6.S890:
Topics in Multiagent Learning

Lecture 11 – Prof. Farina

Equilibrium Perfection

Fall 2024

Are all Nash equilibria equally good? Or should we aim for some but not other?

Not all Nash equilibria are equally sensible, especially in sequential games!

Intuition: Nash equilibria stem from the idea that the opponent is as strong as possible, and might therefore be completely unprepared to handle the case of an imperfect opponent

Very relevant when playing with humans!

Guess-the-Ace game

To make the discussion more concrete, consider the following game (due to Miltersen and Sorensen)

- At the start a standard 52-card deck is perfectly shuffled, face down, by a dealer
- Then, Player 1 decides whether to immediately end the game (no money transfer), 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

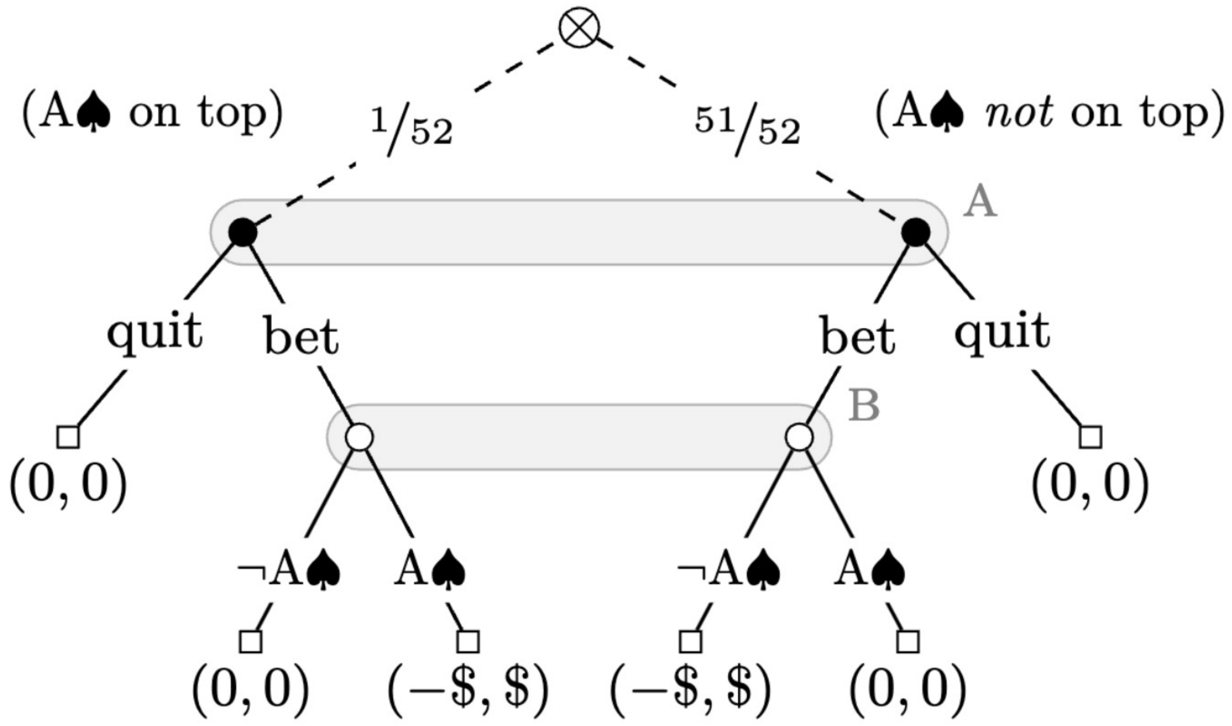

Guess-the-Ace game

To make the discussion more concrete, consider the following game (due to Miltersen and Sorensen)

- At the start a standard 52-card deck is perfectly shuffled, face down, by a dealer
- Then, Player 1 decides whether to immediately end the game (no money transfer), 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 … trans… mo…

Q: As Player 1, what is the only sensible way to play the game?

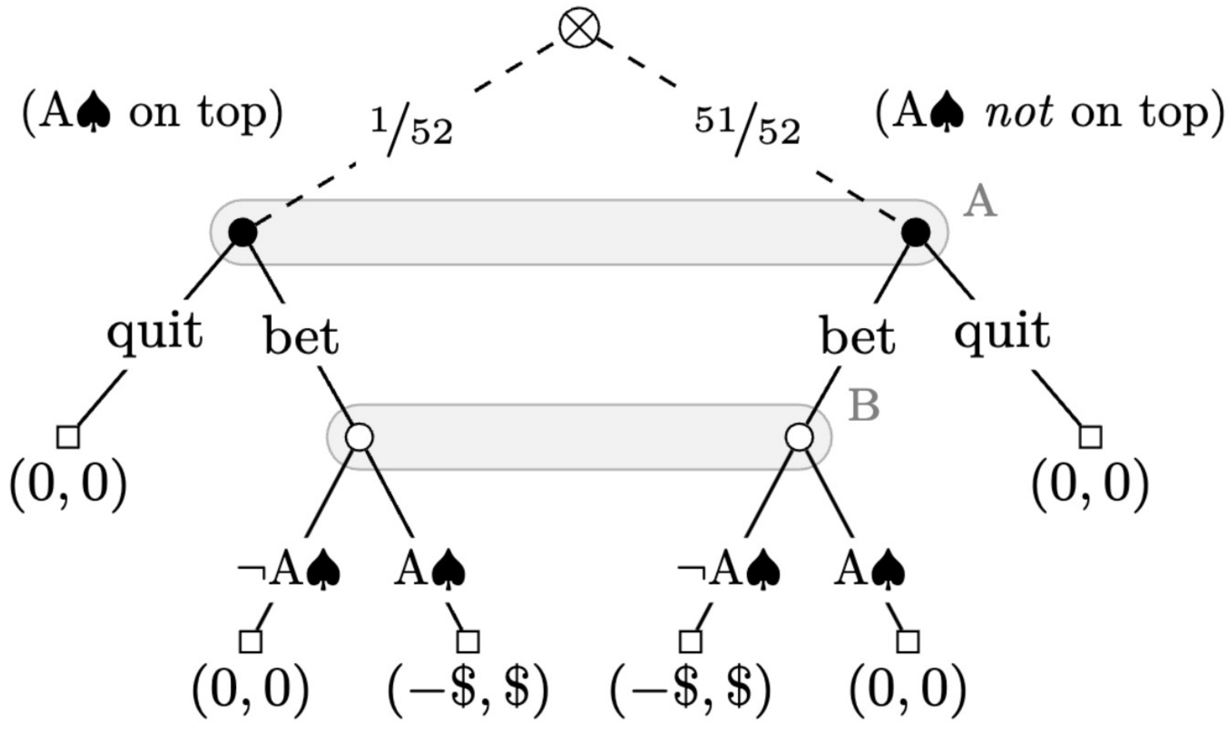

Answer: the only sensible thing for Player 1 to do is to quit immediately
(anything else loses money to Player 1 in expectation)

Indeed, that is the only Nash equilibrium strategy for Player 1.

Guess-the-Ace game

But then, Player 2 does not get to play. From the point of view of the definition of Nash equilibrium, anything that Player 2 does is a Nash equilibrium strategy

Yet, huge difference between the strategies. Only one of the two approaches can be called "rational"

Both of these are Nash equilibria. Nash eq. does not distinguish between the two

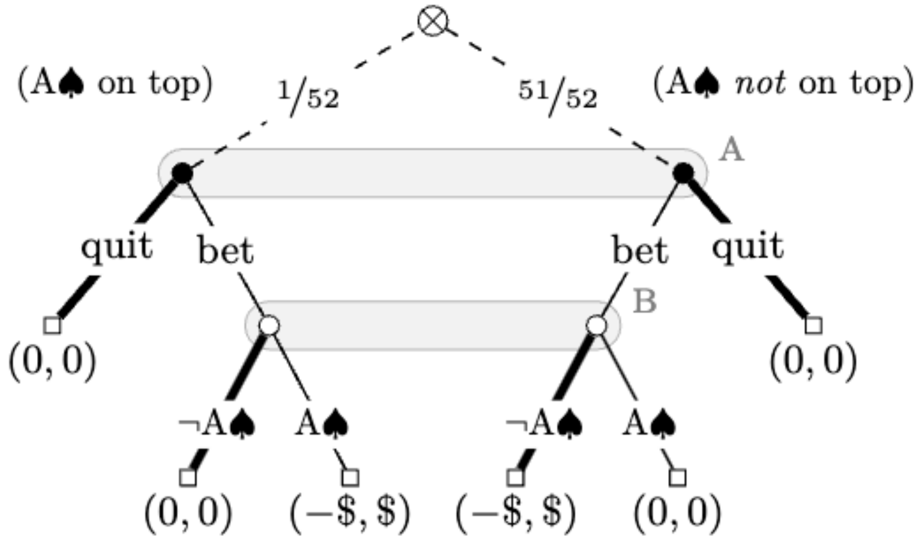

Sensible Nash equilibrium

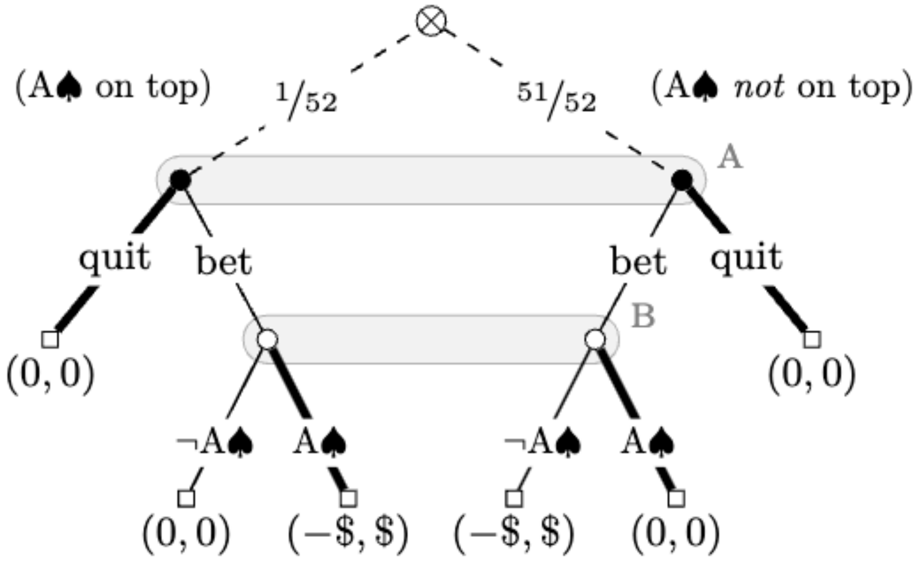

Questionable Nash equilibrium

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.

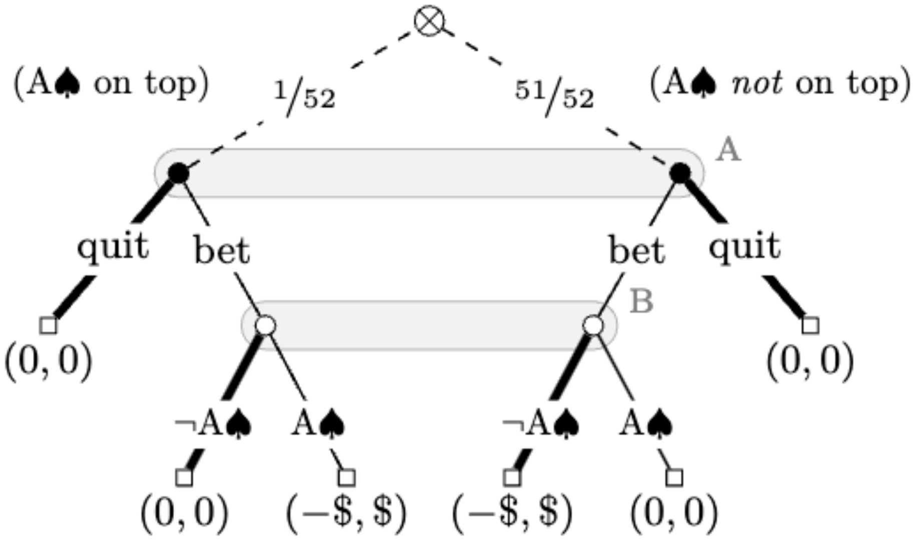

Sensible Nash equilibrium

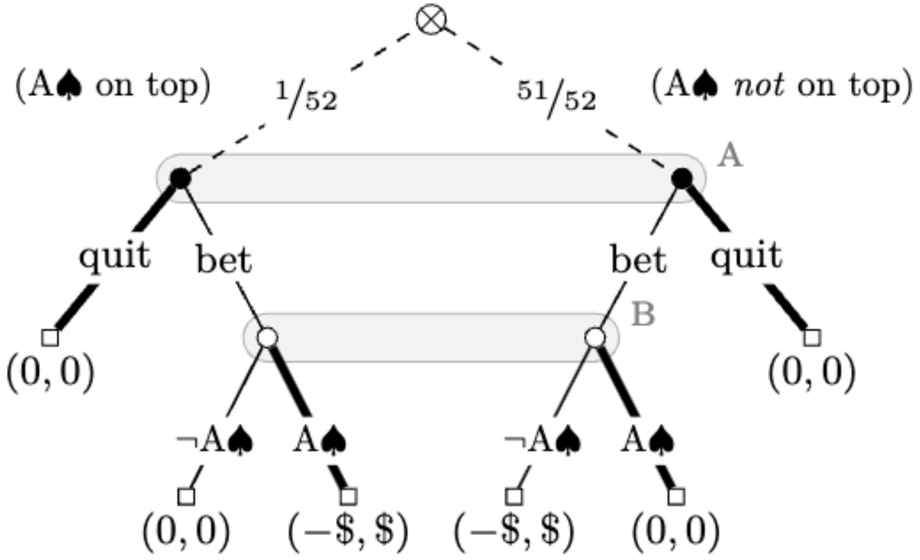

Questionable Nash equilibrium

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.

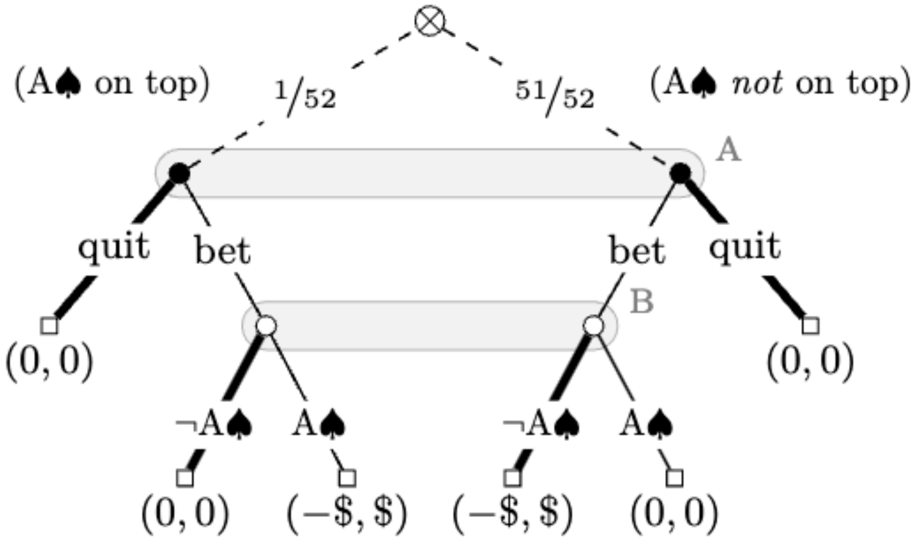

Sensible Nash equilibrium

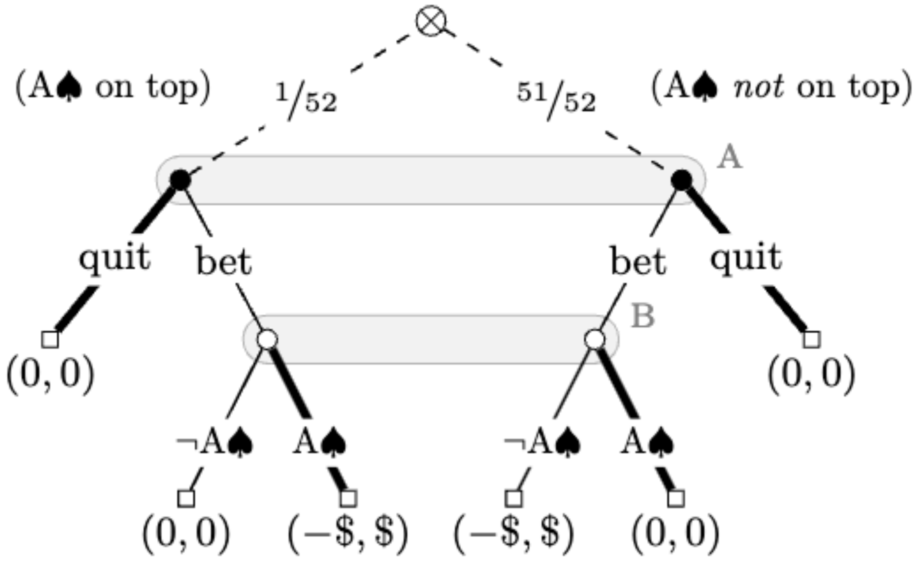

Questionable Nash equilibrium

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.

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

Trivia: this kind of surprising behavior kicked in during the poker tournament with the pros, and people were worried there was possibly a bug in the bot. Instead, it was likely the pro that had made a mistake and entered an off-equilibrium part of the tree

Undomination

One might believe that the problem of sequential irrationality is that of picking **dominated** strategies

While it is true that restricting to undominated strategies fixes the previous example, this is not a general fix!

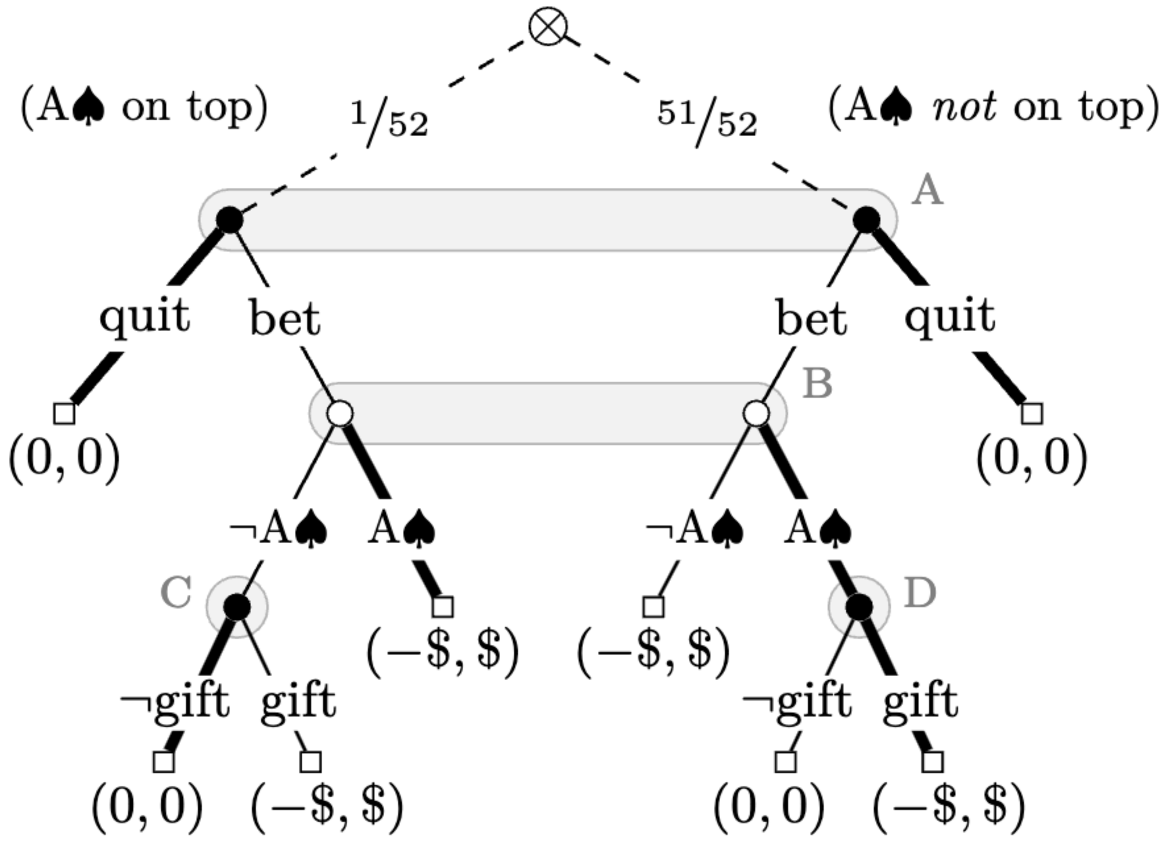

Questionable *undominated* Nash equilibrium

Undomination

One might believe that the problem of sequential irrationality is that of picking **dominated** strategies

While it is true that restricting to undominated strategies fixes the previous example, this is not a general fix!

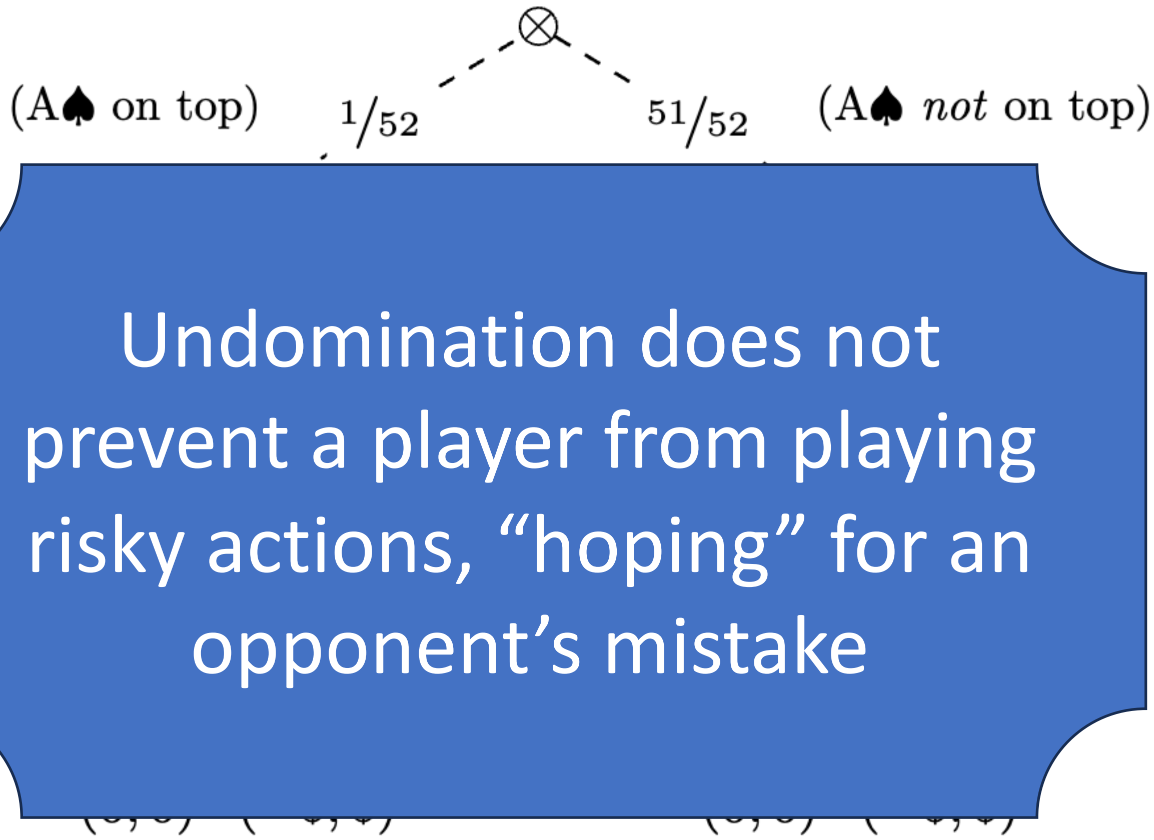

Undomination does not prevent a player from playing risky actions, “hoping” for an opponent’s mistake

Questionable *undominated* Nash equilibrium

Fundamentally, the issue of sequential irrationality stems from the fact that **some parts of the game tree are unreachable at equilibrium**.

IDEA: to avoid sequential irrationality, **force all players to explore the whole game** tree by imposing that they **must pick all actions with some positive, vanishing** probability

"Trembling-hand" equilibria

Two approaches

Trivia: "quasi-perfect" turns out to have stronger properties than "perfect" equilibrium. Naming is hard!

Extensive-Form Perfect Equilibrium

Force **every action in the game** to be played with probability at least $\epsilon > 0$

Take **any limit point** of Nash equilibria of the constrained games as $\epsilon \downarrow 0$

Quasi-Perfect equilibrium

For any d, force **every sequence of d actions from the root** to be played with probability at least ϵ^d for some $\epsilon > 0$

Take **any limit point** of Nash equilibria of the constrained games as $\epsilon \downarrow 0$

Relationship between the equilibria

We have already observed that **undomination does not imply sequential rationality**

Interestingly, the converse is also not true in general, so **sequential rationality and undomination are incomparable** (neither implies the other)

Natural question: is there an equilibrium that achieves both sequential rationality and undomination?

Is there an equilibrium that achieves both sequential rationality and undomination?

YES! A nice property of quasi-perfect equilibrium is that not only it is sequentially rational, but it is also undominated

For this reason, as Mertens noted in 1995, **quasi-perfect equilibrium is nowadays considered superior to extensive-form perfect equilibrium**

> *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.*

Venn diagram of equilibria

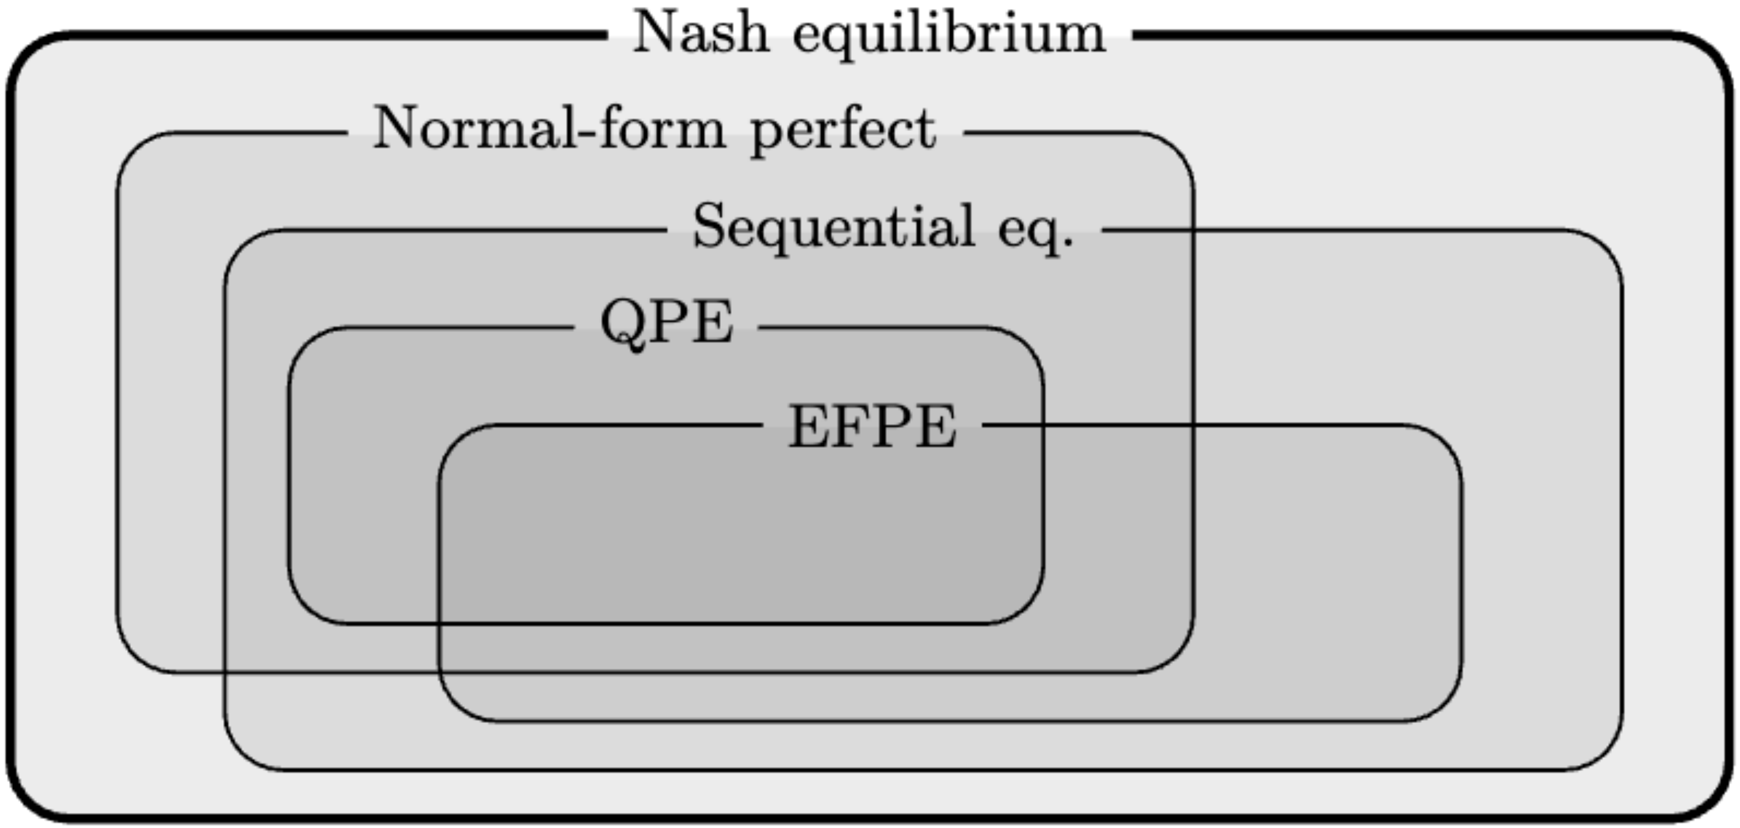

How hard is it compute a sequentially—rational equilibrium?

Cool fact: in theory, it's not harder than Nash!

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]

In practice, unfortunately we are still far behind, due to the added practical intricacy of forcing exploration and taking the limit as exploration vanishes