Agreement and Disagreement in a Non-Classical World

Adam Brandenburger, Patricia Contreras-Tejada, Pierfrancesco La Mura, Giannicola Scarpa, and Kai Steverson

New York University; Instituto de Ciencias Matemáticas, Madrid; Universidad Complutense de Madrid; HHL Leipzig Graduate School of Management; Universidad Politécnica de Madrid; DCI Solutions

Version 07/21/23

“We are all agreed that your theory is crazy. The question that divides us is whether it is crazy enough to have a chance of being correct.”

— Niels Bohr

The Classical Agreement Theorem

Alice and Bob possess a common prior probability distribution on a state space

They each then receive different private information about the true state

They form their conditional (posterior) probabilities q_A and q_B of an underlying event of interest

Theorem (Aumann, 1976): *If these two values* q_A *and* q_B *are common knowledge between Alice and Bob, they must be equal*

Here, an event E is **common knowledge** between Alice and Bob if they both know it, both know they both know it, and so on indefinitely

R. Aumann, "Agreeing to Disagree," *Annals of Statistics*, 4, 1976, 1236-1239

Applications

The agreement theorem is considered a basic requirement in classical epistemics

It has been used to

- show that two risk-neutral agents, starting from a common prior, cannot agree to bet with each other
  (Sebenius and Geanakoplos, 1983)

- prove "no-trade" theorems for efficient markets
  (Milgrom and Stokey, 1982)

- establish epistemic conditions for Nash equilibrium
  (Aumann and Brandenburger, 1995)

J. Sebenius and J. Geanakoplos. "Don't Bet on It: Contingent Agreements with Asymmetric Information, *Journal of the American Statistical Association*, 78, 1983, 424-426; P. Milgrom and N. Stokey, "Information, Trade, and Common Knowledge," *Journal of Economic Theory*, 26, 1982, 17–27; R. Aumann and A. Brandenburger, "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63, 1995, 1161–1180

The Role of Common Knowledge: A "Discontinuity" at Infinity

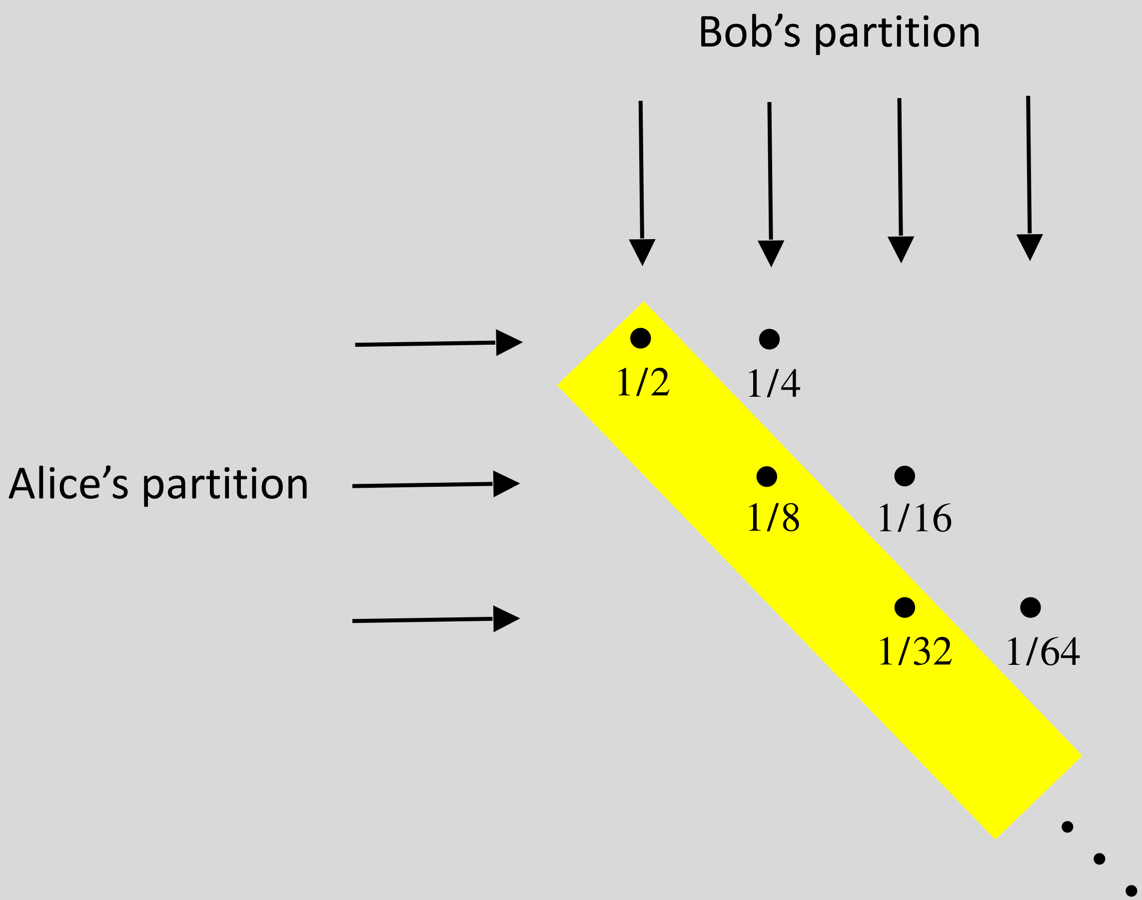

J. Geanakoplos and H. Polemarchakis, "We Can't Disagree Forever," *Journal of Economic Theory*, 28, 1982, 192–200; this variant is due to John Geanakoplos (private communication)

Non-Classical Settings I

What is the status of the Agreement Theorem when classical probability theory does not apply?

In the physical domain, the canonical case is quantum mechanics, where a fundamental result (Bell's Theorem, 1964) says that no "local hidden-variable" theory can model the results of all quantum experiments

As we will see, this implies that the classical Bayesian model does not apply

In probability theory, there is a finite analog to the de Finetti representation theorem for infinite sequences of exchangeable random variables, if mixing is via a signed probability measure (Jaynes, 1986; Kearns and Székely, 2006; Janson, Konstantopoulos, and Yuan, 2016)

This permits an exchangeability derivation of Fermi-Dirac statistics, paralleling an infinite exchangeability derivation of Bose-Einstein statistics (Kearns and Székely, 2006; Bach, Blank, and Francke, 1985)

J. Bell, "On the Einstein Podolsky Rosen Paradox," *Physics*, 1, 1964, 195-200; A. Bach, H. Blank, and H. Francke, "Bose-Einstein Statistics Derived from the Statistics of Classical Particles," *Lettere al Nuovo Cimento*, 43, 1985, 195-198; E. Jaynes, "Some Applications and Extensions of the de Finetti Representation Theorem," in P. Goal and A. Zellner (eds.), *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti,* North-Holland, 1986, 31; G. Kearns and G. Székely, "De Finetti's Theorem for Abstract Finite Exchangeable Sequences," *Journal of Theoretical Probability*, 19, 2006, 589-608; S. Janson, T. Konstantopoulos, and L. Yuan, "On a Representation Theorem for Finitely Exchangeable Random Vectors," *Journal of Mathematical Analysis and Applications*, 442, 2016, 703-714

Non-Classical Settings II

In decision theory, Perea (2022) axiomatizes expected utility theory for conditional preference relations, which assign to every possible probabilistic belief on a set of states, a preference relation over the decision maker's set of actions

The motivation is that, in a game, we typically fix a player's utility function but not beliefs — what, then, does the utility function represent?

The interest is in axiomatizing expected utility for unsigned probability measures, but, for the axioms to bite, one must allow the decision maker to holds signed probability measures on the states

For decision making under ambiguity, Ke and Zhao (2022) include an axiom system on preferences that is representable by maxmin expected utility (as in Gilboa and Schmeidler, 1989) where the minimization is over sets of signed probability measures

A. Perea, "Expected Utility as an Expression of Linear Preference Intensity," working paper, Maastricht University, May 2022; S. Ke and C. Zhao, "From Local Utility to Neural Networks," December 2022, at https://sites.google.com/site/shaoweike/research; I. Gilboa and D. Schmeidler, "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 1989, 141-153

Quantum Theory: 2 x 2 x 2 Boxes

Empirical model:

	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
(a,b)	f_1	f_2	f_3	f_4
(a',b)	f_5	f_6	f_7	f_8
(a,b')	f_9	f_{10}	f_{11}	f_{12}
(a',b')	f_{13}	f_{14}	f_{15}	f_{16}

Bell model:

	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
(a,b)	1/2	0	0	1/2
(a',b)	3/8	1/8	1/8	3/8
(a,b')	3/8	1/8	1/8	3/8
(a',b')	1/8	3/8	3/8	1/8

Phase-Space Representation

	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
(a,b)	f_1	f_2	f_3	f_4
(a',b)	f_5	f_6	f_7	f_8
(a,b')	f_9	f_{10}	f_{11}	f_{12}
(a',b')	f_{13}	f_{14}	f_{15}	f_{16}

	a	a'	b	b'
p_0	0	0	0	0
p_1	0	0	0	1
p_2	0	0	1	0
p_3	0	0	1	1
p_4	0	1	0	0
p_5	0	1	0	1
p_6	0	1	1	0
p_7	0	1	1	1
p_8	1	0	0	0
p_9	1	0	0	1
p_{10}	1	0	1	0
p_{11}	1	0	1	1
p_{12}	1	1	0	0
p_{13}	1	1	0	1
p_{14}	1	1	1	0
p_{15}	1	1	1	1

An Impossibility Result

From phase space and the Bell empirical model, we can calculate

$$
\begin{aligned}
p_0 + p_1 + p_4 + p_5 &= 1/2 \\
p_4 + p_5 + p_{12} + p_{13} &= 1/8 \\
p_1 + p_3 + p_5 + p_7 &= 1/8 \\
p_0 + p_2 + p_8 + p_{10} &= 1/8
\end{aligned}
$$

Adding the second, third, and fourth equations gives

$$
p_0 + p_1 + p_2 + p_3 + p_4 + 2p_5 + p_7 + p_8 + p_{10} + p_{12} + p_{13} = 3/8
$$

which contradicts the first equation

Theorem (Abramsky and Brandenburger, 2011): An empirical model is “no signaling” if and only if there is a phase-space model with a signed probability measure that induces it

S. Abramsky and A. Brandenburger, “The Sheaf-Theoretic Structure of Non-Locality and Contextuality,” *New Journal of Physics*, 13, 2011, 113036

General Set-up

There is a finite abstract state space Ω

Alice and Bob have partitions $\mathscr{P}_A$ and $\mathscr{P}_B$ of Ω representing their private information

There is a common – possibly signed – prior probability measure p on Ω

Observability:

Assume throughout that all members of the partitions $\mathscr{P}_A$ and $\mathscr{P}_B$ receive probability in the interval $(0,1]$

Assume, too, that all events of interest receive probability in $(0,1]$

A Warm-Up Example

Alice's (Bob's) partition is red (blue)

The event of interest is

$$E = \{\omega_1, \omega_3, \omega_4\}$$

The true state is ω_1

At ω_1, Alice assigns (conditional) probability 1 to E

At ω_1, Bob assigns (conditional) probability 0 to E

The event that Bob assigns probability 0 to E is

$$F = \{\omega_1, \omega_2, \omega_3\}$$

At ω_1, Alice assigns probability 1 to F

Call this **singular disagreement**

It is impossible classically!

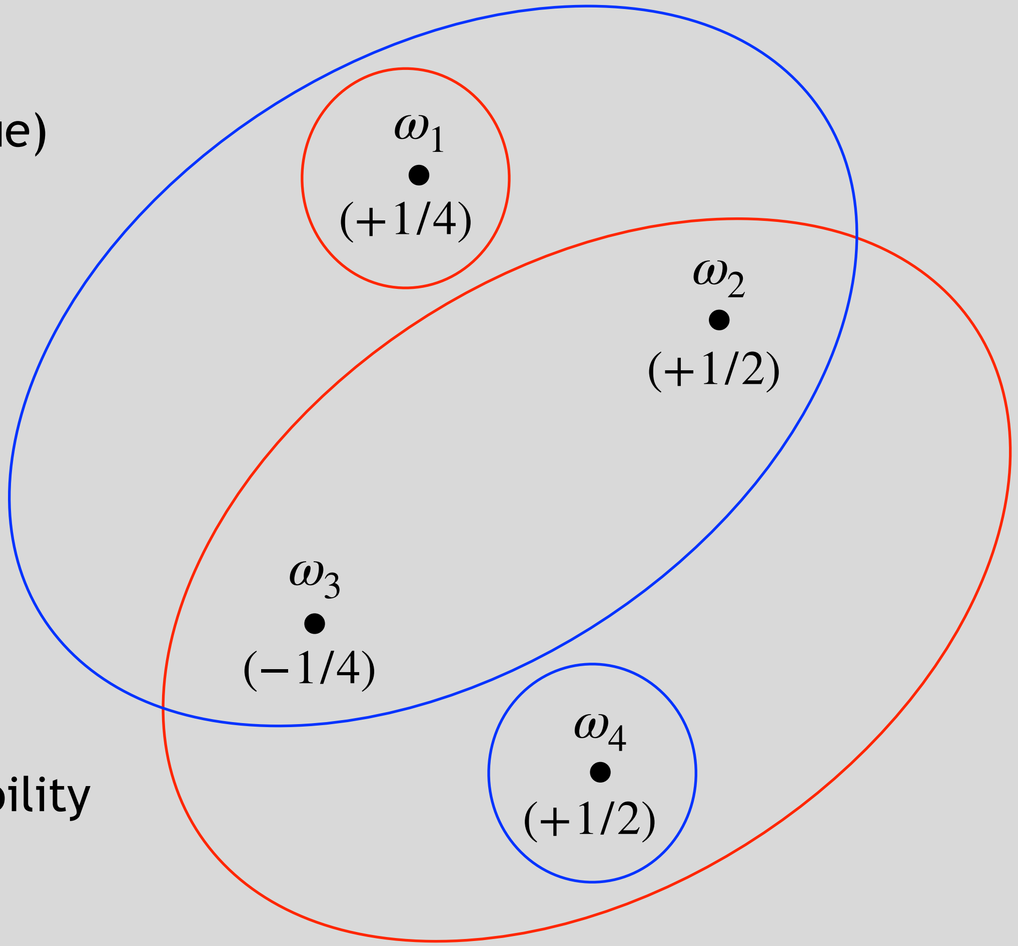

Note: All partition cells and E receive probability in $(0,1]$ and are therefore observable

From Knowledge to Certainty

Definition: Alice **knows** an event E at state ω if $\mathscr{P}_A(\omega) \subseteq E$

Definition: Alice is **certain** of an event E at a state ω if $p(E \,|\, \mathscr{P}_A(\omega)) = 1$

Fix an event E and probabilities q_A and q_B, and let

$$A_0 = \left\{\omega \in \Omega : p(E \mid \mathscr{P}_A(\omega)) = q_A\right\}$$

$$B_0 = \left\{\omega \in \Omega : p(E \mid \mathscr{P}_B(\omega)) = q_B\right\}$$

$$A_{n+1} = A_n \cap \left\{\omega \in \Omega : p(B_n \mid \mathscr{P}_A(\omega)) = 1\right\}$$

$$B_{n+1} = B_n \cap \left\{\omega \in \Omega : p(A_n \mid \mathscr{P}_B(\omega)) = 1\right\}$$

for all $n \geq 0$

Definition: It is **common certainty** at a state ω^* that Alice assigns probability q_A to E and Bob assigns probability q_B to E if $\omega^* \in \cap_{n=0}^{\infty} A_n \cap \cap_{n=0}^{\infty} B_n$

Relationship Between Knowledge and Certainty

If Alice knows an event E at state ω, then she is certain of E at ω

It is also true that common knowledge of E implies common certainty of E

(Proof: If Alice knows Bob knows E, then she knows Bob is certain of E, since knowledge is monotonic. From this, Alice is certain Bob is certain of E. The argument continues to higher levels.)

Arguably, the distinction between these modalities is "small" in the classical domain (arguably, not!)

Also, in the classical domain, there is an Agreement Theorem for common certainty

Theorem (classical): Fix a (non-negative) common prior and an event E. Suppose at a state ω^* it is common certainty that Alice's probability of E is q_A and Bob's probability of E is q_B. Then $q_A = q_B$.

But what happens in the non-classical world?

Non-Classical Agreement with Knowledge

Even without our observability conditions, we get a non-classical analog to the classical Agreement Theorem

Theorem (non-classical): *Fix a signed common prior and an event* E. *Suppose at a state* ω^* *it is common knowledge that Alice's probability of* E *is* q_A *and Bob's probability of* E *is* q_B. *Then* $q_A = q_B$.

Proof: Follow closely the classical argument. Consider the (equal) conditional probabilities q_A for Alice, calculated for members of her partition that are contained in the member of the meet $(\mathscr{P}_A \wedge \mathscr{P}_B)(\omega^*)$. This time, we take an affine rather than convex combination of this constant probability to get $p(E\,|\,(\mathscr{P}_A \wedge \mathscr{P}_B)(\omega^*)) = q_A$. Then run the same argument for Bob.

But let's see what happens with common certainty …

https://www.xixilogic.org/events/wp-content/uploads/2022/05/lfdsn2019-proceedings.pdf; see also M. Leifer and C. Duarte, "Generalizing Aumann's Agreement Theorem," 2022, at https://arxiv.org/abs/2202.02156

Common Certainty of Disagreement

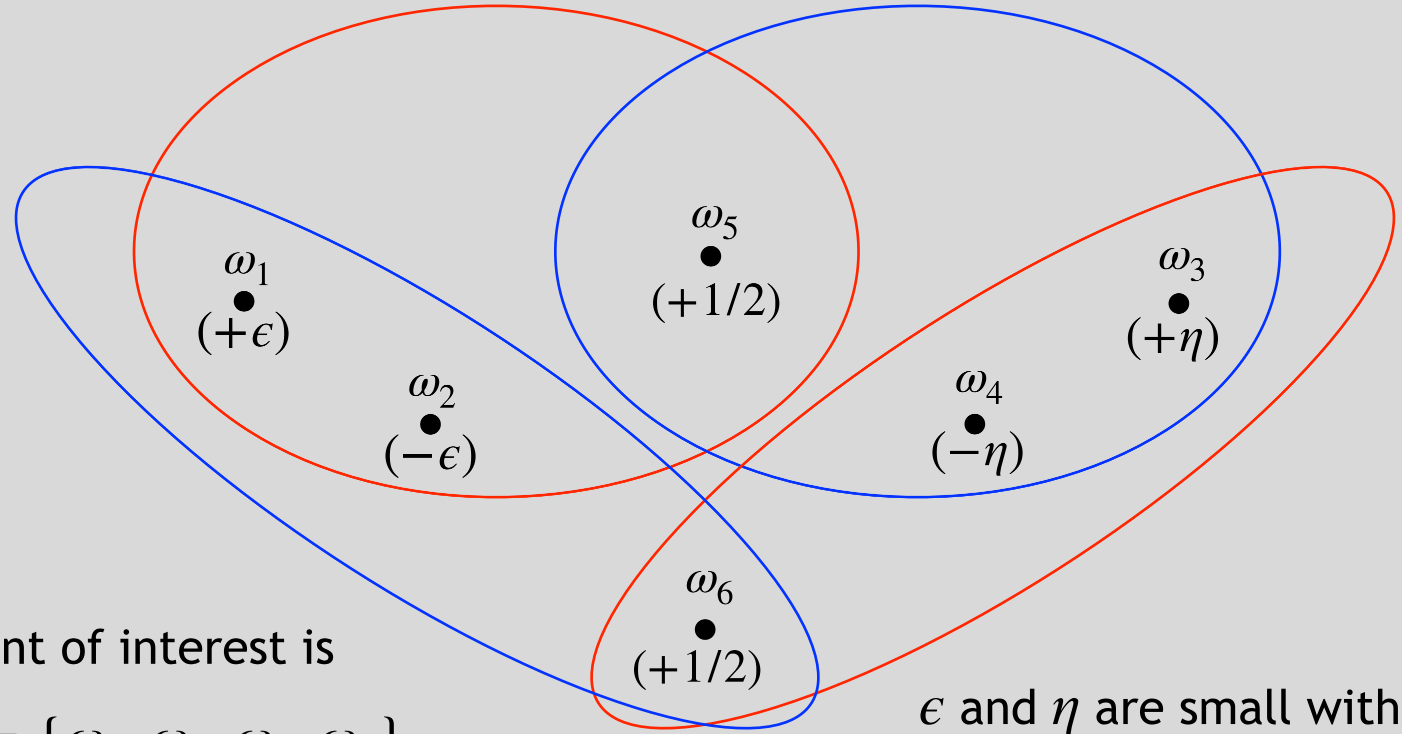

The event of interest is

$$E = \{\omega_2, \omega_4, \omega_5, \omega_6\}$$

ϵ and η are small with $\epsilon \neq \eta$

The true state is ω_5

At ω_5, it is common certainty that Alice assigns probability $1 - 2\epsilon$ to E while Bob assigns probability $1 - 2\eta$ to E

That is, there is common certainty of disagreement!

Communication I

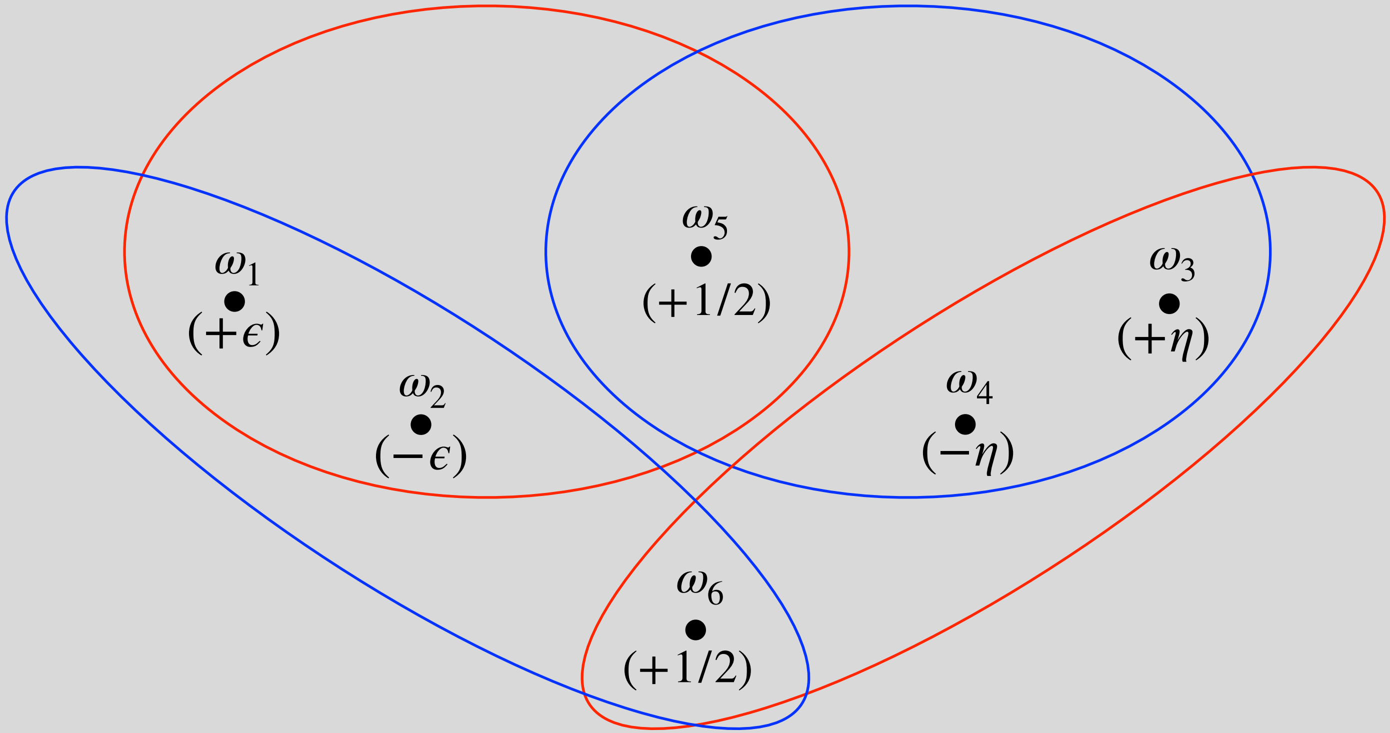

Suppose the true state is ω_1

Alice announces a probability of $1-2\epsilon$ (à la Geanakoplos and Polemarchakis, 1982)

Bob infers that Alice observed $\{\omega_1, \omega_2, \omega_5\}$ and calculates an updated probability of $-\epsilon/0$!

Communication II

In systems where the agents are able to communicate about an event of interest, those communications should lead to well-defined and classical conditional probabilities regarding that event

That is, the resulting conditional probabilities should all lie in the interval $[0,1]$

Communication — even if it concerns a non-classical system — should be considered observable and therefore classical

We next introduce conditions that ensure this is the case (but we go in a different direction from the Geanakoplos-Polemarchakis protocol)

Communication-Enabled Structures

Define a sequence of partitions for Alice, corresponding to announcements she could make about her probability of E, her certainty of Bob's probability, etc., and likewise for Bob

$$\mathscr{M}_A^{(n)} = \{A_n, A_n^c\}$$

$$\mathscr{M}_B^{(n)} = \{B_n, B_n^c\}$$

For any $\pi, E \subseteq \Omega$, say π is **regular with respect to** E if $p(\pi) \geq 0$ and $0 \leq p(\pi \cap E) \leq p(\pi)$

A structure $(\Omega, p, \mathscr{P}_A, \mathscr{P}_B)$ is **communication-enabled with respect to** E if for each $n \geq 0$, each $\pi \in \mathscr{P}_A \vee \mathscr{M}_B^{(n)}$ and each $\pi \in \mathscr{P}_B \vee \mathscr{M}_A^{(n)}$ is regular with respect to E

Note: This property fails in the previous example

A New Agreement Theorem

Theorem: *Fix a structure that is communication-enabled with respect to* E *and suppose at a state* ω^* *it is common certainty that Alice's probability of* E *is* q_A *and Bob's probability of* E *is* q_B. *Then* $q_A = q_B$.

Notice that Alice's potential announcements are made relative to her (initial) partition $\mathscr{P}_A$; and likewise for Bob

In words, the mere ability to "confirm" the epistemic state (here, the state is common certainty of the posteriors) is enough to rule out disagreement — the confirmation need not actually be carried out

Realizability of Common Certainty of Disagreement?

In the physical domain, it can be shown that common certainty of disagreement (CCD) is impossible when observing quantum systems but possible for "superquantum" (no-signaling) systems

The impossibility of CCD can therefore be proposed as a physical axiom

In decision theory, if we equip agents with signed probability measures, it seems we can get highly non-classical behavior, such as betting between risk-neutral agents

Or, should the impossibility of CCD be elevated to an (epistemic) decision-theoretic principle?

If yes, what non-classical behavior is then allowed? This appears to be an open direction …

P. Contreras-Tejada, G. Scarpa, A. Kubicki, A. Brandenburger, and P. La Mura, "Observers of Quantum Systems Cannot Agree to Disagree," *Nature Communications*, 12, 2021, at https://doi.org/10.1038/ s41467-021-27134-6

Two Alternative Models

1. Khrennikov and Basieva (2014) and Khrennikov (2015) consider quantum-like observers of a quantum system who employ either the knowledge or certainty modality

   This approach allows CCD even for quantum systems

2. (With thanks to Miklós Pintér) We could strengthen the belief modality to say:

   Alice is **fully certain** of E if all events in the complement of E receive probability 0

   We could investigate this avenue by developing a preference-based definition of certainty (analogous to defining Savage-null events) from a decision theory with signed probabilities

   This appears to be an interesting open direction

A. Khrennikov and I. Basieva, "Possibility to Agree on Disagree from Quantum Information and Decision Making," *Journal of Mathematical Psychology*, 62, 2014, 1-5; A. Khrennikov, "Quantum Version of Aumann's Approach to Common Knowledge: Sufficient Conditions of Impossibility to Agree on Disagree," *Journal of Mathematical Economics*, 60, 2015, 89-104

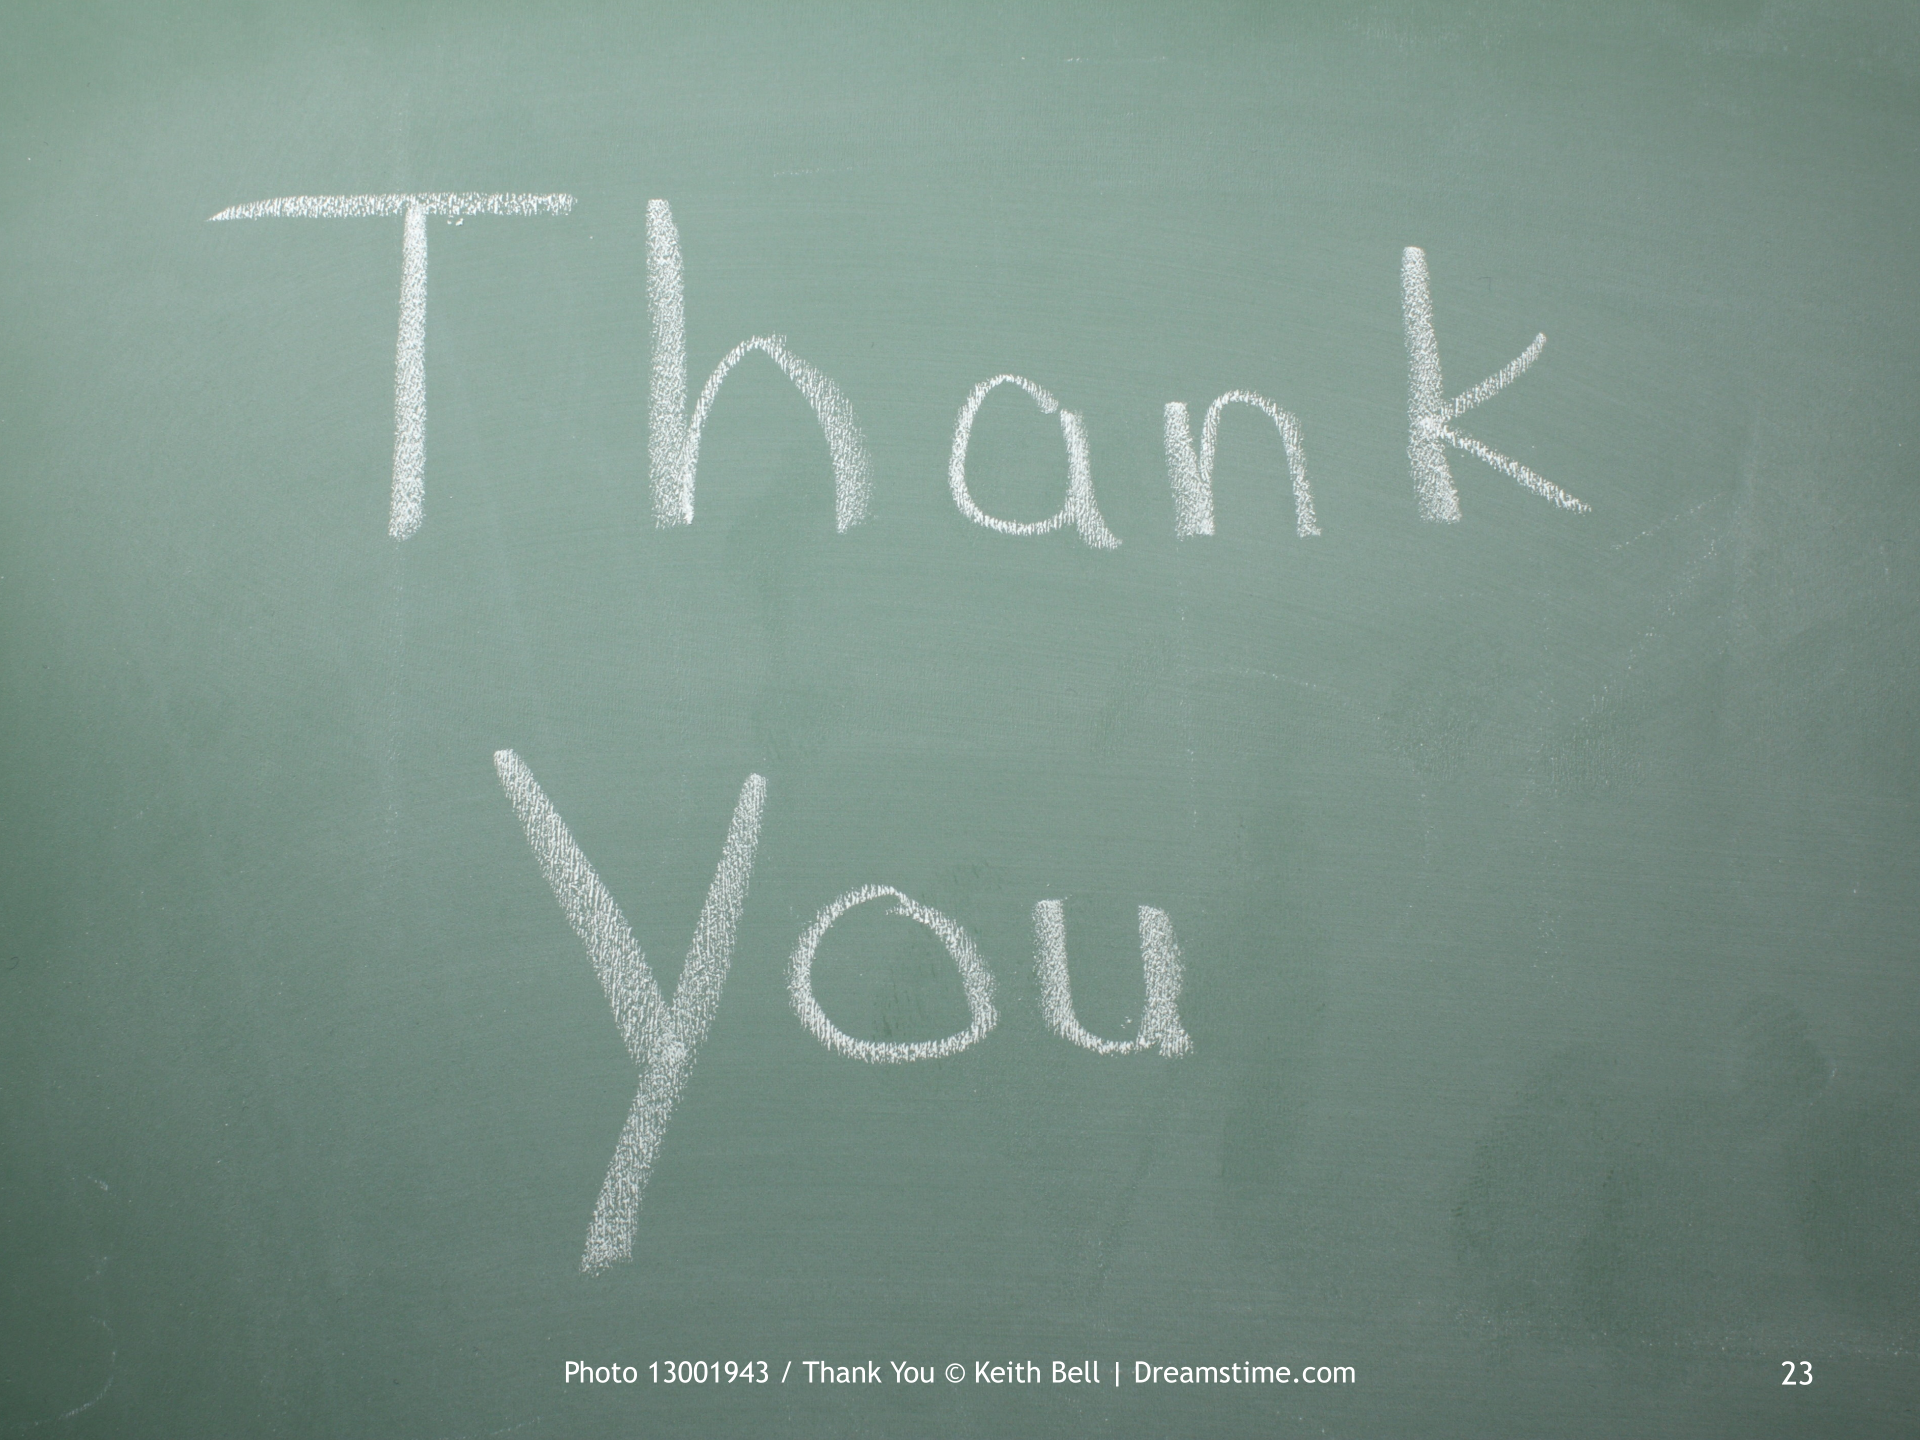

Photo 13001943 / Thank You © Keith Bell | Dreamstime.com