GAME THEORY

László Gulyás
Associate Professor, ELTE, AI Department
lgulyas@inf.elte.hu

Tamás Takács
PhD student, ELTE, AI Department
tamastheactual@inf.elte.hu
tamastheactual.github.io

Game Theory

Lecture 10

Fairness, Pareto Optimality, and Social Welfare

1 Pareto Optimality

2 Fairness

3 Social Welfare

Fictitious Play

Finite normal form game:

$$\mathcal{G} = (N, \{A_i\}_{i=1}^{n}, \{u_i\}_{i=1}^{n})$$

- A_i finite action set for player
- joint action space
- $u_i: A \to R$ payoff function
- Mixed strategies $\Delta(A_i)$, product distributions $\Delta(A) = \prod_i \Delta(A_i)$

Empirical frequencies:

$$\widehat{\pi_{i,t}} = \frac{1}{t}\sum_{\tau=1}^{t} \delta_{a_i^\tau} \in \Delta(A_i)$$

Joint empirical distribution: $\widehat{\mu_t} = \frac{1}{t}\sum_{\tau=1}^{t} \delta_{a^\tau}$

Stochastic Fictitious Play

Perturbation: Add noise or use SBR with temperature $1/\eta$

$$x_i^{t+1} = (1 - \alpha_t)x_i^t + \alpha_t \cdot \mathrm{SBR}_i\left(x_{-i}^t; \eta\right)$$

with decreasing step size $\alpha_t \to 0$

Convergence: Under suitable noise and structure:

- Converges to **quantal response equilibria** (fixed η)
- Converges to NE as $\eta \to \infty$ (under additional assumptions)
- Connection to stochastic approximation theory

Regret Definitions

For sequence of play $a^1, \dots, a^T$, player i:

External regret relative to fixed action a_i':

$$R_T^{\text{ext}}(a_i') = \sum_{t=1}^{T} \left(u_i(a_i', a_{-i}^t) - u_i(a^t) \right)$$

Internal regret for conditional swap $a \to a'$:

$$R_T^{\text{int}}(a \to a') = \sum_{t=1}^{T} \mathbb{1}\{a_i^t = a\} \left(u_i(a', a_{-i}^t) - u_i(a, a_{-i}^t) \right)$$

Swap regret for replacement map $\phi: A_i \to A_i$:

$$R_T^{\text{swap}}(\phi) = \sum_{t=1}^{T} \left(u_i(\phi(a_i^t), a_{-i}^t) - u_i(a^t) \right)$$

Hannan Consistency

Definition: Player i is **Hannan consistent** if

$$\frac{1}{T}\max_{a_i' \in A_i} R_T^{\text{ext}}(a_i') \to 0 \quad \text{as } T \to \infty$$

Interpretation: Average external regret vanishes

Achievability: Many efficient algorithms (Hedge, FTRL, OMD, etc.)

External Regret → CCE

Theorem: If all players use algorithms with $o(T)$external regret:

$$\widehat{\mu_T} \to T \to \infty \text{CCE}(\mathcal{G})$$

Proof sketch:

- CCE requires no profitable deviation before observing recommendation
- External regret $=$ 0 means no profitable unilateral deviation
- Time-average satisfies CCE constraints asymptotically

Rate: With $O(\sqrt{T})$ regret, $O(1/\sqrt{T})$ constraint violation

Internal/Swap Regret → CE

Theorem: If all players use algorithms with $o(T)$ internal regret:

$$\widehat{\mu_T} \to T \to \infty \text{CE}(\mathcal{G})$$

Reduction: Hart & Mas-Colell (2000) show how to convert external regret algorithm to internal regret:

- Maintain m copies of external regret minimizer
- Copy a tracks regret for swap $a \ \to \cdot$
- Aggregate recommendations via master algorithm

Cost: m –fold increase in computation

Game Theory

Lecture 10

Fairness, Pareto Optimality, and Social Welfare

1 Pareto Optimality

2 Fairness

3 Social Welfare

Equilibrium, Efficiency, and Fairness

Equilibrium ≠ Efficiency ≠ Fairness

Consider these scenarios:

- **Prisoner's Dilemma:** Unique NE is (Defect, Defect), but (Cooperate, Cooperate) Pareto dominates it
- **Traffic routing:** Wardrop equilibrium can be inefficient (Pigou's example)
- **Resource allocation:** Equal split is fair but may waste utility if preferences differ
- **Auction revenue:** VCG is efficient but may be unfair to budget-constrained agents

How do we measure and achieve both efficiency and fairness?

Why Fairness and Efficiency Matter

Efficiency concerns:

- Avoid waste: don't leave Pareto improvements on the table
- Welfare theorems: competitive equilibria are efficient under convexity
- Social planner wants to maximize some aggregate objective

Fairness concerns:

- Justice: distributional equity matters
- Stability: unfair outcomes may be rejected or cause conflict
- Incentives: perceived fairness affects participation

The challenge: Efficiency and fairness often conflict - design mechanisms that balance both

Pareto Frontier and Domination

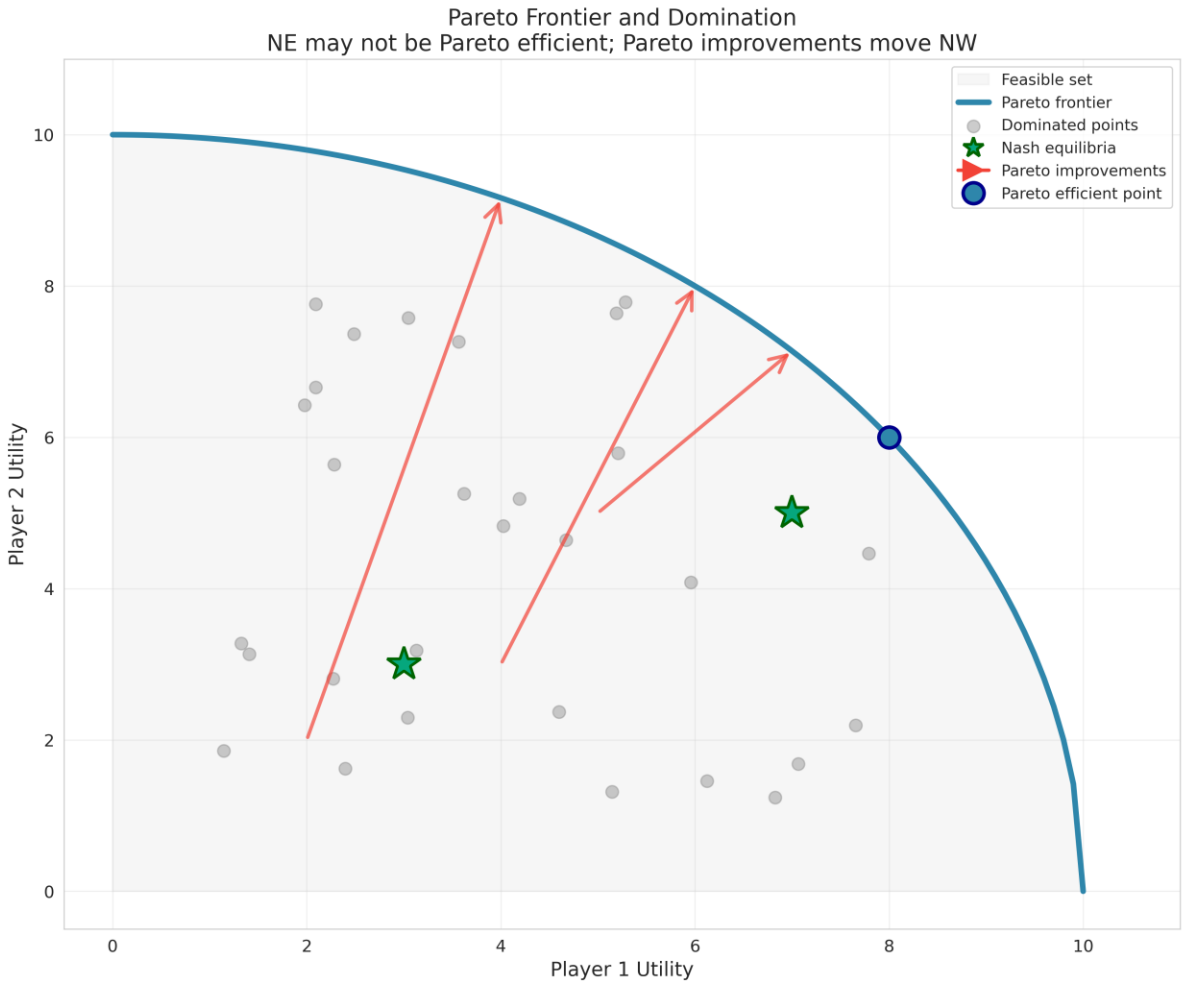

Pareto Optimality

Feasible utility set:*

$$\mathcal{U} \subset R^n$$

where $u = (u_1, \dots, u_n)$ is utility profile for n players

Notation:

- $u \geq v$: component-wise, i.e., $u_i \geq v_i$ for all i
- $u > v$: $u \geq v$ and $u \neq v$ (strict Pareto domination)
- $u \gg v$: $u_i > v_i$ for all i (strongly dominates)

Pareto Domination

Definition (Pareto Domination):

Point $u \in \mathcal{U}$ **Pareto dominates** $v \in \mathcal{U}$ if:

- $u_i \geq v_i$ for all i
- $u_j > v_j$ for at least one j

Interpretation: At least one player strictly better off, no player worse off

Pareto Efficiency

Definition (Pareto Efficient):

Point $u \in \mathcal{U}$ is **Pareto efficient** (or Pareto optimal) if there exists no $v \in \mathcal{U}$ that Pareto dominates u.

Pareto frontier: Set of all Pareto efficient points

$$\mathrm{PE}(\mathcal{U}) = \{u \in \mathcal{U} : \nexists v \in \mathcal{U} \text{ s.t. } v \text{ Pareto dominates } u\}$$

Nash Equilibrium vs Pareto Efficiency

Observation: Nash equilibria need not be Pareto efficient

Example (Prisoner's Dilemma):

	H	T
H	$(1,-1)$	$(-1,1)$
T	$(-1,1)$	$(1,-1)$

- Unique NE: (D,D) with payoff $(1,1)$
- Pareto dominates by (C,C) with payoff $(3,3)$

Zero-sum games: NE **are** Pareto efficient (one player's gain = other's loss)

Welfare Theorems (Informal)

First Welfare Theorem:

In competitive markets with convex preferences and complete markets, every competitive equilibrium is Pareto efficient.

Second Welfare Theorem:

Every Pareto efficient allocation can be decentralized as competitive equilibrium at some prices, after lump-sum redistribution.

Implication: Markets can implement efficient outcomes; fairness requires redistribution

Social Welfare Functions Comparison

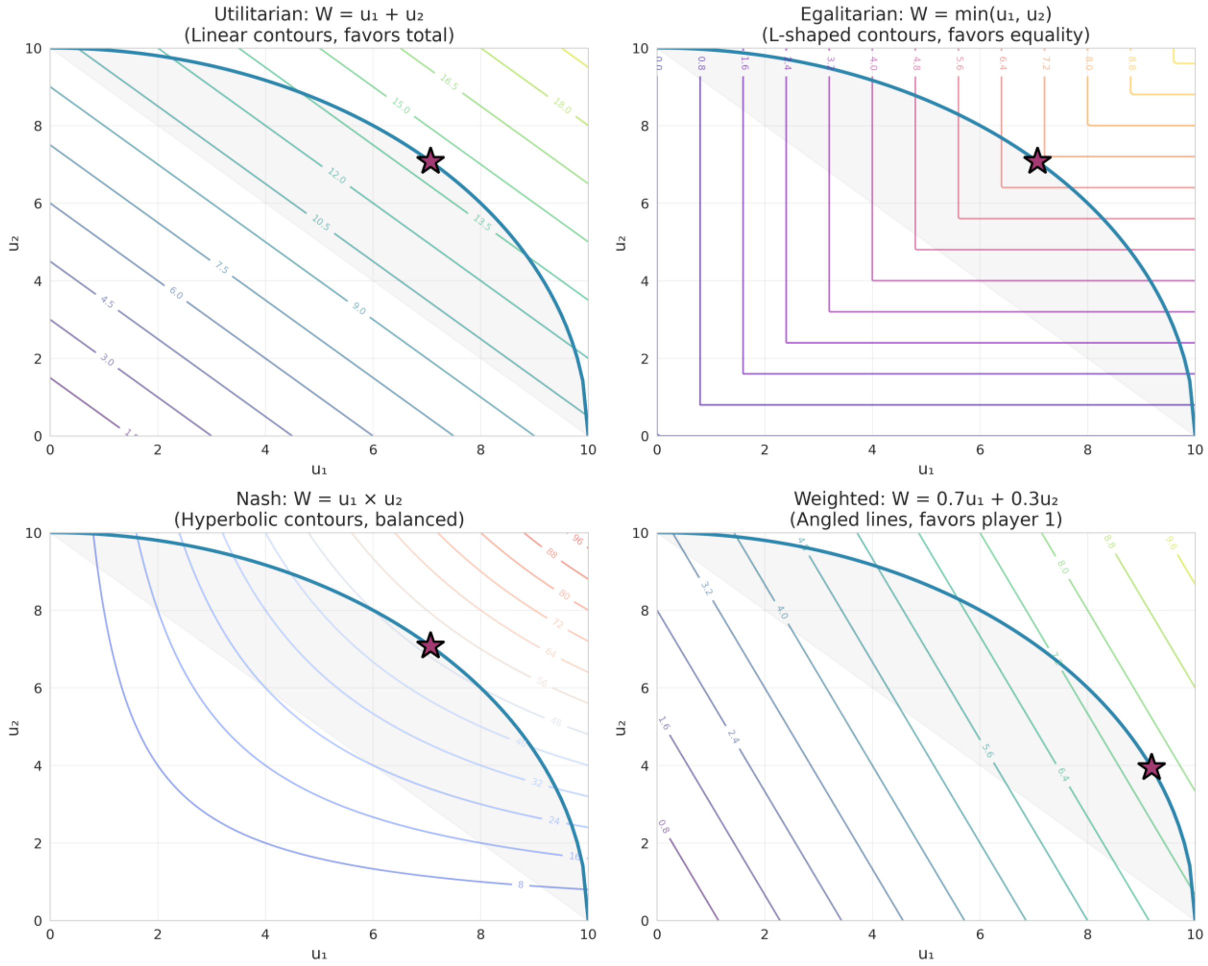

Defining Social Welfare

Social Welfare Function (SWF):

Maps utility profile to scalar:

$$W : R^n \to R$$

Purpose: Rank allocations, select Pareto efficient point

Common choices: Different ethical/efficiency priorities

Note: Basic welfare concepts (Pareto efficiency, social welfare) introduced in Lecture 3. Here we provide comprehensive treatment of welfare functions, fairness criteria, and market mechanisms.

Utilitarian Welfare

Definition:

$$W^{\text{util}}(u) = \sum_{i=1}^{n} u_i$$

Properties:

- Simple, computationally tractable
- Maximizes total surplus
- Ignores distribution: $10 + 0$ same as $5 + 5$

Maximization: Often solved via linear programming

Egalitarian (Maximin) Welfare

Definition:

$$W^{\text{egal}}(u) = \min_{i=1,\ldots,n} u_i$$

Properties:

- Rawlsian: prioritize worst-off agent
- Very inequality-averse
- Can be inefficient: improving rich doesn't help if poor unchanged

Maximization: Non-smooth, use LP or nonlinear methods

Nash Social Welfare

Definition (Geometric Mean):

$$W^{\text{nash}}(u) = \prod_{i=1}^{n} u_i$$

Usually maximize: $\sum_{i=1}^{n} \log u_i$ (equivalent for $u_i > 0$)

Properties:

- Balances efficiency and equity
- Scale invariant, respects Pigou-Dalton transfers
- Concave → optimization tractable

Connection: Arises in Fisher markets (Eisenberg-Gale)

Leximin

Definition:

Order utilities ascending: $u_{(1)} \leq u_{(2)} \leq \cdots \leq u_{(n)}$

Lexicographically maximize: first maximize $u_{(1)}$, then $u_{(2)}$, etc.

Properties:

- Extremely egalitarian
- Refines maximin when ties exist
- Hard to optimize in general

Weighted Sums

Definition:

$$W^{\lambda}(u) = \sum_{i=1}^{n} \lambda_i u_i$$

where $\lambda_i \geq 0, \sum_i \lambda_i = 1$

Properties:

- Generalizes utilitarian (all $\lambda_i = 1/n$)
- Different weights = different priorities
- Convex combination when $\mathcal{U}$ convex

Axiomatic Properties

Common axioms:

- **Anonymity:** Invariant under player permutations
- **Scale invariance:** Ranking preserved under rescaling utilities
- **Pareto monotonicity:** If $u > v$ componentwise, then $W(u) > W(v)$
- **Pigou-Dalton transfer principle:** Mean-preserving transfer from rich to poor increases W

Which SWFs satisfy which?

- Utilitarian: Anonymous, Pareto monotonic; NOT scale invariant (unless cardinal)
- Nash: Anonymous, scale invariant, Pigou-Dalton, Pareto monotonic
- Egalitarian: Anonymous, Pigou-Dalton

Schur Concavity

Definition:

W is **Schur concave** if for any u, v where v is obtained from u by a mean-preserving transfer, $W(v) \geq W(u)$.

Result: Nash social welfare $\prod_i u_i$ is Schur concave on R^n_{++}

Implication: Nash SWF respects inequality aversion

Computing Pareto Sets

Theorem (Supporting Hyperplane):

If $\mathcal{U}$ is convex and compact, every Pareto efficient point solves

$$\max_{u \in \mathcal{U}} \sum_{i=1}^{n} \lambda_i u_i$$

for some weights $\lambda \geq 0, \sum_i \lambda_i = 1$.

Conversely, every optimizer is Pareto efficient.

Algorithm: Vary λ, solve optimization, trace frontier

Limitation: Non-Convex Sets

Problem: Weighted sums only recover **convex hull** of Pareto frontier

Solution: Use ε-constraint method

Epsilon-Constraint Method

Idea: Fix thresholds on $n-1$ objectives, maximize remaining one (using ε-constraints)

$$\max_{u\in\mathcal{U}} u_1 \quad \text{s.t.} \quad u_j \geq \varepsilon_j,, j = 2, \dots, n$$

Procedure:

1. Vary ε_j systematically
2. Solve constrained optimization
3. Collect optimal points → Pareto set

Advantage: Works for non-convex $\mathcal{U}$

Game Theory

Lecture 10

Fairness, Pareto Optimality, and Social Welfare

1 Pareto Optimality

2 Fairness

3 Social Welfare

Envy-Freeness (EF)

Setup: Allocate goods to agents with valuations $v_i: \mathcal{X} \to R$

Definition:

Allocation $(X_1, \ldots, X_n)$ is **envy-free** if for all i, j:

$$v_i(X_i) \geq v_i(X_j)$$

Interpretation: No agent prefers another's bundle

Proportionality (PROP)

Definition:

Allocation is **proportional** if for each agent i:

$$v_i(X_i) \geq \frac{1}{n} v_i(\text{all goods})$$

Relation to EF:

- For additive identical valuations: EF ⇒ PROP
- General case: neither implies the other

Fairness Allocations

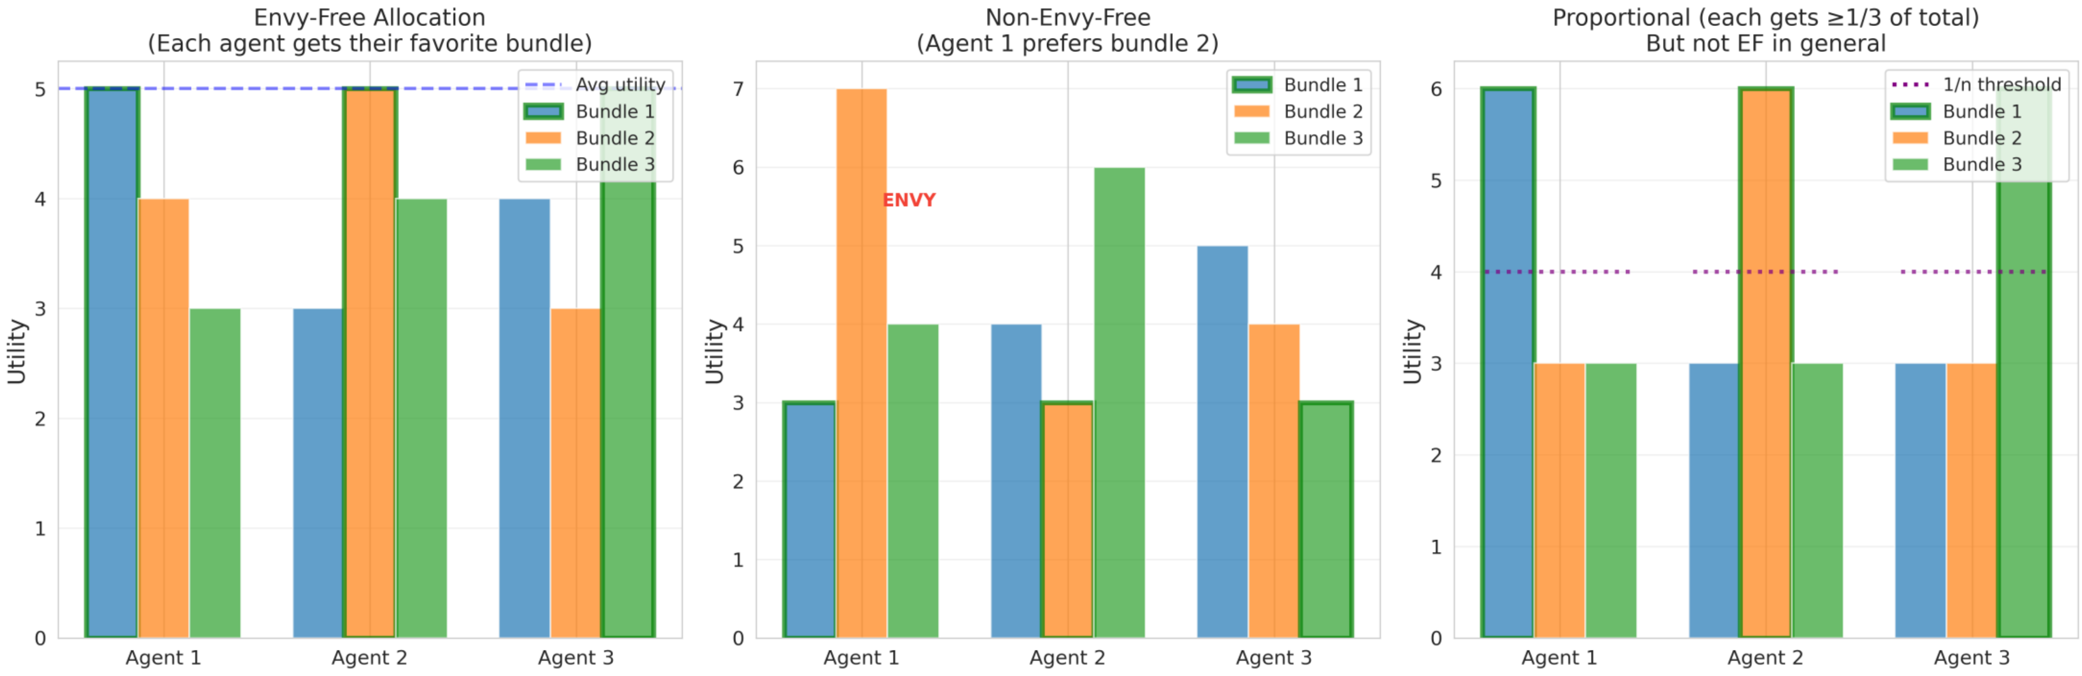

EF1 and EFX (Indivisible Goods)

Problem: Exact EF may be impossible with indivisibles

EF1 (Envy-Free up to 1 good):

For all $i \neq j, \exists g \in X_j$ such that:

$$v_i(X_i) \geq v_i\left(X_j \setminus \{g\}\right)$$

EFX (Envy-Free up to any good):

For all $i \neq j$, for all $g \in X_j$ with $v_i(g) > 0$:

$$v_i(X_i) \geq v_i\left(X_j \setminus \{g\}\right)$$

Existence: EF1 always exists (round-robin). EFX existence open for $n \geq 3$

Maximin Share (MMS)

Definition:

Agent i's **maximin share** is:

$$\text{MMS}_i = \max_{\text{partitions } \mathcal{P} \text{ into } n \text{ bundles}} \min_{B \in \mathcal{P}} v_i(B)$$

Interpretation: Best guarantee if i divides goods and others choose first

MMS fairness: Each agent gets $\geq$ their MMS

Existence: Exact MMS may not exist; approximations (e.g., $2/3$-MMS) do

Fisher Market Setup

Ingredients:

- m divisible goods with total supply S_g (often normalized to 1)
- n agents with budgets $b_i > 0$ (no production, pure exchange)
- Agent utilities $u_i: R_+^m \to R$ (concave, monotone)

Market clears: At prices $p \in R_+^m$, each agent i chooses:

$$x_i \in \arg \max_{x_i : p \cdot x_i \leq b_i} u_i(x_i)$$

and $\sum_i x_{ig} = S_g$ for all goods g

Eisenberg-Gale (EG) Program

Assumption: Utilities homogeneous degree 1 (e.g., linear)

Convex Program:

$$\max_{x_i \geq 0} \sum_{i=1}^{n} b_i \log u_i(x_i) \quad \text{s.t.} \quad \sum_{i=1}^{n} x_{ig} \leq S_g, \forall g$$

KKT conditions give equilibrium prices.

For linear utilities: $u_i(x_i) = \sum_g v_{ig} x_{ig}$, EG is convex

Competitive Equilibrium from Equal Incomes (CEEI)

Setup: Set all budgets equal: $b_i = 1$ for all i

Solve Fisher market → equilibrium allocation and prices

Properties:

- Pareto efficient (First Welfare Theorem)
- Envy-free (equal budgets → equal purchasing power)
- Equal incomes fairness

With indivisibles: Exact CEEI may not exist; use approximate EF1

Fisher Market Equilibrium

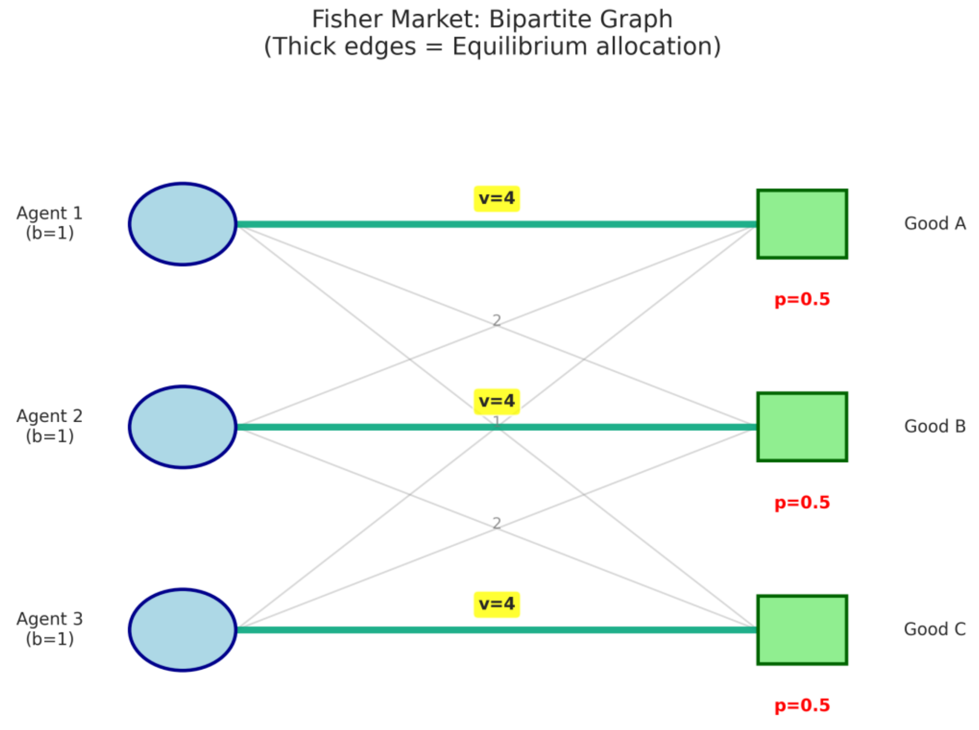

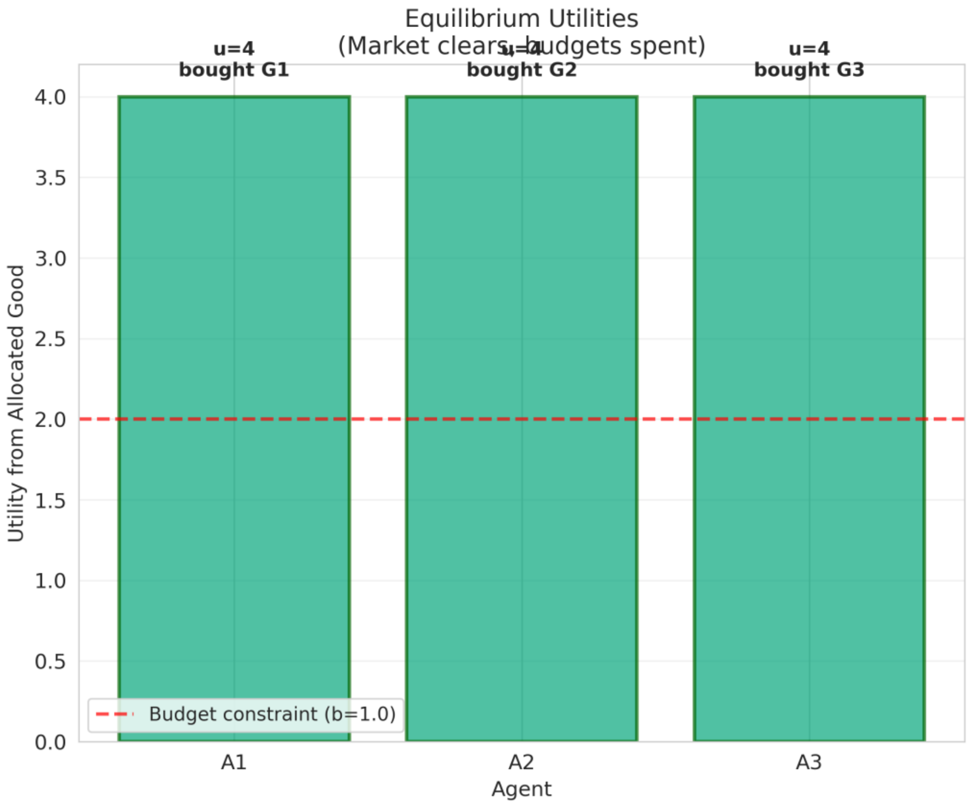

Nash Social Welfare and EG

Connection: For CEEI (equal budgets $b_i = 1$):

$$\text{EG program} = \max \sum_i \log u_i(x_i) = \max \log \prod_i u_i(x_i)$$

This is **Nash social welfare maximization!**

Property: NSW is scale invariant, respects Pigou-Dalton

Result: For indivisible goods with additive valuations, NSW maximization often yields EF1 + Pareto efficiency

Cooperative Game Theory

Setup: Cooperative game (N, v) where $v: 2^N \to R$ is characteristic function

Definition (Core):

$$\text{Core}(v) = \left\{ x \in R^n : \sum_i x_i = v(N),, \sum_{i \in S} x_i \geq v(S), \forall S \subseteq N \right\}$$

Interpretation: No coalition can improve by deviating

Properties:

- May be empty (e.g., some voting games)
- Non-empty for convex games, assignment games
- Forms a lattice in assignment games

Shapley Value

Definition:

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!, (n - |S| - 1)!}{n!} \text{\big}(v(S \cup \{i\}) - v(S)\text{\big})$$

Interpretation: Expected marginal contribution over all orderings

Axioms: Efficiency, symmetry, dummy, additivity

Uniqueness: Shapley value is the unique solution satisfying these axioms

Relation to core: Shapley value in core for convex games, may be outside otherwise

Shapley Value

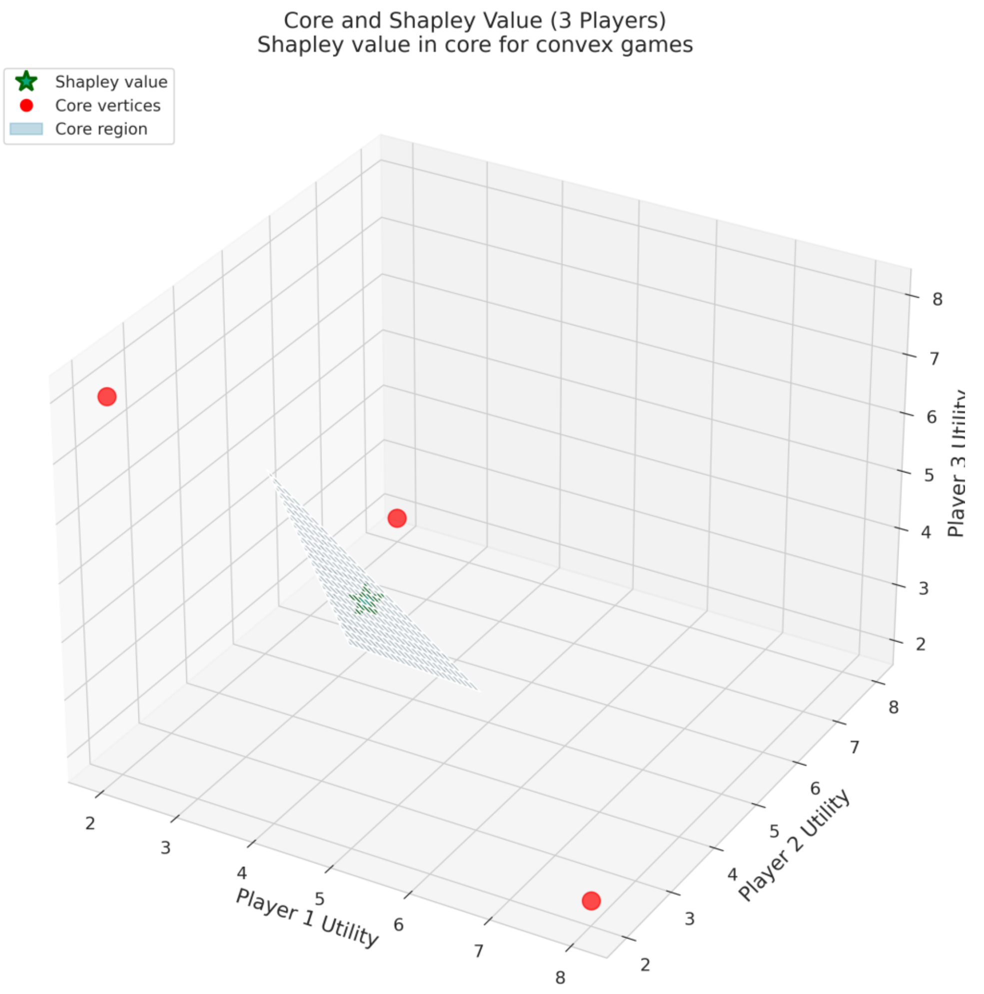

Nash Bargaining Solution (NBS)

Setup: Feasible set $\mathcal{F} \subset R^2$, disagreement point d

Definition:

$$\text{NBS} = \arg \max_{u \in \mathcal{F}, u \geq d} (u_1 - d_1)(u_2 - d_2)$$

Axioms:

- Pareto efficiency
- Symmetry
- Scale invariance
- Independence of irrelevant alternatives (IIA)

Uniqueness: NBS is unique solution satisfying these axioms

Nash Bargaining Solution (NBS)

Setup: Feasible set $\mathcal{F} \subset R^2$, disagreement point d

Definition:

$$\text{NBS} = \arg \max_{u \in \mathcal{F}, u \geq d} (u_1 - d_1)(u_2 - d_2)$$

Axioms:

- Pareto efficiency
- Symmetry
- Scale invariance
- Independence of irrelevant alternatives (IIA)

Uniqueness: NBS is unique solution satisfying these axioms

Kalai-Smorodinsky (KS) Solution

Setup: Same as NBS, plus utopia point $u^{\max} = \left(\max_{u \in \mathcal{F}} u_1 , \max_{u \in \mathcal{F}} u_2\right)$

Definition: Point on Pareto frontier satisfying:

$$\frac{u_1 - d_1}{u_1^{\max} - d_1} = \frac{u_2 - d_2}{u_2^{\max} - d_2}$$

Property: Drops IIA axiom, adds monotonicity w.r.t. utopia point

Comparison: Often more egalitarian than NBS when frontier concave

Nash Bargaining Solution

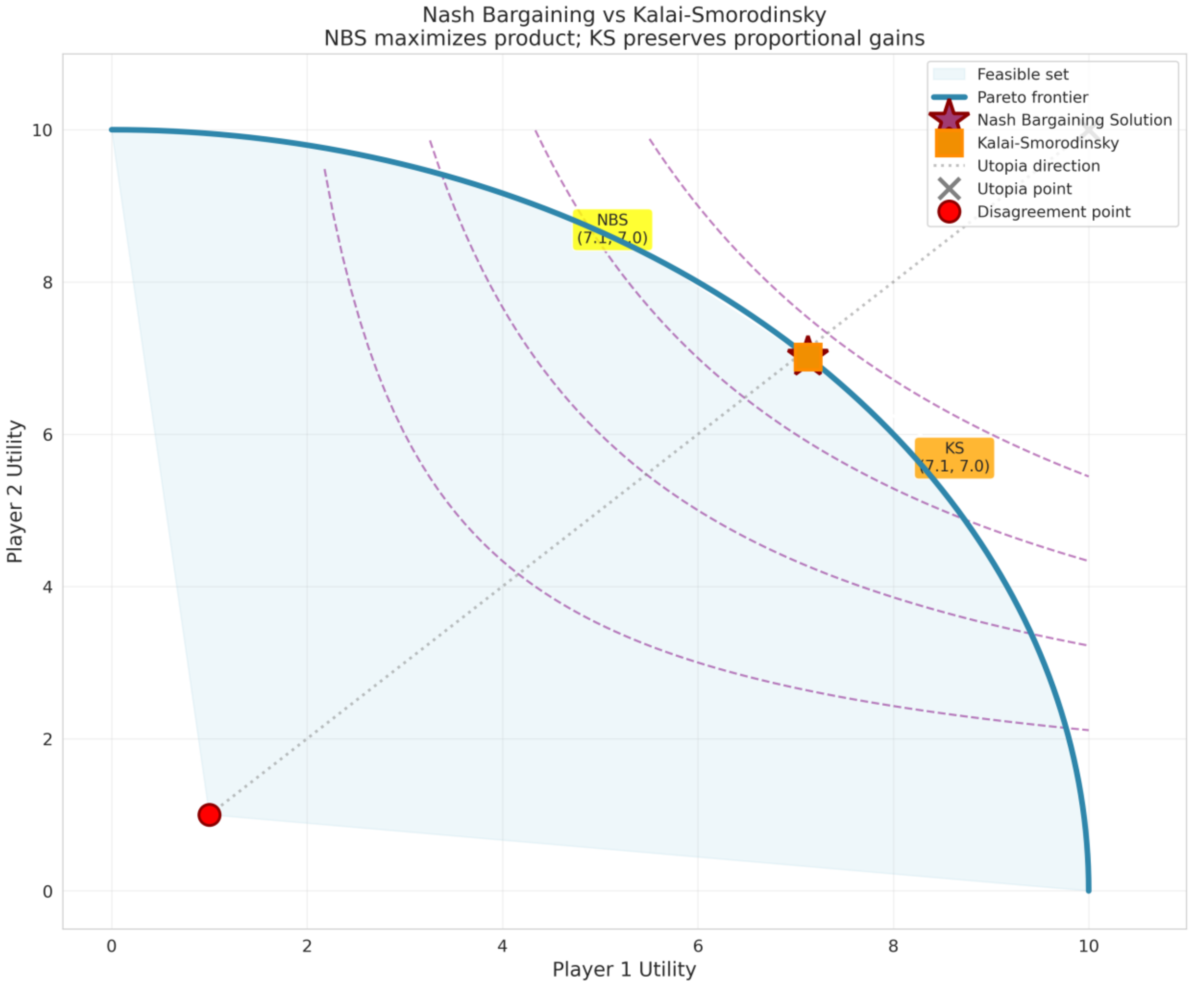

Game Theory

Lecture 10

Fairness, Pareto Optimality, and Social Welfare

1 Pareto Optimality

2 Fairness

3 Social Welfare

Social Choice Theory

Problem: How to aggregate individual preferences into collective decisions?

Settings:
- Elections and voting systems
- Committee decisions
- Resource allocation with heterogeneous preferences
- Constitutional design

Challenge: Individual preferences may conflict - need systematic aggregation method
Example: 3 candidates {A, B, C}, 3 voters with rankings:

- Voter 1: $A \succ B \succ C$
- Voter 2: $B \succ C \succ A$
- Voter 3: $C \succ A \succ B$

Question: Which candidate should win? No clear majority!

Social Choice Function

Definition:

A **social choice function (SCF)** maps preference profiles to outcomes:

$$f: \mathcal{P}^n \to X$$

where:

P = set of preference orderings over alternatives
number of voters/agents
$f(\succ_1, \ldots, \succ_n)$ = chosen alternative

Social welfare function (SWF):

Maps preference profiles to **complete ordering** of alternatives (more general than SCF)

Desirable Properties (Axioms)

Unanimity (Pareto):
If all voters prefer x to y, then social ranking should have x above y

Non-dictatorship:
No single voter determines outcome regardless of others' preferences

Independence of Irrelevant Alternatives (IIA):
Social ranking of x vs y depends only on individual rankings of x vs y, not on rankings involving third alternatives

Unrestricted domain:
SCF must work for **all** possible preference profiles (no restrictions)

Arrow's Impossibility Theorem (1951)

Theorem (Arrow):

For $|X| \geq 3$ alternatives, there exists **no** social welfare function satisfying:

1. Unanimity (Pareto)
2. Non-dictatorship
3. Independence of Irrelevant Alternatives (IIA)
4. Unrestricted domain

Implication: Cannot design "perfect" voting system with these properties!

Proof idea: IIA $+$ Unanimity forces "decisive sets" structure $\rightarrow$ minimal decisive set $=$ single voter (dictator)

Condorcet Paradox

Example demonstrating intransitivity:

- Voter 1: A ≻ B ≻ C
- Voter 2: B ≻ C ≻ A
- Voter 3: C ≻ A ≻ B

Pairwise majority votes:

- A vs B: $2-1$ for A (voters $1, 3$)
- B vs C: $2-1$ for B (voters $1, 2$)
- C vs A: $2-1$ for C (voters $2, 3$)

Result: Cycle $A \succ B \succ C \succ A$ (intransitive!)

Condorcet winner: Alternative that beats all others pairwise (may not exist)

Common Voting Rules

Plurality (First-Past-the-Post):
Each voter votes for one candidate; most votes wins

Borda count:
Rank m alternatives: top gets $m-1$ points, ..., bottom gets 0. Sum points across voters.

Instant Runoff (IRV/RCV):
Eliminate candidate with fewest first-place votes, redistribute; repeat

Approval voting:
Each voter approves any subset; most approvals wins

Pairwise (Condorcet method):
If Condorcet winner exists, select it; otherwise use tiebreaker

None is perfect! Each violates at least one desirable axiom.

Gibbard-Satterthwaite Theorem (1973)

Theorem (Gibbard, Satterthwaite):

For $|X| \geq 3$ alternatives, every social choice function that is:

- **Onto (surjective):** every alternative can win for some profile
- **Non-dictatorial**

is **manipulable** (not strategy-proof): some voter can benefit by misreporting preferences

Implication: Strategic voting is unavoidable in general voting systems!

Connection to mechanism design: Dominant-strategy truthfulness very restrictive

Proof Sketch: Gibbard-Satterthwaite

Steps:

1. Assume SCF f is strategy-proof (truthful) and onto
2. Define social welfare function from f via local independence
3. Show this SWF satisfies Arrow's axioms
4. By Arrow's theorem → dictatorship
5. But SCF onto + dictatorial → contradiction unless $|X| = 2$

Conclusion: Non-dictatorial onto SCF must be manipulable

Restricted Domains: Single-Peaked Preferences

Definition (Single-peaked):

Preferences are **single-peaked** if alternatives can be ordered on a line such that each voter has a unique "peak" with utility decreasing away from peak

Example: Left-right political spectrum

- Voter at position 3: prefers 3, then 4 or 2, then 5 or 1
- Preferences decrease as distance from peak increases

Theorem (Black 1948):

With single-peaked preferences, **median voter theorem** holds: median peak is Condorcet winner and strategy-proof

Implication: Arrow's and Gibbard-Satterthwaite can be **avoided** with domain restrictions!

Strategic Voting and Manipulation

Definition (Manipulation):

SCF is **manipulable** if voter i can benefit by reporting false preference $\succ_i'$ instead of true $\succ_i$:

$$f(\succ_i', \succ_{-i}) \succ_i f(\succ_i, \succ_{-i})$$

Example (Plurality voting):

- True preferences: $C \succ A \succ B$
- Predicted winner: B (plurality)
- Strategic vote for A (compromise) → A wins → better for voter than B

Practical implications:

- Polls influence strategic voting
- "Wasted vote" problem in plurality
- Approval voting reduces (but doesn't eliminate) manipulation

May's Theorem (1952)

Theorem (May):*

For **two alternatives**, majority rule is the **unique** social choice function satisfying:

1. Anonymity (all voters equal)
2. Neutrality (all alternatives treated equally)
3. Positive responsiveness (more support → better outcome)

Implication: Majority rule is optimal for binary decisions!

But: Multi-alternative case much harder (Arrow's and Gibbard-Satterthwaite)

Condorcet vs Borda Tradeoffs

Condorcet method:

- Respects pairwise majority preferences
- But Condorcet winner may not exist (cycles)
- Sensitive to strategic nomination (IIA violation)

Borda count:

- Always produces winner (no cycles)
- Considers full preference rankings (not just top choice)
- But vulnerable to clones and manipulation
- Violates IIA: adding irrelevant candidate changes outcome

Example where they differ:

- 3 voters: $A \succ B \succ C, B \succ C \succ A, C \succ A \succ B$
- Borda: All tied at 3 points each
- Condorcet: Cycle (no winner)

Computational Social Choice

Modern questions:

1. **Winner determination:** Given preference profile, who wins? (Complexity: P to NP-hard depending on rule)
2. **Manipulation detection:** Can voter i manipulate? (NP-hard for many rules)
3. **Bribery:** Cost to change outcome by bribing voters? (Often NP-hard)
4. **Control:** Can chair manipulate by adding/removing candidates or voters?

Computational hardness as barrier:

- If manipulation is NP-hard, maybe it won't happen in practice
- But: heuristics, approximations, learning may still enable manipulation

Liquid Democracy

Concept: Hybrid of direct and representative democracy

Mechanism:
- Each voter can either vote directly OR delegate to another voter
- Delegations are transitive $(A \rightarrow B \rightarrow C)$
- Voter can override delegation on specific issues

Advantages:
- Flexibility: expertise-based delegation
- Reduced voter burden on complex issues

Challenges:
- Cycle prevention in delegation graph
- Power concentration ("super delegates")
- Strategic delegation behavior

Connections to Mechanism Design

Voting as special case of mechanism design:

- Agents have private types (preferences)
- Mechanism selects outcome based on reports
- Goal: Aggregate preferences "well" while incentivizing truthfulness

VCG mechanisms:

- Strategy-proof for multi-dimensional preferences (valuations)
- But: Require transferable utility (money)
- Voting: No transfers, only ordinal preferences

Key difference: Voting harder due to lack of transfers and ordinal-only preferences

Result: General voting more limited than mechanism design with payments

Price of Fairness

Motivation: Fairness constraints may reduce efficiency

Definition:

For fairness notion $\mathcal{F}$ and welfare function W:

$$\mathrm{PoF}(\mathcal{F}) = \inf_{\text{instances}} \frac{\max\{W(u) : u \in \mathcal{F} \cap \mathcal{U}\}}{\max\{W(u) : u \in \mathcal{U}\}}$$

Interpretation: Worst-case ratio of fair welfare to unconstrained optimal welfare

Goal: Bound PoF, design fair mechanisms with small PoF

Examples of PoF

Indivisible allocation:

- EF1 constraint: PoF can be $\Omega(n)$ in worst case
- MMS: Similar bounds

Routing games:

- Max-min fairness vs utilitarian: PoF depends on network topology

Auctions:

- Budget constraints + fairness: Can significantly reduce revenue

Price of Fairness Tradeoff

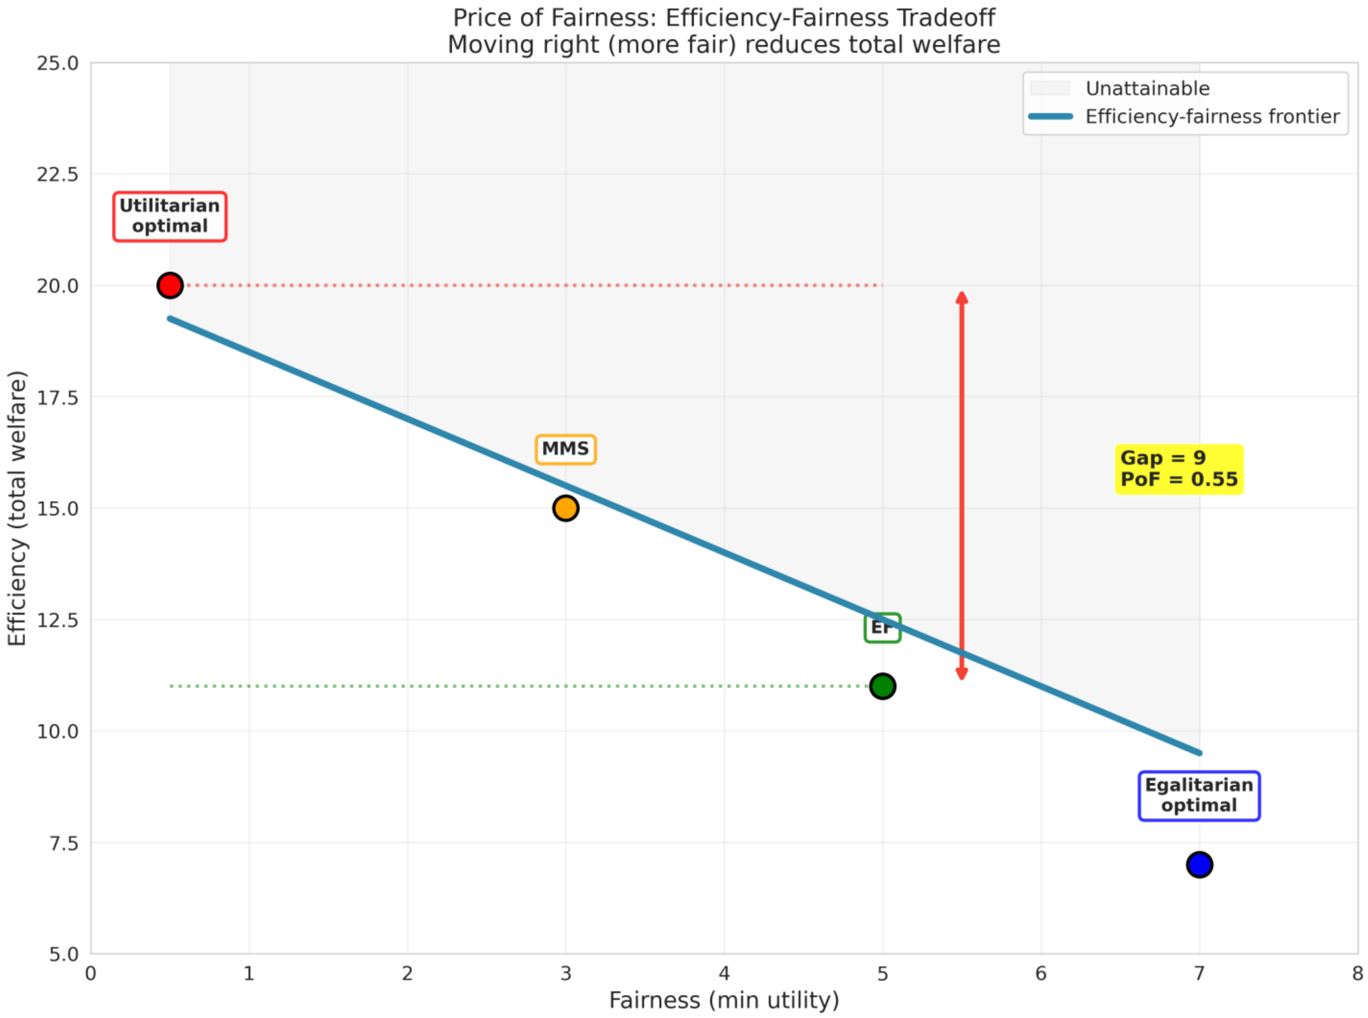

Efficiency-Equity Frontiers

Concept: Plot achievable $\left(W^{\text{egal}}, W^{\text{util}}\right)$ pairs

Frontier: Upper boundary of feasible region

Interpretation: Visualizes tradeoff between total welfare and equality

Applications:
- Network utility maximization
- Fair queuing in networks
- Resource allocation under fairness constraints

Gini Index

Definition:

$$\text{Gini}(u) = \frac{\sum_i \sum_j |u_i - u_j|}{2n \sum_i u_i}$$

Range: $[0,1]$, where $0 \;= $ perfect equality

Properties: Widely used in economics, intuitive

Atkinson Index

Definition:

$$\mathrm{Atk}_{\epsilon}(u) = 1 - \frac{\left(\frac{1}{n}\sum_i u_i^{1-\epsilon}\right)^{\frac{1}{1-\epsilon}}}{\frac{1}{n}\sum_i u_i}$$

Parameter: $\epsilon > 0$ controls inequality aversion

Properties: Axiomatically grounded, respects welfare principles

Theil Index

Definition:

$$\text{Theil}(u) = \frac{1}{n}\sum_i \frac{u_i}{\bar{u}} \log \frac{u_i}{\bar{u}}$$

where $\bar{u} = \frac{1}{n}\sum_i u_i$

Properties: Decomposable across subgroups, related to entropy

Inequality Measures Comparison

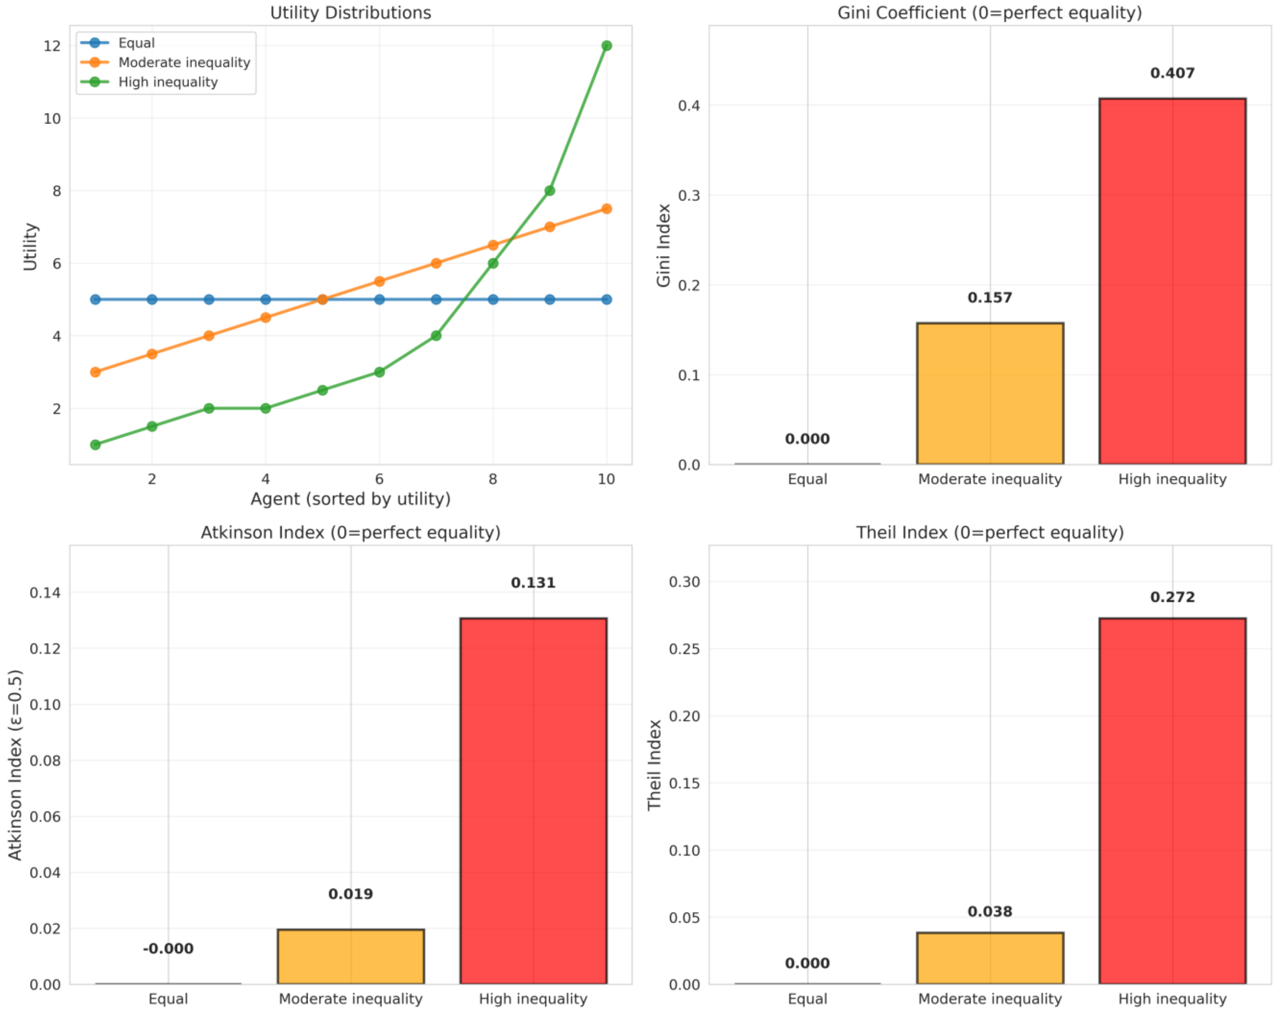

Fair Division: Cake Cutting

Setup: Continuous resource $[0,1]$ with nonatomic valuations

Procedures:

- **Cut and choose** (2 agents): EF and Pareto efficient
- **Dubins-Spanier moving knife** (n agents): Proportional
- **Selfridge-Conway** (3 agents): EF with discrete steps

For general n: Exact EF with finite steps is hard; relaxations exist

Rent Division

Setup: Rooms with base valuations u_{ir}, find prices p_r and assignment

Constraint: Envy-free if assigned room r satisfies:

$$u_{ir} - p_r \geq u_{i\ell} - p_\ell \quad \forall \ell$$

Solution: Linear programming or max-weight matching with price adjustments

Existence: Always exists with quasilinear utilities

Network Routing Fairness

Objectives:

- **Utilitarian:** Minimize total latency $\sum_i \ell_i(f_i)$
- **Max-min fair:** Maximize minimum utility
- **Proportional fair:** Maximize $\sum_i \log u_i$ (equivalent to EG)

Methods: Convex optimization, KKT conditions give prices

Tradeoff: Max-min can reduce throughput significantly

Scheduling Fairness

Setup: Allocate machine time to jobs

Objectives:

- **Utilitarian:** Minimize $\sum_i C_i$ (completion times)
- **Egalitarian:** Minimize $\max_i C_i$ (makespan)

Fair sharing: Processor sharing, weighted fair queuing

Result: Approximate proportional fairness in steady state

Mechanism Design with Fairness

Challenge: Combine incentive compatibility (IC) with fairness

VCG: Efficient and IC, but may be unfair (concentrated payments)

Fair VCG variants: Add fairness constraints or fairness terms in objective

Tradeoff: Strong fairness + IC + budget balance often impossible (impossibility results)

Fairness with Private Information

Setup: Agents have private valuations

Goals: IC mechanism that is also fair

Notions:

- **Ex ante EF:** EF in expectation before types realized
- **Interim EF:** EF given each agent's type
- **Ex post EF:** EF after all types realized

Result: Ex post EF + IC is very restrictive; use ex ante or approximate variants

Exercise 1

Setup: Payoff matrix (row, col):

	L	R
U	$(4, 1)$	$(0, 3)$
D	$(1, 0)$	$(2, 2)$

Tasks:

1. List all pure strategy profiles and payoffs
2. Compute mixed strategy payoffs for p (prob of U) and q (prob of L)
3. Plot feasible region in (u_1, u_2) space
4. Identify Pareto frontier (convex hull)
5. Find Nash equilibria and check if Pareto efficient

Exercise 2

Setup: Feasible set $\mathcal{U} = \{(u_1, u_2): u_1^2 + u_2^2 \leq 100, u_i \geq 0\}$

Tasks:

1. Maximize utilitarian: $u_1 + u_2$
2. Maximize egalitarian: $\min(u_1, u_2)$
3. Maximize Nash: $u_1 \cdot u_2$
4. Compare optimal points. Which is most equal? Most efficient?
5. Plot contours and optimal points

Exercise 3

Setup: 3 agents, 3 goods. Valuations:

$$V = \begin{pmatrix} 4 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 4 \end{pmatrix}$$

Tasks:

1. Allocation $(1,2,3)$: Is it EF? Check all pairs
2. Find an EF allocation (may need mixed allocations for divisible)
3. Compute utilitarian welfare for your EF allocation
4. Is there a more efficient allocation? If so, is it still EF?

Exercise 4

Setup: 2 agents, 2 goods. Linear utilities:

- Agent 1: $u_1 = 2x_{11} + x_{12}$, budget $b_1 = 1$
- Agent 2: $u_2 = x_{21} + 2x_{22}$, budget $b_2 = 1$
- Supply: 1 unit of each good

Tasks:

1. Write Eisenberg-Gale program
2. Solve for equilibrium allocation and prices
3. Verify market clears and budgets spent
4. Check allocation is EF and Pareto efficient

Exercise 5

Setup: 3-player game with characteristic function:

$$v(\{1,2,3\}) = 12v(\{1,2\}) = 7, v(\{1,3\}) = 7, v(\{2,3\}) = 7v(\{1\}) = 2, v(\{2\}) = 2, v(\{3\}) = 2$$

Tasks:

1. Compute Shapley value for each player
2. Verify efficiency: $\sum_i \phi_i = v(N)$
3. Check if Shapley value is in the core
4. Find core constraints and determine if core is non-empty

Summary

- **Pareto efficiency:** No waste, but NE may not be efficient
- **Social welfare functions:** Utilitarian, egalitarian, Nash, weighted - different priorities
- **Fairness:** EF, proportionality, EF1, EFX, MMS - hierarchy of notions
- **Fisher markets:** EG program computes CEEI, which is EF + efficient
- **Nash social welfare:** Balances efficiency and equity, scale invariant
- **Cooperative solutions:** Core, Shapley value, Nash bargaining, KS
- **Social choice theory:** Arrow's impossibility, Gibbard-Satterthwaite, voting systems, strategic manipulation unavoidable
- **Price of fairness:** Quantifies efficiency loss from fairness constraints

Course Textbooks

- Bonanno, G. (2024). ***Game Theory (3rd ed.)***. University of California, Davis. Received from: GT Book
- Axelrod, R. (1984). ***The Evolution of Cooperation***. Basic Books. Received from: Axelrod Article
- Nisan, N., Roughgarden, T., Tardos, É., & Vazirani, V. V. (2007). ***Algorithmic Game Theory***. Cambridge University Press. Received from: AGT Book
- Myerson, R. B. (1991). ***Game Theory: Analysis of Conflict.*** Harvard University Press. Received from: GT Book 2F.
- Christianos et al., ***Multi-Agent Reinforcement Learning: Foundations and Modern Approaches***, 2023. Received from: MARL Book.pdf
- Shoham, Y., & Leyton-Brown, K. (2008). ***Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations***. Cambridge University Press. Received from: MARL Book.pdf
- `***nashpy***` documentation (readthedocs). Link: NashPy Docs

That's All for Today!