

CPS 196.2

Utility theory, normal-form games


Vincent Conitzer
conitzer@cs.duke.edu

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)

What is utility, anyway?

- Function $u: O \rightarrow \mathfrak{R}$ (O is the set of “outcomes” that lotteries randomize over)
 - What are its units?
 - It doesn’t really matter
 - If you replace your utility function by $u'(o) = a + bu(o)$, your behavior will be unchanged
 - Why would you want to maximize expected utility?
 - For two lottery tickets L and L' , let $pL + (1-p)L'$ be the “compound” lottery ticket where you get lottery ticket L with probability p , and L' with probability $1-p$
 - $L \geq L'$ means that L is (weakly) preferred to L'
 - (\geq should be complete, transitive)
 - **Expected utility theorem.** Suppose
 - (continuity axiom) for all L, L', L'' , $\{p: pL + (1-p)L' \geq L''\}$ and $\{p: pL + (1-p)L' \leq L''\}$ are closed sets,
 - (independence axiom – more controversial) for all L, L', L'' , p , we have $L \geq L'$ if and only if $pL + (1-p)L'' \geq pL' + (1-p)L''$
- then there exists a function $u: O \rightarrow \mathfrak{R}$ so that $L \geq L'$ if and only if L gives a higher expected value of u than L'


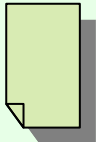


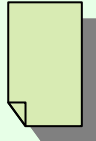

Normal-form games

Rock-paper-scissors

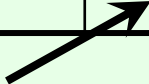
Column player aka.
player 2
(simultaneously)
chooses a column

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






	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5


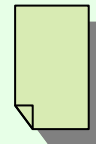


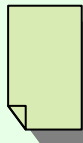

not zero-sum

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.


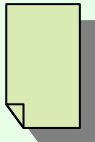


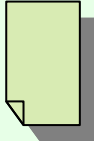
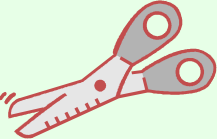


			
	0, 0	1, -1	1, -1
	-1, 1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

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 "

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

strict dominance

weak dominance

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?”

purchasing cost

accident cost



cost: 5

cost: 5



cost: 5



cost: 3

cost: 8




cost: 2

cost: 5


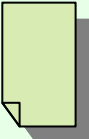
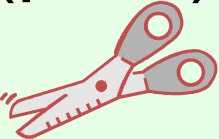


cost: 5



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

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


A blue bracket on the left side of the table groups the first two rows, with a curved arrow pointing from the bracket to the third row, indicating that the third row is dominated by the mixed strategy of the first two rows.

Checking for dominance by mixed strategies

- Linear program for checking whether strategy s_i^* is **strictly** dominated by a mixed strategy:
 - maximize ε
 - such that:
 - for any s_{-i} , $\sum_{s_i} \mathbf{p}_{s_i} u_i(s_i, s_{-i}) \geq u_i(s_i^*, s_{-i}) + \varepsilon$
 - $\sum_{s_i} \mathbf{p}_{s_i} = 1$
- Linear program for checking whether strategy s_i^* is **weakly** dominated by a mixed strategy:
 - maximize $\sum_{s_{-i}} (\sum_{s_i} \mathbf{p}_{s_i} u_i(s_i, s_{-i})) - u_i(s_i^*, s_{-i})$
 - such that:
 - for any s_{-i} , $\sum_{s_i} \mathbf{p}_{s_i} u_i(s_i, s_{-i}) \geq u_i(s_i^*, s_{-i})$
 - $\sum_{s_i} \mathbf{p}_{s_i} = 1$


Iterated dominance

- Iterated dominance: remove (strictly/weakly) dominated strategy, repeat
- Iterated strict dominance on Seinfeld's RPS:



Iterated dominance: path (in)dependence

Iterated weak dominance is **path-dependent**:
sequence of eliminations may determine which
solution we get (if any)
(whether or not dominance by mixed strategies allowed)



Iterated strict dominance is **path-independent**: elimination
process will always terminate at the same point
(whether or not dominance by mixed strategies allowed)


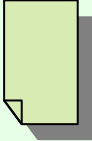


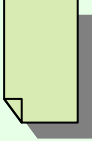

Two computational questions for iterated dominance

- 1. Can a **given strategy** be eliminated using iterated dominance?
- 2. Is there some path of elimination by iterated dominance such that only **one strategy per player remains**?

- For strict dominance (with or without dominance by mixed strategies), both can be solved in polynomial time due to path-independence:
 - Check if any strategy is dominated, remove it, repeat
- For weak dominance, both questions are NP-hard (even when all utilities are 0 or 1), with or without dominance by mixed strategies [Conitzer, Sandholm 05]
 - Weaker version proved by [Gilboa, Kalai, Zemel 93]

Zero-sum games revisited

- Recall: in a zero-sum game, payoffs in each entry sum to zero
 - ... or to a constant: recall that we can subtract a constant from anyone's utility function without affecting their 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

Best-response strategies

- Suppose you know your opponent's mixed strategy
 - E.g. your opponent plays rock 50% of the time and scissors 50%
- What is the best strategy for you to play?
- 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
- There is always some **pure** strategy that is a best response
 - Suppose you have a mixed strategy that is a best response; then every one of the pure strategies that that mixed strategy places positive probability on must also be a best response

Minimax (minmax, maxmin) strategies

- Let us consider 2-player zero-sum games
- Suppose that your opponent can see into your head and thus knows your mixed strategy
- But your opponent does not know your random choice
 - E.g. your opponent knows that you play rock 50% of the time and scissors 50% of the time, but not which one you will actually happen to play this time
 - I.e. your opponent best-responds to your mixed strategy
- What is the best that you (i) can do against such a powerful opponent (-i)?
- $\max_{\sigma_i} \min_{s_{-i}} u_i(\sigma_i, s_{-i})$ (= - $\min_{\sigma_i} \max_{s_{-i}} u_{-i}(\sigma_i, s_{-i})$)
 - Here σ_i is a mixed strategy, s_{-i} 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: $p_{\text{rock}}, p_{\text{paper}}, p_{\text{scissors}}$
- Utility for other player of playing rock is $p_{\text{scissors}} - p_{\text{paper}}$
- Utility for other player of playing paper is $p_{\text{rock}} - p_{\text{scissors}}$
- Utility for other player of playing scissors is $p_{\text{paper}} - p_{\text{rock}}$
- So, we want to minimize $\max\{p_{\text{scissors}} - p_{\text{paper}}, p_{\text{rock}} - p_{\text{scissors}}, p_{\text{paper}} - p_{\text{rock}}\}$
- Minimax strategy: $p_{\text{rock}} = p_{\text{paper}} = p_{\text{scissors}} = 1/3$

Practice games

20, -20	0, 0
0, 0	10, -10

20, -20	0, 0	10, -10
0, 0	10, -10	8, -8

Minimax theorem [von Neumann 1927]

- In general, which one is bigger:
 - $\max_{\sigma_i} \min_{s_{-i}} u_i(\sigma_i, s_{-i})$ (-i gets to look inside i's head), or
 - $\min_{\sigma_{-i}} \max_{s_i} u_i(s_i, \sigma_{-i})$ (i gets to look inside -i's head)?
- Answer: they are always the same!!!
 - This quantity is called the **value** of the game (to player i)
- Closely related to linear programming duality
- Summarizing: 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

- maximize \mathbf{u}_i
- subject to
 - for any s_{-i} , $\sum_{s_i} \mathbf{p}_{s_i} u_i(s_i, s_{-i}) \geq \mathbf{u}_i$
 - $\sum_{s_i} \mathbf{p}_{s_i} = 1$

General-sum games

- You could still play a minimax strategy in general-sum games
 - I.e. pretend that the opponent is only trying to hurt you
- But 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
- In reality, 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 50]

- A vector of strategies (one for each player) is called a **strategy profile**
- A strategy profile $(\sigma_1, \sigma_2, \dots, \sigma_n)$ is a **Nash equilibrium** if each σ_i is a best response to σ_{-i}
 - That is, for any i , for any σ_i' , $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i})$
- Note that this does not say anything about multiple agents changing their strategies at the same time
- In any (finite) game, at least one Nash equilibrium (possibly using mixed strategies) exists [Nash 50]
- (Note - singular: equilibrium, plural: equilibria)

Nash equilibria of “chicken”



	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- (D, S) and (S, 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

Nash equilibria of “chicken” ...

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- Is there a Nash equilibrium that uses mixed strategies? Say, where player 1 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
- So we need to make player 1 **indifferent** between D and S
- Player 1's utility for playing D = $-p^c_S$
- Player 1's utility for playing S = $p^c_D - 5p^c_S = 1 - 6p^c_S$
- So we need $-p^c_S = 1 - 6p^c_S$ which means $p^c_S = 1/5$
- Then, player 2 needs to be indifferent as well
- Mixed-strategy Nash equilibrium: ((4/5 D, 1/5 S), (4/5 D, 1/5 S))
 - People may die! Expected utility -1/5 for each player

The presentation game

		Presenter	
		<i>Put effort into presentation (E)</i>	<i>Do not put effort into presentation (NE)</i>
Audience	<i>Pay attention (A)</i>	4, 4	-16, -14
	<i>Do not pay attention (NA)</i>	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 (and the optimization questions are inapproximable assuming $ZPP \neq NP$), even in 2-player games [Gilboa, Zemel 89; Conitzer & Sandholm IJCAI-03, extended draft]

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 $s_i \in X_i$, $\sum \mathbf{p}_{-i}(\mathbf{s}_{-i}) u_i(s_i, \mathbf{s}_{-i}) = \mathbf{u}_i$
 - for both i , for any $s_i \in S_i - X_i$, $\sum \mathbf{p}_{-i}(\mathbf{s}_{-i}) u_i(s_i, \mathbf{s}_{-i}) \leq \mathbf{u}_i$
- Thus, we can search over possible supports
 - This is the basic idea underlying methods in [Dickhaut & Kaplan 91; Porter, Nudelman, Shoham AAAI04]
- Dominated strategies can be eliminated

Solving for a Nash equilibrium using MIP (2 players)

[Sandholm, Gilpin, Conitzer AAAI05]

- maximize *whatever you like (e.g. social welfare)*
- subject to
 - for both i , for any s_i , $\sum_{s_{-i}} \mathbf{p}_{s_{-i}} u_i(s_i, s_{-i}) = \mathbf{u}_{s_i}$
 - for both i , for any s_i , $\mathbf{u}_i \geq \mathbf{u}_{s_i}$
 - for both i , for any s_i , $\mathbf{p}_{s_i} \leq \mathbf{b}_{s_i}$
 - for both i , for any s_i , $\mathbf{u}_i - \mathbf{u}_{s_i} \leq M(1 - \mathbf{b}_{s_i})$
 - for both i , $\sum_{s_i} \mathbf{p}_{s_i} = 1$
- \mathbf{b}_{s_i} is a binary variable indicating whether s_i is in the support, M is a large number

Correlated equilibrium [Aumann 74]


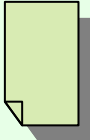


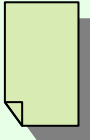

- Suppose there is a **mediator** who has offered to help out the players in the game
- The mediator chooses a profile of pure strategies, perhaps randomly, then tells each player what her strategy is in the profile (but not what the other players' strategies are)
- A **correlated equilibrium** is a 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.)

A correlated equilibrium for “chicken”

	D	S
D	0, 0 20%	-1, 1 40%
S	1, -1 40%	-5, -5 0%

- Why is this a correlated equilibrium?
- Suppose the mediator tells the row player to Dodge
- From Row’s perspective, the conditional probability that Column was told to Dodge is $20\% / (20\% + 40\%) = 1/3$
- So the expected utility of Dodging is $(2/3)*(-1) = -2/3$
- But the expected utility of Straight is $(1/3)*1 + (2/3)*(-5) = -3$
- So Row wants to follow the recommendation
- If Row is told to go Straight, he knows that Column was told to Dodge, 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 $\frac{1}{2}$; playing Paper will give 0; playing Scissors will give $\frac{1}{2}$
- 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
- maximize *whatever you like (e.g. social welfare)*
- subject to
 - for any $i, s_i, s_i', \sum_{s_{-i}} \mathbf{p}_{(s_i, s_{-i})} u_i(s_i, s_{-i}) \geq \sum_{s_{-i}} \mathbf{p}_{(s_i', s_{-i})} u_i(s_i', s_{-i})$
 - $\sum_s \mathbf{p}_s = 1$