

Puzzle: Does Occasional Simulation Enable Cooperation?

(Puzzle in honor of Joe Halpern’s 70th birthday)

VINCENT CONITZER¹

Carnegie Mellon University and University of Oxford

Please send solutions to the author by e-mail, with the title of this puzzle in the subject header. By agreement with the editors, the best solution will be published in the next issue of SIGecom Exchanges, provided that that solution is of sufficiently high quality. Quality is judged by the author, taking into account at least soundness, completeness, and clarity of exposition. (Incidentally, there is another birthday puzzle for which we still need a solution [1]!)

This is a puzzle in honor of Joe Halpern’s 70th birthday and the June 2023 workshop (“Halpernfest”) associated with it. This workshop was held at Cornell University.

Consider the following 2-player game:

2,3	0,4
1,0	1,1

By iterated strict dominance, its only solution is (Bottom, Right) with utilities (1, 1).

However, now consider the following twist to the game: There is a 50% probability that Player 1 (the row player) first gets to *simulate* Player 2, once, before making her own move. Simulating once means that if Player 2 plays the mixed action p Left + $(1 - p)$ Right, then Player 1 will observe one draw from that distribution, and base her decision whether to play Top or Bottom on whether that draw is Left or Right. Player 2’s actual action in the game is determined by a fresh draw. Player 2 does not observe whether Player 1 gets to simulate him or not (but knows that this happens with 50% probability).

Conceptually, there are different ways to think about the game. One is the following: Player 2 needs to choose some value p beforehand, once. This value (not observed by Player 1) is then used to draw Player 2’s simulated play (if there is simulated play), and again (i.i.d.) for Player 2’s actual play.

(a1) Is there an equilibrium in which utilities (2, 3) are always obtained? Explain. (It helps to think of this as an extensive-form game.)

Another way to think about this is that Player 2 does not choose a value p beforehand, but rather at one or two points “wakes up” and has to decide which mixed

¹Thanks to Vojta Kovarik, Caspar Oesterheld, and Emanuel Tewelde for helpful feedback.

action to play (i.e., choose a p), not knowing whether he is in a simulation or not. (But, if he is in a simulation, then he still wants to play in a way that maximizes his expected utility in the real world. Also, if Player 2 plays in the real world after having been simulated, then of course he does not remember being simulated. That is, Player 2 cannot distinguish any of his different kinds of awakening from each other.)

(a2) Waking up as Player 2, what probability would you assign to being in a simulation, and how would you play? Does your answer to the above question change?

(b) Now answer these questions again for the following game:

2,3	0,5
1,0	1,2

REFERENCES

Vincent Conitzer. Puzzle: The AI Circus. (Puzzle in honor of Tuomas Sandholm's 50th birthday.) *SIGecom Exchanges*, 17(2):76–77, October 2019.