


Zolt-Gilburne Imagination Seminar

Knowledge and Games

Sergei Artemov

October 1, 2009

Plato (5-4 Century B.C.)



One of the world's best known and most widely read and studied philosophers, a student of Socrates and teacher of Aristotle. His works on epistemology treat

Knowledge

as

Justified True Belief

Theory of knowledge

Modal logic of knowledge has been developing since 1950s: rich theory, many applications. It took the mathematical apparatus available. Informally, it treats

knowledge as true belief.

Does not capture justifications, hence serious limitations.

Logical Omniscience defect: an agent must know all the logical consequences of his/her assumptions, which is implausible. For example, a logically omniscient agent who knows the rules of Chess would also know whether White has a non-losing strategy.

Justification Logic

Incorporates *justification* into the mathematics of knowledge.

- Provides a mechanism for evidence-tracking.
- Leads to a coherent theory of logical omniscience.
- Furnishes a new foundation for the logic of knowledge and captures the whole of Plato's tripartite account of knowledge. The model is open-ended and more features can be added.

Justification Logic

Target Areas of Application:

- Mathematics;
- Computer Science and Artificial Intelligence;
- Economics and Game Theory;
- Cryptography;
- Philosophy;
- other disciplines.

Justification Logic

Developing since the mid-1990s. Provides simple and applicable answers to certain fundamental questions.

Geography: CUNY, Cornell, Stanford, UC San Diego, Switzerland, Netherlands, UK, Russia, Belgium, Italy, Brazil, Argentina, Japan, China, and others.

Encyclopedia entry: Stanford Encyclopedia of Philosophy (invited).

Grants and fellowships (NSF, DARPA-level): USA, Switzerland, Belgium, Russia.

Why was JL not invented earlier?

Based on Mathematical structure of Logic of Proofs:

K. Goedel 1933, 1938; **S.A.** 1995-2001, which was cutting-edge Mathematics research with many essential participants: **S. Kripke** (Princeton-CUNY), **G. Boolos** (MIT), **R. Solovay** (Berkeley), and many others.

LP technically could have been developed 30 years ago.

Imagination? Sense of direction? Priorities?

Rational decisions, informally


The standard game-theoretical assumption:

the player's rationality yields a payoff maximization given the player's knowledge.

Traditional Game Theory assumes *enough knowledge*

- to avoid uncertainty completely (Aumann)
- to deal with uncertainty probabilistically, i.e., when a player ***knows probability distribution*** of all consequences of his actions and is willing to take chances (von Neumann & Morgenstern)

Game Theory



Robert Aumann.

Alma Mater: City College of
New York,

Nobel Prize of 2005.

Pioneered mathematical studies
of *Rationality and Common
Knowledge*.


Game Theory



John von Neumann was an Hungarian American mathematician who made major contributions to mathematics, quantum mechanics, economics, and computer science. **Oskar Morgenstern** was an Austrian American economist. In 1944, he and von Neumann co-wrote *Theory of Games and Economic Behavior*, recognized as the first book on game theory.

Avoiding uncertainty

Game Tree:




Suppose A knows that B is rational. Then A knows that B will play *across*, thus delivering payoff 0 to A .

Hence the rational choice for A is *down*.

Probabilistic approach

Game Tree:



Suppose B is twice as likely to play *across* than *down*. Then A 's average payoff when A plays *across* is $2/3$ which is less than the payoff of 1 when playing *down*.


Hence the rational choice for A is *down*.

Challenge of uncertainty

There was no theory of making decisions under *uncertainty with unknown probability distribution.*

There is a solution, however, which logically follows from standard postulates of Game Theory and commonly accepted set of knowledge principles.

Example



A is mission control which has the option of sending to space a team which has not yet recovered from its previous mission (payoff 1), or sending a fresh crew B whose captain has been exposed to German measles. If B 's captain does not get sick, the mission will be a success (payoff 2), otherwise the mission should be aborted with failure (payoff 0).

Rational decisions, fresh approach


Knowledge-Based Rationality models decision-making strictly on the basis of players' knowledge:

at each node, rational players choose the best moves known to them.

New features:

- Clear separation of *best move* and *best known move*.
- *Players' knowledge* becomes the key element of game description.

Highest Known Payoff: *HKP*




For player *A*, $HKP(down)=1$.

$HKP(across)=0$, since *A does not know* that he will get any higher payoff than 0.

Solution: *A* plays *down*. (Exactly like in *Apollo 13...*)

Passive manipulation


Game Tree:



A is not sure of B's rational behavior, A plays down, payoffs 3,3.

B does not have the incentive to disclose his rationality since B wants A to move down.

Active manipulation



Suppose A is not aware of B and C 's rationality. Then A moves *left* to secure payoff 2. Actually, A gets 4 which is more than expected. Suppose also that B and C are smart enough to understand this. Then B can manipulate A by leaking the true information that C is rational. A then knows that *right* secures his payoff 3, which is higher than A 's known payoff of *left*: A plays *right* and gets 3 (less), B gets 4 (much more) and C gets 3 (more). C does not have an incentive to disclose that B is rational, hence

B wins without ever making a move!

Full knowledge is power

Model predictions:

- *Every game with rational players has a solution.* Rational players know which moves to make at each node.
- *Those who know the game in full know its solution, i.e., know everybody's moves.*

Partial knowledge can hurt

Model predictions:

- More knowledge yields a *higher known payoff* but not necessarily a *higher actual payoff*. So *nothing but the truth* can be misleading.
- Knowing *the whole truth*, however, yields a higher actual payoff.

When knowledge does not matter

Model predictions:

● In *strictly competitive (e.g. zero-sum) games*, any players' epistemic states lead to the same (maximin) solution. So, for strictly competitive games,

learning is irrelevant.

May be why military actions (normally zero-sum games) do not require sophisticated reasoning about the game: *just do it* normally works.

Imagination, vision, discovery

Vision is not exclusively the result of hard technical work, but rather the product of an independent mind and unfettered spiritual development. Freedom to pursue one's curiosity, a satisfying emotional life, and seeing the big picture, can help.

Discovery is a gift, and feels like a gift: a package arrives, you see it, you even try to guess its contents, but you don't know what it is and you proceed to unwrap it layer by layer, only to find out that it is breathtaking... but not at all what was expected.