

Perfect Conditional ε -Equilibria of Multi-Stage Games with Infinite Sets of Signals and Actions* (Preliminary and Incomplete)

Roger B. Myerson and Philip J. Reny

Department of Economics
University of Chicago

Abstract

Abstract: We extend Kreps and Wilson’s concept of sequential equilibrium to games where the sets of actions that players can choose and the sets of signals that players may observe are infinite. A strategy profile is a conditional ε -equilibrium if, for any player and for any of his positive probability signal events outside a uniformly unlikely set, the player’s conditional expected utility would be within ε of the best that the player could achieve by deviating. Perfect conditional ε -equilibria are defined by testing conditional ε -rationality also under nets of small perturbations of the players’ strategies and of nature’s probability function that can make any finite collection of signals have positive probability. Every perfect conditional ε -equilibrium strategy profile is a subgame perfect ε -equilibrium and admits a finitely consistent conditional belief system that makes it sequentially ε -rational. Nature’s perturbations can produce equilibria that seem unintuitive and so we consider two ways to limit the effects of those perturbations, using topologies on nature’s states and on players’ actions.

1 Introduction

We define perfect conditional ε -equilibrium and perfect conditional equilibrium distributions for multi-stage games with infinite signal sets and infinite action sets and prove their existence for a large class of games.

Kreps and Wilson (1982) defined sequential equilibrium for finite games in which nature’s states have positive probability (henceforth *Kreps-Wilson finite games*). But rigorously defined extensions to infinite games have been lacking. Various formulations of “perfect Bayesian equilibrium” (defined for finite games in Fudenberg and Tirole 1991) have been

*Current Version: December 18, 2017. We thank John Hillas, and Asher Wolinsky for helpful comments. Reny thanks the National Science Foundation (SES-1724747, SES-1227506, SES-0922535) for financial support.

used for infinite games, but no general existence theorem for infinite games is available. (See Watson 2017 for an interesting recent contribution.)

Harris, Stinchcombe and Zame (2000) provided important examples that illustrate some of the difficulties that arise in infinite games and they also introduced a methodology for the analysis of infinite games by way of nonstandard analysis, an approach that they showed is equivalent to considering limits of a class of sufficiently rich sequences (nets, to be precise) of finite game approximations.

It may seem natural to try to define sequential equilibria of an infinite game by taking limits of sequential equilibria of finite games that approximate it. The difficulty is that no general definition of “good finite approximation” has been found. Indeed, it is easy to define sequences of finite games that seem to be converging to an infinite game (in some sense) but have limits of equilibria that seem wrong (e.g., examples 2.2 and 2.3 below).

Instead, we work directly with the infinite game itself. We define a strategy profile to be a conditional ε -equilibrium on a domain of signals if signals outside the domain are uniformly unlikely and, given the strategies of the other players, each player’s continuation strategy is ε -optimal conditional on any positive probability set of signals in that domain.

A central challenge in extending the definition of sequential equilibrium strategies from Kreps-Wilson finite games to infinite games is how to test whether the players’ behavior is rational off the equilibrium path of play. In Kreps-Wilson finite games, rationality is tested by checking whether, for any $\varepsilon > 0$, there is an arbitrarily close strictly mixed strategy profile that is a conditional ε -equilibrium on the entire signal domain. In any Kreps-Wilson finite game, any strategy profile that passes this test is a sequential equilibrium strategy profile (and conversely).

In infinite games, there are two serious difficulties with this approach to testing rational behavior at unreached signals. The first difficulty is that, with uncountably-infinite signal spaces, we cannot make all signal events have positive probability at the same time, no matter how we perturb the players’ strategies. The second difficulty is that, with uncountably many states of nature, nature’s probability function must give all but countably many states probability zero. So, even if rationality could be tested by perturbing the players’ strategies, the players’ resulting conditional probabilities over histories would be biased so as to explain, whenever possible, any probability zero event as being the result of a deviation by some player instead of perhaps being the result of the occurrence of a state of nature that had prior probability zero.

We test for rational behavior by considering, for every possible *finite* set of signals in the game, a slightly perturbed strategy profile *and* a slightly perturbed probability function for nature that, together, give every signal in the finite set positive probability. By considering

a separate perturbation for each finite set of signals, we sidestep the first difficulty of simultaneously giving uncountably many signals positive probability. By perturbing nature’s probability function in addition to perturbing the players’ strategies, we avoid the second difficulty because the likelihood of a deviation from the given strategy profile relative to the occurrence of some probability zero state of nature is endogenously determined by their relative likelihood under the combined perturbation of the strategy profile and of nature’s probability function.

A strategy profile is defined to be a perfect conditional ε -equilibrium if there is a signal domain outside of which signals are uniformly unlikely, and, for any finite set of signals in the game, there is an arbitrarily small perturbation of nature’s probability function and there is an arbitrarily small perturbation of the strategy profile such that, in the game with nature’s perturbed probability function, the perturbed strategy profile gives every signal in the finite set positive probability and is a conditional ε -equilibrium on the given signal domain. A perfect conditional equilibrium distribution is defined as the limit of perfect conditional ε -equilibrium distributions on outcomes as $\varepsilon \rightarrow 0$.¹

For any perfect conditional ε -equilibrium strategy profile, it is shown that there is a *finitely consistent conditional belief system* that makes the strategy profile *sequentially ε -rational*, where these latter two concepts attempt to generalize to infinite games the concepts of consistency of beliefs and sequential rationality introduced in Kreps and Wilson (1982). But in infinite games, finitely consistent conditional belief systems are Bayes-consistent only when considering finite partitions of a player’s signal space. For example, it can happen that a given strategy profile gives probability 1/2 to some history, but a player’s finitely consistent conditional belief system puts conditional probability 1 on that history no matter which of his infinitely many possible signals he observes. In addition, the beliefs in a finitely consistent conditional belief system may be only finitely additive, and they can therefore make a strategy profile sequentially ε -rational even though that strategy profile is not a perfect conditional ε -equilibrium. For these reasons, conditional belief systems do not play a central role here and we instead focus on perfect conditional ε -equilibrium strategy profiles and their limit distributions.

Perfect conditional ε -equilibria and perfect conditional equilibrium distributions are shown to exist for a large class of regular projective games. Examples illustrate properties of these

¹We use the term “perfect” to indicate that behavior is tested for rationality *everywhere* (i.e., at every event both on and off the equilibrium path). But because we allow the players to choose strategies that are conditionally only ε -optimal, we do not, in particular, rule out the use of weakly dominated strategies (which Kreps and Wilson 1982 also do not rule out).

Simon and Stinchcombe (1995) and Bajoori, Flesch, and Vermuelen (2013, 2016) use a topological full-support condition in defining, for infinite normal form games, solutions that they call “perfect.” But “perfect” in English comes from a Latin word meaning “complete,” and so it seems more appropriate for the condition of testing rationality everywhere.

solutions and the difficulties of alternative approaches to the problem of extending sequential equilibrium to infinite games.

The solutions defined here have attractive properties in addition to existence in a large class of games. First, every perfect conditional ε -equilibrium strategy profile is a subgame perfect ε -equilibrium. Second, for any finitely consistent conditional belief system, after any history of play, if any two players have the same information then they must have the same beliefs, a property that holds in any Kreps-Wilson consistent assessment for any Kreps-Wilson finite game. This belief-system property is significant because, as already mentioned, any perfect conditional ε -equilibrium strategy profile induces a finitely consistent conditional belief system that makes the strategy profile sequentially ε -rational. Finally, in any Kreps-Wilson finite game, limits of perfect conditional ε -equilibria as $\varepsilon \rightarrow 0$ are sequential equilibrium strategy profiles and the perfect conditional equilibrium distributions defined here are precisely the distributions over outcomes that arise from sequential equilibria.

The consideration of arbitrary perturbations of nature's probability function allows significant flexibility in the players' updated conditional probabilities after the occurrence of states with prior probability zero, and this can produce perfect conditional ε -equilibria that may seem unintuitive.² It can therefore be useful to limit these effects in applications. We provide two refinements for this purpose, one direct, the other indirect. Both refinements require topological structures to be added to the game.

The direct method is to restrict the class of perturbations of nature. For example, if two states of nature are chosen independently, one might insist that they be chosen independently also in any perturbation of nature. Such a restriction could be interpreted to mean that it is common knowledge among the players that neither state is ever informative about the other. One might also wish to ensure that nature's chosen state is never perturbed very far. We therefore introduce a restricted class of "local" perturbations of nature that require the introduction of a topology on nature's states. Local perturbations of nature try to ensure that the players' signals remain approximately as informative as they are under nature's unperturbed probability function. Perfect conditional ε -equilibria exist in regular projective games even when nature's perturbations are restricted to being local.

The indirect method is to require that the perfect conditional ε -equilibrium strategy profile itself should place positive probability on each event in some large class, so that perturbations (by nature or by any player) are irrelevant for verifying strategic rationality in all of these positive probability events. To operationalize this indirect approach, we need to topologize the players' action spaces. Specifically, each player's set of actions at each date is assumed to be a separable metric space. Then, a strategy profile can be defined to have

²This issue does not arise in games in which nature's probability function puts strictly positive probability on every feasible state after each history.

full support if, for each player and for each of his possible signals at any date, the player’s strategy gives positive probability to every nonempty open subset of his available actions at that date given this signal.³ For any $\varepsilon > 0$, say that a conditional ε -equilibrium is *full* if it has full support. Full perfect conditional ε -equilibria are shown to exist in any regular projective game.

We emphasize ε -equilibria where players are only ε -rational because the $\varepsilon \rightarrow 0$ limits of such strategies may fail to satisfy strategic independence when signal sets are infinite since, as first noted by Börgers (1991), players can be coordinated by infinitesimal details of their signals (strategic entanglement; see Example 2.1).⁴

The remainder of the paper is organized as follows. Section 2 provides a number of examples that motivate our main solution. Section 3 introduces the multi-stage games that we study and provides the notation and concepts that are required for the definition of perfect conditional ε -equilibrium, which is given in Section 4. Section 5 establishes that every perfect conditional ε -equilibrium strategy profile is a subgame perfect ε -equilibrium. Section 6 introduces related concepts such as conditional belief systems and their finite consistency, and sequential ε -rationality. Section 7 applies our definitions to several examples. Section 8 introduces the class of “regular projective games” for which can prove existence. Sections 9 and 10 each provide a topologically-based refinement that limits the effects of nature’s perturbations. All proofs are in Section 12, and an appendix computes perfect conditional ε -equilibria for the motivating examples considered in Section 2.

2 Motivating Examples

In this section, we consider four examples that motivate our solution. Our first example illustrates the need to focus on ε -optimality rather than on exact optimality. The difficulty with exact optimality arises through a phenomenon that we may call “strategic entanglement,” where a sequence of strategy profiles yields a path of randomized play that includes histories with fine details used by later players to correlate their independent actions. When these fine details are lost in the limit because the limit path does not include them, there may be no strategy profile that produces the limit distribution over outcomes.⁵ In our example, this problem is so severe that it precludes the existence of a subgame perfect equilibrium, which motivates our choice to focus attention on ε -optimal strategies and the distributions

³Full support strategy profiles exist because, by separability, each action in a countable dense set can be given positive probability.

⁴See Radner (1980) for a study of ε -rationality in finitely repeated games.

⁵Milgrom and Weber (1985) provided the first example of this kind. The example given here has the stronger property that strategic entanglement is unavoidable: it occurs along any sequence of subgame perfect ε -equilibria (i.e., ε -Nash in every subgame) as ε tends to zero.

over outcomes that they induce in the limit as $\varepsilon \rightarrow 0$.

Example 2.1 *Strategic entanglement in limits of approximate equilibria (Harris-Reny-Robson 1995).*

- On date 1, player 1 chooses $x \in [-1, 1]$ and player 2 chooses $y \in \{-1, 1\}$.
- On date 2, players 3 and 4 each observe (x, y) ; then 3 chooses $z \in \{-1, 1\}$ and 4 chooses $w \in \{-1, 1\}$.
- The payoffs of players 3 and 4 are, $u_3 = xz$ and $u_4 = xw$, so players 3 and 4 each want to match the sign of player 1's choice.
- Player 2's payoff is $u_2 = (y + 1)z$, so player 2 would like to match 3's choice, and 2's payoff (0) when she chooses $y = -1$ is never equal to her payoff (-2 or $+2$) when she chooses $y = +1$, regardless of the others' actions.
- Player 1's payoff is the sum of three terms:
 - (first term) if 3 and 4 match he gets 0, if they mismatch he gets -10 ;
 - plus (second term) if 2 and 3 match he gets $-|x|$, if they mismatch he gets $|x|$;
 - plus (third term) he gets $-|x|^2$.

As shown in Harris et. al. (1995), there is no subgame-perfect equilibrium of this game. But there is an obvious solution that is the limit of strategy profiles where everyone's strategy is arbitrarily close to optimal.

For any $\varepsilon > 0$ and $\alpha > 0$, when players 3 and 4 ε -optimize on $\{x < -\alpha\}$ and on $\{x > \alpha\}$, they must each, with at least probability $1 - \varepsilon/(2\alpha)$, choose -1 on $\{x < -\alpha\}$ and choose $+1$ on $\{x > \alpha\}$.

To prevent player 2 from matching player 3, player 1 should lead 3 to randomize, which 1 can do optimally by randomizing over small positive and negative x .

Any setwise-limit distribution over outcomes is only finitely additive, since, as $\varepsilon > 0$ tends to zero, the probability of at least one of the events $\{x : -\varepsilon < x < 0\}$ or $\{x : 0 < x < \varepsilon\}$ must tend to $1/2$.⁶

The unique weak*-limit distribution over outcomes is $x = 0$ and y, z , and w each putting probability $1/2$ on -1 and $+1$. But in this limit, 3's and 4's actions are perfectly correlated

⁶It is possible that, in the limit, player 1 gives probability as large as $1/2$ to the event $\{x = 0\}$.

independently of 1's and 2's. So no strategy profile can produce this distribution and we may say that players 3 and 4 are strategically entangled in the limit.⁷

The next two examples illustrate why we do not use finite approximating games as a basis for defining sequential equilibrium in infinite games. Instead, our definition is based on strategies that are approximately conditionally optimal among *all* of the infinitely many strategies in the original game. We avoid finite approximations because we have not found any method for providing “good” finite approximations of arbitrary multi-stage games. The finite approximations used in these next examples seem natural but lead to unacceptable results.

Example 2.2 *Problems of spurious signaling in naïve finite approximations.*

This example illustrates a difficulty that can arise when one tries to approximate a game by restricting players to finite subsets of their action spaces. It can happen that no such “approximation” yields sensible equilibria because new signaling opportunities necessarily arise.

- On date 1, nature chooses $\theta \in \{1, 2\}$ with $p(\theta = 1) = 1/4$, and player 1 chooses $x \in [0, 1]$.
- On date 2, player 2 observes a signal $s = x^\theta$ and chooses $y \in \{1, 2\}$.
- Payoffs (u_1, u_2) are as follows:

	$y = 1$	$y = 2$
$\theta = 1$	(1, 1)	(0, 0)
$\theta = 2$	(1, 0)	(0, 1)

Consider subgame perfect equilibria of any finite approximate version of the game where player 1 chooses x in some finite subset of $[0, 1]$ that includes at least one interior point. We shall argue that player 1's expected payoff must be $1/4$.

Player 1 can obtain an expected payoff of at least $1/4$ by choosing the largest feasible $\bar{x} < 1$, as 2 should choose $y = 1$ when $s = \bar{x} > \bar{x}^2$ indicates $\theta = 1$. (In this finite approximation, player 2 has perfect information after the history $\theta = 1, x = \bar{x}$.)

⁷The nonexistence of a strategy supporting the limit outcome distribution can be remedied here by adding an appropriate correlation device between periods as in Harris et. al. (1995). But this approach, which is not at all worked out for general multi-stage games, can add equilibria that are not close to any ε -equilibria of the original game.

Player 1's expected payoff cannot be more than $1/4$, as 1's choice of the smallest $0 < \underline{x} < 1$ in his equilibrium support would lead player 2 to choose $y = 2$ when $s = \underline{x}^2 < \underline{x}$ indicates $\theta = 2$.

But such a scenario cannot be even an approximate equilibrium of the original infinite game, because player 1 could get an expected payoff at least $3/4$ by deviating to $\sqrt{\bar{x}}$ ($> \bar{x}$).

In fact, by reasoning analogous to that in the preceding two sentences, player 1 must receive an expected payoff of 0 in any Nash equilibrium of the infinite game, and so also in any sensibly defined "sequential equilibrium."

Hence, approximating this infinite game by restricting player 1 to any large but finite subset of his actions, would produce subgame perfect equilibria (and hence also sequential equilibria) that are all far from any sensible equilibrium of the original infinite game.

Example 2.3 *More spurious signaling in finite approximating games (Bargaining for Akerlof's lemons).*

Instead of finitely approximating the players' action sets, one might consider restricting the players to finite subsets of their strategy sets. This example makes use of Akerlof's bargaining game to illustrate a difficulty with this approach.

- On date 1, nature chooses θ uniformly from $[0, 1]$.
- On date 2, player 1 observes the signal $s_1 = \theta$ and chooses $x \in [0, 2]$.
- On date 3, player 2 observes the signal $s_2 = x$ and chooses $y \in \{0, 1\}$.
- Payoffs are $u_1 = y(x - \theta)$, $u_2 = y(1.5\theta - x)$.

We can interpret player 1 as the seller of a good and player 2 as the buyer. The good is worth θ to the seller and 1.5θ to the buyer, but only the seller knows θ . The seller offers to sell the good for a price, x . The buyer can accept the offer ($y = 1$) or reject the offer ($y = 0$).

Consider any finite approximate game where player 1 has a given finite set of pure strategies (each being a measurable function from $[0, 1]$ into $[0, 2]$) and player 2 observes a given finite partition of $[0, 2]$ before choosing y (and so player 2 is restricted to the finite set of strategies that are measurable with respect to this partition).

For any $\delta > 0$, we can construct a function $f : [0, 1] \rightarrow [0, 1.5]$ such that: $f(\theta) = 0 \forall \theta \in [0, \delta)$, $f(\cdot)$ takes finitely many values on $[\delta, 1]$ and, for every $\theta \in [\delta, 1]$, it is the case that $\theta < f(\theta) < 1.5\theta$ and $f(\theta)$ has probability 0 under each strategy in 1's given finite set.

Then there is a larger finite game (a "better" approximation) where we add the single strategy f for player 1 and we refine player 2's signal partition by giving player 2 the ability

to recognize each x in the finite range of f . If $\delta > 0$ is small enough, f can be chosen so that this larger finite game has a sequential equilibrium in which player 1 chooses f and player 2 accepts $f(\theta)$ for any θ .

But in the original infinite game this is not even a Nash equilibrium because, when 2 would accept $f(\theta)$ for any θ , player 1 could do strictly better with the strategy that chooses $x = \max_{\theta' \in [0,1]} f(\theta')$ for all signals $s_1 = \theta$ that he observes. Indeed, the event that player 1 makes an offer that 2 accepts has probability zero any Nash equilibrium of this game, and so also in any sensibly defined “sequential equilibrium.”

Thus, restricting players to finite subsets of their strategy spaces can produce equilibria in the finite game that fail to be even approximate equilibria in the infinite game because important strategies may be left out. We eliminate such false equilibria by requiring approximate optimality among *all* strategies in the original infinite game.

Example 2.4 *Why nature must be perturbed to test rational behavior with positive probability in all events.*

Our final example motivates why we perturb both nature’s probability function and the players’ strategies in order that, for every finite set of actions, each action in the set has positive probability in some test for rational behavior at signals reached by those actions. The example shows that such a requirement can be incompatible with the existence of equilibrium if we perturb only the players’ strategies.

- On date 1, nature chooses $\theta = (\theta_1, \theta_2) \in \{1, 2\} \times [0, 1]$, and player 1 chooses $x \in [0, 1]$.
- Nature’s coordinates are independent with $\theta_1 \in \{1, 2\}$ receiving probability $\theta_1/3$, and with θ_2 uniform on $[0, 1]$
- On date 2, player 2 observes the signal $s = x$ if $\theta_1 = 1$ and observes $s = \theta_2$ if $\theta_1 = 2$, and chooses $y \in \{1, 2\}$.
- Payoffs are as follows and are independent of the date-1 choices of θ_2 and x in $[0, 1]$.

	$y = 1$	$y = 2$
$\theta_1 = 1$	(1, 1)	(0, 0)
$\theta_1 = 2$	(1, 0)	(0, 1)

Player 2 wants to correctly guess whether the real number that she observes was chosen by player 1 ($\theta_1 = 1$) or was chosen by nature ($\theta_1 = 2$), and player 1 wants player 2 to always choose $y = 1$.⁸

⁸So the payoffs here are like the payoffs in Example 2.2. But the information structure is significantly different.

This game has a continuum of Nash equilibria, all of which seem reasonable. In one such equilibrium, player 1 chooses x uniformly from $[0, 1]$, and player 2 chooses $y = 2$ regardless of the signal that she observes.⁹ In another, player 1 provides an informative signal to player 2 by choosing $x = 1/2$ with probability 1, and player 2 chooses $y = 1$ if she observes the signal $1/2$ and she chooses $y = 2$ otherwise; so player 2 always guesses correctly.¹⁰ In the first equilibrium, the (expected) payoff vector is $(0, 2/3)$, and in the second it is $(1/3, 1)$.

However, if we required that, for any signal event that can have positive probability under some strategy profile, each player's behavior should pass a conditional rationality test in slightly perturbed strategies that give this event positive probability (so that conditional payoffs can be computed), then there would be no equilibrium at all. Indeed, for any $\alpha \in [0, 1]$, the event $\{s = \alpha\}$ can have positive probability, but only if positive probability is given to the history $(\theta_1, x) = (1, \alpha)$, because the event $\{\theta_2 = \alpha\}$ has probability 0. So, in any scenario where $\{s = \alpha\}$ has positive probability, conditional rationality would require player 2 to choose $y = 1$ when she observes $s = \alpha$ since the resulting conditional probability of the history $(\theta_1, x) = (1, \alpha)$ is one. Applying this same argument to every signal $\alpha \in [0, 1]$ would imply that player 2 must choose $y = 1$ after every signal. But always choosing $y = 1$ is strictly dominated for player 2 by the strategy of always choosing $y = 2$. Hence, there is no strategy for player 2 that is ex-ante rational and that satisfies this positive probability conditional rationality requirement.

This existence problem arises here because, when only the players' strategies are perturbed, the positive probability rationality test biases player 2's conditional beliefs toward explaining prior probability-zero events as always being the result of a deviation by player 1 instead of perhaps being the result of the occurrence of a probability-zero state of nature.

To stay clear of this existence problem, we perturb both the players' strategies and nature's probability function in our tests for rational behavior.

3 Multi-Stage Games

A multi-stage game is played in a finite sequence of dates.¹¹ At each date t , each player receives a private signal, called the player's "signal" at date t , about the history of play. Each player then simultaneously chooses an action from his set of available date- t actions,

⁹There is a whole class of such equilibria in which player 1 chooses x uniformly from any subset C of $[0, 1]$ having positive Lebesgue measure. Player 2 chooses $y = 2$ after any signal s that is not in C . After a signal s in C , 2 chooses $y = 1$ if C has measure less than $1/2$, chooses $y = 2$ if C has measure greater than $1/2$, and can mix in any way between $y = 1$ and $y = 2$ if C has measure exactly $1/2$.

¹⁰There are many equilibria of this sort, where player 1 puts positive probability only on actions that, when observed as signals by player 2 prompt player 2 to choose $y = 1$.

¹¹A countable infinity of dates can be accommodated with some additional notation.

and nature simultaneously chooses a date- t state whose distribution can depend on the entire history of play. Perfect recall is assumed.

Formally, a multi-stage game $\Gamma = (I, T, A, S, \mathcal{M}, p, \sigma, u)$ consists of the following items.

- $\Gamma.1.$** I is the finite set of players, $0 \notin I$. Let $I^* = I \cup \{0\}$, where 0 denotes nature (chance). The finite set of dates of play is $\{1, \dots, T\}$. Let $L = I \times \{1, \dots, T\}$ denote the set of dated players, let $L^* = I^* \times \{1, \dots, T\}$ and write it for (i, t) .
- $\Gamma.2.$** For $it \in L$, A_{it} is the set of all possible *date- t actions* for player i at date t , and A_{0t} is the set of all possible *date- t states* of nature. Let $A_L = \times_{it \in L} A_{it}$ be the set of all possible *action profiles* in the game, let $A_0 = \times_{t \leq T} A_{0t}$ be the set of all possible *states* in the game, and let $A = \times_{t \leq T} \times_{i \in I^*} A_{it} = A_0 \times A_L$ be the set of all possible *outcomes* in the game.¹²
- $\Gamma.3.$** $S = \times_{it \in L} S_{it}$, where S_{it} is the set of possible signals received by player i at date t ; $S_{i1} = \{\emptyset\}$ for all $i \in I$.
- $\Gamma.4.$** Sigma-algebras (closed under complements and countable intersections) of measurable subsets are specified for each S_{it} , A_{it} , and A_{0t} . All one-point sets are measurable and products are given their product sigma-algebras. $\mathcal{M}(X)$ denotes the set of measurable subsets of any set X on which a sigma-algebra is specified.

The subscript, $< t$, will always denote the projection onto dates before t , and $\leq t$ weakly before. e.g., $A_{<t} = \times_{i \in I^*, r < t} A_{ir} = \{\text{possible histories before date } t\}$ ($A_{<1} = \{\emptyset\}$), and for $a \in A$, $a_{<t} = (a_{ir})_{i \in I^*, r < t}$ is the partial sequence of actions and states before date t .

Let $\Delta(X)$ denote the set of countably additive probability measures on $\mathcal{M}(X)$. Recall that, for any two measurable spaces X and Y , a mapping $\lambda : Y \rightarrow \Delta(X)$ is a transition probability iff for every $C \in \mathcal{M}(X)$, $\lambda(C|y)$ is a measurable real-valued function of y on Y .

- $\Gamma.5.$** $p = (p_1, \dots, p_T)$ is nature's probability function where, for each date t , a history-dependent probability measure over nature's date- t states is given by the transition probability $p_t : A_{<t} \rightarrow \Delta(A_{0t})$.

¹²Models with history-dependent action sets can be reduced to equivalent models with history-independent action sets. Indeed, one could specify that A_{it} is the union of all possible feasible action sets for player i at date t , and that any action a_{it} that is outside the feasible set of actions for a given signal s_{it} should be re-interpreted as corresponding to the choice of some $g_{it}(a_{it}, s_{it})$ that is in the feasible set, where $g_{it} : A_{it} \times S_{it} \rightarrow A_{it}$ is measurable. A similar device can be used for models in which nature's possible states are history-dependent. Such history-dependence for nature's possible states can be important to specify since it can affect the set of feasible perturbations of nature.

Γ.6. Player i 's date t information is determined by a measurable signal function $\sigma_{it} : A_{<t} \rightarrow S_{it}$.¹³ Assume perfect recall: $\forall it \in L, \forall r < t$, there are measurable functions $\bar{\Psi}_{itr} : S_{it} \rightarrow S_{ir}$ and $\bar{\psi}_{itr} : S_{it} \rightarrow A_{ir}$ such that $\bar{\Psi}_{itr}(\sigma_{it}(a_{<t})) = \sigma_{ir}(a_{<r})$ and $\bar{\psi}_{itr}(\sigma_{it}(a_{<t})) = a_{ir}, \forall a \in A$. The game's signal function is $\sigma = (\sigma_{it})_{it \in L}$.

Γ.7. Each player i has a bounded measurable utility function $u_i : A \rightarrow \mathbb{R}$, and $u = (u_i)_{i \in I}$.

So, at each date $t \in \{1, \dots, T\}$ starting with date $t = 1$, and given a partial history $a_{<t} \in A_{<t}$, each player i is privately informed of his date- t signal, $s_{it} = \sigma_{it}(a_{<t})$, after which each player i simultaneously chooses an action from his set of date- t actions A_{it} and nature chooses a date- t state $a_{0t} \in A_{0t}$ according to $p_t(\cdot | a_{<t})$. The game then proceeds to the next date. After T dates of play this leads to an outcome $a \in A$ and the game ends with player payoffs $u_i(a), i \in I$.

In the next two subsections, we formally introduce strategies, outcome distributions, payoffs, and conditional payoffs.

3.1 Strategies and Induced Outcome Distributions

A *strategy* for dated player $it \in L$ is any transition probability $b_{it} : S_{it} \rightarrow \Delta(A_{it})$.

Let B_{it} denote it 's set of strategies and let $B_i = \times_{t \in T} B_{it}$ denote i 's (behavioral) strategies. Perfect recall ensures that there is no loss in restricting attention to B_i for each player i . Let $B = \times_{it \in L} B_{it}$ denote the set of all strategy profiles.

For any date t , let $B_{\cdot t} = \times_{i \in I} B_{it}$ denote the set of date- t strategy vectors with typical element $b_{\cdot t} = (b_{it})_{i \in I}$. Let $A_{\cdot t} = \times_{i \in I^*} A_{it}$. Each $b_{\cdot t} \in B_{\cdot t}$ determines a transition probability P_t from $A_{<t}$ to $\mathcal{M}(A_{\cdot t})$ such that, for any measurable product set $H = \times_{i \in I^*} H_{it} \subseteq \times_{i \in I^*} A_{it}$ and for any $a_{<t} \in A_{<t}$,

$$P_t(H | a_{<t}, b_{\cdot t}) = p_t(H_{0t} | a_{<t}) \prod_{i \in I} b_{it}(H_{it} | \sigma_{it}(a_{<t})). \quad (3.1)$$

For any $b \in B$, we inductively define measures $P_{<t}(\cdot | b)$ in $\Delta(A_{<t})$ so that $P_{<1}(\{\emptyset\} | b) = 1$ and, for all $t \in \{1, \dots, T\}$ and for all measurable $H \subseteq A_{<t+1}$,

$$P_{<t+1}(H | b) = \int P_t(\{a_{\cdot t} : (a_{<t}, a_{\cdot t}) \in H\} | a_{<t}, b_{\cdot t}) P_{<t}(da_{<t} | b). \quad (3.2)$$

(Notice that $P_{<t}(\cdot | b)$ depends only on $b_{<t}$.)

¹³It is without loss of generality to assume, for every $r < t$, that σ_{it} does not depend on the date- r signal of any player since earlier signals depend on even earlier states and actions.

Let $P(\cdot|b) = P_{<T+1}(\cdot|b)$ be the probability measure on outcome events in $\mathcal{M}(A)$ that is induced by b . The dependence of $P(\cdot|b)$ on nature's probability function p will sometimes be made explicit by writing $P(\cdot|b; p)$.

For any $b \in B$, we inductively define transition probabilities from $A_{<t}$ to $\Delta(A_{\geq t})$ so that $P_{\geq T}(\cdot|a_{<T}, b) = P_T(\cdot|a_{<T}, b_T)$, and for any date $t < T$ and any measurable $H \subseteq A_{\geq t}$,

$$P_{\geq t}(H|a_{<t}, b) = \int P_{\geq t+1}(\{a_{\geq t+1} : (a_t, a_{\geq t+1}) \in H\}|a_{<t+1}, b) P_t(da_t|a_{<t}, b_t).$$

(Notice that $P_{\geq t}(\cdot|a_{<t}, b)$ does not depend on $b_{<t}$.)

At any date t , the conditional expected utility for player i with strategies b given history $a_{<t}$ is,

$$U_i(b|a_{<t}) = \int u_i(a_{<t}, a_{\geq t}) P_{\geq t}(da_{\geq t}|a_{<t}, b),$$

(notice that $U_i(b|a_{<t})$ does not depend on $b_{<t}$), and, player i 's ex-ante expected utility is

$$U_i(b) = \int u_i(a) P(da|b) = \int U_i(b|a_{<t}) P_{<t}(da_{<t}|b).$$

3.2 Conditional Probabilities

For any $b \in B$, for any $it \in L$ and for any $Z \in \mathcal{M}(S_{it})$, define

$$P_{it}(Z|b) = P_{<t}(\sigma_{it}^{-1}(Z)|b) = P_{<t}(\{a_{<t} : \sigma_{it}(a_{<t}) \in Z\}|b).$$

Then $P_{it}(Z|b)$ is the probability that i 's date t signal is in Z under the strategy profile b . The dependence of $P_{it}(\cdot|b)$ on nature's probability function p will sometimes be made explicit by writing $P_{it}(\cdot|b; p)$.

For any $it \in L$ and for any measurable $Z \subseteq S_{it}$, if $P_{it}(Z|b) > 0$, we may define: *conditional probabilities*,

$$P_{<t}(H|Z, b) = P_{<t}(H \cap \sigma_{it}^{-1}(Z)|b) / P_{it}(Z|b), \quad \forall H \in \mathcal{M}(A_{<t}),$$

and,

$$P(H|Z, b) = P(\{a \in H : \sigma_{it}(a_{<t}) \in Z\}|b) / P_{it}(Z|b), \quad \forall H \in \mathcal{M}(A),$$

and *conditional expected payoffs*,

$$U_i(b|Z) = \int_A u_i(a) P(da|Z, b).$$

(Notice that $P_{<t}(\cdot|Z, b)$ is the marginal of $P(\cdot|Z, b)$ on $A_{<t}$.) The dependence of $P(\cdot|Z, b)$

and $U_i(\cdot|Z)$ on nature's probability function p will sometimes be made explicit by writing $P(\cdot|Z, b; p)$ and $U_i(\cdot|Z; p)$.

4 Perfect Conditional ε -Equilibrium

4.1 Conditional ε -Equilibrium

A *signal domain* is any disjoint union $Y = \cup_{it \in L} Y_{it}$, where each Y_{it} is a measurable subset of S_{it} . A signal domain Y is ε -*sure* iff $P_{it}(Y_{it}|b) \geq 1 - \varepsilon^2$ for all $b \in B$ and for all $it \in L$. Given any event with probability ε or greater under some strategy profile, signals outside an ε -sure domain would have conditional probability less than ε even after a deviation.

For any $it \in L$, and for any $b_i \in B_i$, say that $c_i \in B_i$ is a *date- t continuation* of b_i if $c_{ir} = b_{ir}$ for all $r < t$.

Definition 4.1 For any $\varepsilon \geq 0$, and for any signal domain Y , a strategy profile $b \in B$ is a *conditional ε -equilibrium* on Y iff Y is ε -sure and for every $it \in L$, for every measurable $Z \subseteq Y_{it}$ satisfying $P_{it}(Z|b) > 0$, and for every date- t continuation c_i of b_i ,

$$U_i(c_i, b_{-i}|Z) \leq U_i(b|Z) + \varepsilon.$$

Remark 4.2 In a conditional ε -equilibrium, ε -optimality can fail outside the ε -sure domain. This is sometimes necessary for the existence of a conditional ε -equilibrium. Indeed, Hellman (2014), building on Simon (2003), gives an example of a Bayesian game that, for $\varepsilon > 0$ small enough, has no strategy profile that is almost everywhere interim ε -optimal for each player. Nevertheless, Hellman's game can be shown to possess a conditional ε -equilibrium.¹⁴

4.2 Perfect Conditional ε -Equilibrium

For each date $t \leq T$, let \mathcal{T}_t be the set of transition probabilities $\tau_t : A_{<t} \rightarrow \Delta(A_{0t})$. Let $\mathcal{T} = \times_{t \leq T} \mathcal{T}_t$. Nature's probability function p is an element of \mathcal{T} . For any $\tau = (\tau_1, \dots, \tau_T) \in \mathcal{T}$, let $\Gamma(\tau)$ denote the perturbed game in which nature's probability function is τ instead of p .

For any $\tau', \tau \in \mathcal{T}$, define $\|\tau' - \tau\| = \sup |\tau'_t(C|a_{<t}) - \tau_{it}(C|a_{<t})|$, where the supremum is over all $t \leq T$, $a_{<t} \in A_{<t}$, and $C \in \mathcal{M}(A_{0t})$. For any $b', b \in B$ define $\|b' - b\| = \sup |b'_{it}(C|s_{it}) - b_{it}(C|s_{it})|$, where the supremum is over all $it \in L$, $s_{it} \in S_{it}$, and $C \in \mathcal{M}(A_{it})$.

Say that $\tau' \in \mathcal{T}$ is a δ -*perturbation* of $\tau \in \mathcal{T}$ iff $\|\tau' - \tau\| \leq \delta$ and say that $b' \in B$ is a δ -*perturbation* of $b \in B$ iff $\|b' - b\| \leq \delta$.

¹⁴To construct a conditional ε -equilibrium in Hellman's two-player game, one can begin by seeding a small set of one player's signals to play any action. Then the structure of Hellman's game allows one to iteratively define best replies for all of the other signals of both players.

Definition 4.3 For any $\varepsilon > 0$, a strategy profile $b \in B$ is a *perfect conditional ε -equilibrium* iff there is an ε -sure signal domain Y , and, for every $\delta > 0$ and for every finite $Z \subseteq Y$, there is $b' \in B$ and $p' \in \mathcal{T}$ such that $\|b' - b\| \leq \delta$, $\|p' - p\| \leq \delta$, and, in the perturbed game $\Gamma(p')$, b' is a conditional ε -equilibrium on Y , and $P_{it}(\{s_{it}\}|b'; p') > 0$ for every $s_{it} \in Z \cap Y_{it}$ and for every $it \in L$.

Remark 4.4 If b is a perfect conditional ε -equilibrium, then b is a conditional ε -equilibrium on an ε -sure signal domain. (This is because, when $\|b' - b\| \leq \delta$ and $\|p' - p\| \leq \delta$, the probabilities of all events in the game under b' and p' converge uniformly to their values under b and p as $\delta \rightarrow 0$.) In particular, if $\varepsilon \in [0, 1)$, then b is an ε -Nash equilibrium since the null signal at the start of the game must be included in any ε -sure domain when $\varepsilon \in [0, 1)$.

Remark 4.5 If every feasible state of nature has positive probability at each date, then perturbations of nature are redundant and so, in Definition 4.3, it would be equivalent to always set $p' = p$.¹⁵

Remark 4.6 In any Kreps-Wilson finite multi-stage game, a strategy profile b is part of a Kreps-Wilson sequential equilibrium assessment iff b is the limit of a sequence of perfect conditional ε -equilibria as $\varepsilon \rightarrow 0$.

4.3 Perfect Conditional Equilibrium Distributions

We next define a “perfect conditional equilibrium distribution” as a limit of perfect conditional ε -equilibrium distributions on outcomes as $\varepsilon \rightarrow 0$.

Definition 4.7 A *perfect conditional equilibrium distribution* is a mapping $\mu : \mathcal{M}(A) \rightarrow [0, 1]$ such that, for any finite $\mathcal{F} \subseteq \mathcal{M}(A)$ and for any $\varepsilon > 0$, there is a perfect conditional ε -equilibrium b such that,

$$|\mu(H) - P(H|b)| < \varepsilon, \quad \forall H \in \mathcal{F}.$$

Equivalently, $\mu : \mathcal{M}(A) \rightarrow [0, 1]$ is a perfect conditional equilibrium distribution iff there is a net $\{b^{\varepsilon, \mathcal{F}}\}$ of perfect conditional ε -equilibria, where smaller values of $\varepsilon > 0$ and more inclusive finite subsets \mathcal{F} of $\mathcal{M}(A)$ correspond to larger indices, such that,

$$\mu(H) = \lim_{\varepsilon, \mathcal{F}} P(H|b^{\varepsilon, \mathcal{F}}), \quad \text{for every } H \in \mathcal{M}(A). \quad (4.1)$$

¹⁵See footnote 12 for how to handle games in which nature’s set of feasible states is history-dependent.

Remark 4.8 *In any finite game, the set of perfect conditional equilibrium distributions are precisely the set of distributions over outcomes induced by the set of Kreps-Wilson sequential equilibria.*

The existence of perfect conditional ε -equilibria is taken up in Section 8.1. We record here the simpler result, based on Tychonoff's theorem (Section 12 contains the proof), that a perfect conditional equilibrium distribution exists so long as perfect conditional ε -equilibria always exist.

Theorem 4.9 *If for each $\varepsilon > 0$ there is at least one perfect conditional ε -equilibrium, then a perfect conditional equilibrium distribution exists.*

It follows immediately from (4.1) that if μ is a perfect conditional equilibrium distribution, then μ is a finitely additive probability measure on $\mathcal{M}(A)$.¹⁶

If (4.1) holds, then so long as u_i is bounded and measurable (as we have assumed),

$$\lim_{\varepsilon, \mathcal{F}} \int_A u_i(a) P(da|b^{\varepsilon, \mathcal{F}}) = \int_A u_i(a) \mu(da), \quad (4.2)$$

and so we define i 's *equilibrium expected payoff* (at μ) by

$$\int_A u_i(a) \mu(da).$$

Sometimes μ is only finitely additive, not countably additive (Example 2.1). Even so, in many practical settings there is a natural countably additive probability measure over outcomes that is induced by μ .

Definition 4.10 *Suppose that A is a normal topological space and $\mathcal{M}(A)$ is its Borel sigma-algebra.¹⁷ We say that ν is the regular countably additive distribution induced by μ iff ν is a regular countably additive probability measure on $\mathcal{M}(A)$ such that $\int f(a)\nu(da) = \int f(a)\mu(da)$ for all bounded continuous $f : A \rightarrow \mathbb{R}$.¹⁸*

Remark 4.11 *If A is a compact Hausdorff space with its Borel sigma algebra of measurable sets and μ is a perfect conditional equilibrium distribution, then, by the Riesz representation*

¹⁶For any disjoint sets $G, H \in \mathcal{M}(A)$, (4.1) and $\lim P(G \cup H|b^{\varepsilon, \mathcal{F}}) = \lim[P(G|b^{\varepsilon, \mathcal{F}}) + P(H|b^{\varepsilon, \mathcal{F}})]$ imply that $\mu(G \cup H) = \mu(G) + \mu(H)$.

¹⁷Recall that a topological space is normal if any pair of disjoint closed sets can be separated by disjoint open sets.

¹⁸There can be at most one such Borel measure ν since, by Theorem IV.6.2 in Dunford and Schwartz (1988), any two such measures must agree on all closed sets. Then, by Corollary 1.6.2 in Cohn (1980), the two measures must agree on all Borel sets since the set of closed sets is closed under finite intersections and generates the Borel sigma algebra.

theorem (see, e.g., Theorem IV.6.3 in Dunford and Schwartz 1988), there exists a regular countably additive distribution ν induced by μ .^{19,20} In particular, player i 's equilibrium expected payoff (at μ), namely $\int_A u_i(a)\mu(da)$, will be equal to $\int_A u_i(a)\nu(da)$ whenever u_i is a continuous function.

By Remark 4.6, μ is a perfect conditional equilibrium distribution iff it is the distribution on outcomes generated by some Kreps-Wilson sequential equilibrium strategy profile.

We next introduce the counterpart, for infinite games, of the systems of beliefs that Kreps and Wilson introduced for finite games. We also provide an analogue of their consistency requirement for a system of beliefs given a strategy profile.

5 Subgame Perfection

Given perfect recall, we may say that a date- t history $a_{<t} \in A_{<t}$ is a *subgame* of Γ iff $\sigma_{it}^{-1}(\sigma_{it}(a_{<t})) = \{a_{<t}\}$, for all $i \in I$.

For any $\varepsilon > 0$, a strategy profile $b \in B$ is a *subgame perfect ε -equilibrium* of Γ iff there is an ε -sure signal domain Y such that for every $it \in L$ and for every subgame $a_{<t}$, if $\sigma_{it}(a_{<t}) \in Y_{it}$, then $U_i(c_i, b_{-i}|a_{<t}) \leq U_i(b|a_{<t}) + \varepsilon$ for all $c_i \in B_i$.

Theorem 5.1 *If b is a perfect conditional ε -equilibrium, then b is a subgame perfect ε -equilibrium.*

6 Beliefs and Sequential Rationality

6.1 Conditional Belief Systems

A *conditional belief system* β specifies, for every $it \in L$, and for every $Z \in \mathcal{M}(S_{it})$, a finitely additive probability measure $\beta_{it}(\cdot|Z)$ on the measurable subsets of $A_{<t}$ such that $\beta_{it}(\sigma_{it}^{-1}(Z)|Z) = 1$.²¹

So, a conditional belief system specifies for any measurable set of any dated player's signals, a finitely additive probability measure over histories that gives probability one to the set of all histories that generate signals in the given set.

An *assessment* is any pair (b, β) where $b \in B$ is a strategy profile and β is a conditional belief system.

¹⁹We thank a referee for suggesting the Riesz representation theorem here.

²⁰The existence of a regular countably additive distribution induced by μ can be established under the alternative conditions that for each date t : (i) A_{it} is compact metric for all $i \in I$ and A_{0t} is Polish, and (ii) either A_{0t} is compact or $p_t(\cdot|a_{<t})$ is weak* continuous in $a_{<t}$.

²¹The set $\sigma_{it}^{-1}(Z)$ is nonempty because, in a multi-stage game, each signal function $\sigma_{it} : A_{<t} \rightarrow S_{it}$ is onto, i.e., its range is S_{it} .

6.1.1 Bayes Consistency and Finite Consistency

Definition 6.1 An assessment (b, β) is Bayes consistent on the signal domain Y iff for all $it \in L$, for all $H \in \mathcal{M}(A_{<t})$, and for all measurable $Z \subseteq Y_{it}$ such that $P_{it}(Z|b) > 0$,

$$\beta_{it}(H|Z) = P_{<t}(H|Z, b). \quad (6.1)$$

Bayes' consistency fails to discipline beliefs on signal events that have probability zero. The next definition extends to infinite games Kreps and Wilson's (1982) definition of a belief system that is consistent with a given strategy profile.

For any signal domain Y , define the disjoint union $\mathcal{M}(Y) = \cup_{it \in L} \mathcal{M}(Y_{it})$, where $\mathcal{M}(Y_{it})$ denotes the collection of subsets of Y_{it} that are measurable subsets of S_{it} . So, $\mathcal{M}(Y)$ is the collection of all measurable signal events that are in the signal domain Y .

Definition 6.2 An assessment (b, β) is finitely consistent on the signal domain Y iff for every $\delta > 0$, for every finite collection of signal events $\mathcal{Z} \subseteq \mathcal{M}(Y)$, and for every finite collection of measurable sets of histories $\mathcal{H} \subseteq \cup_{r \leq T} \mathcal{M}(A_{<r})$ (disjoint union), there is $b' \in B$ and $p' \in \mathcal{T}$ such that $\|b' - b\| \leq \delta$, $\|p' - p\| \leq \delta$, and, for every $it \in L$, for every $Z \in \mathcal{Z} \cap \mathcal{M}(Y_{it})$, and for every $H \in \mathcal{H} \cap \mathcal{M}(A_{<t})$,

$$P_{it}(Z|b'; p') > 0 \quad \text{and} \quad |\beta_{it}(H|Z) - P_{<t}(H|Z, b'; p')| \leq \delta. \quad (6.2)$$

Equivalently, (b, β) is finitely consistent on the signal domain Y iff there is a net $\{(b^{\delta, \mathcal{Z}, \mathcal{H}}, p^{\delta, \mathcal{Z}, \mathcal{H}})\}$, where smaller values of $\delta > 0$ and more inclusive finite subsets \mathcal{Z} of $\mathcal{M}(Y)$ and \mathcal{H} of $\cup_{r \leq T} \mathcal{M}(A_{<r})$ correspond to larger indices, such that $\lim_{\delta, \mathcal{Z}, \mathcal{H}} \|b^{\delta, \mathcal{Z}, \mathcal{H}} - b\| = \lim_{\delta, \mathcal{Z}, \mathcal{H}} \|p^{\delta, \mathcal{Z}, \mathcal{H}} - p\| = 0$, $P_{nr}(Z|b^{\delta, \mathcal{Z}, \mathcal{H}}, p^{\delta, \mathcal{Z}, \mathcal{H}}) > 0$ for all $Z \in \mathcal{Z} \cap \mathcal{M}(Y_{nr})$ and all $nr \in L$, and,

$$\beta_{it}(H|Z) = \lim_{\delta, \mathcal{Z}, \mathcal{H}} P_{<t}(H|Z, b^{\delta, \mathcal{Z}, \mathcal{H}}, p^{\delta, \mathcal{Z}, \mathcal{H}}), \quad \forall H \in \mathcal{M}(A_{<t}), \forall Z \in \mathcal{M}(Y_{it}), \forall it \in L. \quad (6.3)$$

Finitely consistent assessments have many consistency properties (see Remarks 6.5-6.7), but they can sometimes seem paradoxical. For example, suppose that on date 1, player 1 chooses $x = (x_1, x_2)$ uniformly from $\{0, 1\} \times [0, 1]$ and that on date 2, player 2 observes player 1's second coordinate x_2 . Then there is a finitely consistent assessment in which player 2's conditional beliefs put probability 1 on $x_1 = 0$ no matter what value of $x_2 \in [0, 1]$ that he observes (i.e., $\beta_{22}(\{(0, x_2)\} | \{x_2\}) = 1, \forall x_2 \in [0, 1]$), even though before observing x_2 , player 2's beliefs put probability 1/2 on $x_1 = 0$ (i.e., $\beta_{22}(\{0\} \times [0, 1] | [0, 1]) = 1/2$).²²

²²For any finite subset F of $[0, 1]$, perturb player 1's strategy to give small positive probability to $x = (0, x_2)$

Thus, with only finite consistency, beliefs on one-point signal events may not be sufficient to determine beliefs more generally. In a finitely consistent conditional belief system, the probability assigned to any set of histories conditional on any particular signal event need not be a convex combination of the probabilities assigned to that set of histories conditional on each element of an arbitrary partition of that signal event. However, that probability can always be obtained as a convex combination of the conditional probabilities given each element of any *finite* partition of that event (see Remark 6.6).

Remark 6.3 *In any Kreps-Wilson finite multi-stage game, any finitely consistent assessment is Kreps-Wilson consistent.*

Remark 6.4 *For any $b \in B$, there is a conditional belief system β such that (b, β) is finitely consistent on the signal domain Y defined by $Y_{it} = S_{it} \forall it \in L$.²³*

Remark 6.5 *Every finitely consistent assessment on any domain is Bayes consistent on that domain.²⁴*

Remark 6.6 *If (b, β) is finitely consistent on Y , then (using 6.3) the conditional belief system β satisfies the following additional Bayes consistency property: For all $it \in L$, for all $W, Z \in \mathcal{M}(Y_{it})$ and for all $H \in \mathcal{M}(A_{<t})$,*

$$\beta_{it}(\sigma_{it}^{-1}(W)|Z)\beta_{it}(H \cap \sigma_{it}^{-1}(Z)|W) = \beta_{it}(\sigma_{it}^{-1}(Z)|W)\beta_{it}(H \cap \sigma_{it}^{-1}(W)|Z).$$

Remark 6.7 *If (b, β) is finitely consistent on Y , then (using equation (6.3) and an analogue of equation (4.2) for expected continuation payoffs and conditional beliefs) any two players that have the same conditional information must agree about the conditional distribution over outcomes. That is, for every $it, nr \in L$, if $Z_{it} \in \mathcal{M}(Y_{it})$ and $Z_{nr} \in \mathcal{M}(Y_{nr})$ are such that $\{a \in A : \sigma_{it}(a_{<t}) \in Z_{it}\} = \{a \in A : \sigma_{nr}(a_{<t}) \in Z_{nr}\}$, then, for every $H \in \mathcal{M}(A)$,*

$$\int_A \mathbf{1}_H(a) P_{\geq t}(da_{\geq t} | a_{<t}, b) \beta_{it}(da_{<t} | Z_{it}) = \int_A \mathbf{1}_H(a) P_{\geq r}(da_{\geq r} | a_{<r}, b) \beta_{nr}(da_{<r} | Z_{nr}).$$

for each $x_2 \in F$. With this perturbation, the Bayes-consistent beliefs put probability 1 on $(0, x_2)$ when conditioning on any $x_2 \in F$. The setwise limit of the resulting net (or possibly subnet) of conditional beliefs for player 2 provides the requisite paradoxical beliefs.

²³For every $\delta > 0$ and for every finite subset of signal events $\mathcal{Z} \subseteq \mathcal{M}(Y)$ choose $b^{\delta, \mathcal{Z}}$ and $p^{\delta, \mathcal{Z}}$ such that $\|b^{\delta, \mathcal{Z}} - b\| \leq \delta$, $\|p^{\delta, \mathcal{Z}} - p\| \leq \delta$, and $P_{it}(Z | b^{\delta, \mathcal{Z}}; p^{\delta, \mathcal{Z}}) > 0 \forall Z \in \mathcal{Z} \cap \mathcal{M}(S_{it}) \forall it \in L$. Then, for each $it \in L$, $\{(P_{it}(\cdot | \cdot, b^{\delta, \mathcal{Z}}, p^{\delta, \mathcal{Z}}))\}$ is a net in the compact (by Tychonoff's theorem) set $[0, 1]^{\mathcal{M}(A_{<t}) \times \mathcal{M}(S_{it})}$, where smaller δ 's and more inclusive finite sets \mathcal{Z} correspond to larger indices. The limit $\beta_{it}(\cdot | \cdot)$ of any convergent subnet that is common to all $it \in L$ (such a convergent subnet exists by compactness) yields the requisite belief system β .

²⁴If 6.3 holds, then, $\beta_{it}(H|Z) = \lim_{\delta, \mathcal{Z}, \mathcal{H}} P_{<t}(H|Z, b^{\delta, \mathcal{Z}, \mathcal{H}}, p^{\delta, \mathcal{Z}, \mathcal{H}}) = P_{<t}(H|Z, b, p)$, where the second equality holds when Z has positive probability under b and p because $\lim_{\delta, \mathcal{Z}, \mathcal{H}} \|b^{\delta, \mathcal{Z}, \mathcal{H}} - b\| = \lim_{\delta, \mathcal{Z}, \mathcal{H}} \|p^{\delta, \mathcal{Z}, \mathcal{H}} - p\| = 0$.

In particular, if $r = t$, then $\beta_{it}(\cdot|Z_{it}) = \beta_{nt}(\cdot|Z_{nt})$, i.e., the two players have the same conditional beliefs (because their conditional beliefs at date $r = t$ are just the marginals on $A_{<r} = A_{<t}$ of their common conditional distribution over outcomes).

6.2 Sequential ε -Rationality

Definition 6.8 For $\varepsilon \geq 0$, an assessment (b, β) is sequentially ε -rational on the signal domain Y iff for every $it \in L$, for every measurable $Z \subseteq Y_{it}$, and for every $c_i \in B_i$,

$$\int U_i(c_i, b_{-i}|a_{<t})\beta_{it}(da_{<t}|Z) \leq \int U_i(b|a_{<t})\beta_{it}(da_{<t}|Z) + \varepsilon.$$

Remark 6.9 By tracing through the definitions, it is straightforward to verify that if (b, β) is Bayes consistent and sequentially ε -rational on an ε -sure signal domain, then b is a conditional ε -equilibrium.

The next result states that for any perfect conditional ε -equilibrium strategy profile, there is a conditional belief system that is finitely consistent and that yields a sequentially ε -rational assessment.

Theorem 6.10 If b is a perfect conditional ε -equilibrium, then there is a conditional belief system β such that (b, β) is finitely consistent and sequentially ε -rational on an ε -sure signal domain.

Remark 6.11 A consequence of Theorem 6.10 is that the plausibility of a perfect conditional ε -equilibrium strategy profile can be tested by considering the plausibility of the finitely consistent beliefs that make it sequentially ε -rational on an ε -sure signal domain. Hence, belief systems can be usefully employed here in much in the same manner as first suggested by Kreps and Wilson (1982).

Remark 6.12 One might think that, instead of requiring a strategy profile to be a perfect conditional ε -equilibrium, it should be enough if, with some finitely consistent system of beliefs, the strategy profile is sequentially ε -rational on an ε -sure signal domain. But this weaker criterion does not seem strong enough (see the first example in the next section).

7 Illustrative Examples

Example 7.1 Problems with sequential rationality and finite consistency.

The following example shows that an assessment that is finitely consistent and sequentially ε -rational on an ε -sure signal domain can be unintuitive because beliefs are only finite additive. This further motivates why we focus on perfect conditional ε -equilibria.

- On date 1, nature chooses θ uniformly from the open interval $(0, 1)$, where neither $\theta = 0$ or $\theta = 1$ is a possible state of nature.
- On date 2, player 1 observes θ and then chooses $x \in \{0, 1\}$.
- On date 3, player 2 observes x and then chooses $y \in [0, 1)$.
- Payoffs are as follows:
 - If $x = 0$ (“out”), then $u_1 = u_2 = 0$.
 - If $x = 1$ (“in”) and $y = 0$ (“out”), then $u_1 = -1$ and $u_2 = 0$.
 - If $x = 1$ and $y > 0$, then $u_1 = 1$, and $u_2 = 1$ if $y \geq \theta$ but $u_2 = -1$ if $y < \theta$.

In any perfect conditional ε -equilibrium of this game with $\varepsilon < 1/3$, player 1 must choose $x = 1$ with probability at least $1 - \varepsilon/(1 - \varepsilon)$, and player 2’s strategy, conditional on the signal $x = 1$, must give the event $\{y \geq \theta\}$ probability at least $1 - \varepsilon/2$. This is because, in any perturbation that gives the signal $x = 1$ positive probability (even if that perturbation involves nature), player 2 can obtain a conditional expected payoff arbitrarily close to 1 by choosing y sufficiently close to 1 (this latter fact is a consequence of the countable additivity of the conditional probability measures defined by the perturbations). Consequently, by choosing $x = 1$, player 1 can obtain a payoff close to 1 when $\varepsilon > 0$ is close to zero.

However, consider the strategy profile b in which player 1 chooses $x = 0$ and player 2 chooses $y = 0$ no matter what signal he observes. In addition, consider the beliefs β_{23} for player 2 such that $\beta_{23}(\cdot|x = 0)$ is uniform on $(0, 1)$ and $\beta_{23}((1 - \delta, 1)|x = 1) = 1$ for every $\delta > 0$. Then $\beta_{23}(\cdot|x = 1)$ is only finitely additive. The resulting assessment (b, β) gives both players a payoff of zero (contrary to every perfect conditional ε -equilibrium when $\varepsilon > 0$ is small), but is nevertheless finitely consistent and sequentially 0-rational.²⁵ To see that this (b, β) is sequentially 0-rational, note first that, given 2’s behavior, it is optimal for player 1 to choose $x = 0$ no matter what θ he observes, and note second that, because player 2’s beliefs after observing $x = 1$ put probability 1 on the event that nature’s θ is greater than each real number less than 1, player 2’s expected payoff from choosing any $y > 0$ is -1 . Hence, player 2’s unique optimal choice is $y = 0$ after $x = 1$. And any choice of y is optimal for player 2 after $x = 0$.

²⁵Finite consistency can be verified with joint perturbations that, for nature, put positive probability on $\theta = 1 - \delta$ and, for player 1, put positive probability on $x = 1$ after observing $\theta = 1 - \delta$.

Example 7.2 *Finite games with probability zero states of nature.*

Perfect conditional ε -equilibrium can be applied to finite games that include states of nature that have probability zero. Kreps and Wilson (1982) only studied finite games in which all of nature's states have strictly positive probability.

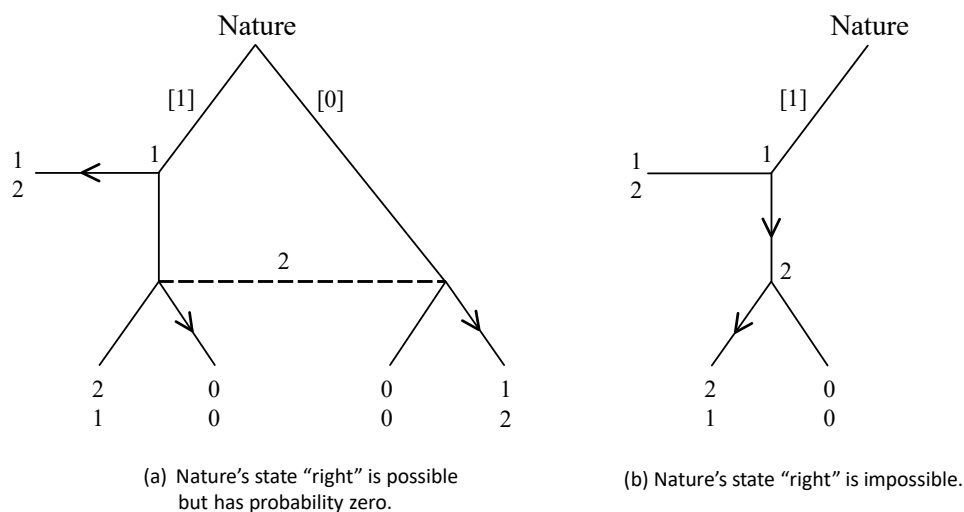


Figure 7.1

There are two multi-stage games depicted in Figure 7.1.²⁶ The games are identical except that, in panel (a), nature's state "right" is possible (although it has probability zero), while in panel (b), nature's state "right" is impossible (and so it is not present). The arrows indicate pure strategies.

The pure strategy profile in panel (b) is the unique subgame perfect equilibrium there and is also the unique limit of perfect conditional ε -equilibrium strategy profiles as $\varepsilon \rightarrow 0$ (because, in perturbed strategies that give all signals positive probability (nature's perturbations are equal to her original probability function), player 2 must choose "left" with probability at least $1 - \varepsilon$ and so player 1 must choose "down" with probability at least $1 - \varepsilon$ any conditional ε -equilibrium).

The pure strategy profile in panel (a) is also the limit of perfect conditional ε -equilibrium strategy profiles that are supported by perturbations in which the positive weight on player 1's action "down" is no more than twice the positive weight on nature's state "right."²⁷ Such

²⁶The players' action sets here are history-dependent (e.g., player 1's date 2 action set depends on nature's date 1 state). See footnote 12

²⁷The perfect conditional ε -equilibrium concept applied to finite games thus fits nicely with the discussion in Myerson (1991, p.189) on how one might extend sequential equilibrium to finite games in which nature's states can have probability zero.

perturbations define beliefs for player 2 in which, conditional on being reached, she believes the probability that nature chose “right” is at least $1/3$. This makes her choice of “right” sequentially rational.

Example 7.3 *Consequences of nonindependent perturbations of independent states of nature.*

The next example illustrates that some perturbations of nature can lead to perfect conditional ε -equilibria that may seem unintuitive.

- On date 1, nature chooses $\theta = (\theta_1, \theta_2)$ uniformly from the square $[-1, 3] \times [-1, 3]$.
- On date 2, player 1 observes θ_1 and chooses $x \in \{-1, 1\}$.
- On date 3, player 2 observes x and chooses $y \in \{-1, 1\}$.
- Payoffs are: $u_1 = xy$ and $u_2 = \theta_2 y$.

Since no player receives any information about θ_2 , and $\mathbb{E}(\theta_2) > 0$, player 2 should choose $y = 1$ regardless of the action of player 1 that she observes. But then player 1 should also choose $x = 1$ regardless of the value of θ_1 that he observes. Hence, the intuitively natural equilibrium expected payoff vector is $(u_1, u_2) = (1, 1)$.

But consider the pure strategy profile (b_{12}, b_{23}) where $b_{12}(\theta_1) = [-1]$ if $\theta_1 > -1$, $b_{12}(-1) = [1]$, and $b_{23}(x) = [-x]$.²⁸

This strategy profile yields the expected payoff vector $(u_1, u_2) = (-1, 1)$, but it is nonetheless a perfect conditional ε -equilibrium for any $\varepsilon > 0$ because it can be supported by a perturbation of nature that puts small positive probability on the event $\{(\theta_1, \theta_2) = (-1, -1)\}$. With this perturbation of nature it would be sequentially rational for player 2 to choose $y = -1$ when she observes $x = 1$ because she would attribute this observation to (θ_1, θ_2) being a mass point on $(-1, -1)$ and therefore would expect the value of θ_2 to be -1 .

This perfect conditional ε -equilibrium may seem unintuitive because, on the one hand θ_2 is observed by no one and is independent of everything in the game, yet, on the other hand player 2 believes that player 1’s action, which depends on the value of θ_1 , is informative about the value of θ_2 .

But one might argue that nothing in the description of the game explicitly says that θ_1 , the state observed by player 1, can never be informative about the value of θ_2 . When the game specifies that θ_1 and θ_2 are jointly uniform on $[-1, 3]^2$, this pins down the distribution of θ_2 conditional on θ_1 only for almost every value of θ_1 .²⁹ In particular, the game model does

²⁸Here, the notation $[c]$ denotes the probability measure that puts probability 1 on the action c .

²⁹This is according to the standard theory of conditional expectations.

not explicitly specify the distribution of θ_2 conditional on the value $\theta_1 = -1$.³⁰ Implicitly then, any conditional distribution is fair game (e.g., a mass point on $\theta_2 = -1$, as in the unintuitive equilibrium). The perfect conditional ε -equilibrium solution concept, with its arbitrary perturbations of nature, merely acknowledges this state of affairs.

If we want to model that it is common knowledge among the players that neither one of two independent coordinates of nature's state can ever be informative about the other, then we should restrict nature's perturbations so that these coordinates are perturbed independently. Such a restriction would eliminate the present unintuitive equilibrium. Sections 9 and 10 introduce two refinements that limit the effects of nature's perturbations.

Example 7.4 *Consequences of large perturbations of nature even with small probability.*

- On date 1, nature chooses $\theta = (\theta_1, \theta_2) \in [0, 1]^2$. With probability $1/2$, the coordinates are independent and uniform on $[0, 1]$, and with probability $1/2$ the coordinates are equal and uniform on $[0, 1]$.
- On date 2, player 1 observes $s_{12} = \theta_1$ and chooses $x \in \{-1, 1\}$.
- On date 3, player 2 observes $s_{23} = x$ and chooses $y \in \{-1, 1\}$.
- Payoffs are, $u_1 = xy$, and $u_2 = y(1/3 + \theta_2 - \theta_1)$.

Thus, player 2 should choose $y = 1$ if she expects $\theta_2 - \theta_1$ to be less than $-1/3$ and she should choose $y = -1$ otherwise. Player 1 wants to choose an action that player 2 will match.

Since for every θ_1 , θ_2 is equally likely to be equal to θ_1 (in which case $\theta_2 - \theta_1 = 0$) as to be uniform on $[0, 1]$ (in which case $\mathbb{E}(\theta_2 - \theta_1 | \theta_1) = 1/2 - \theta_1$), player 2 should expect $\theta_2 - \theta_1$ to be no smaller than $-1/4$, regardless of 1's strategy. So player 2 should choose $y = 1$.

Thus, it seems that all sensible equilibria involve strategies that give probability 1 to $(x, y) = (1, 1)$.

But consider the strategy profile (b_{12}, b_{23}) where $b_{12}(\theta_1) = [-1]$ iff $\theta_1 \neq -1$, and $b_{23}(x) = [1]$ iff $x = -1$.³¹ This profile gives probability 1 to $(x, y) = (-1, 1)$, and is supported in a

³⁰Notice that this would not be true if θ_2 were chosen after θ_1 . Then, the distribution of θ_2 would be specified by nature's transition probability function for any possible value of θ_1 . In this case, the game model could specify that θ_2 is uniform on $[-1, 3]$ for every possible θ_1 , which would eliminate the problem in this example. However, even then, the same problem would arise in a modified example with two additional players, 3 and 4, who, separately from players 1 and 2, play the same game, with player 3 playing the role of player 1 and player 4 playing the role of player 2, and where the roles of θ_1 and θ_2 are reversed, i.e., player 3 observes θ_2 , and player 4's payoff depends on θ_1 . In this modified game, the problem cannot be eliminated by specifying the temporal order in which θ_1 and θ_2 occur because each would have to occur before the other. But the refinements introduced in Sections 9 and 10 can eliminate the problem even in this modified example.

³¹See footnote 28.

perfect conditional ε -equilibrium by the perturbation of nature that never perturbs θ_2 but that with small positive probability perturbs the distribution of θ_1 so that it is a mass point on $\theta_1 = 1$. With this perturbation of nature it is conditionally rational for player 2 to choose $y = -1$ when she observes $x = 1$ because she attributes this observation to θ_1 being a mass point on 1 and therefore expects the value of $\theta_2 - \theta_1$ to be $-1/2$.

Once again, we have an unintuitive equilibrium that can result because the joint distribution of nature's state coordinates pins down the conditionals only almost everywhere. This unintuitive equilibrium can be eliminated if nature's states can be perturbed only to nearby states so as to approximately maintain the informativeness of each coordinate θ_1 and θ_2 about the other, or if we insist that the players strategies given positive probability to a large collection of events. See sections 9 and 10.

8 Regular Projective Games

In this section we introduce a large class of games – regular projective games – for which we can prove the existence of a perfect conditional ε -equilibrium and a perfect conditional equilibrium distribution.

Definition 8.1 *Let $\Gamma = (I, T, A, S, \mathcal{M}, p, \sigma, u)$ be a multi-stage game. Then Γ is a regular projective game iff there is a finite index set J and, for all $(n, r, j) \in I^* \times T \times J$ there are sets A_{nrj} such that, for every $it \in L$*

R.1. $A_{it} = \times_{j \in J} A_{itj}$ and $A_{0t} = \times_{j \in J} A_{0tj}$,

R.2. *if $t > 1$, then there is a nonempty set $M_{it} \subset I^* \times \{1, \dots, t-1\} \times J$ such that $S_{it} = \times_{nrj \in M_{it}} A_{nrj}$ and $\sigma_{it}(a_{<t}) = (a_{nrj})_{nrj \in M_{it}} \forall a_{<t}$ is a projection map; that is, i 's signal at date $t > 1$ is just a list of state coordinates and action coordinates from dates up to t ,*

R.3. A_{itj} and A_{0tj} are nonempty compact metric spaces $\forall j \in J$, and all product spaces, including S_{it} , are given their product topologies, and the measurable subsets of all spaces are their Borel subsets,

R.4. $u_i : A \rightarrow \mathbb{R}$ is continuous,

R.5. *there is a continuous nonnegative conditional density function $f_t : A_{0t} \times A_{<t} \rightarrow [0, \infty)$, and, for each j in J , there is a probability measure ρ_{tj} on $\mathcal{M}(A_{0tj})$ with full support on A_{0tj} such that $p_t(H|a_{<t}) = \int_H f_t(a_{0t}|a_{<t}) \prod_{j \in J} \rho_{tj}(da_{0tj}) \forall H \in \mathcal{M}(A_{0t}), \forall a_{<t} \in A_{<t}$.*

If Γ satisfies R.1 and R.2, we may say that Γ is a *projective game* or a *game with projected signals*.

Remark 8.2 (1) One can always reduce the cardinality of J to $(T+1)^{\#I}$ or less by grouping, for any $(i, t) \in I^* \times \{1, \dots, T\}$, the variables $\{a_{itj}\}_{j \in J}$ according to the $\#I$ -vector of dates at which the players observe them, if ever.

(2) Regular projective multi-stage games can include all finite multi-stage games simply by letting each player's signal be a coordinate of the state.

(3) Since distinct players can observe the same a_{0tj} , nature's probability function in a regular projective multi-stage game need not satisfy the information diffuseness assumption of Milgrom-Weber (1985). Nevertheless, the form of p_t assumed in R.5 of Definition 8.1 is reminiscent of the Milgrom-Weber assumption, and a recent counterexample to the existence of an (ex-ante) Nash ε -equilibrium in a Bayesian game due to Simon and Tomkowicz (2017) shows that some such assumption is necessary for the existence of even a conditional ε -equilibrium.

(4) Perfect recall means here that for all players $i \in I$, for all dates $r < t$ and for all $j \in J$, $M_{ir} \subseteq M_{it}$ and $irj \in M_{it}$.

8.1 Existence

Theorem 8.3 Let Γ be a regular projective game. Then for any $\varepsilon > 0$, Γ has a perfect conditional ε -equilibrium.

By Theorem 5.1, Theorem 8.3 implies the following result.³²

Theorem 8.4 Let Γ be a regular projective game. Then for any $\varepsilon > 0$, Γ has a subgame perfect ε -equilibrium.

An immediate consequence of Theorem 4.9, Remark 4.11, and Theorem 8.3 is the following.

Theorem 8.5 Every regular projective game Γ has a perfect conditional equilibrium distribution μ . Moreover, every perfect conditional equilibrium distribution μ of Γ induces a regular countably additive distribution ν .

In the next two sections, we introduce two refinements that can limit the effects of nature's perturbations. Both refinements can be applied in any multi-stage game, even if it is not

³²See Chakrabarti (1999) for an existence result concerning a related concept, subgame perfect approximate equilibria, for a different class of games.

a regular projective game, but both refinements require the introduction of a topology. In the first refinement a topology is introduced on nature's states. In the second refinement, a topology is introduced on the players' action spaces.

9 Local Perturbations of Nature

The unintuitive perfect conditional ε -equilibria in Examples 7.3 and 7.4 can be eliminated by restricting the class of perturbations of nature that can be used to test for conditional ε -equilibrium behavior in the perturbed games.

For example, in example 7.3 we may wish to capture that nature's two independent coordinates can never be informative about one another, no matter what zero probability event might occur, and, in example 7.4, we may wish to capture that the conditional distribution of θ_2 given *any possible* value of θ_1 (not merely almost every possible value of θ_1) is a mass point on θ_1 with probability 1/2 and is uniform on $[0, 1]$ with probability 1/2. Both of these natural conditions will be satisfied when nature is restricted to a class of "local" perturbations, which we now define.

Let Γ be any multi-stage game. But suppose that there is a finite index set, J , such that, for any date t , nature's set of date t states can be written as is $A_{0t} = \times_{j \in J} A_{0tj}$, where each A_{0tj} is a metric space.³³ For any element x of any metric space, and for any $\delta > 0$, let $\mathbb{B}_\delta(x)$ denote the δ -ball centered at x .

For date $t \leq T$ and for any $j \in J$, let Φ_{tj} denote the set of transition probabilities $\phi_{tj} : A_{0tj} \rightarrow \Delta(A_{0tj})$.

Let $\Phi = \times_{t \leq T} \times_{j \in J} \Phi_{tj}$. For any $\phi \in \Phi$, define the perturbation of nature $p * \phi \in \mathcal{T}$ as follows. For every date $t \leq T$, for every $a_{<t} \in A_{<t}$, and for $C = \times_{j \in J} C_j \in \times_{j \in J} \mathcal{M}(A_{0tj})$,

$$[p * \phi]_t(C|a_{<t}) = \int_{A_{0t}} \prod_{j \in J} \phi_{tj}(C_j|a_{0tj}) p_t(da_{0t}|a_{<t}).$$

The perturbation $p * \phi$ works as follows. At each date t , and after any history $a_{<t} \in A_{<t}$, a provisional state a_{0t} is first drawn according to nature's date- t probability measure $p_t(\cdot|a_{<t})$. Then, independently for each coordinate $j \in J$, the j -th coordinate of the actual date t state is drawn according to the distribution $\phi_{tj}(\cdot|a_{0tj})$, depending only on the j -th coordinate of provisional state.

Say that $\phi \in \Phi$ is δ -constant iff for every date $t \leq T$, for every $j \in J$, and for every

³³Since we can always let J be a singleton set, such an index set J exists for any multi-stage game. But the restrictions on nature's perturbations that are defined below will be tighter when nature's dated state spaces have a nontrivial product structure that is recognized by a nontrivial index set J .

$a_{0tj} \in A_{0tj}$,

$$\phi_{tj}(\{a_{0tj}\}|a_{0tj}) \geq 1 - \frac{\delta}{\#J}.$$

We note that if ϕ is δ -constant, then $p * \phi$ is a δ -perturbation of p (i.e., $\|p * \phi - p\| \leq \delta$).

Say that $\phi \in \Phi$ is δ -local iff for every date $t \leq T$, for every $j \in J$, and for every $a_{0tj} \in A_{0tj}$,

$$\phi_{tj}(\mathbb{B}_\delta(a_{0tj})|a_{0tj}) = 1.$$

When ϕ is δ -local, $p * \phi$ perturbs nature's state at each date to a nearby state, each of whose coordinates is within δ of the original provisional state-coordinate.

So, for small $\delta > 0$, if $\phi \in \Phi$ is δ -constant and δ -local, then $p * \phi$ perturbs nature's coordinates independently to nearby values, and only rarely.

Say that b is a *perfect conditional ε -equilibrium with local perturbations of nature* iff b satisfies the conditions of Definition 4.3 with the additional restriction that the δ -perturbation of nature p' must be of the form $p' = p * \phi$, where ϕ is δ -constant and δ -local.

We can state an existence result for perfect conditional ε -equilibria with local perturbations of nature when Γ is a regular projective game.

Theorem 9.1 *Let Γ be a regular projective game. Then for any $\varepsilon > 0$, Γ has a perfect conditional ε -equilibrium with local perturbations of nature.*

To eliminate the unintuitive perfect conditional equilibria in example 7.3, we should set $J = \{1, 2\}$ and let $A_{011} = A_{012} = [-1, 3]$. Then $\theta = (\theta_1, \theta_2) \in A_{011} \times A_{012}$, and in any local perturbation of nature, the coordinates θ_1 and θ_2 of nature's date-state θ will be perturbed independently. Moreover, in any local perturbation in which some state $\bar{\theta} = (\bar{\theta}_1, \bar{\theta}_2)$ receives positive probability, the conditional distribution of θ_1 given $\theta_2 = \bar{\theta}_2$ will be uniformly close to a uniform distribution on $[-1, 3]$. Consequently, with this specification of the coordinates of nature, the unintuitive equilibrium fails to be a perfect conditional ε -equilibrium with local perturbations of nature. A similar conclusion holds for the unintuitive equilibrium in example 7.4.

Local perturbations of nature appear to be especially effective in regular projective games at eliminating the unintuitive effects of nature's perturbations.³⁴ We consider it a high priority of future research to better understand the effects of perturbing nature in general games and how nature's perturbations should be restricted so as leave the informational content of the players' signals approximately unchanged.

³⁴In a regular projective game, the continuity of nature's density function ensures that, in any local perturbation of nature, no coordinate of nature's state becomes much more informative about any other coordinate, at least on an ε -sure signal domain where nature's density can be bounded away from zero.

10 Full Support Strategies

We now consider another refinement of conditional ε -equilibrium which, when applied together with perfectness, can also limit the effects of nature's perturbations.

Given a strategy profile b , we may say that an observable event $Z \in \mathcal{M}(S_{it})$ could be relevant for a deviation by some player j to an alternative strategy c_j iff the event Z would have positive probability if player j deviated in this way, that is, $P_{it}(Z|c_j, b_{-j}) > 0$. The basic idea of sequential equilibrium is that every player's strategic responses in the equilibrium should be verified as rational in any observable event that could be relevant for any possible deviation by any player.

We would like to reduce the need for nature's perturbations, if not for all events whose probabilities could become positive under any deviations, at least for all events that whose probabilities could become positive for a dense class of deviations. But to talk about a "dense" set of deviation strategies, we need a concept of topology on players' action sets.

So, let Γ be a multi-stage game as defined in Section 3, but suppose now that, for each $it \in L$, the action set A_{it} is a separable metric space, and the measurable sets $\mathcal{M}(A_{it})$ are the Borel sets. We say that a strategy profile b has *full support* iff, $\forall it \in L, \forall s_{it} \in S_{it}, b_{it}(C|s_{it}) > 0$ for every nonempty open subset C of A_{it} .

Full-support strategies exist, by the assumption that the topology on each A_{it} is separable, as each dated player has a countable dense set of actions which could all be given positive probability. Notice that, if $\hat{b} \in B$ has full support then, $\forall b \in B, (1 - \varepsilon)b + \varepsilon\hat{b}$ also has full support when $0 < \varepsilon < 1$. So any strategy has arbitrarily-small-probability perturbations with full support.

Say that a conditional ε -equilibrium b is *full* iff it has full support (with the given topologies).

For any player i , given any strategy $c_i \in B_i$ and any $\varphi = (\varphi_1, \dots, \varphi_T)$ such that $\varphi_t : A_{it} \times S_{it} \rightarrow \Delta(A_{it})$ is a transition probability for each $t \leq T$, we let $c_i * \varphi$ denote the strategy such that, $\forall t \leq T, \forall s_{it} \in S_{it}, \forall C \in \mathcal{M}(A_{it})$,

$$(c_i * \varphi)_t(C|s_{it}) = \int_{A_{it}} \varphi_t(C|a_{it}, s_{it}) c_{it}(da_{it}|s_{it}).$$

This φ is a δ -local iff

$$\varphi_t(\mathbb{B}_\delta(a_{it})|a_{it}, s_{it}) = 1, \forall a_{it} \in A_{it}, \forall s_{it} \in S_{it}, \forall it \in L,$$

where $\mathbb{B}_\delta(a_{it})$ is the ball of radius δ around a_{it} .

The following theorem is a formal statement of the claim that, in any full conditional

ε -equilibrium, perturbations (by nature or by any player) are not needed to verify strategic rationality in all events whose probabilities could become positive for a “dense” set of deviations.

Theorem 10.1 *If the strategy profile b has full support, then for any player $i \in I$, for any deviation $c_i \in B_i$, and for any $\delta > 0$, there is a δ -local φ such that $P(\cdot|c_i * \varphi, b_{-i})$ is absolutely continuous with respect to $P(\cdot|b)$, i.e.,*

$$\{H \in \mathcal{M}(A) : P(H|c_i * \varphi, b_{-i}) > 0\} \subseteq \{H \in \mathcal{M}(A) : P(H|b) > 0\}.$$

That is, any deviation by any player has arbitrarily small local perturbations that do not generate any positive-probability events beyond those events that have positive probability under b . So for a full conditional ε -equilibrium b , any deviation can be locally approximated by perturbed deviations such that, in all positive-probability consequences of the perturbed deviations, ε -rationality of responses under b can be verified without perturbing nature or other strategies.

The full conditional ε -equilibrium concept is attractive because it does not require any perturbations of nature. However, full conditional ε -equilibria need not be subgame perfect ε -equilibria because, even in games without nature, rationality is not tested everywhere.³⁵ To test rationality everywhere, we need to consider nets of perturbations of the players’ strategies and of nature’s probability function, as in our perfectness concept. But then we can limit the effects of nature’s perturbations by requiring the players to use full-support strategies, so that nature’s perturbations will be irrelevant for testing the players’ rationality for the events that have positive probability in the full-support strategy profile.

So, for any multi-stage game Γ in which the players’ action spaces are metrized, we define a *full perfect conditional ε -equilibrium* to be a perfect conditional ε -equilibrium strategy profile that has full support. Our final result says that full perfect conditional ε -equilibria exist in regular projective games.

Theorem 10.2 *Let Γ be a regular projective game. Then for any $\varepsilon > 0$, there is a perfect conditional ε -equilibrium strategy profile that has full support.*

³⁵For example, suppose that player 1 chooses from $[0, 1]$ and then player 2, after observing 1’s choice, also chooses from $[0, 1]$. They each receive a payoff of 1 if they both choose action 1, but they each receive a payoff of 0 otherwise. All subgame perfect ε -equilibria have both players receiving a payoff of at least $1 - 2\varepsilon$. But there is a full conditional ε -equilibrium in which both players choose their actions uniformly from $[0, 1]$, and player 2 does so no matter what action of player 1 she observes. Both players’ payoffs are zero in this equilibrium. Despite the full-support strategies in this equilibrium, player 2’s rationality is not tested, in particular, at the crucial signal that she would observe were player 1 to choose action 1.

11 Conclusion

In order to ensure that *all* off-path behavior is rational in an infinite game, we have been led to perturb not only the players' strategies (as in Kreps and Wilson, 1982), but to perturb nature's probability function as well. Although the effects of nature's perturbations can sometimes seem unintuitive, the strategy profiles that arise as perfect conditional ε -equilibria have been defined so that they possess two fundamental properties, namely, for any finite collection of signal events in the game that are on or off the equilibrium path, (i) (*finite consistency*) all players agree on a common perturbation of nature's probability function and on a common perturbation of the player's equilibrium strategies that together give positive probability to – and so explain the occurrence of – each signal event in the collection, and (ii) (*conditional ε -optimality*) given this common explanation, each player's perturbed equilibrium continuation strategy is ε -optimal conditional on each of his signal events in the collection.

In finite games, the sets of conditional ε -equilibria with full support (introduced in the previous section) and perfect conditional ε -equilibria are essentially equivalent, and their limits yield the set of sequential equilibrium strategy profiles.³⁶ The fact that this coincidence of perfectness and fullness does not extend to infinite games is a basic reason why it has been so difficult to define sequential equilibria for infinite games. An uncountable infinity of outcomes cannot all get positive probability from one strategy profile, and so one must either let the strategy profile satisfy a weaker topological condition of full support, or one must consider a net of perturbed strategies that can test rationality in all events but may yield only finite additivity in the limit. We have emphasized the latter approach here.

12 Proofs. (to be added)

13 Appendix.

In this appendix, we compute perfect conditional ε -equilibria for the motivating examples of Section 2.

Example 2.1. Let b be any perfect conditional ε -equilibrium of the game in Example 2.1. Notice first that, for any $\varepsilon \in [0, 1)$, an ε -sure signal domain must include every signal (x, y)

³⁶Specifically, let Γ^f be any finite extensive-form game such that all alternatives have positive probability at each chance node (as KW assumed). Any full conditional ε -equilibrium (with the discrete topology on the finite A_{it}) of Γ^f is a perfect conditional ε -equilibrium. Conversely, if b is a perfect conditional ε -equilibrium of Γ^f , then for all $\delta > 0$ and for all $\varepsilon' > \varepsilon$, there is a full conditional ε' -equilibrium b' that is a δ -perturbation of b .

in the signal sets of players 3 and 4 (because any particular signal can be made to occur with probability 1 by a suitable choice of pure strategies for players 1 and 2). So for any signal $(x, y) \in [0, 1] \times \{-1, 1\}$, there must be a perturbation b' of b (nature is absent in this example) that gives (x, y) positive probability and against which players 3 and 4 are ε -optimizing conditional on its occurrence. Hence, after observing any $(x, y) \in [0, 1] \times \{-1, 1\}$, if $x \neq 0$, players 3 and 4 must each, with at least probability $1 - \varepsilon/2|x|$, choose -1 if $x < 0$ and choose $+1$ if $x > 0$. To prevent player 2 from matching player 3, player 1 must make player 3 to randomize nearly uniformly, which 1 can do while keeping the $-x^2$ component of his payoff from being very negative, by randomizing over positive and negative x 's that are small in absolute value, giving the positive x 's nearly the same total probability as the negative x 's.

Example 2.2. For this example, let us first show that player 1's expected payoff must converge to zero as $\varepsilon \rightarrow 0$ in any sequence of perfect conditional ε -equilibria. So, consider any such sequence. For each state $\theta = 1, 2$, define q_θ to be the limiting probability that player 2 chooses $y = 1$ conditional on state θ (extract a convergent subsequence if necessary). Since player 1's payoff in this game is nonzero if and only if player 2 chooses $y = 1$, we must show that $q_1 = q_2 = 0$. Player 2 can obtain an expected payoff of $3/4$ by always choosing $y = 2$, and so 2's limit equilibrium payoff, $(1/4)(q_1)1 + (3/4)(1 - q_2)1$, must be at least $3/4$,³⁷ which means that $q_1 \geq 3q_2$. Player 1's limit equilibrium payoff, $(1/4)(q_1)1 + (3/4)(q_2)1$, cannot be less than the limit of the payoffs that he would achieve if, along the sequence, he always deviated to the square root of the action that his equilibrium strategy dictated. Since such deviations along the sequence would yield player 1 a limit payoff of at least $(3/4)q_1$,³⁸ we must have $(1/4)(q_1)1 + (3/4)(q_2)1 \geq (3/4)q_1$, which is equivalent to $3q_2 \geq 2q_1$. But because $q_1 \geq 3q_2$, this means that $q_1 = q_2 = 0$ as desired.

Next we display a strategy profile, b^* , that, for any $\varepsilon > 0$, is a perfect conditional ε -equilibrium, where the signal domain is the entire set of signals in the game. Let player 1's strategy, b_1^* , choose action $x = 0$ with probability 1, and let player 2's strategy b_2^* choose action $y = 2$ regardless of the signal that she observes. For any $\delta \in (0, 1)$ that is less than ε , and for any finite set of signals, $Z \subseteq [0, 1]$ for player 2, let Z' be the smallest set containing Z that is closed under the taking of square roots (i.e., $Z' = \cup_{s \in Z} \{s, \sqrt{s}, \sqrt{\sqrt{s}}, \dots\}$), and consider the perturbations $b' \in B$ of b^* and $p' \in \mathcal{T}$ of p , where we perturb only player 1's strategy b_1^* to b'_1 and set $b'_2 = b_2^*$ and $p' = p$. Define the perturbation b'_1 of b_1^* so that

³⁷Otherwise, for some $\varepsilon > 0$ small enough along the sequence, player 2 would not be ε -optimizing conditional on the null state at date 1.

³⁸Because, under player 1's square root deviation, when the state is $\theta = 2$ player 2 chooses $y = 1$ with limiting probability q_1 , and when the state is $\theta = 1$ player 1's payoff is nonnegative no matter what action player 2 chooses.

$b'_1(\{0\}) = 1 - \delta$, $b'_1(Z') = \delta$, and $b'_1(x) = 3b'_1(\sqrt{x})$ for every $x \in Z'$. Then, $\|b' - b^*\| \leq \delta$, $\|p' - p\| = 0 \leq \delta$, and, in the perturbed game $\Gamma(p') (= \Gamma(p))$, b' gives probability 1 to the set of player 2 signals $Z' \cup \{0\}$. Since conditional on any positive signal in Z' that is less than 1, the states $\theta = 1$ and $\theta = 2$ are equally likely under b' in $\Gamma(p')$, it is conditionally optimal for player 2 to choose $y = 2$ after observing any such signal. And since the signals $s = 0$ and $s = 1$ are uninformative about nature's state, it is also conditionally optimal for player 2 to choose $y = 2$ after observing $s = 0$ or 1. Since player 1's payoff in $\Gamma(p')$ under b' is $1 - \delta \geq 1 - \varepsilon$, player 1 is ε -optimizing. Hence, b' is a perfect conditional ε -equilibrium of $\Gamma(p')$, and so b^* is a perfect conditional ε -equilibrium of $\Gamma(p)$.³⁹

Example 2.3. It is a standard result from information economics that, in any Nash equilibrium of this game, player 2 must reject all strictly positive offers. Similar techniques establish that, for any action $x > 0$ of player 1, the probability that player 2 accepts x (i.e., chooses $y = 1$ given x) tends to zero in any sequence of perfect conditional ε -equilibria as $\varepsilon \rightarrow 0$. So, we will be content here to display one of the many possible perfect conditional ε -equilibria of this game. One such equilibrium, for any $\varepsilon > 0$ and with the entire set of signals as the signal domain, has player 1 put probability one on $x = 2$ for every θ and for player 2 to reject every offer (i.e., choose $y = 0$ given any offer x). For any $\delta \in (0, 1)$ that is less than ε , for any finite subset Z_1 of player 1's signals, and for any finite subset Z_2 of player 2's signals, consider a perturbation of the players' strategies and of nature's probability function where we perturb player 1's strategy to b'_1 and we perturb nature's probability function to p' , but we do not perturb 2's strategy. The perturbations b'_1 and p' are defined as follows. Let $b'_1(\{x = 2\}|\theta) = 1$ for every $\theta \in [\delta, 1]$, and, for every $\theta \in [0, \delta)$, let $b'_1(\{x = 2\}|\theta) = 1 - \delta$ and $b'_1(\{x\}|\theta) = \delta/(\#Z_2)$ for each $x \in Z_2$. Define p' to be uniform on $[0, 1]$ with probability $1 - \delta$ and to be uniform on Z_1 with probability δ . Then, in the perturbed game $\Gamma(p')$, it is ε -optimal for player 2 to choose $y = 0$ after observing any signal (and hence any event) that has positive probability (since, if $\theta \in [0, \delta)$, 2's utility is less than $\theta \leq \varepsilon$), and, any action $x \in [0, 1]$ for player 1 is optimal since his payoff, for any θ , is zero no matter which $x \in [0, 1]$ he chooses. Hence, in the perturbed game, the perturbed strategy profile is conditionally ε -rational and gives every signal in $Z_1 \cup Z_2$ positive probability. Therefore, the original strategy profile is a perfect conditional ε -equilibrium.

Example 2.4. We claimed that all of the Nash equilibria of this game are reasonable and it is not difficult to show that all of them are perfect conditional ε -equilibria for any $\varepsilon > 0$. We shall show this for the the Nash equilibrium in which player 1 chooses x uniformly from $[0, 1]$, and player 2 chooses $y = 2$ regardless of the signal that she observes. For any $\delta > 0$

³⁹There are many perfect conditional ε -equilibria of this game (e.g., player 1 chooses x uniformly from $[0, 1]$ and player 2 always chooses $y = 2$ regardless of the signal that she observes).

and for any finite set of signals $Z \subseteq [0, 1]$ for player 2, let us perturb only nature's probability function to p' , where p' chooses θ_1 and θ_2 independently, with θ_1 uniform on $\{1, 2\}$, as before, but with θ_2 uniform on $[0, 1]$ with probability $1 - \delta$ and uniform on Z with probability δ . In the perturbed game, the only new positive probability signal events are those for player 2 that contain signals in Z . But conditional on any such signal the probability that $\theta_2 = 2$ is one and so it is optimal for player 2 to choose $y = 2$. Hence, in the perturbed game, the given unperturbed strategies are optimal (and so ε -optimal) conditional on all positive probability signal events and so form a conditional ε -equilibrium. Hence, the given strategies constitute a perfect conditional ε -equilibrium of the original game.

References

- Aumann, R. (1961): "Mixed and Behavior Strategies in Infinite Extensive Games," *Economic Research Program*, Research memorandum #32, Princeton University.
- Bertsekas, D. P., and S. E. Shreve (1978): *Stochastic Optimal Control: The Discrete Time Case*, Academic press, London. Available for download at <http://dspace.mit.edu/handle/1721.1/4852>.
- Bajoori, Elnaz, János Flesch, and Dries Vermeulen (2013), "Perfect equilibrium in games with compact action spaces," *Games and Economic Behavior*, 82, 490-502.
- Bajoori, Elnaz, János Flesch, and Dries Vermeulen (2016), "Behavioral perfect equilibrium in Bayesian games," *Games and Economic Behavior* 98, 78-109.
- Billingsley, P., (1968): *Convergence of Probability Measures*, Wiley Series in Probability and Mathematical Statistics, New York.
- Börgers, T. (1991): "Upper Hemicontinuity of the Correspondence of Subgame-Perfect Equilibrium Outcomes," *Journal of Mathematical Economics*, 20, 89–106.
- Chakrabarti, S. K. (1999): "Finite and Infinite Action Dynamic Games with Imperfect Information," *Journal of Mathematical Economics*, 32, 243–266.
- Cohn, D. L., (1980): *Measure Theory*, Birkhauser, Boston.
- Dunford, N., and J. T. Schwartz (1988) *Linear Operators Part I General Theory*, John Wiley and Sons, New York.
- Fudenberg, D. and J. Tirole (1991): "Perfect Bayesian and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236-260.

- Harris, C. J. , P. J. Reny, and A. J. Robson (1995): “The existence of subgame-perfect equilibrium in continuous games with almost perfect information,” *Econometrica* 63, 507-544.
- Harris, C. J., M. B. Stinchcombe, and W. R. Zame (2000): “The Finitistic Theory of Infinite Games,” University of Texas working paper. <http://www.laits.utexas.edu/~maxwell/finsee4.pdf>
- Hellman, Z. (2014): “A game with no Bayesian approximate equilibria,” *Journal of Economic Theory*, 153, 138-151.
- Hellman, Z. and Y. Levy (2013): “Bayesian games with a continuum of states,” Hebrew University working paper.
- Kelley, J. L. (1955), *General Topology*, Springer-Verlag.
- Kreps, D. and R. Wilson (1982): “Sequential Equilibria,” *Econometrica*, 50, 863-894.
- Kuhn, H. (1953): “Extensive games and the problem of information,” in H. Kuhn & A. Tucker, eds. ‘Contributions to the Theory of Games,’ Vol. II, *Annals of Mathematics Study* 28. Princeton University Press, 1953, pp. 193–216.
- Manelli, A. (1996): “Cheap Talk and Sequential Equilibria in Signaling Games,” *Econometrica* 64, 917-942.
- Milgrom, P. and R. Weber (1985): “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research*, 10, 619-32.
- Myerson, R. B. (1991): *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, Massachusetts.
- Radner, R. (1980): “Collusive Behavior in Non-Cooperative Epsilon-Equilibria of Oligopolists with Long But Finite Lives,” *Journal of Economic Theory*, 22, 136–154.
- Simon, Leo K., and Maxwell B. Stinchcombe (1995), “Equilibrium refinement for infinite normal-form games,” *Econometrica*, 63 (6), 1421-1443.
- Simon, R. S. (2003): “Games of incomplete information, ergodic theory, and the measurability of equilibria,” *Israel Journal of Mathematics* 138, 73-92.
- Simon, R. S. and G. Tomkowicz (2017): “A Bayesian Game without ε -Equilibria,” *Israel Journal of Mathematics* (to appear).
- von Neumann, J., and O. Morgenstern (1944): *Theory of Games and Economic Behavior*, Princeton University Press, Princeton.

Watson, J. (2017): “A General, Practicable Definition of Perfect Bayesian Equilibrium,”
UCSD working paper.