

# What Does “Distributed Consensus” Mean?

Adam Brandenburger

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

From: Satoshi Nakamoto  
Subject: **Bitcoin P2P e-cash paper**  
Date: November 13, 2008 at 22:56:55 UTC

---

James A. Donald wrote:

> It is not sufficient that everyone knows X. We also  
> need everyone to know that everyone knows X, and that  
> everyone knows that everyone knows that everyone knows X  
> - which, as in the Byzantine Generals problem, is the  
> classic hard problem of distributed data processing.

... 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. They had to all have seen it because the proof-of-work is proof that they worked on it. If the CPU power exhibited by the proof-of-work chain is sufficient to crack the password, they can safely attack at the agreed time.

# (At Least) Two Notions of Distributed Consensus

1. Common Knowledge

2. Nash Equilibrium

The question is how to achieve each of these states when agents are in a distributed environment

The classic stories that are told assume that all agents are “in the same room”

Thus, it is often claimed that a proposition is common knowledge if it is stated in the presence of everyone

And, it is often claimed that a Nash-equilibrium strategy profile can be achieved if it is proposed in the presence of everyone (and satisfies the requisite inequalities)

# Two Notions of Distributed Consensus ... Partially Achieved?

## 1. Common Knowledge


We will present a protocol that achieves common knowledge with a limited flow of information among agents

## 2. Nash Equilibrium

We will present a protocol that achieves Nash equilibrium with a limited flow of information among agents

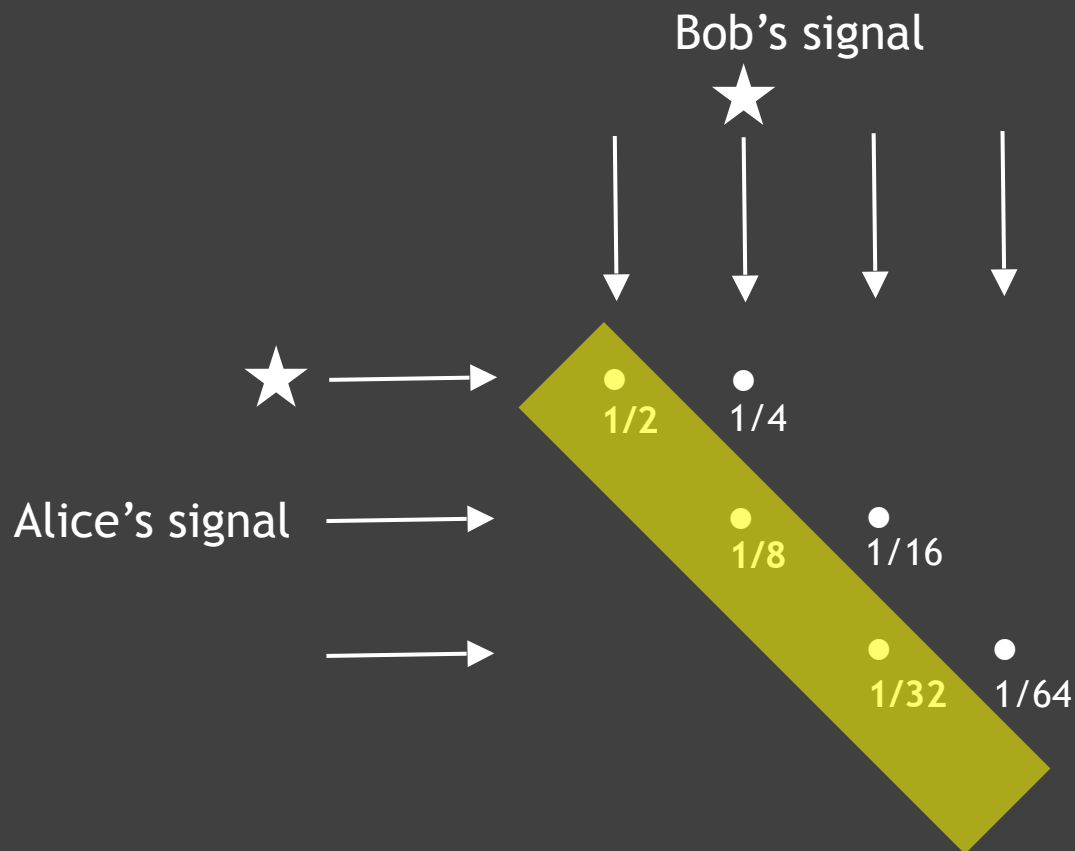
So, we understand a distributed environment as asking for limits on information flow ... this may well be an unorthodox interpretation, but, perhaps, insightful!

# Part 1: Common Knowledge via Communication



Each period Alice and Bob announce their probabilities of the highlighted set

# Common Knowledge via Communication contd.



Start: Alice says probability  $2/3$  and Bob says probability  $1/3$

Next: Alice rules out Scenario 1 and says probability 0

Next: Bob infers Alice's signal and says probability 0

⇒ Agreement and common knowledge of agreement!

# Common Knowledge Formalized

Fix a finite state space  $\Omega$  and a common prior probability measure  $p$  on  $\Omega$

Each agent  $i$  has a partition  $\mathcal{P}_i$  of  $\Omega$  representing their private information

Let  $\mathcal{P}_i(\omega)$  denote the member of  $i$ 's partition that contains state  $\omega$

Definition: Agent  $i$  knows an event  $E$  at  $\omega$  if  $\mathcal{P}_i(\omega) \subseteq E$

Let  $(\bigwedge_i \mathcal{P}_i)(\omega)$  denote the member of the meet that contains  $\omega$

*Theorem (Aumann, 1976): An event  $E$  is common knowledge at  $\omega$  if and only if  $(\bigwedge_i \mathcal{P}_i)(\omega) \subseteq E$*

*Theorem (Aumann, 1976): If the agents' conditional probabilities of an event  $F$  are common knowledge at a state  $\omega$ , then they are equal*

## Common Knowledge with Limited Information Flow

We now suppose that it is not the individual conditional probabilities of  $F$  that are common knowledge, but instead a generalized Kolmogorov-Nagumo average

We also generalize to conditional expectations of a random variable  $X$

So, we consider a “summary statistic”

$$f^{-1} \left[ \frac{\sum_i f_i(e_i)}{n} \right]$$

where  $e_i$  is agent  $i$ 's expectation, and  $f$  and each  $f_i$  are strictly increasing

Theorem (Nielsen, Brandenburger, Geanakoplos, McKelvey, and Page, 1990):  
*If a summary statistic of the agents' conditional expectations of a random variable  $X$  is common knowledge at a state  $\omega$ , then all the conditional expectations are equal*

## Part 2: Nash Equilibrium in a Coordination Game

|             | $c$   | $\emptyset$                           |
|-------------|---|---------------------------------------|
| $c$         | $3\alpha - 1$<br>$3\alpha - 1$<br>$3\alpha - 1$ | $2\alpha - 1$<br>$0$<br>$2\alpha - 1$ |
| $\emptyset$ | $0$<br>$2\alpha - 1$<br>$2\alpha - 1$           | $0$<br>$0$<br>$\alpha - 1$            |

|             | $c$                                   | $\emptyset$                |
|-------------|---------------------------------------|----------------------------|
| $c$         | $2\alpha - 1$<br>$2\alpha - 1$<br>$0$ | $\alpha - 1$<br>$0$<br>$0$ |
| $\emptyset$ | $0$<br>$\alpha - 1$<br>$0$            | $0$<br>$0$<br>$0$          |

Note that a player optimally chooses  $c$  if

$$\alpha \times \text{expected number of players who choose } c \geq 1$$

## Addition of a Computational Puzzle

At time 0, a computational puzzle is posted to an anonymous message board and read by all players present

Each player who is present has a machine that works on the puzzle by randomly guessing, with a Poisson arrival rate  $\lambda$  of finding the solution (independent across machines)

There is a pre-specified time  $T$  such that, if one or more machines find the solution by  $T$ , the solution is posted to the message board at  $T$  — otherwise a null message is posted

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

# Probability Calculations

The probability that  $k$  out of a possible total of  $n$  players are present, conditional on a solution by time  $T$ , is given by

$$\phi(k; T) = \frac{p_k(1 - e^{-\lambda k T})}{\sum_{i=1}^n p_i(1 - e^{-\lambda i T})}$$

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 the expected number of players who are present is sufficiently high

# Three-Player Example



Uniform prior  $p_k$  on number of players present; arrival rate  $\lambda = 0.1$

## Nash Equilibrium with Limited Information Flow

Theorem (Brandenburger and Steverson, 2020): *The expected number of players present, conditional on a solution by time  $T$*

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

is strictly decreasing in  $T$

The proof works by establishing that  $\phi(k; T)$  satisfies the decreasing monotone likelihood ratio principle

Using l'Hôpital's rule, it follows that 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, for a sufficiently small  $T$ , a player might optimally choose  $c$  if a solution is announced at time  $T$ , but not do so in the absence of the computational puzzle

## Some Remaining Questions

Q: Can a distributed ledger be used to achieve — exact — common knowledge among distributed agents?

For answers that relax common knowledge, see Halpern and Pass (2017) and Gonczarowski and Moses (2024)

Q: Can the “proof-of-presence” mechanism be properly operationalized to achieve Nash equilibrium among distributed agents?

Solving the computational puzzle would likely be costly and therefore a reward to posting (first) may be required

The optimal value of the cutoff  $T$  is a problem in mechanism design, perhaps with some idiosyncratic features

... and surely other questions about the notion of “distributed consensus”

Thank You