, and let $\tilde{a}^1$ denote an arbitrary action of player 5 at h_5^1 that does not cancel with player 3.

Clearly, there is a partial strategy profile $\bar{s}_J \in S_J$ that specifies action $\bar{a}^0$ at h_5^0 , along with a partial strategy profile $\bar{s}_{-J} \in S_{-J}$ that specifies action $\bar{a}^1$ at h_5^1 , such that $(\bar{s}_J, \bar{s}_{-J}) \in S(h_3^2)$. Let partial strategy profile $\tilde{s}_J \in S_J$ be identical to $\bar{s}_J$ except that it specifies action $\tilde{a}^0$ at h_5^0 , and likewise let $\tilde{s}_{-J} \in S_{-J}$ be identical to $\bar{s}_{-J}$ except that it specifies action $\tilde{a}^1$ at h_5^1 . We have $(\tilde{s}_J, \tilde{s}_{-J}) \in S(h_3^2)$, because player 5's deviation from the prescribed path in period 0 is not detected by player 3 due to player 5 not cancelling in period 1, and player 3's experience with the other players is the same as on the path of $(\bar{s}_J, \bar{s}_{-J})$. However, $(\tilde{s}_J, \bar{s}_{-J}) \notin S(h_3^2)$, and therefore $S(h_3^2)$ is not a J -product set.¹³

Thus, the information players receive during the game is complicated with respect to combinations of actions at the situations in J and situations in $-J$. Also note that, on various paths through many situations in J , some of the situations in $-J$ are reached later in the game. For instance, at player 3's situation in period 2 on the prescribed path (a member of J), if player 3 were to deviate by cancelling a contract, then the next situation reached by player 3 is not a member of J . Therefore, for analysis of incentives at the situations in J , the players' behavior at the other situations is relevant.

¹³For this demonstration, it suffices to let $\bar{s}_J$ play as described in the partial-equilibrium construction below.




Figure 7: Illustration: $S(h_3^2)$ not a J -product set. Here $h_5^0 \in J$, $h_5^1 \in -J$, and $h_3^2 \in J$.

Partial equilibrium construction

Next let us define the sequence of J -behavior-strategy profiles $(\phi^k)_k$ that determines the J -appraisal system Q . That is, $\phi^k \rightarrow \phi$ and $\phi^k(\cdot | S(h)_J) \rightarrow q^h$ for every $h \in J$. Let ε be a strictly positive number that is sufficiently small to guarantee that the probabilities described next are all feasible. For every $h \in J$ in period 0, let ϕ_h^k put probability ε/k on the action that makes offer $\underline{c}$ to all other players, probability $(\varepsilon/k)^3$ on each other action except the one on the prescribed path, and the remaining probability on the prescribed-path action.

For every $h \in J$ in periods 2 through 5, let ϕ_h^k put probability $(\varepsilon/k)^2$ on every action other than the one prescribed above, and the remaining probability on the prescribed-path action. For example, for any player i whose history is exactly as on the prescribed path, this player chooses to cancel all of her outstanding contracts with probability $(\varepsilon/k)^2$, and, with probability close to 1, to not cancel any contract.

For any $h \in \bar{H}$ for a player i in period 6, let ϕ_h^k put probability ε/k on each action $a \neq a_i^*$ and put the remaining probability on action a_i^* . Finally, for any $h \in \tilde{H}$ for a player i in period 6, let ϕ_h^k put probability ε/k on each action $a \neq \underline{a}_i$ and put the remaining probability on action $\underline{a}_i$.

Confirming partial-equilibrium conditions

By construction, we have $\phi^k \rightarrow \phi$, and $\phi^k(\cdot | S(h)_J)$ converges for every $h \in J$, which defines Q . To understand the J -appraisals and to check sequential r-rationality of the J -appraisal system, consider for example the situation h_5^2 for player 5 in period 2 in which player 5 sees that player 3 has cancelled the contract with her but otherwise player 5 has observed nothing that differs from the prescribed path. Note that $h_5^2 \in \bar{H}$ and, according to q^h , player 5 believes that either player 1 or player 2 (but not both) deviated by sending offer $\underline{c}$ to player 3 in period 0, and that then, in period 1, player 3 cancelled the contracts with both player 5 and either player 1 or player 2 (whichever had not deviated in period 0). Further, player 5 thinks that the behavior of the other players at all other situations in J has or will conform to ϕ , including that players 1, 2, and 3 will choose $\underline{a}_1$, $\underline{a}_2$, and $\underline{a}_3$ in period 6.

Recall that player 5's assurance contracts with players 4 and 6 would then require player 5

to pay them each a penalty of $6\zeta \cdot 3$ for the anticipated underlying-game actions of players 1, 2, and 3. At h_5^2 , player 5 therefore has the incentive to act as $\phi^{h_5^2}$ specifies, which is to immediately cancel these contracts with players 4 and 6, to instead pay the cancellation penalty of 2ζ to each of them. The difference in penalties outweighs any gain player 5 could obtain due to altered actions in the underlying game. Further, player 5 strictly prefers to cancel in period 2 rather than later, because the cancellation penalties rise by at least ζ per period, outweighing any gain player 5 could obtain due to altered actions in the underlying game. These facts depend not on whatever behavior may be specified for the situations in $-J$. Also, $q^{h_5^2} = \phi^{h_5^2}$ by construction. Therefore $q^{h_5^2}$ is r-rational.

The same logic holds for all other situations in $\underline{H}$. The player on the move believes that one of the others deviated by offering $\underline{c}$ to everyone else in period 0, initiating a wave of cancellations that has reached this player, who now strictly prefers to cancel all outstanding contracts. At each situation in $\tilde{H}$, the active player i believes that everyone else will choose their part of action profile $\underline{a}$, and thus $\underline{a}_i$ is r-rational since player i has no outstanding contracts, any cancellation penalties paid are sunk, and $\underline{a}$ is a Nash equilibrium of the underlying game.

Likewise, at situations in $\overline{H}$, the players believe that all past and future play is as specified on the prescribed path, and there is no incentive to deviate. For instance, at such a situation in period 6 for any player i , this player has the strict incentive to choose productive action a_i^* . Because player i has a contract with at least one other, any deviation would trigger payment of assurance penalty 6ζ to at least one other player, outweighing any gain in the underlying game.

Further, at any situation in $\overline{H}$ in periods 1 through 5 for any player i , by continuing to behave as prescribed (not cancelling any contract), the payoff will be $u_i(a^*)$, whereas cancelling triggers a penalty of at least ζ , again outweighing any gains in the underlying game. Note also that cancelling a contract will not cause another player to later cancel with player i , since $\underline{K}$ is a tree (with no loops back to player i).

Finally, in period 0, each player i strictly prefers to behave as prescribed, as it will lead to payoff $u_i(a^*)$. Deviating in any way will lead other players to eventually cancel their contracts. Deviating by offering $\underline{c}$ to everyone else will lead to action profile $\underline{a}$, which is worse for player i . And if player i were to deviate while still forming the prescribed contract with another player, then player i would be on the hook for an assurance penalty or cancellation penalty later.

Applying the theorem

The foregoing analysis confirms that (J, Q, ϕ) is a partial sequential equilibrium. We can also verify that, even though $S(h)$ is not a J -product set for many $h \in J$, in fact (J, Q, ϕ) satisfies the rectangular margin-support condition. For each $h \in J$ in period 0, we have $S(h) = S$ and so trivially $\{s \in S(h) \mid s_J \in \text{supp } q^h\}$ is a J -product set. More substantially, for each $h \in J$ in periods 1 through 6, appraisal q^h puts probability one on play paths that reach only situations in J . That is, at such a situation h , the player on the move believes that prior play conforms to ϕ other than perhaps a single deviation in period 0 (where one player offered $\underline{c}$ to everyone else), which means that the path of play to h has reached only situations in J . With this belief, behavior at situations in $-J$ does not affect whether h is reached, implying that

$\{s \in S(h) \mid s_J \in \text{supp } q^h\}$ is a J -product set.

With (J, Q, ϕ) being a partial sequential equilibrium that satisfies $J \in \mathcal{H}$ and the rectangular margin-support condition, the Theorem applies to establish the existence of a sequential equilibrium (P, σ) satisfying $\sigma_J = \phi$ and $P_J = Q$. The equilibrium path is the prescribed path, in which the efficient action profile a^* is played in the underlying game.

6 Proofs

This section contains proofs of Lemmas 1 and 2 and the Theorem, with the help of two additional lemmas. Each result is restated here.

Lemma 1: *For any given assessment (μ, π) , $P[\mu, \pi]$ is an appraisal system. Further, for each $h \in H_+$, (μ, π) and $P[\mu, \pi]$ induce the same distribution over terminal nodes, and hence the same expected payoffs.*

Proof. Take as given an assessment (μ, π) . To verify $P[\mu, \pi]$ is an appraisal system, note that for any player $i \in N_+$ and situation $h \in H_i$, we have $p^h[\mu, \pi] \in \Delta S(h)$, and $p^h[\mu, \pi]$ exhibits h' -independence for every $h' \in H_i$. To see why the independence claim is true, observe first that $\underline{L}(d) \cap H_i = \underline{L}(d') \cap H_i$ for all $d, d' \in h$ due to perfect recall. Second, observe that for every $d \in h$ and $h' \in \underline{L}(d) \cap H_i$, by definition $\omega_{h'}^d$ puts probability 1 on the action taken at h' on the path to d , and this is the same action for every $d \in h$, again by perfect recall. Then, by construction, $p^h[\mu, \pi]_{h'}$ is degenerate and therefore $p^h[\mu, \pi]$ exhibits h' -independence. Likewise, in the case of $d \in h$ and $h' \in H_i \setminus \underline{L}(d)$, we have that $\omega_{h'}^d = \pi_{h'}$ which, along with ω^d being a behavior-strategy profile, implies that $p^h[\mu, \pi]$ exhibits h' -independence. This proves the first claim.

On the second claim, consider any situation $h \in H_+$ and any node $d \in h$. By perfect recall, no situation can be reached multiple times on any play path. In particular, no situation in $\underline{L}(d)$ could be reached at or after d . By the definition of $P[\mu, \pi]$, for the given h and d , we thus have $\omega_{h'}^d = \pi_{h'}$ for every $h' \in H \setminus \underline{L}(d)$, which includes $h' = h$ and every situation h' that could be reached after h on paths from d to terminal nodes. Therefore $p^h[\mu, \pi]$ captures that π dictates the play from node d .

Continuing with $d \in h \in H_+$, let $\hat{\mu}(d)$ denote the probability of reaching node $d \in h$, conditional on reaching situation h , as given by the implied appraisal system $P[\mu, \pi]$. That is, $\hat{\mu}(d) = \sum_{s \in S(d)} p^h[\mu, \pi](s)$. Using the equality $p^h[\mu, \pi](s) = \sum_{d' \in h} \mu(d') \omega^{d'}(s)$ from the construction of $P[\mu, \pi]$, we have

$$\hat{\mu}(d) = \sum_{s \in S(d)} p^h[\mu, \pi](s) = \sum_{s \in S(d)} \sum_{d' \in h} \mu(d') \omega^{d'}(s).$$

Recall that, by definition, ω^d puts positive probability only on strategy profiles that reach d , and therefore it puts probability zero on strategy profiles that reach other nodes in h . This means that, for any strategy profile $s \in S(d)$, we have $\omega^{d'}(s) = 0$ for every node $d' \in h$ for which $d' \neq d$, implying $\sum_{d' \in h} \mu(d') \omega^{d'}(s) = \mu(d) \omega^d(s)$. As a consequence,

$$\hat{\mu}(d) = \sum_{s \in S(d)} \sum_{d' \in h} \mu(d') \omega^{d'}(s) = \sum_{s \in S(d)} \mu(d) \omega^d(s) = \mu(d) \sum_{s \in S(d)} \omega^d(s).$$

Likewise, $\omega^d(s) > 0$ implies $s \in S(d)$, and so $\sum_{s \in S(d)} \omega^d(s) = \sum_{s \in S} \omega^d(s) = 1$, establishing that $\hat{\mu}(d) = \mu(d)$.

We have shown that the implied appraisal system $P[\mu, \pi]$ specifies, conditional on any situation $h \in H_+$, the same distribution over nodes as does the assessment μ . We have shown also that, from each node in h , $p^h[\mu, \pi]$ specifies continuation play according to π . Thus, (μ, π) and $P[\mu, \pi]$ induce the same distribution over terminal nodes, conditional on any $h \in H_+$. A further implication is that, for any player $i \in N_+$ and situation $h \in H_i$, (μ, π) and $P[\mu, \pi]$ imply the same expected payoff for player i at h , which is $u_i(p^h)$. $\square$

Lemma 2: *Given any appraisal system P and behavior-strategy profile π , (P, π) is fully consistent if and only if there is an assessment (μ, π) that is fully consistent and satisfies $P = P[\mu, \pi]$.*

Proof. Let us first show that, for any fully consistent assessment (μ, π) , the implied appraisal system $P[\mu, \pi]$ is also fully consistent. Let $(\sigma^k)_k$ be a sequence of fully mixed behavior-strategy profiles for which $\sigma^k \rightarrow \pi$ and $\rho[\sigma^k] \rightarrow \mu$. For each positive integer k and decision node $d \in D_+$, let $\omega^d[k]$ be the distribution ω^d in the construction of implied appraisal system $P[\rho[\sigma^k], \sigma^k]$.

From the definition of $\omega^d[k]$ and because σ^k exhibits h -independence for every $h \in H$, we have that $\omega^d[k] = \sigma^k(\cdot | S(d))$. To see why, note that for every $h \in H \setminus \underline{L}(d)$, $S(d) = S(d)_{-h} \times S_h$; and for $h \in \underline{L}(d)$, $S(d) = S(d)_{-h} \times \{\hat{s}_h\}$, where $\hat{s}_h$ is the action taken at h on the path to d . That is, $S(d)$ is an h -product set for every $h \in H$. Since σ^k is fully mixed and exhibits h -independence for every $h \in H$, the conditional probability distribution $\sigma^k(\cdot | S(d))$ is equal to the product of the updated marginals, $\prod_{h \in H} \sigma_h^k(\cdot | S(d)_h)$. The marginal term for each h is exactly as $\omega^d[k]_h$ is defined, and hence $\omega^d[k] = \sigma^k(\cdot | S(d))$.

Because σ^k is fully mixed, the implied belief function is defined and given by $\rho[\sigma^k](d) = \sigma^k(S(d) | S(h))$ for every $d \in h \in H_+$. From the definition of the implied appraisal system, for each $s \in S$, we have

$$p^h[\rho[\sigma^k], \sigma^k](s) = \sum_{d \in h} \rho[\sigma^k](d) \cdot \omega^d[k](s) = \sum_{d \in h} \sigma^k(S(d) | S(h)) \cdot \sigma^k(s | S(d)).$$

For any given $s \in S$, if $s \in S(h)$ then there is exactly one node $d \in h$ for which $s \in S(d)$; call this node $\hat{d}$. This means that $\sigma^k(s | S(d')) = 0$ for every other node $d' \in h$, and we obtain

$$p^h[\rho[\sigma^k], \sigma^k](s) = \sigma^k(S(\hat{d}) | S(h)) \sigma^k(s | S(\hat{d})) = \sigma^k(s | S(h)).$$

Otherwise, $s \notin S(h)$ and $\sigma^k(s | S(d)) = 0$ for every $d \in h$, and $\sigma^k(s | S(h)) = 0$ as well. Thus, $p^h[\rho[\sigma^k], \sigma^k](s) = \sigma^k(s | S(h))$ is an identity in s .

Let ω^d be the distribution used in the construction of $P[\mu, \pi]$, for every decision node d , and note that $\omega^d[k] \rightarrow \omega^d$ due to $\sigma^k \rightarrow \pi$. Using also that $\rho[\sigma^k] \rightarrow \mu$, from the definition of implied appraisal system we conclude that $p^h[\rho[\sigma^k], \sigma^k] \rightarrow p^h$, for every $h \in H_+$.

Summarizing, we started by presuming that (μ, π) is fully consistent, and we derived that $\sigma^k \rightarrow \pi$, $P[\rho[\sigma^k], \sigma^k] \rightarrow P[\mu, \pi]$, and $p^h[\rho[\sigma^k], \sigma^k] = \sigma^k(\cdot | S(h))$ for every $h \in H_+$ and positive integer k . This means that $P[\mu, \pi]$ is fully consistent.

If, alternatively, we start with an appraisal system P and behavior-strategy profile π for which (P, π) is fully consistent, then let $(\sigma^k)_k$ denote a sequence of fully mixed behavior-strategy profiles satisfying $\sigma^k \rightarrow \sigma$ and $\sigma^k(\cdot | S(h)) \rightarrow p^h$ for every $h \in H_+$. The sequence $(\sigma^k)_k$ can be used to define a sequence of assessments $(\rho[\sigma^k], \sigma^k)$. Because the set of decision nodes is finite, there is a convergence subsequence, the limit being a fully consistent assessment (μ, π) . Repeating the steps above establishes that $P[\mu, \pi] = P$. $\square$

Lemma 3: *Take as given a sequence of fully mixed behavior-strategy profiles $(\sigma^k)_k$, sets Y and Z satisfying $Y \subset Z \subset S$, and a distribution $\pi \in \Delta S$. If $\sigma^k(\cdot | Z) \rightarrow \pi$ and $\text{supp } \pi \subset Y$, then $\sigma^k(\cdot | Y) \rightarrow \pi$.*

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

$$\begin{aligned}
\pi(X) &= \lim_{k \rightarrow \infty} \sigma^k(X | Z) = \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Z)}{\sigma^k(Z)} \\
&= \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Y)}{\sigma^k(Z)} + \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Z \setminus Y)}{\sigma^k(Z)} \\
&= \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Y)}{\sigma^k(Y)} \cdot \frac{\sigma^k(Y)}{\sigma^k(Z)} + \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Z \setminus Y)}{\sigma^k(Z)} \\
&= \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Y)}{\sigma^k(Y)} \cdot \lim_{k \rightarrow \infty} \frac{\sigma^k(Y)}{\sigma^k(Z)} + \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Z \setminus Y)}{\sigma^k(Z)} \\
&= \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Y)}{\sigma^k(Y)} \cdot \pi(Y) + \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Z \setminus Y)}{\sigma^k(Z)} \\
&= \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Y)}{\sigma^k(Y)} = \lim_{k \rightarrow \infty} \sigma^k(X | Y).
\end{aligned}$$

These equalities hold for the following reasons, in order: the presumption of the lemma, the definition of conditional probability, separation of $X \cap Z$ into disjoint sets and the two limits existing because they define $\pi(X \cap Y)$ and $\pi(X \cap Z \setminus Y)$, algebra and $\sigma^k(Y) > 0$ because σ^k is fully mixed, the product rule for limits and because $\sigma^k(Y | Z)$ converges, that $\sigma^k(Y | Z) \rightarrow \pi(Y)$, that $\pi(Y) = 1$ and $\sigma^k(X \cap Z \setminus Y | Z) \rightarrow \pi(X \cap Z \setminus Y | Z) = 0$, and the definition of conditional probability. $\square$

Lemma 4: *Take as given a sequence of fully mixed behavior-strategy profiles $(\sigma^k)_k$, a set $J \subset H$, a J -product set $Y \subset S$, a distribution $\pi \in \Delta S$, and a distribution $\psi \in \Delta S_J$. Let $\beta^k \equiv \sigma_J^k$ denote the marginal of σ^k on component J . If $\beta^k(\cdot | Y_J) \rightarrow \psi$ and $\sigma^k(\cdot | Y) \rightarrow \pi$, then π exhibits J -independence and $\pi_J = \psi$.*

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

$$\begin{aligned}
\pi(X) &= \lim_{k \rightarrow \infty} \sigma^k(X | Y) = \lim_{k \rightarrow \infty} \frac{\sigma^k(X \cap Y)}{\sigma^k(Y)} \\
&= \lim_{k \rightarrow \infty} \frac{\sigma_J^k(X_J \cap Y_J) \cdot \sigma_{-J}^k(X_{-J} \cap Y_{-J})}{\sigma_J^k(Y_J) \cdot \sigma_{-J}^k(Y_{-J})} \\
&= \lim_{k \rightarrow \infty} \frac{\sigma_J^k(X_J \cap Y_J)}{\sigma_J^k(Y_J)} \cdot \lim_{k \rightarrow \infty} \frac{\sigma_{-J}^k(X_{-J} \cap Y_{-J})}{\sigma_{-J}^k(Y_{-J})}. \quad (1)
\end{aligned}$$

The third equality is due to σ^k being a behavior-strategy profile and X and Y being J -product sets, implying that $(X \cap Y)_J = X_J \cap Y_J$ and $(X \cap Y)_{-J} = X_{-J} \cap Y_{-J}$. The fourth equality is due to the limit existing, that $\beta^k \equiv \sigma_J^k$, and $\beta^k(\cdot | Y_J) \rightarrow \psi$. Note that the first limit term in the final expression is $\psi(X_J)$. Note that

$$\begin{aligned}
\pi_J(X_J) &= \pi(X_J \times S_{-J}) = \pi((X_J \cap Y_J) \times Y_{-J}) = \lim_{k \rightarrow \infty} \frac{\sigma^k((X_J \cap Y_J) \times (Y_{-J}))}{\sigma^k(Y)} \\
&= \lim_{k \rightarrow \infty} \frac{\sigma_J^k(X_J \cap Y_J) \cdot \sigma_{-J}^k(Y_{-J})}{\sigma_J^k(Y_J) \cdot \sigma_{-J}^k(Y_{-J})} = \lim_{k \rightarrow \infty} \frac{\sigma_J^k(X_J \cap Y_J)}{\sigma_J^k(Y_J)}.
\end{aligned}$$

The first equality is from the definition of marginal probability, the second because Y contains the support of π , the third by presumption, and the fourth because σ^k is a behavior-strategy profile. Applying the same steps for $-J$ yields $\pi_{-J}(X_{-J}) = \lim_{k \rightarrow \infty} \sigma_{-J}^k(X_{-J} \cap Y_{-J}) / \sigma_{-J}^k(Y_{-J})$.

Combining these identities with expression 1 yields $\pi(X) = \pi_J(X_J) \cdot \pi_{-J}(X_{-J})$ for the J -product set X , proving that π exhibits J -independence. That $p_J = \psi$ follows from evaluating $\pi(X_J \times S_{-J})$ for any $X_J \subset S_J$. $\square$

Theorem: For any partial sequential equilibrium (J, Q, ϕ) of a finite extensive-form game, if (J, Q, ϕ) satisfies $J \in \mathcal{H}$ and the rectangular margin-support condition, then (J, Q, ϕ) extends to a sequential equilibrium.

Proof. Take as given a finite game in the extensive form, any set of situations $J \subset H$, a J -appraisal system Q , and a J -behavior-strategy profile $\phi \in \underline{\Delta} S_J$. Suppose $J \in \mathcal{H}$, the rectangular margin-support condition is satisfied, and (J, Q, ϕ) is a partial sequential equilibrium. We must show that there exists a sequential equilibrium (P, σ) with the property that $\sigma_J = \phi$ and $p_J^h = q^h$ for every $h \in J_+$.

Let $(\phi^k)_k$ be a sequence of fully mixed J -behavior-strategy profiles with the properties $\phi^k \rightarrow \phi$ and $\phi^k(\cdot | S(h)_J) \rightarrow q^h$ for every $h \in J_+$. For every $h \in J_+$, define

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

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

Note that for every $h \in J_+$, $\phi^k(Y_J^h | S(h)_J)$ converges to 1, since $\phi^k(\cdot | S(h)_J) \rightarrow q^h$ and $Y_J^h = \text{supp } q^h$. This implies that $\phi^k(S(h)_J \setminus Y_J^h) / \phi^k(Y_J^h)$ converges to 0. We can presume that $(\phi^k)_k$ is defined so that

$$\frac{\phi^k(S(h)_J \setminus Y_J^h)}{\phi^k(Y_J^h)} < \left(\frac{1}{|A|} \right)^{2|H|} \quad (2)$$

for every positive integer k and every $h \in J_+$, because this must hold for k large enough anyway. (Recall that A is the action space of the game, and so $|A|$ is the finite number of different actions. Likewise, $|H|$ is the finite number of situations.) Define sequence $(\delta^k)_k$ by:

$$\delta^k \equiv \max_{h \in J_+} \left(\frac{\phi^k(S(h)_J \setminus Y_J^h)}{\phi^k(Y_J^h)} \right)^{\frac{1}{2}}.$$

For each k , we can define an artificial game in the Bayesian agent-normal form as follows. Chance makes the choices at all situations in J , and Chance's strategy is ϕ^k . This means that at each situation $h \in J$, Chance takes the action according to distribution ϕ_h^k , which is independent of the choices made at the other situations due to ϕ^k being a behavior-strategy profile. In the artificial game, strategic player $i \in N_+$ makes the choices at situations in $H_i \setminus J_i$. Player i is restricted to put probability of at least $(\delta^k)^{1/|H|}$ on each available action, and player i 's random choices are independent across situations. Since $|H|$ is the number of situations in the game, the probability of any particular $s_{-J} \in S_{-J}$ is at least $((\delta^k)^{1/|H|})^{|H|} = \delta^k$. Condition 2 implies $|A|(\delta^k)^{1/|H|} < 1$, ensuring that it is feasible to put at least the probability bound on each action at each situation. Also, recall that $H_0 \subset J$ and therefore every situation in $-J$ is one for a strategic player.

Standard arguments and use of Kakutani's fixed-point theorem imply that the artificial game has a Nash equilibrium in fully mixed behavior-strategy profiles, $\theta^k \in \underline{\Delta}S_{-J}$. Let $\sigma^k \in \underline{\Delta}S$ be defined by $\sigma^k(s) \equiv \phi^k(s_J)\theta^k(s_{-J})$ for every $s \in S$. Note that σ^k is a fully mixed behavior-strategy profile.

Because the pure-strategy space in the artificial games is finite, and the space of feasible behavior-strategy profiles in the artificial games converges to $\underline{\Delta}S_{-J}$ as $k \rightarrow \infty$, we can find a subsequence $(\sigma^{k^m})_m$ and profile $\sigma \in \underline{\Delta}S$ with the properties that $\sigma^{k^m} \rightarrow \sigma$ and that $(\sigma^{k^m}(\cdot | S(h)))_m$ converges for every $h \in H$. Letting $p^h \equiv \lim_{m \rightarrow \infty} \sigma^{k^m}(\cdot | S(h))$ for every $h \in H$, we have thus constructed an appraisal system P and a behavior-strategy profile σ such that (P, σ) is fully consistent.

Because σ^{k^m} is fully mixed, every situation is reached. This implies that, for every $h \in -J$, in the artificial game $\sigma^{k^m}(\cdot | S(h))$ is rational at h . Rationality here is defined subject to the constraint that each action at h is allocated a probability of at least δ^{k^m} . By continuity of payoffs in the action probabilities, that $\delta^{k^m} \rightarrow 0$, and because $\sigma^{k^m}(\cdot | S(h)) \rightarrow p^h$, it must be that p^h is rational at h in the original game, for each $h \in -J$.

We next must establish rationality for the situations in J_+ , and show that for each $h \in J_+$, the appraisal at h satisfies $p_J^h = q^h$. Take any $h \in J_+$. From the definition of Y^h and knowing it is a J -product set owing to the rectangular margin-support condition, $s \in S(h) \setminus Y^h$ implies $s_J \notin Y_J^h$. Therefore $S(h) \setminus Y^h \subset (S(h)_J \setminus Y_J^h) \times S_{-J}$, and we have

$$\sigma^k(S(h) \setminus Y^h) \leq \sigma^k((S(h)_J \setminus Y_J^h) \times S_{-J}) = \phi^k(S(h)_J \setminus Y_J^h). \quad (3)$$

Likewise, we have

$$\sigma^k(Y^h) \geq \phi^k(Y_J^h)\delta^k. \quad (4)$$

This is because $\sigma^k(s) \equiv \phi^k(s_J)\theta^k(s_{-J})$ for every $s \in S$, our construction ensures $\theta^k(s_{-J}) > \delta^k$ for every $s_{-J} \in S_{-J}$, and Y^h is a J -product set.

Combining inequalities 3 and 4, and recalling that $\delta^k \geq (\phi^k(S(h)_J \setminus Y_J^h) / \phi^k(Y_J^h))^{1/2}$, we obtain

$$\begin{aligned} \frac{\sigma^k(S(h) \setminus Y^h)}{\sigma^k(Y^h)} &\leq \frac{\phi^k(S(h)_J \setminus Y_J^h)}{\phi^k(Y_J^h) \cdot \delta^k} \\ &\leq \frac{\phi^k(S(h)_J \setminus Y_J^h)}{\phi^k(Y_J^h)} \cdot \left(\frac{\phi^k(Y_J^h)}{\phi^k(S(h)_J \setminus Y_J^h)} \right)^{\frac{1}{2}} = \left(\frac{\phi^k(S(h)_J \setminus Y_J^h)}{\phi^k(Y_J^h)} \right)^{\frac{1}{2}}. \end{aligned}$$

The last expression on the right side converges to 0 as $k \rightarrow \infty$, because $\phi^k(S(h)_J \setminus Y_J^h) / \phi^k(Y_J^h)$ converges to 0. Therefore $\sigma^k(S(h) \setminus Y^h) / \sigma^k(Y^h) \rightarrow 0$, which implies that $\sigma^k(Y^h) / \sigma^k(S(h)) \rightarrow 1$ since $Y^h \subset S(h)$.

Recalling that $\sigma^{k^m}(\cdot | S(h)) \rightarrow p^h$ as $m \rightarrow \infty$, we have thus established that $\text{supp } p^h \subset Y^h$. Since $Y^h \subset S(h)$, Lemma 3 then implies that $\sigma^{k^m}(\cdot | Y^h) \rightarrow p^h$ as well. Further, noting that $\phi^k = \sigma_J^k$, we can apply Lemma 4 to the subsequence $(\sigma^{k^m})_m$ and set Y^h , which we know is a J -product set from the rectangular margin-support condition, where $\sigma^{k^m}(\cdot | S(h)) \rightarrow p^h$ and $\phi^{k^m}(\cdot | S(h)_J) \rightarrow q^h$. Lemma 4 establishes that p^h exhibits J -independence and $p_J^h = q^h$. This combined with the presumption that Q is sequentially r -rational (because (J, Q, ϕ) is assumed to be a partial sequential equilibrium) implies that P is rational at every $h \in J_+$.

In summary, we have shown that (P, σ) is fully consistent and that P is sequentially rational. Therefore, P is a sequential equilibrium. We have also shown that $p_J^h = q^h$ for every $h \in J_+$, completing the proof. $\square$

Proposition 1: *For any finite multistage game with observed actions and independent types, every AMOD-minimal partial sequential equilibrium extends to a sequential equilibrium, and every sequential equilibrium is an extension of an AMOD-minimal partial sequential equilibrium.*

Proof. On the first claim, take as given any AMOD-minimal partial sequential equilibrium (J, Q, ϕ) . Consider any personal history $h = (b, i, \theta_i) \in J$. Note that no situation $h' = (b', j, \theta_j) \in -J$ could have been reached on the path to h , because in such a case b' contains a simultaneous public deviation and is a truncated version of b , implying that b contains a simultaneous public deviation, which contradicts the AMOD condition. Therefore, whether h is reached cannot depend on actions taken at situations in $-J$. This means that $S(h) = S(h)_J \times S_{-J}$, which trivially implies the rectangular margin-support condition. The Theorem then establishes the existence of a sequential equilibrium to which (J, Q, ϕ) extends.

On the second claim, take as given any sequential equilibrium (P, σ) . Define $L_i \equiv H_i \cup \{h_{0i}\}$ for each $i \in N_+$. As noted above, perfect recall and observed actions imply

$$S(b, i, \theta_i) = S(b, i, \theta_i)_{L_i} \times \left(\times_{j \in N_+ \setminus \{i\}} S(b)_{L_j} \right),$$

for every $b \in B$, $i \in N_+$, and $\theta_i \in \Theta_i$. Further, an implication of this product structure and full consistency is that there exists $\Lambda = (\lambda_j^b)_{b \in B, j \in N_+}$ such that $\lambda_j^b \in \Delta S_{L_j}$ for every $b \in B$ and $j \in N_+$, and for which the following holds: For every $(b, i, \theta_i) \in H_+$ and $j \in N_+ \setminus \{i\}$, it is the case that $p^{(b, i, \theta_i)}$ exhibits L_j -independence and $p_{L_j}^{(b, i, \theta_i)} = \lambda_j^b$. In words, for

public history b , there is a “public belief” about the strategy and type of each strategic player, and player i ’s appraisal specifies exactly the public belief about each of the other players, exhibiting independence across them.

We can define the AMOD property in relation to P just as we did for Q above. That is, let us say that (b, b') has the AMOD property for P if b and b' are consecutive public histories and $\lambda_j^b(S(b')_{L_j}) = 0$ for at most one $j \in N_+$. We can then find the set $J \in \hat{\mathcal{H}}$ that is minimal with respect to the following conditions: $(\emptyset, i, \theta_i) \in J$ for every $i \in N_+$ and $\theta_i \in \Theta_i$, and $[(b, b') \text{ has the AMOD property for } P \text{ and } (b, i, \theta_i) \in J] \text{ implies } (b', i, \theta_i) \in J$. Thus, J is the set of personal histories that entail no simultaneous public deviations, in the sequential equilibrium (P, σ) .

Define J -behavior-strategy profile $\phi \equiv \sigma_J$ and define J -appraisal system Q by $q^h \equiv p^h$ for every $h \in J_+$. Let $(\sigma^k)_k$ be a sequence of fully mixed behavior-strategy profiles satisfying $\sigma^k \rightarrow \sigma$ and $\sigma^k(\cdot | S(h)) \rightarrow p^h$ for every $h \in H_+$, guaranteed to exist by the full-consistency condition of sequential equilibrium. Consider any personal history $h = (b, i, \theta_i) \in J$. For the same reason provided in the first paragraph of this proof, we have $S(h) = S(h)_J \times S_{-J}$. As noted above, updating from a fully mixed prior distribution satisfying independence across a component of the strategy profile, conditional on an event that is conjunctive with respect to this component, is equivalent to calculating the product of the updated marginals. This implies that $q^h = p^h = \lim_{k \rightarrow \infty} \sigma_J^k(\cdot | S(h)_J)$. Noting also that $(\sigma_J^k)_k \rightarrow \sigma_J = \phi$, we have established that (J, Q, ϕ) is fully consistent.

In fact, also Q must be sequentially r-rational. To see this, note first that for every $h \in J_+$, appraisal p^h is rational at h . The rationality condition examines any h -deviation from p^h , under the presumption that play will otherwise conform to p^h_{-h} , meaning that the deviation at h is the only deviation from Λ in the continuation of the game from h . Second, notice that the expected payoff vectors of p^h and of h -deviations do not depend on p^h_{-J} because p^h_{-h} does not reach any situations in $-J$ (reaching any of them requires simultaneous public deviations). This implies that q^h is r-rational for every $h \in J_+$, and so Q is sequentially r-rational.

Finally, we see that, for P and our constructed J and Q , it must be that (J, Q, ϕ) is AMOD-minimal. This follows from an induction argument starting with the personal histories of the form (b, i, θ_i) for $b = \emptyset$. Note that for such a personal history and any $b' \in B$ such that b and b' are consecutive public histories, $\lambda_j^b(S(b')_{L_j}) = 0$ if and only if $\lambda_j^b(S(b')_{L_j(J)}) = 0$, because the former expression is the probability that player j will choose the action a_j specified by b' in the first stage, and in the first stage player j acts at a situation in J_j . We can repeat the argument for every public history b at the beginning of the second stage that entailed no simultaneous public deviation in the first stage, and continue to the later stages. In each step, we find that the situations in J for each stage are exactly those that would be included in the construction using the public beliefs $R = (r_j^b)_{b \in B, j \in N_+}$ for which $q_{L_j(J)}^{(b, i, \theta_i)} = r_j^b$ and $q_{L_j(J)}^{(b, i, \theta_i)} = p_{L_j(J)}^{(b, i, \theta_i)}$ for every $b \in B$ and $j \in N_+$.

In summary, we have found an AMOD-minimal partial sequential equilibrium (J, Q, ϕ) that extends to (P, σ) . This completes the proof of the second claim. $\square$

References

- Battigalli, P. (1996): “Strategic Independence and Perfect Bayesian Equilibria,” *Journal of Economic Theory* 70: 201–234.
- Battigalli, P. (1997): “On Rationalizability in Extensive Form Games,” *Journal of Economic Theory* 74: 40–61.
- Battigalli, P., R. Corrao, and M. Dufwenberg (2019): “Incorporating Belief-Dependent Motivation in Games,” *Journal of Economic Behavior and Organization* 167: 185–218.
- Battigalli, P. and N. De Vito (2021): “Beliefs, Plans, and Perceived Intentions in Dynamic Games,” *Journal of Economic Theory* 195: 105283.
- Battigalli, P. and N. Generoso (2024): “Information Flows and Memory in Games,” *Games and Economic Behavior* 145: 356–376.
- Govindan, S. and R. Wilson (2009): “On Forward Induction,” *Econometrica* 77 (1, January): 1–28.
- Hendon, E., H. J. Jacobsen, and B. Sloth (1996): “The One-Deviation Principle for Sequential Rationality,” *Games and Economic Behavior* 12: 274–282.
- Kreps, D. and R. Wilson (1982): “Sequential Equilibrium,” *Econometrica* 50 (4): 863–894.
- Mailath, G., L. Samuelson, and J. Swinkels (1993): “Extensive Form Reasoning in Normal Form Games,” *Econometrica* 61 (March): 273–302.
- Osborne, M. J. and A. Rubinstein (1994): *A Course in Game Theory* Cambridge: MIT press.
- Pearce, D. G. (1984), “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica* 52 (July): 1029–1050.
- Perea, A. (2002): “A Note on the One-Deviation Property in Extensive-Form Games,” *Games and Economic Behavior* 40: 322–338.
- Shimoji, M. and J. Watson (1998): “Conditional Dominance, Rationalizability, and Game Forms,” *Journal of Economic Theory* 83: 161–195.
- Watson, J. (2025): “Perfect Bayesian Equilibrium: Consistency Conditions for Practical Definitions,” unpublished draft, UC San Diego.
- Watson, J. (2024): “Contractual Chains,” *Econometrica* 92 (No. 5, September): 1735–1774.