EconS 503 – Advanced Microeconomics II
Recitation #1 – Game Theory (Basics)[1]

1. **MWG 7.E.1 – Behavior vs. mixed strategies and perfect recall.** Consider the two-player extensive-form game in Figure 1. Player 1 moves at three information sets (the root, plus two later sets), Player 2 moves at a single information set that contains two nodes (so Player 2 does not know whether Player 1 played M or R).

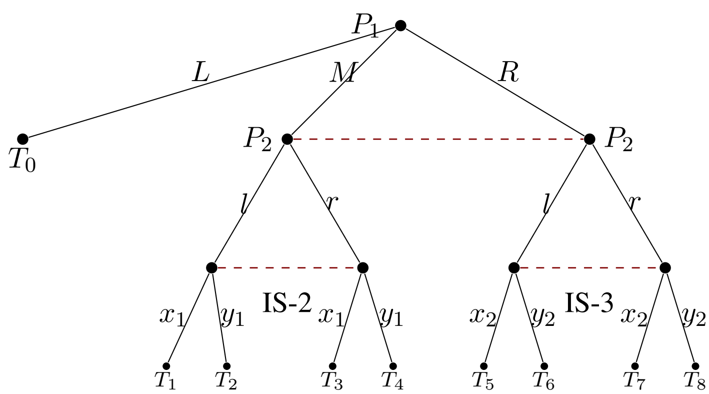

Figure 1: Extensive-form representation of the game in MWG 7.E.1. Player 1 moves at three information sets: the root, IS-2 (both nodes reached after Player 1 chose M), and IS-3 (both nodes reached after Player 1 chose R). Player 2 moves at a single information set that contains both nodes reached after M and R (dashed red connectors indicate nodes in the same information set). Terminal nodes are $T_0, \ldots, T_8$.

   (a) What are the possible strategies of Player 1 and Player 2?

   (b) Show that for *any* behavior strategy Player 1 might play, there is a realization-equivalent mixed strategy — one that generates the same probability distribution over terminal nodes for every mixed strategy of Player 2.

   (c) Show the converse: for any mixed strategy of Player 1, there is a realization-equivalent behavior strategy.

   (d) Suppose we merge IS-2 and IS-3 into a single information set. Argue the game no longer has perfect recall. Which of (b) and (c) still holds?

   Recall two definitions we will use throughout.

   - A *mixed strategy* for a player is a probability distribution over her *pure* strategies (each pure strategy specifies one action at every information set the player owns). Randomization is done once, at the beginning of the game, over the entire sequence of actions.

[1] Felix Munoz-Garcia, School of Economic Sciences, Washington State University, 103H Hulbert Hall, Pullman, WA. Tel. 509-335-8402. Email: fmunoz@wsu.edu.

- A *behavior strategy* specifies, independently at each information set, a probability distribution over the actions available there. Randomization is local to each information set.
- Two strategies are *realization equivalent* if, against every strategy of the opponent, they induce the same probability distribution over terminal nodes. This is the operational test of "do these strategies do the same thing" from an outsider's viewpoint.

Kuhn's theorem (1953) says: in a game of *perfect recall* (each player remembers her own past actions), behavior and mixed strategies are realization-equivalent classes of each other. Parts (b) and (c) walk through both directions of this equivalence in the specific tree of Figure 1; part (d) shows what breaks when perfect recall fails.

Answer. *Part (a): the strategy sets.*

Step 1. Count Player 1's information sets. Player 1 owns three information sets:

- The root: three actions available, $\{L, M, R\}$.
- IS-2 (both nodes reached only after M): two actions available, $\{x_1, y_1\}$.
- IS-3 (both nodes reached only after R): two actions available, $\{x_2, y_2\}$.

A pure strategy for Player 1 must prescribe an action at *every* information set she owns, even those she may never actually reach. That is a triple $(a_{\text{root}}, a_{\text{IS-2}}, a_{\text{IS-3}})$, giving

$$|S_1| = 3 \times 2 \times 2 = 12.$$

Step 2. Enumerate.

$$\begin{aligned} S_1 = \{&(L, x_1, x_2),\ (L, x_1, y_2),\ (L, y_1, x_2),\ (L, y_1, y_2),\\ &(M, x_1, x_2),\ (M, x_1, y_2),\ (M, y_1, x_2),\ (M, y_1, y_2),\\ &(R, x_1, x_2),\ (R, x_1, y_2),\ (R, y_1, x_2),\ (R, y_1, y_2)\}. \end{aligned}$$

For example, (L, x_1, y_2) says: "play L at the root; *if we ever reached IS-2*, play x_1; *if we ever reached IS-3*, play y_2." The last two components are *contingency plans* for information sets that may never be reached (they aren't reached under L).

Step 3. Player 2. Player 2 has only one information set with two actions:

$$S_2 = \{l, r\}.$$

Answer. *Part (b): every behavior strategy has a realization-equivalent mixed strategy.*

Step 1. Set up the behavior strategy. A behavior strategy for Player 1 randomizes *independently* at each of her three information sets. Denote:

- At the root, L, M, R with probabilities p_1, p_2, p_3 (with $p_1 + p_2 + p_3 = 1$).
- At IS-2, x_1, y_1 with probabilities q_1, q_2 (with $q_1 + q_2 = 1$).
- At IS-3, x_2, y_2 with probabilities s_1, s_2 (with $s_1 + s_2 = 1$).

Let Player 2 play l, r with probabilities $\sigma(l), \sigma(r)$.

Step 2. Compute the terminal-node distribution under the behavior strategy. Reading off Figure 1 branch by branch,

$$\begin{aligned}
&\Pr(T_0) = p_1, && \\
&\Pr(T_1) = p_2\,\sigma(l)\,q_1, && \Pr(T_2) = p_2\,\sigma(l)\,q_2, \\
&\Pr(T_3) = p_2\,\sigma(r)\,q_1, && \Pr(T_4) = p_2\,\sigma(r)\,q_2, \\
&\Pr(T_5) = p_3\,\sigma(l)\,s_1, && \Pr(T_6) = p_3\,\sigma(l)\,s_2, \\
&\Pr(T_7) = p_3\,\sigma(r)\,s_1, && \Pr(T_8) = p_3\,\sigma(r)\,s_2.
\end{aligned}$$

Every terminal probability factors into (root-choice) $\times$ (Player-2 choice) $\times$ (IS-2 or IS-3 choice), because the three information sets randomize independently.

Step 3. Construct a candidate mixed strategy. Consider the following distribution over Player 1's pure strategies (assign zero probability to any pure strategy not listed):

$$\begin{array}{ll}
(L, x_1, x_2) & \text{with probability } p_1, \\
(M, x_1, x_2) & \text{with probability } p_2 q_1, \\
(M, y_1, x_2) & \text{with probability } p_2 q_2, \\
(R, x_1, x_2) & \text{with probability } p_3 s_1, \\
(R, x_1, y_2) & \text{with probability } p_3 s_2.
\end{array}$$

Probabilities sum to $p_1 + p_2(q_1 + q_2) + p_3(s_1 + s_2) = p_1 + p_2 + p_3 = 1$, so this is a well-defined mixed strategy.

Step 4. Why these specific pure strategies? Consider any pure strategy in the support, e.g., (M, x_1, x_2): it prescribes M at the root and x_1 at IS-2. Under this strategy the only reachable terminal nodes are T_1 (if Player 2 plays l) and T_3 (if Player 2 plays r); nothing in IS-3 is reached, so the third component (x_2) is irrelevant. That's why we chose (M, x_1, x_2) but could equally have chosen (M, x_1, y_2) — both deliver the same terminal-node probabilities for this cell.

Step 5. Compute the terminal distribution under the mixed strategy. Under our mixed strategy and Player 2's σ,

$$\begin{aligned}
\Pr(T_0) &= \Pr(L,\cdot,\cdot) = p_1, \\
\Pr(T_1) &= \Pr(M,x_1,\cdot)\cdot\sigma(l) = p_2q_1\,\sigma(l), && \text{compare to Step 2: match.} \\
\Pr(T_2) &= \Pr(M,y_1,\cdot)\cdot\sigma(l) = p_2q_2\,\sigma(l), && \text{match.} \\
\Pr(T_3) &= \Pr(M,x_1,\cdot)\cdot\sigma(r) = p_2q_1\,\sigma(r), && \text{match.} \\
\Pr(T_4) &= \Pr(M,y_1,\cdot)\cdot\sigma(r) = p_2q_2\,\sigma(r), && \text{match.} \\
\Pr(T_5) &= \Pr(R,\cdot,x_2)\cdot\sigma(l) = p_3s_1\,\sigma(l), && \text{match.} \\
\Pr(T_6) &= \Pr(R,\cdot,y_2)\cdot\sigma(l) = p_3s_2\,\sigma(l), && \text{match.} \\
\Pr(T_7) &= \Pr(R,\cdot,x_2)\cdot\sigma(r) = p_3s_1\,\sigma(r), && \text{match.} \\
\Pr(T_8) &= \Pr(R,\cdot,y_2)\cdot\sigma(r) = p_3s_2\,\sigma(r), && \text{match.}
\end{aligned}$$

The distributions agree at every terminal node, for every strategy of Player 2. So the constructed mixed strategy is realization-equivalent to the given behavior strategy.

Intuitively, the behavior strategy tosses independent coins at each information set. A realization-equivalent mixed strategy tosses one big compound coin at the start of the game, over pure strategies whose contingent choices *happen to line up* with what those independent tosses would have produced. The specific probabilities p_2q_1, p_2q_2, etc. are simply the joint probabilities of the sequential draws.

Answer. *Part (c): every mixed strategy has a realization-equivalent behavior strategy.*

Step 1. Set up the mixed strategy. A mixed strategy for Player 1 is a distribution over her 12 pure strategies; write the probabilities as $p_1,\dots,p_{12}$ (with $p_i \geq 0$ and $\sum p_i = 1$), one for each of the 12 pure strategies listed in Part (a).

Step 2. Compute the terminal-node probabilities under the mixed strategy. Group pure strategies by their root action:

- L-pure strategies (labels 1–4): all lead to T_0 regardless of anything else, so $\Pr(T_0) = p_1 + p_2 + p_3 + p_4$.
- M-pure strategies (labels 5–8): terminal depends on (x_1 vs y_1) and Player 2's action. The x_1-at-IS-2 group is $\{(M,x_1,x_2),(M,x_1,y_2)\}$ with total probability $p_5 + p_6$; the y_1 group is $p_7 + p_8$. So

  $$\Pr(T_1) = (p_5{+}p_6)\sigma(l),\ \Pr(T_2) = (p_7{+}p_8)\sigma(l),\ \Pr(T_3) = (p_5{+}p_6)\sigma(r),\ \Pr(T_4) = (p_7{+}p_8)\sigma(r).$$

- R-pure strategies (labels 9–12): analogously,

  $$\Pr(T_5) = (p_9{+}p_{10})\sigma(l),\ \Pr(T_6) = (p_{11}{+}p_{12})\sigma(l),\ \Pr(T_7) = (p_9{+}p_{10})\sigma(r),\ \Pr(T_8) = (p_{11}{+}p_{12})\sigma(r)$$

Notice the two facts we will exploit: the third component (x_2 vs y_2) of an M-pure strategy is irrelevant (never reached under M), so pairs like (M,x_1,x_2) and (M,x_1,y_2) contribute together; likewise for R-pure strategies and IS-2.

Step 3. Construct the behavior strategy by "marginalizing" at each information set. Define:

- At the root: play L, M, R with probabilities

$$P_L = p_1 + p_2 + p_3 + p_4, \qquad P_M = p_5 + p_6 + p_7 + p_8, \qquad P_R = p_9 + p_{10} + p_{11} + p_{12}.$$

- At IS-2 (conditional on having reached it, i.e., conditional on having played M): play x_1 with probability

$$Q_{x_1} = \frac{p_5 + p_6}{P_M}$$

(the conditional probability, under the mixed strategy, that Player 1 plays x_1 at IS-2 given she played M), and y_1 with the complementary probability.

- At IS-3 analogously: x_2 with probability $Q_{x_2} = (p_9 + p_{10})/P_R$.

This is well-defined provided $P_M, P_R > 0$; if a root action has probability zero, the corresponding information set is never reached and any conditional probabilities are irrelevant.

Step 4. Verify equivalence. Under this behavior strategy, the terminal probabilities are

$$\begin{aligned}\Pr(T_4) &= P_M \cdot \sigma(r) \cdot (1 - Q_{x_1}) = (p_5 + p_6 + p_7 + p_8) \cdot \sigma(r) \cdot \frac{p_7 + p_8}{p_5 + p_6 + p_7 + p_8} \\ &= (p_7 + p_8)\, \sigma(r),\end{aligned}$$

which matches the value from Step 2 exactly. The eight other terminal probabilities check similarly. The behavior strategy is realization-equivalent to the mixed strategy.

Intuitively, to go from a mixed strategy (one big coin) to a behavior strategy (many small independent coins), *condition* on reaching each information set. The behavior probability at IS-2 is the mixed-strategy conditional probability that the player takes each action there, given she got there. Under perfect recall, these conditional distributions can be defined consistently at every information set.

Answer. *Part (d): merging IS-2 and IS-3 breaks perfect recall.*

Step 1. What perfect recall means. A game has *perfect recall* if every player remembers her own past actions (and all information sets she has previously visited). In the original game, when Player 1 reaches her later information set she remembers whether she just came from M (in which case she's in IS-2) or R (in which case she's in IS-3), because the two information sets are separate.

Step 2. After the merge. If IS-2 and IS-3 are merged into one information set (call it IS'), Player 1 can no longer tell whether she previously chose M or R. Yet she *did* choose one of them — so she has forgotten her own past action. Perfect recall fails.

Step 3. Part (b) still holds: any behavior strategy has an equivalent mixed strategy. With the merge, Player 1's behavior strategy at IS' is a single distribution over $\{x, y\}$ (the two actions at the merged set); call the probabilities q_1, q_2. The terminal-node probabilities under this behavior strategy are still products of the local independent randomizations:

$$\Pr(T_1) = p_2\,\sigma(l)\,q_1, \qquad \Pr(T_2) = p_2\,\sigma(l)\,q_2,$$
$$\Pr(T_5) = p_3\,\sigma(l)\,q_1, \qquad \Pr(T_6) = p_3\,\sigma(l)\,q_2,$$

(and similarly for $\sigma(r)$). The realization-equivalent mixed strategy is: (L, x) with probability p_1, (M, x) with probability p_2q_1, (M, y) with probability p_2q_2, (R, x) with probability p_3q_1, (R, y) with probability p_3q_2. Verification is the same tabulation as in Part (b).

Step 4. Part (c) can fail: some mixed strategies have NO realization-equivalent behavior strategy. Consider the specific mixed strategy in which Player 1 plays (M, x) with probability 1/2 and (R, y) with probability 1/2. Suppose Player 2 plays the pure strategy l ($\sigma(l) = 1$). Then the mixed strategy generates

$$\Pr(T_1) = \tfrac{1}{2},\ \Pr(T_6) = \tfrac{1}{2},\ \Pr(T_k) = 0 \text{ for all other } k.$$

Now attempt to find a behavior strategy that produces the same distribution. It would randomize at the root over L, M, R with some p_1, p_2, p_3 and at IS' over x, y with some q_1, q_2. The terminal probabilities under this behavior strategy would be

$$\Pr(T_1) = p_2q_1, \quad \Pr(T_2) = p_2q_2, \quad \Pr(T_5) = p_3q_1, \quad \Pr(T_6) = p_3q_2.$$

We require the following four constraints simultaneously:

$$\begin{aligned}
p_2q_1 &= \tfrac{1}{2} && \Rightarrow p_2 > 0 \text{ and } q_1 > 0,\\
p_2q_2 &= 0 && \Rightarrow q_2 = 0 \text{ (since } p_2 > 0),\\
p_3q_1 &= 0 && \Rightarrow p_3 = 0 \text{ (since } q_1 > 0),\\
p_3q_2 &= \tfrac{1}{2} && \Rightarrow p_3 > 0 \text{ AND } q_2 > 0.
\end{aligned}$$

The second and fourth conditions are contradictory ($p_3 = 0$ vs. $p_3 > 0$), so no behavior strategy can produce this distribution.

What went wrong? With perfect recall, Player 1's behavior at IS-2 could be independent of her behavior at IS-3 because they were different information sets. In the merged game, Player 1 must use a *single* distribution at IS' regardless of whether she came from M or R. That kills her ability to *correlate* the two decisions — but the mixed strategy $\{(M, x)$ w.p. 1/2, (R, y) w.p. $1/2\}$ requires exactly such correlation: pick x if the previous action was M, pick y if it was R. Behavior strategies at a single information set cannot capture that correlation, so no behavior strategy is realization-equivalent to this mixed strategy.

Intuitively, Kuhn's theorem tells us: under perfect recall, the two randomization languages (behavior vs. mixed) are equally expressive — (b) and (c) both hold. Once perfect recall fails, mixed strategies can still "remember" their own earlier tosses through the pure-strategy support, but behavior strategies (which randomize afresh at each information set) cannot. The two languages become asymmetric: behavior $\subsetneq$ mixed. (For a general proof, see Fudenberg and Tirole, *Game Theory*, MIT Press 1991, p. 87.)

2. **MWG 8.C.4 – Best response, strict dominance, and correlation.** A three-player game Γ_N with $S_1 = \{L, M, R\}$, $S_2 = \{U, D\}$, $S_3 = \{l, r\}$. Player 1's payoffs from each of her three strategies, conditional on the strategies of Players 2 and 3, are shown as triples (u_L, u_M, u_R) in the four cells below, with parameters $(\pi, \varepsilon, \delta) \gg 0$ and $\delta < 4\varepsilon$:

	Player 3: l	Player 3: r
Player 2: U	$\pi + 4\varepsilon,\ \pi - \delta,\ \pi - 4\varepsilon$	$\pi - 4\varepsilon,\ \pi + \frac{\delta}{2},\ \pi + 4\varepsilon$
Player 2: D	$\pi + 4\varepsilon,\ \pi + \frac{\delta}{2},\ \pi - 4\varepsilon$	$\pi - 4\varepsilon,\ \pi - \delta,\ \pi + 4\varepsilon$

Player 2 plays U with probability σ and D with probability $1-\sigma$; Player 3 plays l with probability τ and r with probability $1 - \tau$. Denote Player 1's expected payoff from action $A \in \{L, M, R\}$ (given σ, τ) by P_A. Direct calculation:

$$P_L = \pi + 4\varepsilon(2\tau - 1),$$
$$P_M = \pi + \delta\Big[\frac{-3\sigma + 3}{2}\sigma - 3\sigma\tau - 1\Big] \text{ (as printed; see comment in Part (a) below)},$$
$$P_R = \pi + 4\varepsilon(1 - 2\tau).$$

(a) Show that M is never a best response to any *independent* randomization by Players 2 and 3.

(b) Show that M is not strictly dominated.

(c) Show that M *can* be a best response if Players 2 and 3 are allowed to randomize in a *correlated* way.

Answer. *Part (a): M is never a best response to independent randomization.*

The key observation: P_L and P_R depend only on τ (not on σ), while P_M depends on both. We case-split on τ.

Step 1. Case 1: $\tau > \frac{1}{2}$. $\partial P_M/\partial\sigma < 0$: raising σ hurts P_M (for these parameter values). So the largest P_M over $\sigma \in [0, 1]$ is at $\sigma = 0$. Evaluating at $\sigma = 0$:

$$P_M(\sigma = 0) = \pi + \delta\big[\tfrac{3}{2}\tau - 1\big] < \pi + 4\varepsilon\big[\tfrac{3}{2}\tau - 1\big] < \pi + 4\varepsilon(2\tau - 1) = P_L,$$

where the first inequality uses $\delta < 4\varepsilon$ (given) and the second uses $\tau > 1/2$ (so $2\tau - 1 > 3/2\tau - 1$, wait let me check: $2\tau - 1 = 2\tau - 1$ and $3\tau/2 - 1$; for $\tau > 1/2$, $2\tau > 3\tau/2$ iff $\tau/2 > 0$, always true). So $P_M < P_L$ throughout, meaning M is dominated by L in the maximum, and hence is never a best response.

Step 2. Case 2: $\tau < \frac{1}{2}$. $\partial P_M/\partial\sigma > 0$: raising σ helps P_M. So the largest P_M is at $\sigma = 1$. Evaluating:

$$P_M(\sigma = 1) = \pi + \delta\left[\tfrac{1}{2} - \tfrac{3}{2}\tau\right] < \pi + 4\varepsilon\left[\tfrac{1}{2} - \tfrac{3}{2}\tau\right] < \pi + 4\varepsilon(1 - 2\tau) = P_R,$$

where the first inequality uses $\delta < 4\varepsilon$ and the second uses $\tau < 1/2$. So M is dominated by R in the maximum, again never a best response.

Step 3. Case 3: $\tau = \frac{1}{2}$. Here $P_L = P_R = \pi$ (the $4\varepsilon(2\tau - 1)$ term vanishes for $\tau = 1/2$), and $P_M = \pi - \delta/4 < \pi = P_L = P_R$. So M is strictly worse than both L and R.

Step 4. Conclusion. In every case, M is strictly worse than one of the other two pure strategies given any (σ, τ). So M is never a best response to any independent randomization.

Answer. *Part (b): M is not strictly dominated.*

We show that no mixed strategy over $\{L, R\}$ strictly dominates M. Let $\rho \in [0, 1]$ be the probability of playing R in the candidate dominating mixture (with $1 - \rho$ on L). For each of the four opponent-strategy profiles $(P_2, P_3) \in \{U, D\} \times \{l, r\}$, Player 1's payoff from playing M is (from the table): $\pi - \delta$, $\pi + \delta/2$, $\pi + \delta/2$, $\pi - \delta$. Her payoff from playing the mixture $(1 - \rho)L + \rho R$ is the corresponding weighted average of the L and R payoffs.

Step 1. Set up the four dominance inequalities. Strict dominance requires the mixture's payoff to exceed M's at every opponent profile:

$$\begin{aligned}
(U, l):&\ \ (1 - \rho)(\pi + 4\varepsilon) + \rho(\pi - 4\varepsilon) > \pi - \delta,\\
(U, r):&\ \ (1 - \rho)(\pi - 4\varepsilon) + \rho(\pi + 4\varepsilon) > \pi + \tfrac{\delta}{2},\\
(D, l):&\ \ (1 - \rho)(\pi + 4\varepsilon) + \rho(\pi - 4\varepsilon) > \pi + \tfrac{\delta}{2},\\
(D, r):&\ \ (1 - \rho)(\pi - 4\varepsilon) + \rho(\pi + 4\varepsilon) > \pi - \delta.
\end{aligned}$$

Simplifying,

$$\begin{aligned}
(U, l):&\ \ \pi + 4\varepsilon - 8\varepsilon\rho > \pi - \delta &\iff&\ \ \rho < \tfrac{1}{2} + \tfrac{\delta}{8\varepsilon},\\
(U, r):&\ \ \pi - 4\varepsilon + 8\varepsilon\rho > \pi + \tfrac{\delta}{2} &\iff&\ \ \rho > \tfrac{1}{2} + \tfrac{\delta}{16\varepsilon},\\
(D, l):&\ \ \pi + 4\varepsilon - 8\varepsilon\rho > \pi + \tfrac{\delta}{2} &\iff&\ \ \rho < \tfrac{1}{2} - \tfrac{\delta}{16\varepsilon},\\
(D, r):&\ \ \pi - 4\varepsilon + 8\varepsilon\rho > \pi - \delta &\iff&\ \ \rho > \tfrac{1}{2} - \tfrac{\delta}{8\varepsilon}.
\end{aligned}$$

Step 2. Find a contradiction. Combine the (U, r) and (D, l) constraints:

$$\rho > \tfrac{1}{2} + \tfrac{\delta}{16\varepsilon} \quad \text{and} \quad \rho < \tfrac{1}{2} - \tfrac{\delta}{16\varepsilon}.$$

These say $\rho > \rho_{U,r} \equiv 1/2 + \delta/(16\varepsilon)$ and $\rho < \rho_{D,l} \equiv 1/2 - \delta/(16\varepsilon)$. Since $\delta > 0$, $\rho_{U,r} > \rho_{D,l}$, so no ρ can satisfy both. Contradiction.

Step 3. Conclusion. No mixed strategy over $\{L, R\}$ strictly dominates M. Together with Part (a), this gives a striking result: M is *never a best response* to any independent randomization, yet is *not strictly dominated.* Being never-best-response is strictly weaker than being strictly dominated.

Intuitively, the standard equivalence "never best response $\Leftrightarrow$ strictly dominated" holds for two-player games but breaks in three-plus-player games when opponents randomize *independently.* In two-player games, the set of possible opponent behaviors (mixed strategies) is convex, and standard separating-hyperplane arguments deliver the equivalence. With three or more opponents, the set of independent-randomization products is not convex — and M can exploit correlations that don't exist in that set, but do exist in the larger set of joint distributions.

Answer. *Part (c): M can be a best response to a correlated randomization.*

Step 1. The correlated strategy. Suppose Players 2 and 3 correlate as follows: play (U, r) with probability 1/2 and (D, l) with probability 1/2. Under this correlation, when Player 2 plays U, Player 3 always plays r; when Player 2 plays D, Player 3 always plays l.

Step 2. Compute expected payoffs.

- From L: $\frac{1}{2}(\pi - 4\varepsilon) + \frac{1}{2}(\pi + 4\varepsilon) = \pi$. (The good and bad states of L balance out.)
- From R: $\frac{1}{2}(\pi + 4\varepsilon) + \frac{1}{2}(\pi - 4\varepsilon) = \pi$. (Same balance.)
- From M: $\frac{1}{2}(\pi + \frac{\delta}{2}) + \frac{1}{2}(\pi + \frac{\delta}{2}) = \pi + \frac{\delta}{2}$. (Both realized states put M in a high-payoff cell.)

Any mixture over L and R gives π; M gives $\pi + \delta/2 > \pi$. Hence M is the unique best response.

Step 3. Why does this correlation help M? Look at the payoff table for Player 1. The strategy M earns a *high* payoff at (U, r) and at (D, l) — the two "anti-diagonal" cells. It earns a *low* payoff at (U, l) and at (D, r) — the "main-diagonal" cells. Independent randomization puts probability $\sigma\tau$ on (U, l) and $(1-\sigma)(1-\tau)$ on (D, r) (both diagonal, low for M), diluting M's expected payoff. A correlated distribution that puts weight only on the anti-diagonal turns M into a winner.

Intuitively, M is a strategy specialized to a particular kind of anti-correlated opponent behavior. Independent mixtures can never generate that anti-correlation, but correlated ones can. This is precisely the observation that motivates the concept of *correlated equilibrium* (Aumann 1974), which extends Nash equilibrium to allow players to condition their strategies on a common public signal.

3. **Iterated deletion of strictly dominated strategies in a three-player game.** Three firms A, B, C simultaneously and independently decide whether to enter a new market or stay out (e.g., remain in their existing personal-computer software business rather than move to smartphone apps). Payoff triples are (u_A, u_B, u_C); each cell shows Firm A's payoff first (row player), Firm B's second (column player), and Firm C's third (matrix player). The payoff matrix is:

Matrix 1: Firm C chooses Enter.

	B: Enter	B: Stay Out
A: Enter	14, 24, 32	8, 30, 27
A: Stay Out	30, 16, 24	13, 12, 50

Matrix 2: Firm C chooses Stay Out.

	B: Enter	B: Stay Out
A: Enter	16, 26, 30	31, 16, 24
A: Stay Out	31, 23, 14	14, 26, 32

Find the set of strategy profiles that survive IDSDS. Is the surviving profile unique?

Answer.

Step 1. Check whether Firm C has a strictly dominant strategy. Compare Firm C's payoff (the third component) across the two matrices, cell by cell. For each (s_A, s_B) profile of Firms A and B, we ask: does Firm C prefer to Enter (Matrix 1) or to Stay Out (Matrix 2)?

(s_A, s_B)	C: Enter	C: Stay Out
(E, E)	**32**	30
(E, O)	**27**	24
(O, E)	**24**	14
(O, O)	**50**	32

In every one of the four cells, Enter strictly dominates Stay Out for Firm C. So Firm C's strategy Stay Out is strictly dominated, and we can eliminate Matrix 2 entirely.

Step 2. Reduced game after eliminating C's Stay Out. We are left with Matrix 1:

	B: Enter	B: Stay Out
A: Enter	14, 24, 32	8, 30, 27
A: Stay Out	30, 16, 24	13, 12, 50

Step 3. Check Firm A's dominant strategy in the reduced game. Compare Firm A's payoff (first component) across her two rows, for each of Firm B's choices:

- If B plays Enter: A gets 14 from Enter vs. 30 from Stay Out. Stay Out wins $(30 > 14)$.
- If B plays Stay Out: A gets 8 from Enter vs. 13 from Stay Out. Stay Out wins $(13 > 8)$.

So for Firm A, Stay Out strictly dominates Enter in the reduced game. Eliminate A's Enter row.

Step 4. Further reduced game. With A committed to Stay Out (and C to Enter), only one row remains:

	B: Enter	B: Stay Out
A: Stay Out	30, **16**, 24	13, **12**, 50

Firm B now faces an individual decision problem: choose Enter (payoff 16) or Stay Out (payoff 12). Enter is strictly better.

Step 5. Surviving strategy profile. The unique profile surviving IDSDS is

$$(s_A, s_B, s_C) = (\text{Stay Out, Enter, Enter}),$$

with payoffs $(30, 16, 24)$.

Uniqueness. Because at each round of deletion the eliminated strategy was *strictly* dominated (not just weakly), no other profile can survive. Moreover, IDSDS is known to be order-independent for strict dominance: no matter which player's dominated strategy we delete first, we reach the same surviving set. Here that surviving set is a singleton — an unambiguous prediction of the game.

Intuitively, this is an "anti-coordination" outcome: exactly two firms enter the market (Firms B and C), and one stays out (Firm A). The parameters of the payoff table are calibrated so that Firm C's market-share advantage from entering is so large that it always enters; given C's entry, Firm A prefers to stay out (avoid over-crowding); and given the other two, Firm B finds it profitable to squeeze in. The IDSDS analysis exposes this asymmetric outcome as the only rationalizable equilibrium.