Multiselves equilibria: Are you thinking what I'm thinking?

Paul W. Goldberg
JECCO, Edinburgh University
16th June 2025

My intro to game theory: *Scientific American* 1983

Superrationality

文A 3 languages

Article Talk Read Edit View history Tools

From Wikipedia, the free encyclopedia

> This article **possibly contains original research**. Please improve it by verifying the claims made and adding inline citations. Statements consisting only of original research should be removed. *(January 2019)* (*Learn how and when to remove this message*)

In economics and game theory, a participant is considered to have **superrationality** (or **renormalized rationality**) if they have perfect rationality (and thus maximize their utility) but assume that all other players are superrational too and that a superrational individual will always come up with the same strategy as any other superrational thinker when facing the same problem. Applying this definition, a superrational player who assumes they are playing against a superrational opponent in a prisoner's dilemma will cooperate while a rationally self-interested player would defect.

This decision rule is not a mainstream model in game theory and was suggested by Douglas Hofstadter in his article, series, and book *Metamagical Themas*[1] as an alternative type of rational decision making different from the widely accepted game-theoretic one. Hofstadter provided this definition: "Superrational thinkers, by recursive definition, include in their calculations the fact that they are in a group of superrational thinkers."[1]

Douglas Hofstadter's Luring Lottery

Douglas Hofstadter writes: The Luring Lottery, proposed in my June column, created quite a stir. Let me remind you that it was open to anyone; all you had to do was send a postcard with a clearly specified positive integer on it stating how many entries you wanted to make. This integer was to be, in effect, your "weight" in the final drawing, so that if you wrote "100," your name would be 100 times as likely to be drawn as that of someone who wrote "1." The only catch was that the cash value of the prize was inversely proportional to the sum of all the weights received by June 30. Specifically, the prize to be awarded was $1,000,000/*W*, where *W* is the sum of all the weights sent in.

The Luring Lottery was set up as an

Douglas Hofstadter's Luring Lottery

them. What was the breakdown of entries? Part of it is given in this table:

1: 1,133
2: 31
3: 16
4: 8
5: 16
6: 0
7: 9
8: 1
9: 1
10: 49
100: 61
1,000: 46
1,000,000: 33
1,000,000,000: 11
602,300,000,000,000,000,000,000 (Avogadro's number): 1
Googol (10^{100}): 9
Googolplex (10^{googol}): 14

Curiously, many if not most of the people who submitted just one entry patted themselves on the back for being "cooperators." Stuff and nonsense! The *real* cooperators were those among the

this. Scores of readers strained their hardest to come up with inconceivably large numbers. Some filled their entire postcard with tiny 9's, others filled their card with rows of exclamation points, creating iterated factorials of gigantic size. A handful of people carried this game much further, recognizing that the optimal solution avoids all pattern (for the reason see Gregory J. Chaitin's article "Randomness and Mathematical Proof," *Scientific American,* Vol. 232, No. 5, pages 47–52; May, 1975). It consists simply of a "dense pack" of definitions built on definitions, followed by a final line in which the "fanciest" of the definitions is applied to a relatively small number such as 2 or, better yet, 9.

I received, as I say, a few such entries. Some of them exploited such powerful concepts of mathematical logic and set theory that to evaluate which one was the largest became a serious problem, and it is not even clear that I, or for that matter anyone else, would be able to determine which is the largest integer submitted. I was strongly reminded of

Multiselves in Learning in games, and elsewhere

Example: Fictitious play. Robinson (1951): if both players use FP in a 2-player zero-sum game, payoffs converge to the game's value. Note: assumes both players use FP. FP plays against itself.

Similar kinds of results in other contexts, e.g. replicator dynamics.

Multiselves in Learning in games, and elsewhere

Example: Fictitious play. Robinson (1951): if both players use FP in a 2-player zero-sum game, payoffs converge to the game's value. Note: assumes both players use FP. FP plays against itself.

Similar kinds of results in other contexts, e.g. replicator dynamics.

Uniform algorithms, in setting of distributed algorithms:
Each processor gets the same algorithm, they have to achieve something such as leader election.

Multiselves in Learning in games, and elsewhere

Example: Fictitious play. Robinson (1951): if both players use FP in a 2-player zero-sum game, payoffs converge to the game's value. Note: assumes both players use FP. FP plays against itself.

Similar kinds of results in other contexts, e.g. replicator dynamics.

Uniform algorithms, in setting of distributed algorithms:
Each processor gets the same algorithm, they have to achieve something such as leader election.

Rendezvous problem (or dilemma): 2 robots want to end up in the same place.

An obscure problem in ML theory

Introductory PAC learning problem:

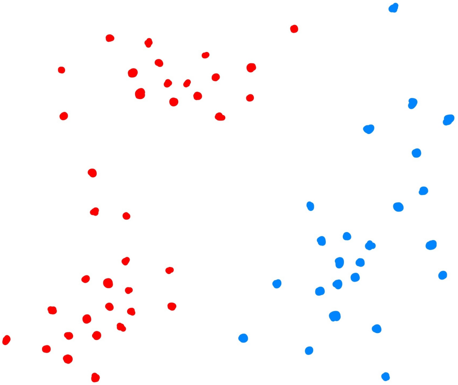

Want to classify linearly separable classes of points

Fact: any consistent linear separator works

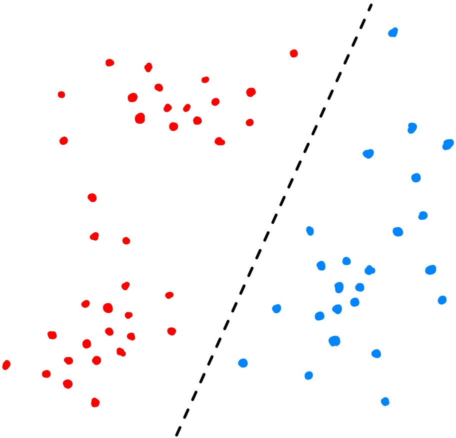

An obscure problem in ML theory

Fit distributions to data sets separately, use them to classify:

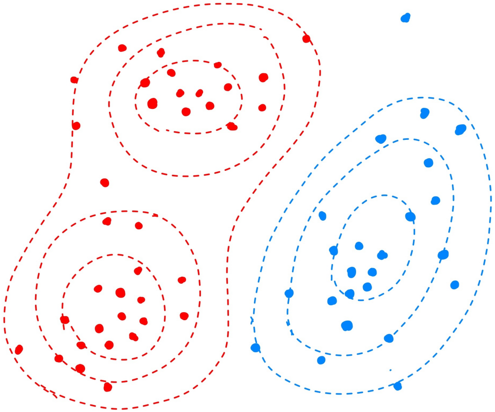

Some Discriminant-based PAC Algorithms, *JMLR* (2006)

An obscure problem in ML theory

Design algorithm $\mathscr{A}$ mapping point set to distribution such that:

If we run $\mathscr{A}$ *twice*, once on the red points and once on the blue points, you get PAC learning when you combine the results in the obvious way.

$\mathscr{A}$, "seeing" the red points, knows that another copy of $\mathscr{A}$ is processing the blue points, and vice versa.

An obscure problem in ML theory

Design algorithm $\mathscr{A}$ mapping point set to distribution such that:

If we run $\mathscr{A}$ *twice*, once on the red points and once on the blue points, you get PAC learning when you combine the results in the obvious way.
$\mathscr{A}$, "seeing" the red points, knows that another copy of $\mathscr{A}$ is processing the blue points, and vice versa.

Problems: $\mathscr{A}$ is complicated/artificial, and I never figured out how to learn propositional disjunctions...

It's a "multi-selves" solution concept, of sorts

Symmetric equilibria

In a contest, agents incur (non-refundable) disutility or cost; the more you pay the more likely you are to win some prize.
E.g. all-pay auctions

Symmetric equilibria

In a contest, agents incur (non-refundable) disutility or cost; the more you pay the more likely you are to win some prize.
E.g. all-pay auctions

Tullock contest: prize is awarded by lottery; in general each player i has a non-decreasing production function f_i mapping effort to output.

Homogenous players: all the f_i's are the same. Notice that there's a symmetric equilibrium. Although in general there may be other equilibria.

Symmetric equilibria

In a contest, agents incur (non-refundable) disutility or cost; the more you pay the more likely you are to win some prize.
E.g. all-pay auctions

Tullock contest: prize is awarded by lottery; in general each player i has a non-decreasing production function f_i mapping effort to output.

Homogenous players: all the f_i's are the same. Notice that there's a symmetric equilibrium. Although in general there may be other equilibria.

Symmetric equilibrium: as a player, I take the view that my opponents are thinking what I'm thinking.

Symmetric equilibria

Objection: if you count on your opponents doing what you do, best to make no effort (incur no cost) and keep same probability to win the prize...

Symmetric equilibria

Objection: if you count on your opponents doing what you do, best to make no effort (incur no cost) and keep same probability to win the prize...

Literature on contests like this one tends to focus on the symmetric equilibrium, since it's know to exist. Arguably we should also consider other equilibria, since ths symmetric equilibrium does not result from players counting on their opponents to do the same thing...

Symmetric equilibria

Objection: if you count on your opponents doing what you do, best to make no effort (incur no cost) and keep same probability to win the prize...

Literature on contests like this one tends to focus on the symmetric equilibrium, since it's know to exist. Arguably we should also consider other equilibria, since ths symmetric equilibrium does not result from players counting on their opponents to do the same thing...

Is there a more meaningful scenario of counting on our opponents to do the same thing?

A warm-up: the Keynesian beauty contest

"It is not a case of choosing those [faces] that, to the best of one's judgment, are really the prettiest, nor even those that average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees." (Keynes, General Theory of Employment, Interest and Money, 1936).

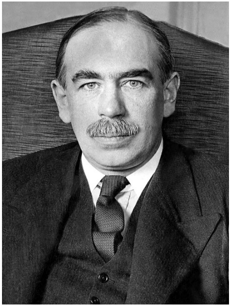

A warm-up: the Keynesian beauty contest

Multiselves decision theory

- Imperfect-Recall Games: Equilibrium Concepts and Their Complexity. *IJCAI* 2024
- The Computational Complexity of Single-Player Imperfect-Recall Games. *IJCAI* 2023

General set-up (examples to follow): extensive-form "game" with one player.

Decision tree is partitioned into information sets; within each info set all nodes have the same outdegree/labels.

The absentminded driver

Sequence of exits off road, some exit other than the first is desired. How to proceed?

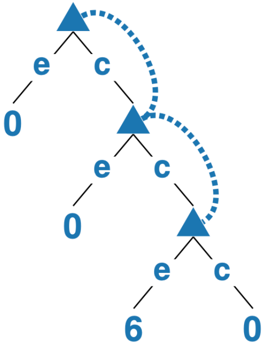

Absentmindedness: an info set may contain 2 nodes one of which follows from the other.
Note link to multiselves idea

Sleeping Beauty

1. SB goes to sleep
2. Flip a fair coin, H: awaken SB once, T: awaken SB twice
3. SB forgets each awakening
4. At each awakening, SB is asked to give degree of belief the coin outcome is H

Sleeping Beauty

1. SB goes to sleep
2. Flip a fair coin, H: awaken SB once, T: awaken SB twice
3. SB forgets each awakening
4. At each awakening, SB is asked to give degree of belief the coin outcome is H

Answer 1: $\frac{1}{2}$: awaking provides no new information on the coin

Answer 2: $\frac{1}{3}$: if repeated, about $\frac{1}{3}$ of wakings occur when heads is shown

The set-up: game tree with nodes $\mathscr{H}$, terminal nodes $\mathscr{Z} \subseteq \mathscr{H}$. Terminal nodes have real-valued *utilities*. There may be *chance nodes* that select one of their descendents with given probabilities. The other nodes are the player's *decision nodes*, partitioned into *info sets*.

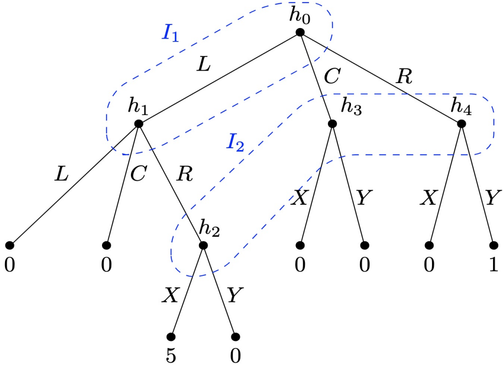

Solution concept: for each info set I, a distribution over labels of actions that can be taken at nodes in I.

Given a game tree, utility maximisation is a problem of maximising a polynomial function subject to a collection of constraints.
NP-hard to compute an optimal strategy; results date back to Koller and Megiddo (1992).
$\exists\mathbb{R}$-complete (Gimbert et al, 2020).

So, we look for solution concept that may be easier to find...

Causal vs. Evidential Decision Theory

Two alternative solution concepts to 1-player extensive-form games with absentmindedness.

- Causal Decision Theory: Generalised Thirding (GT)
- Evidential Decision Theory: Generalized Double Halving (GDH)

Strategy: maps each info set I to a distribution D_I over labels of its outgoing arcs.
Any strategy leads to an expected utility.

Deviation:

- (EDT,GDH): replace D_I with some D'_I
- (CDT,GT): there's a distribution P_I over nodes in I. Replace D_I with D'_I at individual node sampled from P_I, assuming you leave D_I alone elsewhere.

Equilibrium: strategy with no profitable deviation.

(I think (CDT,GT) assumes SB always has opportunity to deviate, (EDT,GDH) assumes it's infrequent...)

Example

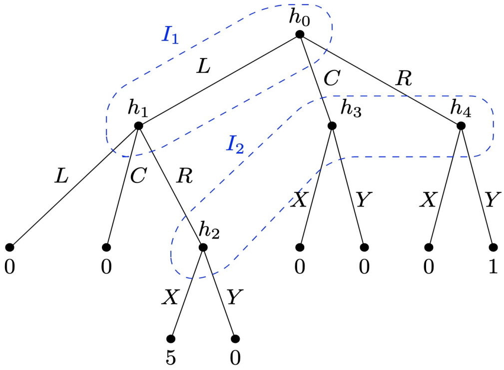

Suppose player plays (R, X). At info set I_1, consider deviation to L.
CDT: player will not deviate elsewhere in I_1, may win 5
EDT: deviation will lead to action history (L, L).

Suppose player is in info set I_1, playing (C, X).
No (EDT,GDH) deviation to a pure strategy, but to $\frac{1}{2}L + \frac{1}{2}R$.

Some results, TOCG'23

In games without absentmindedness, both solution concepts are the same.

Moving from global to local maximisation, takes to the class CLS... Also, (CDT,GT)-equilibrium is sort-of CLS-complete. In fact (CDT,GT)-equilibria are KKT points of the utility function. Containment requires a lower bound on positive visit frequencies to be easily obtainable. Hardness holds even for no chance nodes, tree depth 5, and one one info set.

CLS-completeness extends to (EDT,GDH) equilibria since it applies to games without absentmindedness.

Various NP-hardness results for search for equilibria having specific properties.

Multi-player version (Tewolde et al, IJCAI'24)

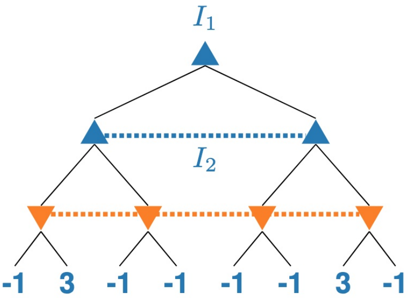

Forgetful penalty shoot-out. There is no NE

Results

Multi-player

	Nash (D)	EDT (D)	CDT (S)
exact	$\exists\mathbb{R}$-hard and in $\exists\forall\mathbb{R}$ (Thms. 1 & 3)		—
1/exp	Σ_2^{P}-complete (Thms. 2 and 4)		PPAD-complete (Thm. 6)
1/poly			

Single-player

	Optimal (D)	EDT (S)	CDT (S)
exact	$\exists\mathbb{R}$-complete [Gimbert *et al.*, 2020]	—	—
1/exp	NP-complete [Koller and Megiddo, 1992; Tewolde *et al.*, 2023]	PLS-complete (Thm. 5*)	CLS-complete [Tewolde *et al.*, 2023]
1/poly		P (Cor. 22*)	P (Cor. 17)

green: new results; (S): search problem; (D): decision problem

Back to voting

Ongoing work with Tomasz Wąs.
Idea: in large-scale elections, we vote because we reckon others think like us. "social projection"
Model: n voters, m candidates. Each voter gets a signal about the candidates, drawn from a commonly-known distribution D.
Question: can we design distribution D that presents voters with a hard task? Variants:

- Voters are rewarded for voting for winner, à la KBC.
- ...or maybe, want to maximise welfare (seems challenging in the multi-winner setting)
- ...or maybe, have diverse utilities for candidates

Back to voting

Ongoing work with Tomasz Wąs.
Idea: in large-scale elections, we vote because we reckon others think like us. "social projection"
Model: n voters, m candidates. Each voter gets a signal about the candidates, drawn from a commonly-known distribution D.
Question: can we design distribution D that presents voters with a hard task? Variants:

- Voters are rewarded for voting for winner, à la KBC.
- ...or maybe, want to maximise welfare (seems challenging in the multi-winner setting)
- ...or maybe, have diverse utilities for candidates

Thanks!