NashForge — Flashcards

A Spark & Anvil printable flashcard deck. Quiz yourself on the NashForge cast's curriculum — one card at a time.

How to use these cards

Fold each sheet down the middle line so the answers are hidden. Read the **question** on the left, say your answer out loud, then **unfold** to check the right side. Testing yourself — then checking — is one of the most powerful ways to remember. Get one wrong? That's the best kind of practice: try it again tomorrow.

Want real flashcards? Cut along the fold line and the rows into cards — question on one side, answer on the other.

---

1. Game theory foundations — players, strategies, payoffs

✂ Question (fold →)	Answer
A situation with only ONE decision-maker whose outcome depends on nobody else is best called:	**a decision problem, not a game** — With no strategic interdependence it's a single-agent decision problem, not a game.
If you change a game so a player has one MORE strategy, the matrix gains:	**an extra row (or column) of cells** — Adding a strategy to one player adds a whole row or column of new strategy combinations.
Which statement about payoffs is TRUE?	**a higher payoff simply means the player prefers that outcome more** — Payoffs encode preference ordering; they can be negative and differ across players.
Before you can find a "best strategy," you first need to know:	**the strategies available to each player and the payoffs for every combination** — Strategic analysis requires the full specification first: everyone's strategy set and every combination's payoff.
The FIRST question a strategist asks about any new situation is:	**who are the players, what can they choose, and what does each outcome pay them?** — Framing any situation as players/strategies/payoffs is the strategist's opening move.
The "players" in a game are:	**the decision-makers whose choices affect the outcomes** — Players are the agents making strategic choices; a player can be a person, firm, team, or country.
"Common knowledge" in a game means:	**each player knows the game, knows the others know it, and so on** — Common knowledge = the structure is known, known to be known, etc. — a standard rationality assumption.
In game theory, a "game" is any situation where the outcome for each participant depends on:	**the choices of ALL the participants, not just their own** — A game is defined by strategic interdependence: each player's payoff depends on the combination of everyone's choices.
Two firms each choose "High price" or "Low price." This is a game mainly because:	**each firm's profit depends on what the OTHER firm also chooses** — It's a game when one firm's best move depends on the rival's move — strategic interdependence.
A "strategy" for a player is:	**a complete plan for what to choose in the situations they might face** — A strategy is a complete plan of action — distinct from the payoff (result) and from beliefs about the opponent.
In "rock-paper-scissors," the game is best represented by:	**a 3×3 payoff matrix played simultaneously** — Two players, three simultaneous strategies each → a 3×3 simultaneous payoff matrix.
The main reason a payoff matrix is used for simultaneous games (rather than a tree) is:	**no one observes the other's move, so order doesn't matter — a grid of combinations fits** — Simultaneous play has no move order to depict, so a grid of strategy combinations is the natural form.
If a player has strategies {A, B, C} and the other has {X, Y}, the payoff matrix has how many cells?	**6** — The number of cells equals the product of the two strategy-set sizes: 3 × 2 = 6.
A game where players move at the	**a simultaneous game** — In a simultaneous game players choose without

SAME time, without seeing each other's choice, is called:	observing others' moves (a payoff matrix models it).
In a payoff pair written (3, 5) inside a cell, by the usual convention the numbers are:	**the row player's payoff first, then the column player's payoff** — By convention the first entry is the row player's payoff and the second is the column player's.
The payoff numbers in a game represent:	**how much each player values that outcome (utility), higher = better for them** — Payoffs are utilities — each player has their own, so an outcome great for one can be poor for another.
The three basic ingredients that define any game are:	**players, strategies, and payoffs** — Every game is specified by its players, their strategies, and the payoffs for each combination of choices.
Which of these is NOT a strategic game in the game-theory sense?	**rolling a single die to see who gets the higher number** — Rolling a die is pure chance — no player chooses a strategy that affects another's payoff.
Reading a cell (2, 2) versus a cell (5, 0): the row player prefers the second only if:	**they care about their own payoff (5 > 2), regardless of the column player getting 0** — Each player ranks cells by their OWN payoff; 5 > 2 for the row player, whatever the column player gets.
In a payoff matrix for a two-player game, each CELL shows:	**the payoff to each player for that specific pair of strategies** — A cell sits at the intersection of one row-strategy and one column-strategy and lists each player's payoff there.
Modeling a real situation as a game is useful mainly because it:	**forces you to name the players, choices, and payoffs, then reason about interdependence** — The modeling discipline surfaces the players, strategies, and payoffs so you can reason about best responses.
A game is "well-defined" once you can state, for every combination of strategies:	**the payoff each player receives** — A well-defined game maps every strategy combination to a payoff for each player.
Which is the clearest example of strategic interdependence?	**two gas stations across the street setting prices, each watching the other** — The two stations' best prices each depend on the rival's price — textbook interdependence.
A "two-player, two-strategy" game is often called a:	**2×2 game** — A 2×2 game has two players each with two strategies — the smallest interesting matrix.
An outcome is a "possible result" of a game; it corresponds to:	**one specific combination of strategies, one per player** — An outcome is one full strategy profile (a choice for each player) — one cell of the matrix.

2. Strictly and weakly dominant and dominated strategies

✂ Question (fold →)	Answer
The value of finding a dominant strategy is that you:	**can predict a rational player's move without knowing the opponent's plan** — Dominance lets you pin down a rational choice regardless of the opponent — a powerful shortcut.
Eliminating dominated strategies is useful because it:	**shrinks the game, sometimes revealing a unique rational outcome** — Removing dominated options simplifies the game and can pin down a single predicted outcome.
A dominant strategy is defined WITHOUT reference to:	**what the opponent actually chooses** — Dominance holds against every opponent move, so it needs no prediction of the opponent's actual choice.
If neither player has any dominant strategy, you should next look for:	**best responses and a Nash equilibrium** — With no dominance, the next tool is best-response analysis to locate Nash equilibria.
A strategy that is dominant is ALSO always:	**a best response to every opponent strategy** — If a strategy is best against every opponent move, it is by definition a best response to each.
A STRICTLY dominant strategy gives a payoff that is:	**strictly higher than every other option, whatever the opponent does** — Strict dominance means a strictly higher payoff against each opponent strategy — never a tie.
A common MISTAKE is to assume every game has a dominant strategy. In truth:	**many games have no dominant strategy for either player** — Plenty of games (e.g. coordination games) have no dominant strategy; dominance is a bonus, not a guarantee.
The quickest first step when solving a matrix game is to:	**check each player for a dominant strategy** — Checking for dominant strategies is the fastest route to a solution when one exists.
A strategy is a DOMINANT strategy for a player when it:	**gives that player at least as good a payoff no matter what the others do** — A dominant strategy is best (weakly or strictly) for the player against every possible opponent choice.
Row payoffs — A: (5 vs L, 2 vs R); B: (5 vs L, 4 vs R). For the row player, B is:	**weakly dominant** — B ties A against L (5=5) but beats it against R (4>2): at least as good, strictly better once → weakly dominant.
If a player has a strictly dominant strategy, a rational player will:	**always play it** — Since it's best against everything, a rational player always plays a strictly dominant strategy.
Row — Up:(6,?),(1,?); Down:(5,?),(2,?) against L,R. Up vs Down for the row player:	**Up strictly dominates (6>5 and 1<2 is false — recheck): actually neither, since 1<2** — Up gives 6>5 vs L but 1<2 vs R, so Up is NOT better everywhere — neither strategy dominates.
The safest way to describe a dominant strategy to a friend is:	**"the move that's best for me no matter what you do"** — Dominance = best for me against any opponent choice — independent of what the opponent does.
Row payoffs — Up: (4 vs L, 3 vs R); Down: (2 vs L, 1 vs R). For the row player, Up is:	**strictly dominant** — Up beats Down against L (4>2) and against R (3>1): strictly higher everywhere → strictly dominant.
You should NOT eliminate a strategy during dominance	**not dominated by any of the player's other strategies** — Only strategies

analysis if it is:	dominated by another of the player's own options may be eliminated.
Row — Left: (3 vs a, 3 vs b); Right: (3 vs a, 3 vs b). Between Left and Right, the row player is:	**indifferent — neither dominates** — Identical payoffs in every column mean neither strategy dominates; the player is indifferent.
If a strategy is strictly dominated, iterated elimination will:	**safely remove it — a rational player never uses it** — Strictly dominated strategies can always be removed without loss, since no rational player would use them.
"Iterated elimination of dominated strategies" means:	**remove dominated strategies, then re-check, repeating as new ones become dominated** — You delete dominated strategies, which can make others dominated, and repeat until none remain.
A DOMINATED strategy is one that:	**is worse than some other strategy against every opponent move** — A dominated strategy is beaten by another of the player's own strategies against every opponent choice.
If BOTH players have a strictly dominant strategy, the outcome where both play it is:	**a dominant-strategy equilibrium (and also a Nash equilibrium)** — Both playing their dominant strategy is a dominant-strategy equilibrium, which is always a Nash equilibrium.
A weakly dominant strategy is trickier to eliminate against because:	**the order of elimination can change which equilibria survive** — With weak dominance, different elimination orders can yield different surviving outcomes — a known subtlety.
In the Prisoner's Dilemma, "Confess" is a dominant strategy because:	**it yields a better payoff whether the other confesses or stays silent** — Confessing beats staying silent against both of the opponent's choices — that's what makes it dominant (and the dilemma).
A WEAKLY dominant strategy is:	**at least as good against every opponent move, and strictly better against at least one** — Weak dominance ties in some cases but is strictly better in at least one, and never worse.
If Down is dominated for the row player, a rational row player will:	**never play Down** — A rational player never plays a dominated strategy — another option is always at least as good and sometimes better.
Which pair of concepts is correctly matched?	**dominant = best against all opponent moves; dominated = worst against all** — Dominant is best against every opponent move; dominated is beaten against every opponent move.

3. The PD; individual vs collective outcome

✂ Question (fold →)	Answer
If both prisoners could make a BINDING agreement to stay silent, they would:	**both cooperate and reach the better mutual outcome** — A binding, enforceable agreement removes the temptation to defect, enabling the better cooperative outcome.
A key insight of the PD is that:	**individually rational behavior can produce collectively bad results** — The PD shows self-interested rationality can lead everyone to a worse joint outcome.
A student says "the PD proves cooperation is impossible." The best correction is:	**"in a one-shot game defection is rational, but repetition and institutions can sustain cooperation"** — The PD explains WHY cooperation is hard and points to the mechanisms (repetition/enforcement) that enable it.
The reason players defect in a one-shot PD is:	**defection pays more regardless of the other's choice** — Defection is dominant — better against cooperation and against defection — so rational one-shot play defects.
The PD is a NON-zero-sum game because:	**the players' payoffs don't simply sum to a constant — both can gain or both can lose** — In the PD both can do better (cooperate) or both worse (defect) — the total isn't fixed, so it's non-zero-sum.
A vaccine or clean-water public good resembles the PD in that:	**everyone benefits if all contribute, but each is tempted to let others pay** — Public goods create the free-rider temptation central to the multi-player PD.
The single most important takeaway from the PD is:	**aligning individual incentives with the group outcome is the central challenge of strategy design** — The PD frames mechanism/institution design as the task of making the group-good choice individually rational.
Which real situation is best modeled as a PD?	**two countries deciding whether to limit emissions, each tempted to pollute** — Each country gains from others cutting emissions while it pollutes — a classic collective-action PD.
In a one-shot PD, communication WITHOUT enforcement usually:	**does not change the outcome — talk is cheap, defection stays dominant** — Non-binding promises don't alter payoffs, so defection remains dominant — 'cheap talk.'
Why is the PD outcome called a "dilemma"?	**because rational self-interest leads to an outcome both players regret** — The dilemma is that doing the individually smart thing makes both worse off — a regretful trap.
In the PD, mutual cooperation is:	**better for both than mutual defection, but not individually stable** — Both cooperating pays more than both defecting, but each is tempted to defect — so it isn't stable.
In a standard PD, each player's dominant strategy is to:	**defect (not cooperate)** — Defection pays more against both cooperation and defection, so it's dominant — and the trap.
The outcome (both defect) in the PD is:	**a Nash equilibrium that is not Pareto-optimal** — Mutual defection is stable (Nash) yet both would be better off cooperating (not Pareto-optimal).
The Prisoner's Dilemma is famous because:	**each player's rational (dominant) choice leads to an outcome worse for both than cooperating** — Individually rational defection yields mutual defection — worse for both than mutual cooperation.
	both would profit more by NOT advertising, yet each has a dominant

Two firms deciding to advertise heavily illustrate a PD when:	**reason to advertise** — If advertising is individually dominant but jointly wasteful, the firms are trapped in a PD.
In the PD payoff order (best→worst) for one player, the ranking is:	**I defect while you cooperate > we both cooperate > we both defect > I cooperate while you defect** — The classic ranking (T>R>P>S) makes defection dominant and mutual defection the trap.
Mutual defection is "stable" in the PD because:	**given the other defects, switching to cooperate would only lower your payoff** — If the other defects, your best response is also to defect — neither gains by switching, so it's a Nash equilibrium.
Which change would turn a PD into a game where cooperation is rational?	**repeating the game many times so reputation matters** — Repetition lets strategies like tit-for-tat make cooperation rational via the shadow of the future.
The "collective action problem" generalizes the PD to:	**many players who each benefit from free-riding, hurting the group** — Public-goods and common-pool dilemmas are multi-player PDs where individual incentives undercut the group.
The tit-for-tat strategy works in a REPEATED PD because it:	**rewards cooperation and punishes defection, making cooperation pay over time** — Tit-for-tat cooperates first then mirrors, so cooperating becomes the rational long-run choice.
The PD teaches that to sustain cooperation, societies often need:	**institutions, repetition, or enforcement that change the incentives** — Repetition, reputation, contracts, and institutions realign incentives so cooperation becomes rational.
"Free-riding" in a group PD means:	**enjoying the shared benefit without contributing to it** — A free-rider takes the public benefit while defecting on the cost — the multi-player defect move.
If defecting were NOT dominant (e.g., cooperating paid more against a cooperator), the game:	**would no longer be a Prisoner's Dilemma** — The PD is defined by defection being dominant; remove that and it's a different game.
"Pareto-optimal" means an outcome where:	**no one can be made better off without making someone else worse off** — Pareto optimality: you can't improve one player without hurting another.
The gap between the Nash outcome and the cooperative outcome in a PD is often called:	**the price of anarchy (inefficiency of self-interested play)** — The loss from everyone playing selfishly versus cooperating is a measure of the inefficiency (informally, the price of anarchy).

4. Best response; the stable point

✂ Question (fold →)	Answer
A Nash equilibrium is a set of strategies where:	**no player can improve their payoff by changing their own strategy alone** — At a Nash equilibrium each player's strategy is a best response to the others; no unilateral deviation helps.
Every dominant-strategy equilibrium is:	**also a Nash equilibrium** — If each plays a dominant strategy, each is best-responding, so it's a Nash equilibrium (the stronger implies the weaker).
A game can have:	**zero, one, or several pure-strategy Nash equilibria** — Games may have none in pure strategies, one, or multiple; a mixed Nash always exists in finite games.
The underlining method is reliable for finding PURE Nash equilibria but:	**does not directly find mixed-strategy equilibria** — Underlining locates pure Nash cells; mixed equilibria need the indifference (best-response) calculation.
A Nash equilibrium in which players randomize is called a:	**mixed-strategy Nash equilibrium** — When equilibrium play requires randomizing, it's a mixed-strategy Nash equilibrium.
The underlining method finds Nash equilibria by:	**underlining each player's best response in every case; a cell underlined by both is Nash** — Mark each player's best responses; cells that are a mutual best response (both marks) are Nash equilibria.
A "best response" is:	**the strategy that gives you the highest payoff GIVEN what the others are doing** — A best response maximizes your payoff conditional on the opponents' strategies.
A strategist checks for Nash equilibria mainly to predict:	**where strategic interaction will settle** — Nash analysis predicts the resting point(s) of strategic interaction.
A Nash equilibrium is stable but:	**not necessarily the best outcome for the players (e.g. the PD)** — Nash means no one deviates, not that the outcome is jointly best — the PD is the counterexample.
If both players are already best-responding, the strategic situation is:	**at rest — a Nash equilibrium, no one has reason to move** — Mutual best response means no one has an incentive to change — the game is at rest.
If a cell is a best response for the row player but NOT the column player, it is:	**not a Nash equilibrium** — Nash requires BOTH players to be best-responding; one-sided best response is not enough.
The difference between "best outcome" and "Nash outcome" matters most when:	**the stable outcome is inefficient, as in the Prisoner's Dilemma** — The gap is largest when self-interested stability produces a jointly inferior result (the PD).
The most common misconception about Nash equilibrium is that it:	**must be the outcome that is best for the group** — Nash is about stability, not joint optimality; the PD's Nash is worse for both than cooperation.
Which statement is TRUE about the PD's Nash equilibrium?	**it exists (both defect) and is worse for both than cooperating** — Both-defect is the PD's pure Nash equilibrium and is Pareto-dominated by mutual cooperation.

Two Nash equilibria in one game means:	**players may face a coordination problem about which to reach** — Multiple equilibria raise the question of which one players coordinate on — the focal-point issue.
The reason Nash equilibrium is central to game theory is that it:	**predicts stable outcomes of strategic interaction without needing dominance** — Nash gives a general prediction of stable play even when no dominant strategies exist.
When advising a client, "reach a Nash equilibrium" is useful because it means:	**a self-enforcing arrangement no one wants to break alone** — A Nash arrangement is self-enforcing: given everyone else, no one profits by breaking it — useful for stable deals.
In "matching pennies," there is no PURE Nash equilibrium because:	**whatever one player picks, the other wants to switch — only randomizing is stable** — Matching pennies is a cycling zero-sum game; only a mixed (50/50) strategy is a stable equilibrium.
If you find a cell where BOTH players have a strictly dominant strategy, you have found:	**the unique Nash equilibrium of that game** — When both have strictly dominant strategies, their intersection is the unique Nash equilibrium.
To CHECK whether a cell is a Nash equilibrium, ask:	**can either player raise their payoff by switching alone? If no, it is Nash** — The Nash test is the no-profitable-unilateral-deviation check for each player.
At a Nash equilibrium, each player's strategy is:	**a best response to the others' strategies** — Nash equilibrium = mutual best responses; that's exactly the no-profitable-deviation condition.
John Nash's key theorem states that every finite game has:	**at least one Nash equilibrium (possibly in mixed strategies)** — Nash proved every finite game has at least one equilibrium, allowing mixed strategies.
Matrix — Row best responses: Up to L, Down to R. Column best responses: L to Up, R to Down. The Nash equilibria are:	**(Up, L) and (Down, R)** — (Up,L): Row best-responds Up to L and Column best-responds L to Up — mutual; likewise (Down,R). Two pure Nash.
A firm at a Nash equilibrium price will change its price only if:	**a rival changes first (its best response depends on the rival's move)** — At equilibrium the firm is best-responding; it moves only if the rival's move makes a different price a better response.
In a coordination game (e.g. both pick the same meeting spot), the Nash equilibria are:	**the outcomes where both choose the same spot** — Coordinating (both choose the same) is mutually best-responding, so those matching cells are Nash equilibria.

5. Mixed strategies; randomize to be unexploitable

✂ Question (fold →)	Answer
A mixed strategy is written as, e.g., (0.7 Attack, 0.3 Defend), meaning:	**attack 70% of the time and defend 30%** — The numbers are the probabilities placed on each action.
A commentator says "great teams are unpredictable." In game terms this means they:	**play a mixed strategy the opponent cannot best-respond to** — Unpredictable = mixing so no single counter dominates.
If both players in matching pennies deviate from 50/50, the deviator:	**becomes exploitable by the other** — Any tilt away from 50/50 lets the opponent exploit the imbalance.
If action A always beats action B for you regardless of the opponent, then:	**you should play A purely, not mix** — With a dominant pure action there's no reason to randomize.
Why would a rational player randomize at all?	**to stop the opponent from predicting and exploiting a fixed choice** — Randomizing removes the pattern an opponent could exploit.
Randomization is unnecessary in games that:	**have a pure-strategy Nash equilibrium the players are content to play** — If a satisfactory pure equilibrium exists, mixing isn't required.
A team that becomes predictable on offense is vulnerable because:	**the defense can pre-commit to the best counter** — Predictability lets the opponent best-respond in advance.
If the payoff to Attack rises, a player's equilibrium mix will generally:	**shift, and the OPPONENT's mix adjusts to restore indifference** — Changing payoffs shifts the indifference point, so the equilibrium mixes adjust.
Which is a real setting where mixing is rational?	**a tax auditor deciding which returns to check so filers can't predict audits** — Unpredictable auditing deters gaming the system — a mixed-strategy use.
The equilibrium mix in a 2×2 game is found by:	**setting the opponent's expected payoffs from each action equal (indifference)** — Solve the indifference equation: opponent's expected payoffs across actions equal.
A mixed strategy differs from "playing randomly with no thought" because:	**the probabilities are chosen deliberately to be optimal** — An optimal mix uses specific, computed probabilities — not aimless randomness.
In rock-paper-scissors, the unexploitable strategy is to play each option:	**with probability 1/3** — Playing 1/3 each makes you unpredictable and unexploitable.
A pure strategy is a special case of a mixed strategy where:	**one action has probability 1** — A pure strategy is a degenerate mix putting all weight on one action.
A pure-strategy player facing a smart adaptive opponent should expect to:	**be exploited unless they switch to a mix** — Against adaptation, a pure strategy is exploitable; mixing is the defense.
If a soccer kicker always shoots left, the goalie will:	**always dive left, so the kicker should mix** — A predictable kicker is exploited; mixing restores unpredictability.

A MIXED strategy is one where a player:	**chooses among actions according to probabilities** — A mixed strategy assigns probabilities to actions and randomizes.
The point of the opponent-indifference condition is that at equilibrium:	**your mix is chosen so the opponent can't gain by favoring any one action** — You set YOUR probabilities to make the OPPONENT indifferent, pinning their play.
In a mixed equilibrium, changing YOUR probabilities (while opponent holds theirs) will:	**not change your expected payoff (you're indifferent across your mixed actions)** — At a mixed NE you're indifferent, so reweighting your mix doesn't change your expected payoff.
Nash proved that allowing mixed strategies guarantees:	**at least one equilibrium exists in every finite game** — Mixed strategies are what make Nash's existence theorem hold for all finite games.
A mixed-strategy Nash equilibrium makes each player:	**indifferent among the actions they mix over, given the other's mix** — At a mixed NE each player is indifferent across their mixed actions given the opponent's probabilities.
The core idea of mixed strategies is:	**sometimes the best plan is to be deliberately unpredictable** — When predictability is punished, deliberate randomization is optimal.
A PURE strategy is one where a player:	**always chooses the same single action** — A pure strategy picks one action with certainty.
"Matching pennies" has no pure Nash equilibrium, so its equilibrium is:	**each player mixing 50/50** — The unique equilibrium is each randomizing 50/50.
The best way to describe a mixed strategy to a teammate is:	**"vary your choice on purpose so they can't read you"** — A mix = deliberate variation to stay unreadable.
The reason a fixed, known routine loses in competitive games is:	**the opponent adapts to exploit the routine** — A known routine is a pure strategy the opponent can best-respond to.

6. Fixed vs growing pie

✂ Question (fold →)	Answer
Framing a negotiation as zero-sum when it is actually non-zero-sum tends to:	**destroy value both sides could have shared** — Treating a growable pie as fixed forecloses mutual gains.
A negotiator who only ever claims value (never creates it) assumes the game is:	**zero-sum** — Pure value-claiming behavior treats the pie as fixed — zero-sum framing.
If two firms both cut costs and both profit more, the situation is:	**non-zero-sum** — Both gaining means value was created — non-zero-sum.
Chess (win/lose/draw as +1/−1/0) is:	**zero-sum** — One player's win is the other's loss — a zero-sum contest.
A tug-of-war for a fixed rope length is which kind of game?	**zero-sum** — One side's gain in rope is exactly the other's loss — a fixed total, so zero-sum.
A trade where both parties are better off afterward shows the deal was:	**non-zero-sum (value was created)** — Voluntary mutually-beneficial trade creates value — non-zero-sum.
The FIRST question a strategist asks about a new situation is often:	**is the pie fixed (zero-sum) or can it grow (non-zero-sum)?** — Whether value is fixed or expandable reshapes the whole strategy.
A ZERO-SUM game is one where:	**one player's gain exactly equals the other's loss** — In zero-sum games the payoffs sum to a constant; my gain is your loss.
Recognizing a non-zero-sum game lets you:	**search for cooperative moves that make everyone better off** — Spotting non-zero-sum structure opens mutually-beneficial strategies.
If total payoffs across all outcomes are always the same constant, the game is:	**constant-sum (a generalization of zero-sum)** — Constant-sum games (fixed total) are strategically equivalent to zero-sum.
The strategic danger of assuming every game is zero-sum is:	**you may compete when cooperating would grow the pie for everyone** — Zero-sum bias makes you miss cooperative value-creation.
In a purely zero-sum game, cooperation:	**cannot help both players simultaneously** — Zero-sum means gains are exactly opposed, so both can't benefit at once.
The concept most closely tied to zero-sum two-player games is:	**the minimax theorem** — Von Neumann's minimax theorem is the foundational result for zero-sum games.
"The pie is fixed" is a good assumption for:	**dividing a set prize with no way to increase it** — A truly fixed prize (nothing you do grows it) is genuinely zero-sum.
In a zero-sum two-player game, a Nash equilibrium coincides with each player's:	**minimax (securing the best guaranteed outcome)** — Von Neumann's minimax theorem: zero-sum equilibria are minimax/maximin strategies.
A game can be PARTLY zero-sum and partly not when:	**players share some common interests and oppose on others** — Mixed-motive games blend cooperation and conflict (e.g. the PD).

Which pair is correctly classified?	**poker for a fixed pot = zero-sum; a startup partnership = non-zero-sum** — Fixed-pot poker is zero-sum; a partnership can grow value for both.
A poker hand for a fixed pot among players is essentially:	**zero-sum (chips won equal chips lost)** — For a fixed pot, one player's winnings equal others' losses — zero-sum.
The deepest lesson of the zero/non-zero distinction is:	**know whether you're dividing value or creating it before you choose tactics** — Diagnose the pie (fixed vs growable) first; it dictates whether to compete or cooperate.
A NON-zero-sum game is one where:	**players can both gain or both lose — the total isn't fixed** — Non-zero-sum: the pie can grow or shrink; outcomes aren't purely opposed.
Sports leagues are non-zero-sum overall because:	**the whole league can grow revenue even as teams compete** — Individual games are zero-sum, but the league's total value can grow — non-zero-sum at the system level.
"Win-win" outcomes are only possible in:	**non-zero-sum games** — Win-win requires a growable pie, i.e. a non-zero-sum structure.
The Prisoner's Dilemma is:	**non-zero-sum (both can be better or worse off together)** — Both cooperating beats both defecting for both — the total isn't fixed, so non-zero-sum.
Non-zero-sum thinking encourages negotiators to:	**look for trades that create value before dividing it** — Expand the pie (create value) before splitting it — the non-zero-sum move.
In a zero-sum game, "helping the opponent" is:	**exactly equal to hurting yourself** — By definition, their gain is your loss in zero-sum play.

7. Coordination games; focal points

✂ Question (fold →)	Answer
In coordination games, communication BEFORE play usually:	**helps, because it lets players agree which equilibrium to reach** — Unlike the PD, pre-play talk helps coordination by selecting an equilibrium.
A focal point works WITHOUT communication because it is:	**salient in a way both players expect the other to notice** — Shared salience lets each expect the other to pick it — enabling silent coordination.
The core difficulty in a coordination game is:	**choosing WHICH equilibrium to reach when several are good** — With multiple good equilibria, the problem is coordinating on one.
Driving on the same side of the road as everyone else is a coordination game because:	**you do better matching others' choice, whichever side that is** — Left or right both work; what matters is everyone matching — pure coordination.
Language itself is a coordination equilibrium because:	**a word means what it does only because everyone uses it the same way** — Shared meaning is a self-sustaining coordination equilibrium.
A "focal point" (Schelling point) is:	**a choice that stands out and helps players coordinate without communicating** — A focal point is a naturally salient option people converge on absent communication.
A focal point can be SUBOPTIMAL when:	**the salient equilibrium isn't the highest-payoff one** — Salience may pull players to an obvious-but-inferior equilibrium.
If a coordination game has equilibria (A,A) and (B,B), a mismatch (A,B) is:	**not an equilibrium — at least one player wants to switch** — A mismatch isn't stable: someone gains by switching to match.
A meeting scheduled for "noon" rather than "sometime" uses:	**a focal time to coordinate arrival** — A specific salient time is a focal point that coordinates arrivals.
If a group must all pick the same messaging app to talk, that is a:	**coordination game** — Everyone benefits from converging on one app — pure coordination.
A coordination game is one where players:	**both do better when they match choices** — In coordination games matching choices yields the best mutual outcomes.
If everyone expects everyone else to pick option X, then picking X is:	**a best response — expectations become self-fulfilling** — Shared expectations make matching the crowd a best response — self-fulfilling coordination.
Why can two good equilibria be a PROBLEM?	**players may fail to coordinate and land on a mismatch** — Without a selection device, players can miscoordinate onto a worse mismatch.
A traffic light solves a coordination problem by:	**creating a shared signal that tells everyone which equilibrium to play** — The light is a public signal selecting one equilibrium (who goes now).
The tool that best resolves a coordination game is:	**a shared signal or convention selecting one equilibrium** — A common signal/convention coordinates everyone onto one equilibrium.
Two friends who lose their phones in a city will often meet at:	**an obvious landmark both would think of (a focal point)** — They converge on a mutually obvious landmark — the classic focal-point result.

A company picking the industry-standard file format is choosing:	**the focal equilibrium to ensure compatibility** — Adopting the prevailing standard is picking the focal, compatible equilibrium.
The synthesis of coordination games is:	**when interests align, the challenge is agreeing on which good outcome to reach** — Coordination shifts the problem from conflict to equilibrium selection.
Coordination games typically have:	**multiple Nash equilibria (each matching outcome)** — Each matched outcome is a Nash equilibrium, so there are several.
The "Battle of the Sexes" game differs from pure coordination because:	**players agree they want to coordinate but disagree on which equilibrium** — Both prefer coordinating, but each prefers a different matched outcome — a coordination + conflict blend.
Focal points often come from:	**shared culture, precedent, or salience rather than payoffs** — Salience arises from context — culture, history, prominence — not just numbers.
Which is a coordination game rather than a PD?	**two firms choosing whether to adopt the same charging-plug standard** — Adopting a common standard is pure coordination (match = win), not a defect-temptation PD.
Standards (like plug shapes or file formats) are coordination successes because:	**everyone benefits from converging on one common option** — Common standards are coordinated equilibria that let everyone interoperate.
"Schelling" is associated with the idea of:	**focal points in coordination without communication** — Thomas Schelling introduced focal points for tacit coordination.
The reason focal points are studied is that they explain:	**how people coordinate successfully even without any agreement** — Focal points explain silent, agreement-free coordination.

8. Repeated games; tit-for-tat; reputation

✂ Question (fold →)	Answer
A repeated game is one where:	**the same players play the same game many times** — Repeated games have the same players facing the same stage game over time.
The "tit-for-tat" strategy:	**cooperates first, then copies the opponent's last move** — Tit-for-tat starts nice and then mirrors — the famous robust repeated-PD strategy.
The discount factor in a repeated game measures:	**how much players value future payoffs relative to now** — A higher discount factor = future payoffs matter more = cooperation easier.
If defection is punished immediately and cooperation restored quickly, cooperation is:	**more likely to be sustained** — Prompt punishment plus forgiveness stabilizes cooperation (the tit-for-tat recipe).
The "folk theorem" broadly says that in indefinitely repeated games:	**many outcomes, including cooperation, can be sustained as equilibria** — The folk theorem: with enough patience, a wide range of outcomes (incl. cooperation) are equilibria.
Which change makes defection MORE tempting in a repeated game?	**the game is about to end for certain** — A certain, near end removes future punishment, tempting defection.
The strategic value of being predictably retaliatory is that it:	**deters opponents from defecting in the first place** — A credible retaliatory reputation deters defection before it happens.
Reputation systems (ratings, reviews) work by:	**making one-shot-like interactions behave like repeated games** — Public track records import future consequences into otherwise one-shot deals.
Tit-for-tat is "forgiving" because:	**it returns to cooperation as soon as the opponent does** — After punishing once, tit-for-tat forgives and cooperates again if the opponent does.
A repeated game turns the PD's lesson from "defect" into:	**"cooperation can be rational when the future casts a long shadow"** — Repetition can make cooperation the rational choice — the key repeated-game insight.
A "grim trigger" strategy:	**cooperates until the opponent defects once, then defects forever** — Grim trigger cooperates until any defection, then punishes permanently.
A country that breaks treaties repeatedly develops:	**a bad reputation that makes future cooperation harder to obtain** — Repeated defection damages reputation, reducing others' willingness to cooperate.
The "shadow of the future" means:	**valuing future payoffs enough that today's defection isn't worth it** — When future payoffs matter, the long-run cost of defection deters it.
"Reputation" matters in repeated games because:	**a player's past behavior signals how they'll act, shaping others' responses** — Track records let others predict and condition their play — reputation has strategic value.
A repeated interaction with an UNKNOWN number of future rounds most helps cooperation because:	**players can't unravel it by backward induction from a known last round** — Without a known final round there's no last-round defection to unravel backward, so cooperation can hold.

The main reason a one-shot game and a repeated game differ is:	**future rounds create incentives absent in a single play** — Repetition adds future-oriented incentives (reward, punishment, reputation) missing in one-shot play.
The synthesis of repeated games is:	**time and reputation can transform selfish incentives into cooperative ones** — Repetition + reputation realign incentives toward cooperation.
A business that values repeat customers tends to:	**treat them well because the relationship is a repeated game** — Repeat interaction makes honest treatment rational — cheating loses future value.
"Generous tit-for-tat" improves on tit-for-tat by:	**occasionally forgiving a defection to break punishment cycles** — Occasional forgiveness prevents error-driven retaliation loops.
Axelrod's tournaments found the winning repeated-PD strategies were:	**nice, retaliatory, forgiving, and clear** — Successful strategies were nice, punished defection, forgave, and were simple/clear.
Repetition can sustain cooperation in a PD because:	**future rounds let players reward cooperation and punish defection** — The shadow of future play makes cooperating rational when defection is punished later.
In repeated business dealings, honesty is often rational because:	**the long-run value of the relationship exceeds a one-time gain from cheating** — When future dealings are valuable, preserving trust beats a one-shot cheat.
If a repeated game has a KNOWN final round, backward induction suggests:	**players defect in the last round, unraveling cooperation** — A known last round removes future punishment, and defection unravels backward.
Cooperation is easier to sustain when the game:	**is repeated indefinitely (no known end)** — An uncertain/indefinite horizon preserves the shadow of the future.
Tit-for-tat can get stuck in a cycle of mutual punishment if:	**a single mistaken defection triggers alternating retaliation** — One error can start an echo of retaliation — a known weakness addressed by forgiving variants.

9. Extensive form; sequential moves

✂ Question (fold →)	Answer
In a game tree, a NODE represents:	**a decision point for a player** — Each decision node is where a specific player chooses an action.
A "subgame" is:	**a portion of the tree that itself forms a smaller complete game** — A subgame starts at a node (with a singleton information set) and includes everything after it.
Converting a tree to a matrix (normal form) is possible but often:	**loses the visible move order and information structure** — The matrix keeps strategies+payoffs but hides the sequence and information that the tree shows.
The end of each path in a game tree shows:	**the payoffs to all players for that sequence of moves** — Terminal nodes list the payoffs resulting from that full sequence of choices.
A player's STRATEGY in extensive form specifies:	**an action at every decision node they could reach** — An extensive-form strategy is a complete contingent plan — an action at each of the player's nodes.
The solution concept most associated with game trees is:	**backward induction / subgame-perfect equilibrium** — Trees are solved by backward induction, yielding subgame-perfect equilibria.
The extensive form is essential when analyzing:	**threats, commitments, and first-mover advantages** — Sequence-dependent ideas (threats, commitment, first-mover) require the extensive form.
The first decision node of a game tree is the:	**root** — The root is the starting decision node of the tree.
A tree with all singleton information sets is a game of:	**perfect information** — Every node distinguishable → perfect information (like chess).
An extensive-form game where a later player does NOT observe an earlier move is:	**a game of imperfect information** — Unobserved earlier moves (grouped in information sets) = imperfect information.
An "information set" grouping two nodes means the player:	**cannot tell which of those nodes they are at** — An information set marks nodes a player can't distinguish — modeling imperfect information.
First-mover advantage can only be modeled in:	**the extensive (sequential) form** — You need an explicit move order — the extensive form — to model moving first.
A key advantage of the extensive form over the matrix is that it:	**shows the ORDER of moves and what each player observes** — Trees capture sequence and information — things a simultaneous matrix can't show.
At a node, "it is Player 2's turn" tells you:	**Player 2 chooses among the branches leaving that node** — The node's label marks which player selects among its branches.
The number of terminal nodes equals:	**the number of distinct complete paths through the tree** — Each full path from root to a leaf is one outcome; leaves count the outcomes.
A game TREE (extensive form) is used to represent:	**games where players move in sequence** — Extensive-form game trees depict sequential moves and who knows what when.

The synthesis of extensive form is:	**when order and information matter, draw the tree before you solve** — Model sequence and information explicitly with a tree, then apply backward induction.
The main reason to DRAW a game tree is to:	**reason clearly about who moves when and what they know** — The tree organizes sequence and information so you can reason (e.g. backward-induct).
A simultaneous game can be drawn as a tree by:	**using an information set so the second mover doesn't see the first move** — Hiding the first move inside an information set models simultaneity in a tree.
The branches from a node represent:	**the actions available at that decision point** — Branches are the possible actions a player can take at that node.
A game with both sequential AND simultaneous elements is drawn with:	**a tree using information sets for the simultaneous parts** — Information sets let one tree capture both sequential and simultaneous moves.
Reading a game tree from ROOT to LEAF traces:	**one possible sequence of moves and its outcome** — A root-to-leaf path is one full play of the game.
Chess is naturally an extensive-form game because:	**players alternate moves, observing the board each turn** — Alternating, observed moves are the definition of a sequential/extensive-form game.
In a game tree, two branches leaving the same node must represent:	**different actions the same player can take there** — All branches out of one node are the alternative actions available to whoever moves at that node.
A sequential game where the second player sees the first player's move has:	**perfect information at the second node** — If the mover observes prior moves, that node has perfect information.

10. Subgame-perfect; reason from the end

✂ Question (fold →)	Answer
Why does a NON-credible threat fail under backward induction?	**the threatened player knows the threatener wouldn't actually carry it out** — If executing the threat hurts the threatener, the target anticipates they won't — so it's ignored.
The "centipede game" is famous because backward induction predicts:	**players stop immediately, though cooperating longer would pay more** — Backward induction unravels the centipede to an early stop, despite mutual gains from continuing.
The synthesis of backward induction is:	**in sequential games, reason from the end and only make threats you'd actually carry out** — Solve from the end; credibility of threats/commitments is decided by future payoffs.
Backward induction can fail to predict real behavior when:	**players are not perfectly rational or care about fairness** — Real players (e.g. in the centipede/ultimatum) deviate due to bounded rationality or fairness.
Backward induction requires the game to have:	**a finite horizon (a known last move)** — You need a defined end to start reasoning backward from it.
Backward induction solves a sequential game by:	**starting at the LAST decision and reasoning toward the first** — You determine the last mover's best choice, then the prior mover's, etc.
In the finitely-repeated PD, backward induction predicts:	**defection in every round** — With a known last round, defection unravels backward to every round.
The ultimatum game challenges backward induction because responders:	**often reject unfair offers that backward induction says they should accept** — Backward induction predicts accepting any positive offer, but people reject unfair ones.
The first step of backward induction is to:	**find each final-node player's best action** — Solve the last movers first — their best actions are known with certainty.
The practical lesson of backward induction for planning is:	**decide your first move by imagining how the game ends** — Reason from the endgame back to the present decision.
A chess engine using "look-ahead" is doing a form of:	**backward induction over possible move sequences** — Evaluating from future positions back to the present move is backward induction.
A "credible commitment" changes a sequential game by:	**altering future payoffs so a formerly empty threat becomes rational to carry out** — Commitment devices change payoffs so the threat is now in the threatener's interest — credible.
The single most useful habit from backward induction is:	**anticipating others' future best responses before you commit** — Anticipate the rational endgame, then choose your current move accordingly.
Backward induction assumes every player is:	**rational and expects others to be rational** — It relies on common knowledge of rationality at every node.
Given the rival Accommodates (2,1), your Entry choice In vs Out(0,2) is:	**In (2 > 0)** — Anticipating Accommodate, you compare In→2 vs Out→0 and enter.
If two actions tie at a final node,	**may allow multiple subgame-perfect outcomes** — Ties at a node can

backward induction:	produce multiple subgame-perfect equilibria.
A firm that builds a big factory to signal it will fight entrants is using:	**a commitment device to make a threat credible** — The sunk investment changes payoffs so fighting becomes rational — a credible commitment.
A player anticipates a later player's response by:	**assuming the later player will choose their own best action** — Backward induction assumes rational later play, which the earlier player anticipates.
To find the subgame-perfect equilibrium you should:	**solve each subgame from the smallest/last up to the whole game** — Solve the deepest subgames first, then fold their solutions upward.
Backward induction on a 2-move tree means you first solve the move made:	**last (deepest in the tree)** — You determine the final mover's best action first, then fold it back to the earlier move.
Tree: You choose In/Out; if In, rival chooses Fight(−1,−1) or Accommodate(2,1). Rival will:	**Accommodate (1 > −1)** — At the last node the rival compares 1 (Accommodate) vs −1 (Fight) and picks Accommodate.
The "entry deterrence" game shows that a threat to Fight is:	**not credible if fighting hurts the threatener** — Backward induction reveals a threat only works if carrying it out is in the threatener's interest.
Backward induction is most reliable in games that are:	**finite with perfect information** — It cleanly applies to finite perfect-information games (one action best at each node).
An equilibrium found by backward induction is called:	**subgame-perfect** — Backward induction yields a subgame-perfect equilibrium (optimal in every subgame).
Subgame-perfection is STRONGER than Nash equilibrium because it:	**rules out equilibria relying on non-credible threats** — SPE requires optimal play in every subgame, eliminating Nash equilibria propped up by empty threats.

11. Credible threats & commitment

✂ Question (fold →)	Answer
The reason credibility matters is that:	**opponents best-respond to what you'll ACTUALLY do, not what you say** — Rational opponents anticipate your real future action, so only credible commitments move them.
A store's "price-match guarantee" can be a commitment that:	**signals it won't be undercut, discouraging rivals from cutting prices** — The guarantee credibly commits to matching, reducing rivals' incentive to undercut.
A commitment device works by:	**changing your own future payoffs so a desired action becomes rational later** — By binding your future self, you make an otherwise-empty promise or threat credible.
A peace treaty with automatic monitoring is credible because it:	**makes violations detectable and costly, aligning incentives to comply** — Monitoring + penalties make compliance the rational choice — a credible commitment.
The paradox of commitment is that:	**limiting your future freedom can increase your power now** — Binding yourself can improve outcomes by making threats/promises credible.
An empty threat can still work if:	**the opponent mistakenly believes it is credible** — Perceived credibility (even if false) can deter — but it's risky if tested.
A promise is credible when:	**keeping it is in the promiser's own interest (or is enforced)** — Like threats, promises need aligned incentives or enforcement to be believed.
A threat that would HURT the threatener to carry out is:	**not credible** — If executing the threat is costly to the threatener, opponents discount it.
The synthesis of commitment is:	**sometimes tying your own hands is the strongest move you can make** — Credible self-binding can beat flexibility by shaping opponents' choices.
Backward induction is the tool used to CHECK whether a threat is credible because it:	**asks whether executing the threat is optimal at that future node** — You verify credibility by solving the future node: is executing the threat the best response there?
Reducing your OWN options can improve your outcome when it:	**convinces opponents you'll take a specific action** — Fewer options can be a credible signal, shifting the opponent's best response.
A manager who publicly stakes their reputation on a deadline is:	**using a commitment to make their resolve credible** — Publicly staking reputation raises the cost of backing down, making resolve credible.
Committing to a strategy BEFORE the opponent moves can create:	**a first-mover advantage** — Visible commitment before the rival moves can shift their best response in your favor.
A threat becomes MORE credible if the threatener:	**has arranged to benefit (or avoid loss) from carrying it out** — Aligning execution with the threatener's interest is what makes it credible.
A country that passes a law auto-triggering tariffs is using:	**a commitment device to make its trade threat credible** — The automatic trigger removes discretion, making the threat credible.
A student says "just threaten harder." The better advice is:	**"make the threat credible by aligning your incentives to carry it out"** — Credibility, not intensity, is what makes a threat effective.

Delegating a decision to a tough negotiator can be a commitment because:	**it credibly ties your hands to a firm position** — A committed agent limits your ability to soften, making the hard line credible.
Commitment and flexibility trade off because:	**commitment gains credibility but sacrifices the ability to adapt later** — You buy credibility with reduced future flexibility — a real strategic trade-off.
A credible threat is one that:	**the threatener would actually want to carry out if tested** — Credibility comes from it being in the threatener's interest to execute, not from how it's stated.
Cheap talk (non-binding words) is credible only when:	**the speaker's interests align with telling the truth** — Words carry weight only if the speaker gains from being truthful.
Which is a genuine commitment device?	**an automatic contract penalty that triggers without your discretion** — A device removes your discretion (automatic penalty), unlike mere intentions or announcements.
"Burning your bridges" can be strategically smart because it:	**removes your retreat option, committing you credibly** — Eliminating retreat makes your commitment to advance credible to opponents.
"Doomsday device" logic (from strategy) shows a threat is useless if:	**the opponent doesn't know about it or doubts it will trigger** — A commitment deters only if it's known and believed to trigger.
A signed contract with an automatic penalty for breaking it is a commitment device because it:	**changes your future payoff so keeping the deal is now your best move** — The penalty alters your payoffs so compliance becomes rational — making the promise credible.
The difference between a bluff and a credible threat is:	**a credible threat is backed by the incentive to actually carry it out** — A bluff lacks the incentive to execute; a credible threat has it.

12. Bargaining; ultimatum; fairness

✂ Question (fold →)	Answer
A stronger outside option makes a bargainer:	**able to demand a larger share** — A better alternative raises your floor and your leverage for a bigger share.
Repeated bargaining between the same parties encourages:	**fairer offers to protect the ongoing relationship** — Reputation and future dealings push toward fairer, sustainable offers.
"Splitting the difference" is popular because it:	**is a fair, simple focal point that ends disputes quickly** — It's a salient fairness heuristic that resolves bargaining fast.
The bargaining "surplus" is:	**the total value to be divided if the parties reach a deal** — The surplus is the joint gain available from agreement (over disagreement).
If both parties have strong, equal outside options, the deal tends toward:	**a roughly even split of the surplus** — Balanced alternatives lead to a balanced division of the surplus.
The "dictator game" differs from the ultimatum game because:	**the responder cannot reject — the proposer simply decides** — In the dictator game the recipient has no power to reject; giving reveals pure other-regard.
If your BATNA (walk-away option) improves during a negotiation, you can now:	**demand a larger share of the surplus** — A better alternative raises your floor and your leverage, so you can hold out for more.
In real ultimatum experiments, responders often:	**reject offers seen as unfair, even at a cost to themselves** — People reject unfair offers, sacrificing money to punish unfairness — contradicting pure self-interest.
The lesson of ultimatum experiments for negotiators is:	**offers perceived as unfair may be rejected even when accepting is the "rational" choice** — Account for fairness perceptions; lowball offers can blow up a deal.
Rejecting a 90/10 offer (keeping nothing) can be RATIONAL if the responder:	**values fairness or building a reputation for not being exploited** — If fairness or reputation enters payoffs, rejecting an unfair split is rational.
A party's "outside option" (BATNA) sets:	**the minimum they'll accept — they won't take less than walking away** — Your best alternative to a deal is your floor; you reject anything worse than it.
The ultimatum and dictator games together separate:	**strategic fairness (fear of rejection) from pure generosity** — Comparing the two isolates how much giving is strategic vs genuinely other-regarding.
Improving your BATNA before negotiating is smart because it:	**raises your floor and your share of the surplus** — A stronger alternative directly increases your bargaining power.
A fair reputation in bargaining is valuable because it:	**makes others more willing to deal with you in the future** — Being known as fair lowers others' guard and sustains beneficial future deals.
A party with NO outside option is:	**in a weak bargaining position (must take almost any deal)** — Without an alternative, your floor is near zero, so you have little leverage.
Alternating-offer bargaining (Rubinstein) shows that patience:	**is a source of bargaining power** — The more patient party (lower discounting) captures a larger share.

An "equal split" is often observed because it is:	**a fair focal point both sides can accept** — 50/50 is a salient, fair focal point people readily coordinate on.
In the ULTIMATUM game, one player proposes a split and the other:	**accepts (both get the split) or rejects (both get nothing)** — The responder can accept the proposed split or reject, leaving both with zero.
The bargaining range (zone of possible agreement) exists when:	**the buyer's max exceeds the seller's min** — A deal is possible only when the buyer's ceiling is above the seller's floor.
First-mover (proposer) advantage in the ultimatum game comes from:	**setting the terms the responder must accept or reject** — The proposer frames the split, capturing more if the responder is purely self-interested.
If delay is costly to both parties, they are incentivized to:	**reach agreement sooner rather than later** — Costly delay pushes both toward a quicker deal.
Backward induction predicts the proposer will offer:	**the smallest positive amount, expecting acceptance** — A rational responder accepts any positive amount, so the proposer offers the minimum.
The synthesis of bargaining is:	**leverage comes from your alternatives and patience, but fairness shapes what deals actually hold** — BATNA + patience set leverage; fairness perceptions determine which deals stick.
This gap between theory and behavior shows people value:	**fairness, not just their own monetary payoff** — Rejections reveal fairness is part of people's actual payoffs.
The concept that best explains WHY people reject unfair splits is:	**fairness preferences (people's payoffs include fairness)** — Fairness preferences extend payoffs beyond money, explaining costly rejection.

13. Auction formats; bidding strategy

✂ Question (fold →)	Answer
PRIVATE-value vs COMMON-value auctions differ in that:	**private values differ per bidder; common value is the same true value for all** — Private-value: each bidder has their own worth; common-value: one true worth all estimate.
To avoid the winner's curse, a bidder should:	**shade bids down to account for the bad news of winning** — Adjust bids downward because winning itself signals you over-estimated.
The "winner's curse" happens in COMMON-value auctions when:	**the winner is the one who most overestimated the item's value** — Winning means you bid highest — often because you over-estimated the common value, so you overpay.
In an ENGLISH (ascending) auction, bidders:	**openly raise bids until only one remains** — Prices rise openly until a single highest bidder remains.
An auctioneer wanting to discourage collusion among bidders might prefer:	**a sealed-bid format so bidders can't signal each other** — Sealed bids prevent the open signaling that facilitates bidder collusion.
In a FIRST-price auction, rational bidders typically:	**shade their bids below their true value** — You bid below value to leave a profit margin, trading win-probability against surplus.
Auctions are a form of MECHANISM DESIGN because the format:	**is a rule chosen to produce desired bidding behavior and outcomes** — The auction rules are designed to elicit particular (often truthful) behavior — mechanism design.
A DUTCH (descending) auction lowers the price until:	**a bidder accepts the current price** — The price falls until the first bidder claims the item at that price.
The strategic complexity of a FIRST-price auction comes from:	**having to guess others' bids to set your shading** — Optimal shading depends on beliefs about rivals' bids — a genuine strategic guess.
In a first-price SEALED-bid auction, the winner pays:	**their own bid (the highest bid)** — The highest sealed bid wins and pays exactly that bid.
An auction is a game because each bidder's outcome depends on:	**the bids of the other bidders too** — Who wins and what they pay depends on all bids — strategic interdependence.
In a second-price auction, the dominant strategy is to bid:	**your true value** — Bidding your true value is (weakly) dominant in a second-price auction — you never regret it.
The synthesis of auctions is:	**the format's rules determine the optimal bidding strategy — so read the rules first** — Strategy is format-dependent (truthful in Vickrey, shaded in first-price) — diagnose the rules first.
A reserve price in an auction is:	**a minimum below which the seller won't sell** — The reserve is the seller's floor; bids below it don't win.
Bidding your true value is safe in a Vickrey auction because:	**the price you pay doesn't depend on your own bid, only whether you win** — Since you pay the second price, shading your bid only risks losing without lowering the price.
The revenue-equivalence idea says that, under certain conditions,	**yield the same expected revenue to the seller** — Under standard assumptions, first-price, second-price, English, and Dutch auctions raise the

different auction formats:	same expected revenue.
Which format makes truthful bidding a dominant strategy?	**the second-price (Vickrey) auction** — Only the second-price/Vickrey (and equivalently English) auction makes truth-telling dominant.
Bidding ABOVE your value in a second-price auction risks:	**winning and paying more than the item is worth to you** — Overbidding can win at a second price above your value — a loss.
If you know your PRIVATE value is $50, in a second-price auction you should bid:	**$50** — Bid your true value ($50) — it's the weakly dominant strategy.
You value an item at $30. In a SECOND-price sealed auction, your best bid is:	**$30 (your true value)** — Bidding your true value is weakly dominant in a second-price auction.
The strategic goal of a bidder in any auction is to:	**win the item only when doing so is profitable, at the lowest feasible price** — Win only value-positive items, minimizing what you pay — the core bidding objective.
The winner's curse is WORSE when:	**bidders have very uncertain, noisy estimates of a common value** — More estimate noise means the winner more likely over-estimated — a bigger curse.
A government selling radio spectrum uses auctions mainly to:	**allocate it efficiently to those who value it most while raising revenue** — Auctions channel scarce resources to high-value users and raise public revenue.
The English and second-price auctions are strategically similar because both:	**lead bidders to effectively reveal their true value** — In both, staying in until your value / bidding your value is optimal — truthful revelation.
In a SECOND-price (Vickrey) auction, the winner pays:	**the second-highest bid** — The highest bidder wins but pays the runner-up's bid.

14. Designing rules so honesty pays; incentive compatibility

✂ Question (fold →)	Answer
Why can't every desirable outcome be implemented?	**impossibility theorems show some goal combinations conflict** — Results (e.g. Gibbard-Satterthwaite, Myerson-Satterthwaite) prove some combinations are impossible.
The designer's core question is:	**"what rules make the behavior I want the players' own best choice?"** — Mechanism design asks how to make the desired behavior individually rational.
A mechanism is "individually rational" when:	**every player is at least as well off participating as not** — Individual rationality (participation constraint) ensures no one is worse off for joining.
"Split the cake": one child cuts, the other chooses. This mechanism ensures:	**the cutter cuts fairly, because they get the piece the chooser leaves** — Cut-and-choose makes fair division the cutter's self-interest — an incentive-compatible mechanism.
The synthesis of mechanism design is:	**if you can't change the players, change the rules so honesty and cooperation pay** — Design incentives so self-interested play yields honesty, efficiency, and cooperation.
A matching mechanism (like assigning students to schools) aims for:	**stable, incentive-compatible assignments people won't want to game** — Good matching mechanisms (e.g. deferred acceptance) are stable and hard to game.
The connection between auctions and mechanism design is that an auction is:	**a mechanism whose rules are chosen to elicit desired bidding** — Auction formats are engineered mechanisms to produce efficient, often truthful, bidding.
The main challenge of mechanism design is that players have:	**private information and will act strategically on it** — Designers must elicit truthful behavior despite players' private info and strategic incentives.
A toll that charges drivers the cost they impose on others is designed to:	**align private choices with the socially efficient outcome** — Charging the external cost makes drivers internalize it — aligning private and social incentives.
A mechanism designer must balance:	**efficiency, incentive compatibility, and participation (you can't always have all)** — Impossibility results show tensions among efficiency, truthfulness, budget balance, and participation.
A mechanism that is easy to game will:	**produce outcomes based on manipulation rather than true preferences** — If gaming pays, reported behavior reflects manipulation, not real preferences.
"Deferred acceptance" (the Gale-Shapley algorithm) is famous for producing:	**stable matches where no pair would both prefer to break away** — Gale-Shapley yields stable matchings — no blocking pair wants to defect.
A mechanism is "incentive-compatible" when:	**players do best by acting honestly / revealing their true preferences** — Incentive compatibility means truthful behavior is each player's best strategy.

The "revelation principle" says any outcome achievable by a mechanism can be achieved by one where:	**players truthfully report their preferences** — Any implementable outcome has an equivalent incentive-compatible (truthful) direct mechanism.
A mechanism where reporting a HIGHER need always gets you more will be:	**exploited by everyone over-reporting their need** — If over-reporting always pays, everyone inflates claims — a non-incentive-compatible design.
The cut-and-choose method generalizes to fair division because it:	**uses each participant's self-interest to enforce fairness** — Self-interest polices fairness (the cutter fears the small piece) — a general mechanism idea.
A good mechanism aligns:	**individual incentives with the designer's desired social outcome** — The aim is to make self-interested play produce the intended collective result.
An incentive-compatible tax form would be one where taxpayers:	**have no advantage from misreporting their situation** — Incentive compatibility removes any gain from misreporting — honest filing is best.
A voting rule that makes voters better off voting sincerely is:	**strategy-proof (a form of incentive compatibility)** — Strategy-proof rules make sincere voting a dominant strategy — no gain from strategic voting.
Designing rules so "honesty is the best policy" prevents:	**players from gaming the system by misrepresenting themselves** — Incentive compatibility removes the payoff to strategic lying/gaming.
A mechanism where reporting a LOWER income always lowers your tax is likely NOT incentive-compatible because:	**people can gain by under-reporting their true income** — If under-reporting always pays, truthful reporting isn't a best response — so it's not incentive-compatible.
The reason mechanism design matters for society is that it:	**lets us build institutions where self-interest serves the common good** — Well-designed rules harness self-interest for collective benefit — the promise of mechanism design.
A referral or reward scheme that rewards honest reporting is applying:	**incentive-compatible mechanism design** — Designing rewards so honesty pays is exactly incentive-compatible mechanism design.
Mechanism design is often called "reverse game theory" because it:	**designs the RULES to produce a desired outcome, rather than analyzing a given game** — You start with the goal and design rules/incentives that make players' rational play achieve it.
The Vickrey (second-price) auction is a classic mechanism because it makes:	**truthful bidding a dominant strategy** — Its rules make revealing your true value dominant — a canonical incentive-compatible mechanism.

15. Common-pool / public-goods dilemmas

✂ Question (fold →)	Answer
Overfishing a shared ocean is an example of:	**a common-pool resource dilemma** — Each boat's incentive to fish more depletes the shared stock — a common-pool dilemma.
The "tragedy of the commons" describes:	**a shared resource being overused because each user ignores the cost to others** — Common-pool resources get depleted when each user maximizes personal use.
Clear boundaries around a shared resource help because they:	**define who may use it, reducing outsider overuse** — Defined access limits the user pool, easing management of the commons.
The free-rider problem is that people:	**enjoy a public good without paying for it, so it's under-provided** — Since you benefit whether or not you pay, individuals under-contribute — under-provision.
Elinor Ostrom showed that communities CAN manage commons through:	**shared rules, monitoring, and graduated sanctions** — Ostrom documented self-governing institutions (rules, monitoring, sanctions) that avert the tragedy.
A shared office fridge that everyone uses but no one cleans illustrates:	**a small tragedy of the commons** — Each person benefits from a clean fridge but is tempted to free-ride on cleaning — a mini commons dilemma.
Which is the BEST example of a common-pool resource dilemma?	**a shared groundwater basin that multiple farms pump from** — Shared, depletable groundwater with many users is a textbook common-pool resource.
The difference between a public good and a common-pool resource is:	**a public good is non-rival; a common-pool resource is rival (can be used up)** — Public goods aren't depleted by use; common-pool resources are (rival) but hard to exclude.
A collective-action problem arises when:	**individually rational choices lead to a collectively bad outcome for the group** — Like a multi-player PD: each person's rational choice harms the group.
Community monitoring works against the tragedy of the commons by:	**making defection visible so it can be sanctioned** — Visibility enables sanctions, which deter overuse — a core Ostrom mechanism.
The reason pure self-interest under-provides public goods is:	**the benefit is shared but the cost is individual, so each waits for others** — Shared benefit + individual cost = each defers, so the good is under-supplied.
A "matching grant" that funds a project only if enough people contribute uses:	**a threshold to overcome free-riding by making each contribution pivotal** — Thresholds make each contribution feel decisive, countering free-riding.
The key insight linking collective action to the PD is:	**both feature individually rational choices producing collectively worse outcomes** — Collective-action problems are multi-player prisoner's dilemmas.
One classic solution to the free-rider problem is:	**making contributions mandatory (e.g. taxes) to fund the public good** — Compulsory contributions overcome free-riding for public-good provision.
A vaccination program faces a collective-action problem because:	**individuals may rely on others' immunity while avoiding their own small cost/risk** — Herd immunity is a public good; each person is tempted to free-ride on others' vaccination.

Collective action is EASIER to sustain when:	**the group is small, repeated, and can monitor each other** — Small, repeated, observable groups can enforce cooperation (reputation + monitoring).
A "public good" is one that is:	**non-excludable and non-rival — everyone can benefit without using it up** — Public goods (clean air, national defense) are non-excludable and non-rival.
The larger the group, the collective-action problem usually becomes:	**harder, because each person's contribution feels negligible** — In big groups individual impact seems tiny, worsening free-riding.
A "cap-and-trade" system addresses a commons (emissions) by:	**limiting total use and letting the right to use be traded** — A cap sets total use; tradable permits allocate it efficiently — a commons mechanism.
Social norms help solve collective-action problems by:	**creating informal rewards for cooperating and costs for defecting** — Norms add reputational rewards/sanctions that make cooperating rational.
The overarching lesson of collective action is:	**shared benefits often require institutions to overcome individual free-riding** — Sustaining public goods/commons typically needs rules, monitoring, or incentives.
Voting is a collective-action problem because:	**each vote feels negligible, so individuals may free-ride on others voting** — The individual impact is tiny, so rational-choice logic predicts under-participation.
Graduated sanctions (Ostrom) mean:	**small penalties for first violations, larger for repeat offenders** — Escalating penalties deter defection while remaining fair — an Ostrom design principle.
Assigning private property rights can solve some commons problems by:	**giving owners an incentive to conserve their own resource** — Owners internalize conservation incentives, reducing overuse.
The synthesis of collective action is:	**the group's good needs mechanisms that make each person's cooperation individually worthwhile** — Design incentives/institutions so contributing is each individual's rational choice.

16. Apply strategy to real cases (climate/pricing/misinformation)

✂ Question (fold →)	Answer
A vaccine-uptake campaign can be improved by mechanism-design thinking such as:	**reducing the individual cost/effort so cooperating is each person's easy choice** — Lowering the personal cost aligns individual incentive with the public good.
To find where a real strategic situation will "settle," look for:	**the Nash equilibrium (no player gains by deviating alone)** — The Nash equilibrium predicts the stable resting point of the interaction.
When presenting a game-theoretic analysis of a values-laden issue, avoid:	**declaring one side morally "correct" via a payoff score** — Map incentives and trade-offs; do not convert a values dispute into a scored verdict.
To improve a bad equilibrium in the real world, you can try to:	**change the payoffs, add repetition, or design new rules** — Alter incentives, extend the horizon, or redesign the mechanism to shift the equilibrium.
To analyze a real situation strategically, the FIRST step is to:	**identify the players, their strategies, and their payoffs** — Frame any real case as players/strategies/payoffs before applying tools.
Climate-emissions negotiations resemble which game structure?	**a multi-player collective-action problem (each tempted to free-ride)** — Each country benefits from others cutting while it pollutes — a global collective-action problem.
A credible commitment in real policy (e.g. an automatic trigger law) works because it:	**removes discretion, making the threat or promise believable** — Removing discretion makes the commitment credible to other players.
Before recommending action on a real strategic dispute, the responsible final step is to:	**state the trade-offs and stakeholders' incentives rather than declare a single winner** — On value-laden cases, illuminate incentives and trade-offs; leave the value judgment to the deliberators (anti-evangelism).
The single biggest takeaway of NashForge is:	**outcomes emerge from interacting incentives — change the incentives to change the outcome** — Strategic outcomes flow from incentive structures; the lever for change is the incentives themselves.
The capstone habit for any real situation is to:	**map players, strategies, payoffs, and interdependence before judging what to do** — Set up the game (players/strategies/payoffs/interdependence) first, then reason about action.
A strategist warns against "zero-sum bias" in real disputes because it:	**makes people miss cooperative solutions that would help everyone** — Assuming a fixed pie forecloses win-win options — a common real-world error.
Repeated interaction in ongoing rivalries (firms, nations) means:	**reputation and future consequences shape today's strategic choices** — Ongoing relationships bring the shadow of the future into current decisions.
The difference between the Nash outcome and the cooperative outcome in a real dilemma tells you:	**how much value is lost to individually-rational behavior** — That gap measures the inefficiency of self-interested play — the target for intervention.
Applying strategy responsibly means remembering that payoffs in real cases:	**may include fairness, reputation, and values, not just money** — Real payoffs incorporate fairness/reputation/values — model them, don't reduce everything to cash.

Modeling misinformation spread as a game highlights that:	**sharing sensational content can be individually rewarding but collectively harmful** — Individual engagement rewards vs collective harm mirror a collective-action tension.
The value of game theory for real decisions is that it:	**clarifies incentives and likely responses so you can choose wisely** — It's a lens for anticipating strategic responses, not a moral oracle or crystal ball.
A social-media platform reducing "outrage" rewards is applying:	**mechanism design to realign individual incentives with collective well-being** — Changing what behavior is rewarded is mechanism design aimed at a better collective outcome.
A monopolist deterring entry with a credible capacity commitment illustrates:	**commitment as strategic advantage** — Sunk capacity credibly commits to fighting entrants — a real commitment strategy.
A strategist analyzing an arms race would emphasize that:	**both sides may build up despite both preferring mutual restraint (a PD)** — Arms races are PDs: each builds up defensively though mutual restraint is better for both.
A price war between two firms is best modeled as:	**a prisoner's-dilemma-like game (both cut, both worse off)** — Mutual price-cutting is individually tempting but jointly harmful — PD-like.
Two allies deciding how much to invest in a shared defense face:	**a collective-action problem (each tempted to under-contribute)** — Shared defense is a public good vulnerable to free-riding by each ally.
When two countries can both gain from trade, the situation is:	**non-zero-sum — cooperation can grow the pie** — Mutually beneficial trade creates value — a non-zero-sum opportunity.
When analyzing a contested real-world case, a strategist should present:	**the trade-offs and the different stakeholders' payoffs, not a single "right" verdict** — On value-laden cases, lay out incentives and trade-offs; don't score a 'winning' position.
A traffic-congestion problem is a commons dilemma addressable by:	**congestion pricing that makes drivers internalize the delay they impose** — Pricing the external cost aligns private driving choices with social efficiency.
When a case has genuine ethical disagreement, the strategist's job is to:	**illuminate the incentives and trade-offs each side faces** — Clarify the strategic landscape; leave the value judgment to the deliberators.

Keep going

You practiced **400 cards** from NashForge. Come back to the ones you missed — spaced-out practice makes them stick.

Keep exploring online

Play it now	More free books
`https://spark-and-anvil.org/play/nashforge`	`https://spark-and-anvil.org/books`

Scan a code with any phone camera, or type the address. Every app, story, and book is free, forever — no account, no tracking.