CSCE 631-D: Intelligent Agents: Computational Game Solving

Lecture 4: Dominance, Rationalizability, and Correlated Equilibria

Alan Kuhnle

Fall 2026

Texas A&M University

Today's Plan

Where we are. Lectures 1–3 introduced strategic-form games, mixed strategies, and Nash equilibrium. Today we add two ideas: ways to simplify games (dominance and iterated elimination) and a richer equilibrium concept (correlated equilibrium) that

Review of Maxmin

Maxmin and Security Level (Review)

Recall from Lecture 3:

Maxmin Value

$$\overline{v}_i = \max_{\sigma_i \in \Delta(S_i)} \min_{\sigma_{-i} \in \Delta(S_{-i})} u_i(\sigma_i, \sigma_{-i}).$$

- $\overline{v}_i$ is the **security level**: the payoff player i can guarantee regardless of opponents.
- A maxmin strategy is "pessimistic": it assumes the worst.
- In zero-sum games: maxmin strategy $=$ NE strategy.
- In general-sum games: maxmin $\neq$ NE in general; maxmin is typically more conservative.

Individual Rationality

A natural question: can a Nash equilibrium ever force a player below the payoff they could guarantee on their own? The answer is no.

Definition

A payoff vector $(v_1, \ldots, v_n)$ is **individually rational** if $v_i \geq \overline{v}_i$ for all i.

Proposition

In any Nash equilibrium σ^*, $u_i(\sigma^*) \geq \overline{v}_i$ for all i.

Proof: In a NE, σ_i^* is a best response. Its payoff cannot be below the maxmin value, since player i could deviate to the maxmin strategy and guarantee $\overline{v}_i$. $\square$

Takeaway. NE payoffs are always individually rational. The converse is false: not every individually rational payoff is a NE.

Dominated Strategies

Strict Dominance

Before analyzing equilibria, we can often simplify a game by eliminating strategies that a rational player would never choose.

Definition

Strategy s_i **strictly dominates** s_i' if for all $s_{-i} \in S_{-i}$:

$$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}).$$

A strategy is **strictly dominated** if some other strategy strictly dominates it.

Proposition

A rational player never plays a strictly dominated strategy. Moreover, no strictly dominated strategy is in the support of any Nash equilibrium.

Proof: If s_i' is strictly dominated by s_i, then for any $s_{-i} \in S_{-i}$: $u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$.

Weak Dominance

Strict dominance requires improvement against *every* opponent profile. A weaker notion only requires "at least as good always, strictly better sometimes."

Definition

Strategy s_i **weakly dominates** s_i' if for all $s_{-i} \in S_{-i}$:

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}),$$

with strict inequality for at least one s_{-i}.

- Weakly dominated strategies *can* appear in NE (unlike strictly dominated ones).
- Whether to eliminate weakly dominated strategies is debatable; it can change the set of NE.
- Weak dominance is important in mechanism design (e.g., "dominant-strategy

Dominance by Mixed Strategies

Checking only pure strategies for dominance is not enough. A pure strategy s_i can be strictly dominated by a **mixed strategy** σ_i, even if no single pure strategy dominates it.

Example

	L	R
T	3	0
M	0	3
B	1	1

Neither T nor M dominates B. But the mixed strategy $\sigma_1 = (1/2, 1/2, 0)$ gives expected payoff $3/2$ against both L and R, strictly dominating B (which gives 1).

Takeaway. When checking for dominance, we must consider mixed strategies, not just pure ones.

Iterated Elimination of Dominated Strategies

IEDS: The Idea

If all players are rational and this is common knowledge, we can iteratively simplify the game by removing strategies that no rational player would ever use.

Iterated Elimination of Strictly Dominated Strategies (IEDS)

1. Find and remove all strictly dominated strategies.
2. In the reduced game, repeat.
3. Continue until no dominated strategies remain.

Key property: IEDS never eliminates a Nash equilibrium.

Why does iterative reasoning work?

- A rational player never plays a strictly dominated strategy.
- If all players are rational, and this is *common knowledge*, then iterative reasoning applies.
- Round k of elimination corresponds to k levels of "I know that you know that I know..." rationality.

IEDS: Example

We apply IEDS to the following 3×3 game, checking each player's strategies systematically.

	L	C	R
T	$(1,0)$	$(1,2)$	$(0,1)$
M	$(0,3)$	$(0,1)$	$(2,0)$
B	$(0,1)$	$(2,0)$	$(3,1)$

Round 1: No pure strategy strictly dominates another for either player. For Player 1, B weakly dominates M ($u_1(B,c) \geq u_1(M,c)$ for every column, with equality at $c = L$), but strict dominance fails because $u_1(B,L) = 0 = u_1(M,L)$.

However, as in the Dominance by Mixed Strategies example earlier, a *mixed* strategy can strictly dominate a pure one. Consider the mixture $\sigma_1 = \frac{1}{4}T + \frac{3}{4}B$. Its expected payoffs against each column are

$$\tfrac{1}{4}(1) + \tfrac{3}{4}(0) = \tfrac{1}{4}, \quad \tfrac{1}{4}(1) + \tfrac{3}{4}(2) = \tfrac{7}{4}, \quad \tfrac{1}{4}(0) + \tfrac{3}{4}(3) = \tfrac{9}{4},$$

while M yields (0, 0, 2). Every component is strictly larger ($\frac{1}{4} > 0$, $\frac{7}{4} > 0$, $\frac{9}{4} > 2$), so σ_1 strictly dominates M. Eliminate M.

IEDS: Example (cont.)

After eliminating M:

	L	C	R
T	$(1, 0)$	$(1, 2)$	$(0, 1)$
B	$(0, 1)$	$(2, 0)$	$(3, 1)$

Round 2: Player 2's payoffs: $L = (0, 1)$, $C = (2, 0)$, $R = (1, 1)$. R weakly dominates L ($(1, 1) \geq (0, 1)$, strict in first coordinate). If we do strict IEDS, L is not strictly dominated.

But: the mixed strategy $\frac{1}{2}C + \frac{1}{2}R$ gives $(3/2, 1/2)$, which strictly dominates L's $(0, 1)$? Check: $3/2 > 0$ yes, $1/2 < 1$ no. So L survives.

IEDS terminates here. The surviving game may still have multiple NE.

IEDS: Formal Properties

Three properties make IEDS a reliable tool for simplifying games.

Theorem

1. **Order independence:** The set of strategies surviving IEDS (strict dominance) is the same regardless of the elimination order.
2. **NE preservation:** Every NE of the original game survives IEDS.
3. **Solvability:** If IEDS reduces the game to a single strategy profile, that profile is the unique NE.

Warning: Order independence fails for iterated elimination of *weakly* dominated strategies. The set of surviving strategies can depend on which order you eliminate.

Rationalizability

Rationalizability

IEDS removes strategies that are dominated. A related but distinct question: which strategies can be justified by *some* reasonable belief about the opponents?

Definition (Bernheim, 1984; Pearce, 1984)

A strategy s_i is **rationalizable** if it survives iterated elimination of strategies that are never a best response.

Equivalently: s_i is rationalizable if there exist beliefs σ_{-i} (a probability distribution over opponents' strategies) such that s_i is a best response, and the beliefs are "consistent" (they assign positive probability only to rationalizable strategies of opponents).

Theorem

For two-player games: rationalizable strategies = strategies surviving IEDS.

For $n > 2$ players: rationalizability is (weakly) more permissive than IEDS.

Relationship: Dominance, Rationalizability, NE

The following diagram summarizes the containment relationships. Each level is a subset of the one above it.

All Strategies

$\downarrow \subseteq$

Strategies surviving IEDS

$\downarrow =$ (2-player)

Rationalizable Strategies

$\downarrow \subseteq$

Nash Equilibria

Every NE strategy is rationalizable, so NE strategies form a subset of rationalizable strategies. Not every rationalizable strategy is part of a NE.

Correlated Equilibrium

Motivation: Beyond Nash

We have so far been refining the set of strategies a rational player might use. We now ask a different question: can we expand the set of achievable outcomes by allowing coordination through a mediator?

We now turn to the central new concept of this lecture. Nash equilibrium has limitations:

- Multiple NE $\Rightarrow$ coordination problem.
- Mixed NE require players to randomize independently.
- The NE payoff can be Pareto-dominated by achievable outcomes.

Idea (Aumann, 1974): allow players to receive **correlated signals** from a trusted mediator before choosing actions.

The Mediator / Signal Device

Think of a traffic light at an intersection: it privately tells each driver what to do, and each driver follows because deviating would be worse. This is the structure of correlated equilibrium.

Setup

1. A **mediator** draws a strategy profile $s = (s_1, \ldots, s_n)$ from a joint distribution $p \in \Delta(S)$.
2. The mediator privately tells each player i their recommended action s_i (but *not* others' actions).
3. Each player decides whether to follow the recommendation.

A **correlated equilibrium** is a distribution p such that every player finds it optimal to **follow** the recommendation, given their signal.

Correlated Equilibrium: Formal Definition

We now state the CE condition precisely. The key requirement is **obedience**: no player gains by deviating from the mediator's recommendation.

Definition (Aumann, 1974)
A probability distribution $p \in \Delta(S_1 \times \cdots \times S_n)$ is a **correlated equilibrium** (CE) if for every player i and every pair of strategies $s_i, s_i' \in S_i$:

$$\sum_{s_{-i} \in S_{-i}} p(s_i, s_{-i}) \big[u_i(s_i, s_{-i}) - u_i(s_i', s_{-i}) \big] \geq 0.$$

Interpretation: Conditional on being told to play s_i, player i has no incentive to deviate to any other s_i'.

The condition can be rewritten, for each s_i with $p(s_i) > 0$ and all i, s_i':

$$\mathbb{E}_{s_{-i} \sim p(\cdot|s_i)}[u_i(s_i, s_{-i})] \geq \mathbb{E}_{s_{-i} \sim p(\cdot|s_i)}[u_i(s_i', s_{-i})].$$

The left-hand side sums over all profiles consistent with the recommendation s_i; this is equivalent to conditioning on the signal.

Example: Traffic Light as a Correlated Equilibrium

To see CE in action, consider the **Game of Chicken** (two cars at an intersection):

	Go	Stop
Go	$(-10, -10)$	$(5, 0)$
Stop	$(0, 5)$	$(0, 0)$

NE: (Go, Stop) and (Stop, Go), each with payoff sum 5.

Mixed NE exists but has positive probability of the crash outcome.

Traffic light (mediator): signal (Go, Stop) with prob $1/2$ and (Stop, Go) with prob $1/2$.

Check CE condition: Told "Go" $\Rightarrow$ know opponent stops $\Rightarrow$ $u(\text{Go}) = 5 > u(\text{Stop}) = 0$. ✓

Told "Stop" $\Rightarrow$ know opponent goes $\Rightarrow$ $u(\text{Stop}) = 0 > u(\text{Go}) = -10$. ✓

Expected payoff: $(5/2, 5/2)$, **better than any single NE payoff for the disadvantaged player**.

Question: Is the uniform distribution over all four outcomes a CE of Chicken? Why or why not?

Another CE: Game of Chicken

The previous CE only mixed over Nash equilibrium outcomes. Can a CE place positive weight on a non-NE outcome like (Stop, Stop)?

The following distribution does exactly that while preserving obedience:

$p(s_1, s_2)$	Go	Stop
Go	0	1/3
Stop	1/3	1/3

Check: If told "Go," player knows opponent plays Stop (only possibility).
$u(\text{Go}, \text{Stop}) = 5 > u(\text{Stop}, \text{Stop}) = 0$. ✓

If told "Stop," opponent plays Go with prob $1/2$, Stop with prob $1/2$.

$$u(\text{Stop}) = \tfrac{1}{2}(0) + \tfrac{1}{2}(0) = 0.$$
$$u(\text{Go}) = \tfrac{1}{2}(-10) + \tfrac{1}{2}(5) = -5/2 < 0.$$

✓ No incentive to deviate.

Expected payoff: $(5/3, 5/3)$; total welfare $10/3$.

This distribution is outside the convex hull of NE *distributions*, even though its payoff lies inside the convex hull of NE payoffs.

Shared Prompts as Correlating Devices

With the CE definition and examples in hand, we can connect this concept to multi-agent LLM systems, where two forms of correlation arise:

Private Mediator (CE Model)
A mediator draws a joint recommendation and sends each agent its action **privately**. Each agent observes only its own signal.

Shared Public Context
A common system prompt or orchestrator message is observed by **all** agents. Every agent sees the same context.

Both correlate play, but the information structures differ:

- Private recommendations support CE: each agent conditions on its own signal, not others'.
- A shared prompt is common knowledge; agents coordinate, but the deviation incentives change because each agent knows the full signal.

Correlated Equilibrium as a Linear Program

CE as an LP

A major computational advantage of CE over NE: the CE constraints are **linear** in the distribution p, so finding a CE reduces to solving a linear program.

Linear Program for CE

$$\text{Find } p \in \Delta(S) \text{ such that:}$$

$$\sum_{s_{-i}} p(s_i, s_{-i})[u_i(s_i, s_{-i}) - u_i(s_i', s_{-i})] \geq 0 \quad \forall\, i,\ \forall\, s_i,\ s_i' \in S_i.$$

- Variables: $p(s)$ for each $s \in S$ ($\prod_i |S_i|$ variables).
- Constraints: $\sum_i |S_i|(|S_i| - 1)$ incentive constraints, plus $p \geq 0$ and $\sum_s p(s) = 1$.
- **All constraints are linear** $\Rightarrow$ CE is an LP feasibility problem.

Can optimize any linear objective over the CE polytope (e.g., maximize social welfare, fairness, etc.). Computation of CE and its variants is deepened in Lecture 8 (M2).

CE vs. NE: Computational Comparison

The following table highlights the structural and computational differences between NE and CE.

	Nash Equilibrium	Correlated Equilibrium
Definition	Fixed point of best response	Linear incentive constrain
Structure	Nonlinear (product of independent marginals)	Linear (polytope)
Computation	PPAD-complete	Polynomial (LP)
Set of equilibria	Finite (generically)	Convex polytope
Uniqueness	Rarely	Rarely

Key insight. CE is a **convex polytope** (intersection of half-spaces) while NE is the intersection of polynomial constraints. This is why CE is computationally tractable.

Optimizing over CE

Because the CE set is a polytope, we can optimize any linear objective over it. Given a finite game:

Maximize social welfare:

$$\max_{p \in \mathrm{CE}} \sum_{s \in S} p(s)\big[u_1(s) + u_2(s)\big].$$

Maximize fairness (egalitarian):

$$\max_{p \in \mathrm{CE}} \min_{i \in N} \sum_{s \in S} p(s)\, u_i(s).$$

Maximize Player *i*'s payoff:

$$\max_{p \in \mathrm{CE}} \sum_{s \in S} p(s)\, u_i(s).$$

All are linear programs (the egalitarian version needs a standard transformation with an auxiliary variable).

Relationship Between NE and CE

Every NE Is a CE

We claimed that CE generalizes NE. Here is the precise statement.

Theorem
If $\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*)$ is a Nash equilibrium, then the product distribution $p(s) = \prod_{i \in N} \sigma_i^*(s_i)$ is a correlated equilibrium.

Proof: In a NE, σ_i^* is a best response to σ_{-i}^*. Under the product distribution, the conditional distribution over s_{-i} given s_i is exactly σ_{-i}^* (by independence). So the CE incentive constraint reduces to the NE best-response condition. □

Thus: NE $\subseteq$ CE (viewing NE as product distributions).

The inclusion is generally **strict**: CE allows correlated randomization that NE does not.

The CE Polytope Contains the NE Convex Hull

The following diagram illustrates the relationship in payoff space. Each NE is a point inside the CE polytope; their convex hull is contained in but generally smaller than the full polytope.

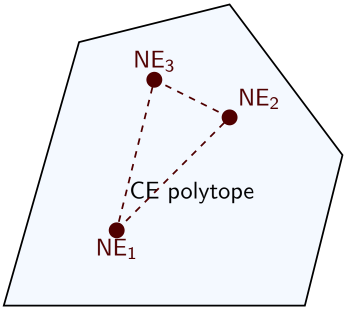

Schematic projection in distribution space. The CE polytope contains every NE product distribution and therefore their convex hull; the inclusion can be strict.

Aumann's Result

We can now summarize the structural properties of the CE set.

Theorem (Aumann, 1974)

The set of correlated equilibria of a finite game is:

1. A **convex polytope** (intersection of linear constraints).
2. A **superset** of the convex hull of Nash product distributions.
3. Non-empty (since every finite game has a NE, hence a CE).

Aumann's broader philosophical point: if players share a common prior and have common knowledge of rationality, the appropriate solution concept is CE, not NE.

NE implicitly assumes **independent** randomization. CE relaxes this to allow any joint distribution consistent with incentive compatibility.

Comparison of Solution Concepts

Coarse Correlated Equilibrium (CCE)

Before comparing solution concepts, we define the last equilibrium notion in the hierarchy.

Definition

A distribution $p \in \Delta(S)$ is a **coarse correlated equilibrium** (CCE) if for every player i and every deviation $s_i' \in S_i$:

$$\sum_{s \in S} p(s)\, u_i(s_i', s_{-i}) \;\le\; \sum_{s \in S} p(s)\, u_i(s).$$

Interpretation. In a CCE, a player must decide whether to follow or deviate *before* observing the recommendation. Unlike CE, the player cannot condition their deviation on what was recommended.

Check Your Understanding

Question. Every CE is also a CCE (since the CE constraints are stronger). Is the converse true: is every CCE also a CE? Why or why not?

Solution Concepts: Two Related Hierarchies

We have now seen several solution concepts. They organize into two separate hierarchies that should not be conflated, because they concern different mathematical objects: one is about *distributions over profiles*, the other about *sets of pure strategies*.

Equilibrium distributions

$$\text{dominant-strategy equilibria} \subseteq \text{NE} \subseteq \text{CE} \subseteq \text{CCE}.$$

A strategy profile is a **dominant-strategy equilibrium** if each player's strategy weakly dominates all alternatives. CE permits private correlated recommendations; CCE permits only deviations chosen before the recommendation is observed.

Rationality restrictions on strategies

$$\{\text{strategies used in some NE}\} \subseteq \{\text{rationalizable strategies}\}.$$

For two-player games, rationalizable strategies are those surviving iterated elimination of strictly dominated strategies (allowing mixed dominance).

These hierarchies are related through NE, but neither CCE distributions nor rationalizable strategies contain the other: they are different object types.

When to Use Which Concept?

Choosing a solution concept depends on what assumptions you make about the players and the computational setting.

Concept	When appropriate
Dominant strategy	Mechanism design (goal: make truth-telling dominant)
IEDS	When common knowledge of rationality is plausible
Nash equilibrium	Standard prediction; independently deployed agents
CE	When a mediator (traffic light, protocol, orchestrator) exists
CCE	When agents learn via no-regret dynamics (Module 3)
Maxmin / security	Adversarial settings; worst-case guarantees
Multi-agent LLM	CE when an orchestrator or shared prompt coordinates agents; NE when agents are independently deployed

The choice of solution concept often depends on the **computational** question: can we compute it, learn it, or enforce it?

Computational Complexity Summary

The following table collects the computational complexity results for each solution concept we have studied.

Concept	Repr.	Exact/Approx	Complexity
Dom.-strat. eq.	Explicit	Exact	Poly
IEDS	Explicit	Exact	Poly
NE (2-player, zero-sum)	Expl. matrix	Exact	**P** (LP)
NE (2-player, general)	Expl. bimatrix	Exact	PPAD-complete
CE	Explicit	Exact	**P** (LP)
CCE	Explicit	Exact	**P** (LP)
Maxmin	Expl. bimatrix	Exact	**P** (LP)

Main contrast. Finding a general-sum NE is PPAD-complete, while CE and CCE can both be computed in polynomial time via linear programming. Structured games

Extended Examples

CE of Battle of the Sexes

Payoffs (P1 prefers Opera, P2 prefers Football):

	Opera	Football
Opera	$(3, 2)$	$(0, 0)$
Football	$(0, 0)$	$(2, 3)$

CE distribution:

p	Opera	Football
Opera	p_{OO}	p_{OF}
Football	p_{FO}	p_{FF}

CE constraints for Player 1:

$$\text{If told Opera:} \quad 3p_{OO} + 0 \cdot p_{OF} \geq 0 \cdot p_{OO} + 2p_{OF} \iff 3p_{OO} \geq 2p_{OF}.$$

$$\text{If told Football:} \quad 0 \cdot p_{FO} + 2p_{FF} \geq 3p_{FO} + 0 \cdot p_{FF} \iff 2p_{FF} \geq 3p_{FO}.$$

Similarly for Player 2:

$$2p_{OO} \geq 3p_{FO}, \qquad 3p_{FF} \geq 2p_{OF}.$$

Plus: $p_{OO} + p_{OF} + p_{FO} + p_{FF} = 1$, all $p \geq 0$.

CE of Battle of the Sexes: Optimal

We can now optimize over the CE polytope defined by the constraints on the previous slide.

Maximizing social welfare $\sum_s p(s)[u_1(s) + u_2(s)]$:

$$\max \; 5p_{OO} + 5p_{FF} \quad \text{s.t. CE constraints.}$$

One symmetric welfare-optimal solution is $p_{OO} = p_{FF} = 1/2$, $p_{OF} = p_{FO} = 0$.

This is the "fair coin flip" CE: go to Opera together or Football together, each with prob $1/2$.

Expected payoffs: $u_1 = 5/2$, $u_2 = 5/2$.

Compare to the three NE payoffs:

- $(3, 2)$ at (Opera, Opera).
- $(2, 3)$ at (Football, Football).
- $(6/5, 6/5)$ at the mixed NE.

The CE achieves higher total payoff than the mixed NE and is **fairer** than either pure NE.

Check Your Understanding

Question. Can you find a CE of Battle of the Sexes with total expected payoff greater than 5?

Coordination in Multi-Agent LLM Systems

Coordination in Multi-Agent LLM Pipelines

We return to the question of how CE applies to LLM systems. Common multi-agent LLM architectures serve as correlation devices:

Debate: Two agents argue opposing sides. The protocol (turn order, judge, scoring rule) is the correlating device.

Orchestrator-worker: A central agent assigns roles and aggregates outputs. The orchestrator's dispatch is a private signal to each worker.

Role assignment: Agents receive role prompts ("you are the critic"). The role specification correlates play.

The incentive question: if an agent gains by ignoring its assigned role, the architecture is not incentive-compatible.

Module 7 formalizes this: when does a multi-agent pipeline constitute an equilibrium

No-Regret Learning and CCE (Preview)

Finally, we preview a connection that will become central in Module 3: correlated equilibria arise naturally from learning dynamics.

External-regret theorem (see Cesa-Bianchi–Lugosi, 2006)

If all players use **no-regret learning** algorithms (defined in Module 3) in repeated play, the empirical distribution of play converges to the set of **coarse correlated equilibria (CCE)**.

CCE is the convergence target: NE $\subseteq$ CE $\subseteq$ CCE.

With **no-swap-regret** (a stronger condition), convergence is to CE itself.

- No external mediator needed: the history of play serves as the correlating device.
- No-regret dynamics are computationally efficient, unlike exact NE computation.
- Module 3 develops regret minimization and proves this convergence result

Summary

Today's key ideas:

- **Maxmin review:** security level $\overline{v}_i$; NE payoffs are individually rational.
- **Dominated strategies and IEDS:** strictly dominated strategies are never in NE support; IEDS preserves all NE (order-independent for strict dominance).
- **Rationalizability:** consistent with common knowledge of rationality; equals IEDS outcome in two-player games.
- **Correlated equilibrium (CE):** obedience to a mediator's private signal. Convex polytope, computable by LP.
- **Coarse correlated equilibrium (CCE):** NE $\subseteq$ CE $\subseteq$ CCE. No-regret dynamics converge to CCE (Module 3).
- **Two hierarchies:** dominant-strategy equilibria $\subseteq$ NE $\subseteq$ CE $\subseteq$ CCE; separately, NE-supported strategies are rationalizable. These are different object types.
- **Multi-agent LLM systems:** shared prompts and orchestrators are correlating