Appendix B

Solutions to Exercises

Solutions are deliberately compact: they give the key steps and answers, leaving routine algebra to the reader. Attempt each exercise seriously before reading.

Chapter 1

Exercise 1.1. (a) Two players, three actions each, *simultaneous*, complete information. (b) Seller and buyer; seller moves first (set price / describe car), buyer accepts or not; *sequential* with *incomplete* information (quality is the seller's private type)—a signaling/lemons situation (Chapters 10–11). (c) Two players, *sequential*, *perfect* information, finite (with repetition rules); by Zermelo's theorem it is determined (Chapter 6). (d) n bidders move *simultaneously* (sealed bids); *incomplete* information (private valuations): a Bayesian game (Chapter 10).

Exercise 1.2. Any answer with: two parties, mutual cooperation better than mutual defection, but defection individually profitable at each opponent action. Typical sustaining mechanisms: repetition and retaliation (Chapter 9), reputation, contracts, norms with sanctions.

Exercise 1.3. With money-only preferences, backward induction gives the younger sibling accepting any positive amount, so the older offers the smallest possible slice. With the anger term, accepting less than a third yields (slice -10) < 0 for small slices, so rejection is credible; the older sibling optimally offers exactly one third. Preferences, not the game tree, changed—and so did the prediction.

Chapter 2

Exercise 2.1. Start her with z. She pays ε to swap $z \to y$ (since $y \succ z$), ε to swap $y \to x$, then—by the cycle $z \succ x$—ε to swap $x \to z$. She is back at z having paid 3ε; repeat forever. Cyclic preferences are exploitable, hence "irrational."

Exercise 2.2. Expected utility of action a_i under belief μ is linear in payoffs, so under $v_1 = 3u_1 + 7$ every expected payoff is transformed by the same increasing affine map; rankings, hence arg max sets, are unchanged. For the nonlinear part: let action S (safe) pay 0 always and action R pay $+1$ with probability $2/3$ and -2 with probability $1/3$ (probabilities induced by the opponent's mixture). Under u: $\mathbb{E}[R] = \frac{2}{3} - \frac{2}{3} = 0 = \mathbb{E}[S]$ (indifferent). Under u^3:

$\mathbb{E}[R] = \frac{2}{3}(1) + \frac{1}{3}(-8) = -2 < 0 = \mathbb{E}[S]$. The cube preserves the ordinal ranking of sure outcomes but changes the best response.

Exercise 2.3. $\mathbb{E}[W] = 50$. $\mathbb{E}[u] = \frac{1}{2}\sqrt{100} = 5$. Certainty equivalent: $u(CE) = 5 \Rightarrow CE = 25$. Risk premium $= 50 - 25 = 25$.

Exercise 2.4. Write A and B as mixtures involving the common component: with $P =$ (1M for sure) and $Q =$ ($\frac{1}{11}$ chance 0, $\frac{10}{11}$ chance $5M$): $A = 0.11P + 0.89P$ and $B = 0.11Q + 0.89P$, while $C = 0.11P + 0.89\,\delta_0$ and $D = 0.11Q + 0.89\,\delta_0$ ($\delta_0 =$ zero for sure). Independence applied to the common 0.89-components forces $A \succsim B \iff P \succsim Q \iff C \succsim D$. The modal pattern $A \succ B$, $D \succ C$ contradicts this.

Chapter 3

Exercise 3.1. Step 1: for player 1, $T = (2, 1, 4)$ strictly dominates $B = (1, 0, 3)$; delete B. Step 2: in the reduced game, for player 2, R (payoffs 2, 3 vs T, M) strictly dominates C $(1, 2)$; delete C. No further strict domination: the survivors are $\{T, M\} \times \{L, R\}$. Within the reduced game, best-response marking gives two pure Nash equilibria, (M, L) and (T, R) (plus one mixed). Moral: IESDS need not reach a single profile; it brackets the predictions.

Exercise 3.2. Order 1: delete T (weakly dominated by M); then L (weakly dominated by R: $1 = 1$ vs M, $0 < 1$ vs B); outcome payoffs $(2, 1)$. Order 2: delete B (weakly dominated by M); then R (now weakly dominated by L: $1 = 1$ vs M, $0 < 1$ vs T); remaining outcomes pay $(1, 1)$. Different orders, different predictions—why iteration uses *strict* dominance.

Exercise 3.3. $u_i = q_i(12 - q_i - q_j)$. Best response $\mathrm{BR}(q_j) = (12 - q_j)/2 = 6 - q_j/2$. Any $q_i > 6$ is strictly dominated by $q_i = 6$ (marginal profit $12 - 2q_i - q_j < 0$ beyond 6 for all $q_j \geq 0$). Given $q_j \leq 6$, BR ≥ 3: outputs below 3 are dominated. Iterating: the surviving interval $[a_k, b_k]$ satisfies $a_{k+1} = 6 - b_k/2$, $b_{k+1} = 6 - a_k/2$, contracting to the fixed point $q = 6 - q/2 \Rightarrow q = 4$. The game is dominance solvable at the Cournot–Nash output $q_1 = q_2 = 4$.

Exercise 3.4. Claiming 100 is weakly dominated by 99 (if the rival claims 100 you get $99 + 2$ instead of 100; other cases tie or win). Once 100 is removed, 99 is dominated by 98, and so on down to the unique survivor $(2, 2)$. Most people (and experimental subjects) claim 95–100: with a small penalty/bonus parameter, the IESDS logic feels implausibly aggressive—a benchmark case for level-k thinking (Chapter 18).

Chapter 4

Exercise 4.1. Best responses — row: vs X: A; vs Y: B; vs Z: C. Column: vs A: X; vs B: Y; vs C: Z. Pure equilibria: (A, X) with $(3, 2)$, (B, Y) with $(2, 3)$, (C, Z) with $(2, 2)$.

Exercise 4.2. FOC for i: $\frac{1}{2}(1 + \frac{1}{2}e_j) - 2e_i = 0 \Rightarrow e_i = \frac{1}{4} + e_j/8$. Symmetric solution: $e^* = \frac{1}{4} + e^*/8 \Rightarrow e^* = \frac{2}{7}$. Joint surplus maximization: $1 + \frac{1}{2}e_j - 2e_i = 0$ symmetric gives $e^{\text{opt}} = \frac{2}{3} > \frac{2}{7}$. Each partner bears the full cost of effort but captures only half the extra output—classic team free-riding (Example 7.2 in miniature).

Exercise 4.3. At $(\frac{1}{2}, \frac{1}{2})$ each gets $\frac{1}{2}$; deviating to $x \neq \frac{1}{2}$ yields market share $\frac{x+1/2}{2} < \frac{1}{2}$ (for $x < \frac{1}{2}$)—no gain. If the profile has some $x_i \neq \frac{1}{2}$, the rival can move adjacent to x_i on the longer side and capture more than half, so no other profile is an equilibrium. Political reading: with single-peaked voters, two vote-maximizing candidates converge to the median voter's position.

Exercise 4.4. Impossible. If (a_1^*, a_2^*) is Nash, then $u_1(a^*) \geq u_1(a_1, a_2^*)$ for the other row(s): player 1's equilibrium payoff is (weakly) the best *in its column*, so it cannot be the strict minimum of all four cells. Strongest possible statement: the unique equilibrium can be strictly Pareto-dominated by another cell, and can give each player their worst payoff *outside their own deviation set*—the Prisoner's Dilemma achieves both players' second-worst payoffs, which is the worst possible.

Chapter 5

Exercise 5.1. Kicker mixes p on Left to equalize keeper's sides: $0.6p + 0.9(1-p) = 0.95p + 0.6(1-p) \Rightarrow p = \frac{6}{13}$. Keeper mixes q on Left to equalize kicker: $0.6q + 0.95(1-q) = 0.9q + 0.6(1-q) \Rightarrow q = \frac{7}{13}$. Equilibrium scoring probability $= 0.6 + 0.3 \cdot \frac{7}{13} = \frac{9.9}{13} \approx 0.762$.

Exercise 5.2. Standard RPS: by symmetry the value is 0; if any action had probability $> \frac{1}{3}$, its counter would earn strictly positive payoff, contradiction—so uniform is the unique equilibrium. Modified game (Rock beats Scissors for ± 2): indifference conditions $-p+2s = r-s = -2r+p = 0$ with $r+p+s = 1$ give $(r, p, s) = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$. *Rock's* probability falls (and Paper's rises): making Rock's win bigger makes opponents avoid Scissors and play Paper, which punishes Rock. You exploit an action's strength by playing its *counter's counter*, not the action itself.

Exercise 5.3. Symmetric equilibrium: calling gives $v - c$; not calling gives $v\,[1-(1-p)^{n-1}]$. Indifference: $(1-p)^{n-1} = c/v$, so $p^* = 1 - (c/v)^{1/(n-1)}$, decreasing in n. Probability nobody calls: $(1-p^*)^n = (c/v)^{n/(n-1)}$, *increasing* in n (exponent falls toward 1 and $c/v < 1$). More bystanders, more diffusion of responsibility, less help—the equilibrium logic of the bystander effect.

Exercise 5.4. Let q = probability of L. Row payoffs: $u(T) = 6q$, $u(M) = 4 - 2q$, $u(B) = 3$. There is no pure equilibrium (marking fails everywhere). All three row payoffs are equal only at $q = \frac{1}{2}$ (all give 3). Column is indifferent when $2t + 4m + 3b = 6t + 2m + 3b \Rightarrow m = 2t$. Hence a *component* of equilibria: $\sigma_2 = (\frac{1}{2}, \frac{1}{2})$ and $\sigma_1 = (t,\ 2t,\ 1-3t)$ for any $t \in [0, \frac{1}{3}]$ (all give payoffs $(3, 3)$). This degenerate game illustrates why "generic" games (Theorem 5.2) have finitely many equilibria but special payoff ties create continua.

Chapter 6

Exercise 6.1. Backward induction: after L, player 2 picks r $(2 > 1)$; after R, ℓ' $(1 > 0)$. Player 1 compares $L \to (1, 2)$ with $R \to (2, 1)$ and plays R: outcome $(2, 1)$. Normal form (player 2's strategies $\ell\ell', \ell r', r\ell', rr'$): pure Nash equilibria are $(R, r\ell')$—the SPE—and (L, rr'), which is sustained by the incredible plan to play r' after R (at that unreached node r' pays $0 < 1$).

Exercise 6.2. After burning, row's payoffs are shifted down by 1. Key steps of iterated weak dominance on the 4×4 normal form: (i) for row, "burn then play F" is dominated (best case 0); (ii) given (i), column infers from observed burning that row will play B, so column plays B after burning; (iii) then "burn, B" guarantees row 2, which dominates "don't burn, F" (at most 1); (iv) column, anticipating (iii), plays B even when row does *not* burn; (v) so row plays "don't burn, B" and gets 3. The unused option to burn selects row's favorite equilibrium—forward induction reasoning in action.

Exercise 6.3. Sequential BoS: row commits to B; column's best response is B; outcome $(3, 1)$ —a *first-mover advantage*, because the game rewards being predictable (coordination). In Matching Pennies the first mover loses for the same reason reversed: pure conflict rewards being *un*predictable, creating a second-mover advantage. The dividing line: does your opponent want to match you (move first) or mismatch you (move second)?

Exercise 6.4. Convention: strict majority; indifferent pirates vote against. With 2 pirates (4 proposes to 5): 5 rejects everything (he then gets all 100 *and* sees 4 overboard), so 4 dies: with 2 left, proposer 4 cannot win. Hence with 3 pirates, 4 votes for *anything* 3 proposes (voting no means the 2-pirate subgame where he dies): 3 proposes $(100, 0, 0)$ and wins with votes $\{3, 4\}$. With 4 pirates, proposer 2 needs 3 of 4 votes: buy the two pirates who get 0 in the 3-pirate continuation—4 and 5—with 1 coin each: $(98, 0, 1, 1)$. With 5 pirates, proposer 1 needs 3 of 5: continuation (from the 4-pirate case) gives $(-, 98, 0, 1, 1)$ to pirates $2, 3, 4, 5$; the two cheapest votes are pirate 3 (offer 1) and pirate 4 or 5 (offer 2): **proposal** $(97, 0, 1, 2, 0)$ (or $(97, 0, 1, 0, 2)$). Seniority plus foresight is worth 97 coins.

Chapter 7

Exercise 7.1. FOCs: $q_i = \frac{a - c_i - q_j}{2}$. Solving: $q_i^* = \frac{a - 2c_i + c_j}{3}$, $u_i^* = (q_i^*)^2$. Lower cost $\Rightarrow$ higher output and profit. Firm 2 is driven to zero when $a - 2c_2 + c_1 \leq 0$, i.e. $c_2 \geq \frac{a + c_1}{2}$.

Exercise 7.2. With the convention that the efficient (low-cost) firm wins ties: equilibrium $p_1 = p_2 = c_2$; firm 1 serves the market earning margin $c_2 - c_1$; firm 2 earns 0. (Without the tie-break convention, no equilibrium exists in continuous prices: firm 1 wants to price "just below" c_2; on a discrete grid it prices one tick below.) Prices are pinned by the *second-lowest* cost—the Bertrand analogue of the second-price auction logic.

Exercise 7.3. FOC for i: $\frac{x_j}{(x_1 + x_2)^2} V = 1$. Symmetry ($x_1 = x_2 = x$): $\frac{x}{4x^2} V = 1 \Rightarrow x^* = V/4$ each. Total rent dissipation $x_1^* + x_2^* = V/2$: half the prize is burned; each contestant nets $V/4$. (With n contestants, dissipation $\to V$: rent-seeking can consume the entire rent.)

Exercise 7.4. Consistent numbers: $D = 100$, both wait $\to$ 150 each; one withdraws $\to$ $(100, 60)$; both withdraw $\to$ 80 each.

	Wait	Withdraw
Wait	$150, 150$	$60, 100$
Withdraw	$100, 60$	$80, 80$

Both (Wait, Wait) and (Withdraw, Withdraw) are strict Nash equilibria: a run is a self-fulfilling coordination failure, not irrationality. Deposit insurance guarantees the waiter $100 \geq$ any

withdrawal payoff, making Wait a (weakly) dominant strategy and deleting the run equilibrium—which is why insurance works largely *without being used*.

Chapter 8

Exercise 8.1. Pareto frontier: $s_2 = 2 - s_1/2$. Maximize $(s_1 - 1)(\frac{3}{2} - \frac{s_1}{2})$: derivative $2 - s_1 = 0 \Rightarrow s^* = (2, 1)$. On $S' = S \cap \{s_1 \leq 3\}$ the point $(2, 1)$ is still feasible and the same maximization applies: unchanged, confirming IIA.

Exercise 8.2. $t = 2$: proposer 1 takes 1 (offers 0). $t = 1$: player 2 must offer player 1 at least $\delta \cdot 1$; keeps $1 - \delta$. $t = 0$: player 1 must offer player 2 $\delta(1 - \delta)$; keeps $x_3 = 1 - \delta + \delta^2$. Generally $x_T = 1 - \delta + \delta^2 - \cdots \to \frac{1}{1+\delta}$, the Rubinstein share.

Exercise 8.3. Stationary conditions: $1 - x = (1 - \pi)b_2 + \pi y$ and $1 - y = (1 - \pi)b_1 + \pi x$. Solving: $x = \frac{1-b_2+\pi b_1}{1+\pi}$ (and symmetrically for y). As $\pi \to 1$: $x \to b_1 + \frac{1-b_1-b_2}{2}$—the Nash bargaining solution with the *breakdown payoffs* as the disagreement point. What matters for threat-point purposes is what happens if bargaining *fails*, not impatience per se.

Exercise 8.4. With a hard deadline at T and no discounting, backward induction gives the period-T proposer the whole pie (the responder accepts anything over 0); at $T - 1$ the responder can secure that by waiting, so the $T - 1$ proposer concedes everything, and so on: the equilibrium division is determined entirely by who proposes *last*, and agreement (in the frictionless model) is immediate at that division. Adding any small per-period cost or risk makes early agreement strict—but the *terms* still reflect deadline power, which is why real negotiations cluster near deadlines: delay is costly, and bargaining power shifts discontinuously at T.

Chapter 9

Exercise 9.1. (a) Grim: $5(1 - \delta) + \delta \leq 2 \Rightarrow \delta \geq \frac{3}{4}$. (b) One-period punishment: the deviation gains 3 today and costs only $(2 - 1)\delta$ next period: $3(1 - \delta) \leq \delta(1 - \delta) \Rightarrow \delta \geq 3$—impossible. With these payoffs a single punishment period can never deter: the punishment must last several periods (roughly k with $3 \leq \sum_{j=1}^{k} \delta^j$). Shorter, milder punishments demand more patience; sometimes no feasible patience suffices.

Exercise 9.2. After 1's unilateral defection, mutual TFT prescribes the alternating path $(C, D), (D, C), (C, D), \ldots$ Check player 2's one-shot deviation at the start of the punishment phase: following TFT (play D) yields the alternation worth $\frac{4}{1+\delta}$; deviating to C restores (C, C) forever, worth 3. TFT requires $\frac{4}{1+\delta} \geq 3 \Rightarrow \delta \leq \frac{1}{3}$. But sustaining cooperation on-path requires $\delta \geq \frac{1}{3}$. Both hold only at the knife-edge $\delta = \frac{1}{3}$: mutual TFT is essentially never subgame perfect, despite being a Nash equilibrium for all $\delta \geq \frac{1}{3}$.

Exercise 9.3. Feasible set: convex hull of $(3, 1)$, $(1, 3)$, $(0, 0)$. Minmax for player 1: player 2 chooses q on B to minimize $\max(3q, 1 - q)$: at $3q = 1 - q$, $q = \frac{1}{4}$, value $\frac{3}{4}$; symmetrically $\underline{v}_2 = \frac{3}{4}$. Sustainable set: feasible payoffs dominating $(\frac{3}{4}, \frac{3}{4})$. The alternating scheme's average payoff $\to (2, 2)$ as $\delta \to 1$: feasible, individually rational, hence sustainable —and indeed the natural "fair" convention.

Exercise 9.4. Rule: "cooperate iff your partner is labeled good; defecting on a good partner makes you bad (forever, or for k periods)." A bad-labeled player is defected on by all future partners, so the one-shot gain g is weighed against a permanent per-period loss, exactly as in grim trigger: cooperation is an equilibrium for δ above a threshold. With noisy labels, permanent exclusion destroys too much surplus by accident; finite punishments, forgiveness, and label-repair (apologies, paying fines, rebuilding scores) become optimal—which is what credit systems and platform reputation mechanisms implement.

Chapter 10

Exercise 10.1. If the rival bids $v_j/2$, bidding b wins with probability $\mathbb{P}(v_j/2 < b) = 2b$ (for $b \le \frac{1}{2}$); expected profit $(v-b)\,2b$ is maximized at $b = v/2$. Revenue: $\mathbb{E}[\max(v_1, v_2)/2] = \frac{1}{2}\cdot\frac{2}{3} = \frac{1}{3} = \mathbb{E}[\min(v_1, v_2)]$, the second-price revenue—revenue equivalence verified.

Exercise 10.2. Imitating type z against a rival using $b(v) = v^2/2$: profit $= zv - z^2/2$ (win probability z, pay always). FOC: $v - z = 0$ at $z = v$: equilibrium confirmed. Revenue $= \mathbb{E}[b(v_1)] + \mathbb{E}[b(v_2)] = \mathbb{E}[v^2] = \frac{1}{3}$: equivalent again. Real-world all-pay contests: lobbying, patent races, sports training, wars of attrition, crowdfunded prize competitions.

Exercise 10.3. Cutoff θ^*: contributing yields $\theta - 1$; abstaining yields $\theta \cdot \mathbb{P}(\text{other contributes}) = \theta\frac{2-\theta^*}{2}$. Indifference at θ^*: $\theta^* - 1 = \theta^*\frac{2-\theta^*}{2} \Rightarrow (\theta^*)^2 = 2 \Rightarrow \theta^* = \sqrt{2} \approx 1.41$. First-best: build iff $\max_i \theta_i \ge 1$ with one contributor (or $\theta_1 + \theta_2 \ge 1$ with cost sharing)—the equilibrium underprovides: types in $(1, \sqrt{2})$ free-ride hoping the other contributes.

Exercise 10.4. The informed buyer offers p only when the asset is worth more than p to her *and* the offer is profitable: she buys exactly when $1.5\theta > p$, i.e. your acceptance selects against you. Accepting all $p \ge 50$: offers arrive (at $p = 50$) only when $\theta > 100/3$; you receive 50 for assets worth on average $\mathbb{E}[\theta \mid \theta > 33.3] = 66.7$. Expected payoff $\approx \frac{1}{3}(16.7) + \frac{2}{3}(50) = 38.9 < 50 =$ never trading. Like the lemons market: the willingness of an informed counterparty to trade at p is itself bad news about p.

Chapter 11

Exercise 11.1. Pooling on Beer: on-path belief = prior (0.9 strong) $\Rightarrow$ no duel; off-path (Quiche) belief "weak" $\Rightarrow$ duel. Strong: 3 vs 0; weak: 2 vs 1: both stay—PBE. Pooling on Quiche: symmetric construction works (2 vs 1 for strong, 3 vs 0 for weak)—also a PBE. Intuitive criterion on the Quiche equilibrium: deviating to Beer is *equilibrium-dominated* for the weak type (best conceivable payoff $0 + 2 = 2 < 3$, his equilibrium payoff) but not for the strong type $(1 + 2 = 3 > 2)$. Beliefs after Beer must then assign probability 1 to strong $\Rightarrow$ no duel $\Rightarrow$ the strong type profitably deviates: the all-Quiche equilibrium fails the criterion; all-Beer survives.

Exercise 11.2. Pure signaling, cost e^2/θ: low type: $e_L = 0$, $w_L = 1$. Separation requires the low type not to imitate: $3 - (e^s)^2 \le 1 \Rightarrow e^s \ge \sqrt{2}$; and the high type to comply: $3 - (e^s)^2/3 \ge 1 \Rightarrow e^s \le \sqrt{6}$. Riley outcome: $e^s = \sqrt{2}$. Productive education (wage $\theta + e/2$): full-information optima $e_\theta = \theta/4$ (from $\frac{1}{2} = 2e/\theta$). Separation now requires $3 + e/2 - e^2 \le$ (low type's honest payoff ≈ 1.06), giving $e^s \approx 1.66$, far above the high type's efficient 0.75: the

high type *over*-invests to separate. Education is productive but still partly wasteful at the margin—the signaling distortion survives.

Exercise 11.3. $b = \frac{1}{10}$: condition $2K(K-1)b < 1 \iff K(K-1) < 5 \Rightarrow K = 2$. Two intervals $[0, a], [a, 1]$ with actions $a/2$ and $(1+a)/2$; the boundary type a must be indifferent between the two actions, which gives $a/2 + b = (1-a)/2 - b$, hence $a = \frac{1}{2} - 2b = 0.3$. Receiver's expected loss: babbling $= \mathrm{Var}(U[0,1]) = \frac{1}{12} \approx 0.083$; two-interval $= \frac{0.3^3 + 0.7^3}{12} \approx 0.031$. Communication is coarse but valuable.

Exercise 11.4. Pooling fails: a contract slightly better for low-risks only (lower premium, lower coverage) attracts them and leaves the pool high-risk, making the pooling contract unprofitable—cream-skimming. Separating equilibrium: high-risks get full insurance at their fair premium (their IC binds, no distortion for them); low-risks get partial coverage at their own fair rate—coverage limited exactly enough that high-risks prefer their full-coverage contract. If high-risks are too rare, even this fails and no equilibrium exists (Rothschild–Stiglitz)—an early sign that competitive markets with adverse selection are fragile.

Chapter 12

Exercise 12.1. Row minima: $T : 2$, $C : 0$, $B : 1 \Rightarrow \underline{v} = 2$. Column maxima: $L : 4$, $M : 2$, $R : 5 \Rightarrow \overline{v} = 2$. Saddle point at (T, M), value 2 (unique: no other cell is simultaneously a row-max of its column and column-min of its row).

Exercise 12.2. Payoff matrix (row's payoff): $u(1,\cdot) = (0, 2, -1)$, $u(2,\cdot) = (-2, 0, 2)$, $u(3,\cdot) = (1, -2, 0)$. Antisymmetry ($u(i,j) = -u(j,i)$) confirms value 0. Equalizing: $2x_2 - x_3 = 0$, $-2x_1 + 2x_3 = 0$, $x_1 - 2x_2 = 0$ give $x^* = (\frac{2}{5}, \frac{1}{5}, \frac{2}{5})$. The "middle" action is played least—it never wins big.

Exercise 12.3. Primal: $\max v$ s.t. $x_H - x_T \geq v$ (vs H), $-x_H + x_T \geq v$ (vs T), $x_H + x_T = 1$, $x \geq 0$. Adding the constraints: $0 \geq 2v$, so $v \leq 0$; $x = (\frac{1}{2}, \frac{1}{2})$ achieves $v = 0$. The dual is the column player's identical program; both optima are 0 with uniform strategies.

Exercise 12.4. With the stated tie-breaking, one obtains (one legitimate run): $(H,H), (H,T), (H,T), (T,T), (T$ each player chases the other's empirical frequency, overshoots, and reverses. Frequencies after 6 rounds are already near $(\frac{1}{2}, \frac{1}{2})$; play cycles forever but the *averages* converge to optimal play, exactly as Robinson's theorem promises (and the value estimate converges to 0).

Chapter 13

Exercise 13.1. $\mathrm{BR}_i(x_j) = x_j/2$ (unique, continuous). $A_i = [0,1]$ compact convex; u_i continuous and strictly concave in own action: Theorem 13.4 applies. Fixed point: $x_1 = x_2/2$, $x_2 = x_1/2 \Rightarrow (0,0)$, the unique equilibrium.

Exercise 13.2. Pure equilibria: (T,L), (B,R), and (M,C). The strict equilibria (T,L) and (B,R) are trivially perfect. (M,C): against the *symmetric* totally mixed tremble $\sigma_2 = (\varepsilon, 1-2\varepsilon, \varepsilon)$, all three of row's actions pay exactly 1, so M is a best response, and likewise for C: (M,C) is also perfect (though not strict). Perfection kills weakly dominated behavior—but

M is not dominated, merely payoff-equivalent to $\frac{1}{2}T + \frac{1}{2}B$. The example shows perfection is weaker than one might hope in non-generic games.

Exercise 13.3. Standard example structure: players 1 and 2 move in sequence, player 3 moves at an information set reached only if *both* deviate. PBE permits player 3 to hold *any* belief there—including beliefs that treat 1's and 2's deviations as perfectly correlated in an arbitrary way. Consistency (limits of independent trembles) forces 3's belief to be the *product* of the two small deviation probabilities, pinning the conditional to specific ratios. Sequential rationality against those pinned-down beliefs can then eliminate equilibria PBE allows. Moral: consistency disciplines off-path beliefs by requiring a single coherent "theory of mistakes."

Exercise 13.4. Suppose a bet occurs with positive probability at some pair of signals. Willingness to bet is common knowledge at that point (betting is voluntary and observed). Player 1 bets only if his conditional expected gain is positive; player 2 likewise—but the bet is zero-sum, so the two conditional expectations, taken with respect to the *common prior* conditional on the same public event, cannot both be positive (iterated expectations). Aumann's agreement logic turns "I'm happy to bet" into information that destroys the other's willingness. Without the common prior, players can knowingly disagree, and betting on differences of *opinion* (not information) becomes possible.

Chapter 14

Exercise 14.1. Efficient allocation: $A \to 2$, $B \to 3$ (total $12 > 10$). Payments: without bidder 2, the best the others can do is 10 (bidder 1 takes the pair); others' value at the actual allocation is 6, so bidder 2 pays 4; symmetrically bidder 3 pays 4. Revenue $= 8$. VCG revenue can be fragile: raise $v_2(A) = v_3(B) = 10$ and each pays $10 - 10 = 0$—efficient allocation, *zero* revenue. Adding bidders can *reduce* VCG revenue; this non-monotonicity is why high-stakes combinatorial auctions modify VCG (core-selecting designs).

Exercise 14.2. Exponential: $\frac{1-F}{f} = 1$, so $\psi(v) = v - 1$ (increasing: regular); $r^* = 1$. Pareto $F = 1 - v^{-2}$ on $[1, \infty)$: $\frac{1-F}{f} = \frac{v^{-2}}{2v^{-3}} = \frac{v}{2}$, so $\psi(v) = v/2 > 0$ everywhere: the "reserve" is the support minimum (always sell). Indeed, for one bidder a posted price p earns $p\,(1 - F(p)) = p \cdot p^{-2} = p^{-1}$, maximized at the corner $p = 1$. Heavy tails make revenue maximization degenerate—and for tail index 1, expected values are infinite and every price earns the same revenue. Optimal auction theory quietly assumes tails are thin enough.

Exercise 14.3. With $c, v \sim U[0, 1]$, efficient trade (iff $v > c$) requires expected transfers that violate budget balance under IC and IR; running the Myerson–Satterthwaite calculation, the *second-best* mechanism trades iff $v \geq c + \frac{1}{4}$: one maximizes expected gains from trade subject to IC, IR, and budget balance; the binding budget constraint enters as a Lagrange multiplier on the virtual surplus, and solving for the multiplier that exactly balances the budget yields the $\frac{1}{4}$ wedge (this is also the outcome of the linear equilibrium of the Chatterjee–Samuelson double auction). First-best trade region: the triangle above the diagonal; second-best: the smaller triangle shifted up by $\frac{1}{4}$—trades with gains below $\frac{1}{4}$ are sacrificed to control information rents.

Exercise 14.4. Median voting: a voter left of the median can change the outcome only by reporting on the other side of the median, moving the chosen point *further* from her peak;

misreporting on her own side changes nothing. Single-peakedness makes "further in the wrong direction" unambiguously worse: strategy-proof. Phantoms preserve the argument verbatim (the median of true peaks plus constants). Gibbard–Satterthwaite does not apply because single-peaked preferences are a *restricted domain*—the theorem needs all preference profiles to be possible.

Chapter 15

Exercise 15.1. Core: $\sum x_i = 135$, $x_1 + x_2 \geq 90$, $x_1 + x_3 \geq 80$, $x_2 + x_3 \geq 70$ (equivalently $x_3 \leq 45$, $x_2 \leq 55$, $x_1 \leq 65$), $x_i \geq 0$: a nonempty polytope; e.g. $(65, 45, 25)$. Shapley value: averaging marginal contributions over the 6 orders gives $\phi = (50, 45, 40)$. Check core membership: pairwise sums $95, 90, 85$ all clear their thresholds: yes, $\phi \in C(v)$.

Exercise 15.2. Core: if a worker received more than $g(w) - g(w-1)$, the coalition of L with the other $w-1$ workers could object; so worker pay is capped by the marginal product. Shapley with $g = \sqrt{k}$, $w = 3$: in a uniformly random order, the number of workers preceding L is uniform on $\{0, 1, 2, 3\}$, so $\phi_L = \frac{1}{4}(0+1+\sqrt{2}+\sqrt{3}) \approx 1.04$; each worker gets $(\sqrt{3} - \phi_L)/3 \approx 0.23$ (below the marginal product $\sqrt{3} - \sqrt{2} \approx 0.32$, consistent with the core cap).

Exercise 15.3. Enumerating pivots over the 24 orders of $(A{:}4, B{:}3, C{:}2, D{:}1)$ with quota 6: A pivots 10 times, B 6, C 6, D 2. Index: $(\frac{5}{12}, \frac{1}{4}, \frac{1}{4}, \frac{1}{12})$. Note B and C have *equal power* despite different weights—power is about pivotal position in coalitions, not vote count (the deep lesson of power indices).

Exercise 15.4. Players $\{1, 2, 3, 4\}$; preferences: $1 : 2 \succ 3 \succ 4$; $2 : 3 \succ 1 \succ 4$; $3 : 1 \succ 2 \succ 4$; 4: anyone. In any matching, someone among $\{1, 2, 3\}$ is paired with 4; that person is the *favorite* of one of the other two, who is currently matched to their least-preferred partner among $\{1, 2, 3\}$—the two form a blocking pair. All three matchings are blocked: no stable matching exists. Two-sidedness matters because it prevents exactly these odd "cyclic envy" configurations within one side.

Chapter 16

Exercise 16.1. $u(A, B) = 5 > 3 = u(B, B)$ and $u(A, A) = 2 > 1 = u(B, A)$: A strictly dominates B. The only symmetric equilibrium is (A, A), which is strict, hence the unique ESS. (The indifference equation $5 - 3p = 3 - 2p$ has no solution in $[0, 1]$, confirming no mixed equilibrium.)

Exercise 16.2. Payoffs: $u_H = -x + 2(1 - x) = 2 - 3x$, $u_D = 1 - x$. Replicator: $\dot{x} = x(1 - x)\,(u_H - u_D) = x(1 - x)(1 - 2x)$. Rest points $0, \frac{1}{2}, 1$. Sign: positive on $(0, \frac{1}{2})$, negative on $(\frac{1}{2}, 1)$: interior point $x^* = \frac{1}{2} = V/C$ is asymptotically stable; both monomorphic states are unstable. The population converges to the ESS mix from any interior start.

Exercise 16.3. $\frac{d}{dt} \log \frac{x_a}{x_b} = \frac{\dot{x}_a}{x_a} - \frac{\dot{x}_b}{x_b} = u(a, x) - u(b, x) \leq -\varepsilon$, where $\varepsilon = \min_x [u(b, x) - u(a, x)] > 0$ by strict domination and compactness. Hence $x_a / x_b \leq e^{-\varepsilon t} \cdot x_a(0)/x_b(0) \to 0$, and since $x_b \leq 1$, $x_a \to 0$.

Exercise 16.4. With mutant share ε: $u_{\text{Stag}} = 4\varepsilon + 0(1-\varepsilon) = 4\varepsilon$; $u_{\text{Hare}} = 2$. Stag grows iff $4\varepsilon > 2$, i.e. $\varepsilon > \frac{1}{2}$: the good convention needs the *majority* to switch at once—an enormous invasion barrier (the mixed equilibrium share). This is why inferior conventions persist and why coordinated "big push" interventions (assurance contracts, standards bodies announcing a switch date, critical-mass subsidies) are the practical remedy.

Chapter 17

Exercise 17.1. Take $\Phi(q_1, q_2) = (a-c)(q_1+q_2) - q_1^2 - q_2^2 - q_1 q_2$. Then $\partial\Phi/\partial q_i = a - c - 2q_i - q_j = \partial u_i/\partial q_i$: an exact (continuous) potential. Maximizing Φ: FOCs give $q_1 = q_2 = \frac{a-c}{3}$, the Cournot equilibrium—and concavity of Φ proves uniqueness and gives convergence of continuous best-response dynamics.

Exercise 17.2. Equilibrium: all traffic on road 2 (cost $x^d \le 1$ with equality at $x = 1$): social cost 1. Optimum: minimize $1 - x + x^{d+1}$: $x^{\text{opt}} = (d+1)^{-1/d}$, cost $1 - \frac{d}{d+1}(d+1)^{-1/d} \to 0$ like $\frac{\ln d}{d}$. Hence PoA $\sim \frac{d}{\ln d} \to \infty$: with very convex congestion, selfish routing can be arbitrarily bad—and the affine bound $\frac{4}{3}$ of Theorem 17.3 is seen to be a genuinely structural fact about *linear* congestion.

Exercise 17.3. Rock–Paper–Scissors; let μ be uniform on the three *diagonal* profiles $(R,R), (P,P), (S,S)$, giving each player payoff 0. CCE: any constant deviation faces a uniform opponent marginal, earning 0—no gain. Not CE: conditional on the recommendation "Rock," the opponent plays Rock for sure, so swapping to Paper earns $+1 > 0$. The extra power of swap deviations (using the recommendation as information) is exactly what separates the two concepts.

Exercise 17.4. Congestion game with $c(k) = k$; potential $\Phi = \sum_r \frac{\ell_r(\ell_r+1)}{2}$. Pure equilibria: profiles where no job can move to a strictly less loaded machine, i.e. $|\ell_1 - \ell_2| \le 1$: loads are balanced. Every equilibrium makespan is $\lceil n/2 \rceil$, which is also optimal: PoA $= 1$ here. (With m *unrelated* machines or weighted jobs, equilibria can be unbalanced and the PoA grows—the classic Koutsoupias–Papadimitriou setting.)