GAME REPRESENTATIONS
&
REFINEMENTS OF NASH EQUILIBRIUM

Game representations

Extensive form
(aka tree form)

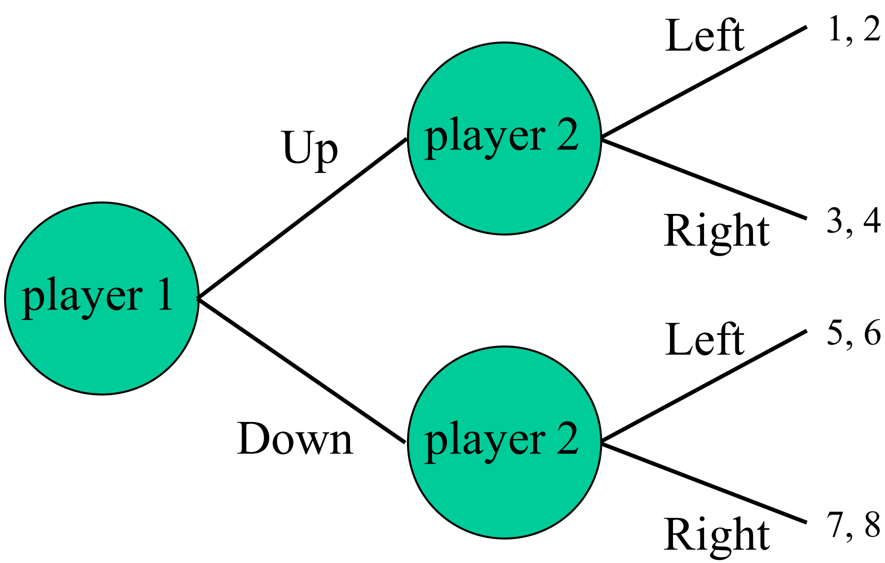

Normal form
aka strategic form
aka matrix form / bimatrix form

player 2's strategy

player 1's strategy	Left, Left	Left, Right	Right, Left	Right, Right
Up	1, 2	1, 2	3, 4	3, 4
Down	5, 6	7, 8	5, 6	7, 8

Potential combinatorial explosion →

Rock-scissors-paper game

Sequential moves

Rock-scissors-paper game

Simultaneous moves

Imperfect-information extensive-form games

Mixed strategy **= agent's chosen probability distribution over pure strategies from its strategy set**

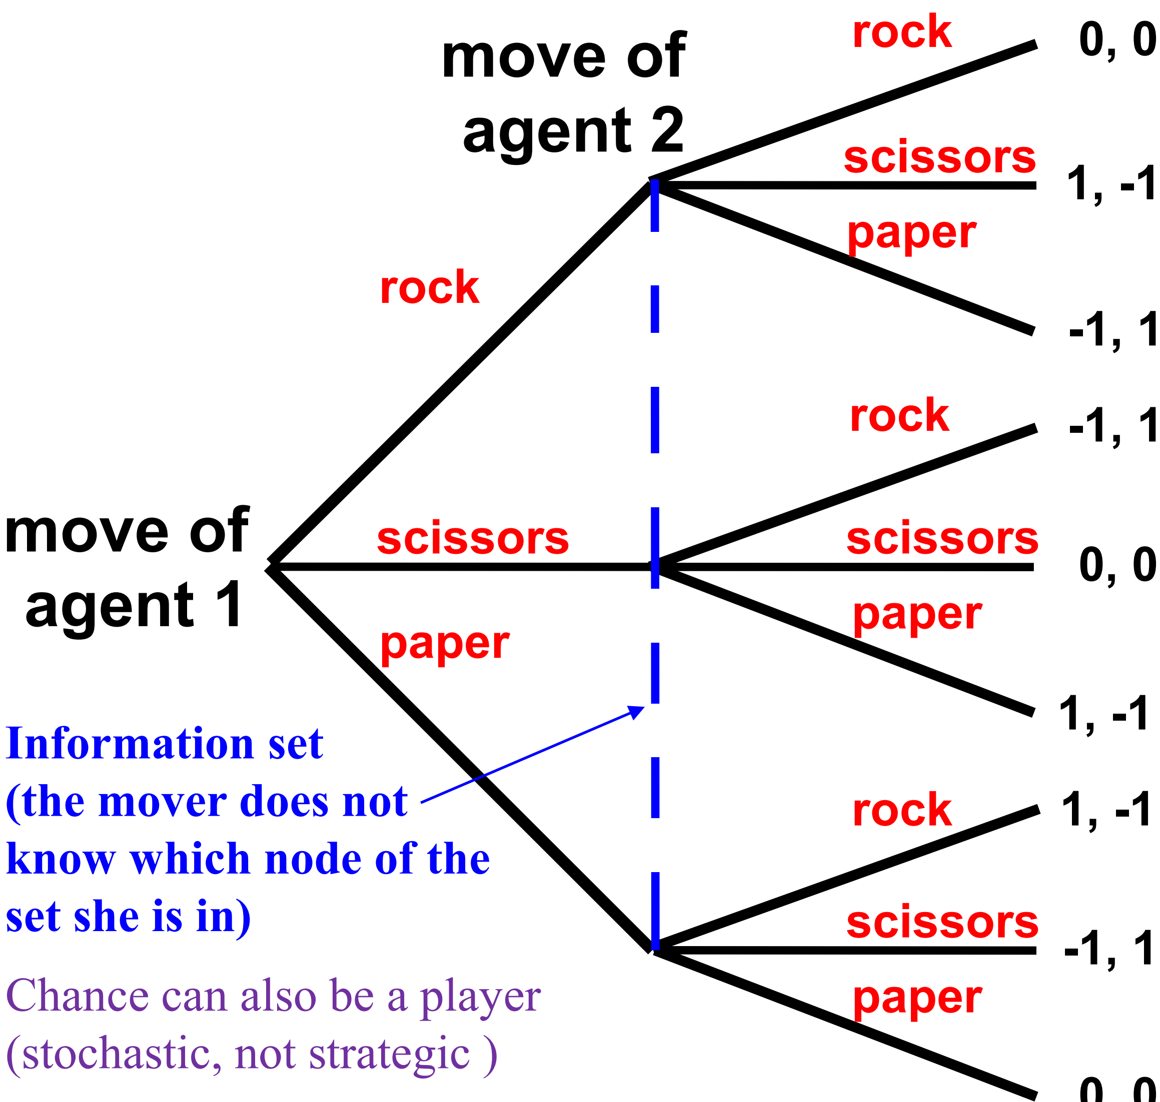

Information set (the mover does not know which node of the set she is in)

Chance can also be a player (stochastic, not strategic)

(Bayes-)Nash equilibrium: Each agent uses a best-response strategy and has consistent beliefs

Rock-paper-scissors game has a symmetric mixed-strategy Nash equilibrium where each player plays each pure strategy with probability 1/3

Fact: In mixed-strategy equilibrium, each strategy that occurs in the mix of agent i has equal expected utility to i

Behavioral strategy

- Agent has a probability distribution over her actions at each of her information sets
- **Kuhn's theorem:** If an agent has perfect recall, for every mixed strategy there is a behavioral strategy that has an equivalent payoff (i.e., the strategies are equivalent)
  – Applies also to infinite games

Existence of pure-strategy Nash equilibria

- **Thrm.**
  - **Any finite game,**
  - **where each action node is alone in its information set**
    - (i.e., at every point in the game, the agent whose turn it is to move knows what moves have been played so far)
  - **is dominance solvable by backward induction (at least as long as ties are ruled out)**
- **Constructive proof: Multi-player minimax search**
- **Lots of interesting work has been done on computer chess and Go to tackle the computational complexity**
  - **See Prof. Sandholm's lecture in an earlier course (pptx, video)**
  - **We won't cover that work in this course because most real-world games are imperfect-information games**

Existence of mixed-strategy Nash equilibria

- **Every finite player, finite strategy game has at least one Nash equilibrium if we admit mixed-strategy equilibria as well as pure [Nash 50]**
  - **(Proof is based on Kakutani's fix point theorem)**

REFINEMENTS OF NASH EQUILIBRIUM

Ultimatum game

(for distributional bargaining)

Subgame perfect equilibrium [Selten 72] & credible threats

- Proper subgame = subtree (of the game tree) whose root is alone in its information set
- Subgame perfect equilibrium = strategy profile that is in Nash equilibrium in every proper subgame (including the root), whether or not that subgame is reached along the equilibrium path of play
- E.g. Cuban missile crisis

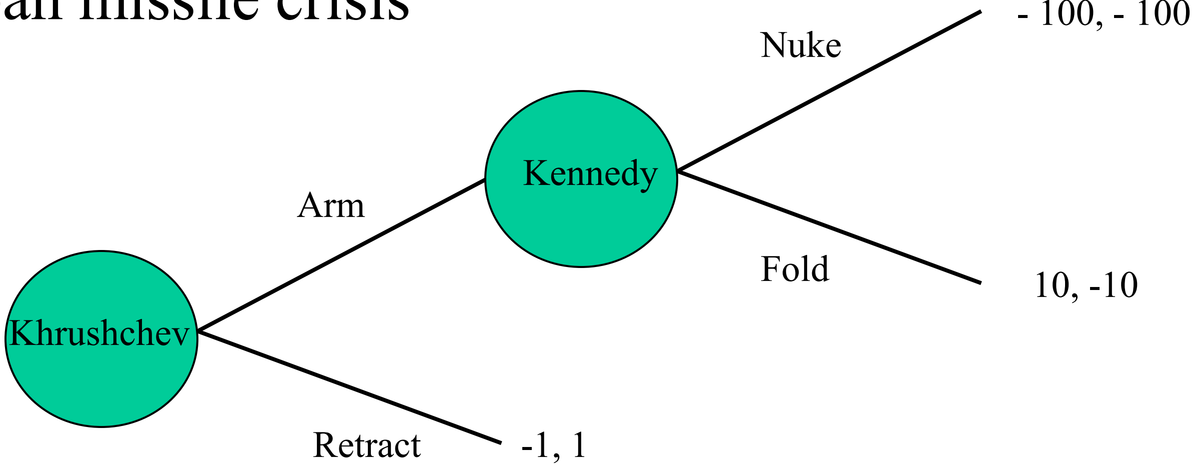

- Pure strategy Nash equilibria: (Arm,Fold), (Retract,Nuke)
- Pure strategy subgame perfect equilibria: (Arm,Fold)
- Conclusion: Kennedy's Nuke threat was not *credible*

Ultimatum game, again

Thoughts on credible threats

- Could use software as a commitment device
  - If one can credibly convince others that one cannot change one's software agent, then revealing the agent's code acts as a credible commitment to one's strategy
  - E.g. nuke in the missile crisis
  - E.g. accept no less than 60% as the second mover in the ultimatum game
- Restricting one's strategy set can increase one's utility
  - This cannot occur in single-agent settings
- Social welfare can increase or decrease

Solution concepts

Strength against collusion

Strong Nash eq [Auman 1959]

Coalition-Proof Nash eq [Bernheim, Peleg & Whinston 1987]

Nash eq
Bayes-Nash eq

Subgame perfect equilibrium

Perfect Bayesian equilibrium

Sequential equilibrium

Dominant strategy eq

Strength

Ex post equilibrium = Nash equilibrium for all priors

There are other equilibrium refinements too (see, e.g., following slides & wikipedia)

Example from the ***Brains vs AI***
Heads'Up No-Limit Texas Hold'em poker competition that I organized in April-May 2015

- *Claudico* made a bad move (not in the beginning of a hand)
- "How can that mistake be part of a GTO strategy?"

Guess-the-Ace game

[Miltersen & Sørensen, 2006]

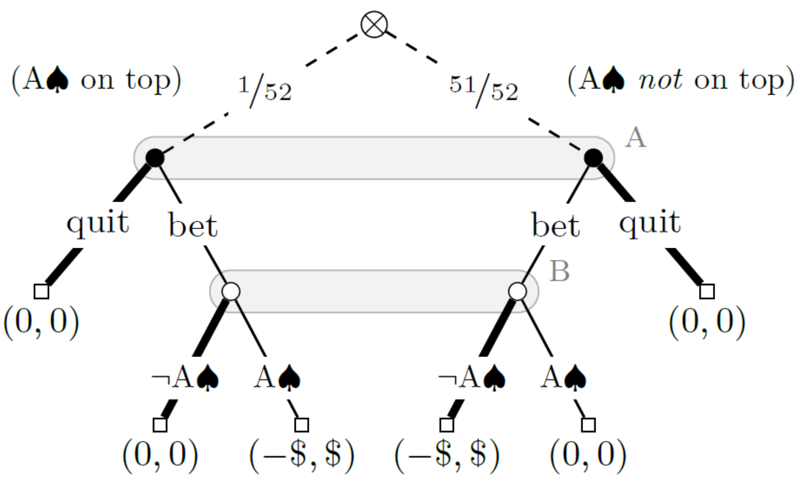

Sensible Nash equilibrium

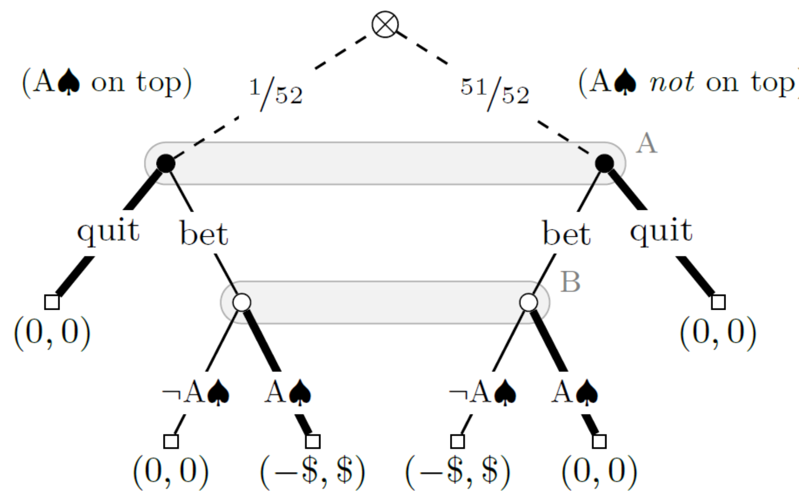

Questionable Nash equilibrium

Solution concepts for extensive-form imperfect-information games (slide 1 of 3)

More refined

- A player's *beliefs* consist of probability distributions over nodes occurring in her information sets
- A (Bayesian) Nash equilibrium is a strategy profile where each player maximizes her expected utility given the strategies played by the other players
- Perfect Bayesian equilibrium (PBE)
  - A belief system is *consistent* for a given strategy profile if the probability assigned by the system to every node is computed as the probability of that node being reached given the strategy profile, i.e., by Bayes' rule
  - A strategy profile is *sequentially rational* at a particular information set for a particular *belief system* if the expected utility of the player whose information set it is is maximal given the strategies played by the other players
    - A strategy profile is sequentially rational for a particular belief system if it satisfies the above for every information set
  - A PBE is a strategy profile and a belief system such that the strategies are sequentially rational given the belief system and the belief system is *consistent*, wherever possible, given the strategy profile
    - 'wherever possible' clause is necessary: some information sets might be reached with zero probability given the strategy profile; hence Bayes' rule cannot be employed to calculate the probability of nodes in those sets. Such information sets are said to be *off the equilibrium path* and any beliefs can be assigned to them
    - In Guess-the-Ace, the questionable Nash equilibrium is also a PBE, so PBE does not solve the issue
- Sequential equilibrium [Kreps and Wilson 82]. Refinement of PBE that specifies constraints on beliefs in such zero-probability information sets. Strategies and beliefs must be a limit point of a sequence of totally mixed strategy profiles and associated sensible (in PBE sense) beliefs.
  - In Guess-the-Ace, the questionable Nash equilibrium is not a sequential equilibrium, so that issue is solved

More detail about sequential equilibrium

(slide content from Yiling Chen's course)

Restrictions to Beliefs?

- We want beliefs for information sets that are off the equilibrium path to be reasonable. But what is reasonable?

Consider the NE (L, r) again. Player 2's information set will not be reached at the equilibrium, because player 1 will play L with probability 1. But assume that player 1 plays a completely mixed strategy, playing L, M, and R with probabilities $1-\epsilon$, $\frac{3\epsilon}{4}$, and $\frac{\epsilon}{4}$. Then, the belief on player 2's information set is well defined. Now, if $\epsilon \to 0$, it's still well defined.

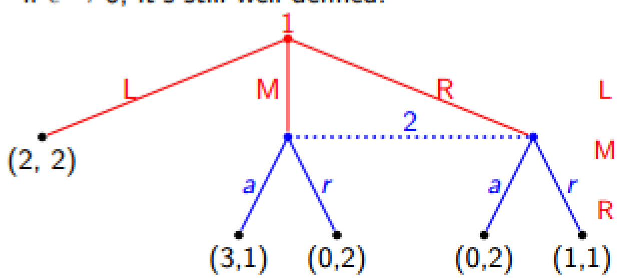

	a	r
L	2,2	2,2
M	3,1	0,2
R	0,2	1,1

The Nash equilibrium is also a sequential equilibrium in this case.

Consistent Assessment

- An assessment (s, μ) is consistent if there is a sequence $((s^n, \mu^n))_{n=1}^{\infty}$ of assessments that converges to (s, μ) and has the properties that each strategy profile s^n is completely mixed and that each belief system μ^n is derived from s^n using Bayes rule.

Sequential Equilibrium

- An assessment (s, μ) is a sequential equilibrium of a finite extensive-form game with perfect recall if it is sequentially rational and consistent.
- **Thm**: Every finite extensive-form game with perfect recall has a sequential equilibrium.
- A sequential equilibrium is a Nash equilibrium.
- With perfect information, a subgame perfect equilibrium is a sequential equilibrium.

Solution concepts for extensive-form imperfect-information games (slide 2 of 3)

More refined

- A player's *beliefs* consist of probability distributions over nodes occurring in her information sets
- A (Bayesian) Nash equilibrium is a strategy profile where each player maximizes her expected utility given the strategies played by the other players
- Perfect Bayesian equilibrium (PBE)
  - A belief system is *consistent* for a given strategy profile if the probability assigned by the system to every node is computed as the probability of that node being reached given the strategy profile, i.e., by Bayes' rule
  - A strategy profile is *sequentially rational* at a particular information set for a particular *belief system* if the expected utility of the player whose information set it is is maximal given the strategies played by the other players
    - A strategy profile is sequentially rational for a particular belief system if it satisfies the above for every information set
  - A PBE is a strategy profile and a belief system such that the strategies are sequentially rational given the belief system and the belief system is *consistent*, wherever possible, given the strategy profile
    - 'wherever possible' clause is necessary: some information sets might be reached with zero probability given the strategy profile; hence Bayes' rule cannot be employed to calculate the probability of nodes in those sets. Such information sets are said to be *off the equilibrium path* and any beliefs can be assigned to them
    - In Guess-the-Ace, the questionable Nash equilibrium is also a PBE, so PBE does not solve the issue
- Sequential equilibrium [Kreps and Wilson 82]. Refinement of PBE that specifies constraints on beliefs in such zero-probability information sets. Strategies and beliefs must be a limit point of a sequence of totally mixed strategy profiles and associated sensible (in PBE sense) beliefs.
  - In Guess-the-Ace, the questionable Nash equilibrium is not a sequential equilibrium, so that issue is solved
- Extensive-form trembling hand ***perfect*** equilibrium (EFPE) [Selten 75]. Require every move at every information set to be taken with non-zero probability. Take limit as tremble probability $\rightarrow 0$
- Extensive-form ***proper*** equilibrium [Myerson 78]. Idea: Costly trembles much less likely. At any information set, for any two actions A and B, if the mover's utility from B is less than from A, then prob(B) $\leq \varepsilon$ prob(A). Take limit as $\varepsilon \rightarrow 0$

Solution concepts for extensive-form imperfect-information games (slide 3 of 3)

More refined ↓

- Extensive-form perfect/proper equilibrium can involve playing weakly dominated strategies => argument for other solution concepts:
- *Normal-form* perfect equilibrium
  - Defined analogously to extensive-form trembling hand perfect equilibrium, but now for normal-form games
  - Normal- and extensive-form perfect equilibria are incomparable
  - A normal-form perfect equilibrium of an extensive-form game may or may not be sequential (and might not even be subgame perfect)
- Quasi-perfect equilibrium (QPE) [van Damme 84]
  - Original definition, informally: “A player takes observed as well as potential future mistakes of his opponents into account but assumes that he himself will not make a mistake in the future, even if he observes that he has done so in the past.”
  - Can be defined like EFPE, but the trembling constraints state that each sequence must have probability ≥ $\varepsilon^{\text{length(sequence)}}$ [Miltersen & Sørensen 2010; Gatti *et al.* 2020]
  - Incomparable to extensive-form perfect/proper
  - Admissible, unlike EFPE
- *Normal-form* proper equilibrium
  - Defined analogously to extensive-form proper equilibrium, but now for normal-form games
  - Always sequential
  - For 0-sum games, provides a strategy that maximizes the conditional utility (among minmax strategies), conditioned on the opponent making a mistake
    - Here a mistake is defined as a pure strategy that doesn’t achieve the value of the game against all minmax strategies

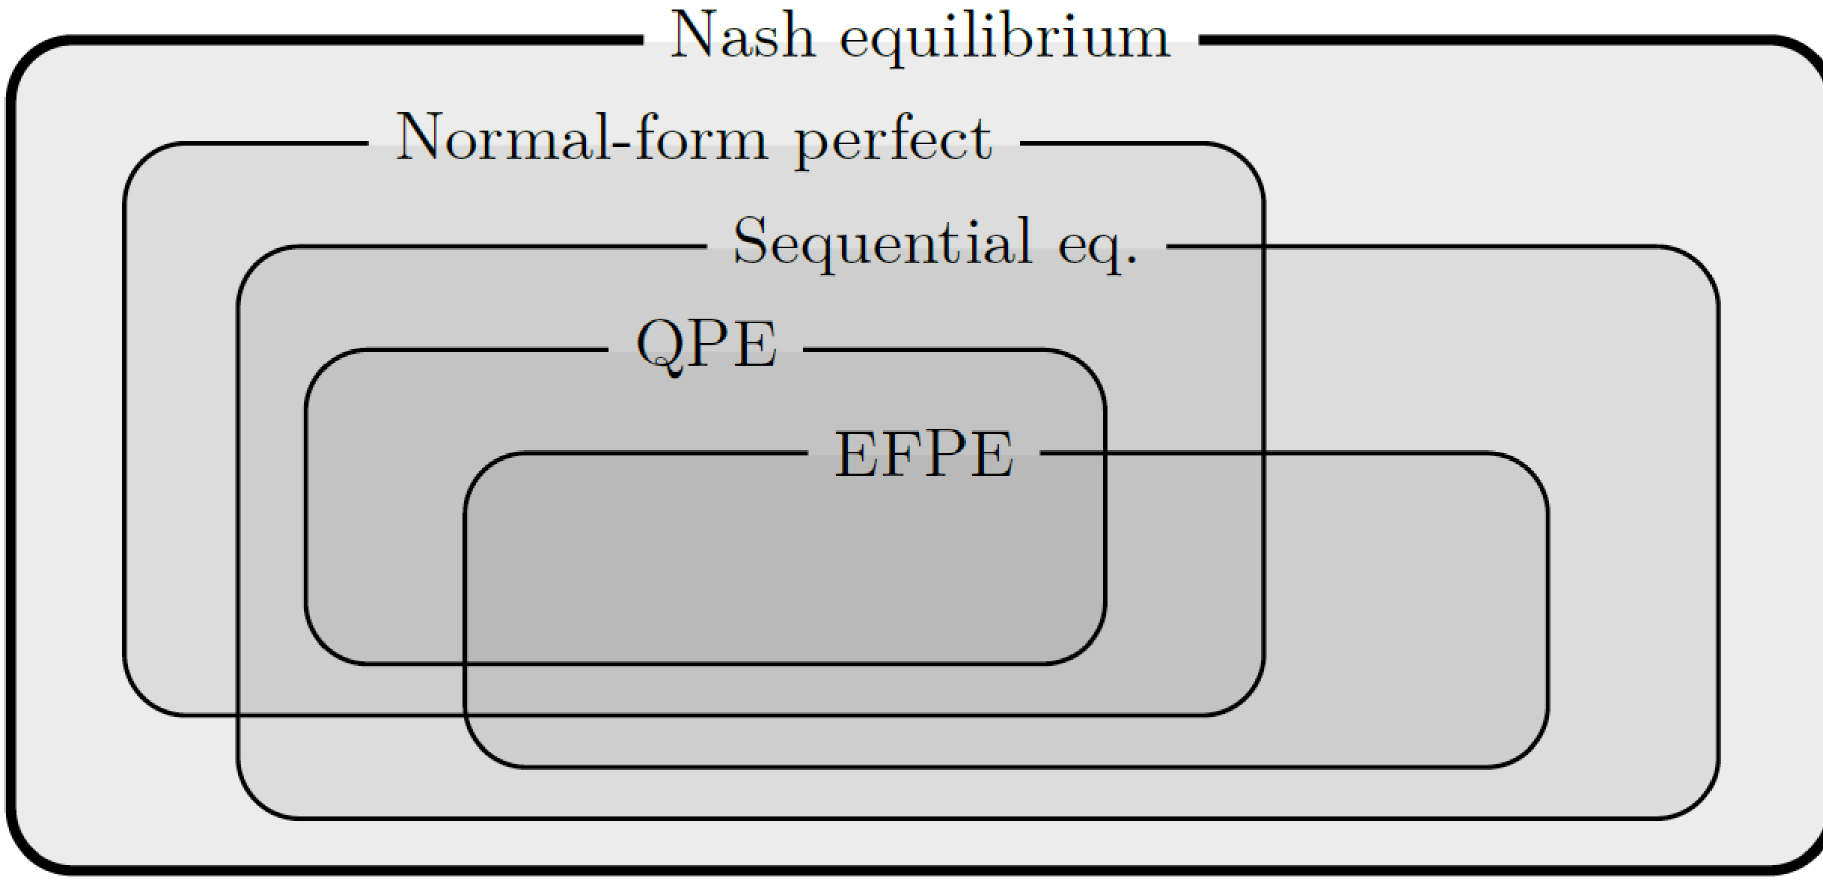
Nash equilibrium
Normal-form perfect
Sequential eq.
QPE
EFPE

Algorithms for equilibrium refinements in 2-player 0-sum extensive-form imperfect-information games

Sequential equilibrium [Kreps and Wilson, 1982]
- **Extensive-form trembling hand perfect equilibrium (EFPE)** [Selten, 1975]
- **Quasi-perfect equilibrium (QPE)** [van Damme, 1984]

- Even in 2-player 0-sum setting, these had been too complex to compute beyond small games
  - Best prior algorithm [Miltersen & Sørensen 2010] can solve games with 1,000 leaves [Farina, Gatti, and Sandholm, NeurIPS-18, draft-22]
- New algorithm finds exact EFPEs and QPEs for games with 100,000,000's of leaves on a laptop in a day
- EFPE-style and QPE-style trembling can be modeled as *trembling linear programs* for a given trembling magnitude $\varepsilon > 0$

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

- As $\varepsilon \to 0$, the LP optimum approaches an EFPE or QPE, respectively
- **Theorem** [Farina, Gatti & Sandholm, NeurIPS-18, draft-22]. ∃ finite game-specific $\varepsilon^* > 0$ s.t. for all $0 < \varepsilon \leq \varepsilon^*$, the optimal LP basis is stable
- **Algorithm** [Farina, Gatti & Sandholm, NeurIPS-18, draft-22]

  Initialize $\varepsilon > 0$

  Repeat

  Solve $P(\varepsilon)$

  if basis is stable

  Compute the limit of the optimal LP solution as $\varepsilon^* \to 0$ and return

  else $\varepsilon \leftarrow \varepsilon / 1000$
- Newest version runs LP sparsification as a preprocessor before running the above algorithm [Farina, Gatti & Sandholm, 2021]