Computational Game Theory

Slides prepared by Vincent Conitzer (Duke U)
Modified by Michal Feldman and Amos Fiat

What is game theory?

- Game theory studies settings where multiple parties (agents) each have
  - different preferences (utility functions),
  - different actions that they can take
- Each agent's utility (potentially) depends on all agents' actions
  - What is optimal for one agent depends on what other agents do
    - Very circular!
- Game theory studies how agents can rationally form beliefs over what other agents will do, and (hence) how agents should act
  - Useful for acting as well as predicting behavior of others

Where is game theory used?

- Economics (& business)
  - Auctions, exchanges, price/quantity setting by firms, bargaining, funding public goods, …
- Political science
  - Voting, candidate positioning, …
- Biology
  - Stable proportions of species, sexes, behaviors, …
- Computer science
  - Electronic marketplaces, networked systems, …
  - Computing the solutions that game theory prescribes

A brief history of game theory

- Some isolated early instances of what we would now call game-theoretic reasoning
  - e.g. Cournot 1838
- Von Neumann wrote a key paper in 1928
- 1944: *Theory of Games and Economic Behavior* by von Neumann and Morgenstern
- 1950: Nash invents concept of Nash equilibrium
- Game theory booms after this…

- Nobel Laureates in Economics for game theory work
  - 1994: Harsanyi, Nash, and Selten
  - 2005: Aumann and Schelling
  - 2007: Hurwicz, Maskin, and Myerson

What is mechanism design?

- In mechanism design, we get to design the game (or mechanism)
  - e.g. the rules of the auction, marketplace, election, …
- Goal is to obtain good outcomes when agents behave strategically (game-theoretically)
- Mechanism design often considered part of game theory
- Sometimes called “inverse game theory”
  - In game theory the game is given and we have to figure out how to act
  - In mechanism design we know how we would like the agents to act and have to figure out the game

Example: (single-item) auctions

- Sealed-bid auction: every bidder submits bid in a sealed envelope
  - First-price sealed-bid auction: highest bid wins, pays amount of own bid
  - Second-price sealed-bid auction: highest bid wins, pays amount of second-highest bid

bid 1: $10
bid 2: $5
bid 3: $1
0

first-price: bid 1 wins, pays $10
second-price: bid 1 wins, pays $5

Which auction generates more revenue?

- Each bid depends on
  - bidder’s true valuation for the item (utility = valuation - payment),
  - bidder’s beliefs over what others will bid (→ game theory),
  - and... the auction mechanism used
- In a first-price auction, it does not make sense to bid your true valuation
  - Even if you win, your utility will be 0…
- In a second-price auction, it always makes sense to bid your true valuation

a likely outcome for the first-price mechanism

bid 1: $5
bid 2: $4
bid 3: $1
0

a likely outcome for the second-price mechanism

bid 1: $10
bid 2: $5
bid 3: $1
0

Combinatorial auctions

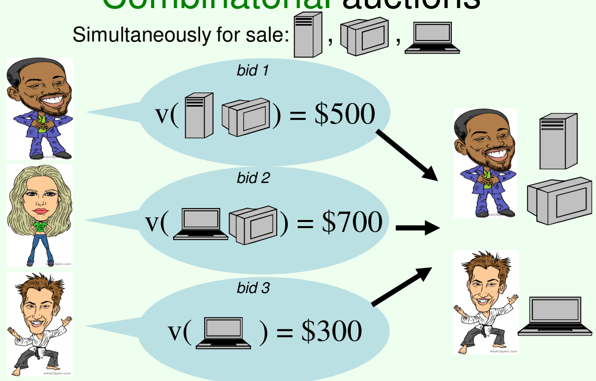

used in truckload transportation, industrial procurement, radio spectrum allocation, …

Combinatorial auction problems

- Winner determination problem
  - Deciding which bids win is (in general) a hard computational problem (NP-hard)
- Preference elicitation (communication) problem
  - In general, each bidder may have a different value for each bundle
  - But it may be impractical to bid on every bundle (there are exponentially many bundles)
- Mechanism design problem
  - How do we get the bidders to behave so that we get good outcomes?
- These problems interact in nontrivial ways
  - E.g. limited computational or communication capacity limits mechanism design options

Risk attitudes

- Which would you prefer?
  - A lottery ticket that pays out $10 with probability .5 and $0 otherwise, or
  - A lottery ticket that pays out $3 with probability 1
- How about:
  - A lottery ticket that pays out $100,000,000 with probability .5 and $0 otherwise, or
  - A lottery ticket that pays out $30,000,000 with probability 1
- Usually, people do not simply go by expected value
- An agent is risk-neutral if she only cares about the expected value of the lottery ticket
- An agent is risk-averse if she always prefers the expected value of the lottery ticket to the lottery ticket
  - Most people are like this
- An agent is risk-seeking if she always prefers the lottery ticket to the expected value of the lottery ticket

Decreasing marginal utility

- Typically, at some point, having an extra dollar does not make people much happier (decreasing marginal utility)

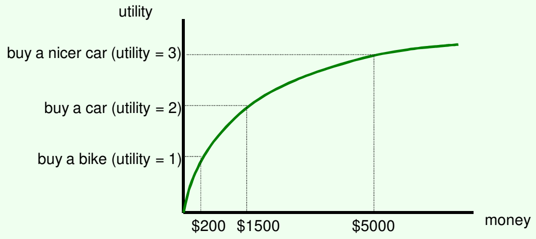

Maximizing expected utility

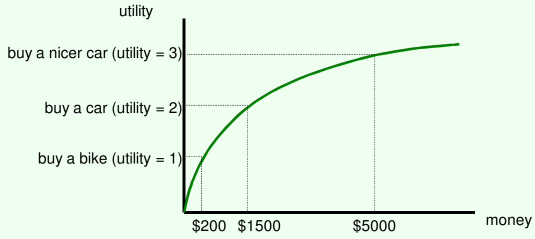

- Lottery 1: get $1500 with probability 1
  - gives expected utility 2
- Lottery 2: get $5000 with probability .4, $200 otherwise
  - gives expected utility .4*3 + .6*1 = 1.8
  - (expected amount of money = .4*$5000 + .6*$200 = $2120 > $1500)
- So: maximizing expected utility is consistent with risk aversion

Different possible risk attitudes under expected utility maximization

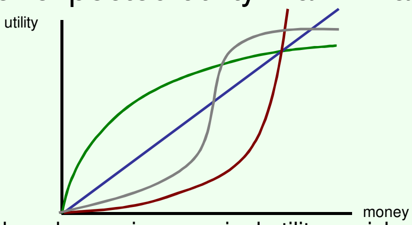

- Green has decreasing marginal utility → risk-averse
- Blue has constant marginal utility → risk-neutral
- Red has increasing marginal utility → risk-seeking
- Grey's marginal utility is sometimes increasing, sometimes decreasing → neither risk-averse (everywhere) nor risk-seeking (everywhere)

Rock-paper-scissors

Column player aka. player 2 (simultaneously) chooses a column

Row player aka. player 1 chooses a row

A row or column is called an action or (pure) strategy

	0, 0	-1, 1	1, -1
	1, -1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

Row player's utility is always listed first, column player's second

Zero-sum game: the utilities in each entry sum to 0 (or a constant)
Three-player game would be a 3D table with 3 utilities per entry, etc.

"Chicken"

- Two players drive cars towards each other
- If one player goes straight, that player wins
- If both go straight, they both die

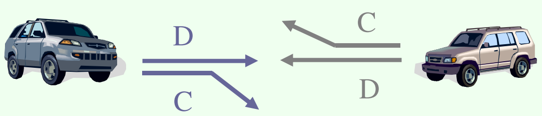

	C	D
Chicken	0, 0	-1, 1
Dare	1, -1	-5, -5

not zero-sum

Rock-paper-scissors – Seinfeld variant

MICKEY: All right, rock beats paper!
(Mickey smacks Kramer's hand for losing)
KRAMER: I thought paper covered rock.
MICKEY: Nah, rock flies right through paper.
KRAMER: What beats rock?
MICKEY: (looks at hand) Nothing beats rock.

	0, 0	1, -1	1, -1
	-1, 1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

Dominance

- Player i's strategy s_i strictly dominates s_i' if
  - for any s_{-i}, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$
- s_i weakly dominates s_i' if
  - for any s_{-i}, $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$; and
  - for some s_{-i}, $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$

-i = "the player(s) other than i"

strict dominance

weak dominance

	0, 0	1, -1	1, -1
	-1, 1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

Mixed strategies

- Mixed strategy for player i = probability distribution over player i's (pure) strategies
- E.g. 1/3 , 1/3 , 1/3
- Example of dominance by a mixed strategy:

1/2	3, 0	0, 0
1/2	0, 0	3, 0
	1, 0	1, 0

Prisoner's Dilemma

- Pair of criminals has been caught
- District attorney has evidence to convict them of a minor crime (1 year in jail); knows that they committed a major crime together (3 years in jail) but cannot prove it
- Offers them a deal:
  - If both confess to the major crime, they each get a 1 year reduction
  - If only one confesses, that one gets 3 years reduction

	confess	don't confess
confess	-2, -2	0, -3
don't confess	-3, 0	-1, -1

"Should I buy an SUV?"

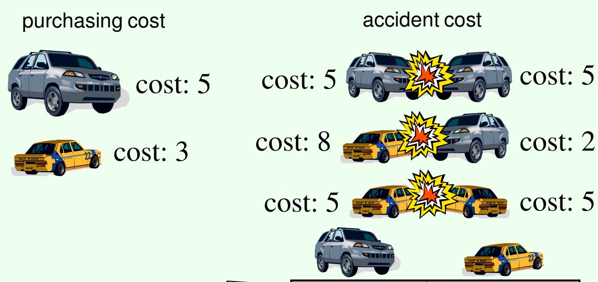

	-10, -10	-7, -11
	-11, -7	-8, -8

MaxMin Strategies

- What can the row player guarantee to herself?
  - Without assuming anything about her opponent?
- If she has a dominant strategy?
- What if she doesn't?
- $Max_i\ Min_j\ u(i,j)$

Zero-sum games

- In a zero-sum game, payoffs in each entry sum to zero
  - … or to a constant: we can subtract a constant without affecting behavior
- What the one player gains, the other player loses

	0, 0	-1, 1	1, -1
	1, -1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

Minimax theorem [von Neumann 1927]

- 2-player zero-sum games
- Which one is bigger:
  - $max_{s_{row}}\ min_{s_{col}}\ u_{row}(s_{row}, s_{col})$ (col knows s_{row}), or
  - $min_{s_{col}}\ max_{s_{row}}\ u_{row}(s_{row}, s_{col})$ (row knows s_{col})?
    - Here, s_{row} and s_{col} are pure strategies

Player row plays first

Player column plays first

- Example:

	Left	Middle	Right
Up	1	4	5
Down	3	2	6

- If row plays first, he chooses $max_i\ min_j\ u(i,j)$
- If column plays first, he chooses $min_j\ max_i\ u(i,j)$
- Lemma: $max_i\ min_j\ u(i,j) <= min_j\ max_i\ u(i,j)$ for every u
- Proof: sufficient to prove that for every j',i' $min_j u(i',j) <= max_i u(i,j')$ this holds since $min_j u(i',j) <= u(i',j') <= max_i u(i,j')$

Best-response strategies

- Suppose you know your opponent plays rock 50% of the time and scissors 50%
- What is the best strategy?
  (utility functions are extended to mixed strategies by taking the expectation of the utility over pure strategies)

	0, 0	-1, 1	1, -1
	1, -1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

- Rock gives .5*0 + .5*1 = .5
- Paper gives .5*1 + .5*(-1) = 0
- Scissors gives .5*(-1) + .5*0 = -.5
- So the best response to this opponent strategy is to (always) play rock

Best-response strategies

- There is always some pure strategy that is a best response
- Suppose you have a mixed strategy that is a best response; then every pure strategy with prob >0 in the mixed strategy must *also* be a best response

Minimax (minmax, maxmin) strategies

- 2-player zero-sum games with mixed strategies:
- Suppose that your opponent can see into your head and thus knows your mixed strategy
- But does not know the random choices
  - E.g. your opponent knows that you play rock 50% of the time and scissors 50% of the time, but not which you will actually play
  - I.e. your opponent best-responds to your mixed strategy
- What is the best that the row player can do against such a powerful column player?
- $\max_{\sigma_{row}} \min_{s_{col}} u_{row}(\sigma_{row}, s_{col})$
  - Here σ_{row} is a mixed strategy, s_{col} is a pure strategy, and utility functions are extended to mixed strategies by taking the expectation of the utility over pure strategies

Computing a minimax strategy for rock-paper-scissors

- Need to set: p_{rock}, p_{paper}, $p_{scissors}$
- Row player utility if column plays rock is $p_{paper} - p_{scissors}$
- Row player utility if column plays paper is $p_{scissors} - p_{rock}$
- Row player utility if column plays scissors is $p_{rock} - p_{paper}$
- So, $\max_{\sigma_{row}} \min_{scol} u_{row}(\sigma_{row}, s_{col}) = \max_{\{p1,p2,p3\}} \min \{p_{paper} - p_{scissors}, p_{scissors} - p_{rock}, p_{rock} - p_{paper}\}$
- Minimax (maxmin) strategy: $p_{rock} = p_{paper} = p_{scissors} = 1/3$

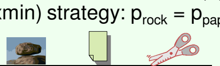

p_{rock} 	0, 0	-1, 1	1, -1
p_{paper} 	1, -1	0, 0	-1, 1
$p_{scissors}$ 	-1, 1	1, -1	0, 0

Minimax theorem [von Neumann 1927]

- In general, which one is bigger:
  - $\max_{\sigma_{row}} \min_{s_{col}} u_{row}(\sigma_{row}, s_{col})$ (col knows σ_{row}), or
  - $\min_{\sigma_{col}} \max_{s_{row}} u_{row}(s_{row}, \sigma_{col})$ (row knows σ_{col})?
- Minimax Theorem:

  $\max_{\sigma_{row}} \min_{s_{col}} u_{row}(\sigma_{row}, s_{col}) = \min_{\sigma_{col}} \max_{s_{row}} u_{row}(s_{row}, \sigma_{col})$
  - This quantity is called the value of the game (to row player)
- Follows from (predates) linear programming duality
- Ergo: if you can look into the other player's head (but the other player anticipates that), you will do no better than if the roles were reversed
- Only true if we allow for mixed strategies
  - If you know the other player's pure strategy in rock-paper-scissors, you will always win

Solving for minimax strategies using linear programming (Row player viewpoint)

- maximize **r**
- subject to
  - for any pure strategy of the column player, s_{col},
    - $\Sigma_{s_{row}} \mathbf{p}_{s_{row}} u_{row}(s_{row}, s_{col}) \geq \mathbf{r}$
  - $\Sigma_{s_{row}} \mathbf{p}_{s_{row}} = 1$
  - $\mathbf{p}_{s_{row}} \geq 0$
  - Variables: $\mathbf{p}_{s_{row}}$ for every pure strategy of the row player, s_{row}

Note: linear programs can be solved in polynomial time

Solving for minimax strategies using linear programming (Column player viewpoint)

- minimize **c**
- subject to
  - for any pure strategy of the row player, s_{row},
    - $\Sigma_{s\ col} \mathbf{q}_{s\ col} u_{row}(s_{row}, s_{col}) \leq \mathbf{c}$
  - $\Sigma_{s\ col} \mathbf{q}_{s\ col} = 1$
  - $\mathbf{q}_{s\ col} \geq 0$
  - Variables: $\mathbf{q}_{s\ col}$ for every pure strategy of the column player, s_{col}

Note: Linear programming duality: the solution to the primal is equal to the solution to the dual if both are finite

Yao's Principle, Zero Sum Games, and the Minimax Theorem

- Yao's principle is used to give lower bounds (or, even better, tight bounds) on the time complexity of randomized algorithms
- The Yao principle gives a lower bound on the time complexity of randomized algorithms as follows:
  - Choose a "nasty" distribution on inputs
  - Prove that every deterministic algorithm behaves badly on this input distribution
- Ergo, bound the expected running time of a randomized algorithm on worst case inputs by the expected running time of any deterministic algorithm on random inputs

Yao's Principle, Zero Sum Games, and the Minimax Theorem

- Fix the length of the input – if the algorithm terminates then it uses a finite number of coin tosses (say 1017) and can be viewed as a probability distribution on 2^{1017} deterministic algorithms.
- Consider the following zero sum game:
  - The (pure) strategy for the column player is to choose some (deterministic) algorithm
  - The (pure) strategy for the row player is to choose some (deterministic) input
  - The *(i,j)* payoff matrix entry is *T(i,j): th*e time it takes (det.) algorithm *j* to run on (det.) input *i*
- Mixed strategies:
  - A mixed strategy for the column player is a distribution on deterministic algorithms – another name for a randomized algorithm
  - A mixed strategy for the row player is a distribution on inputs.

Yao's Principle, Zero Sum Games, and the Minimax Theorem

- Minimax Theorem:

  $\max_{\sigma_{row}} \min_{s_{col}} u_{row}(\sigma_{row}, s_{col}) = \min_{\sigma_{col}} \max_{s_{row}} u_{row}(s_{row}, \sigma_{col})$

Or,

$\max_{\text{input dist } D} \min_{\text{det algo. } A} E_{I \text{ from } D}(\text{Time}(I,A)) = \min_{\text{rand. alg. } R} \max_{\text{det. input } I} E_{R \text{ coins}}(\text{Time}(I,R))$

Not zero sum

- One could pretend that the column player is only trying to hurt the row player (and vice versa). This is not rational:

0, 0	3, 1
1, 0	2, 1

- If Column was trying to hurt Row, Column would play Left, so Row should play Down
- If Column only cares about his own payoff, Column will play Right (strictly dominant), so Row should play Up
- Is there a better generalization of minimax strategies in zero-sum games to general-sum games?

Nash equilibrium [Nash 1950]

- A vector of strategies (one for each player) is called a strategy profile, strategies may be mixed
- A strategy profile $(\sigma_1, \sigma_2, \ldots, \sigma_n)$ is a Nash equilibrium if each σ_i is a best response to σ_{-i}
  - That is, for any i, for any σ_i', $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i})$
- This does not say anything about multiple agents changing their strategies at the same time
- (Note - singular: equilibrium, plural: equilibria)

Nash Equilibrium: some examples

Prisoner's dilemma

	cooperate	defect
cooperate	-1 , -1	-5 , 0
defect	0 , -5	-3 , -3

A unique Nash equilibrium

Battle of the sexes

	opera	boxing
opera	3 , 1	0 , 0
boxing	0 , 0	1 , 3

2 Nash equilibria

Rock , scissors, paper

	rock	scissors	paper
rock	0 , 0	1 , -1	-1 , 1
scissors	-1 , 1	0 , 0	1 , -1
paper	1 , -1	-1 , 1	0 , 0

No Nash equilibrium (in pure strategies), but 1 Nash equilibrium in mixed strategies (what is it?)

Nash's Theorem

- Theorem: Every finite game admits a NE
- Proof (sketch):
  - Based on Brower's fixed point theorem: let $C \in R^t$ be a compact convex set, and $f:C \rightarrow C$ a continuous function. Then $\exists x \in C$ *s.t.* $x=f(x)$
  - $C = \Delta(S_1)X\ldots X\Delta(S_n)$ where S_i is the strategy space of player i
  - First attempt: $f(x^1,\ldots,x^n)=(BR(x^{-1}),\ldots,BR(x^{-n}))$
    - *What is the problem?*

Nash's Theorem

- Instead of going "all the way" to BR, we'll go in that direction
- Denote the weight of pure strategy j in x^i by x^i_j
- Define $c^i_j = u^i(j,x^{-i}) - u^i(x^i,x^{-i})$, $c^{i+}_j=max(0,c^i_j)$
- Let $f:(x^1,\ldots,x^n) \rightarrow \hat{x}^1,\ldots,\hat{x}^n, where: \hat{x}^i_j = \frac{x^i_j + c^{i+}_j}{1+\sum_j c^{i+}_j}$
  - $\forall i \sum_j x^i_j c^i_j = 0$ easy to see (exercise)
  - f admits a fixed point (by Brower)
  - left to show: at the fixed point every player best responds; i.e., $c_j^{i+}=0$ for all i,j. (in class)

Nash equilibria of "chicken"

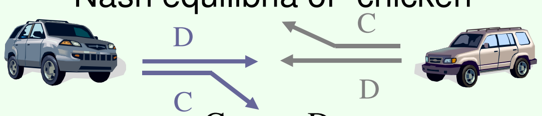

	C	D
C	0, 0	-1, 1
D	1, -1	-5, -5

D = Dare
C = Chicken

- (D, C) and (C, D) are Nash equilibria
  - They are pure-strategy Nash equilibria: nobody randomizes
  - They are also strict Nash equilibria: changing your strategy will make you strictly worse off
- No other pure-strategy Nash equilibria

		p^2_C C	p^2_D D
p^1_C	C	0, 0	-1, 1
	D	1, -1	-5, -5

- Is there a Nash equilibrium that uses mixed strategies? Say, where player 1 (row player) uses a mixed strategy?
- Recall: if a mixed strategy is a best response, then all of the pure strategies that it randomizes over must also be best responses
- Player 1's utility for playing C = $-p^2_D$
- Player 1's utility for playing D = $p^2_C - 5p^2_D = 1 - 6p^2_D$
- So we need $-p^2_D = 1 - 6p^2_D$ which means $p^2_D = 1/5$
- Then, player 2 needs to be indifferent as well
- Mixed-strategy Nash equilibrium: **((4/5 C, 1/5 D), (4/5 C, 1/5 D))**
  - People may die! Expected utility -1/5 for each player

The presentation game

		Presenter: *Put effort into presentation (E)*	Presenter: *Do not put effort into presentation (NE)*
Audience	*Pay attention (A)*	4, 4	-16, -14
Audience	*Do not pay attention (NA)*	0, -2	0, 0

- Pure-strategy Nash equilibria: (A, E), (NA, NE)
- Mixed-strategy Nash equilibrium:
  ((1/10 A, 9/10 NA), (4/5 E, 1/5 NE))
  - Utility 0 for audience, -14/10 for presenter
  - Can see that some equilibria are strictly better for both players than other equilibria, i.e. some equilibria Pareto-dominate other equilibria

The "equilibrium selection problem"

- You are about to play a game that you have never played before with a person that you have never met
- According to which equilibrium should you play?
- Possible answers:
  - Equilibrium that maximizes the sum of utilities (social welfare)
  - Or, at least not a Pareto-dominated equilibrium
  - So-called focal equilibria
    - "Meet in Paris" game - you and a friend were supposed to meet in Paris at noon on Sunday, but you forgot to discuss where and you cannot communicate. All you care about is meeting your friend. Where will you go?
  - Equilibrium that is the convergence point of some learning process
  - An equilibrium that is easy to compute
  - …
- Equilibrium selection is a difficult problem

Some properties of Nash equilibria

- If you can eliminate a strategy using strict dominance or even iterated strict dominance, it will not occur (i.e. it will be played with probability 0) in every Nash equilibrium
  - Weakly dominated strategies may still be played in some Nash equilibrium
- In 2-player zero-sum games, a profile is a Nash equilibrium if and only if both players play minimax strategies
  - Hence, in such games, if (σ_1, σ_2) and (σ_1', σ_2') are Nash equilibria, then so are (σ_1, σ_2') and (σ_1', σ_2)
    - No equilibrium selection problem here!

How hard is it to compute *one* (any) Nash equilibrium?

- Complexity was open for a long time
  - [Papadimitriou STOC01]: "together with factoring […] the most important concrete open question on the boundary of P today"
- Recent sequence of papers shows that computing one (any) Nash equilibrium is PPAD-complete (even in 2-player games) [Daskalakis, Goldberg, Papadimitriou 05; Chen, Deng 05]
- All known algorithms require exponential time (in the worst case)

What if we want to compute a Nash equilibrium with a specific property?

- For example:
  - An equilibrium that is not Pareto-dominated
  - An equilibrium that maximizes the expected social welfare (= the sum of the agents' utilities)
  - An equilibrium that maximizes the expected utility of a given player
  - An equilibrium that maximizes the expected utility of the worst-off player
  - An equilibrium in which a given pure strategy is played with positive probability
  - An equilibrium in which a given pure strategy is played with zero probability
  - …
- All of these are NP-hard, even in 2-player games [Gilboa, Zemel 89; Conitzer & Sandholm IJCAI-03]

Example: finding equilibrium that maximizes social welfare is NPC

- The *bi-clique* problem:
  - Given a bipartite graph G and a number k
  - Are there subsets of V_{up} and V_{down} (of size k each) that form a bi-clique ?
  - E.g., G admits a 3-biclique but not a 4-biclique

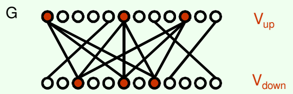

Example: finding equilibrium that maximizes social welfare is NPC

- Reduction from bi-clique:

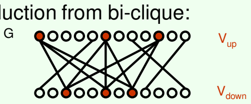

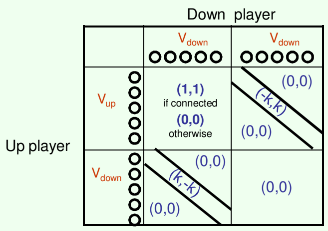

Lemma: There exists a Nash equilibrium with social welfare = 2 iff G admits a k-biclique

Proof: in class

Search-based approaches (for 2 players)

- Suppose we know the support X_i of each player i's mixed strategy in equilibrium
  - That is, which pure strategies receive positive probability
- Then, we have a linear feasibility problem:
  - for both i, for any $s_i \in X_i$, $\Sigma \mathbf{p_{-i}(s_{-i})} u_i(s_i, s_{-i}) = \mathbf{u_i}$
  - for both i, for any $s_i \in S_i - X_i$, $\Sigma \mathbf{p_{-i}(s_{-i})} u_i(s_i, s_{-i}) \leq \mathbf{u_i}$
- Thus, we can search over possible supports
  - This is the basic idea underlying methods in [Dickhaut & Kaplan 91; Porter, Nudelman, Shoham AAAI04; Sandholm, Gilpin, Conitzer AAAI05]
- Dominated strategies can be eliminated

Correlated equilibrium [Aumann 74]

- Suppose there is a mediator who has offered to help out the players in the game
- The mediator chooses a profile of pure strategies, perhaps randomly, then tells each player what her strategy is in the profile (but not what the other players' strategies are)
- A correlated equilibrium is a distribution over pure-strategy profiles for the mediator, so that every player wants to follow the recommendation of the mediator (if she assumes that the others do so as well)
- Every Nash equilibrium is also a correlated equilibrium
  - Corresponds to mediator choosing players' recommendations independently
- … but not vice versa
- (Note: there are more general definitions of correlated equilibrium, but it can be shown that they do not allow you to do anything more than this definition.)

New version of "Chicken"

	C	D
C	8,8	1,9
D	9,1	0,0

- Two pure NE: (D,C),(C,D)
  - Social welfare (sum of payoffs) = 10
- One mixed NE: (½ C, ½ D),(½ C, ½ D)
  - Expected social welfare = 9
- Can sum of payoffs be improved by a correlated equilibrium?

A correlated equilibrium for "chicken"

	C	D
C	8,8 1/3	1,9 1/3
D	9,1 1/3	0,0 0

Expected social welfare = 12

- Why is this a correlated equilibrium?
- Suppose the mediator tells the row player to Chicken
- From Row's perspective, the conditional probability that Column was told to Chicken is (1/3) / (1/3 + 1/3) = 1/2
- So the expected utility of Chicken is (1/2)*(8)+ (1/2)*1= 4.5
- But the expected utility of Dare is (1/2)*9 + (1/2)*0 = 4.5
- So Row wants to follow the recommendation
- If Row is told to Dare, he knows that Column was told to Chicken, so again Row wants to follow the recommendation
- Similar for Column

A nonzero-sum variant of rock-paper-scissors (Shapley's game [Shapley 64])

			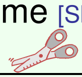
	0, 0 0	0, 1 1/6	1, 0 1/6
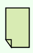	1, 0 1/6	0, 0 0	0, 1 1/6
	0, 1 1/6	1, 0 1/6	0, 0 0

- If both choose the same pure strategy, both lose
- These probabilities give a correlated equilibrium:
- E.g. suppose Row is told to play Rock
- Row knows Column is playing either paper or scissors (50-50)
  - Playing Rock will give ½; playing Paper will give 0; playing Scissors will give ½
- So Rock is optimal (not uniquely)

Solving for a correlated equilibrium using linear programming (n players!)

- Variables are now $\mathbf{p_s}$ where s is a profile of pure strategies (note the difference between mixed eq. and correlated eq.)
- maximize *whatever you like (e.g. social welfare)*
- subject to
  - for any i, s_i, s_i', $\Sigma_{s_{-i}} \mathbf{p}_{(s_i, s_{-i})} u_i(s_i, s_{-i}) \geq \Sigma_{s_{-i}} \mathbf{p}_{(s_i, s_{-i})} u_i(s_i', s_{-i})$
  - $\Sigma_s \mathbf{p_s} = 1$