

Discussion Paper Series – CRC TR 224

Discussion Paper No. 703
Project B 02

Sequential Equilibria in Mixed Strategies

Francesc Dilmé¹

October 2025

¹ University of Bonn. Email: fdilme@uni-bonn.de.

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

Sequential Equilibria in Mixed Strategies

Francesc Dilmé*

A Nash equilibrium of a game in extensive form is a *sequential equilibrium in mixed strategies* if it can be approximated through equilibria of close-by games with slightly perturbed payoffs and small-probability behavioral types. We show that sequential equilibria in mixed strategies are equivalent to (i) weakly sequential equilibria (Reny, 1992), (ii) normal-form perfect equilibria (Selten, 1975) in games with generic payoffs, and (iii) purifiable Nash equilibria (Harsanyi, 1973). A corollary of our results is that extensive-form perfect equilibria are normal-form perfect equilibria in games with generic payoffs.

Since its introduction by Kreps and Wilson (1982), sequential equilibrium has been influential in both theoretical and applied work. Its definition requires considering a belief system along with a strategy profile, and demands both consistency and sequential rationality. Nevertheless, Kreps and Wilson offer an insightful equivalent characterization: a strategy profile is a sequential equilibrium if and only if it is the limit of Nash equilibria of similar games with perturbed payoffs and mistakes. This characterization lends plausibility to sequential equilibria, as small perturbations of payoffs and mistakes are likely in practice (and possibly difficult for the econometrician to observe); sequential equilibria are those approximated by equilibria of such perturbed games. Importantly, it also permits extending the concept of sequential equilibria to strategy spaces beyond behavior strategies.

*September 30, 2025. University of Bonn. fdilme@uni-bonn.de. This work was funded by a grant from the European Research Council (ERC 949465). Support from the German Research Foundation (DFG) through CRC TR 224 (Project B02) and under Germany's Excellence Strategy (EXC 2126/1-390838866 and EXC 2047-390685813) is gratefully acknowledged.

In this paper, we define and analyze sequential equilibria in mixed strategies, using Kreps and Wilson’s (1982) characterization, but perturbing the space of mixed strategies instead of behavior strategies. These perturbations are interpreted as small probabilities of a player’s irrationality (“behavioral types”). We find that sequential equilibria in mixed strategies are closely related to weakly sequential equilibria; we characterize their relationship with perfect equilibria; and we show that, remarkably, they coincide with the set of purifiable Nash equilibria. Let us elaborate.

For our analysis, we fix a game in extensive form. A strategy profile is a *sequential equilibrium in mixed strategies* if it is the limit of Nash equilibria of a sequence of perturbed games converging to the unperturbed game, where players “tremble” in their choice of strategy and where payoffs are also perturbed. A natural interpretation of the tremble involves behavioral types: at the outset, nature independently decides whether each player is rational (with the corresponding perturbed payoffs) or behavioral (and plays according to a prespecified fully mixed strategy).¹ Sequential equilibria in mixed strategies differ from sequential equilibria in that a deviation from equilibrium behavior signals that a player is irrational rather than that she is rational but has made a mistake.

We first show the equivalence between outcomes of sequential equilibria in mixed strategies and outcomes of weakly sequential equilibria (Reny, 1992).² Recall that weakly sequential equilibria are consistent assessments where a player’s sequential rationality is not required at her information sets that her strategy cannot reach. Our result is intuitive given our interpretation of trembles as behavioral types, as Reny’s main motivation is that deviators should not be perceived as rational players (even though he does not explicitly consider trembles). It then follows that outcomes of sequential equilibria are outcomes of sequential equilibria in mixed strategies. We also provide a characterization of sequential equilibria in mixed strategies that mimics Kreps and Wilson’s (1982) definition: a strategy profile is a sequential equilibrium in mixed strategies if and only if it is part of a consistent and weakly sequentially rational *conditional assessment*, where beliefs at information sets

¹Equivalently, we could have many behavioral types (one for each pure strategy of the player) and nature choosing one of them with a small probability.

²Because weakly sequential equilibria are defined using behavioral strategies and belief systems, it is appropriate to compare the two concepts through the distributions over terminal histories they induce.

concern the continuation strategy profile rather than histories.

Next, we define perfect equilibria in mixed strategies as those approximated by Nash equilibria where players' strategies are perturbed (but not their payoffs). Their outcomes coincide with those of perfect equilibria in the normal form of the game. We show that perfect equilibria in mixed strategies are sequential equilibria in mixed strategies, and that the two concepts coincide for generic payoffs. This result extends the analogous result by Kreps and Wilson (1982) for behavior strategies.

A corollary of our results is that outcomes of extensive-form perfect equilibria are outcomes of perfect equilibria in mixed strategies. This result is intuitive because perfect equilibria are approximated by perturbed strategy profiles that are sequentially rational at all information sets, whereas perfect equilibria in mixed strategies require only ex-ante optimality. However, because the set of perturbations that define them differs, these concepts are not universally ranked: outcomes of perfect equilibria in mixed strategies are not always outcomes of perfect equilibria, and vice versa.

Finally, we show that sequential equilibria in mixed strategies coincide with the set of purifiable Nash equilibria (Harsanyi, 1973). Recall that a Nash equilibrium is purifiable if it can be approximated by Nash equilibria in a sequence of games where players' payoffs are independently perturbed according to an atomless distribution with a large support.³ Hence, while purifiability has a limited effect in normal-form games (where all Nash equilibria are purifiable), it exerts significant selection power in extensive games, as all information sets are reached with positive probability.

Overall, we argue that sequential equilibria in mixed strategies are the natural analog of sequential equilibria, both in terms of their definition and properties. They are equilibria of similar games with small-probability behavioral types; generically coincide with perfect equilibria; and can be purified by payoff perturbations. Sequential equilibria in mixed strategies are thus well suited as an equilibrium concept for environments with small uncertainty about other players' rationality or payoffs.

³By a "large support" we mean that, for each strategy of a player, there is a payoff in the support where this strategy is strictly dominant.

The paper is organized as follows. Section 1 provides the notation and basic definitions for the paper. Section 2 introduces sequential equilibria in mixed strategies and Section 3 characterizes their relationship with weakly sequential equilibria. Section 4 introduces perfect equilibria in mixed strategies and relates them to sequential equilibria in mixed strategies. Section 5 establishes the equivalence between sequential equilibria in mixed strategies and purifiable Nash equilibria. Section 6 concludes. Appendix A introduces ε -sequential equilibria in mixed strategies, a concept useful for the proofs of the previous results, which are placed in Appendix B. Finally, Appendix C relates our concepts to quasi-perfect equilibria (van Damme, 1984) and normal-form sequential equilibria (Mailath et al., 1993).

1 Games in extensive form

We define and introduce notation for an extensive-form game with perfect recall. The definitions are standard and the notation is summarized in Table 1.

A (finite) *game* $G := \langle A, H, \mathcal{I}, N, \iota, \pi, u \rangle$ has the following components: (1) A finite set of *actions* A . (2) A finite set of *histories* H . Here, a history is a finite sequence of actions $h \equiv (h_j)_{j=1}^{|h|}$ (where $|h|$ denotes the length of history h), and the set H has the property that if $h \equiv (h_j)_{j=1}^{|h|} \in H$ with $|h| > 0$, then $(h_j)_{j=1}^{|h|-1} \in H$ as well. (In particular, $\emptyset =: (h_j)_{j=1}^0 \in H$.) The set of *terminal histories* is Z . (3) An *information partition* \mathcal{I} , that is, a partition of $H \setminus Z$ such that there is a partition $\{A^I | I \in \mathcal{I}\}$ of A with the property that, for each $I \in \mathcal{I}$ and $h \in H$, we have $(h, a) \in H$ for some $a \in A^I$ if and only if $h \in I$. The elements of \mathcal{I} are called *information sets*.⁴ (4) A finite set of *players* $N \not\equiv \emptyset$. (5) A *player assignment* $\iota : \mathcal{I} \rightarrow N \cup \{0\}$, assigning each information set to a player or to nature (represented by 0), and satisfying perfect recall.⁵ (6) A *nature's probability assignment* $\pi : \bigcup_{I \in \iota^{-1}(\{0\})} A^I \rightarrow (0, 1]$ satisfying $\sum_{a \in A^I} \pi(a) = 1$ for each $I \in \iota^{-1}(\{0\})$. (7) For each player $i \in N$, a (von Neumann–Morgenstern) *payoff function* $u_i : Z \rightarrow \mathbb{R}$. For convenience, we set $u_0(z) = 0$ for all $z \in Z$.

⁴We assume, without loss of generality, that each action is available at only one information set; actions can always be renamed to ensure this.

⁵Perfect recall implies that for all $I, I' \in \mathcal{I}$ with $\iota(I) = \iota(I')$ and all $h, \hat{h} \in I$, if $(h', a) \preceq h$ for some $h' \in I'$ and $a \in A$, then $(\hat{h}', a) \preceq \hat{h}$ for some $\hat{h}' \in I'$. Here, $(h', a) \preceq h$ indicates that (h', a) precedes or equals h .

$i \in N$	player i	$s_i \in S_i$	i 's pure strategies	$G(\xi, u)$	perturbed game
$a_i \in A_i$	i 's actions	$\sigma_i \in \Sigma_i$	i 's mixed strategies	$\mathcal{I}_i^{s_i}$	s_i -relevant inf. sets
$h \in H$	histories	$\beta_i \in \mathcal{B}_i$	i 's behavior strat.	S_i^I	i 's strategies for
$z \in Z$	terminal histories	μ	belief system		which I is relevant
$I \in \mathcal{I}_i$	i 's info. sets	ξ	tremble	γ	cond. belief syst.
u_i	i 's payoff	\mathbb{P}^σ	prob. under σ	ν	payoff prob. measure

Table 1 Notation

A *pure strategy* for $i \in N \cup \{0\}$ is a map s_i that assigns an action $s_{i,I} \in A^I$ to each of i 's information sets $I_i \in \mathcal{I}_i$. We denote the set of i 's pure strategies by S_i , and use S to denote $\times_{i \in N} S_i$. A *mixed strategy* for player $i \in N$ is a distribution over her pure strategies, $\sigma_i \in \Delta(S_i)$. We use Σ_i to denote the set of i 's mixed strategies. For all $s_0 \in S_0$, we let $\pi(s_0)$ denote $\prod_{I \in \mathcal{I}^{-1}(\{0\})} \pi(s_{0,I})$ and let Σ_0 denote the set with the only mixed strategy by nature, $\{\pi\}$. We let $\Sigma := \times_{i \in \{0\} \cup N} \Sigma_i$ to denote the set of mixed strategy profiles.

Each mixed strategy profile $\sigma \in \Sigma$ induces an *outcome* $\mathbb{P}^\sigma \in \Delta(Z)$, that is, a probability distribution over terminal histories. We often abuse notation by using $u(\sigma)$ to denote $u(\mathbb{P}^\sigma) \equiv \sum_{z \in Z} \mathbb{P}^\sigma(z) u(z)$.

A *behavior strategy* for player $i \in N \cup \{0\}$ is a map β_i that assigns a distribution over the actions in A^I to each of i 's information sets $I_i \in \mathcal{I}_i$. We use \mathcal{B}_i to denote the set of i 's behavior strategies and $\mathcal{B} := \times_{i \in \{0\} \cup N} \mathcal{B}_i$ to denote the set of behavior strategy profiles. A behavior strategy profile β is *fully mixed* if $\beta_i(a_i) > 0$ for all $i \in N$ and $a_i \in A_i$, where $A_i := \cup_{I \in \mathcal{I}_i} A^I$ is the set of actions played by i . Note that pure strategies can be naturally seen as both mixed and behavior strategies. The outcome generated by a behavior strategy profile β is denoted \mathbb{P}^β .

As we shall see, some equilibrium concepts are defined using mixed strategies and others using behavior strategies. Because these are objects living in different spaces, it is convenient to use their outcomes to compare them. It will be useful to use the following classical result.

Lemma 1.1 (Kuhn, 1953). *The sets of outcomes of mixed and behavior strategies coincide.*

1.1 Sequential equilibria

As we will provide a mixed-strategy analog of Kreps and Wilson's (1982) concept of sequential equilibria, we now revisit it.

Belief systems and consistency: A *belief system* is a map μ assigning a probability $\mu(h) \in [0, 1]$ to each non-terminal history $h \in H \setminus Z$, in such a way that $\sum_{h \in I} \mu(h) = 1$ for all $I \in \mathcal{I}$. An *assessment* is then a pair (β, μ) . An assessment (β, μ) is *consistent* if there is a fully-mixed sequence $(\beta_n) \rightarrow \beta$ such that

$$\mu(h) = \lim_{n \rightarrow \infty} \frac{\mathbb{P}^{\beta_n}(h)}{\sum_{h' \in I} \mathbb{P}^{\beta_n}(h')}$$

for all $I \in \mathcal{I}$ and $h \in I$. In this case, (β_n) is said to *support* (β, μ) .

Conditional payoffs and sequential rationality: For a given assessment (β, μ) , define player i 's continuation payoff at $I \in \mathcal{I}_i$ as

$$u_i(\beta, \mu|I) = \sum_{h \in I} \sum_{z \in Z^h} \mu^I(h) \mathbb{P}^\beta(z|h) u_i(z),$$

where Z^h is the set of terminal histories that follow h , and where, for each $z \in Z^h$, we have $\mathbb{P}^\beta(z|h) := \prod_{j=J}^{|z|} \beta(z_j)$, with J denoting the index such that $(z_j)_{j=1}^J = h$. For a given $i \in N$ and $I \in \mathcal{I}_i$, we say that (β, μ) is *sequentially rational at I* if

$$\beta_i \in \operatorname{argmax}_{\hat{\beta}_i \in B_i} u_i(\hat{\beta}_i, \beta_{-i}, \mu|I).$$

We say that (β, μ) is *sequentially rational* if it is sequentially rational at all information sets.

Sequential equilibria: We say that (β, μ) is a *sequential equilibrium assessment* if it is consistent and sequentially rational. We say that β is a *sequential equilibrium* if there is some belief system μ such that (β, μ) is a sequential equilibrium assessment.

2 Sequential equilibria in mixed strategies

Nash equilibria under a perturbation

On our way to defining sequential equilibria in mixed strategies, we define trembles and Nash equilibria under a perturbed version of a game following Selten (1975).

Definition 2.1. A *tremble* ξ assigns to each $i \in N$ and $s_i \in S_i$ a value $\xi_i(s_i) \in (0, 1]$ such that $\sum_{\hat{s}_i \in S_i} \xi_i(\hat{s}_i) \leq 1$. A *tremble sequence* is a sequence $(\xi_n) \rightarrow 0$.

We follow Selten (1975) in defining Nash equilibria of perturbed games. As we will consider both perturbations of mixed strategies and payoffs, we define Nash equilibria for any game perturbed according to some tremble ξ and payoff function \hat{u} .

Definition 2.2. $\sigma \in \Sigma$ is a *Nash equilibrium* of $G(\xi, \hat{u})$ if, for all $i \in N$ and $s_i \in S_i$, (i) $\sigma_i(s_i) \geq \xi_i(s_i)$, and (ii) if $\sigma_i(s_i) > \xi_i(s_i)$, then $\hat{u}_i(s_i, \sigma_{-i}) \geq \hat{u}_i(\hat{s}_i, \sigma_{-i})$ for all $\hat{s}_i \in S_i$.

Selten's interpretation is that in the game perturbed according to ξ , players tremble and make mistakes, sometimes choosing a dominated strategy. Analogously to standard Nash equilibria, a player chooses a strategy with a probability above the tremble only if the strategy is a best response. As we are perturbing the set of mixed strategies, a natural interpretation of a mixed tremble is as a small probability that each given player is a behavioral action type. In this case, a Nash equilibrium of $G(\xi, u)$ can be seen as a Nash equilibrium of a game with small uncertainty about each player's rationality.

Sequential equilibria in mixed strategies

An important result in Kreps and Wilson (1982) is that sequential equilibria constitute the set of strategy profiles that can be approximated by Nash equilibria of games with perturbed strategies and payoffs.⁶ We now state an analogous result for mixed strategies.

Definition 2.3. $\sigma \in \Sigma$ is a *sequential equilibrium in mixed strategies* if there is a sequence $(\sigma_n, \xi_n, u_n) \rightarrow (\sigma, 0, u)$ such that each σ_n is a Nash equilibrium of $G(\xi_n, u_n)$.

Like the analogous result of Kreps and Wilson (1982) (their Proposition 6), Definition 2.3 gives plausibility to sequential equilibria in mixed strategies. It defines σ to be a sequential equilibrium in mixed strategies if, for all $\varepsilon > 0$, there is a ε -perturbation of the game (in terms of payoffs and behavioral types) with a Nash equilibrium ε -close to σ . Hence, assuming that the econometrician cannot fully determine payoffs or identify

⁶See also discussions in Blume and Zame (1994) (Proposition A) and Dilmé (2025).




Figure 1

small probabilities of irrationality, Definition 2.3 requires that the prediction is close to the prediction of an environment close to the one observed.

The converse implication of Definition 2.3 is that if σ is *not* a sequential equilibrium in mixed strategies, then it is fragile: no similar game (with slightly perturbed strategies and payoffs) has a Nash equilibrium close to σ . Hence, focusing on sequential equilibria in mixed strategies rules out behavior that ceases to be a Nash equilibrium for any similar perturbed game. The following result formalizes this observation.

Proposition 2.1. *σ is not a sequential equilibrium in mixed strategies if and only if, for all $\varepsilon > 0$, there is some $\delta > 0$ such that for all ξ and \hat{u} with $\|\xi\| + \|\hat{u} - u\| < \delta$, no Nash equilibrium of $G(\xi, \hat{u})$ is at a distance less than ε from σ .⁷*

Example 2.1. Throughout, we will use the simple game in Figure 1, which corresponds to the game in Figure 1 in Reny (1992). This game has two outcomes of Nash equilibria: one is D_1 and the other is (A_1, A_2, A'_1) . We will now explicitly show D_1 is an outcome of sequential equilibria in mixed strategies (it is easy to verify that (A_1, A_2, A'_1) is too).

Let (a_1, a'_1) indicate player 1's pure strategy specifying playing $a_1 \in \{D_1, A_1\}$ and $a'_1 \in \{D'_1, A'_1\}$ at the respective information sets. Let $(\xi_n) \rightarrow 0$ be such that $\xi_{1,n}(A_1, D'_1) = n^{-1}$ and $\xi_{1,n}(s_1) = \xi_{2,n}(s_2) = n^{-2}$ for all $s_1 \neq (A_1, D'_1)$ and s_2 . Then, it is easy to see that, for each n , the strategy profile $(\sigma_{1,n}, \sigma_{2,n})$ given by

$$\begin{aligned} & \left((1 - 2n^{-2} - n^{-1})(D_1, D'_1) + n^{-2}(D_1, A'_1) + n^{-1}(A_1, D'_1) + n^{-2}(A_1, A'_1), \right. \\ & \left. (1 - n^{-2})D_2 + n^{-2}A_2 \right) \end{aligned}$$

⁷Here and in the rest of the paper, $\|\cdot\|$ denotes the sup norm of the corresponding space.

is a Nash equilibrium of $G(\xi_n, u)$ for $n \geq 2$.⁸ Also, it is clear that the outcome of $(\sigma_{1,n}, \sigma_{2,n})$ converges to D_1 .

While D_1 is not the outcome of a subgame perfect Nash equilibrium, it is plausible in close-by games with small uncertainty about player 1's rationality. Indeed, in the game perturbed according to ξ_n , the probability of 1's irrational type playing (A_1, D'_1) is n times as high as the probability of any of its other irrational types. As a result, in an equilibrium where the rational type of player 1 is known to choose D_1 in the first information set, player 2 gets to play only if player 1 is a behavioral type, in which case she is more likely to choose D'_1 than A_1 , making D_2 an optimal strategy for player 2. This makes it optimal for the rational player 1 to choose D_1 .

3 Relationship to weakly sequential equilibria

3.1 Weakly sequential equilibria

Reny (1992) introduced weakly sequential equilibria, a weakening of sequential equilibria obtained by relaxing the requirement of sequential rationality at all information sets. In this section, we revisit his approach and show that weakly sequential equilibria are equivalent to sequential equilibria in mixed strategies.

We begin with the concept of relevant information sets for a given strategy, a concept introduced in Reny (1992). We say that $I \in \mathcal{I}$ is s_i -relevant if $\mathbb{P}^{s_i, s_{-i}}(I) > 0$ for some $s_{-i} \in S_{-i}$. Equivalently, I is s_i -relevant if it is reached with positive probability under all fully-mixed strategies by players other than i . We use $\mathcal{I}_i^{s_i}$ to denote the set of s_i -relevant information sets in \mathcal{I}_i , respectively. Conversely, we use S_i^I to denote the set of i 's pure strategies s_i for which I is s_i -relevant, that is, the i 's pure strategies that do not preclude I . For a given strategy profile σ , we say that $I \in \mathcal{I}_i$ is σ -relevant if it is relevant for some pure strategy s_i played with positive probability under σ .

⁸In most games, that a strategy profile is a sequential equilibrium in mixed strategies can be shown by considering sequences of games perturbed only according to trembles (and not payoffs). We will see in Section 4 that, in fact, this is always possible in games with generic payoffs. Figure 2 provides an example with non-generic payoffs where this is not possible (see Footnote 12).

Reny (1992) defines an assessment (β, μ) to be *weakly sequentially rational* if it is sequentially rational at all information sets that are β -relevant. That is, the main difference between sequential rationality and weak sequential rationality is that, in the second case, the condition that the continuation strategy is optimal is only imposed at β -relevant information sets. The motivation for this definition is that information sets $I \in \mathcal{I}_i$ that are *not* β_i -relevant can only be reached if player i has deviated, so the other players need not assume i is rational anymore. Instead, information sets $I \in \mathcal{I}_i$ that are β_i -relevant can be reached if players other than i have deviated (but not i), so players should continue to believe that i is rational (and plays according to the prescribed strategy).

Definition 3.1 (Reny, 1992). (β, μ) is a *weakly sequential equilibrium* if it is consistent and weakly sequentially rational.

Proposition 3.1. *The sets of outcomes of sequential equilibria in mixed strategies and weakly sequential equilibria coincide.*

Proposition 3.1 establishes that sequential equilibria in mixed strategies and weakly sequential equilibria predict the same behavior. It provides additional foundation for the use of sequential equilibria in mixed strategies, as it establishes that it can be obtained by considering equilibria where players are known to be rational, but where deviations are perceived as a sign of irrationality.

An important step in the proof of Proposition 3.1, which is key to understanding the result, is the following lemma that connects Nash equilibria of perturbed games with the condition of weak sequential rationality. In the statement, “ s_i is sequentially rational given σ_{-i} at I ” means that $u_i(s_i, \sigma_{-i}|I) \geq u_i(\hat{s}_i, \sigma_{-i}|I)$ for all $\hat{s} \in S_i^I$.

Lemma 3.1. *σ is a Nash equilibrium of $G(\xi, u)$ if and only if for all $i \in N$, $s_i \in S_i$ with $\sigma_i(s_i) > \xi_i(s_i)$, and $I \in \mathcal{I}_i^{s_i}$, we have that s_i is sequentially rational at I given σ_{-i} .*

Lemma 3.1 establishes that s_i is a best response against σ_{-i} if and only if its continuation is a best response against σ_{-i} at all s_i -relevant information sets. The “if” direction is clear because the payoff from playing s_i is determined by the continuation payoffs at all information sets where i first plays, which are relevant under all her strategies. The “only

if” condition follows from a standard argument: if s_i is not a best response against σ_{-i} , there must be an s_i -relevant information set where player i can benefit from choosing a different continuation play.

From Lemma 3.1, it is straightforward to see that outcomes of sequential equilibria in mixed strategies are outcomes of weakly sequential equilibria. Intuitively, if σ is sequential equilibrium in mixed strategies supported by some sequence (σ_n, ξ_n, u_n) , then there is a corresponding sequence of $(\beta_n)_n$ (where each β_n has the same outcome as σ_n) supporting a consistent and weakly sequentially rational assessment (β, μ) .

That outcomes of weakly sequential equilibria are outcomes of sequential equilibria in mixed strategies is more involved. The reason is that weakly sequential equilibria are required to be supported by a sequence (β_n) to ensure consistency, but weak sequential rationality is only imposed in the limit. Sequential equilibria in mixed strategies are instead supported by a sequence (σ_n, ξ_n, u_n) where now each σ_n is a Nash equilibrium of $G(\xi_n, u_n)$, hence (by Lemma 3.1) each s_i with $\sigma_{i,n}(s_i) > 0$ is *exactly* sequentially rational given $\sigma_{-i,n}$ at each $I \in \mathcal{I}_i^{s_i}$. We then find an algorithm to obtain a convenient sequence of trembles and payoffs which, for each n , iteratively adjusts the payoff at histories passing through the different information sets of each player.

Proposition 3.1 aligns with the interpretation of mixed-strategy trembles: perturbing the strategy space makes deviations signal irrationality. Indeed, perturbations of the space of mixed strategies naturally lead to interpreting deviations as a sign of irrationality. For sequential rationality, instead, actions are perturbed, so deviations are perceived as one-time mistakes by players who continue to be rational.

Example 3.1. Example 2.1 shows that D_1 is the outcome of a sequential equilibrium in mixed strategies of the game in Figure 1. As Reny (1992) explains, D_1 is the outcome of a weakly sequential equilibrium too. To see this, consider the behavior assessment (β, μ) where $\beta_1(D_1) = \beta_1(D'_1) = \beta_2(D_2) = 1$ (the rest of the assessment is uniquely pinned down). Note that only the first information set by player 1 is β_1 -relevant, and that player 2's information set is β_2 -relevant. It is then clear that (β, μ) is weakly sequentially rational (and consistent, as all information sets are singletons).⁹

⁹As Reny (1992) poses it, “the reason for this is that once player 1 deviates from $((D_1, D'_1), D_2)$ by playing A_1 , player 2 is allowed to

3.2 Conditional assessments

We now provide a characterization of sequential equilibria in mixed strategies in terms of conditional assessments, which is an analog of the original definition of sequential equilibria by Kreps and Wilson (1982). Our approach is similar to that of Govindan and Wilson (2009) in defining weakly sequential equilibria.

A *conditional belief system* is a map γ that assigns a distribution $\gamma^I \in \Delta(S_i^I) \times \Delta(S_{-i}^I)$ to each information set $I \in \mathcal{I}_i$ with $\mathbb{P}^{\gamma^I}(I) = 1$. We interpret γ^I as an external observer's assessment about the players' and nature's choices of pure strategies given that the information set has been reached. Similarly, we interpret γ_i^I as i 's continuation strategy at I . Finally, we interpret γ_{-i}^I as the belief that player i holds at information set I about the other players' pure strategies. Note that γ^I may not be a strategy profile, as randomizations of players other than i may not be independent across players.

A *conditional assessment* is a pair (σ, γ) formed by a strategy profile and a conditional belief system. We say that (σ, γ) is *consistent* if there is a fully-mixed sequence $(\sigma_n) \rightarrow \sigma$ such that

$$\gamma^I(s) = \lim_{n \rightarrow \infty} \frac{\mathbb{P}^{\sigma_n}(I) \sigma_n(s)}{\mathbb{P}^{\sigma_n}(I)}$$

for all $I \in \mathcal{I}$ and $s \in S$. In this case, we say that (σ_n) *supports* γ (note that $\gamma^{\{\emptyset\}} = \sigma$). Paralleling Reny (1992), we say that (σ, γ) is *weakly sequentially rational* if for all $i \in N$ and $I \in \mathcal{I}_i$ that is σ -relevant, we have

$$\gamma_i^I \in \operatorname{argmax}_{\hat{\sigma}_i \in \Delta(S_i^I)} u_i(\hat{\sigma}_i, \gamma_{-i}^I | I).$$

An equivalent and often useful definition of weak sequential rationality is the following: (σ, γ) is sequentially rational if for all $i \in N$ and $s_i \in S_i$ with $\sigma_i(s_i) > 0$, we have $u_i(s_i, \gamma_{-i}^I | I) \geq u_i(\hat{s}_i, \gamma_{-i}^I | I)$ for all $I \in \mathcal{I}_i^{s_i}$, and $\hat{s}_i \in S_i^I$. In words, a strategy s_i receives a positive probability only if it is credible, in the sense that it delivers the highest continuation payoff to player i at each s_i -information set I given γ_{-i}^I .

believe, for instance, that player 1 is 'crazy' and if given the chance will choose D_2 " (p. 634). We believe that the explicit definition of sequential equilibria in mixed strategies in terms of small probabilities of behavioral types in Example 2.1 makes this intuition clearer.

Corollary 3.1. *A strategy profile is a sequential equilibrium in mixed strategies if and only if it is part of a conditional assessment that is consistent and weakly sequentially rational.*

Corollary 3.1 establishes that we can use conditional assessments to work with sequential equilibria in mixed strategies. Doing so is often easier than directly using Definition 2.3, as characterizing Nash equilibria along sequences of games with perturbed strategies and payoffs is often difficult.

Remark 3.1. Unlike us, Govindan and Wilson (2009) define “beliefs” as a map from each information set $I \in \mathcal{I}_i$ to the distribution over strategies in S_{-i} , that is, only use γ_{-i}^I . As we shall see, even though γ_i^I is not necessary to compute player i ’s continuation payoff of a strategy $s_i \in S_i^I$ at I , γ_i^I will be useful to verify γ ’s consistency through Bayes’ rule (see Example 3.2). Note that γ_i^I can be interpreted as the “continuation strategy” of player i at I .

Example 3.2. Examples 2.1 and 3.1 explicitly show that D_1 is the outcome of both a sequential equilibrium in mixed strategies and a weakly sequential equilibrium of the game in Figure 1. We now provide a consistent and weakly sequentially rational assessment supporting it. Let $\sigma = ((D_1, D'_1), D_2)$ and consider the following conditional belief system¹⁰

$$\gamma^{\{\emptyset\}} = \delta_{(D_1, D'_1)} \delta_{D_2}, \quad \gamma^{\{A_1\}} = \delta_{(A_1, D'_1)} \delta_{D_2}, \quad \text{and} \quad \gamma^{\{(A_1, A_2)\}} = \delta_{(A_1, D'_1)} \delta_{A_2},$$

where δ_{s_i} is the distribution degenerate at s_i (i.e., Dirac’s delta). It is straightforward to see that (σ, γ) is supported by the sequence (σ_n) used in Example 2.1. It is also straightforward to see that it is weakly sequentially rational.

We make two observations. The first is that it is useful to have $\gamma^{\{A_1\}}$ have information about player 1’s continuation strategy. The reason is that it helps to assess the consistency of γ : because $\{A_1\}$ does not belong to player 1, it requires that $\gamma_1^{\{A_1\}} = \gamma_1^{\{(A_1, A_2)\}}$ (in our case, equal to $\delta_{(A_1, D'_1)}$). This conforms to the “no signaling what you do not know” condition often required in perfect Bayesian equilibria.

The second observation is that conditional beliefs often fail the “never dissuaded once convinced” condition, which requires the belief about an event that is believed to have

¹⁰Henceforth, for a statement P , $\mathbb{I}_P = 1$ if P is true and $\mathbb{I}_P = 0$ if P is false.

probability zero at some history to remain to be assigned probability zero in the future. For example, $\gamma_2^{\{\emptyset\}}$ assigns probability zero to A_2 , but then $\gamma_2^{\{(A_1, A_2)\}}$ assigns probability one to A_2 . This is natural because the information set $\{(A_1, A_2)\}$ can only be reached if player 2 has deviated and played A_2 .

Relationship to sequential equilibria

Ranking tremble-based equilibrium concepts defined in terms of behavior and mixed strategies is often difficult. For example, there are perfect equilibria that are not perfect equilibria in mixed strategies, and vice versa (in Section 4.1 we show that, nonetheless, they can be ranked for generic payoffs). The reason is that the sets of trembles against which robustness is required is different. The following result establishes that such a ranking is possible when sequential equilibria in behavior and mixed strategies are compared.

Corollary 3.2. *Outcomes of sequential equilibria are outcomes of sequential equilibria in mixed strategies. The reverse holds if each player plays at most once along any history.*

The first part of Corollary 3.2 follows from Proposition 3.1 and the fact that weakly sequential equilibria are a weakening of sequential equilibria. The second statement follows because, if each player plays at most once along any history, all information sets are relevant under any strategy (no own earlier move can preclude reaching a later own information set). Hence, in this case, weak sequential optimality coincides with sequential optimality.

4 Perfect equilibria in mixed strategies

Selten (1975) introduced the concept of (trembling hand) perfect equilibria, both for a game in extensive form and a game in normal form. For games in extensive form, Kreps and Wilson (1982) showed that perfect equilibria and sequential equilibria are closely related: sequential equilibria are perfect equilibria, and the two concepts coincide in games with generic payoffs. We will now show that similar results hold when we use mixed instead of behavior strategies.

We begin extending the concept of perfect equilibria to mixed strategies.

Definition 4.1. σ is a *perfect equilibrium in mixed strategies* if there is a tremble sequence $(\xi_n) \rightarrow 0$ and a sequence $(\sigma_n) \rightarrow \sigma$ such that each σ_n is a Nash equilibrium of $G(\xi_n, u)$.

It is easy to see that perfect equilibria in mixed strategies and perfect equilibria of the normal form of the game coincide. That is, the set of perfect equilibria in mixed strategies of a game coincides with the set of perfect equilibria of the simultaneous-move game where the set of actions of each player i is S_i .

Lemma 4.1. *Perfect equilibria in mixed strategies are sequential equilibria in mixed strategies.*

This result trivially follows from using the constant sequence (u, u, \dots) in the definition of sequential equilibria in mixed strategies. That is, sequential equilibria in mixed strategies are strategy profiles approximable through both sequences of trembles and payoff perturbations, while perfect equilibria in mixed strategies are only approximable through sequences of trembles. Note that the result is also implied by the result in Reny (1992) that outcomes of normal-form perfect equilibria are outcomes of weakly sequential equilibria (see his Proposition 1).

Lemma 4.1 allows our characterizations of sequential equilibria in mixed strategies in Proposition 3.1 and Corollary 3.1 to identify candidates to perfect equilibria in mixed strategies or, alternatively, to rule out other candidates. Indeed, in the same way that Kreps and Wilson (1982) point out, “It is vastly easier to verify that a given equilibrium is sequential than that it is perfect” (p. 264), our characterizations of sequential equilibria in mixed strategies using assessments are easier to use than using sequences of trembles and corresponding Nash equilibria.

4.1 Generic equivalence with perfect equilibria in mixed strategies

Lemma 4.1 establishes that sequential equilibria in mixed strategies is a weakening of the concept of perfect equilibria in mixed strategies. An important result in Kreps and Wilson (1982) is that this weakening is “minimal” for behavioral strategies: the sets of perfect and

sequential equilibria coincide in games with generic payoffs. We now extend this result to mixed strategies.

Proposition 4.1. *Generically in payoffs, the sets of perfect equilibria in mixed strategies and sequential equilibria in mixed strategies coincide.*

The proof follows Blume and Zame (1994)'s analogous result for behavior strategies. In their proof, a key step is to establish that sequential equilibria can be approximated by sequences of Nash equilibria of close-by games perturbed according to extensive-form trembles and payoff perturbations. For sequential equilibria in mixed strategies, such a property is guaranteed by definition (Definition 2.3).¹¹ Remarkably, the rest of the proofs turn out to be identical because the graphs of the equilibrium correspondences can be described using the same expressions, but using mixed instead of behavioral strategies and trembles.

4.2 Relationship to perfect equilibria

Selten (1975) defined two versions of perfect equilibria. The first is normal-form perfect equilibria, as described above, which is equivalent to perfect equilibria in mixed strategies. The second is extensive-form perfect equilibria (often referred to as just perfect equilibria), which are limits of sequences of Nash equilibria along sequences of perturbed games where players tremble in their choices of the actions instead of strategies. It is not difficult to see that the two concepts coincide when one considers the agent-normal form game.

As we explain in Section 3.2, there are (extensive-form) perfect equilibria that are not perfect equilibria in mixed strategies, and vice versa. For example, the arguments used in Example 2.1 can be used to show that D_1 is the outcome of a perfect equilibrium in mixed strategies. Nevertheless, it is clearly not the outcome of a perfect equilibrium, as it is not subgame perfect. Conversely, consider the game in Figure 2, which corresponds to Figure 6.4.2 in van Damme (1991). From van Damme's discussion it follows that (B_1, T_2) is the outcome of a perfect equilibrium, but not the outcome of a perfect equilibrium in mixed

¹¹As explained after Proposition 3.1, the difficulty in our approach is to then establish the equivalence between sequential equilibria in mixed strategies and their characterizations as assessments (such as Proposition 3.1 and Corollary 3.1).




Figure 2

strategies. Indeed, note that this is the unique limit of Nash outcomes along a sequence of perturbed games where player 2 trembles less than player 1. Still, for any fully mixed strategy $\sigma_2 \in \Sigma_2$,

$$u_1((T_1, T'_1), \sigma_2) = 1 > \sigma_2(T_2) = u_1((B_1, T'_1), \sigma_2) = \max_{s_1 \in S_1 \setminus \{(T_1, T'_1)\}} u_1(s_1, \sigma_2) .$$

Intuitively, in an extensive-form perfect equilibrium, player 1 takes into account that she may tremble in a subsequent information set of hers, but she does not do that in a normal-form perfect equilibrium. Such a conclusion is nevertheless fragile to small changes in the payoffs.¹²

As the following result states, an implication of Proposition 4.1 and Corollary 3.2 is that, generically in payoffs, the set of outcomes of perfect equilibria is a subset of the set of outcomes of perfect equilibria in mixed strategies.¹³

Corollary 4.1. *Generically in payoffs, outcomes of perfect equilibria are outcomes of perfect equilibria in mixed strategies.*

5 Purification

Harsanyi (1973) asked which strategy profiles could be purified, that is, obtained as limits of Nash equilibria of nearby games with a small amount of payoff uncertainty. His answer was that, in games in normal form, all and only Nash equilibria could be obtained that

¹²All outcomes assigning probability one to $\{(T_1, T'_1), (B_1, T_2)\}$ are the outcomes of sequential equilibria and the outcomes of sequential equilibria in mixed strategies. Indeed, any such outcomes can be approximated through Nash outcomes of $G(\xi_n, u_n)$ where each ξ_n is a uniform tremble $\xi_{i,n}(s_i) = n^{-1}$ for all i and s_i and the payoff sequence $u_n(z) = u(z)$ for all $z \neq (B_1, T_2)$ and $u_n(B_1, T_2) = (1 - n^{-1})^{-1}$.

¹³Note that van Damme (1984) introduces the concept of quasi-perfect equilibria, which is defined like extensive-form perfect equilibria except that players do not take into account their own future tremble probabilities.

way, providing a justification for mixed strategy equilibria. We now show that if one asks the same question for games in extensive form, all and only sequential equilibria in mixed strategies can be purified.

To state and prove our result, we first provide a bit of notation. For each $i \in N$, let $U_i \subset \mathbb{R}^Z$ be an open and bounded set of payoffs that includes u_i and, for each s_i , a payoff function that makes s_i a strictly dominant strategy.

We let $\Delta^*(U_i)$ be the space of probability measures on U_i without atoms and with full support.¹⁴ For each $s_i \in S_i$, let $S_i^{s_i}$ be the set of player i 's strategies \hat{s}_i satisfying that $\hat{s}_{i,I} = s_{i,I}$ for all $I \in \mathcal{I}_i^{s_i}$. Note that all strategies in $S_i^{s_i}$ are strategically equivalent:

$$\mathbb{P}^{s_i, \sigma_{-i}}(z) = \mathbb{P}^{\hat{s}_i, \sigma_{-i}}(z)$$

for all $z \in Z$, $\hat{s}_i \in S_i^{s_i}$, and $\sigma_{-i} \in \Sigma_{-i}$. We let $\times_{i \in N} \Delta^*(U_i)$ be the space of product probability measures. Note that for all $\nu_i \in \Delta^*(U_i)$, $s_i, \hat{s}_i \in S_i$, and $\sigma_{-i} \in \Sigma_{-i}$,

$$\nu_i(\{\hat{u}_i \in U_i \mid \hat{u}_i(s_i, \sigma_{-i}) = \hat{u}_i(\hat{s}_i, \sigma_{-i})\}) = 0$$

if and only if $\hat{s}_i \notin S_i^{s_i}$. That is, the probability that player i is indifferent between two of its pure non-strategically equivalent strategies is zero.

We now introduce the notion of Nash equilibria perturbed according to some payoff distribution. It requires that, for each player i and $s_i \in S_i$, the value $\sigma_i(S_i^{s_i})$ corresponds to the mass of payoff types for whom s_i (or any strategy strategically equivalent to s_i) is a best response against σ_{-i} .

Definition 5.1. Fix some $\nu \in \times_{i \in N} \Delta^*(U_i)$. Then, $\sigma \in \Sigma$ is a *Nash equilibrium* of $G(\nu)$ if for all $i \in N$ and $s_i \in S_i$,

$$\sigma_i(S_i^{s_i}) = \nu_i(\{\hat{u}_i \in U_i \mid \hat{u}_i(s_i, \sigma_{-i}) \geq \hat{u}_i(\hat{s}_i, \sigma_{-i}) \ \forall \hat{s}_i \in S_i\}).$$

We use $(\nu_n) \rightarrow \delta_u$ to denote that (ν_n) converges in probability to the probability measure degenerate at u (i.e., the Dirac measure δ_u). That is, $(\nu_n) \rightarrow \delta_u$ if, for all $\varepsilon > 0$ and

¹⁴Following Harsanyi (1973), we assume that the distribution of payoffs in the perturbed game has no atoms. This assumption simplifies the analysis, but it is not necessary: similar results can be obtained as long as the assumption of full support is maintained.

$i \in N$, we have

$$\nu_{i,n}(\{\hat{u}_i \in U_i \mid \|\hat{u}_i - u_i\| > \varepsilon\}) \rightarrow 0 \text{ as } n \rightarrow \infty$$

(recall that we use the sup-norm). We now define the concept of purifiable strategy profile, following Harsanyi (1973)'s definition.

Definition 5.2. We say that σ is *purifiable* if there exists a sequence $(\nu_n) \rightarrow \delta_u$ and corresponding sequence $(\sigma_n) \rightarrow \sigma$, where each σ_n is a Nash equilibrium of $G(\nu_n)$.

Proposition 5.1. σ is purifiable if and only if it is a sequential equilibrium in mixed strategies.

Remarkably, Proposition 5.1 characterizes sequential equilibria in mixed strategies in terms of only payoff perturbations. This characterization is conceptually different from the robustness property required in the definition of sequential equilibrium in mixed strategies. Indeed, we define sequential equilibria in mixed strategies by requiring approximability through close-by games satisfying that (i) there is a small probability that each player is irrational, and (ii) players know with certainty the payoffs of the other players' rational types (for some payoffs close to the non-perturbed ones). Instead, Proposition 5.1 establishes approximability through close-by games satisfying that (i) players are known to be rational, and (ii) players do not know the payoffs of the other players. It then bridges two important sources of strategic uncertainty: uncertainty about the rationality of other players and uncertainty about their payoffs.

Proposition 5.1 can also be viewed as providing a simple way of obtaining behavior robust to payoff perturbations. Even in simple examples such as the game in Figure 1, explicitly proving that a given equilibrium is purifiable is difficult. Hence, using sequential equilibria in mixed strategies simplifies the identification of sensible behavior in settings where players are slightly uncertain about other players' payoffs.

It is important to note that we are considering payoff perturbations of the payoff function, that is, the players' payoffs at the terminal nodes. This is different from perturbing their payoffs in the normal form of the game, where Harsanyi (1973)'s result establishes that all and only Nash equilibria are purifiable. Such a difference highlights the difference between working with mixed strategies in the game in extensive form and the corresponding normal form.

Remark 5.1. In recent work, Bhaskar and Stinchcombe (2024) study purification and strong purification (requiring purification along all payoff perturbations) in extensive games with perfect information and signaling games. Consistent with our results, they obtain that in simple trees with generic payoffs (each player moves at most once on any path), only the unique backward-induction equilibrium is strongly purifiable, and is the only one that can be purified. When a player moves more than once, the unique backward-induction equilibrium is not strongly purifiable, and purifiable equilibria may not be subgame perfect. They also study symmetric purifiability in signaling games (with i.i.d. payoff shocks across a player’s terminal nodes), where they show that the quiche outcome is not symmetrically purifiable. Our work complements theirs in showing that all and only sequential equilibria in mixed strategies are purifiable.

6 Conclusions

Our paper contributes to the study of games by characterizing behavior that is robust to a small uncertainty about other players’ rationality. We parallel the approaches in seminal work of Selten (1975) and Kreps and Wilson (1982), focusing on mixed instead of behavior strategies.

Our results show the equivalence between behavior that is robust to a small uncertainty about other players’ rationality and behavior that is robust to low-probability payoff perturbations. Such equivalence adds plausibility to sequential equilibria in mixed strategies, especially in environments where such small perturbations are likely to affect equilibrium outcomes. The equivalence with weakly sequential equilibria simplifies obtaining and manipulating such robust behavior in practice.

While Reny (1992) and Govindan and Wilson (2009) suggest refinements to the concept of weakly sequential equilibria (termed explicable equilibria and outcomes that satisfy forward induction), we abstain from doing so in this paper. In a companion paper, we propose the concept of sequentially stable in mixed strategies outcomes, which refines the set of outcomes of sequential equilibria in mixed strategies (and satisfies forward induction).

There is some additional work related to our analysis. In Section C.1, we relate our

analysis to the concept of quasi-perfect equilibria (van Damme, 1984), and we establish the equivalence between perfect equilibria in mixed strategies and weakly quasi-perfect equilibria. In Section C.2, we relate our analysis to the analysis of sequential equilibria in normal-form games in Mailath et al. (1993).

A ε -Sequential equilibria in mixed strategies

In this section, we introduce the concept of ε -sequential equilibria in mixed strategies. This concept is analogous to the concept of sequential ε -equilibria (Dilmé, 2024). We will then provide a characterization of sequential equilibria in mixed strategies in terms of sequences of ε -sequential equilibria in mixed strategies as $\varepsilon \rightarrow 0$. This characterization will be very useful in the proofs of the previous results, and so we place its proof before them.

Definition A.1. σ is an ε -sequential equilibrium in mixed strategies under ξ if, for all $i \in N$ and $s_i \in S_i$, (i) $\sigma_i(s_i) \geq \xi_i(s_i)$, and (ii) if $\sigma_i(s_i) > \xi_i(s_i)$, then $u_i(s_i, \sigma_{-i}|I) \geq u_i(\hat{s}_i, \sigma_{-i}|I) - \varepsilon_n$ for all $I \in \mathcal{I}_i^{s_i}$ and $\hat{s}_i \in S_i^I$.

The following result is a mixed-strategies analog to Proposition 3.1 in Dilmé (2024), which characterizes sequential equilibria in terms of sequences of sequential ε -equilibria.

Proposition A.1. σ is a sequential equilibrium in mixed strategies if and only if there is some $(\sigma_n, \varepsilon_n, \xi_n) \rightarrow (\sigma, 0, 0)$ such that each σ_n is a ε_n -sequential equilibrium in mixed strategies under ξ_n .

Note the difference between Proposition A.1 and the definition of sequential equilibria in mixed strategies (Definition 2.3). The first involves limits of ε_n -sequential equilibria in mixed strategies, while the second involves limits of 0-sequential equilibria in mixed strategies along games with perturbed payoffs (note that Lemma A.1 below shows that 0-sequential equilibria in mixed strategies under ξ are Nash equilibria of $G(\xi, u)$). Then, the “if” part of Proposition A.1 is straightforward: by setting $\varepsilon_n := \|u - u_n\|/2$, a 0-sequential equilibrium in mixed strategies under ξ_n of $G(u_n)$ is an ε_n -sequential equilibrium in mixed strategies under ξ_n (of $G(u)$).

The proof of the “only if” part of Proposition A.1 is more involved. For each (σ_n) such that each σ_n is an ε_n -sequential equilibrium in mixed strategies under ξ_n , we want to find some $(u_n) \rightarrow u$ such that each σ_n is a Nash equilibrium of $G(\xi_n, u_n)$. Nevertheless, ensuring the “exact” optimality of all s_i with $\sigma_n(s_i) > 0$ is difficult, as changes in the payoff of a terminal history affect the payoff from different strategies in different ways. In the proof, we construct an algorithm that constructs each desired payoff sequence by iteratively adjusting the payoff at histories passing through the different information sets of each player (such an algorithm is also useful to prove the equivalence between weakly sequential equilibria and sequential equilibria in mixed strategies).

Proposition A.1 is useful because, in many cases, proving the existence of $(\sigma_n, \varepsilon_n, \xi_n) \rightarrow (\sigma, 0, 0)$ with the above properties is easier than proving the existence of $(\sigma_n, u_n, \xi_n) \rightarrow (\sigma, u, 0)$ with the properties in Definition 2.3. The reason is that the constraints in the first case are inequalities, while the constraints in the second case are equalities.

Proof of Proposition A.1

Proof. We begin the proof with a result establishing the equivalence between Nash equilibria with 0-sequential equilibria in mixed strategies.

Lemma A.1. *Fix a tremble ξ . Then, Nash equilibria of $G(\xi, u)$ and 0-sequential equilibria in mixed strategies under ξ coincide.*

Proof. Assume first that σ is a 0-sequential equilibrium in mixed strategies under ξ . This implies that $\sigma_i(s_i) \geq \xi_i(s_i)$ for all $i \in N$ and $s_i \in S_i$. Fix some $s_i \in S_i$ with $\sigma_i(s_i) > \xi_i(s_i)$ and assume for a contradiction that there is some $\hat{s}_i \in S_i$ such that $u_i(\hat{s}_i, \sigma_{-i}) > u_i(s_i, \sigma_{-i})$. By the same argument as in the proof of Lemma 3.1, there must be some $I \in \mathcal{I}_i^{s_i} \cap \mathcal{I}_i^{\hat{s}_i}$ such that $u_i(s_i, \sigma_{-i}|I) < u_i(\hat{s}_i, \sigma_{-i}|I)$. This contradicts that σ is a 0-sequential equilibrium in mixed strategies under ξ .

Assume now that σ is a Nash equilibrium of $G(\xi, u)$. Assume for a contradiction that there is some s_i with $\sigma_i(s_i) > \xi_i(s_i)$, some $I \in \mathcal{I}_i^{s_i}$, and some $\hat{s}_i \in S_i^I$ such that $u_i(s_i, \sigma_{-i}|I) < u_i(\hat{s}_i, \sigma_{-i}|I)$. Let \hat{s}'_i be the strategy coinciding with s_i in all information sets except I and all that succeed it, where it coincides with \hat{s}_i . Note that $\hat{s}'_i \in S_i^I$ and that $u_i(\hat{s}'_i, \sigma_{-i}|I) =$

$u_i(\hat{s}_i, \sigma_{-i}|I)$. Now we have

$$\begin{aligned} u_i(s_i, \sigma_{-i}) - u_i(\hat{s}_i', \sigma_{-i}) &= \sum_{z \in Z^I} \mathbb{P}^{s_i, \sigma_{-i}}(z) u_i(z) - \sum_{z \in Z^I} \mathbb{P}^{\hat{s}_i', \sigma_{-i}}(z) u_i(z) \\ &= \mathbb{P}^{s_i, \sigma_{-i}}(I) (u_i(s_i, \sigma_{-i}|I) - u_i(\hat{s}_i', \sigma_{-i}|I)) < 0, \end{aligned}$$

where Z^I is the set terminal history that pass through I . This contradicts that $\sigma_i(s_i) > 0$ and σ is a Nash equilibrium of $G(\xi, u)$. \square

(End of the proof of Lemma A.1. The proof of Proposition A.1 continues.)

Proof of the “if” part of Proposition A.1. Assume that σ is such that there is a sequence $(\sigma_n, \xi_n, u_n) \rightarrow (\sigma, 0, u)$ such that each σ_n is a 0-sequential equilibrium in mixed strategies under ξ_n of the game $G(u_n)$. For each $n \in \mathbb{N}$ define $\varepsilon_n := \|u_n - u\|/2$, and note that $\varepsilon_n \rightarrow 0$. Fix some $n \in \mathbb{N}$, $i \in N$, $s_i \in S_i$ such that $\sigma_{i,n}(s_i) > \xi_{i,n}(s_i)$, $I \in \mathcal{I}_i^{s_i}$, and $\hat{s}_i \in S_i^I$. Note that

$$0 \leq u_{i,n}(s_i, \sigma_{-i,n}|I) - u_{i,n}(\hat{s}_i, \sigma_{-i,n}|I) \leq u_i(s_i, \sigma_{-i,n}|I) - u_i(\hat{s}_i, \sigma_{-i,n}|I) + \varepsilon_n.$$

It is then clear that σ_n is a ε_n -sequential equilibrium in mixed strategies under ξ_n (in G), and so σ is a sequential equilibrium in mixed strategies.

Proof of the “only if” part. Using Lemma A.1, this part of the proof follows from the following lemma.

Lemma A.2. *Let $(\sigma_n, \varepsilon_n, \xi_n) \rightarrow (\sigma, 0, 0)$ be such that each σ_n is a ε_n -sequential equilibrium in mixed strategies under ξ_n . Then, there is a sequence $(u_n) \rightarrow u$ such that each σ_n is a 0-sequential equilibrium in mixed strategies under ξ_n in $G(u_n)$.*

Proof. Let $(\sigma_n, \varepsilon_n, \xi_n) \rightarrow (\sigma, 0, 0)$ be such that each σ_n is a ε_n -sequential equilibrium in mixed strategies under ξ_n . We assume without loss (i.e., taking a subsequence if necessary) that the set S_i^* of s_i such that $\sigma_{i,n}(s_i) > \xi_{i,n}(s_i)$ remains fixed along the sequence.¹⁵ We assume that (σ_n) supports some conditional belief system γ . Note that γ is consistent. Note also that if $I \in S_i^{s_i}$ for some $s_i \in S_i^*$, then for all $\hat{s}_i \in S_i^I$ we have

$$\lim_{n \rightarrow \infty} (u_i(s_i, \sigma_{-i,n}|I) - u_i(\hat{s}_i, \sigma_{-i,n}|I)) \geq \lim_{n \rightarrow \infty} \varepsilon_n = 0$$

¹⁵Note that the statement of Lemma A.2 applies to all sequences (not subsequences). Nevertheless, the claim can then be applied to each subsequence defined by the different possible values of S_i^* .

Hence, we have that

$$u_i(s_i, \gamma_{-i}^I | I) = \max_{\hat{s}_i \in S_i^I} u_i(\hat{s}_i, \gamma_{-i}^I | I) .$$

Let \mathcal{I}_i^* denote $\cup_{s_i \in S_i^*} \mathcal{I}_i^{s_i}$.

Note that there is a transitive “precedence” relationship \prec for the elements in \mathcal{I}_i defined by

$$I \prec I' \iff \text{for all } h' \in I' \text{ there is some } h \in I \text{ such that } h \prec h' .$$

For all pairs of distinct information sets $I, I' \in \mathcal{I}_i$, either $I \prec I'$, or $I' \prec I$, or there is no terminal history that passes through both of them. Note that if $I \prec I'$ then there is some $a_i \in A^I$ such that for all $h' \in I'$ there is some $h \in I$ such that $(h, a_i) \leq h'$.

We now suggest an algorithm to obtain the desired sequence of payoff functions. We initialize $\hat{u}^0 \equiv u$. In the first step, we let \mathcal{I}_i^1 be the set containing all information sets $I \in \mathcal{I}_i^*$ with the property that there is no $I' \succ I$ with $I' \in \mathcal{I}_i^*$. Fix some $I \in \mathcal{I}_i^1$ and some $a_i \in A^I$ such that $s_{i,I} = a_i$ for some $s_i \in S_i^* \cap S_i^I$. Then, take some $s_i \in S_i^I$ such that $s_{i,I} = a_i$ and $\sigma_{i,n}(s_i) > \xi_{i,n}(s_i)$. For all $z \in Z^{a_i}$ (where recall Z^{a_i} denotes the set of terminal histories containing a_i), define¹⁶

$$\hat{u}_{i,n}^1(z) := \hat{u}_{i,n}^0(z) + u_i(s_i, \gamma_{-i}^I | I) - \hat{u}_{i,n}^0(s_i, \sigma_{-i,n} | I) ,$$

If instead $a_i \in A^I$ is such that there is no $s_i \in S_i^* \cap S_i^I$ with $s_{i,I} = a_i$, we define

$$\hat{u}_{i,n}^1(z) := \hat{u}_{i,n}^0(z) + \max_{s_i \in S_i^I \setminus S_i^*} (u_i(s_i, \gamma_{-i}^I | I) - \hat{u}_{i,n}^0(s_i, \sigma_{-i,n} | I)) \quad (\text{A.1})$$

for all $z \in Z^{a_i}$. For the rest of the terminal histories z (i.e., those that do not pass through any of the information sets in \mathcal{I}_i^1), we define $\hat{u}_{i,n}^1(z) := \hat{u}_{i,n}^0(z)$. Note that, for all $s_i \in S_i^* \cap S_i^I$ and $\hat{s}_i \in S_i^I$, we have

$$\hat{u}_{i,n}^1(s_i, \sigma_{-i,n} | I) = u_i(s_i, \gamma_{-i}^I | I) \geq u_i(\hat{s}_i, \gamma_{-i}^I | I) = \hat{u}_{i,n}^1(\hat{s}_i, \sigma_{-i,n} | I) .$$

Clearly, $\|\hat{u}_{i,n}^1 - \hat{u}_{i,n}^0\| \rightarrow 0$ as $n \rightarrow \infty$.

¹⁶Importantly, because player i does not play after I , the right side of expression (A.1) only depends on s_i through $s_{i,I}$ (i.e., if $\hat{s}_i \in S_i^I$ and $\hat{s}_{i,I} = s_{i,I}$, then the right side of (A.1) is the same for both s_i and \hat{s}_i). Hence, the value of $\hat{u}_{i,n}^1(z)$ is independent of the choice of s_i .

We now proceed iteratively in j . Let \mathcal{I}_i^j be the set containing all information sets $I \in \mathcal{I}_i^*$ with the property that there is no $I' \succ I$ with $I' \in \mathcal{I}_i^* \setminus \bigcup_{j=1}^{j-1} \mathcal{I}_i^j$. Fix some $I \in \mathcal{I}_i^j$. Note first that if $s_i, \hat{s}_i \in S_i^* \cap S_i^I$ and $s_{i,I} = \hat{s}_{i,I}$, then

$$\hat{u}_{i,n}^{j-1}(s_i, \sigma_{-i,n}|I) = \hat{u}_{i,n}^{j-1}(\hat{s}_i, \sigma_{-i,n}|I) .$$

The reason is that both induce the same distribution over information sets in \mathcal{I}_i^{j-1} and, by the definition of $\hat{u}_{i,n}^{j-1}$, they achieve the same continuation payoff in each of them. Fix some $a_i \in A^I$ such that $s_{i,I} = a_i$ for some $s_i \in S_i^* \cap S_i^I$. Then, take some $s_i \in S_i^I$ such that $\sigma_{i,n}(s_i) > \xi_{i,n}(s_i)$. For all $z \in Z^{a_i}$, define¹⁷

$$\hat{u}_{i,n}^j(z) := \hat{u}_{i,n}^{j-1}(z) + u_i(s_i, \gamma_{-i}^I|I) - \hat{u}_{i,n}^{j-1}(s_i, \sigma_{-i,n}|I) , \quad (\text{A.2})$$

If instead $a_i \in A^I$ is such that there is no $s_i \in S_i^* \cap S_i^I$ with $s_{i,I} = a_i$, we define

$$\hat{u}_{i,n}^j(z) := \hat{u}_{i,n}^{j-1}(z) + \max_{s_i \in S_i^I \setminus S_i^*} (u_i(s_i, \gamma_{-i}^I|I) - \hat{u}_{i,n}^{j-1}(s_i, \sigma_{-i,n}|I))$$

for all $z \in Z^{a_i}$. For the rest of the terminal histories z (i.e., those that do not pass through any of the information sets in \mathcal{I}_i^j), we define $\hat{u}_{i,n}^j(z) := \hat{u}_{i,n}^{j-1}(z)$. Note also that, for all $s_i, \hat{s}_i \in S_i^* \cap S_i^I$, we have

$$\hat{u}_{i,n}^j(s_i, \sigma_{-i,n}|I) = \hat{u}_{i,n}^j(\hat{s}_i, \sigma_{-i,n}|I) .$$

Hence, for all $s_i \in S_i^* \cap S_i^I$ and $\hat{s}_i \in S_i^I$, we have

$$\hat{u}_{i,n}^j(s_i, \sigma_{-i,n}|I) = u_i(s_i, \gamma_{-i}^I|I) \geq u_i(\hat{s}_i, \gamma_{-i}^I|I) = \hat{u}_{i,n}^j(\hat{s}_i, \sigma_{-i,n}|I) .$$

Clearly, $\|\hat{u}_{i,n}^j - \hat{u}_{i,n}^{j-1}\| \rightarrow 0$ as $n \rightarrow \infty$.

Let u_n be the result of applying the previous algorithm to all players. Note that for all $i \in N$, $I \in \mathcal{I}_i^*$, $s_i \in S_i^* \cap S_i^I$, and $\hat{s}_i \in S_i^I$, we have

$$u_{i,n}(s_i, \sigma_{-i,n}|I) = u_i(s_i, \gamma_{-i}^I|I) \geq u_i(\hat{s}_i, \gamma_{-i}^I|I) = u_{i,n}(\hat{s}_i, \sigma_{-i,n}|I) .$$

Hence, σ_n is a 0-sequential equilibrium in mixed strategies under ξ_n of $G(u_n)$. \square

Then, by Lemma A.1, the “only if” part of the statement of Proposition A.1 holds. \square

¹⁷Note that, again, the right side of expression (A.2) only depends on s_i through $s_{i,I}$. Hence, the value of $\hat{u}_{i,n}^j(z)$ is independent of the choice of s_i .

B Proofs

This appendix contains the proofs. Note that some of the proofs follow from Proposition A.1 (and Lemmas A.1 and A.2), which does not follow from any of the other results.

Proof of Proposition 2.1

Proof. Proposition 2.1 follows straightforwardly from Definition 2.3. \square

Proof of Proposition 3.1 and Lemma 3.1

Proof. Proof of Lemma 3.1. Assume σ is a Nash equilibrium of $G(\xi, u)$. Assume for a contradiction that there is some s_i with $\sigma_i(s_i) > \xi_i(s_i)$, some $I \in \mathcal{I}^{s_i}$, and some $\hat{s}_i \in S_i^I$ such that $u_i(s_i, \sigma_{-i}|I) < u_i(\hat{s}_i, \sigma_{-i}|I)$. Let \hat{s}'_i be the strategy coinciding with s_i in all information sets except I and all that succeed it, where it coincides with \hat{s} . Note that $\hat{s}'_i \in S_i^I$ and that $u_i(\hat{s}'_i|I, \sigma_{-i}) = u_i(\hat{s}_i, \sigma_{-i}|I)$. Now we have

$$\begin{aligned} u_i(s_i, \sigma_{-i}) - u_i(\hat{s}'_i, \sigma_{-i}) &= \sum_{z \in Z^I} \mathbb{P}^{s_i, \sigma_{-i}}(z) u_i(z) - \sum_{z \in Z^I} \mathbb{P}^{\hat{s}'_i, \sigma_{-i}}(z) u_i(z) \\ &= \mathbb{P}^{s_i, \sigma_{-i}}(I) (u_i(s_i, \sigma_{-i}|I) - u_i(\hat{s}'_i, \sigma_{-i}|I)) < 0. \end{aligned}$$

This contradicts that $\sigma_i(s_i) > \xi_i(s_i)$ and σ is a Nash equilibrium of $G(\xi, u)$.

Assume now that for all $i \in N$, $s_i \in S_i$ with $\sigma_i(s_i) > \xi_i(s_i)$, and $I \in \mathcal{I}_i^{s_i}$, we have that s_i is sequentially rational at I . Let \mathcal{I}_i^* be the set of player i 's information sets that are not preceded by any other information set by player i . Note that $\mathcal{I}_i^* \subset \mathcal{I}_i^{s_i}$ for all $s_i \in S_i$. Then, we have that, for all $s_i \in S_i$,

$$u_i(s_i, \sigma_{-i}) = \sum_{I \in \mathcal{I}_i^*} \mathbb{P}^{s_i, \sigma_{-i}}(I) u_i(s_i, \sigma_{-i}|I).$$

Note that $\mathbb{P}^{s_i, \sigma_{-i}}(I)$ is independent of s_i . Assume for a contradiction that σ is not a Nash equilibrium of $G(\xi, u)$. Fix some $s_i, \hat{s}_i \in S_i$ with $\sigma_i(s_i) > \xi_i(s_i)$ and $u_i(s_i, \sigma_{-i}) < u_i(\hat{s}_i, \sigma_{-i})$. Note that, by our assumption, s_i is sequentially rational at all $I \in \mathcal{I}_i^{s_i}$. It then follows that

$$u_i(s_i, \sigma_{-i}) - u_i(\hat{s}_i, \sigma_{-i}) = \sum_{I \in \mathcal{I}_i^*} \mathbb{P}^{s_i, \sigma_{-i}}(I) (u_i(s_i, \sigma_{-i}|I) - u_i(\hat{s}_i, \sigma_{-i}|I)).$$

Then, the assumption that $u_i(s_i, \sigma_{-i}) < u_i(\hat{s}_i, \sigma_{-i})$ implies that $u_i(s_i, \sigma_{-i}|I) < u_i(\hat{s}_i, \sigma_{-i}|I)$ for some $I \in \mathcal{I}_i^*$, but this contradicts that s_i is sequentially rational for all $I \in \mathcal{I}_i^{s_i}$.

Proof that outcomes of sequential equilibria in mixed strategies are outcomes of weakly sequential equilibria. Let σ be a sequential equilibrium in mixed strategies. Let $(\sigma_n, u_n, \xi_n) \rightarrow (\sigma, u, 0)$ such that each σ_n is a Nash equilibrium for $G(\xi_n, u_n)$. Let β denote the behavior strategy with the same outcome as σ . Similarly, for each n , let β_n denote the behavior strategy with the same outcome as σ_n . Without loss of generality, assume that (σ_n) supports an assessment (β, μ) .

Assume for a contradiction that (β, μ) is not a weakly sequential equilibrium. Let $i \in N$, $I \in \mathcal{I}_i$, and $\hat{s}_i \in S_i^I$, be such that I is β_i -relevant and $u_i(\hat{s}_i, \beta_{-i}, \mu|I) > u_i(\beta, \mu|I)$. Note that, for all $s_i \in S_i^I$,

$$u_{i,n}(s_i, \sigma_{-i,n}|I) = u_{i,n}(s_i, \beta_{-i,n}|I) \xrightarrow{n \rightarrow \infty} u_i(s_i, \beta_{-i}, \mu|I).$$

Hence, there is some $s_i \in S_i$ such that $\sigma_i(s_i) > 0$ and $u_i(s_i, \beta_{-i}, \mu|I) \leq u_i(\beta, \mu|I)$. Then, we have

$$\begin{aligned} 0 &< u_i(\hat{s}_i, \beta_{-i}, \mu|I) - u_i(\beta, \mu|I) \leq u_i(\hat{s}_i, \beta_{-i}, \mu|I) - u_i(s_i, \beta_{-i}, \mu|I) \\ &= \lim_{n \rightarrow \infty} (u_{i,n}(\hat{s}_i, \sigma_{-i,n}|I) - u_{i,n}(s_i, \sigma_{-i,n}|I)) \leq 0, \end{aligned}$$

where the last inequality holds because, since $\sigma_{i,n}(s_i) > \xi_{i,n}(s_i)$ for n large enough, we have that $u_{i,n}(\hat{s}_i, \sigma_{-i,n}|I) \leq u_{i,n}(s_i, \sigma_{-i,n}|I)$ for n large enough (by Lemma 3.1).

Proof that outcomes of weakly sequential equilibria are outcomes of sequential equilibria in mixed strategies. Let (β, μ) be a weakly sequential equilibrium and let (β_n) support it. Let σ denote a mixed strategy with the outcome as β . Similarly, for each n , let σ_n denote a mixed strategy with the outcome as β_n such that $\sigma_n \rightarrow \sigma$. Define

$$\xi_{i,n}(s_i) := \begin{cases} \sigma_{i,n}(s_i) & \text{if } \sigma_i(s_i) = 0, \\ 2^{-n} \sigma_{i,n}(s_i) & \text{otherwise,} \end{cases}$$

for all $n \in \mathbb{N}$, $i \in N$, and $s_i \in S_i$. We will show that there is a sequence $(\varepsilon_n) \rightarrow 0$ such that each σ_n is a ε_n -sequential equilibrium in mixed strategies under ξ_n (as defined in Section

A). By Proposition A.1, we will then deduce that σ is a sequential equilibrium in mixed strategies.

Fix some $i \in N$, $s_i \in S_i$ with $\sigma_i(s_i) > 0$, and $I \in \mathcal{I}_i^{s_i}$. Note that

$$u_i(s_i, \sigma_{-i,n} | I) = u_i(s_i, \beta_{-i,n} | I) \xrightarrow{n \rightarrow \infty} u_i(s_i, \beta_{-i}, \mu | I).$$

Then, for all $\hat{s}_i \in S_i^I$, we have

$$\lim_{n \rightarrow \infty} (u_i(s_i, \sigma_{-i,n} | I) - u_i(\hat{s}_i, \sigma_{-i,n} | I)) = u_i(s_i, \beta_{-i}, \mu | I) - u_i(\hat{s}_i, \beta_{-i}, \mu | I) \geq 0,$$

where the last inequality holds because (β, μ) is a weakly sequential equilibrium. It then follows that there is a sequence $(\varepsilon_n) \rightarrow 0$ such that each σ_n is a ε_n -sequential equilibrium in mixed strategies under ξ_n . \square

Proof of Corollary 3.1

Proof. **Proof of the “if” part of Corollary 3.1.** Let σ be part of a consistent and weakly sequentially rational conditional assessment (σ, γ) . Let (σ_n) support (σ, γ) . For all $n \in \mathbb{N}$, $i \in N$ and $s_i \in S_i$, define

$$\xi_{i,n}(s_i) := \begin{cases} \sigma_{i,n}(s_i) & \text{if } \sigma_i(s_i) = 0, \\ 2^{-n} \sigma_{i,n}(s_i) & \text{otherwise.} \end{cases}$$

Note that $\sigma_{i,n}(s_i) \geq \xi_{i,n}(s_i)$, and $\sigma_{i,n}(s_i) > \xi_{i,n}(s_i)$ if and only if $\sigma_i(s_i) > 0$. Let $i \in N$ and $s_i \in S_i$ be such that $\sigma_i(s_i) > 0$. Then, for any $I \in \mathcal{I}_i^{s_i}$ and $\hat{s}_i \in S_i^I$, we have

$$\lim_{n \rightarrow \infty} (u_i(s_i, \sigma_{-i,n} | I) - u_i(\hat{s}_i, \sigma_{-i,n} | I)) = u_i(s_i, \gamma_{-i}^I | I) - u_i(\hat{s}_i, \gamma_{-i}^I | I) \geq 0.$$

As a result, there is a sequence $(\varepsilon_n) \rightarrow 0$ such that, for all $n \in \mathbb{N}$, if $\sigma_{i,n}(s_i) > \xi_{i,n}(s_i)$ then $u_i(s_i, \sigma_{-i,n} | I) \geq u_i(\hat{s}_i, \sigma_{-i,n} | I) - \varepsilon_n$ for all $I \in \mathcal{I}_i^{s_i}$ and $\hat{s}_i \in S_i^I$. This implies that each σ_n is an ε_n -sequential equilibrium in mixed strategies under ξ_n . By Proposition A.1, σ is a sequential equilibrium in mixed strategies.

Proof of the “only if” part of Corollary 3.1. Let σ be a sequential equilibrium in mixed strategies. Let $(\sigma_n, u_n, \xi_n) \rightarrow (\sigma, u, 0)$ such that each σ_n is a Nash equilibrium of $G(\xi_n, u_n)$.

Assume without loss that (σ_n) supports a conditional assessment (σ, γ) . Note that (σ, γ) is consistent. Assume for a contradiction that (σ, γ) is not weakly sequentially rational. Then, there is some $i \in N$, $s_i \in S_i$, $I \in \mathcal{I}_i^{s_i}$, and $\hat{s}_i \in S_i^I$ such that

$$u_i(s_i, \gamma_{-i}^I | I) < u_i(\hat{s}_i, \gamma_{-i}^I | I).$$

This implies that

$$0 > u_i(s_i, \gamma_{-i}^I | I) - u_i(\hat{s}_i, \gamma_{-i}^I | I) = \lim_{n \rightarrow \infty} (u_{i,n}(s_i, \sigma_{-i,n} | I) - u_{i,n}(\hat{s}_i, \sigma_{-i,n} | I)),$$

and so $u_{i,n}(s_i, \sigma_{-i,n} | I) < u_{i,n}(\hat{s}_i, \sigma_{-i,n} | I)$ for some n . By Lemma 3.1, this contradicts that σ_n is a Nash equilibrium of $G(\xi_n, u_n)$. \square

Proof of Corollary 3.2

Proof. As explained in the main text, the proof is straightforward. \square

Proof of Lemma 4.1

Proof. As explained in the main text, the proof is straightforward. \square

Proof of Proposition 4.1

Proof. The key step in Blume and Zame (1994) is to show that for any sequential equilibrium β there is a sequence $(\beta_n, \eta_n, u_n) \rightarrow (\beta, 0, u)$ such that each β_n is a Nash equilibrium of $G(\eta_n, u_n)$ (their Proposition B). Given Definition 2.3, the analogous property holds trivially for sequential equilibria in mixed strategies.

We first define the graph of the *perturbed game equilibrium in mixed strategies correspondence (GPNE)* as¹⁸

$$\begin{aligned} \text{GPNE} = \{ (u, \xi, \sigma) \in U \times \mathbb{R}_{++}^S \times \Sigma \mid \sigma \in \Sigma(\xi) \ \forall i \ \forall \hat{\sigma}_i \in \Sigma_i(\xi), \\ v_i(\hat{\sigma}_i, \sigma_{-i}, u) \leq v_i(\sigma_i, \sigma_{-i}, u) \}, \end{aligned}$$

¹⁸Blume and Zame (1994) also define the perturbed game equilibrium correspondence (PNE), and then GPNE is its graph. Nevertheless, for the argument here, we only need the graph.

where we use $v_i(\sigma, u) := u_i(\sigma)$ to follow Blume and Zame's notation more closely, and where $U \equiv \mathbb{R}^{N \times Z}$ and $\Sigma(\xi) \equiv \{\sigma \in \Sigma \mid \sigma \geq \xi\}$. Note that GPNE is the set of triples (u, ξ, σ) where σ is a Nash equilibrium of $G(\xi, u)$. We now define the *graph of the perfect equilibrium in mixed strategies correspondence (GPE)* as

$$\text{GPE} = \left\{ (u, \sigma) \in U \times \Sigma \mid \forall \varepsilon > 0 \ \forall \delta > 0 \ \exists \xi \in \mathbb{R}_{++}^S \ \exists \sigma' \in \Sigma(\xi) \right. \\ \left. (u, \xi, \sigma') \in \text{GPNE} \wedge \|\xi\| < \delta \wedge \|\sigma - \sigma'\| < \varepsilon \right\}.$$

Now, from Definition 4.1, we have that GPE is the set of pairs (u, σ) where σ is a perfect equilibrium in mixed strategies when the payoff is u . Finally, we define the *graph of the sequential equilibrium in mixed strategies correspondence (GSE)* as

$$\text{GSE} = \left\{ (u, \sigma) \in U \times \Sigma \mid \forall \varepsilon > 0 \ \forall \delta > 0 \ \exists u' \in U \ \exists \xi \in \mathbb{R}_{++}^S \ \exists \sigma' \in \Sigma(\xi) \right. \\ \left. (u, \xi, \sigma') \in \text{GPNE} \wedge \|\xi\| < \delta \wedge \|u - u'\| < \delta \wedge \|\sigma - \sigma'\| < \varepsilon \right\}.$$

Now, from Definition 2.3, we have that GSE is the set of pairs (u, σ) where σ is a sequential equilibrium in mixed strategies when the payoff is u .

The key for the rest of the proof is to realize that the expressions GPNE, GPE, and GSE coincide *exactly* with the analogs in Blume and Zame (1994). The difference is that our spaces of strategies and trembles are set in mixed strategies, while their spaces of strategies and trembles are those set in behavior strategies. Nonetheless, this difference is irrelevant in the proof of their result establishing the generic equivalence between perfect and sequential equilibria for generic payoffs (their Theorem 4). Our proof is then complete. \square

Proof of Proposition 5.1

Proof. The proof is divided into two parts.

Part 1. Let σ be purifiable. Let $(v_n) \rightarrow \delta_u$ and $(\sigma_n) \rightarrow \sigma$ be such that each σ_n is a Nash equilibrium of $G(v_n)$. Let $(\varepsilon_n, \varepsilon'_n) \rightarrow (0, 0)$ be such that

$$v_{i,n}(\{\hat{u}_i \in U_i \mid \|\hat{u}_{i,n} - u_i\| < \varepsilon_n\}) \geq 1 - \varepsilon'_n$$

(which exist because $(\nu_n) \rightarrow \delta_u$). For each $i \in N$ and $s_i \in S_i$, let (ξ_n) be such that

$$\xi_{i,n}(s_i) = \begin{cases} \sigma_{i,n}(s_i) & \text{if } \sigma_i(s_i) = 0, \\ 2^{-n} \sigma_{i,n}(s_i) & \text{if } \sigma_i(s_i) > 0, \end{cases}$$

for all $i \in N$, $s_i \in S_i$, and $n \in \mathbb{N}$. We want to show that there is some ε_n'' such that $\sigma_i(s_i) > 0$ only if $u_i(s_i, \sigma_{-i,n}) \geq u_i(\hat{s}_i, \sigma_{-i,n}) - \varepsilon_n''$ for all $\hat{s}_i \in S_i$ and then use Proposition A.1.

Take some $i \in N$ and $s_i \in S_i$ such that $\sigma_i(s_i) > 0$. Note that if n is high enough, then $\sigma_{i,n}(s_i) > \varepsilon_n'$. For each s_i , let $U_{i,n}(s_i)$ be the set of payoffs $\hat{u}_{i,n}$ such that s_i is optimal against $\sigma_{-i,n}$. Then, we have that there is some $\hat{u}_i \in U_{i,n}(s_i)$ such that $\|\hat{u}_i - u_i\| \leq \varepsilon_n$. Then, an argument similar to the one in the proof of Lemma 3.1 implies that

$$\hat{u}_i(s_i, \sigma_{-i,n}|I) \geq \hat{u}_i(\hat{s}_i, \sigma_{-i,n}|I)$$

for all $I \in \mathcal{I}_i^{s_i}$ and $\hat{s}_i \in S_i^I$. Using that $\|\hat{u}_i - u_i\| \leq \varepsilon_n$ twice, we have

$$u_i(s_i, \sigma_{-i,n}|I) \geq u_i(\hat{s}_i, \sigma_{-i,n}|I) - 2\varepsilon_n$$

for all $I \in \mathcal{I}_i^{s_i}$ and $\hat{s}_i \in S_i^I$; that is, $\varepsilon_n'' = 2\varepsilon_n$. Then, by Proposition A.1, σ is a sequential equilibrium in mixed strategies.

Part 2. Let σ be a sequential equilibrium in mixed strategies. Let $(\sigma_n, \xi_n, u_n) \rightarrow (\sigma, 0, u)$ be such that each σ_n is a Nash equilibrium of $G(\xi_n, u_n)$. We assume, without loss of generality, that (i) $\sigma_{i,n}(s_i) > \xi_{i,n}(s_i)$ for all n whenever $\sigma_i(s_i) > 0$, and (ii) $\sigma_{i,n}(s_i) = \xi_{i,n}(s_i)$ for all n whenever $\sigma_i(s_i) = 0$. As before, for each s_i , let $U_{i,n}(s_i)$ be the set of payoffs $\hat{u}_{i,n}$ such that s_i is optimal against $\sigma_{-i,n}$. We will treat $\{U_{i,n}(s_i) | s_i \in S_i\}$ as a partition of U_i .¹⁹

We now define the distribution $\nu_{i,n}$. We assume that $\nu_{i,n}$ is absolutely continuous, and we let $f_{i,n}$ denote its density. We will define $f_{i,n}$ in each of the elements of U_i , and then argue it integrates to 1. We consider two types of elements $U_i(s_i)$:

1. Consider first some s_i such that $\sigma_i(S_i^{s_i}) = 0$. In this case, for all $\hat{u}_{i,n} \in U_{i,n}(s_i)$, we define

$$f_{i,n}(\hat{u}_{i,n}) := \frac{\sigma_{i,n}(s_i)}{\|U_{i,n}(s_i)\|}$$

¹⁹Technically, $\{U_{i,n}(s_i) | s_i \in S_i\}$ is not a partition of U_i because its elements may intersect for payoff functions where player i is indifferent between two or more strategies (note also that it may be $U_{i,n}(s_i) = U_{i,n}(\hat{s}_i)$ for some $s_i \neq \hat{s}_i$ (when $\hat{s}_i \in S_i^{s_i}$), but these are the same element of the partition). Nevertheless, because these intersections will have probability zero, we will dismiss them in our analysis. Note also that $U_{i,n}(s_i) = U_{i,n}(\hat{s}_i)$ whenever $\hat{s}_i \in S_i^{s_i}$.

for all $\hat{u}_{i,n} \in U_{i,n}(s_i)$, where $\|U_{i,n}(s_i)\|$ indicates the Lebesgue measure of $U_{i,n}(s_i) \subset \mathbb{R}^Z$.

2. Consider now the case where $\sigma_i(S_i^{s_i}) > 0$. Because σ_n is a Nash equilibrium of $G(\xi_n, u_n)$ and $\sigma_{i,n}(\check{s}_i) > \xi_{i,n}(\check{s}_i)$ for some $\check{s}_i \in S_i^{s_i}$, we have that $u_{i,n}(s_i, \sigma_{-i,n}) \geq u_{i,n}(\hat{s}_i, \sigma_{-i,n})$ for all $\hat{s}_i \in S_i$. This implies that, for all $\varepsilon > 0$, we have

$$\text{int}(U_{i,n}(s_i) \cap B_\varepsilon(u_{i,n})) \neq \emptyset,$$

where $B_\varepsilon(u_{i,n}) \subset \mathbb{R}^Z$ indicates the ball of radius ε around $u_{i,n}$. To see this, let $Z^{s_i, \sigma_{-i,n}}$ indicate the set of terminal histories achieved with positive probability under $(s_i, \sigma_{-i,n})$. Note that, for each $z \in Z^{s_i, \sigma_{-i,n}}$, we have that $\mathbb{P}^{s_i, \sigma_{-i,n}}(z) \geq \mathbb{P}^{\hat{s}_i, \sigma_{-i,n}}(z)$ for all $\hat{s}_i \in S_i$. Then, let $v_i \in \mathbb{R}^Z$ denote a vector satisfying that $v_{i,z} > 0$ if $z \in Z^{s_i, \sigma_{-i,n}}$ and $v_{i,z} < 0$ if $z \notin Z^{s_i, \sigma_{-i,n}}$. Hence, we have that

$$(u_{i,n} + \varepsilon' v_i)(s_i, \sigma_{-i,n}) \geq (u_{i,n} + \varepsilon' v_i)(\hat{s}_i, \sigma_{-i,n})$$

for all $\varepsilon' > 0$ and $\hat{s}_i \in S_i$. In particular, if ε' is small enough, we have $(u_{i,n} + \varepsilon' v_i) \in B_\varepsilon(u_{i,n})$. Take now a sequence $(\varepsilon_n) \rightarrow 0$. Define

$$f_{i,n}(\hat{u}_{i,n}) := (1 - e^{-n}) \frac{\sigma_{i,n}(s_i)}{\|U_{i,n}(s_i) \cap B_{\varepsilon_n}(u_{i,n})\|}$$

if $\hat{u}_{i,n} \in U_{i,n}(s_i) \cap B_{\varepsilon_n}(u_{i,n})$, and

$$f_{i,n}(\hat{u}_{i,n}) := e^{-n} \frac{\sigma_{i,n}(s_i)}{\|U_{i,n}(s_i) \setminus B_{\varepsilon_n}(u_{i,n})\|}$$

if $\hat{u}_{i,n} \in U_{i,n}(s_i) \setminus B_{\varepsilon_n}(u_{i,n})$, where each $\varepsilon_n > 0$ is taken small enough that the denominators in the previous expressions are positive.

It is clear that $f_{i,n}$ integrates to 1 over U_i . Note also that, by construction, σ_n is a Nash equilibrium for $v_{i,n}$. Furthermore, note that the probability under $v_{i,n}$ that \hat{u}_i is at distance higher than ε_n from u_n is

$$\sum_{s_i | \sigma_{i,n}(s_i) = 0} \sigma_i(s_i) + \sum_{s_i | \sigma_{i,n}(s_i) > 0} e^{-n} \sigma_{i,n}(s_i),$$

which tends to 0 as $n \rightarrow \infty$. Because $u_{i,n} \rightarrow u_i$ as $n \rightarrow \infty$, we have that $(v_{i,n}) \rightarrow \delta_{u_i}$, hence the proof is complete. \square

C Other derivations

C.1 Quasi- and weakly quasi-perfect equilibria

In this section, we provide a characterization of perfectness in mixed strategies using behavioral strategies. To do so, we first recall the concept of quasi-perfect equilibria, introduced in van Damme (1984).²⁰

Definition C.1 (van Damme, 1984). β is a *quasi-perfect equilibrium* if there is a fully-mixed sequence $(\beta_n) \rightarrow \beta$ such that, for each $i \in N$ and $I \in \mathcal{I}_i$, β_i is sequentially rational at I given $\beta_{-i,n}$.

As explained by van Damme (1984), “the difference between this concept and Selten’s perfectness concept is that the latter requires that each player at every information set takes a choice which is optimal against mistakes of all players (including the player himself), whereas the quasi-perfectness concept requires that at every information set a choice is taken which is optimal against mistakes of the other players” (p. 2). That is, in the perturbed games used to approximate perfect equilibria, players take into account their own future mistakes when assessing the optimality of playing a given action (recall the discussion of Figure 2 in Section 4.2, while players only take into account other players’ mistakes when approximating quasi-perfect equilibria. van Damme (1984) shows that a proper equilibrium of a normal form game induces a quasi-perfect equilibrium in every extensive form game having this normal form.

The following proposition shows that quasi-perfectness is stronger than perfectness in mixed strategies in all games (recall that Corollary 4.1 establishes that an analogous result for perfectness for games with generic payoffs).

Proposition C.1. *Outcomes of quasi-perfect equilibria are outcomes of perfect equilibria in mixed strategies.*

We now define the concept of weakly perfect equilibria, which is analogous to quasi-perfect equilibria but requires sequential rationality only on relevant information sets.

²⁰In the statement, recall that $(\beta_i, \beta_{-i,n})$ is sequentially rational at I if β_i maximizes $u(\hat{\beta}_i, \beta_{-i,n}|I)$ for all $\hat{\beta}_i \in \mathcal{B}_i$, where such conditional payoff is well defined because $\beta_{-i,n}$ has full support.

Definition C.2. β is a *weakly quasi-perfect equilibrium* if there is a fully-mixed sequence $\beta_n \rightarrow \beta$ such that, for each $i \in N$ and $I \in \mathcal{I}_i$ that is β_i -relevant, we have $(\beta_i, \beta_{-i,n})$ is sequentially rational at I .

It is clear that quasi-perfect equilibria are weakly quasi-perfect, because they only require sequential rationality in some information sets instead of all. The following proposition provides a characterization of perfect equilibria in mixed strategies in terms of behavior strategies. Note that it is analogous to Proposition 3.1.

Proposition C.2. *The sets of outcomes of weakly quasi-perfect equilibria and perfect equilibria in mixed strategies coincide.*

Proofs of Propositions C.1 and C.2

Proof of Propositions C.1 and C.2. Because quasi-perfect equilibria are weakly quasi-perfect, Proposition C.1 follows from Proposition C.2.

Proof that outcomes of weakly quasi-perfect equilibria are outcomes of perfect equilibria in mixed strategies. Fix a weakly quasi-perfect equilibrium β . Let (β_n) support β . Let σ and σ_n have the same outcome as β and β_n , respectively. Define the sequence (ξ_n) as

$$\xi_{i,n}(s_i) := \begin{cases} \sigma_{i,n}(s_i) & \text{if } \sigma_i(s_i) = 0, \\ 2^{-n} \sigma_{i,n}(s_i) & \text{otherwise.} \end{cases}$$

Fix some $s_i \in S_i$ with $\sigma(s_i) > 0$. Assume for a contradiction that there is some n and $\hat{s}_i \in S_i$ such that $u_i(s_i, \sigma_{-i,n}) < u_i(\hat{s}_i, \sigma_{-i,n})$. From the proof of Lemma 3.1, we have that there must be some $I \in \mathcal{I}_i^{s_i} \cap \mathcal{I}_i^{\hat{s}_i}$ where

$$u_i(s_i, \sigma_{-i,n} | I) < u_i(\hat{s}_i, \sigma_{-i,n} | I) .$$

We then have that σ_i is not sequentially at I given $\sigma_{-i,n}$, which implies β_i is not sequentially rational at I given $\beta_{-i,n}$, a contradiction.

Proof that outcomes of perfect equilibria in mixed strategies are outcomes of weakly quasi-perfect equilibria. Fix a perfect equilibrium in mixed strategies σ_n . Let (σ_n, ξ_n)

support σ . Let β and β_n have the same outcome as σ and σ_n , respectively, and note that $\beta_n \rightarrow \beta$. Assume for a contradiction that there is some β_i -relevant information set $I \in \mathcal{I}_i$ and some $\hat{s}_i \in S_i^I$ such that

$$u_i(\hat{s}_i, \beta_{-i,n} | I) > u_i(\beta_i, \beta_{-i,n} | I) .$$

This implies that there is some $s_i \in S_i^I$ such that $\sigma_i(s_i) > 0$ and

$$u_i(s_i, \sigma_{-i,n} | I) = u_i(s_i, \beta_{-i,n} | I) \leq u_i(\sigma_i, \beta_{-i,n} | I) < u_i(\hat{s}_i, \beta_{-i,n} | I) = u_i(\hat{s}_i, \sigma_{-i,n} | I) .$$

By the same argument in the proof of Lemma 3.1, this implies that σ_n is not a Nash equilibrium of $G(\xi_n, u)$, a contradiction. \square

C.2 Relationship to Mailath et al. (1993)

Mailath et al. (1993) investigate whether we can reproduce extensive-form reasoning (like sequential equilibrium) inside the normal form. In their answer, they define normal-form sequential equilibria, a concept we will now describe and compare to sequential equilibria in mixed strategies.

Mailath et al. first define *normal-form information sets* as product blocks of the strategy space that capture strategic independence. Beliefs on each block come from conditional limits of completely mixed profiles (conditionally convergent sequences). A *normal-form sequential equilibrium* (NFSE) requires best responses on every normal-form information set with respect to those conditional beliefs.

They prove existence by linking proper equilibria to their concept: a proper equilibrium is an NFSE. They also prove a bridge to the extensive form: any normal-form sequential equilibrium of a pure-strategy reduced normal form (PRNF, where obtained after collapsing duplicate pure strategies) induces a sequential equilibrium in every extensive-form game with that PRNF. The converse holds if the same supporting sequence induces sequential equilibrium in every such extensive form. Extending these equivalences to MRNF (mixed-strategy reduced normal form, where obtained after collapsing strategies that are equivalent to mixed combinations of others) is delicate: a single MRNF profile need not work across all extensive forms with that MRNF.

Mailath et al. and our approaches differ in primitives. Our construction stays in the extensive form and selects limits of Nash equilibria of nearby perturbed games, highlighting the role of uncertainty about rationality and payoffs in equilibrium selection. Their construction, instead, allows sequential reasoning inside the normal form via blockwise optimality and conditional limits. The former emphasizes approximability and selection through perturbations; the latter emphasizes strategic independence and mapping between PRNF and extensive-form.

References

- BHASKAR, V. AND M. B. STINCHCOMBE (2024): “Extensive games with randomly disturbed payoffs: A new approach to equilibrium refinements,” Mimeo.
- BLUME, L. AND W. R. ZAME (1994): “The algebraic geometry of perfect and sequential equilibrium,” *Econometrica*, 62, 783–794.
- DILMÉ, F. (2024): “Sequentially stable outcomes,” *Econometrica*, 92, 1097–1134.
- (2025): “Equilibrium destabilization,” Mimeo.
- GOVINDAN, S. AND R. WILSON (2009): “On forward induction,” *Econometrica*, 77, 1–28.
- HARSANYI, J. C. (1973): “Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points,” *International Journal of Game Theory*, 2, 1–23.
- KREPS, D. M. AND R. WILSON (1982): “Sequential equilibria,” *Econometrica*, 50, 863–894.
- KUHN, H. W. (1953): “Extensive games and the problem of information,” *Contributions to the Theory of Games*, 2, 193–216.
- MAILATH, G. J., L. SAMUELSON, AND J. M. SWINKELS (1993): “Extensive form reasoning in normal form games,” *Econometrica*, 61, 273–302.
- RENY, P. J. (1992): “Backward induction, normal form perfection and explicable equilibria,” *Econometrica*, 627–649.

SELTEN, R. (1975): "Reexamination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory*, 4, 25–55.

VAN DAMME, E. (1984): "A relation between perfect equilibria in extensive form games and proper equilibria in normal form games," *International Journal of Game Theory*, 13, 1–13.

——— (1991): *Stability and perfection of Nash equilibria*, vol. 339, Springer.