Chapter 10

Bayesian Games: Incomplete Information

Until now, everyone knew everyone's payoffs. In reality, a bidder doesn't know rivals' valuations; an insurer doesn't know the applicant's risk; a poker player doesn't see opponents' cards. Modeling such *incomplete information* coherently stumped the field until Harsanyi's breakthrough.

10.1 Harsanyi's framework: types and beliefs

Harsanyi's idea (1967–68): convert incomplete information into *imperfect* information by introducing a chance move. "Nature" first draws each player's private characteristics—her *type*—from a commonly known prior distribution; each player observes her own type only, then actions are chosen. Infinite hierarchies of beliefs (what I believe about your payoff, what I believe you believe about mine, ...) collapse into a tractable object.

Definition 10.1: Bayesian game

A *Bayesian game* is $\langle N, (A_i), (\Theta_i), p, (u_i) \rangle$ where Θ_i is player i's *type space*, $p \in \Delta(\Theta_1 \times \cdots \times \Theta_n)$ is the common prior, and payoffs $u_i(a; \theta)$ may depend on everyone's types. A *(pure) strategy* is a function $s_i : \Theta_i \to A_i$—a full plan of what to do for each type you might be.

Definition 10.2: Bayesian Nash equilibrium

A strategy profile $(s_1, \ldots, s_n)$ is a *Bayesian Nash equilibrium* (BNE) if for every player i and every type θ_i, the action $s_i(\theta_i)$ maximizes i's expected payoff given her belief $p(\theta_{-i} \mid \theta_i)$ about others' types and their strategies:

$$s_i(\theta_i) \in \arg\max_{a_i \in A_i} \; \mathbb{E}_{\theta_{-i}}\big[\, u_i(a_i, s_{-i}(\theta_{-i}); \, \theta_i, \theta_{-i}) \,\big|\, \theta_i \big].$$

Operationally: treat each *type* as a separate player and find an ordinary Nash equilibrium of the expanded game.

Example 10.1: Battle of the Sexes with uncertain interest

Row loves Column, but is unsure whether Column wants to meet (type θ^m, probability $\frac{1}{2}$) or avoid her (type θ^a): meeting types have BoS payoffs; avoiding types get the BoS payoffs with Column's coordination payoffs flipped to miscoordination. Solving type-by-type: the avoiding type has a dominant tendency to mismatch, and Row optimally goes to her favorite event regardless; the meeting type follows Row. Working through the best responses gives a strict BNE: Row plays B; Column plays B if θ^m, F if θ^a. The point: equilibrium play of *every* type must be specified, even types that do not occur in the realized play.

Example 10.2: Cournot with unknown costs

Firm 1's cost is c (common knowledge); firm 2's cost is c_L w.p. μ or c_H w.p. $1-\mu$, privately known. Demand $P = a - q_1 - q_2$. Each firm-2 type sets $q_2(c_2) = \frac{a-c_2-q_1}{2}$; firm 1 best-responds to the *expected* output $\bar{q}_2 = \mu q_2(c_L) + (1-\mu)q_2(c_H)$. Solving the three linear equations gives, with $\bar{c}_2 = \mu c_L + (1-\mu)c_H$:

$$q_1^* = \frac{a - 2c + \bar{c}_2}{3}, \qquad q_2^*(c_2) = \frac{a - 2c_2 + c}{3} + \frac{c_2 - \bar{c}_2}{6}.$$

The high-cost type produces more than it would under complete information (its rival, not knowing costs are high, produces less aggressively than it otherwise would). Information asymmetries change strategic behavior even for the *uninformed* party.

10.2 Auctions: private values

Auctions are the flagship application of Bayesian games—and of game theory in practice, running trillions of dollars of ad, spectrum, treasury, and electricity markets.

Setup (*independent private values*, IPV): one object; n bidders; valuations v_i drawn i.i.d. from a distribution F on $[0, \bar{v}]$, each privately known.

Example 10.3: Second-price (Vickrey) auction

Highest bidder wins, pays the second-highest bid. As shown in Example 3.5, bidding $b_i = v_i$ is weakly dominant. Truthful bidding requires no knowledge of F, of n, or of others' rationality—an ideal of *robust* mechanism design. Revenue: the expected second-highest valuation.

Example 10.4: First-price auction

Highest bidder wins, pays *her own* bid. Now bidding v_i guarantees zero profit; optimal bids *shade* below the valuation, trading off win probability against margin. Look for a symmetric increasing equilibrium $b(v)$. Bidder i with value v choosing bid $b(z)$ (imitating type z) wins iff all rivals' values are below z, probability $F(z)^{n-1}$; expected profit $\pi(z) = F(z)^{n-1}(v - b(z))$. The FOC at $z = v$ plus the boundary condition $b(0) = 0$ yields

$$b(v) \;=\; v - \frac{\int_0^v F(x)^{n-1}\,dx}{F(v)^{n-1}} \;=\; \mathbb{E}\big[\max_{j\neq i} v_j \,\big|\, \max_{j\neq i} v_j < v\big].$$

With uniform values on $[0,1]$: $b(v) = \frac{n-1}{n}v$. You bid the *expected highest rival value*

conditional on winning—shading exactly to the point where you'd tie with your strongest rival.

Theorem 10.1: Revenue Equivalence (Vickrey 1961; Myerson 1981)

In the IPV setting with risk-neutral bidders, *any* auction in which (i) the object always goes to the highest-value bidder in equilibrium and (ii) a bidder with the lowest possible value gets zero expected payoff, yields the *same expected revenue* to the seller. In particular, first-price, second-price, English (ascending), and Dutch (descending) auctions are all revenue-equivalent.

The proof is a corollary of the mechanism-design machinery in Chapter 14. The theorem's force is in its failure modes: risk aversion, correlated values, asymmetric bidders, and collusion break the equivalence in predictable directions, and auction *design* becomes consequential (Chapter 14).

10.3 Common values and the winner's curse

Example 10.5: The winner's curse

An oil tract has the same unknown value V to all bidders; each observes a noisy private signal $s_i = V + \varepsilon_i$. If everyone bids their unbiased estimate, the *winner* is the bidder with the largest positive error: winning is bad news. Rational bidders shade for this selection effect, bidding as if their signal were the highest—i.e., conditioning on the event of winning. Inexperienced bidders systematically overpay (documented in offshore oil leases, book publishing advances, free agency in sports, and corporate takeovers).

Key Idea: Condition on winning

In any common-value interaction—auctions, trading against possibly better-informed counterparties, accepting a deal others declined—evaluate the object *conditional on your action succeeding*. "If my bid wins / my offer is accepted, what does that tell me about what others know?" Ignoring this is among the most costly strategic errors in practice.

A related classic: *trade with adverse selection.*

Example 10.6: Akerlof's lemons market

A seller knows her car's quality $q \sim U[0,1]$; the buyer values quality q at $\frac{3}{2}q$, the seller at q. At price p, only sellers with $q \le p$ sell; average quality traded is $p/2$, worth $\frac{3}{4}p < p$ to the buyer. The market unravels: only $q = 0$ trades in equilibrium, though *every* car has positive gains from trade. Asymmetric information alone can destroy a market—the founding result of information economics (Akerlof 1970, Nobel 2001), and the reason for warranties, inspections, and certification (signaling: next chapter).

10.4 Global games: equilibrium selection through uncertainty

Incomplete information can *reduce* the multiplicity of equilibria.

Example 10.7: Currency attack (Morris–Shin 1998, sketch)

Speculators decide whether to attack a currency peg; attacking costs t and pays 1 if the peg breaks, which happens if enough speculators attack relative to fundamentals θ. With θ common knowledge, there is a range of θ with two equilibria (attack/no attack)—self-fulfilling crises. Now let each speculator observe θ with tiny idiosyncratic noise. Iterated dominance from the extreme regions (very bad θ: attacking dominant; very good: abstaining dominant) propagates inward and selects a *unique* equilibrium: attack iff the signal is below a cutoff θ^*. Vanishingly small uncertainty destroys the multiplicity, and the selected cutoff corresponds to the *risk-dominant* behavior. Global games are now a standard tool for bank runs, debt crises, and regime change.

Exercises

Exercise 10.1: First-price auction with two uniform bidders

Two bidders, $v_i \sim U[0,1]$ i.i.d. Verify directly that $b(v) = v/2$ is a symmetric BNE of the first-price auction. Compute the seller's expected revenue and check it equals the expected second-highest value $\mathbb{E}[\min(v_1, v_2)] = \frac{1}{3}$.

Exercise 10.2: All-pay auction

Two bidders with $v_i \sim U[0,1]$; both pay their bids, highest wins. Show $b(v) = v^2/2$ is a symmetric BNE, and verify revenue equivalence with the first- and second-price auctions. Name two real-world contests that are effectively all-pay.

Exercise 10.3: Public good with private values

Two neighbors decide simultaneously whether to contribute to a streetlight costing $c = 1$ to each contributor. The light is built if at least one contributes; enjoying it is worth $\theta_i \sim U[0,2]$, privately known. Find the symmetric cutoff BNE (contribute iff $\theta_i \geq \theta^*$) and compare total surplus with the first-best rule.

Exercise 10.4: Trading against the informed

You hold an asset worth $\theta \sim U[0,100]$ to you; a counterparty who *knows* θ offers to buy at price p of her choosing, which you accept or reject; the asset is worth 1.5θ to her. Characterize which offers get made and accepted in equilibrium, and compute your expected payoff from a policy of accepting any $p \geq 50$. Relate to the lemons example: why should an attractive-looking offer from an informed counterparty make you suspicious?