# Module 10 — Strategic Systems Integration

Game Theory

hbar.university

## Games Embedded in Institutions

### Position in Knowledge Graph

**System:** Game Theory **Structural Domain:** Strategic Systems Integration **Node:** 10.1

Up to this point in the course, a "game" has been an object specified in isolation: a triple $(N, S, u)$, or an extensive form, or a mechanism. But no actual strategic situation is ever specified in isolation. Every real game sits inside an **institution** — a set of rules, norms, legal constraints, enforcement mechanisms, and custom — that determines which game is being played at all. The institution specifies who the players are, what actions are legally available, what the payoffs look like after taxes and transfers, which strategies are prohibited, and what happens when someone breaks a rule. In this sense, institutions are a **meta-game** whose outcome is the selection of an object-level game. 10.1 opens Module 10 by making this observation precise: the formalism of games within games, the institutional rules as a higher-order game-selector, and the way institutions both constrain and create strategic possibilities. Everything in Module 10 — layered mechanisms (10.2), contracts (10.3), commitment devices (10.4), collective choice (10.5), strategic ecosystems (10.6) — depends on this reframing.

---

### Formal Definition

An **institution** $I$ is a tuple

$$I = (\mathcal{G}, R, E, \phi)$$

where:

- $\mathcal{G}$ is a set of candidate object-level games over a common player set $N$. Each $G \in \mathcal{G}$ is a strategic-form or extensive-form game as defined in Modules 1–3.
- $R$ is a set of **institutional rules**: constraints on player actions, admissibility conditions, and payoff transformations (taxes, rake, fines, transfers).
- $E$ is an **enforcement mechanism** specifying which deviations from $R$ are detectable, which are punished, and with what cost.
- $\phi : R \times E \to \mathcal{G}$ is the **game-selection function**: it takes the rules and the enforcement regime and returns the effective object-level game that the players face.

Players move in two phases:

1. **Institutional phase (meta-game).** A prior process — legislation, custom formation, casino licensing, civil-law development, market-platform design — selects $R$ and $E$. This is itself a game (possibly a collective-choice game of the form studied in 10.5).
2. **Object-level phase.** Given $\phi(R, E) = G$, the players play $G$ as an ordinary game.

The effective payoff function for the object-level game is

$$u_i^{\text{eff}}(s) = u_i^{\text{raw}}(s) - t_i(s) - \pi_i(s) \cdot \mathbb{1}[s \notin R_i]$$

where $t_i(s)$ is the institutional transfer (taxes, fees, rake), $R_i$ is the rule-admissible action set for player $i$, and $\pi_i(s)$ is the punishment if player $i$ violates rules and is caught by $E$.

A **well-defined institution** is one where $\phi$ is deterministic and the enforcement mechanism $E$ is credible (in the sense of 10.4): the prescribed punishments are actually carried out when violations occur. An **institutional equilibrium** is a pair $(s^*, R^*)$ where $s^*$ is a Nash equilibrium of $\phi(R^*, E)$ and $R^*$ is stable under the meta-game process that generates rules.

---

## Intuition

No one plays "poker" in the abstract. You play poker-at-the-Bellagio, or poker-on-PokerStars, or poker-at-a-home-game. These are structurally different games even though the card-ranking rules are identical, because the **institutions** are different: rake structures, player pools, enforcement of collusion norms, legal status of winnings, tax treatment, availability of real-time assistance, tournament structures, table selection rules. Each institution selects a different effective game.

The same is true of every other strategic situation. You do not play "oligopoly" — you play oligopoly-under-EU-competition-law, or oligopoly-under-US-antitrust, or oligopoly-in-an-unregulated-grey-market. You do not play "the dating game" — you play dating-in-1950s-America, or dating-on-Tinder, or dating-in-rural-Japan. The institutional context is not flavor. It determines the action set, the payoff function, the information structure, the credibility of threats, and the identity of the players.

The structural point is that **treating the institution as fixed is a modeling choice, not a fact**. For short-run questions (how should I play this hand?) the institution is rightly taken as given. For medium-run questions (should I play at the Bellagio or on PokerStars?) the institution is a strategic choice made by the player. For long-run questions (should this kind of poker be legal, taxed, licensed?) the institution is the object of strategic design. Module 10 is the module where the boundary of "the game" becomes a strategic variable.

Three observations make this precise:

1. **Institutions create games.** Without the Bellagio's rake structure, dealer, security, and licensing, the game "poker at the Bellagio" would not exist. The institution is not a modifier on a preexisting game; it is the condition of possibility for the game.
2. **Institutions are equilibria of higher-order games.** Rules do not fall from the sky. They are selected by a collective-choice process (legislation, custom, market competition among platforms) that is itself a game. This nesting is potentially unbounded: the rules of the legislative process are themselves an institution, etc.
3. **Institutional rules transform payoffs, strategies, and players.** A tax changes payoffs. A licensing rule changes who can play. An enforcement mechanism changes which strategies are credible. A disclosure rule changes the information structure. Every dimension of the game triple $(N, S, u)$ is institutionally contingent.

---

## Structural Role

10.1 is the opening reframing of Module 10. Every subsequent node in the module is a specialization of the institutional embedding:

- **10.2 Layered mechanisms** treats institutions as nested mechanism-design problems: the object-level mechanism is embedded in a higher-level regulatory mechanism.
- **10.3 Contracts and principal-agent** treats contracts as private institutions that selectively constrain actions and transform payoffs for a bilateral relationship.
- **10.4 Commitment and credibility** asks when institutional rules are self-enforcing — when the meta-game actually produces the intended object-level game.
- **10.5 Political economy and collective choice** opens the black box of the institutional-selection meta-game and analyzes it with its own tools (voting, bargaining, constitutional design).
- **10.6 Strategic ecosystems** treats the whole stack — institutions, object-level games, players, strategies — as a coevolving dynamical system.

Looking backward, 10.1 is where many threads from earlier modules are retied. Repeated-game folk theorems (Module 5) gave us a large set of possible equilibria; institutions are what select among them. Mechanism design (Module 6) gave us tools to shape outcomes by choosing rules; institutions are the *context* in which a mechanism designer operates. Evolutionary dynamics (Module 7) gave us adjustment processes; institutional change is one such process operating on a longer timescale. Bayesian games (Module 4) gave us type-contingent strategies; institutional rules often work precisely by constraining how types can be revealed or concealed.

---

## Mathematical Development

### Institutional transformation of games

Let $G^{\text{raw}} = (N, S^{\text{raw}}, u^{\text{raw}})$ be a "bare" game before any institutional overlay. An institution $I = (R, E, \phi)$ produces a transformed game $G^{\text{eff}} = \phi(R, E)$ by three operations:

1. **Strategy restriction.** $S_i^{\text{eff}} = S_i^{\text{raw}} \cap R_i$. Institutional rules delete some strategies from the action set (prohibition: collusion in auctions, insider trading, use of solvers during live hands).

2. **Payoff transformation.** $u_i^{\text{eff}}(s) = u_i^{\text{raw}}(s) - t_i(s)$. Taxes, fees, rake, transfers are subtracted; subsidies are added.

3. **Enforcement overlay.** For strategies that violate rules but are not physically prevented, an enforcement term is added: $u_i^{\text{eff}}(s) = u_i^{\text{raw}}(s) - t_i(s) - p_i(s)\pi_i(s)$ where $p_i(s)$ is the detection probability and $\pi_i(s)$ is the punishment if caught.

The institutional transformation is generally **not a bijection on game space**: many raw games can map to the same effective game (informational indifference at the object level), and one raw game can map to many effective games (institutional heterogeneity). The transformation is also not invertible from observing play: you cannot in general recover $u^{\text{raw}}$ from observing $u^{\text{eff}}$ unless you separately observe the institution.

### The nested-game formalism

Let $\mathcal{I}$ be a set of candidate institutions and $\mu : \mathcal{I} \to \Delta(\mathcal{G})$ be a **meta-game strategy** for the institutional phase: given a state of the meta-game, what distribution over institutions is the output? An **institutional equilibrium** is a fixed point

$$(I^*, s^*) \quad \text{with} \quad I^* \in \arg\max_{\text{meta}} \sum_i u_i^{\text{meta}}(I, s^*(I)), \qquad s^*(I) \in \text{NE}(\phi(I))$$

where the meta-game players (legislators, platform designers, custom-makers) anticipate the object-level equilibrium $s^*(I)$ when choosing their institution $I$. This is a two-level fixed-point problem: the institution is chosen taking the object-level equilibrium as given, and the object-level equilibrium is computed taking the institution as given. Solving both levels simultaneously is the conceptual heart of institutional analysis.

### Institutional stability: self-enforcement

An institution is **self-enforcing** if, once in place, no coalition of players has both the incentive and the ability to change it. Formally, let $C \subseteq N$ be a coalition and $u_C^{\text{meta}}$ be their meta-game payoff. $I^*$ is self-enforcing if

$$\forall C \subseteq N, \quad u_C^{\text{meta}}(I^*) \geq \max_{I' \in \text{reach}(I^*, C)} u_C^{\text{meta}}(I')$$

where $\text{reach}(I^*, C)$ is the set of institutions that coalition $C$ could unilaterally bring about given the current state. This condition ties 10.1 to the folk-theorem machinery of 5.4 and to the core of cooperative game theory: a self-enforcing institution is analogous to a core allocation in cooperative games.

### The constitutional layer

Above any specific institution sits the **constitutional** layer that governs how institutions can be changed. This is a meta-meta-game. The sequence can in principle continue indefinitely, but in practice it terminates at a layer that is **evolutionarily stable** (7.4) rather than strictly self-enforcing: the constitutional layer is held in place not by incentive compatibility at its own level but by the cost of coordinated change. This is why constitutions are so rarely rewritten even when no one argues they are optimal.

---

## Worked Example

### Casino-rake-as-institution

Consider heads-up no-limit hold'em cash as the object game. Raw game: zero-sum two-player poker, payoff $u_i^{\text{raw}}(s)$ is expected winnings.

Institution: a Las Vegas casino with a 10% rake capped at $5 per hand, a minimum buy-in, a maximum buy-in, a prohibition on real-time assistance (no solvers, no phones at the table), a prohibition on collusion, a dealer providing neutral administration, and an enforcement mechanism consisting of floor staff, surveillance cameras, and ejection / banning / prosecution for violations.

**Effective game transformation:**

1. **Strategy restriction.** $S_i^{\text{eff}}$ excludes strategies involving phone consultation, explicit communication with other players in a multiway game, and marked cards. The action space is narrower.
2. **Payoff transformation.** $u_i^{\text{eff}}(s) = u_i^{\text{raw}}(s) - \text{rake}(s)$. The game is no longer zero-sum between the two players — the house extracts a third-party payoff. The effective game is **constant-negative-sum for the players**.
3. **Enforcement overlay.** Phone use during a hand is physically not prevented (you could pull out your phone), but the enforcement term $-p(s)\pi(s)$ where $p$ is the detection probability (essentially 1, cameras and floor staff) and $\pi$ is the punishment (ejection, loss of comp, possible forfeiture of winnings) is large. So the effective payoff is large-negative for phone-use strategies even though they are technically in $S^{\text{raw}}$.

**Strategic consequences.** The rake means that weak players lose faster than in a zero-sum game; the edge required for a winning player rises. The prohibition on real-time assistance means that the relevant notion of "strategy" is a human's memorized strategy, not a solver's dynamically computed strategy — the game is about human implementation, not pure game-theoretic optimality. The buy-in cap limits the effective stack size and the scope of deep-stack strategy. The collusion prohibition rules out multiway coordinated play that would be possible in an unregulated game.

**Comparison with online poker.** The same raw game under PokerStars' institution has a different rake structure (smaller per hand, percentage of pot up to a cap), different enforcement (software-detected collusion, but no floor staff for angle shooting), different strategy restrictions (HUDs allowed on some sites, banned on others), and different player-pool dynamics. The effective game is structurally distinct even though the card-ranking rules are the same. A player who is winning at PokerStars may be losing at the Bellagio or vice versa, not because they are "worse at poker" but because they are playing a different effective game.

**Meta-game.** The casino selects its rake structure, rule set, and enforcement regime to maximize its own payoff (casino profit) subject to the constraint that players must voluntarily participate. This is an institutional design problem — a mechanism-design problem in the sense of Module 6, but with the meta-game feature that the casino must anticipate the object-level equilibrium to know whether players will actually show up. Too much rake and the game dies; too little and the casino loses money. The observed rake structures across major card rooms are a revealed equilibrium of this meta-game.

---

## Poker Anchor

### Concept mapping

| Formal object | Poker instantiation |
|---|---|
| Raw game $G^{\text{raw}}$ | Abstract poker: card dealing, betting rounds, hand rankings, zero-sum transfer |
| Institution $I$ | The venue: Bellagio, PokerStars, a home game, Macau high-stakes, a bar tournament |
| Rules $R$ | Rake, buy-in caps, no phones, no collusion, tournament structure, seating |
| Enforcement $E$ | Dealer, floor staff, surveillance, software detection, community norms |
| Transfer $t_i(s)$ | The rake — casino's cut per hand |
| Effective game $G^{\text{eff}}$ | The actual game that sits in front of you when you sit down |
| Meta-game | Casino licensing, legislation, platform competition, venue selection by players |

### Concrete situation: rake structure and winning threshold

At a \$2/\$5 NLH cash game with 10% rake capped at \$5, the house's per-hand extraction is significant in small-to-medium pots and negligible in large pots. A break-even player in the zero-sum abstract game loses at the rake rate; a small winner becomes a break-even player; a large winner becomes a small winner. The winning threshold — how good you must be to profit — is shifted by the institution. Move the same game to an online site with 5% rake capped at \$3 and the threshold drops. Move it to Macau where the rake is nominally 2.5% and the threshold drops further. Move it to a home game with no rake and the threshold is exactly the zero-sum threshold.

The same abstract skill (ability to find the Nash equilibrium strategy and execute it) translates to very different bankrolls under different institutions. This is why serious players select venues as carefully as they select spots within a hand.

### Concrete situation: online vs live as different games

Live poker and online poker are structurally different games not because the cards are different but because the institutions are different. Live poker has physical tells, table-captain dynamics, live-ratio pacing (slower play, fewer hands per hour), and physical presence norms. Online poker has multi-tabling, HUD availability, faster pacing, and weaker enforcement of some norms. A live specialist who is crushing \$5/\$10 live can easily be losing at \$5/\$10 online because the effective games are different even though the raw game is the same. The transition between the two is a change of institution, not a change of difficulty.

**What this understanding buys you**

**The choice of where to play is a strategic choice on par with how to play.** Many serious poker players treat venue selection as a tactical question (availability, comfort, convenience). Recognizing that venues are institutions that select different effective games promotes this to a strategic question. The right venue for a given player is the one whose institutional transformation is most favorable to that player's actual skill profile. A player who reads physical tells well should play live; a player who exploits HUD data well should play online; a player who specializes in exploiting weak recreational players should play in a venue with a high recreational-player ratio.

**Complaining about the institution is a category error.** When a player says "the rake is too high" or "PokerStars is rigged" they are complaining about the institutional terms of the game. These terms are not a feature of the raw game; they are the institution. The correct response is not to complain but to either (a) accept the institution and recalibrate your expected value, or (b) switch institutions, or (c) participate in the meta-game that sets institutional terms (venue competition, community pressure, regulation). These are three different strategic responses to the same observation.

**Institutional arbitrage is a real edge.** Differences in rake structure, player-pool quality, and enforcement across venues create arbitrage opportunities for players who are willing to travel, play online, or specialize in a particular venue type. The very best professionals know the institutional landscape of poker across continents and select venues accordingly. This is 10.1 in action.

---

## Cross-Links

- **6.1 Mechanism design** — institutions as large-scale mechanism-design problems.
- **5.4 Folk theorem** — which equilibria of the raw game survive institutional selection.
- **10.2 Layered mechanisms** — the nested-mechanism structure of institutional design.
- **10.4 Commitment and credibility** — why enforcement mechanisms work (or fail).
- **10.5 Political economy and collective choice** — the meta-game that produces institutions.
- **7.4 Evolutionarily stable strategies** — the analogous notion of institutional stability.
- **Economics / institutional economics** — North, Ostrom, Acemoglu: institutions as rules of the game.

---

## Structural Compression

**Invariant:** Every real strategic situation is the object-level component of an institutional embedding; the institution determines which game is being played by restricting strategies, transforming payoffs, and enforcing rules; analysis must account for both the object-level game and the meta-game that produces the institution.

**Mechanism:** The institution transforms a raw game $G^{\text{raw}} = (N, S^{\text{raw}}, u^{\text{raw}})$ into an effective game $G^{\text{eff}}$ by deleting prohibited strategies, subtracting institutional transfers, and adding enforcement-cost terms for rule-violating strategies; the institution itself is selected by a meta-game whose equilibrium is a self-enforcing institutional fixed point.

**Cross-System Echo:** This is structurally identical to the **layered protocol stacks** of computer networking (each layer's "game" is defined by the rules of the layer below it) and to the **nested gauge groups** of theoretical physics (each level of symmetry is a constraint inherited from a higher level). In both cases, the object at a given level is only well-defined given a choice of the level above. Institutions are the gauge-fixing of strategic situations: they select a particular representative from an equivalence class of formally indistinguishable object-level games. This theme will recur in every remaining node of Module 10.

---

## Exercises

1. Give three examples of economic situations where the same raw game is embedded in two different institutions, and describe how the effective game differs under each. At least one example should be from poker.
2. Let $G^{\mathrm{raw}}$ be the prisoner's dilemma with payoffs $(C, C) = (3, 3), (D, C) = (5, 0), (C, D) = (0, 5), (D, D) = (1, 1)$. Define an institution that imposes a fine of $f$ on any player who plays $D$, detected with probability $p$. Write the effective payoff matrix. Find the smallest $pf$ such that $(C, C)$ is a Nash equilibrium of the effective game.
3. Suppose the meta-game of institutional choice has two candidate institutions $I_1$ and $I_2$. The object-level equilibrium under $I_1$ gives players payoffs $(4, 4)$ and under $I_2$ gives $(5, 2)$. Under simple majority voting among the two players, which institution is chosen? Under unanimity? What does this tell you about the institutional-selection meta-game's sensitivity to voting rules?
4. Consider online poker with a HUD (heads-up display) allowed versus banned. Write down the effective strategy space in each case and argue that the allowed-HUD institution is not Pareto-dominant even though it gives each player more information.
5. Define "self-enforcing institution" precisely and show that grim-trigger-based enforcement of a rule against collusion can fail to be self-enforcing if the detection probability is below a threshold. Compute the threshold in a simple two-player setting.
6. A casino charges rake $r$ per pot. Show that there is a critical $r^*$ above which the highest-skill player at the table prefers to leave (because their post-rake edge goes negative). What does $r^*$ depend on, and why does the casino want to set $r < r^*$?
7. **Argue, structurally,** why the distinction between "the game" and "the institution in which the game is embedded" is not merely a labeling convenience but a genuine ontological feature of strategic situations. Give one example from outside game theory where failing to distinguish the two leads to a substantive mistake.

# Layered Mechanisms

## Position in Knowledge Graph

**System:** Game Theory **Structural Domain:** Strategic Systems Integration **Node:** 10.2

Module 6 introduced mechanism design as a problem of choosing the rules of a game to achieve a desired outcome. 10.1 showed that every game sits inside an institution. 10.2 brings these together: in most real-world settings, a mechanism designer does not operate in a vacuum but within an **outer mechanism** imposed by a regulator, a market structure, or a set of social norms. The result is **mechanism design within mechanism design** — a layered structure where the inner mechanism is optimized subject to the constraints imposed by the outer mechanism, and the outer mechanism is chosen anticipating the inner mechanism's behavior. This layered structure is ubiquitous: auctions run inside regulated markets, poker tournaments run inside casino licensing regimes, corporate governance structures sit inside securities law, and electoral systems sit inside constitutional frameworks. 10.2 formalizes the nesting, characterizes optimality at each layer, and identifies the new phenomena that arise from the interaction between layers.

---

## Formal Definition

A **layered mechanism** is a pair $(M^{\text{out}}, M^{\text{in}})$ where:

- $M^{\mathbf{out}} = (N^{\mathbf{out}}, \Sigma^{\mathbf{out}}, g^{\mathbf{out}})$ is the **outer mechanism** (regulator, platform, constitution). Here $N^{\text{out}}$ is the set of outer-level agents (regulators, platform designers), $\Sigma^{\text{out}}$ is their message space, and $g^{\text{out}} : \Sigma^{\text{out}} \to \mathcal{M}$ maps outer messages to a set of **permitted inner mechanisms** $\mathcal{M}$.
- $M^{\mathbf{in}} = (N^{\mathbf{in}}, \Sigma^{\mathbf{in}}, g^{\mathbf{in}})$ is the **inner mechanism** (auction, poker game, contract), with $M^{\text{in}} \in \mathcal{M}$ constrained by the outer mechanism's output.

The **sequential structure** is:

1. Outer-level agents choose $\sigma^{\text{out}} \in \Sigma^{\text{out}}$, producing the regulatory framework $g^{\text{out}}(\sigma^{\text{out}}) = \mathcal{M}' \subseteq \mathcal{M}$.
2. Inner-level designer chooses $M^{\text{in}} \in \mathcal{M}'$ to maximize their objective subject to outer constraints.
3. Inner-level agents play the game induced by $M^{\text{in}}$.

The **layered equilibrium** is a triple $(\sigma^{*\text{out}}, M^{*\text{in}}, \sigma^{*\text{in}})$ where:

$$\sigma^{*\text{in}} \in \text{BNE}(M^{*\text{in}}), \qquad M^{*\text{in}} \in \arg\max_{M \in \mathcal{M}'} V^{\text{in}}(M, \sigma^{*\text{in}}(M)), \qquad \sigma^{*\text{out}} \in \arg\max V^{\text{out}}(\mathcal{M}', M^{*\text{in}}(\mathcal{M}'), \sigma^{*\text{in}})$$

where $V^{\text{in}}$ and $V^{\text{out}}$ are the inner designer's and outer designer's objective functions, and each level optimizes taking the lower levels' equilibrium behavior as given. This is a **three-level principal-agent problem**: the outer principal (regulator) designs constraints, the inner principal (mechanism designer) designs the mechanism, and the agents play.

---

## Intuition

Consider an auction house that sells spectrum licenses under government regulation. The government (outer mechanism) sets rules: reserve prices, anti-collusion provisions, eligibility criteria, spectrum caps. The auction house (inner mechanism) designs the specific auction format — sealed-bid, ascending, combinatorial — subject to these constraints. The bidders (agents) play the resulting auction game.

Neither layer alone captures the full strategic picture. Analyzing the auction in isolation misses the regulatory constraints that restrict the auction designer's choices. Analyzing the regulation in isolation misses the strategic behavior of the auction designer who optimizes within the regulatory framework. The layered structure is essential to understanding the outcome.

The same nesting appears everywhere:

- **Poker:** The casino (outer mechanism) sets rake, buy-in caps, game-selection rules, and enforcement. The tournament director (inner mechanism) designs blind structures, payout tables, and ICM-relevant rules. The players play within both layers.
- **Markets:** The regulator (outer) sets disclosure requirements, capital adequacy rules, and anti-manipulation provisions. The exchange (inner) designs the order book, matching algorithm, and fee structure. Traders trade.
- **Organizations:** The legal system (outer) sets employment law, fiduciary duties, and tax rules. The firm (inner) designs incentive contracts, reporting structures, and promotion criteria. Employees work.

The key phenomenon of layered mechanisms is **strategic propagation between layers**. A change at the outer layer propagates through the inner designer's optimization to the agents' equilibrium behavior. Conversely, information about agent behavior propagates upward through the inner designer's reports (or regulatory audits) to the outer layer. This bidirectional flow means that optimizing at one layer in isolation is typically suboptimal: the outer designer must anticipate the inner designer's strategic response, and the inner designer must anticipate the agents' strategic response. This is the **nested principal-agent problem** that 10.3 will specialize to bilateral contracts.

---

## Structural Role

10.2 occupies the structural position between 10.1 (the general concept of institutional embedding) and 10.3 (the specific bilateral case of contracts). It is the module's theory of mechanism-design interaction:

- **From 6.1–6.3:** Mechanism design gave us the tools for a single layer — VCG, revenue equivalence, optimal auctions. 10.2 shows what happens when you stack these tools.
- **To 10.3:** Contracts are the simplest layered mechanism — an outer legal system constraining an inner bilateral agreement. The principal-agent theory of 10.3 is a special case of the layered formalism.
- **To 10.6:** Strategic ecosystems are the limit of many interacting layers. 10.2 provides the two-layer case that the ecosystem generalizes.

The structural novelty of 10.2 is that **the inner mechanism's optimality is constrained optimality**, and the binding constraints are themselves the output of a strategic process. This means that the "optimal mechanism" at the inner layer depends on the outer mechanism, and the "optimal regulation" at the outer layer depends on the inner mechanism's response. The fixed-point nature of this interaction is what makes layered mechanism design harder than single-layer design.

---

## Mathematical Development

### The constrained mechanism-design problem

The inner designer solves:

$$\max_{M^{\text{in}} \in \mathcal{M}'} V^{\text{in}}(M^{\text{in}}, \sigma^*(M^{\text{in}}))$$

subject to: - $M^{\text{in}} \in \mathcal{M}'$ (outer-mechanism constraint). - $\sigma^*(M^{\text{in}}) \in \text{BNE}(M^{\text{in}})$ (incentive compatibility at the agent level). - Participation constraints: agents must voluntarily participate.

Without the outer constraint, this is the standard Myerson optimal auction problem (6.3). With the outer constraint, the solution may be distorted. For example, if the outer mechanism imposes a minimum reserve price above the Myerson optimal reserve, the inner mechanism must use the suboptimal reserve, and both expected revenue and allocative efficiency may suffer.

**Regulatory design: the outer layer's problem**

The outer designer (regulator) solves:

$$\max_{\mathcal{M}' \subseteq \mathcal{M}} W(\mathcal{M}', M^{*\text{in}}(\mathcal{M}'), \sigma^*(M^{*\text{in}}))$$

where $W$ is the regulator's welfare function (social welfare, consumer surplus, revenue, etc.) and $M^{*\text{in}}(\mathcal{M}')$ is the inner designer's best response to the constraint set $\mathcal{M}'$.

**Theorem (Layered Revenue-Welfare Tradeoff).** Let $W$ be a weighted social welfare function and $V^{\text{in}}$ be the inner designer's revenue. For any outer constraint set $\mathcal{M}'$, if the inner designer maximizes revenue subject to $\mathcal{M}'$, and the outer designer maximizes welfare, then the equilibrium outcome satisfies:

$$W(\mathcal{M}'^*, M^{*\text{in}}) \leq W^{\text{first-best}}$$

with equality only if the inner designer's revenue-maximizing mechanism within $\mathcal{M}'^*$ happens to coincide with the welfare-maximizing mechanism. Generically, there is a strict welfare loss from the layered structure — the "cost of delegation."

**Information flow between layers**

Define the **information partition** at each layer: - Agents observe their own types $\theta_i$ and the mechanism rules $M^{\text{in}}$. - The inner designer observes the outer constraint $\mathcal{M}'$ and (possibly) aggregate statistics of agent behavior. - The outer designer observes the inner mechanism choice $M^{\text{in}}$ and (possibly) the allocation outcome.

The flow is:

$$\theta_i \xrightarrow{\text{agents}} \sigma_i^{\text{in}} \xrightarrow{\text{inner mechanism}} \text{outcome} \xrightarrow{\text{reports}} \text{outer designer}$$

Each arrow is a strategic communication channel: agents may lie to the inner mechanism (standard IC problem), and the inner designer may distort reports to the outer designer (regulatory evasion). The layered mechanism must be incentive-compatible at **both** levels simultaneously, which is a strictly harder requirement than IC at either level alone.

**The tournament-within-casino layered structure**

Formalize a poker tournament as a layered mechanism. Let the casino be the outer mechanism with rules $R^{\text{out}} = (\text{rake } r, \text{max entries } \bar{n}, \text{guarantee } G)$. The tournament director is the inner mechanism choosing blind structure $B = (b_1, b_2, \ldots, b_T)$, payout structure $P = (p_1, \ldots, p_k)$, and starting stack $c_0$, subject to $R^{\text{out}}$.

The inner designer's objective: maximize expected entries (which determines the inner mechanism's revenue $n \cdot$ buy-in) subject to the casino's rake and guarantee constraints. The agents (players) choose whether to enter and how to play, given the full specification $(B, P, c_0)$.

The ICM (Independent Chip Model) introduces a nonlinear payoff transformation $u_i^{\text{ICM}}(c_i, \mathbf{c}_{-i}, P)$ where $c_i$ is player $i$'s chip count and $P$ is the payout structure. This transformation is a feature of the inner mechanism (the tournament structure), not of the raw poker game. It creates strategic incentives — tighter play near the bubble, risk aversion even for risk-neutral players — that exist only because of the layered mechanism.

---

## Worked Example

### Spectrum auction under regulatory caps

**Outer mechanism (regulator).** The FCC imposes: (i) a spectrum cap of 100 MHz per bidder, (ii) a minimum reserve price of $10M per license, (iii) a prohibition on bidder collusion, (iv) a requirement of simultaneous ascending auction format.

**Inner mechanism (auctioneer).** Subject to these constraints, the auctioneer chooses: activity rules, bid increments, package-bidding provisions, and information-revelation policy (which bids are publicly visible during the auction).

**Agents (telecom firms).** Each firm has a private valuation $\theta_i$ for each license and bids strategically in the ascending auction.

**Without the outer mechanism.** The auctioneer would run an unconstrained Myerson-optimal auction: set individual reserve prices at $r_i^* = \theta_0 + (1 - F(\theta_0))/f(\theta_0)$ (virtual-value formula from 6.3), use a sealed-bid format for revenue maximization, and impose no spectrum caps.

**With the outer mechanism.** The spectrum cap constrains the allocation space (no firm can win too much). The minimum reserve may be above or below the Myerson optimal reserve. The format constraint (ascending auction) sacrifices some revenue relative to optimal sealed-bid but has transparency advantages the regulator values. The anti-collusion constraint affects information policy: showing too many bids may facilitate tacit collusion even without explicit communication.

**Layered equilibrium.** The auctioneer, constrained to ascending format with caps, chooses tight activity rules and limited information revelation to maximize revenue. The firms bid, adjusting for the caps (which reduce demand for very large firms and create opportunities for smaller firms). The regulator, anticipating this, chose the caps to balance efficiency (large firms may be more efficient) against competition (market concentration is socially costly). The outcome is neither the revenue-maximizing outcome nor the welfare-maximizing outcome, but the equilibrium of the layered system.

**Welfare loss from layering.** Suppose the Myerson optimal auction without caps generates expected revenue $500M and welfare $800M. The constrained auction generates revenue $450M and welfare $750M. The $50M revenue loss is the "cost of regulation" borne by the auctioneer; the $50M welfare loss is the cost of the caps. But the regulator values competition, which the caps promote, so the regulator's augmented welfare may be higher. The tradeoff is intrinsic to the layered structure.

---

## Poker Anchor

### Concept mapping

| Formal object | Poker instantiation |
|---|---|
| Outer mechanism $M^{\text{out}}$ | Casino rules: rake, game-spread rules, licensing, tournament overlay/guarantee |
| Inner mechanism $M^{\text{in}}$ | Tournament structure: blind schedule, payout table, starting stacks, clock, late registration |
| Outer designer | Casino management, gaming commission |
| Inner designer | Tournament director |
| Agent-level game | The poker hands played within the tournament |
| Layered constraint | Casino rake constrains tournament prize pool; blind structure constrains hand-level strategy |
| ICM transformation | Payout structure creates nonlinear chip-to-dollar mapping — a meta-game over the hand game |

#### Concrete situation: blind structure as inner mechanism design

A tournament director designs a 12-level blind structure for a $200 buy-in daily tournament. The casino's outer constraints: the tournament must finish within 5 hours (table rental cost), the rake is 15% ($30 of the $200 goes to the house), and a $10,000 guarantee must be met. Given these constraints, the director chooses starting stack (15,000 chips), blind levels (20 minutes each), and payout structure (top 15% of field).

The blind structure is a mechanism that determines the pace of elimination and the relative importance of skill versus variance. Fast blinds (turbo) increase variance, attract recreational players (who enjoy the "lottery" feel), and finish within the time constraint. Slow blinds increase the skill component, attract regulars, but risk running over the time limit.

The inner designer's tradeoff: fast blinds fill the guarantee (more recreational entries) but reduce the skill edge that attracts regulars; slow blinds attract regulars but may not fill the guarantee if the recreational pool is thin. This is a constrained mechanism-design problem — the inner designer optimizes within the outer constraints.

The ICM layer on top of the payout structure creates a second mechanism layer. Near the money bubble, a player's optimal strategy depends not on chip EV (the hand-level game) but on dollar EV (the tournament-level game), which is a nonlinear function of all stack sizes and the payout table. This is mechanism design within mechanism design: the payout structure transforms hand-level poker into a different strategic game with different equilibria.

#### What this understanding buys you

**Tournament selection is mechanism analysis.** When choosing which tournament to enter, a skilled player should analyze the layered mechanism: how does the blind structure interact with the payout table? How does the rake constrain the overlay? How does the starting stack interact with the blind schedule to determine how many "playable" levels exist before the game becomes a push-fold exercise? These are not fuzzy qualitative questions — they are quantifiable properties of the layered mechanism.

**ICM awareness is layer awareness.** Most intermediate players learn ICM as a "tool" — plug in stack sizes, get a payout equity. The deeper insight is that ICM is a **layer** — it is the transformation that the payout mechanism imposes on the hand game. Understanding this lets you see when ICM matters (near pay jumps, short stacks, final table) and when it does not (deep in the tournament, far from the money). It also lets you see that ICM is not a law of nature but a consequence of a specific payout structure — change the payout structure (winner-take-all vs flat) and ICM changes correspondingly.

**Regulatory arbitrage between layers.** Different casinos impose different outer constraints, leading to different inner mechanisms. A player who understands the layered structure can identify which casino's tournament program produces the most favorable inner mechanism for their skill profile. This is 10.2 applied to bankroll optimization.

---

## Cross-Links

- **6.1–6.3 Mechanism design** — the single-layer theory that 10.2 stacks.
- **10.1 Games embedded in institutions** — the general nesting framework; 10.2 specializes to mechanisms within mechanisms.
- **10.3 Contracts and principal-agent** — the bilateral special case of layered mechanisms.
- **4.3 Bayesian Nash equilibrium** — the solution concept at the agent level of each layer.
- **10.6 Strategic ecosystems** — the many-layer generalization.
- **Economics / regulation theory** — Laffont-Tirole regulatory economics as layered mechanism design.

---

## Structural Compression

**Invariant:** Real mechanisms are never designed in isolation; every mechanism sits inside an outer mechanism that constrains the designer's choices, creating a layered optimization problem where each layer's equilibrium depends on the others, and the overall outcome is the fixed point of the nested strategic interaction.

**Mechanism:** The outer mechanism constrains the set of available inner mechanisms; the inner designer optimizes within this constraint set; agents play the resulting game; information and incentives propagate bidirectionally between layers; the layered equilibrium is the simultaneous solution of all layers' optimization problems.

**Cross-System Echo:** Layered mechanisms are structurally identical to **compiler optimization within hardware constraints** in computer science (the compiler optimizes code subject to instruction-set constraints imposed by the chip architecture, which is itself designed anticipating compiler behavior) and to **gene regulation within cellular constraints** in molecular biology (transcription factors optimize gene expression within the regulatory architecture of the cell, which is itself the product of evolutionary selection on the regulatory architecture). In each case, the optimization at the inner level is meaningful only given the constraints at the outer level, and the constraints at the outer level are meaningful only given the response at the inner level.

---

## Exercises

1. Formalize a layered mechanism where the outer mechanism is a government tax on auction revenue and the inner mechanism is a first-price sealed-bid auction. Derive the optimal reserve price as a function of the tax rate and show that it differs from the Myerson optimal reserve.

2. A casino sets rake $r$ and a tournament director designs payout structure $P$. Show that for a fixed buy-in $B$, the prize pool is $(1 - r) \cdot B \cdot n$ where $n$ is the number of entrants, and that the payout structure $P$ is constrained by this total. Derive the ICM equity of a player with fraction $f$ of total chips in a winner-take-all versus a flat-payout structure.

3. Consider a two-layer mechanism where the outer layer imposes a maximum price (price cap) and the inner layer is a monopolist choosing quantity. Show that the layered equilibrium differs from both the unconstrained monopoly outcome and the competitive outcome. Characterize the welfare as a function of the price cap.

4. Prove that adding an outer-layer constraint to a revenue-maximizing inner mechanism can never increase inner revenue (the inner designer's problem is weakly harder with more constraints). State conditions under which it strictly decreases revenue.

5. In a poker tournament with payout structure $P = (50\%, 30\%, 20\%)$ for three players, compute the ICM equity for each player as a function of chip stacks $(c_1, c_2, c_3)$ with $c_1 + c_2 + c_3 = T$. Show that ICM equity is concave in own chip count and derive the implication for risk-taking at the final table.

6. Two mechanisms are layered: the outer mechanism is a Vickrey auction selecting which of two firms gets the right to run a monopoly, and the inner mechanism is the winning firm's pricing decision. Compare the layered equilibrium to a direct mechanism where the regulator runs both the auction and the pricing. Under what conditions does the direct mechanism dominate?

7. **Argue, structurally,** why ICM is not a feature of poker but a feature of the tournament mechanism layered on top of poker. Identify what changes about ICM if the payout structure changes, and explain why this shows that ICM is a property of the layer, not the game.

# Contracts and Principal-Agent Problems

## Position in Knowledge Graph

**System:** Game Theory **Structural Domain:** Strategic Systems Integration **Node:** 10.3

A contract is a **private institution** (10.1) that binds two parties to a set of conditional obligations, transforming their payoffs and constraining their strategies. The principal-agent problem is the canonical strategic situation in which contracts arise: one party (the principal) wants another party (the agent) to take an action, but the agent's action is imperfectly observable and the agent's interests may diverge from the principal's. This creates two distinct information problems — **moral hazard** (hidden action: the principal cannot observe the agent's effort) and **adverse selection** (hidden type: the principal cannot observe the agent's quality). The theory of optimal contracts, developed by Holmstrom (1979), Grossman and Hart (1983), and the mechanism-design tradition of Module 6, asks: what contract maximizes the principal's expected payoff subject to the agent's incentive-compatibility and participation constraints? 10.3 develops this theory formally, derives the key results (Holmstrom's informativeness principle, optimal linear contracts, the tradeoff between risk-sharing and incentive provision), and connects them to the layered-mechanism framework of 10.2 — a contract is the simplest layered mechanism, with the legal system as the outer layer and the bilateral agreement as the inner layer.

---

## Formal Definition

A **principal-agent problem** is defined by:

- **Principal** $P$ and **agent** $A$, with outside options $\bar{u}_P$ and $\bar{u}_A$ (reservation utilities).
- **Agent's action** $e \in E \subseteq \mathbb{R}_+$ (effort), chosen by the agent after the contract is signed.
- **Output** $x \in X$, a random variable with distribution $F(x|e)$ depending on effort. Output is observable by both parties; effort is **not** observable by the principal (moral hazard).
- **Contract** $w : X \to \mathbb{R}$, a compensation schedule specifying the agent's pay as a function of output.
- **Payoffs:** Principal receives $x - w(x)$; agent receives $w(x) - c(e)$, where $c(e)$ is the cost of effort with $c'(e) > 0, c''(e) > 0$.
- **Agent's preferences:** $U_A = \mathbb{E}[v(w(x))|e] - c(e)$ where $v$ is a von Neumann-Morgenstern utility (possibly concave: risk-averse agent).
- **Principal's preferences:** $U_P = \mathbb{E}[x - w(x)|e]$ (risk-neutral principal, WLOG by standard arguments).

The **moral hazard problem** is:

$$\max_{w(\cdot), e} \ \mathbb{E}[x - w(x)|e]$$

subject to: - **Incentive compatibility (IC):** $e \in \arg\max_{e'} \mathbb{E}[v(w(x))|e'] - c(e')$. - **Participation (IR):** $\mathbb{E}[v(w(x))|e] - c(e) \geq \bar{u}_A$.

The **adverse selection** variant replaces hidden action with hidden type: the agent has a privately known type $\theta \in \Theta$ that determines the cost function $c(e, \theta)$, and the principal designs a menu of contracts $\{(w_\theta(\cdot), e_\theta)\}_{\theta \in \Theta}$ such that each type self-selects the intended contract (IC across types) and participates (IR for each type). This is a screening problem in the sense of Module 6.

---

## Intuition

Every employment relationship, every staking deal, every consulting contract, every delegation of authority is a principal-agent problem. The principal wants something done. The agent does it. The problem is that the principal cannot perfectly monitor the agent's effort, and the agent — rationally — will shirk unless the contract provides sufficient incentive not to.

The central tension of contract design is **risk versus incentives**. A risk-neutral agent could simply be "sold the firm" (pay a fixed fee to the principal, keep all output): this provides perfect incentives (the agent bears all the consequences of their effort) but also transfers all risk to the agent. If the agent is risk-averse, this is costly — the agent demands a risk premium to accept the contract, and the principal ends up paying for both effort and risk-bearing. The optimal contract trades off these two forces: some risk is shifted to the agent to provide incentives, but not all risk, because the agent's risk premium would be too expensive.

**Holmstrom's informativeness principle** (1979) says: a performance measure should be included in the contract if and only if it is **informative** about the agent's effort — that is, if observing the measure changes the posterior distribution over effort levels. Including uninformative signals adds noise (increasing the risk premium) without providing incentive value (no information about effort). Including informative signals tightens the inference about effort, allowing the principal to provide incentives with less risk.

**Adverse selection** adds a second layer: even before the effort choice, the principal does not know the agent's type (quality, ability, cost structure). The principal must design a menu of contracts that induces truthful revelation of type, which typically requires leaving **information rents** to high-ability types (they could mimic low-ability types and do less work for the same pay). The tradeoff here is between extraction (getting value from each type) and incentives (getting each type to reveal themselves truthfully).

---

## Structural Role

10.3 sits at the intersection of mechanism design (Module 6) and institutional embedding (10.1). A contract is:

- A **mechanism** in the sense of 6.1: it specifies a message space (reports, output) and an outcome function (payment).
- A **private institution** in the sense of 10.1: it transforms the raw bilateral interaction into a constrained game with legally enforced payoff transfers.
- A **layered mechanism** in the sense of 10.2: the outer layer is the legal system that makes contracts enforceable, and the inner layer is the specific contract terms.

The principal-agent framework also underlies: - **10.4 Commitment and credibility:** Whether a contract is enforceable — whether the parties will actually follow through — is a commitment problem. - **10.5 Political economy:** Many political relationships (voters and politicians, citizens and bureaucrats) are principal-agent problems where the "contract" is the electoral mechanism. - **10.6 Strategic ecosystems:** Staking relationships in poker are principal-agent problems embedded in the broader poker ecosystem.

---

## Mathematical Development

### First-best (observable effort)

If the principal can observe and contract on effort directly, the problem reduces to choosing $(e^*, w^*)$ to maximize:

$$\mathbb{E}[x|e] - w^* \quad \text{subject to} \quad v(w^*) - c(e^*) \geq \bar{u}_A$$

The IR constraint binds: $w^* = v^{-1}(\bar{u}_A + c(e^*))$. The optimal effort satisfies $\mathbb{E}[x'|e^*] = c'(e^*)/v'(w^*)$ — marginal product of effort equals marginal cost, adjusted for risk. With a risk-neutral agent ($v(w) = w$): $\mathbb{E}[x'|e^*] = c'(e^*)$, and the agent is paid a fixed wage equal to their reservation utility plus effort cost. No performance pay is needed because effort is directly observable.

**Second-best (moral hazard)**

With hidden effort, the principal cannot contract on $e$ and must use the output $x$ as a noisy signal. The Lagrangian of the principal's problem yields the first-order condition for the optimal contract:

$$\frac{1}{v'(w(x))} = \lambda + \mu \frac{f_e(x|e)}{f(x|e)}$$

where $\lambda$ is the multiplier on IR, $\mu$ is the multiplier on IC, and $f_e/f$ is the **likelihood ratio** — the sensitivity of the output density to effort. This is the **Holmstrom condition**: optimal pay at output $x$ is increasing in the likelihood ratio $f_e(x|e)/f(x|e)$. Outputs that are more "indicative" of high effort receive higher pay.

**Theorem (Holmstrom Informativeness Principle, 1979).** Let $x$ and $y$ be two observable signals. If $y$ is informative about $e$ conditional on $x$ — that is, $f(y|x,e)$ depends on $e$ — then the optimal contract uses both $x$ and $y$. If $y$ is uninformative ($f(y|x,e) = f(y|x)$ for all $e$), then the optimal contract does not depend on $y$.

**Corollary.** Relative performance evaluation (comparing an agent's output to peers') is optimal when peer output is informative about common shocks but not about the agent's effort. This filters out common noise, tightening the signal-to-noise ratio on effort.

**Linear contracts**

In many applications (and for tractability), attention is restricted to **linear contracts**:

$$w(x) = \alpha + \beta x$$

where $\alpha$ is a fixed salary and $\beta \in [0,1]$ is the **incentive intensity** (piece rate, bonus coefficient, equity share).

Assume $x = e + \varepsilon$ with $\varepsilon \sim N(0, \sigma^2)$, $c(e) = e^2/2$, and agent utility $U_A = \mathbb{E}[w] - (r/2)\mathrm{Var}(w) - c(e)$ (mean-variance with risk aversion $r$). Then:

- Agent's optimal effort: $e^*(\beta) = \beta$ (from FOC of agent's problem).
- Agent's certainty equivalent: $\alpha + \beta^2 - \beta^2/2 - (r/2)\beta^2\sigma^2 = \alpha + \beta^2/2 - (r/2)\beta^2\sigma^2$.
- Principal's expected payoff: $e^*(\beta) - \alpha - \beta e^*(\beta) = \beta - \alpha - \beta^2$.

Substituting the binding IR constraint $\alpha = \bar{u}_A - \beta^2/2 + (r/2)\beta^2\sigma^2$:

$$U_P = \beta - \beta^2 - \bar{u}_A + \beta^2/2 - (r/2)\beta^2\sigma^2 = \beta - \beta^2/2 - (r/2)\beta^2\sigma^2 - \bar{u}_A$$

Maximizing over $\beta$:

$$\beta^* = \frac{1}{1 + r\sigma^2}$$

This is the **canonical incentive-intensity formula**. Incentive intensity is: - **Decreasing in risk aversion** $r$: more risk-averse agents get lower-powered incentives. - **Decreasing in noise** $\sigma^2$: noisier environments get lower-powered incentives. - Equal to 1 (sell the firm to the agent) only when $r = 0$ or $\sigma^2 = 0$.

**Adverse selection: the screening problem**

When the agent's type $\theta$ is private (e.g., $\theta$ determines the cost function $c(e,\theta)$ with $c_\theta < 0$ — higher types have lower cost), the principal designs a menu $\{(w(\theta), e(\theta))\}_\theta$ subject to:

$$\text{IC:} \quad v(w(\theta)) - c(e(\theta),\theta) \geq v(w(\theta')) - c(e(\theta'),\theta) \quad \forall \theta, \theta'$$
$$\text{IR:} \quad v(w(\theta)) - c(e(\theta),\theta) \geq \bar{u}_A \quad \forall \theta$$

Standard results (Mirrlees, Mussa-Rosen): IC implies monotonicity ($e(\theta)$ is nondecreasing in $\theta$), the IR constraint binds only for the lowest type, and all higher types earn **information rents** $\int_{\underline{\theta}}^{\theta}(-c_\theta(e(\tilde{\theta}),\tilde{\theta}))d\tilde{\theta}$. The principal distorts effort downward for all types below the top (to reduce information rents) and assigns efficient effort only to the highest type. This "no distortion at the top" result is the screening analogue of Myerson's optimal auction (6.3).

---

## Worked Example

**Poker staking as a principal-agent problem**

**Setup.** A backer (principal) stakes a player (agent) for a 100-session online poker campaign. The backer puts up the buy-ins; the player keeps a share of the profits.

**Moral hazard.** The backer cannot observe the player's effort directly — they cannot see whether the player is studying, table-selecting carefully, quitting when tilted, or playing their A-game. They observe only the output: session results $x_1, x_2, \ldots, x_{100}$.

**Contract structure.** A standard staking contract is linear: the player receives fraction $\beta$ of net profits above a "makeup" threshold (cumulative losses that must be repaid before profit-sharing begins). The backer receives $1-\beta$. Typical splits are $\beta \in [0.4, 0.6]$.

**Applying the linear-contract formula.** Let $r$ be the player's risk aversion, $\sigma^2$ be the variance of session results. The optimal split is:

$$\beta^* = \frac{1}{1 + r\sigma^2}$$

For a high-variance game (PLO, high stakes) with large $\sigma^2$, the optimal $\beta$ is lower — the player bears less of the variance and the backer bears more, but this weakens the player's incentive to exert effort. For a low-variance game (tight NLH, low stakes) with small $\sigma^2$, the optimal $\beta$ is higher — stronger incentives with less risk.

**Informativeness principle applied.** Should the backer include non-financial performance measures in the contract? For instance: hours played, hands per hour, number of tables, study-session logs. If these are informative about effort conditional on results (a player who plays 40 hours and runs bad looks different from a player who plays 5 hours and runs bad), then yes — the optimal contract should depend on both results and hours. If hours are uninformative (all players play the same hours regardless of effort, and effort manifests only in decision quality), then no — adding hours to the contract adds noise without information.

**Adverse selection.** Before signing the contract, the backer does not know the player's true skill level $\theta$. A high-skill player ($\theta$ high) produces higher expected output for the same effort. The backer must screen: offer a menu of contracts — e.g., "50/50 split with no makeup" for players who claim high skill (willing to accept because they expect high profits) and "70/30 split with makeup" for players who claim low skill (need protection). If the menu is designed correctly, each type self-selects the intended contract. The high-skill player earns an information rent (they could have taken the low-skill contract and gotten an easy deal). The backer accepts this rent as the cost of learning the player's type.

**Moral hazard in session selection.** The most insidious moral hazard in staking is not lazy play but **adverse session selection**: the player, staked for cash games, might choose softer games where they have a smaller edge but lower variance (comfortable but low-EV for the backer) rather than tougher games where their edge is larger but variance is higher. The backer's contract should, by the informativeness principle, include measures that are informative about game selection: average stakes played, player-pool quality metrics, table-selection data. Without these, the player's optimal response to a profit-sharing contract is to minimize their own risk (play soft games) rather than maximize expected profit (play the highest-EV games).

---

## Poker Anchor

### Concept mapping

| Formal object | Poker instantiation |
|---|---|
| Principal | Backer, staking stable owner, coaching investor |
| Agent | Staked player, horse |
| Action $e$ | Effort: study hours, A-game discipline, table selection, tilt control |
| Output $x$ | Session results (profit/loss in BB or dollars) |
| Contract $w(x)$ | Staking deal: split, makeup, bonuses, volume requirements |
| Moral hazard | Backer cannot observe effort quality — only results |
| Adverse selection | Backer cannot observe player's true skill level before signing |
| Incentive intensity $\beta$ | Player's profit share (typically 40–60%) |
| Informativeness principle | Include non-result metrics (hours, game selection) only if informative about effort |

### Concrete situation: the makeup trap

A player is down \$20,000 in makeup (cumulative losses the backer has absorbed). Under a standard makeup contract, the player must earn back \$20,000 before any profit-sharing kicks in. The player's marginal incentive for the next session is nearly zero: even if they win \$500, it goes entirely to reducing makeup, not to their pocket. The rational (if cynical) response: reduce effort, play carelessly, or quit the staking deal and find a new backer who will offer a clean slate.

This is a textbook moral hazard problem created by the contract structure. The makeup provision is designed to protect the backer (ensuring the player repays losses), but it destroys incentives when losses are large. The contract-design fix: either (a) cap makeup at a level where the player's incentive is still meaningful, (b) offer periodic "makeup resets" conditional on demonstrated effort (a stick-and-carrot structure from 5.3), or (c) switch to a "no-makeup" contract with a lower split, accepting more backer risk in exchange for preserving player incentives.

### What this understanding buys you

**Every staking deal is a contract-design problem.** Negotiating a staking deal without understanding principal-agent theory is like playing poker without understanding pot odds — you might get it right by intuition, but you cannot diagnose mistakes. The informativeness principle tells you which metrics to include. The risk-incentive tradeoff tells you how to set the split. The adverse-selection framework tells you how to screen players before staking them.

**Makeup destroys incentives at scale.** The deeper the makeup hole, the weaker the player's incentive. This is a direct consequence of the moral-hazard model: when the agent's marginal return to effort goes to zero (because all incremental gains go to repaying debt), effort goes to zero. Backers who do not understand this will attribute the player's deteriorating performance to "tilt" or "bad attitude" when it is actually a rational response to a broken contract.

**Relative performance evaluation applies to staking stables.** A backer with multiple horses can apply the informativeness principle by comparing results across players in similar games. If all players in the stable run bad simultaneously, it is a common shock (bad run), not individual shirking. If one player runs bad while others run well, it is more likely individual — the backer should investigate.

## Cross-Links

- **6.1–6.3 Mechanism design** — contracts as mechanisms; adverse selection as screening.
- **4.1–4.3 Bayesian games** — the agent's type is private information; BNE is the solution concept.
- **10.1 Games embedded in institutions** — the legal system as the outer institution that makes contracts enforceable.
- **10.2 Layered mechanisms** — contracts as the simplest layered mechanism (legal system + bilateral terms).
- **10.4 Commitment and credibility** — contract enforcement requires credible commitment by both parties.
- **5.3 Trigger strategies** — informal "contracts" sustained by repeated-game punishment rather than legal enforcement.
- **Economics / contract theory** — Holmstrom, Hart, Tirole: the full contract-theory canon.

---

## Structural Compression

**Invariant:** A principal-agent problem arises whenever one party delegates an action to another under asymmetric information; the optimal contract balances incentive provision (motivating effort) against risk-sharing (insuring the agent), with the informativeness principle determining which signals belong in the contract, and adverse selection requiring menus that trade information rents for truthful type revelation.

**Mechanism:** The principal designs a contract $w(x)$ mapping observable output to compensation; the agent chooses effort to maximize their expected utility net of effort cost; the optimal contract satisfies the Holmstrom condition (pay proportional to likelihood ratio) under moral hazard and the monotonicity/no-distortion-at-the-top conditions under adverse selection; the incentive intensity of linear contracts is $\beta^* = 1/(1 + r\sigma^2)$, trading off risk aversion and noise against incentive strength.

**Cross-System Echo:** The principal-agent problem is structurally identical to the **signal-extraction problem** in information theory (the "agent's effort" is the signal, the "output" is the noisy observation, and the "contract" is the decoder that maps observations to actions). Shannon's separation theorem (source coding and channel coding can be optimized independently) fails here: the "encoder" (agent) is strategic and chooses the signal (effort) to maximize their own utility, not the principal's. This strategic encoding is what makes contract theory harder than information theory and what connects it to the mechanism-design tradition.

---

## Exercises

1. Derive the first-best effort level when both principal and agent are risk-neutral and effort is observable. Show that the first-best contract is a "franchise fee" — the agent pays a fixed fee to the principal and keeps all output.

2. For the linear contract model with $x = e + \varepsilon$, $\varepsilon \sim N(0, \sigma^2)$, $c(e) = e^2/2$, and CARA utility with risk aversion $r$, derive the optimal incentive intensity $\beta^*$ and the principal's expected payoff. Show that the principal's payoff is decreasing in both $r$ and $\sigma^2$.

3. A backer stakes two players with identical observable characteristics. Player 1 generates results $x_1 = e_1 + \varepsilon_1 + \eta$ and Player 2 generates $x_2 = e_2 + \varepsilon_2 + \eta$, where $\eta$ is a common shock. Show that the optimal contract for Player 1 depends on $x_2$ (relative performance evaluation) and derive the optimal linear contract $w_1 = \alpha + \beta_1 x_1 - \gamma x_2$.

4. In the adverse-selection staking problem, the backer faces two types: high-skill ($\theta_H$, win rate 5 BB/100) and low-skill ($\theta_L$, win rate 1 BB/100). Design a separating menu of contracts $\{(\beta_H, \alpha_H), (\beta_L, \alpha_L)\}$ such that each type self-selects. Compute the information rent earned by the high-skill type.

5. Show that in the moral-hazard model, a contract that pays the agent a fixed wage regardless of output (pure insurance) leads to zero effort. Then show that a contract that gives the agent 100% of output (pure incentives) provides first-best effort but requires the agent to bear all risk. Conclude that the second-best contract is strictly interior: $0 < \beta^* < 1$ whenever $r > 0$ and $\sigma^2 > 0$.
6. A staking contract has a makeup provision: the player receives $\beta \cdot \max(X - M, 0)$ where $X$ is cumulative profit and $M$ is cumulative losses. Show that as $M$ grows large, the player's marginal incentive to exert effort in the next session approaches zero. Propose a contract modification that preserves incentives.
7. **Argue, structurally,** why the informativeness principle implies that a staking contract should include hours-played data if and only if hours are informative about effort conditional on results. Give a concrete scenario where hours are informative and one where they are not, and explain what the optimal contract looks like in each case.

# Commitment and Credibility

## Position in Knowledge Graph

**System:** Game Theory **Structural Domain:** Strategic Systems Integration **Node:** 10.4

A threat, a promise, a contract, or an institutional rule is only as powerful as its credibility. Throughout the course — from subgame perfection (3.3) to trigger strategies (5.3) to mechanism enforcement (6.1) — we have encountered the same constraint: a strategy that prescribes an action at some node is only part of an equilibrium if the player would actually carry out that action when the node is reached. This is the problem of **credibility**. The complementary problem is **commitment**: how to make a non-credible strategy credible by altering the structure of the game so that the player has no choice but to follow through. 10.4 unifies the credibility analysis scattered across earlier modules into a single formal treatment. It develops Schelling's theory of strategic commitment, Nash bargaining as a commitment baseline, hands-tying devices, strategic delegation, and the connection between commitment and institutional design (10.1). This node is the theoretical spine of Module 10: every institution (10.1), every mechanism layer (10.2), every contract (10.3), every collective-choice rule (10.5), and every ecosystem equilibrium (10.6) depends on whether the players' strategies are credible and whether they can commit.

---

## Formal Definition

Let $\Gamma$ be an extensive-form game (Module 3) with player set $N$ and strategy profiles $s \in S$.

A **commitment device** for player $i$ is a transformation of $\Gamma$ into a new game $\Gamma'$ in which player $i$'s action set at some information set is restricted to a singleton (the player's choice is removed; they are forced to take a particular action). Formally, if $h$ is an information set of player $i$ and $A_i(h)$ is the action set at $h$, the commitment device replaces $A_i(h)$ with $\{a_i^*\} \subseteq A_i(h)$.

A strategy $s_i$ is **credible** at information set $h$ if $s_i(h) \in B_i(h, s_{-i})$, the best-response set at $h$ given the opponents' strategies. A strategy profile $s$ is **subgame-perfect** if every $s_i(h)$ is credible at every $h$ — the definition of SPE from 3.3.

**Schelling's commitment principle** (1960): A player can improve their equilibrium payoff by *reducing* their own options — specifically, by committing to an action that would otherwise not be credible. The commitment is valuable precisely because it is irreversible: the opponent, observing the commitment, updates their best response and the resulting equilibrium shifts in the committer's favor.

**Nash bargaining solution.** Two players bargain over a surplus $S$ with disagreement payoffs $d = (d_1, d_2)$. The Nash bargaining solution is

$$(x_1^*, x_2^*) = \arg \max_{(x_1, x_2) \geq d,\ x_1 + x_2 = S} (x_1 - d_1)(x_2 - d_2)$$

which yields $x_i^* = d_i + (S - d_1 - d_2)/2$. Each player gets their disagreement payoff plus half the surplus. The disagreement point is the **commitment baseline**: a player who can credibly threaten a worse disagreement outcome for the opponent improves their own bargaining position.

**Strategic delegation.** Player $i$ appoints an agent (delegate) whose preferences differ from $i$'s own, and the delegate plays the game on $i$'s behalf. If the delegation is observable and irreversible (a commitment device), the opponent best-responds to the delegate's preferences, not to $i$'s. The optimal delegate is chosen to maximize $i$'s payoff given the opponent's best response to the delegate.

---

## Intuition

Commitment is the act of voluntarily destroying your own future options. This sounds irrational — and in a single-agent optimization problem it is. But in a game, destroying options can be rational because it changes the opponents' beliefs about your future play, and hence their best responses.

Schelling's paradigmatic example: two cars heading toward each other on a one-lane road. If one driver visibly removes their steering wheel and throws it out the window, the other driver must swerve — the first driver has no option to swerve, and both drivers know this. The commitment (removing the steering wheel) converts a non-credible threat ("I will not swerve") into a credible one ("I cannot swerve"). The driver who commits first wins the game of chicken.

Three structural observations:

1. **Commitment requires observability.** If the opponent does not know you have committed, the commitment has no strategic effect. This is why commitment devices must be public: burning bridges must be visible to the enemy; contracts must be verifiable by both parties; institutional rules must be common knowledge.
2. **Commitment requires irreversibility.** If the commitment can be cheaply reversed, it is not credible. A player who "commits" to a tough bargaining position but can back down at any time has not committed. True commitment requires either physical irreversibility (the steering wheel is gone), legal enforcement (the contract is binding), or reputational cost (backing down destroys reputation, which is a repeated-game mechanism from 5.3).
3. **Too much commitment is costly.** A player who commits to a completely rigid strategy loses the ability to adapt to new information. The optimal commitment is selective: commit on the dimensions where commitment shifts the equilibrium in your favor, but retain flexibility on all other dimensions. This is why real-world commitment devices — contracts, constitutions, reputations — are always partial.

**Hands-tying** is the general term for commitment by restricting your own future options. The canonical forms are:

- **Contractual commitment:** signing a legally binding agreement (10.3).
- **Reputational commitment:** building a public track record that would be costly to violate (5.5).
- **Institutional commitment:** creating or entering an institution that restricts your future actions (10.1).
- **Physical commitment:** burning bridges, sinking costs, making irreversible investments.

---

## Structural Role

10.4 is the theoretical core of Module 10. Every other node in the module presupposes a theory of commitment and credibility:

- **10.1 Institutions** are credible only if their enforcement mechanisms are self-enforcing (a credibility condition).
- **10.2 Layered mechanisms** require credible regulatory commitment — the regulator must credibly commit to enforcing the outer-layer rules.
- **10.3 Contracts** require credible enforcement by the legal system and credible effort by the agent.
- **10.5 Collective choice** requires credible commitment by voters and representatives to the outcomes of the voting process.
- **10.6 Strategic ecosystems** are stable only if the interaction norms within the ecosystem are credibly maintained.

Looking backward, 10.4 unifies: - **3.3 Subgame perfection:** the formal requirement of credibility in extensive-form games. - **5.3 Trigger strategies:** credibility of punishment as the binding constraint on cooperation. - **5.5 Reputation:** commitment through repeated-game reputation building. - **6.1 Mechanism design:** the designer's credible commitment to the mechanism's rules.

---

## Mathematical Development

### The value of commitment in sequential games

Consider a two-player sequential game where player 1 moves first and player 2 observes player 1's move. Let $s_1^{\mathrm{NE}}$ be player 1's Nash equilibrium action (without commitment) and $s_1^{\mathrm{SPE}}$ be their SPE action. Define the **commitment value** as:

$$\Delta_{\text{commit}} = u_1(s_1^C, B_2(s_1^C)) - u_1(s_1^{\mathrm{SPE}}, B_2(s_1^{\mathrm{SPE}}))$$

where $s_1^C$ is the action player 1 would choose if they could commit (i.e., move first with the opponent best-responding), and $B_2$ is player 2's best response. In a simultaneous game, there is no first-mover advantage and $\Delta_{\text{commit}} = 0$. In a sequential game where commitment shifts the equilibrium, $\Delta_{\text{commit}} > 0$.

**Example: Stackelberg competition.** Two firms choose quantities. Simultaneous (Cournot) NE: $q_1 = q_2 = (a-c)/3$ with profit $(a-c)^2/9$ each. Sequential (Stackelberg) with firm 1 as leader: $q_1^S = (a-c)/2, q_2^S = (a-c)/4$, profit $(a-c)^2/8$ for the leader (higher than Cournot). The commitment value is:

$$\Delta_{\text{commit}} = \frac{(a-c)^2}{8} - \frac{(a-c)^2}{9} = \frac{(a-c)^2}{72} > 0$$

The leader's commitment to a large quantity is credible because the quantity is sunk (produced before the follower moves). Without sunk production, the threat to produce a large quantity is not credible (the leader would prefer to reduce output after seeing the follower's quantity).

### Nash bargaining and outside options

The Nash bargaining solution $x_i^* = d_i + (S - d_1 - d_2)/2$ has two comparative statics relevant to commitment:

1. **Improving your disagreement payoff** $d_i \uparrow$ increases your bargaining share. Any action that credibly raises your outside option — threatening to walk away, developing a BATNA (best alternative to negotiated agreement) — improves your position.
2. **Worsening the opponent's disagreement payoff** $d_j \downarrow$ also increases your share. Any action that credibly worsens the opponent's outside option — making their alternatives worse, creating switching costs — improves your position.

The commitment dimension is that these threats must be credible: you must actually be willing to walk away, and the opponent must believe it.

### Hands-tying: formal model

Player $i$ can, before the game begins, choose a **commitment level** $k_i \in [0,1]$, where $k_i = 0$ means full flexibility and $k_i = 1$ means complete rigidity. The commitment restricts $i$'s action set to $A_i(k_i) \subseteq A_i(0)$ with $A_i(1) = \{a_i^*\}$.

**Theorem (Value of Hands-Tying).** In a two-player game where player 1 wants to credibly commit to an action $a_1^*$ that is not a best response to player 2's equilibrium action in the uncommitted game, the optimal commitment level $k_1^*$ is the minimum level that makes $a_1^*$ the only available action, provided:

$$u_1(a_1^*, B_2(a_1^*)) > u_1(s_1^{\mathrm{NE}}, s_2^{\mathrm{NE}})$$

That is, the committed outcome is better for player 1 than the uncommitted equilibrium. If this condition fails, commitment has no value (committing to a threat that does not improve your equilibrium payoff is pointless).

### Strategic delegation

Player $i$ delegates to an agent with utility function $\tilde{u}_i \neq u_i$. The opponent observes the delegation and best-responds to $\tilde{u}_i$. Player $i$ receives payoff $u_i(s_{\tilde{i}}, B_j(s_{\tilde{i}}))$ where $s_{\tilde{i}}$ is the delegate's equilibrium action under $\tilde{u}_i$.

**Theorem (Optimal Delegation in Cournot Duopoly, Fershtman-Judd 1987).** In a linear Cournot duopoly, the owner of firm $i$ delegates to a manager with utility $\tilde{u}_i = \alpha u_i + (1-\alpha) \cdot \text{revenue}$. The optimal delegation parameter is $\alpha^* < 1$: the manager is instructed to be "more aggressive" than profit-maximizing (partially maximizing revenue, which leads to higher output). This is credible because the manager's contract is observable. The result: the delegating firm produces more, the rival produces less, and the delegating firm earns higher profit than in the symmetric Cournot equilibrium — precisely the Stackelberg outcome achieved through delegation rather than sequential timing.

### Credibility and time consistency

A commitment is **time-consistent** if the player, given the opportunity to revise their commitment at a later date, would choose to maintain it. Most interesting commitment devices are **time-inconsistent**: the player would prefer to deviate from the commitment after the opponent has already responded to it. This is why commitment must be physically, legally, or reputationally irreversible — if it were time-consistent, no commitment device would be needed.

**Kydland-Prescott (1977):** A government that commits to low inflation (to reduce inflationary expectations) has an incentive to deviate from this commitment once expectations are set (the Phillips curve tempts them to exploit low expectations with a surprise inflation). The commitment to low inflation is time-inconsistent. Central bank independence is the institutional commitment device that makes the low-inflation policy credible.

---

## Worked Example

### Commitment in a poker context: the image of unpredictability

Consider a regular at a weekly home game who announces: "I never fold top pair. If you bet into my top pair, I call. Always." This is a commitment strategy — by publicly restricting their action set (removing "fold" from the available actions when holding top pair), they change opponents' optimal play.

**Without commitment.** In the uncommitted equilibrium, opponents bluff into this player at the game-theoretically optimal frequency. The player folds top pair some fraction of the time (against large bets on scary boards). The equilibrium is mixed.

**With commitment.** If the player credibly commits to never folding top pair, opponents' optimal response is to never bluff into this player when the player could have top pair — the expected value of bluffing is negative against a player who always calls. The player gains: they capture all the value from opponents' bluffs (which now never come) but also from opponents' thin value bets (which must now beat top pair to be profitable).

**Credibility analysis.** Is this commitment credible? Only if: 1. **Observability:** Opponents know about the commitment (the player has announced it publicly and has a track record of following through). 2. **Irreversibility:** The player actually follows through even when it is costly (e.g., top pair on a four-flush board against a large river bet). If the player folds top pair once, the commitment is broken and opponents return to bluffing. 3. **The repeated-game structure makes it self-enforcing:** the player maintains the commitment because the long-run value of the reputation (opponents never bluffing into them) exceeds the short-run cost of occasionally calling with a losing hand.

**Formal model.** Let $v^{\text{commit}}$ be the per-session value of the commitment (opponents stop bluffing, player captures extra value) and $c^{\text{commit}}$ be the per-session cost (occasionally calling with a losing hand). The commitment is self-enforcing if $v^{\text{commit}} > c^{\text{commit}}$. In a repeated game with discount factor $\delta$, the commitment is SPE if:

$$\frac{v^{\text{commit}} - c^{\text{commit}}}{1-\delta} > \frac{v^{\text{deviate}}}{1-\delta} \cdot \delta + v^{\text{no-commit}} \cdot \frac{1}{1-\delta}$$

where $v^{\text{deviate}}$ is the one-shot gain from deviating (folding top pair when you should call to maintain the commitment). This is exactly the trigger-strategy analysis of 5.3 applied to a commitment strategy.

---

## Poker Anchor

### Concept mapping

| Formal object | Poker instantiation |
|---|---|
| Commitment device | Public statement: "I never fold top pair," or a pattern of play that functions as a commitment |
| Credibility | Track record of actually following through; reputation at the table |
| Hands-tying | Announcing a rigid strategy that restricts your own future options |
| Strategic delegation | Telling a friend at the table "if I try to leave before midnight, stop me" (committing to session length) |
| Nash bargaining | Negotiating a table change, a save in a tournament, or a chop deal |
| Disagreement point | What each player gets if the chop negotiation fails (continuing to play) |
| Time inconsistency | The temptation to abandon the commitment when it becomes costly |

### Concrete situation: final-table chop negotiation

Five players remain in a tournament. The payouts are \$10K, \$7K, \$5K, \$3.5K, \$2.5K. Current chip stacks: A has 45%, B has 25%, C has 15%, D has 10%, E has 5%. A deal is proposed to redistribute the remaining prize money.

**Nash bargaining framework.** The disagreement point for each player is their ICM equity (the expected payout if they play it out). The surplus from a deal is the risk reduction — each player avoids the variance of playing it out. The Nash bargaining solution splits the surplus equally (adjusted for the number of players):

$$x_i^* = d_i^{\text{ICM}} + \frac{S - \sum_j d_j^{\text{ICM}}}{5}$$

where $S$ is the total remaining prize pool and $d_i^{\text{ICM}}$ is player $i$'s ICM equity.

**Commitment in chop negotiations.** Player A (chip leader) can improve their deal by committing to "I will not accept any deal that gives me less than \$9K." If this is credible (A has a reputation for walking away from bad deals, or A genuinely prefers to play it out), opponents must accommodate. If it is not credible (A obviously wants to lock up money and everyone knows it), the threat is ignored.

Player E (short stack) has the weakest negotiating position: their disagreement payoff is lowest (high probability of busting in 5th). But E can improve their position by committing to "I will shove every hand if we don't deal" — making the game riskier for everyone. This worsens the disagreement payoff for all other players, which improves E's bargaining share. The commitment is credible if E's stack is short enough that shoving is nearly optimal anyway.

#### What this understanding buys you

**Commitment is a strategic tool, not a character trait.** When a player says "I never bluff" or "I always defend my blinds" or "I never fold to a 3-bet," they are deploying a commitment strategy — whether they know it or not. The game-theoretic question is whether the commitment is credible and whether it improves their equilibrium payoff. Some commitments are valuable (never folding top pair against known bluffers); others are costly (never folding when the opponent has the nuts). The theory tells you which is which.

**Credibility is the binding constraint.** Most poker "commitment strategies" fail not because the strategy is wrong but because the commitment is not credible. A player who says "I never fold top pair" but folds top pair when the pot is large has no commitment — opponents learn this and adjust. The player who actually never folds top pair, even when it costs them, has a genuine commitment device.

**Chop negotiations are Nash bargaining.** Every final-table deal is a Nash bargaining problem. Understanding the disagreement point (ICM equity), the surplus (risk reduction), and the role of commitment (credible threats to walk away or play aggressively) gives you a structural framework for these negotiations rather than relying on intuition or social pressure.

---

## Cross-Links

- **3.3 Subgame perfection** — the formal credibility requirement for extensive-form strategies.
- **5.3 Trigger strategies** — commitment through repeated-game reputation.
- **5.5 Reputation and chain-store paradox** — reputation as a commitment device.
- **10.1 Games embedded in institutions** — institutions as credible commitment mechanisms.
- **10.3 Contracts and principal-agent** — contracts as legally enforced commitment.
- **10.5 Political economy** — constitutional commitment and credibility of collective-choice outcomes.
- **Economics / Schelling** — *The Strategy of Conflict* (1960): the original theory of strategic commitment.

---

## Structural Compression

**Invariant:** A strategy is credible if and only if the player would actually execute it when the relevant decision node is reached; commitment is the act of altering the game's structure to make a non-credible strategy credible, either by removing alternative actions (hands-tying), delegating to an agent with different preferences, or building a reputation whose destruction is more costly than following through; the value of commitment equals the difference between the committed and uncommitted equilibrium payoffs.

**Mechanism:** Commitment transforms the game tree by restricting the committer's action set at future nodes; the opponent, observing the commitment, updates their best response; the resulting SPE of the modified game gives the committer a higher payoff than the SPE of the original game; credibility requires observability, irreversibility, and either physical constraint, legal enforcement, or reputational cost sufficient to deter deviation.

**Cross-System Echo:** Commitment and credibility in game theory are structurally identical to **precommitment in dynamical systems** — a controller that "commits" to a feedback policy by fixing certain gains (making them physically unmodifiable) achieves robustness at the cost of adaptability, exactly as a game-theoretic committer achieves favorable equilibrium outcomes at the cost of flexibility. The same tradeoff appears in **constitutional design** (rigidity vs. adaptability), **software architecture** (immutable interfaces vs. flexible implementations), and **biological development** (cell differentiation as irreversible commitment to a developmental pathway).

## Exercises

1. In the game of Chicken with payoffs $(0, 0)$ for crash, $(2, 6)$ if player 1 swerves and player 2 goes straight, $(6, 2)$ if reversed, and $(4, 4)$ if both swerve: compute the mixed NE and show that player 1's expected payoff in the mixed NE is less than 6. Then show that if player 1 commits to going straight, the unique SPE gives player 1 payoff 6. Compute the commitment value.
2. In a Nash bargaining problem with $S = 100, d = (20, 10)$, compute the Nash bargaining solution. Then show that if player 1 can credibly raise their disagreement payoff to 40 (by developing an outside option), their bargaining share increases. Compute the increase.
3. In the Stackelberg duopoly with inverse demand $P = a - Q$ and constant marginal cost $c$, derive the leader's commitment value as a function of $a$ and $c$. Show that it is positive for all $a > c$.
4. A player in a repeated game wants to commit to grim trigger. Show that the commitment is credible (self-enforcing) if and only if the punishment phase is itself an SPE continuation, which holds when the punishment is a stage Nash equilibrium. Relate this to 5.3.
5. In the Fershtman-Judd delegation model, the owner of a Cournot duopoly firm delegates to a manager who maximizes $\alpha \cdot \text{profit} + (1 - \alpha) \cdot \text{revenue}$. Derive the equilibrium quantities as a function of $\alpha$, the optimal $\alpha^*$ for the delegating owner (when the rival owner does not delegate), and the equilibrium when both owners delegate.
6. A poker player commits to "always 3-betting with pocket aces preflop" (never slow-playing). Model the table's best response to this commitment. Under what conditions does this commitment increase the player's EV? Under what conditions does it decrease EV?
7. **Argue, structurally,** why the problem of credible commitment is fundamentally different from the problem of optimal strategy selection. Show that a player can have a strategy that is optimal-if-believed but not credible, and a strategy that is credible but not optimal-if-believed. Give one example of each from poker.

# Political Economy and Collective Choice

## Position in Knowledge Graph

**System:** Game Theory **Structural Domain:** Strategic Systems Integration **Node:** 10.5

10.1 showed that institutions are selected by a meta-game. 10.5 opens the black box of that meta-game: when the institution is chosen by a group of agents rather than a single designer, the selection mechanism is a **collective-choice problem** — voting, bargaining, or some hybrid. The theory of collective choice, developed by Condorcet (1785), Black (1948), Arrow (1951), and Downs (1957), asks: given a set of alternatives and a set of voters with preferences over those alternatives, what outcome does a given voting rule produce? The central results are the **Condorcet winner** (the alternative that beats every other in pairwise majority vote), the **median voter theorem** (under single-peaked preferences, the Condorcet winner is the median voter's ideal point), **Black's theorem** (single-peaked preferences guarantee a Condorcet winner exists), and **Arrow's impossibility theorem** (no aggregation rule simultaneously satisfies a small set of reasonable axioms). These results determine which institutions can be stably selected by a collective, which connects directly to the institutional stability question of 10.1 and the credibility question of 10.4.

---

## Formal Definition

A **collective-choice problem** is a tuple $(N, X, \{\succeq_i\}_{i \in N}, f)$ where:

- $N = \{1, \dots, n\}$ is a finite set of **voters** (agents, citizens, players).
- $X$ is a set of **alternatives** (policies, candidates, institutions).
- $\succeq_i$ is voter $i$'s **preference ordering** over $X$: a complete, transitive binary relation. Write $x \succ_i y$ for strict preference and $x \sim_i y$ for indifference.
- $f : \{\succeq_i\}_{i \in N} \to X$ is the **social choice function** (voting rule): it maps a profile of preferences to an outcome.

**Condorcet winner.** An alternative $x^* \in X$ is a **Condorcet winner** if it defeats every other alternative in pairwise majority voting:

$$\forall y \in X, y \neq x^* : \; |\{i \in N : x^* \succ_i y\}| > n/2$$

A Condorcet winner need not exist (the Condorcet paradox). When it does exist, it is unique.

**Single-peaked preferences.** Preferences are **single-peaked** over a linearly ordered set $X \subseteq \mathbb{R}$ if each voter $i$ has an ideal point $x_i^* \in X$ such that $u_i(x)$ is strictly increasing for $x < x_i^*$ and strictly decreasing for $x > x_i^*$. Equivalently: each voter has a unique most-preferred alternative, and moving away from it in either direction is always worse.

**Theorem (Black, 1948).** If preferences are single-peaked over a linearly ordered set, then a Condorcet winner exists and it is the ideal point of the **median voter**: the voter whose ideal point $x_m^*$ satisfies $|\{i : x_i^* \leq x_m^*\}| \geq n/2$ and $|\{i : x_i^* \geq x_m^*\}| \geq n/2$.

**Theorem (Median Voter Theorem, Downs 1957).** In a two-candidate election where candidates choose platforms to maximize votes and voters have single-peaked preferences, both candidates converge to the median voter's ideal point in equilibrium. This is a Nash equilibrium of the platform-selection game: deviating from the median loses the election.

**Spatial voting.** When $X \subseteq \mathbb{R}^d$ (multidimensional policy space), each voter has an ideal point $x_i^* \in \mathbb{R}^d$ and Euclidean preferences $u_i(x) = -\|x - x_i^*\|^2$. In one dimension, single-peakedness holds automatically. In $d \geq 2$ dimensions, generic preference profiles have no Condorcet winner — majority rule cycles through all alternatives (McKelvey's chaos theorem, 1976).

**Arrow's Impossibility Theorem (1951).** No social welfare function $f$ mapping preference profiles to a social ordering simultaneously satisfies: 1. **Universal domain:** every logically possible preference profile is admissible. 2. **Pareto:** if all voters prefer $x$ to $y$, so does society. 3. **Independence of irrelevant alternatives (IIA):** the social ranking of $x$ vs $y$ depends only on individual rankings of $x$ vs $y$. 4. **Non-dictatorship:** there is no single voter whose preferences always determine the social outcome.

---

## Intuition

Every group decision — which poker rules to adopt at a home game, which regulation to impose on online gambling, which candidate to elect, which payout structure to use in a tournament — is a collective-choice problem. The theory asks: does a "fair" outcome exist, and if so, can a voting rule find it?

The Condorcet winner is the most natural notion of a "fair" outcome: the alternative that a majority prefers to every other alternative. When it exists, it has strong normative appeal — no majority coalition can overturn it. But the Condorcet paradox shows it need not exist: with three voters and three alternatives, preferences can cycle ($A \succ B \succ C \succ A$ under pairwise majority), and no alternative is a Condorcet winner.

Black's theorem rescues the situation under single-peaked preferences: if alternatives are ordered on a line and each voter's utility has a single peak, Condorcet cycles cannot occur, and the median voter's ideal point wins. This is an enormously powerful result because many real policy dimensions are naturally single-peaked: how much to tax, how strict to regulate, how high to set the minimum bet.

The median voter theorem adds a strategic layer: even when the "voters" are candidates (or firms, or platform designers) choosing where to position, the equilibrium position is the median of the electorate. This explains centrist convergence in two-party elections and moderate institutional design in competitive markets.

Arrow's theorem is the deepest result: it shows that the problem of collective choice has no "ideal" solution. Every aggregation rule sacrifices something — either it is dictatorial, or it violates IIA, or it violates Pareto, or it restricts the domain. This is not a technical curiosity; it is a fundamental structural constraint on any collective decision-making process, including the meta-games that produce institutions (10.1).

---

## Structural Role

10.5 provides the strategic analysis of the institutional-selection meta-game introduced in 10.1. Where 10.1 said "institutions are selected by a higher-order game," 10.5 says "and here is what that higher-order game looks like when the selection is collective."

The connections are:

- **10.1 Institutional embedding:** The institution is the output of the collective-choice process analyzed here.
- **10.4 Commitment and credibility:** The collective-choice outcome is only stable if it is credible — if no majority coalition can profitably deviate. The Condorcet winner is credible in this sense; when no Condorcet winner exists, the outcome depends on the agenda (which introduces the commitment problem of who controls the agenda).
- **10.6 Strategic ecosystems:** The ecosystem's norms are a collective-choice equilibrium among its participants.
- **Module 6 Mechanism design:** Arrow's theorem shows the limits of mechanism design for collective choice; the Gibbard-Satterthwaite theorem (a cousin of Arrow's) shows that every non-dictatorial voting rule is susceptible to strategic manipulation.

Looking backward, 10.5 also connects to: - **2.1 Nash equilibrium:** The median voter theorem is a Nash equilibrium of the platform-selection game. - **7.1–7.4 Evolutionary dynamics:** Institutional change can be modeled as evolutionary dynamics over collective-choice rules.

---

## Mathematical Development

### Proof sketch of Black's theorem

Let $N = \{1, \ldots, n\}$ with ideal points $x_1^* \leq x_2^* \leq \cdots \leq x_n^*$. Suppose $n$ is odd for simplicity. The median voter is voter $m = (n+1)/2$ with ideal point $x_m^*$.

**Claim:** $x_m^*$ is a Condorcet winner. Take any $y \neq x_m^*$. WLOG $y > x_m^*$. By single-peakedness, every voter $i$ with $x_i^* \leq x_m^*$ prefers $x_m^*$ to $y$ (since $x_m^*$ is closer to their ideal point). There are at least $m = (n+1)/2$ such voters (voters $1, \ldots, m$). So a strict majority prefers $x_m^*$ to $y$. Since $y$ was arbitrary, $x_m^*$ is a Condorcet winner.

### The median voter theorem as Nash equilibrium

Two candidates $L, R$ choose platforms $p_L, p_R \in \mathbb{R}$. Voters have single-peaked preferences with ideal points distributed on $\mathbb{R}$ according to CDF $F$. Each voter votes for the closer candidate. Candidate $L$ wins if $F((p_L + p_R)/2) > 1/2$ (the midpoint between platforms determines the vote split). Each candidate maximizes vote share.

**Nash equilibrium.** If $p_L < p_R$, candidate $L$ captures all voters with ideal points below $(p_L + p_R)/2$. Candidate $L$ wants to move $p_L$ rightward (toward the median) to capture more votes; similarly $R$ wants to move leftward. The only stable pair is $p_L^* = p_R^* = x_m^*$, the median of $F$. At this point, neither candidate can gain by deviating: moving away from the median loses more votes on one side than it gains on the other.

### Condorcet cycles and McKelvey's chaos theorem

When $X \subseteq \mathbb{R}^d$ with $d \geq 2$, a Condorcet winner generically does not exist for $n \geq 3$ voters with Euclidean preferences. Worse, McKelvey (1976) showed:

**Theorem (McKelvey).** In $\mathbb{R}^d$ with $d \geq 2$, if there is no Condorcet winner, then the **top cycle** (the set of alternatives reachable from any starting point by a sequence of majority-preferred moves) is **all of** $X$. That is, from any starting policy, an agenda-setter can construct a sequence of pairwise votes leading to any other policy.

This means that in multidimensional policy spaces, majority rule is maximally unstable: the outcome depends entirely on the **agenda** — the order in which alternatives are voted on. Agenda control is a commitment device (10.4): the player who sets the agenda effectively commits the group to a particular sequence of comparisons, and the final outcome reflects the agenda-setter's preferences as much as (or more than) the voters' preferences.

### The Gibbard-Satterthwaite theorem

**Theorem (Gibbard 1973, Satterthwaite 1975).** If $|X| \geq 3$ and the social choice function $f$ is surjective (every alternative can be chosen) and strategy-proof (no voter can benefit by misrepresenting their preferences), then $f$ is dictatorial.

This is the mechanism-design analogue of Arrow's theorem: no non-dictatorial voting rule is immune to strategic manipulation. Every real voting system — plurality, runoff, Borda count, approval voting — is manipulable. The practical question is not "which rule is strategy-proof?" (none are) but "which rule is least manipulable in the strategic environment at hand?"

### Structure-induced equilibrium

When a Condorcet winner does not exist, institutions **induce** stability by restricting the set of alternatives or the agenda. A **structure-induced equilibrium (SIE)** is an outcome that is stable not because it is a Condorcet winner of the full alternative space but because the institutional rules (committee structure, amendment procedures, germaneness rules) restrict the feasible alternatives to a subset on which a Condorcet winner does exist.

Formally, let $X' \subseteq X$ be the set of alternatives that survive institutional filtering. If the voters' preferences restricted to $X'$ are single-peaked, Black's theorem applies and the median voter's ideal point (projected onto $X'$) is the SIE. This connects collective choice back to institutional design (10.1): the institution is the meta-game that selects $X'$, and the SIE is the outcome of the object-level game on the restricted set.

---

## Worked Example

### Home-game rule selection as collective choice

Five players in a weekly home game must collectively decide on the game format. The alternatives, ordered by "looseness" (action level):

$$X = \{\text{Limit Hold'em (LHE)}, \text{No-Limit Hold'em (NLH)}, \text{Pot-Limit Omaha (PLO)}, \text{PLO with a bomb pot every orbit}\}$$

Each player has single-peaked preferences over this ordering (preferring a certain action level and disliking both more and less action). Ideal points:

| Player | Ideal game |
|---|---|
| Alice | LHE (conservative, learning) |
| Bob | NLH (standard preference) |
| Carol | NLH (standard preference) |
| Dave | PLO (action player) |
| Eve | PLO-bomb (maximum action) |

The median voter is Carol (the 3rd of 5 players, ordered by ideal point). Carol's ideal point is NLH. By Black's theorem, NLH is the Condorcet winner: it beats LHE (Bob, Carol, Dave, Eve prefer more action than LHE), beats PLO (Alice, Bob, Carol prefer less action than PLO), and beats PLO-bomb (Alice, Bob, Carol, Dave prefer less action than PLO-bomb). The home game plays NLH.

**Strategic manipulation.** Eve knows the median voter theorem and tries to manipulate by proposing: "Let's vote between PLO and NLH only" (eliminating LHE and PLO-bomb from the agenda). Under this restricted vote, NLH still wins (Alice, Bob, Carol prefer it to PLO). Eve cannot manipulate the outcome under single-peaked preferences because the median voter's ideal point is robust to agenda restriction on the same ordered set.

**But if Eve can change the ordering.** Suppose Eve proposes a new dimension: "We should rotate games each week." Now the alternative space is two-dimensional (game type and rotation policy), preferences may no longer be single-peaked on a single dimension, and McKelvey's chaos theorem applies. Eve, as agenda-setter, can potentially steer the outcome toward PLO by controlling the sequence of votes. This is why home games often have a "house rules" document that fixes the game format — it is an institutional commitment (10.4) that prevents agenda manipulation.

---

## Poker Anchor

### Concept mapping

| Formal object | Poker instantiation |
|---|---|
| Voters $N$ | Players at the table, members of a poker community, regulators |
| Alternatives $X$ | Game formats, table rules, etiquette norms, payout structures |
| Preference ordering $\succeq_i$ | Each player's preferred rules (action level, stakes, game type) |
| Social choice function $f$ | The process by which the group decides: majority vote, house rules, floor decision |
| Condorcet winner | The rule set that a majority prefers to every alternative |
| Median voter | The player whose preferences are "in the middle" of the group |
| Agenda control | Who proposes the alternatives and in what order |

### Concrete situation: table etiquette as collective choice

At a live poker table, norms about behavior — slow-rolling, excessive celebration, phone use, English-only rules, shot clock — are collective-choice outcomes. They are not legislated from above but emerge from the preferences of the players and the enforcement mechanism (social pressure, floor calls).

Consider the norm "no slow-rolling." Most players dislike being slow-rolled (it is perceived as disrespectful). A few players enjoy the power move. If preferences are single-peaked over a "politeness" dimension, the median voter's preference determines the norm. At a table of regulars, the median voter's preference is "no slow-rolling," and this norm is enforced by social pressure (repeated-game punishment from 5.3). At a table of anonymous online players, the median voter may be indifferent (no social pressure, no repeated interaction), and the norm does not hold.

The structural point: table norms are not fixed features of "poker culture" — they are equilibria of a collective-choice game among the current players, sustained by the repeated-game structure and the enforcement mechanisms available. Change the players, change the enforcement mechanism, and the norms change.

### Concrete situation: community enforcement of norms

A poker community (forum, Discord, local card room regulars) collectively enforces norms against collusion, soft play, and angle shooting. The enforcement mechanism is community-level collective choice: if a majority of regulars believe a player has violated the norm, the player is ostracized (excluded from games, refused staking, publicly called out).

This is a Condorcet-winner outcome: the anti-collusion norm beats the alternative (no enforcement) in a pairwise majority vote among honest players. It is sustained by Black's theorem (preferences over "how strictly to enforce" are single-peaked, and the median voter wants moderate-to-strict enforcement). It fails when the community is too fragmented (no single ordering of alternatives), when the median voter shifts (too many colluders), or when enforcement is not credible (10.4 — the community cannot actually carry out the ostracism threat).

### What this understanding buys you

**Norms are not immutable.** They are collective-choice equilibria. When you move from one poker community to another, the norms change because the electorate (the set of regulars) changes. Understanding this lets you predict which norms will hold at which tables and in which communities, rather than treating norms as arbitrary cultural facts.

**The median player determines the culture.** At any table, the player in the middle of the preference distribution on any norm dimension determines what the group will tolerate. If you want to shift the table culture (toward more or less action, more or less etiquette), you need to shift the median — which means changing the player composition. This is why table selection is cultural as well as strategic.

**Agenda control matters.** In any group decision at the poker table — what game to play, whether to allow straddles, how to handle a disputed pot — the player who proposes the options and the order of discussion has disproportionate influence. This is McKelvey's chaos theorem in miniature: when multiple dimensions are in play, the agenda-setter can steer the outcome. Being the person who frames the discussion is a strategic advantage beyond the cards.

---

## Cross-Links

- **10.1 Games embedded in institutions** — the institutional meta-game whose structure is analyzed here.
- **10.4 Commitment and credibility** — credibility of collective-choice outcomes and agenda commitment.
- **6.1 Mechanism design** — voting rules as mechanisms; Gibbard-Satterthwaite as the impossibility of strategy-proof non-dictatorial voting.
- **5.3 Trigger strategies** — community enforcement as repeated-game trigger strategies sustaining collective-choice norms.
- **10.6 Strategic ecosystems** — ecosystem norms as collective-choice equilibria.
- **7.4 Evolutionary stability** — the evolutionary dynamics of norm change.
- **Political science / social choice theory** — Arrow, Black, Condorcet, Downs, McKelvey.

---

## Structural Compression

**Invariant:** Collective choice aggregates individual preferences into a group decision via a voting rule; under single-peaked preferences, the median voter's ideal point is the Condorcet winner (Black's theorem) and the equilibrium platform in electoral competition (median voter theorem); without single-peakedness, Condorcet cycles generically exist (McKelvey's chaos theorem) and outcomes depend on agenda control; no non-dictatorial voting rule with three or more alternatives is both strategy-proof and surjective (Gibbard-Satterthwaite).

**Mechanism:** Voters rank alternatives; the social choice function maps preference profiles to outcomes; single-peaked preferences guarantee existence of a Condorcet winner at the median; multidimensional policy spaces destroy single-peakedness and produce cycling, recoverable only by institutional restrictions on the alternative set (structure-induced equilibrium); the Gibbard-Satterthwaite theorem ensures that all real voting rules are strategically manipulable.

**Cross-System Echo:** Arrow's impossibility theorem is structurally parallel to the **no-free-lunch theorems** of optimization and machine learning: no single procedure dominates all others across all environments. Just as no optimization algorithm outperforms all others on all possible objective functions (Wolpert-Macready), no voting rule satisfies all desiderata across all preference profiles. The impossibility is not a flaw of specific voting rules but a structural constraint on aggregation itself — a fundamental limit analogous to Godel's incompleteness in logic and Heisenberg's uncertainty in physics: systems of sufficient richness cannot be fully resolved by any single formal device.

---

## Exercises

1. Three voters rank four candidates: Voter 1: $A \succ B \succ C \succ D$. Voter 2: $B \succ C \succ D \succ A$. Voter 3: $C \succ D \succ A \succ B$. Does a Condorcet winner exist? If so, identify it. If not, identify the Condorcet cycle.

2. Five voters have single-peaked preferences on $X = [0, 100]$ with ideal points $\{10, 30, 50, 70, 90\}$. Verify that 50 is the Condorcet winner. Then show that if voter 5 strategically misreports their ideal point as

50 (instead of 90), the Condorcet winner shifts to 50 (unchanged). Interpret: under what conditions can strategic misreporting change the outcome?

3. In a two-candidate election with voters' ideal points uniformly distributed on $[0, 1]$, derive the Nash equilibrium platforms. Show that both candidates converge to 1/2.
4. Prove that single-peaked preferences are sufficient but not necessary for the existence of a Condorcet winner. Give an example of a preference profile that is not single-peaked but has a Condorcet winner.
5. State and prove (or sketch) Arrow's impossibility theorem for three voters and three alternatives. Identify which axiom is violated by (a) majority rule, (b) Borda count, (c) dictatorial rule.
6. A poker community of 7 regulars must decide whether to (i) allow straddles, (ii) require English-only at the table, (iii) enforce a shot clock. Each issue is voted on separately by majority rule. Show that even if each individual issue has single-peaked preferences, the bundled outcome may not be a Condorcet winner of the product space, because the three dimensions interact.
7. **Argue, structurally,** why McKelvey's chaos theorem implies that agenda control is the most powerful strategic instrument in multidimensional collective choice, more powerful than having extreme preferences or a large coalition. Give an example from poker governance (tournament rule changes, community policy) where the agenda-setter determined the outcome.

# Strategic Ecosystems and Meta-Games

## Position in Knowledge Graph

**System:** Game Theory **Structural Domain:** Strategic Systems Integration **Node:** 10.6

This is the capstone node of the entire game theory course. Every preceding module has studied games in isolation: a fixed set of players, a fixed strategy space, a fixed payoff function, a fixed information structure. 10.1–10.5 began to relax these assumptions: institutions select which game is played (10.1), mechanisms layer on mechanisms (10.2), contracts structure bilateral relationships (10.3), commitment devices alter the game tree (10.4), collective choice determines institutional rules (10.5). 10.6 takes the final step: it considers the **entire ecosystem** of interacting games, players, strategies, and institutions as a single dynamical system. In a strategic ecosystem, players participate in multiple games simultaneously, strategies in one game affect outcomes in another, populations of strategies evolve over time across games, and the games themselves change in response to the population dynamics. The poker ecosystem — solvers, human players, training sites, coaches, staking operations, casinos, online platforms, tournament circuits, regulatory bodies — is the paradigmatic example. This node develops the formal framework for analyzing such systems, proves results about ecosystem equilibria and dynamics, and shows how everything in the course comes together.

---

## Formal Definition

A **strategic ecosystem** is a tuple

$$\mathcal{E} = (\mathcal{G}, N, \Theta, \Sigma, \Pi, \Phi, \delta)$$

where:

- $\mathcal{G} = \{G_1, G_2, \ldots, G_K\}$ is a finite set of **games** (the constituent strategic interactions of the ecosystem). Each $G_k = (N_k, S_k, u_k)$ is a game in the sense of 1.1, possibly extensive-form, repeated, or Bayesian.
- $N$ is a finite set of **ecosystem agents**. Each agent $i \in N$ participates in a subset $\mathcal{G}_i \subseteq \mathcal{G}$ of games.
- $\Theta = \{\theta_i\}_{i \in N}$ is the **type profile**: each agent has a type (skill, risk preference, bankroll, information set) that affects their play across all games.
- $\Sigma = \prod_i \prod_{k \in \mathcal{G}_i} S_k^i$ is the **ecosystem strategy space**: each agent's ecosystem strategy specifies their strategy in every game they participate in. An ecosystem strategy profile is $\sigma = (\sigma_i)_{i \in N} \in \Sigma$.
- $\Pi : \Sigma \to \mathbb{R}^N$ is the **ecosystem payoff function**: $\Pi_i(\sigma) = \sum_{k \in \mathcal{G}_i} w_{ik} \cdot u_k^i(\sigma_k)$, where $w_{ik}$ is the weight (frequency, stakes, importance) of game $k$ for agent $i$, and $\sigma_k$ is the restriction of $\sigma$ to game $k$.
- $\Phi : \Sigma \times \mathcal{G} \to \mathcal{G}'$ is the **ecosystem dynamics**: the strategy population at time $t$ determines the games that exist at time $t+1$. This captures institutional evolution (10.1), mechanism adaptation (10.2), and technological change.
- $\delta \in (0, 1)$ is the ecosystem discount factor, governing the relative importance of current vs. future payoffs across the ecosystem.

An **ecosystem equilibrium** is a pair $(\sigma^*, \mathcal{G}^*)$ such that:

1. **Within-game equilibrium.** For each game $G_k^* \in \mathcal{G}^*$, the strategy profile $\sigma_k^*$ restricted to $G_k^*$ is a Nash equilibrium (or BNE, or SPE, as appropriate).
2. **Cross-game consistency.** No agent can improve their ecosystem payoff $\Pi_i$ by reallocating effort, time, or capital across games: $\sigma_i^* \in \arg\max_{\sigma_i} \Pi_i(\sigma_i, \sigma_{-i}^*)$.
3. **Institutional stability.** The game set $\mathcal{G}^*$ is a fixed point of the dynamics $\Phi$: $\Phi(\sigma^*, \mathcal{G}^*) = \mathcal{G}^*$. No new games emerge and no existing games disappear at the equilibrium.

An ecosystem equilibrium is **evolutionarily stable** if, in addition, it is robust to small perturbations in the population strategy (per 7.4): any small invasion by a mutant strategy earns lower ecosystem payoff than the incumbent.

---

## Intuition

No player exists in a single game. A professional poker player simultaneously participates in: the hand game (the actual cards), the session game (selecting tables, managing tilt, quitting at the right time), the career game (studying, coaching, networking), the financial game (bankroll management, staking deals), the institutional game (choosing venues, responding to regulatory changes), and the social game (reputation, community standing, information sharing). Each of these "games" has its own strategy space, its own equilibrium, and its own payoff structure. But the games interact: your reputation in the social game affects your opponents' strategies in the hand game; your staking deal in the financial game constrains your risk tolerance in the session game; the regulatory environment in the institutional game determines which hand games are available.

A **strategic ecosystem** is the formalization of this interacting collection of games. The ecosystem is not just a big game (that would be a standard multiplayer game, analyzed with the tools of Modules 1–2). It is a system in which:

1. **Games are objects, not just settings.** Games are created, modified, and destroyed by the players' collective behavior. A new poker format catches on; an old format dies; a regulation changes the game; a solver changes the strategy population.
2. **Strategies span multiple games.** An agent's ecosystem strategy is not a strategy in a single game but a portfolio of strategies across multiple games. Optimizing in one game while ignoring others is suboptimal if the games interact (which they always do through shared resources — time, bankroll, attention, reputation).
3. **Populations evolve.** The distribution of strategies in the ecosystem changes over time. Weak strategies are exploited and exit; strong strategies attract imitators; the resulting population dynamics determine the long-run state of the ecosystem. This is evolutionary dynamics (Module 7) applied to a multi-game setting.
4. **The ecosystem itself evolves.** Not only do strategies change; the games themselves change. A casino introduces a new tournament format; a training site publishes a new solver output; a government regulates online poker. These changes alter the game set $\mathcal{G}$, the payoff functions, the strategy spaces. The ecosystem dynamics $\Phi$ captures this.

The key insight of 10.6 is that **no game in the ecosystem can be fully understood in isolation**. The equilibrium of the hand game depends on who is playing (determined by the session game and the institutional game); the equilibrium of the staking game depends on the hand game's variance and edge; the equilibrium of the training market depends on the gap between solver-optimal and human-optimal play. Everything is connected.

---

## Structural Role

10.6 is the terminal node of the course. Every module feeds into it:

- **Module 1 (Strategic form):** The atomic unit of the ecosystem is a game $G_k$.
- **Module 2 (Equilibrium):** Ecosystem equilibrium requires within-game NE at every constituent game.
- **Module 3 (Extensive form):** Many constituent games are sequential (the hand game, the session game).
- **Module 4 (Bayesian games):** Players have private types across the ecosystem.
- **Module 5 (Repeated games):** Many constituent games are repeated (the career game, the reputation game).
- **Module 6 (Mechanism design):** Institutions and platforms are mechanisms designed to govern constituent games.
- **Module 7 (Evolutionary games):** Population dynamics over strategies across the ecosystem.

- **Module 8 (Networks):** The ecosystem has a network structure — games connect to games, players connect to players.
- **Module 9 (Learning):** Players learn and adapt across the ecosystem.
- **Module 10.1–10.5:** Institutions, layered mechanisms, contracts, commitment, and collective choice are all features of the ecosystem.

The capstone status of 10.6 means it does not enable future nodes — it is the destination. But it also means it requires the conceptual machinery of the entire course. A student who understands 10.6 understands how all the pieces fit together.

---

## Mathematical Development

### Cross-game interaction effects

Let agent $i$ participate in games $G_1$ and $G_2$. The ecosystem payoff is

$$\Pi_i(\sigma_i^1, \sigma_i^2) = w_1 u_1^i(\sigma_i^1, \sigma_{-i}^1) + w_2 u_2^i(\sigma_i^2, \sigma_{-i}^2)$$

If the games are independent (separate opponents, no shared resources), the ecosystem equilibrium decomposes: $\sigma_i^{1*} \in \mathrm{NE}(G_1), \sigma_i^{2*} \in \mathrm{NE}(G_2)$ independently. But if the games interact — through shared resources (time, bankroll), shared opponents (reputation), or shared information — the decomposition fails.

**Shared bankroll.** If the agent has a fixed bankroll $B$ allocated between games, the payoff functions are coupled:

$$\Pi_i = w_1 u_1^i(\sigma_i^1, b_1) + w_2 u_2^i(\sigma_i^2, B - b_1)$$

where $b_1$ is the bankroll allocated to game 1. The optimal allocation satisfies the equimarginal principle:

$$w_1 \frac{\partial u_1^i}{\partial b_1} = w_2 \frac{\partial u_2^i}{\partial b_2}$$

Marginal returns per dollar of bankroll must be equal across games — otherwise the agent should reallocate.

**Shared reputation.** If the agent's reputation in game 1 (e.g., being known as a tough player) affects opponents' strategies in game 2 (e.g., opponents avoid the agent in game 2), then the games are informationally coupled. The ecosystem equilibrium must account for this cross-game information flow.

### Population dynamics in the ecosystem

Let $x_k(t) \in \Delta(S_k)$ be the population strategy distribution in game $G_k$ at time $t$. The **ecosystem replicator dynamics** extend the single-game replicator (7.2) to the multi-game setting:

$$\dot{x}_{k,s}(t) = x_{k,s}(t) \left[\Pi_s(\mathbf{x}(t)) - \bar{\Pi}_k(\mathbf{x}(t))\right]$$

where $\Pi_s(\mathbf{x})$ is the ecosystem fitness of strategy $s$ in game $k$ (accounting for cross-game interactions) and $\bar{\Pi}_k$ is the average ecosystem fitness in game $k$.

**Theorem (Ecosystem ESS).** A strategy profile $\sigma^*$ is an ecosystem ESS if it is an ESS in every constituent game and the cross-game interactions do not create profitable multi-game mutations. Formally: for all mutant ecosystem strategies $\sigma' \neq \sigma^*$, there exists $\bar{\varepsilon} > 0$ such that for all $\varepsilon < \bar{\varepsilon}$:

$$\Pi(\sigma^*, (1-\varepsilon)\sigma^* + \varepsilon\sigma') > \Pi(\sigma', (1-\varepsilon)\sigma^* + \varepsilon\sigma')$$

This is strictly stronger than game-by-game ESS because a mutant might be weak in each game individually but strong in the ecosystem due to cross-game synergies.

#### The meta-game over games

Define the **meta-game** $\Gamma^{\text{meta}}$ as the game played by ecosystem architects (casino operators, platform designers, regulators, training-site operators) who choose which games to offer:

$$\Gamma^{\text{meta}} = (N^{\text{meta}}, \mathcal{G}, U^{\text{meta}})$$

where $N^{\text{meta}}$ is the set of architects, each choosing a game $G_k \in \mathcal{G}$ to offer, and $U_j^{\text{meta}}$ is the architect's payoff (platform revenue, rake, subscription fees) given the equilibrium play in the offered games.

The **meta-game equilibrium** satisfies:

$$G_k^* \in \arg\max_{G \in \mathcal{G}} U_j^{\text{meta}}(G, \sigma^*(G), \mathbf{G}_{-j}^*)$$

where $\sigma^*(G)$ is the agent-level equilibrium in game $G$. This is a layered mechanism (10.2) at the ecosystem scale: the architects design the games, the agents play them, and the outcome is the fixed point of both levels.

#### Ecosystem stability and phase transitions

An ecosystem equilibrium can lose stability when parameters change. Define the **ecosystem stability index** $\lambda$ as the largest eigenvalue of the Jacobian of the ecosystem dynamics at equilibrium. The ecosystem is stable if $\lambda < 0$, unstable if $\lambda > 0$.

A **phase transition** occurs when a parameter change (e.g., the introduction of solvers, a regulatory change, a new platform) pushes $\lambda$ through zero, destabilizing the existing equilibrium and driving the ecosystem to a new one. The poker ecosystem has experienced several such phase transitions:

1. **Pre-online era (pre-2003):** Small, local ecosystems. Slow strategy evolution. Stable equilibrium among live regulars.
2. **Online boom (2003–2011):** Massive expansion of player pool, dramatic strategy evolution, new games (SNGs, MTTs, HU hypers), training sites as a new game type.
3. **Post-Black-Friday contraction (2011–2015):** Regulatory shock, player pool shrinks, recreational players exit, strategy arms race accelerates.
4. **Solver era (2015–present):** Solvers as a new "species" in the ecosystem, fundamentally changing the population dynamics. Human strategies converge toward solver-optimal play; the gap between human and solver narrows; training shifts from intuition to solver study.

Each transition is a change in $\lambda$ — the old equilibrium becomes unstable and the ecosystem transitions to a new one. Understanding these dynamics is the ecosystem-level analogue of understanding equilibrium selection in a single game.

---

## Worked Example

#### The poker ecosystem as a strategic ecosystem

#### Games in the ecosystem:

| Game $G_k$ | Players | Strategy space | Payoff |
|---|---|---|---|
| Hand game | Two or more players at a table | Betting strategies (Module 1–4) | Chip EV |

| Game $G_k$ | Players | Strategy space | Payoff |
|---|---|---|---|
| Session game | A player choosing which hands/tables to play | Table selection, session length, tilt management | Session profit |
| Training game | Player vs. study material (solvers, courses, coaches) | Study allocation, method selection | Skill improvement rate |
| Staking game | Backer and player | Contract terms (10.3) | Expected return on staking deal |
| Platform game | Casino/site vs. players | Rake, game spread, rewards (10.1, 10.2) | Platform revenue / player value |
| Regulatory game | Government vs. platforms | Legal framework, licensing, taxation (10.5) | Tax revenue / market structure |
| Social game | Players in a community | Reputation, information sharing, norm enforcement (5.3, 10.5) | Social capital, access to games |

**Cross-game interactions:**

- **Training → Hand:** Better training produces better hand-game strategies; the gap between trained and untrained players determines the hand-game's exploitability landscape.
- **Platform → Session:** Platform rules (rake, game spread) determine which session games are available and profitable.
- **Staking → Session:** Staking contracts affect risk tolerance and game selection (10.3: moral hazard in session selection).
- **Social → Hand:** Reputation at the table affects opponents' strategies in the hand game.
- **Regulatory → Platform:** Regulation constrains platform design (10.2: layered mechanisms).
- **Hand → Training:** Hand-game results generate data that feeds back into training.

**Ecosystem equilibrium analysis.** At the current ecosystem equilibrium:

1. **Solver-trained players dominate** the hand game at mid-to-high stakes. The population strategy in the hand game is converging toward the solver equilibrium (Nash equilibrium of the zero-sum hand game).
2. **Training sites and coaches** compete for students. The training game's equilibrium determines the rate at which the player population improves, which in turn determines the hand game's equilibrium (as more players approach solver play, exploitability decreases and the game becomes tighter).
3. **Staking operations** provide bankroll to players whose skill exceeds their capital. The staking game's equilibrium determines which players can access which stakes, which determines the player-pool composition at each stake level.
4. **Platforms compete** for players by adjusting rake, rewards, and game spread. The platform game's equilibrium determines which games exist and at what price (rake).
5. **Regulators** set the legal framework. The regulatory game's equilibrium determines which platforms operate, which affects everything downstream.

**Ecosystem dynamics.** The introduction of solvers (circa 2015) was a phase transition. Before solvers, the training game equilibrium was "study through hand review and intuition." After solvers, the training game equilibrium shifted to "study through solver analysis." This changed the population dynamics of the hand game: the strategy distribution shifted toward GTO play, which reduced the profitability of exploitative play, which reduced the edge of live-read specialists, which shifted the session game toward online play (where solver study translates more directly), which shifted the platform game toward sites that accommodate solver-trained players.

This cascade — a change in one game propagating through the ecosystem — is the defining feature of a strategic ecosystem. It cannot be understood by analyzing any single game in isolation.

## Poker Anchor

### Concept mapping

| Formal object | Poker instantiation |
|---|---|
| Ecosystem $\mathcal{E}$ | The poker world: all games, all players, all institutions, all technologies |
| Constituent game $G_k$ | Hand game, session game, training game, staking game, platform game, regulatory game |
| Ecosystem agent $i$ | A poker professional, a recreational player, a backer, a casino, a regulator, a coach |
| Ecosystem strategy $\sigma_i$ | How the agent plays across all games: hand strategy + table selection + study method + staking deals + venue choice |
| Ecosystem payoff $\Pi_i$ | Total career value: winnings + staking returns + coaching income + reputation capital |
| Ecosystem dynamics $\Phi$ | The co-evolution of strategies, games, institutions, and technologies over time |
| Phase transition | The online boom, Black Friday, the solver revolution |
| Ecosystem ESS | A population state that resists invasion by any multi-game mutant strategy |

#### Concrete situation: the solver revolution as ecosystem dynamics

Before solvers, the poker ecosystem's hand-game equilibrium was a mixed population: tight-aggressive regs, loose-aggressive maniacs, nitty grinders, recreational calling stations. Each "species" had a niche. The training game favored intuition and hand-reading. Coaching was a cottage industry.

Solvers changed the training game's technology parameter. Players who adopted solver study improved faster. The solver-trained population grew. As it grew, the hand-game equilibrium shifted: exploitative strategies that worked against unbalanced opponents became less effective against solver-trained opponents. Non-solver-trained players lost edge. Some exited; some adapted; some moved to live games where solver training translates less directly.

Simultaneously, the training game's equilibrium shifted: coaching operations that did not use solvers lost students; solver-based training sites gained market share; the "study meta" converged on solver analysis as the dominant approach. The staking game shifted: backers began requiring solver study as a condition of staking deals (10.3: informativeness principle — solver study hours are informative about future performance). The platform game shifted: sites that attracted solver-trained players saw tighter, lower-variance games and lower rake tolerance.

This is a complete ecosystem phase transition: one technological change (solvers) propagated through every game in the ecosystem, shifted every population distribution, and produced a new equilibrium qualitatively different from the old one. The old equilibrium was not "wrong" — it was stable under the old conditions. The new equilibrium is stable under the new conditions. Understanding this dynamics is the integration of everything in the course: Nash equilibrium (Module 2), evolutionary dynamics (Module 7), mechanism design (Module 6), institutional embedding (10.1), layered mechanisms (10.2), and collective choice (10.5).

#### What this understanding buys you

**You are not playing a game. You are living in an ecosystem.** Every decision you make — which hand to play, which table to sit at, which game to study, which backer to partner with, which platform to use — is a move in a multi-game ecosystem strategy. Optimizing any one of these in isolation is suboptimal if the games interact (and they always do). The player who understands the ecosystem optimizes across all games simultaneously: they study the format with the highest marginal return, play at the venue with the best institutional terms, structure their staking deal to align incentives, and build a reputation that pays dividends across multiple games.

**Ecosystem awareness is the ultimate edge.** At the hand-game level, edges between competent players are small and shrinking (the solver revolution is compressing the strategy distribution). At the ecosystem level, edges can be enormous: most players optimize the hand game and neglect the session game, the training game, the staking game, and the institutional game. A player who is 95th percentile at the hand game and 50th percentile at ecosystem management will underperform a player who is 85th percentile at the hand game and 95th percentile at ecosystem management. The ecosystem is where the real edge lives.

**Phase transitions are opportunities.** When the ecosystem transitions — a new regulation, a new technology, a new platform — the old equilibrium dissolves and a new one forms. During the transition, players who recognize the new equilibrium before others have an enormous edge: they adapt while others are still playing the old game. The online boom rewarded early adopters of online play. The solver revolution rewarded early adopters of solver study. The next phase transition will reward those who recognize it first. This course gives you the conceptual toolkit to recognize it.

---

## Cross-Links

- **1.1 Players, strategies, payoffs** — the atomic components that ecosystem games are built from.
- **2.1 Nash equilibrium** — the within-game equilibrium concept at each constituent game.
- **5.4 Folk theorem** — the multiplicity of equilibria in repeated games, resolved by ecosystem selection.
- **6.1 Mechanism design** — platforms and institutions as mechanisms within the ecosystem.
- **7.1–7.4 Evolutionary games** — population dynamics over strategies within the ecosystem.
- **8.1–8.3 Network games** — the network structure connecting games and agents in the ecosystem.
- **9.1–9.5 Learning** — how agents adapt and learn across the ecosystem.
- **10.1–10.5** — all preceding Module 10 nodes: institutions, layers, contracts, commitment, collective choice.

---

## Structural Compression

**Invariant:** A strategic ecosystem is a collection of interacting games played by overlapping populations of agents, where strategies in one game affect outcomes in others, the games themselves evolve in response to the population dynamics, and the ecosystem equilibrium is a simultaneous fixed point of within-game equilibrium, cross-game optimality, and institutional stability; no constituent game can be fully understood in isolation from the ecosystem.

**Mechanism:** Agents choose ecosystem strategies spanning multiple constituent games; cross-game interactions (shared resources, reputation, information) couple the equilibria of the constituent games; population dynamics (replicator dynamics generalized to the multi-game setting) govern the evolution of strategy distributions; ecosystem architects (platforms, regulators) design the games themselves, creating a meta-game whose equilibrium determines which games exist; phase transitions occur when parameter changes destabilize the current ecosystem equilibrium and drive the system to a qualitatively different one.

**Cross-System Echo:** Strategic ecosystems are the game-theoretic analogue of **ecological systems** in biology — multiple interacting species (strategies) competing across multiple niches (games), with the niches themselves shaped by the species' behavior (niche construction), and phase transitions corresponding to mass extinctions and adaptive radiations. They are also the analogue of **economic general equilibrium** — multiple interacting markets (games) with shared resources (factors of production), where the equilibrium is a simultaneous clearing of all markets. And they are the analogue of **complex adaptive systems** in the Santa Fe tradition — populations of heterogeneous agents interacting across multiple levels, with emergent macroscopic patterns arising from microscopic strategic interactions. The poker ecosystem is a laboratory-scale instance of all three.

## Exercises

1. Define a two-game ecosystem where agent $i$ plays both a Prisoner's Dilemma and a coordination game with the same opponent. Show that cooperation can be sustained in the Prisoner's Dilemma in the ecosystem equilibrium even though it cannot be sustained in isolation (without repetition), by conditioning behavior in the coordination game on the opponent's Prisoner's Dilemma play.

2. A poker player allocates 40 hours per week between playing (hand game) and studying (training game). Playing earns expected winrate $w(h_p, s)$ where $h_p$ is hours played and $s$ is skill level. Studying increases skill: $s(t+1) = s(t) + g(h_s)$ where $h_s = 40 - h_p$. Derive the optimal allocation $h_p^*$ as a function of current skill and the marginal returns to playing vs. studying.

3. Two platforms compete for players by setting rake $r_1, r_2$. Players choose which platform to play on based on rake and player-pool quality. Show that the platform game has a Nash equilibrium in rake and characterize it. Under what conditions does a new entrant platform destabilize the equilibrium?

4. Model the solver revolution as a population dynamics problem: let $x(t)$ be the fraction of solver-trained players. Solver-trained players earn $\pi_S(x) = a - bx$ (edge decreases as more players adopt solvers) and non-solver players earn $\pi_N(x) = c - dx$ with $a > c, b < d$ (solver players' edge shrinks slower). Find the evolutionary equilibrium $x^*$ and determine whether it involves full adoption, partial adoption, or extinction of non-solver players.

5. Show that an ecosystem ESS can fail to be a Nash equilibrium of the constituent games taken individually. Construct a two-game example where the ecosystem ESS prescribes a non-NE strategy in one of the constituent games because the cross-game interaction makes the non-NE strategy ecosystem-optimal.

6. A poker community has three types of agents: recreational players (provide the money), regulars (extract the money), and platforms (extract rake from both). Model this as a three-species ecosystem and derive the conditions under which each species survives. What happens to the ecosystem when recreational-player influx drops below a threshold?

7. **Argue, structurally,** why the claim "poker is getting tougher" is an ecosystem-level statement, not a game-level statement. Identify which constituent games in the poker ecosystem have gotten harder, which have gotten easier, and which have changed qualitatively rather than quantitatively. Explain why ecosystem-level reasoning is necessary to evaluate the claim, and why game-level reasoning alone is insufficient.