Using Proof-of-Work to Coordinate

Adam Brandenburger* and Kai Steverson‡

* J.P. Valles Professor, NYU Stern School of Business
Distinguished Professor, NYU Tandon School of Engineering
Faculty Director, NYU Shanghai Program on Creativity + Innovation
Global Network Professor
New York University

‡ Postdoctoral Associate, Center for Neural Science
New York University

Research support provided by Marilyn Tsaih

Version 09/10/18

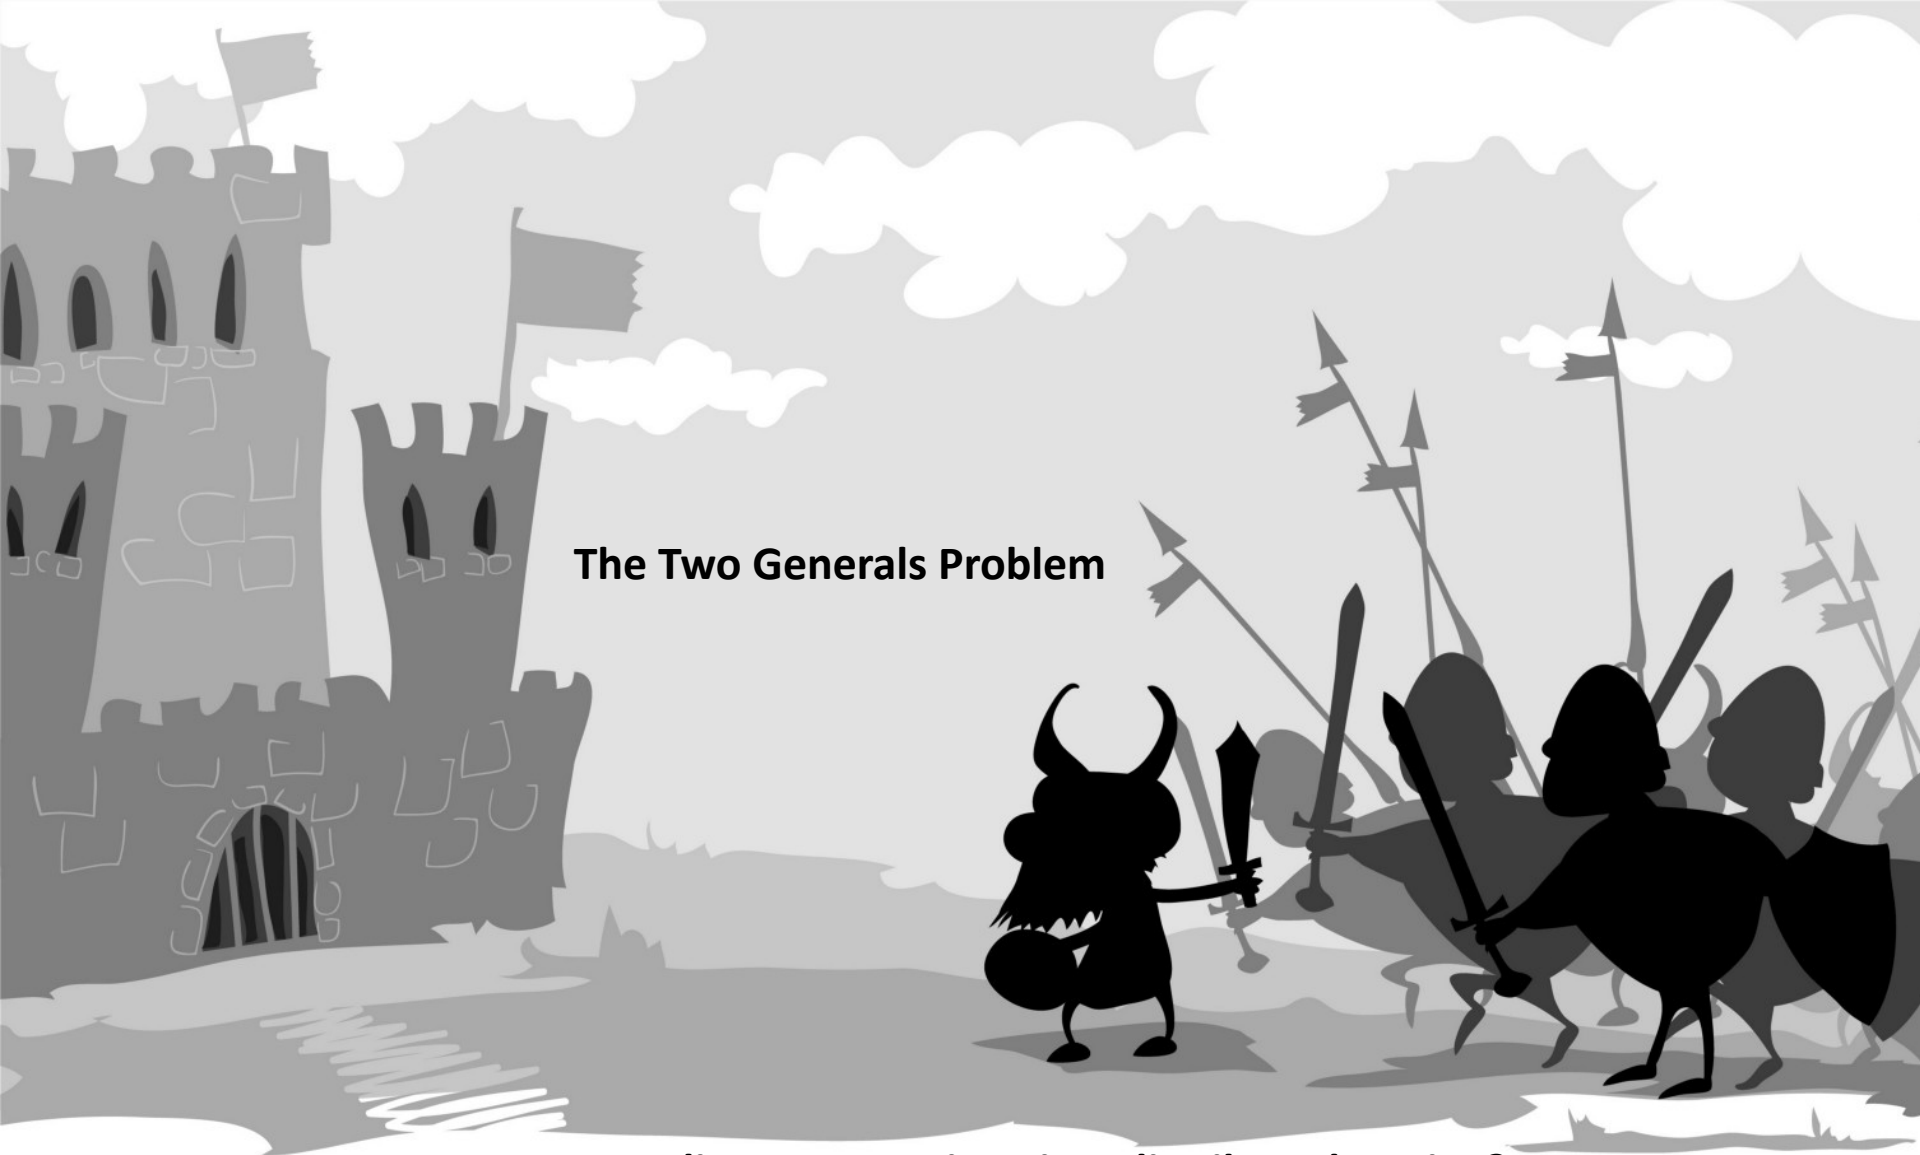

How can we coordinate our actions in a distributed setting?

E. Akkoyunlu, K. Ekanadham, and R. Huber, “Some Constraints and Tradeoffs in the Design of Network Communications,” 1975; J. Gray, “Notes on Data Base Operating Systems,” 1977; image from https://medium.com/coinmonks/a-note-from-anthony-if-you-havent-already-please-read-the-article-gaining-clarity-on-key-787989107969

Game-Theory Perspective: A First Take

If messenger #1 arrives safely

then both generals know the plan is to attack at dawn

If messenger #2 arrives safely

then both generals know that both generals know the plan

If messenger #3 arrives safely

then both generals know that both generals know that both generals know the plan

…

No finite sequence of messages will achieve common knowledge of the plan

Implicitly, the view is that if the generals could achieve common knowledge of the plan, then they would attack --- but that even high-order mutual knowledge of the plan does not suffice

Common Knowledge: A 'Discontinuity at Infinity'

In game theory, the sensitivity of behavior to high-order mutual knowledge vs. common knowledge was first observed by Geanakoplos and Polemarchakis (1982) in the setting of the Agreement Theorem (Aumann, 1976)

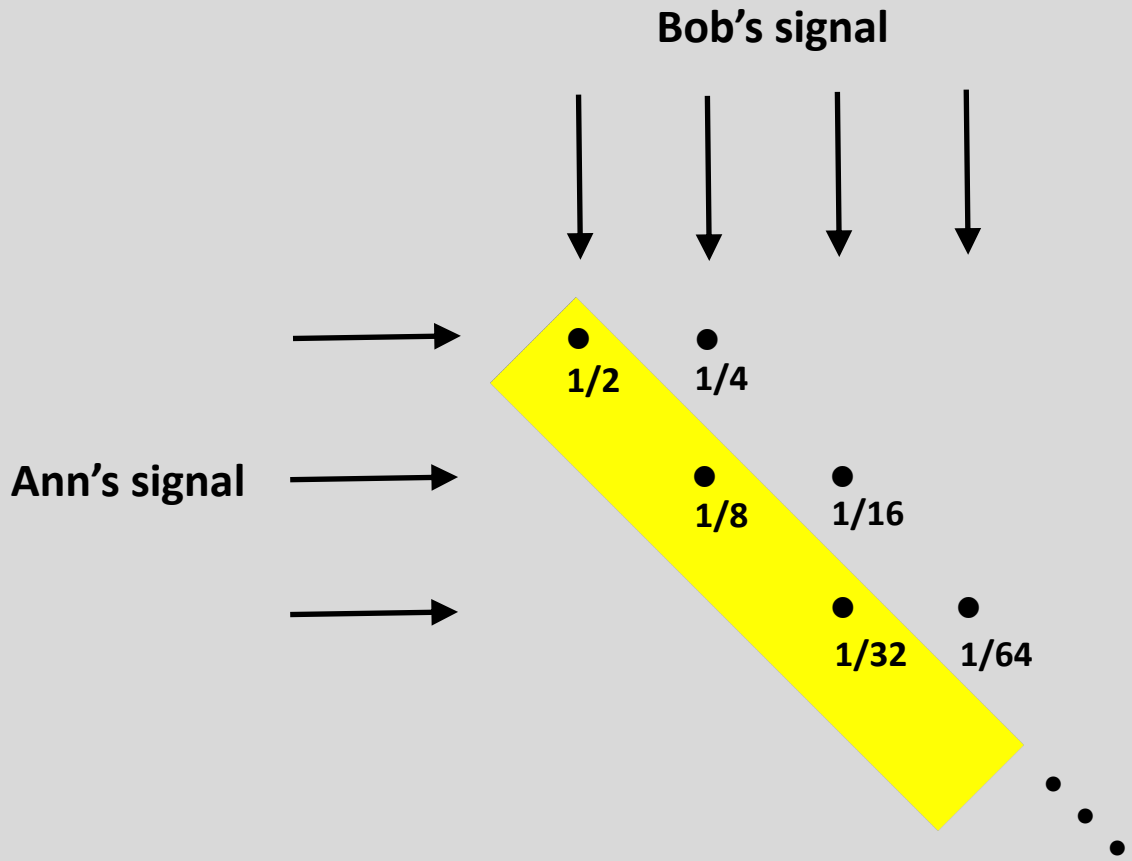

J. Geanakoplos and H. Polemarchakis, "We Can't Disagree Forever," 1982; R. Aumann, "Agreeing to Disagree," 1976; this variant is due to John Geanakoplos (private communication)

Common Knowledge contd.

Aumann and Brandenburger (1995) showed that common knowledge of the players' conjectures in a game (in the presence of other assumptions) yields Nash equilibrium, but high-order mutual knowledge does not

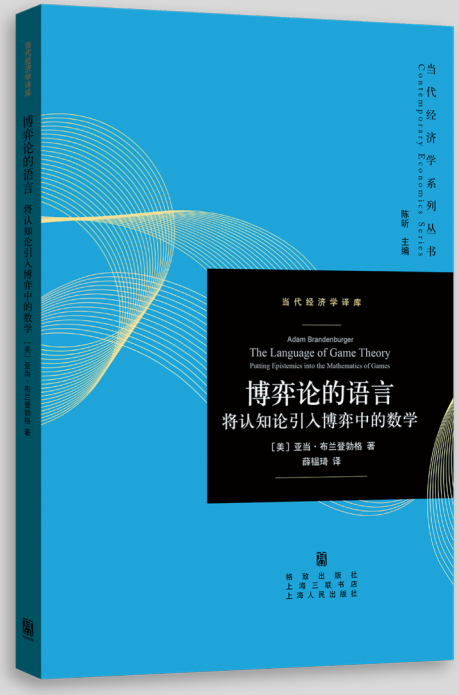

Rubinstein (1989) formalized the inadequacy of high-order mutual knowledge in a version of the Two Generals Problem (with uncertainty over the payoff functions)

R. Aumann and A. Brandenburger, "Epistemic Conditions for Nash Equilibrium," 1995;
A. Rubinstein, "The Electronic Mail Game: Strategic Behavior Under Almost Common Knowledge," 1989

Game-Theory Perspective: A Second Take

In this talk, we present another game-theory formulation of the Two Generals Problem

This version is inspired by the emergence of blockchain (specifically, proof-of-work)

We will look for an approximate solution rather than an impossibility argument

> "Every general, just by verifying the difficulty of the proof-of-work chain, can estimate how much parallel CPU power per hour was expended on it and see that it must have required the majority of the computers to produce that much proof-of-work in the allotted time."
>
> -- Satoshi Nakamoto

From: "How Satoshi Nakamoto explained the solution to the Byzantine Generals' Problem," at http://satoshi.nakamotoinstitute.org/emails/cryptography/11/

A Coordination Game with an Uncertain Number of Active Players

Players may be inactive (and always choose Ø) or active (and can choose *c* or Ø)

Coordination is positive only if there is a sufficient (expected) number of active players

	c	Ø
c	3α-1 3α-1 3α-1	2α-1 0 2α-1
Ø	0 2α-1 2α-1	0 0 α-1

c

	c	Ø
c	2α-1 2α-1 0	α-1 0 0
Ø	0 α-1 0	0 0 0

Ø

The idea is that action *c* will be chosen if and only if

α × expected number of active players ≥ 1

Adding a Computational Puzzle to the Game

A computational puzzle is distributed to each active player at time 0

Each active player has a machine that works on the puzzle and finds the solution with Poisson arrival rate λ (independent across machines)

If a machine solves the puzzle, there is a delay until time *T*, when the solution is transmitted to all players

If no machine solves the puzzle by time *T*, a null message is transmitted to all players

The puzzle can be solved only by guesswork but the solution can be immediately verified

Probability Calculations

Write the probability that *k* players are active, conditional on a solution by time *T*, as

$$\phi(k; T) = \frac{p_k \left[1 - \exp\left(-\lambda k T\right)\right]}{\sum_{i=1}^{n} p_i \left[1 - \exp\left(-\lambda i T\right)\right]}$$

We are interested in cases where

$$\alpha \sum_{k=1}^{n} k \cdot p_k < 1$$

but there is a (finite) *T* such that

$$\alpha \sum_{k=1}^{n} k \cdot \phi(k; T) \geq 1$$

The idea is that we can choose a time *T* so that, if a solution is found by *T*, then there is a good chance that a good number of players are active

Calculations with Three Players

of players n = 3; arrival rate λ = 0.1; uniform prior on # of active players

Probability that k players are active, conditional on a solution by T:

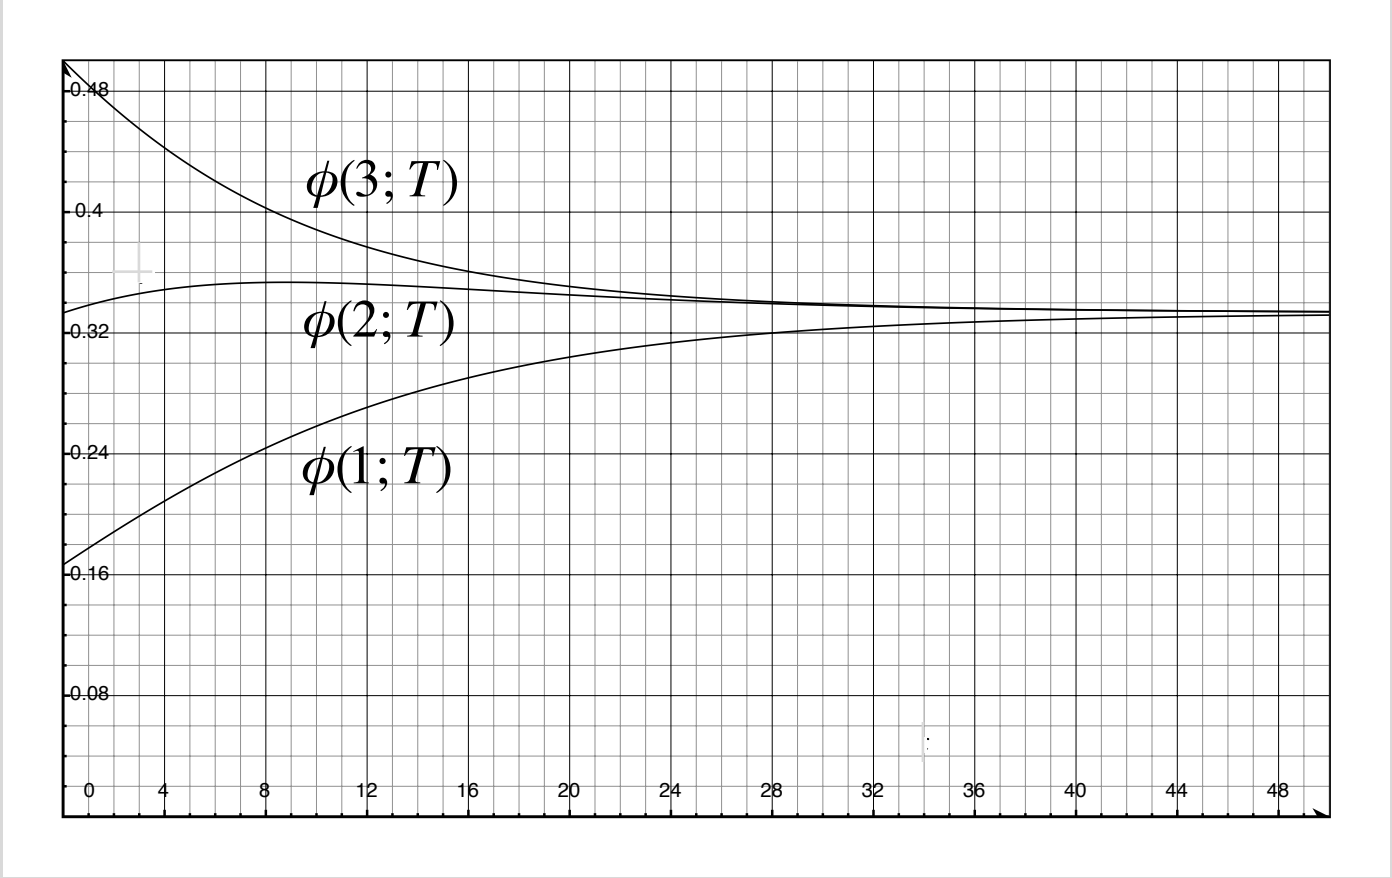

Calculations with Three Players contd.

Expected number of active players, conditional on a solution by *T*:

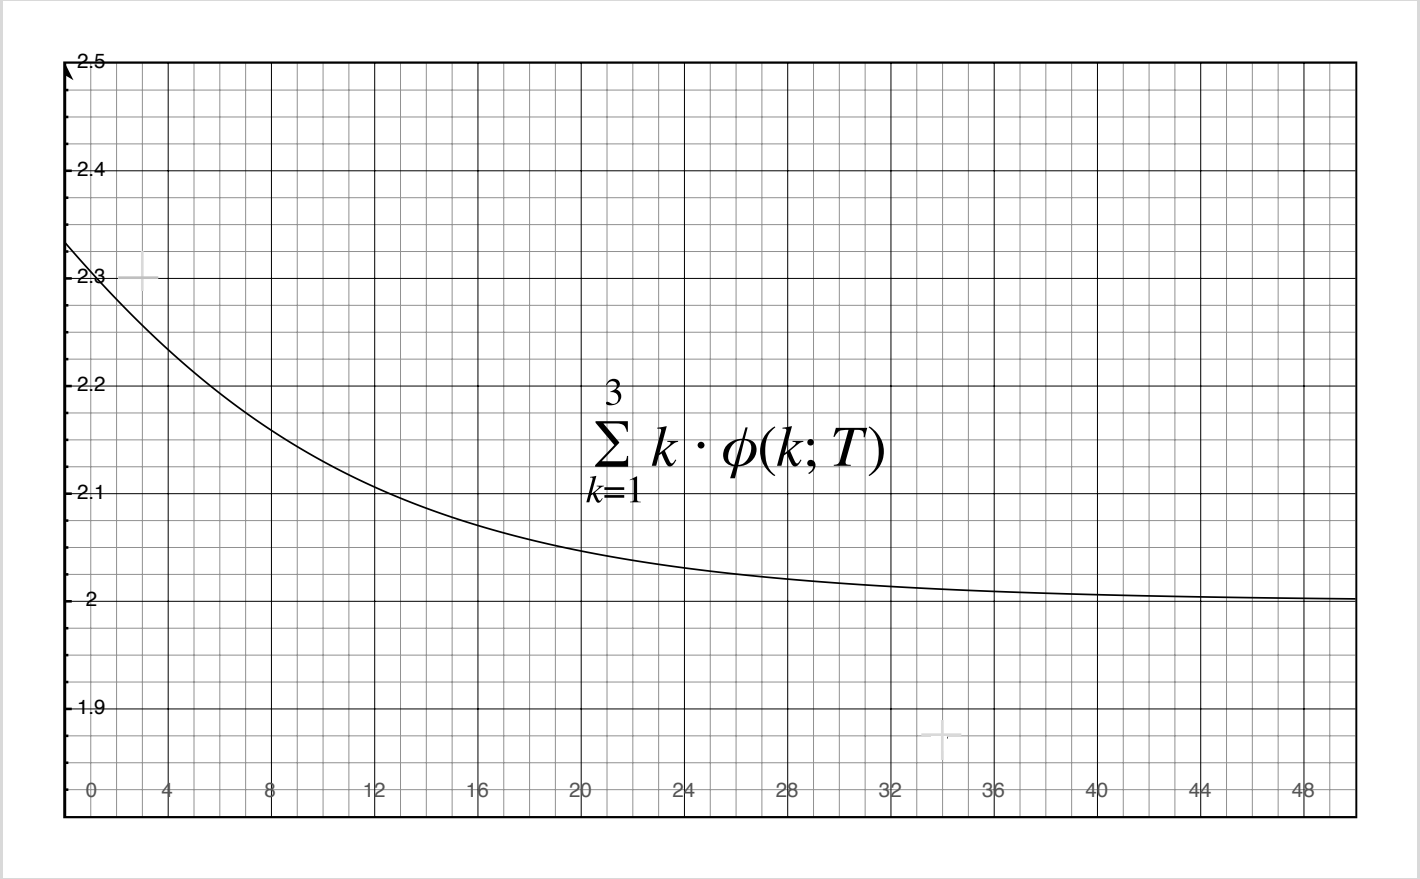

Proposition

If

$$\frac{\sum_{k=1}^{n} k\, p_k}{\sum_{k=1}^{n} k^2\, p_k} < \alpha < \frac{1}{\sum_{k=1}^{n} k\, p_k}$$

then coordination does not happen without the computational puzzle but can happen, for sufficiently small *T*, with the computational puzzle

To do:

- Design a protocol that ensures a player can be active if and only if that player works on the puzzle
- Design a mechanism that balances the benefits and costs of different choices of time duration *T* (and λ)

Back to Common Knowledge

At time *T*, if the computational puzzle is solved, then each (active) player assigns some --- possibly, large --- probability to the event that the puzzle was distributed to a large number of other players

Suppose the puzzle came attached to an underlying statement *S* of interest

Then, depending on how long *T* is, we may be able to say that each player 'approximately knows' that a large number of players know *S*

We can iterate this process by next distributing a message consisting of *S*, the original puzzle, the solution, and a new puzzle

If, at a time *T**, the new puzzle is solved, we may be able to say that each player approximately knows that a large number of players approximately know that a large number of players know *S*

…

Such higher-order (approximate) knowledge did not play a formal role in the coordination game, but could be important in other applications

S. DeDeo, "The Bitcoin Paradox: Why Cryptocurrency Will Always Be Political," 2017, at http://nautil.us/issue/55/trust/the-bitcoin-paradox

Image from https://medium.com/coinmonks/a-note-from-anthony-if-you-havent-already-please-read-the-article-gaining-clarity-on-key-787989107969