

# Lecture 19

## Sequential irrationality and Nash equilibrium refinements

Instructor: Gabriele Farina\*

A major focus so far has been around the computation of Nash equilibrium strategies in two-player imperfect-information games. As we discussed on multiple occasions, Nash equilibrium strategies encode the idea of playing optimally against the strongest possible opponent. Even when the opponent is only close to optimal (for example, in the poker competitions where the opponent were top professional poker players), playing a Nash equilibrium is often the safe choice, as professional players are very quick at exploiting suboptimal strategies, making opponent modeling risky.

However, the dichotomy between playing a Nash equilibrium and doing risky opponent modeling is not quite accurate. In the previous class, we saw that it is possible to start from playing a Nash equilibrium strategy. Then, as the game progresses, we can quantify the amount of mistakes (“gifts”) that were made by the opponent, and use those as capital to spend to do safe opponent exploitation.

In this lecture, we talk about a different angle of the problem: what initial Nash equilibrium strategy to pick to start with, so as to maximize the possibility of capitalizing on opponents’ mistakes.

### 1 Sequential irrationality

Not all Nash equilibria are equally sensible. Intuitively, the problem lies in the fact that Nash equilibrium strategies are only optimized for the strongest possible opponent. Because of that, they are completely indifferent to what happens in parts of the game tree that are reached only if a player makes a mistake.

To make the discussion more concrete, consider the *Guess-the-Ace* game, introduced by [Miltersen and Sørensen \[2006\]](#).

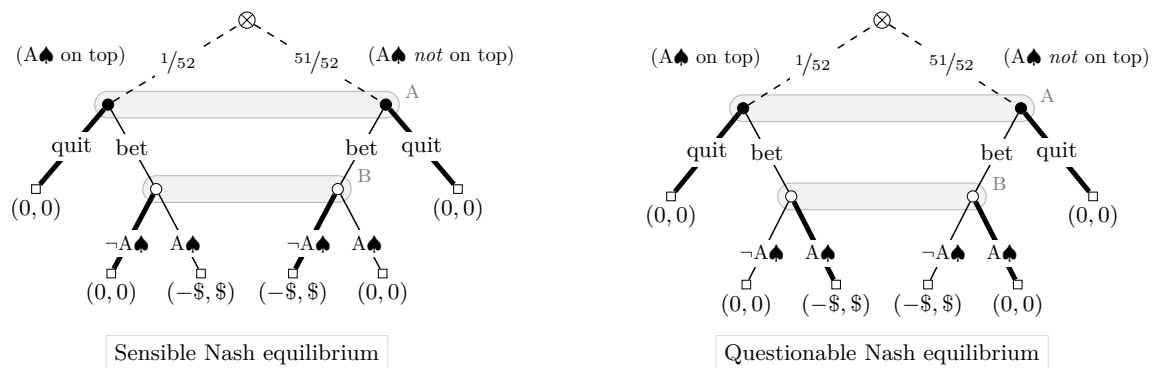


Figure 1: Guess-the-Ace game. On the left, a sequentially-rational Nash equilibrium is highlighted. On the right, a non sequentially-rational equilibrium is highlighted. The first payoff of each outcome is assigned to the black player, the second to the white player. The payoff ‘\$’ refers to some amount of money, e.g., \$1000.

\*Computer Science Department, Carnegie Mellon University. ✉ [gfarina@cs.cmu.edu](mailto:gfarina@cs.cmu.edu).

In Guess-the-Ace, at the start a standard 52-card deck is perfectly shuffled, face down, by a dealer. Then, Player 1 can decide whether to immediately end the game, at which point no money is transferred between the players, or offer \$1000 to Player 2 if they can correctly guess whether the top card of the shuffled deck is the ace of spaces or not. If Player 2 guesses correctly, the \$1000 get transferred from Player 1 to Player 2; if not, no money is transferred. The game tree is summarized in Figure 1.

Clearly, the only Nash equilibrium strategy for Player 1 is to quit immediately, or they are guaranteed to *lose* money. Since Player 2 does not get to play, any strategy for Player 2 is a Nash equilibrium strategy.

In particular, both highlighted equilibria in Figure 1 are Nash equilibria. However, the two equilibria are significantly different from a practical point of view. Imagine that Player 2 is a bot playing against opponents in the real world, blindly following the Nash equilibrium strategy it has precomputed. If Player 1 makes a mistake and decides to offer the \$1000 instead of immediately quitting, the Nash equilibrium that bets that the top card is *not* the ace of space has an expected utility of  $> \$980$  whereas the Nash equilibrium that bets that the top card is the ace of spades only has an expected utility of  $< \$20$ .

So, while both strategy profiles in Figure 1 are Nash equilibria, only one of the two is “*sensible*”.


Formalizing this subtle notion of rationality within the set of Nash equilibria has been a major endeavor for the game-theoretic literature in the 70s and 80s. Today, we say that the equilibrium in Figure 1 (Left) is *sequentially irrational*, while the one on the right is *sequentially rational*. The takeaway lesson is the following.

**Observation 1.1.** Not all Nash equilibria are equally “good” when the agents can make mistakes. Specifically, sequentially-irrational Nash equilibria might leave value on the table, by being incapable of capitalizing on opponents’ mistakes.

The goal of this lecture is to investigate how one can rule out sequential irrationality and compute a sequentially-rational Nash equilibrium in a two-player zero-sum imperfect-information game.

## 2 Undomination is not the solution

One might believe that the problem of sequential irrationality is that of picking dominated strategies. So, one might be inclined to look into the problem of finding a Nash equilibrium whose support does not include any (weakly) dominated strategy (the concept is not immediately well defined, but for the purposes of this discussion let’s restrict ourselves to Nash equilibria in deterministic strategies).



Questionable *undominated* Nash equilibrium

Figure 2: Modified *Guess-the-Ace* games. The highlighted equilibrium is undominated, but still questionable.

Unfortunately, domination of strategies is not the root cause of sequential irrationality, and therefore undomination is not its solution. Indeed, as much as undomination *does* get rid of the undesirable behavior

of Figure 1 (Right), since action ‘A♠’ is strictly dominated by action ‘-A♠’, it does not prevent sequential irrationality in more complex settings, such as Figure 2.

In Figure 2, again due to [Miltersen and Sørensen \[2006\]](#), the Guess-the-Ace game is slightly modified in that, when Player 2 guesses wrong, Player 1 can decide whether they still want to give \$1000 to Player 2 out of the kindness of their heart or not. By introducing that possibility, action ‘-A♠’ is not strictly dominating anymore, because Player 2 might still hope that the second gift of \$1000 is given only when the insensible guess ‘A♠’ is made. So, the second takeaway lesson for today’s class is the following.

**Observation 2.1.** Undomination does not prevent a player from playing risky actions, “hoping” for an opponent’s mistake.

### 3 Trembling-hand refinements

The issue of sequential irrationality stems from the fact that some parts of the game tree are unreachable at equilibrium. For those excluded parts of the game tree, any strategy can be picked without affecting the equilibrium. The idea behind trembling-hand refinements is simple: to avoid sequential irrationality, it forces all players to explore the whole game tree. It does so by forcing the players to *tremble*, that is, by constraining them to play all actions at all decision points with a strictly positive lower bound probability that grows as a function of a hyperparameter  $\epsilon > 0$ . For each  $\epsilon > 0$ , a Nash equilibrium subject to the trembling constraints is found. A trembling-hand refinements is then any limit points of such Nash equilibria as  $\epsilon \rightarrow 0^+$ .

Different equilibrium notions differ as to how the lower bounds are set as a function of  $\epsilon$ . We will see two, which are the two best known: extensive-form perfect equilibrium and quasi-perfect equilibrium.

#### 3.1 Extensive-form perfect equilibrium (EFPE)

*Extensive-form perfect equilibrium (EFPE)*, due to [Selten \[1975\]](#), is conceptually the simplest of the two. In an EFPE, the trembles are *behavioral*: given  $\epsilon > 0$ , the perturbed game simply mandates that every action at every decision point must be picked with probability at least  $\epsilon$ .

Since our game solving formalism is based around the sequence-form representation of strategies, it is important to check that those behavioral trembling constraints can be expressed in the sequence form. That is the case: asking that action  $a$  at decision point  $j$  of Player 1 be selected with probability at least  $\epsilon$  corresponds to the sequence-form constraint

$$\mathbf{x}[ja] \geq \begin{cases} \epsilon & \text{if } p_j = \emptyset \\ \epsilon \cdot \mathbf{x}[p_j] & \text{otherwise.} \end{cases} \quad (1)$$

Collecting all sequence-form trembling constraints (1) across all decision points  $j \in \mathcal{J}$  and actions  $a \in A_j$  of Player 1, we can express the whole set of trembling constraints in matrix form as  $\mathbf{M}_1(\epsilon) \mathbf{x} \geq \mathbf{m}_1(\epsilon)$ . (An analogous statement holds for Player 2). So, given any  $\epsilon > 0$ , a Nash equilibrium strategy for Player 1 under the trembling constraints can be expressed as the saddle point problem

$$\left\{ \begin{array}{l} \max_{\mathbf{x}} \left\{ \begin{array}{l} \min_{\mathbf{y}} \quad \mathbf{x}^\top \mathbf{U} \mathbf{y} \\ \text{s.t.} \quad \textcircled{1} \quad \mathbf{F}_2 \mathbf{y} = \mathbf{f}_2 \\ \quad \quad \textcircled{2} \quad \mathbf{M}_2(\epsilon) \mathbf{y} \geq \mathbf{m}_2(\epsilon) \end{array} \right. \\ \text{s.t.} \quad \textcircled{3} \quad \mathbf{F}_1 \mathbf{x} = \mathbf{f}_2 \\ \quad \quad \textcircled{4} \quad \mathbf{M}_1(\epsilon) \mathbf{x} \geq \mathbf{m}_1(\epsilon). \end{array} \right. \quad (\text{EFPE})$$

We will look into how to compute a limit point of solutions to (EFPE) as  $\epsilon \rightarrow 0^+$  in Section 4.

### 3.2 Quasi-perfect equilibrium (QPE)

*Quasi-perfected equilibrium (QPE)*, introduced by [van Damme \[1984\]](#), is a bit more intricate than EFPE. Specifically, while in an EFPE each trembling constraints mandates a lower bound of  $\epsilon$  on the probability of playing each *action*, in the case of a QPE the lower bounds are given on the probability of each *sequence* of actions. More precisely, for any  $\epsilon > 0$  and for each player  $i \in \{1, 2\}$ , let  $\ell_i : \epsilon \rightarrow \mathbb{R}_{>0}^{|\Sigma_i|}$  denote the vector parametrized on  $\epsilon$  and indexed on the sequences  $\Sigma_i$  of Player  $i$ , whose entries are defined as

$$\ell_i(\epsilon)[\sigma] = \epsilon^{|\sigma|} \quad \forall \sigma \in \Sigma_i, \quad (2)$$

where  $|\sigma|$  denotes the number of actions for Player  $i$  in the sequence  $\sigma$ . [Miltersen and Sørensen \[2010\]](#) proved that any limit point of the solution to the perturbed optimization problem

$$\begin{cases} \max_{\mathbf{x}} & \begin{cases} \min_{\mathbf{y}} & \mathbf{x}^\top \mathbf{U} \mathbf{y} \\ \text{s.t.} & \textcircled{1} \mathbf{F}_2 \mathbf{y} = \mathbf{f}_2 \\ & \textcircled{2} \mathbf{y} \geq \ell_2(\epsilon) \end{cases} \\ \text{s.t.} & \begin{cases} \textcircled{3} \mathbf{F}_1 \mathbf{x} = \mathbf{f}_1 \\ \textcircled{4} \mathbf{x} \geq \ell_1(\epsilon). \end{cases} \end{cases} \quad (\text{QPE})$$

is a QPE. (Recently, [Gatti et al. \[2020\]](#) took this construction further, and showed that *any* QPE can be expressed as a limit point of solutions to (QPE), as long as more general vectors of polynomials  $\ell_1, \ell_2$  are used than in (2).) In this paper we will focus on Miltersen-Sørensen-style perturbation as defined in (2).

Once again, we will discuss how to compute a limit point of solutions to (QPE) as  $\epsilon \rightarrow 0^+$  in Section 4.

### 3.3 Relationship between the equilibria

We already know from Section 2 that undomination does not imply sequential rationality. Interestingly, the converse also is not true in general. So, undomination and sequential rationality are actually incomparable concepts, in the sense that neither implies the other.

At this point, one might naturally wonder whether a refinement that is both undominated and sequentially-rational can be devised. The answer is yes: a nice property of QPE is that not only it is sequentially rational, but it is also undominated! The same cannot be said of EFPE. So, as [Mertens \[1995\]](#) noted, a *quasi-perfect equilibrium* is nowadays considered superior to EFCE.

*Observe that the “quasi-perfect” equilibria [...] are still sequential—and sequential equilibria have all backward-induction properties (e.g., Kohlberg and Mertens, 1986)—but are at the same time normal form perfect—which can be viewed as the strong version of undominated. (And every proper equilibrium is quasi-perfect.) Thus, by some irony of terminology, the “quasi”-concept seems in fact far superior to the original unqualified perfection itself.*

The relationship between the different refinements is summarized in the Venn diagram of Figure 3.

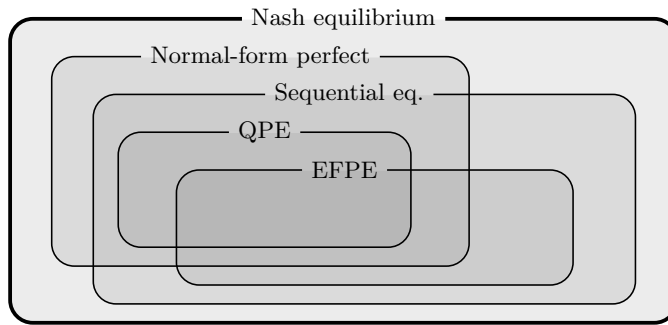


Figure 3: Relationship between the different Nash equilibrium refinements.

### 3.4 Computational complexity

Perhaps surprisingly, finding an EFPE or a QPE in a two-player game is not harder than finding a Nash equilibrium. In particular, in zero-sum games, an EFPE and a QPE can be found in polynomial time in the size of the input game. Table 1 summarizes the computational complexity of computing the Nash equilibrium refinements mentioned so far in two-player games.

Solution concept	General-sum	Zero-sum
Nash (NE)	PPAD-complete [Daskalakis et al., 2009]	FP [Romanovskii, 1962] [von Stengel, 1996]
Quasi Perfect (QPE)	PPAD-complete [Miltersen and Sørensen, 2010]	FP [Miltersen and Sørensen, 2010]
Extensive-Form Perfect (EFPE)	PPAD-complete [Farina and Gatti, 2017]	FP [Farina and Gatti, 2017]

Table 1: Complexity of computing different Nash equilibrium refinements in two-player games.

## 4 Trembling linear programs and computation of QPE and EFPE

We can compute a limit point of solutions to (EFPE) and (QPE) using the same machinery. As a first step, just like what we did for the Nash equilibrium, we convert the bilinear saddle-point formulations (EFPE), (QPE) into linear programs by dualizing the internal minimization problems. This gives us a linear program where the constraints matrix and the objective function depend polynomially on  $\epsilon$ . In particular, for both QPE and EFPE we end up with a linear program of the form

$$P(\epsilon) : \begin{cases} \max & \mathbf{c}(\epsilon)^\top \mathbf{x} \\ \text{s.t.} & \mathbf{A}(\epsilon) \mathbf{x} = \mathbf{b}(\epsilon) \\ & \mathbf{x} \geq \mathbf{0}, \end{cases}$$

where  $\mathbf{c}$ ,  $\mathbf{A}$  and  $\mathbf{b}$  are *polynomial* functions of  $\epsilon$  with rational coefficients. We will call an object of that form a *trembling linear program (TLP)*, and a limit point of solutions to  $P(\epsilon)$  as  $\epsilon \rightarrow 0^+$  a *limit solution* of the TLP.

With this formalism, we can reframe the computation of an EFPE or a QPE as the problem of finding a limit solution to their corresponding TLPs.

We will now discuss the complexity of solving a TLP, and two different computational approaches. Both of them are based on the concept of *basis stability* (Recall that a *basis* of an LP is a subset of the program's variables such that when only those columns of matrix  $\mathbf{A}$  that correspond to those variables are included in a new matrix  $\mathbf{A}'$ , the new matrix  $\mathbf{A}'$  is invertible [Bertsimas and Tsitsiklis, 1997, page 55].)

**Definition 4.1** (Stable basis). Let  $P(\epsilon)$  be a TLP. The LP basis  $\mathcal{B}$  is said to be *stable* if there exists  $\bar{\epsilon} > 0$  such that  $\mathcal{B}$  is optimal for  $P(\epsilon)$  for all  $\epsilon : 0 < \epsilon \leq \bar{\epsilon}$ .

If a stable basis were to be found, from there a limit solution of  $P(\epsilon)$  could be computed in polynomial time. As it turns out, a stable basis always exists, and can be computed in polynomial time.

### 4.1 Negligible Positive Perturbations (NPP)

Farina et al. [2018], extending prior work by Miltersen and Sørensen [2010] and Farina and Gatti [2017], showed the following.

**Theorem 4.1** (Farina et al. [2018]). Given as input a TLP  $P(\epsilon)$ , there exists  $\epsilon^* > 0$ —called a *negligible positive perturbation (NPP)*—such that for all  $0 < \bar{\epsilon} \leq \epsilon^*$ , any optimal basis for the numerical LP  $P(\bar{\epsilon})$  is stable. Furthermore, such a value  $\epsilon^*$  can be computed in polynomial time in the input size, assuming that a polynomial of degree  $d$  requires  $\Omega(d)$  space in the input.<sup>a</sup>

<sup>a</sup>If this were not the case, evaluating a polynomial in an integer  $n$  would not be an efficient operation, since it requires  $\Omega(d \log n)$  bits to represent the output.

So, at least in principle, a solution to a TLP  $P(\epsilon)$  could be computed as follows:

- First, compute the value of the NPP  $\epsilon^*$  using the constructive proof of Theorem 4.1.
- Then, solve the numerical linear program  $P(\epsilon^*)$  to optimality. Since the bit complexity of  $\epsilon^*$  is polynomial in the size of the TLP, the numerical LP can be solved to optimality in polynomial time, and a basis  $\mathcal{B}$  can be extracted. From Theorem 4.1, such a basis is stable (Definition 4.1)
- Finally, extract the limit solution to the TLP from the stable basis.

The algorithm just described has polynomial complexity in the TLP size. In the case of the TLP arising from QPE and EFPE, that translates into a polynomial-time algorithm to find an exact EFPE and QPE in a two-player zero-sum game (see also Table 1).

## 4.2 A significantly more scalable approach

While technically polynomial, the NPP-based algorithm described in the previous subsection is mostly of conceptual interest. In practice, because the value of the NPP is so small, any linear programming solver that wants to have a chance at solving the numerical linear program  $P(\epsilon^*)$  must—as a minimum—use rational arithmetic, rendering the algorithm extremely slow.

A significantly more scalable algorithm for solving TLPs, due to Farina et al. [2018], avoids the pessimistically small numerical NPP  $\epsilon^*$  of Theorem 4.1 by using an efficient stability-checking oracle for checking if a basis is stable or not. The workflow is summarized in Figure 4.




Figure 4: Overview of the practical algorithm for solving TLPs.

The iterative algorithm repeatedly picks a numerical perturbation  $\bar{\epsilon}$ , computes an optimal basis for the perturbed LP  $P(\bar{\epsilon})$ , and queries the basis-stability oracle. If the basis is not stable, the algorithm concludes that the perturbation value  $\bar{\epsilon}$  was too optimistic, and a new iteration is performed with a smaller perturbation reduced by a multiplicative constant (for example, divide it by 1000). On the other hand, if the basis is stable, the algorithm takes the limit of the LP solution and returns it as the limit solution of the TLP. Correctness and termination are guaranteed by the following observation.

**Observation 4.1.** Any value of  $\bar{\epsilon}$  in the range  $(0, \epsilon^*]$  guarantees termination of the algorithm. Indeed, by Theorem 4.1, any optimal basis for  $P(\bar{\epsilon})$  is stable and makes our iterative algorithm terminate. Furthermore, if after every negative stability test the value of  $\bar{\epsilon}$  is reduced by a constant multiplicative factor (e.g., halved), then since  $\epsilon^*$  only has a polynomial number of bits, the algorithm terminates after trying at most a polynomial number of different values for  $\bar{\epsilon}$ .


The practical algorithm is 3-4 orders of magnitude faster than the conceptual algorithm described in Section 4.1, and is the current state-of-the-art algorithm for computing QPE and EFPE.

## References

- Peter Bro Miltersen and Troels Bjerre Sørensen. Computing sequential equilibria for two-player games. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2006.
- Reinhard Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International journal of game theory*, 1975.
- Eric van Damme. A relation between perfect equilibria in extensive form games and proper equilibria in normal form games. *International Journal of Game Theory*, 1984.
- Peter Bro Miltersen and Troels Bjerre Sørensen. Computing a quasi-perfect equilibrium of a two-player game. *Economic Theory*, 42(1), 2010.
- Nicola Gatti, Mario Gilli, and Alberto Marchesi. A characterization of quasi-perfect equilibria. *Games and Economic Behavior*, 122:240–255, 2020.
- J-F Mertens. Two examples of strategic equilibrium. *Games and Economic Behavior*, 8(2):378–388, 1995.
- Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. *Commun. ACM*, 52(2):89–97, 2009.
- I. Romanovskii. Reduction of a game with complete memory to a matrix game. *Soviet Mathematics*, 3, 1962.
- Bernhard von Stengel. Efficient computation of behavior strategies. *Games and Economic Behavior*, 14(2): 220–246, 1996.
- Gabriele Farina and Nicola Gatti. Extensive-form perfect equilibrium computation in two-player games. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- D. Bertsimas and J.N. Tsitsiklis. Introduction to linear optimization. 1997.
- Gabriele Farina, Nicola Gatti, and Tuomas Sandholm. Practical exact algorithm for trembling-hand equilibrium refinements in games. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
- David M. Kreps and Robert Wilson. Sequential equilibria. *Econometrica*, 50(4):863–894, 1982.

## A Why not uniform lower bounds in QPE?

Not all vanishing perturbations  $\ell_1(\epsilon)$ ,  $\ell_2(\epsilon)$  in the QPE formulation (QPE) lead to a sequentially-rational equilibrium. For example, it is natural to wonder whether it is *really* necessary to consider lower bounds of the form  $\epsilon^{|\sigma|}$  instead of, for example, the uniform lower bound  $\epsilon$  for all sequences. After all, surely a uniform lower bound of  $\epsilon$  would still force the whole game to be explored, wouldn't it? While appealing on the surface, such a uniform lower bound might result in a solution that is not even subgame perfect, much less sequentially rational! In particular, consider the perfect-recall game in Figure 5.



Player	Action	Probability
Player 1 (black)	$a$	$1 - 4\epsilon$
Player 1 (black)	$b$	$4\epsilon$
Player 1 (black)	$c, d, p, q$	$1/2$
Player 2 (white)	$r$	$1 - \epsilon$
Player 2 (white)	$s$	$\epsilon$

Figure 5: Small perfect-information game that illustrates that uniform  $\epsilon$  lower bounds can induce irrational behavior. Black nodes belong to Player 1, the white node belongs to Player 2.

For any choice of  $\epsilon \in [0, 1/4]$ , we now argue that the only Nash equilibrium of the perturbed game assigns probability  $1 - \epsilon$  to action  $r$  of Player 2, and probability  $1/2$  to actions  $c$  and  $d$  of Player 1. Indeed, action  $a$  strictly dominates  $b$ , since all payoffs for the black player (Player 1) are strictly lower in the subtree rooted at  $b$ . Hence, the black player must minimize the probability mass put on the sequences that contain action  $b$ , compatibly with lower bounds. Because we are using uniform lower bounds  $\epsilon$  on the probability of each sequence, the black player will need to put at least probability  $\epsilon$  on the four sequences  $bc, bd, bp, bq$ . This can be achieved when  $c, d, p, q$  are each selected with probability  $1/2$  and action  $b$  with probability  $4\epsilon$ . From the point of view of the white player (Player 2), information set C guarantees an expected utility of  $-1 \cdot 1/2 + 2 \cdot 1/2 = 1/2$ , while information set D guarantees an expected utility of 0. So, it is rational for the white player to put as much probability mass as allowed by the lower bounds to action  $r$ . This is achieved when action  $r$  is selected with probability  $1 - \epsilon$ , and action  $s$  with probability  $\epsilon$ .

So, as  $\epsilon \rightarrow 0^+$ , any limit point sees Player 2 pick action  $r$  with probability 1 and Player 1 randomizing uniformly between actions  $c$  and  $d$ , despite action  $d$  being strictly dominated. Thus, both players will act irrationally (with Player 1 not even playing a best response in the subtree rooted at C) should Player 1 make the mistake of picking action  $b$  instead of  $a$  at the root A. The resulting equilibrium is not sequentially rational. (In fact, it's not even subgame perfect, which is even stronger [Kreps and Wilson, 1982].)

## B Definition of Quasi-Perfection

We give one of the multiple known equivalent definitions, presented for the special case of two-player games only. Several equivalent definitions that apply to more general games can be found in the original work by van Damme, as well as in the work by Miltersen and Sørensen [2010] and Gatti et al. [2020].

**Definition B.1** ( $j$ -local purification). Let  $i \in \{1, 2\}$  be a player,  $\mathbf{q} \in Q_i$  be a strategy for Player  $i$ , and let  $j \in \mathcal{J}_i$  be an information set. We say that a strategy  $\mathbf{q}' \in Q_i$  for Player  $i$  is a  $j$ -local purification of  $\mathbf{q}$  if  $\mathbf{q}'$  is deterministic at any information set  $j' \succeq j$ , and coincides with  $\mathbf{q}$  at any other information set. When  $\mathbf{q}'$  is an  $j$ -local purification of  $\mathbf{q}$ , we further say that

- $\mathbf{q}'$  is  $\epsilon$ -consistent with  $\mathbf{q}$  if, for all  $j' \succeq j$ ,  $\mathbf{q}'$  assigns probability 1 only to actions that have probability  $\geq \epsilon$  in  $\mathbf{q}$ ;
- $\mathbf{q}'$  is optimal against a given strategy of the opponent if no other  $j$ -local purification of  $\mathbf{q}$  achieves (strictly) higher expected utility against said strategy of the opponent.

**Definition B.2** ( $\epsilon$ -quasi-perfect best response). A strategy  $\mathbf{q}_i$  is an  $\epsilon$ -quasi-perfect best response to the opponent strategy  $\mathbf{q}_{-i}$  if (i)  $\mathbf{q}_i$  assigns strictly positive probability to all actions of Player  $i$ ; and (ii) for all information sets  $j \in \mathcal{J}_i$  of Player  $i$ , every  $\epsilon$ -consistent  $j$ -local purifications of  $\mathbf{q}_i$  (Definition B.1) is



optimal for  $\mathbf{q}_{-i}$ . A strategy profile  $(\mathbf{q}_1, \mathbf{q}_2)$  where each strategy is an  $\epsilon$ -quasi-perfect best response to the opponent's strategy is called an  $\epsilon$ -quasi-perfect strategy profile.

**Definition B.3** (Quasi-perfect equilibrium). A *quasi-perfect equilibrium* is any limit point of  $\epsilon$ -quasi-perfect strategy profiles as  $\epsilon \rightarrow 0^+$ .