



Article submitted to journal

Subject Areas:

Bayesian agreement, Signed probability theory

Keywords:

Bayesian agents, Common certainty, Signed probabilities, Quasi-probabilities, Agreement and disagreement, Communication

Author for correspondence:

Adam Brandenburger

e-mail:

adam.brandenburger@stern.nyu.eduAgreement and Disagreement
in a Non-Classical WorldAdam Brandenburger¹, Patricia
Contreras-Tejada², Pierfrancesco La
Mura³, Giannicola Scarpa⁴, and Kai
Steverson⁵¹Stern School of Business, Tandon School of
Engineering, NYU Shanghai, New York University, New
York, NY 10012, U.S.A.,adam.brandenburger@stern.nyu.edu²Instituto de Ciencias Matemáticas, 28049 Madrid,
Spain, p.con.tejada@gmail.com³HHL Leipzig Graduate School of Management, 04109
Leipzig, Germany, plamura@hhl.de⁴Escuela Técnica Superior de Ingeniería de Sistemas
Informáticos, Universidad Politécnica de Madrid,
28031, Madrid, Spain, g.scarpa@upm.es⁵DCI Solutions, Aberdeen, MD 21005, U.S.A.,
ksteverson@dcι-solutions.com

The Agreement Theorem (Aumann, 1976) states that if two Bayesian agents start with a common prior, then they cannot have common knowledge that they hold different posterior probabilities of some underlying event of interest. In short, the two agents cannot “agree to disagree.” This result applies in the classical domain where classical probability theory applies. But in non-classical domains, such as the quantum world, classical probability theory does not apply. Inspired principally by their use in quantum mechanics, we employ signed probabilities to investigate the epistemics of the non-classical world. We find that here, too, it cannot be common knowledge that two agents assign different probabilities to an event of interest. However, in a non-classical domain, unlike the classical case, it can be common certainty that two agents assign different probabilities to an event of interest. Finally, in a non-classical domain, it cannot be common certainty that two agents assign different probabilities, if communication of their common certainty is possible – even if communication does not take place.

© The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, provided the original author and source are credited.

1. Introduction

In the domain of classical probability theory, Aumann (1976) proved the fundamental result that Bayesian agents cannot agree to disagree. Two agents Alice and Bob begin with a common prior probability distribution on a state space. Next, they each receive different private information about the true state and form their conditional (posterior) probabilities q_A and q_B of an underlying event of interest. Then, if these two values q_A and q_B are common knowledge between Alice and Bob, they must be equal: $q_A = q_B$. By “common knowledge” is meant that Alice knows Bob’s probability is q_B , Bob knows Alice’s probability is q_A , Alice knows Bob knows her probability is q_A , Bob knows Alice knows his probability is q_B , and so on indefinitely. This is the Agreement Theorem.

The role of common knowledge in this result is crucial. To conclude that $q_A = q_B$, it is not sufficient that Alice and Bob know each other’s probabilities. It is not even enough that they know these probabilities, and know they know them to some high finite order. Examples in which this condition allows $q_A \neq q_B$ are well known in the interactive epistemology literature (Geanakoplos and Polemarchakis, 1982; Aumann and Brandenburger, 1995). The condition of common knowledge is tight.

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), to prove “no-trade” theorems for efficient markets (Milgrom and Stokey, 1982), and to establish epistemic conditions for Nash equilibrium (Aumann and Brandenburger, 1995).

In this paper we ask 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 the fundamental Bell’s Theorem (Bell, 1964) says that no “local hidden-variable” theory can model the results of all quantum experiments. However, it can be shown (Abramsky and Brandenburger, 2011) that if signed probabilities on states are allowed, then there is a phase-space representation for all no-signaling models (Popescu and Rohrlich, 1994), which is a family of models that (strictly) includes those realizable within quantum mechanics. A phase-space model can be thought of as a canonical hidden-variable model, where the different states are precisely the different values the hidden variable can take. A signed probability measure is a measure that can assign negative values to certain events, while still assigning probability 1 to the whole space.

In fact, the use of signed probabilities in quantum mechanics goes back even earlier, to the Wigner “quasi-probability distribution” (Wigner, 1932), which is widely used in quantum mechanics – for example, in the field of quantum optics (Kenfack and Życzkowski, 2004). Dirac (1942) and Feynman (1987) also promoted the use of quasi- or signed probabilities in quantum mechanical calculations. Kaszlikowski and Kurzyński (2023) continue this tradition with their proposal for treating the “nebit,” which is a negative bit, as the basic unit of stochastic negativity.

Bell’s Theorem applies to a two-qubit system. Signed probabilities also arise in phase-space representations of a one-qubit system, under certain conditions. Brandenburger, La Mura, and Zoble (2022) derive the qubit from an entropic uncertainty principle stated on phase space. In their framework, there are quantum states whose only phase-space representations that respect the uncertainty bound involve negative probabilities. Onggadinata, Kurzyński, and Kaszlikowski (2023) derive the qubit with its full dynamics via an entropic invariance principle involving signed probabilities.

In probability theory, there is a finite analog to the de Finetti representation theorem for infinite sequences of exchangeable random variables (de Finetti, 1931), if mixing is via a signed probability measure (Dellacherie and Meyer, 1982, pp.46-47; Jaynes, 1986; Kearns and Székely, 2006; Janson, Konstantopoulos, and Yuan, 2016). Turning again to physics, this permits an exchangeability derivation of Fermi-Dirac statistics, paralleling an infinite exchangeability derivation of Bose-Einstein statistics (Bach, Blank, and Francke, 1985; Kearns and Székely, 2006).

Decision theory is another area in which signed probabilities have emerged. Perea (2022) axiomatizes expected utility theory for conditional preference relations, where such a relation assigns to every possible probability measure on a (finite) set of states that the decision maker might hold, a preference relation over the decision maker's (finite) set of actions. The motivation is game theory, where one typically specifies the game matrix, and hence the players' payoff or utility functions, but one thinks of a player as contemplating different beliefs (probability measures) they might hold concerning the actions chosen by other players. The question is when is such a conditional preference relation representable by a single utility function, with the expectation of utility taken under the given probability measure. Perea (2022) proposes a set of axioms that yields such a representation, but the axiomatization requires the decision maker to entertain signed as well as ordinary probability measures on the states.

Ke and Zhao (2022) includes new representation results for decision making under ambiguity. In their set-up, a preference relation satisfies the Gilboa and Schmeidler (1989) certainty-independence axiom if and only if it has a "cautiously optimistic linear utility" representation (equivalent to a utility function that is locally exactly linear almost everywhere). The representation involves a possibly signed (subjective) probability measure over states.

A crucial common feature in the quantum mechanical and statistical mechanical applications of signed probabilities is that all observable events must receive probability between 0 and 1. It is less clear-cut in the decision-theoretic application that this condition must hold, since the setting might be a one-off decision and probabilities could be subjective rather than frequentist. This said, in the subjective case, too, it would be decidedly unorthodox to allow negative probabilities on observable events. To be "conservative," we shall require all observable events in the formalism of this paper to receive probability in $[0, 1]$.

Returning to the classical Agreement Theorem as our starting point, we establish three results, where the second and third use the "common certainty" modality (to be defined later) in place of common knowledge:

- (i) In a non-classical domain, with signed probabilities, and as in the classical domain, it cannot be common knowledge that two agents assign different probabilities to an event of interest.
- (ii) In a non-classical domain, and unlike the classical domain, it can be common certainty that two agents assign different probabilities to an event of interest.
- (iii) In a non-classical domain, it cannot be common certainty that two agents assign different probabilities to an event of interest, if communication of their common certainty is possible – even if communication does not take place.

We formulate and prove these results in the following sections of the paper.

2. Example

Figure 1 depicts an epistemic state space that contains a non-classical component. There are four states, labeled ω_1 through ω_4 . There is a common prior, and the prior probabilities of the states are given in parentheses. Notice that the (prior) probability of state ω_3 is negative, which cannot, of course, happen in a classical setting. There are two agents, Alice and Bob. Alice receives private information about the true state as represented by the red sets partitioning the state space, while Bob receives private information as represented by the blue sets. Finally, we are interested in the agents' respective (conditional) probabilities of the event $E = \{\omega_1, \omega_3, \omega_4\}$ when the true state of the world is ω_1 .




Figure 1: Singular disagreement in the non-classical world

Let p denote the common prior (which is a signed probability measure). When the true state is ω_1 , Alice's information is $\{\omega_1\}$, so that her conditional probability of E is equal to

$$\frac{p(\{\omega_1, \omega_3, \omega_4\} \cap \{\omega_1\})}{p(\{\omega_1\})} = \frac{p(\{\omega_1\})}{p(\{\omega_1\})} = \frac{+\frac{1}{4}}{+\frac{1}{4}} = 1.$$

Bob's information is $\{\omega_1, \omega_2, \omega_3\}$, so that his conditional probability of E is equal to

$$\frac{p(\{\omega_1, \omega_3, \omega_4\} \cap \{\omega_1, \omega_2, \omega_3\})}{p(\{\omega_1, \omega_2, \omega_3\})} = \frac{p(\{\omega_1, \omega_3\})}{p(\{\omega_1, \omega_2, \omega_3\})} = \frac{+\frac{1}{4} - \frac{1}{4}}{+\frac{1}{4} + \frac{1}{2} - \frac{1}{4}} = \frac{0}{+\frac{1}{2}} = 0.$$

Thus, at state ω_1 , Alice assigns probability 1 to E and Bob assigns probability 0 to E . Next, we find the event, which we depict F , that Bob assigns probability 0 to E . We know that $\omega_1 \in F$. Bob's probability of E is again 0 at states ω_2 or ω_3 . At state ω_4 , Bob's probability of E is

$$\frac{p(\{\omega_1, \omega_3, \omega_4\} \cap \{\omega_4\})}{p(\{\omega_4\})} = \frac{p(\{\omega_4\})}{p(\{\omega_4\})} = \frac{+\frac{1}{2}}{+\frac{1}{2}} = 1,$$

so that $F = \{\omega_1, \omega_2, \omega_3\}$. At state ω_1 , Alice's probability of F is

$$\frac{p(\{\omega_1, \omega_2, \omega_3\} \cap \{\omega_1\})}{p(\{\omega_1\})} = \frac{p(\{\omega_1\})}{p(\{\omega_1\})} = \frac{+\frac{1}{4}}{+\frac{1}{4}} = 1.$$

Thus, at state ω_1 , Alice assigns probability 1 to E while at the same time she assigns probability 1 to Bob's assigning probability 0 to E . Say Alice is **certain** of an event E at a state ω_i if she assigns probability 1 to E , conditional on the information she has at state ω_i . Then the scenario we have just constructed is one where there is a state at which Alice is certain of an event E , and Alice is certain Bob is certain of the complementary event E^c . Call this a situation of **singular disagreement** (like calling two probability measures mutually singular). Evidently, this phenomenon can arise in a non-classical environment. In the next section, we will verify that singular disagreement cannot arise in a classical environment.

In a model with negative probabilities, events that receive probability in $[0, 1]$ are observable in the sense that they can be associated to frequencies. In the example, all partition cells, namely $\{\omega_1\}$, $\{\omega_2, \omega_3, \omega_4\}$, $\{\omega_1, \omega_2, \omega_3\}$, and $\{\omega_4\}$, receive probability in $(0, 1]$. So, they are observable and, in fact, strict positivity of these events ensures that the agents can also condition on them. The event E receives probability in $(0, 1]$, which means it is observable and non-trivial. We will consider observability further in Section 4.

3. General Formulation

For the general case, let the state space be a finite set Ω , and let Alice and Bob have partitions of Ω denoted by \mathcal{P}_A and \mathcal{P}_B , respectively. For $\omega \in \Omega$, the event $\mathcal{P}_A(\omega)$ consists of the member of Alice's partition that contains ω ; similarly for $\mathcal{P}_B(\omega)$. Let p denote the common (possibly signed) prior probability measure on Ω . We will assume throughout that all members of the partitions \mathcal{P}_A and \mathcal{P}_B receive non-zero probability, so that conditioning is well defined.

We begin with a remark about the classical domain.

Remark 3.1. Suppose that p is non-negative and fix an event E . Let F be the event that Bob assigns probability 0 to E , i.e.,

$$F = \{\omega' \in \Omega : p(E | \mathcal{P}_B(\omega')) = 0\}.$$

Then there is no state ω at which Alice assigns probability 1 to $E \cap F$.

Proof. Suppose there is such a state ω . Then $p(E | \mathcal{P}_A(\omega)) = 1$ and $p(F | \mathcal{P}_A(\omega)) = 1$. Note that we can write $F = \bigcup_{i \in I} \pi_i$ where each $\pi_i \in \mathcal{P}_B$ and I is a (finite) index set. In particular, there is a $\pi_i \in \mathcal{P}_B$ such that $p(E | \pi_i) = 0$ and $p(\pi_i | \mathcal{P}_A(\omega)) > 0$.

We now have three events $A (= \mathcal{P}_A(\omega))$, $B (= E)$, and $C (= \pi_i)$ such that $p(B | A) = 1$, $p(B | C) = 0$, and $p(C | A) > 0$. From $p(B | A) = 1$ we get $p(A \cap (C \setminus B)) = 0$. From $p(B | C) = 0$ we get $p(A \cap (B \cap C)) = 0$. It follows that $p(A \cap C) = 0$, contradicting $p(C | A) > 0$. \square

Remark 3.1 says that singular disagreement is impossible in the classical domain, verifying that this phenomenon is non-classical. The example of the previous section makes use of the fact that signed probability measures do not satisfy monotonicity. Specifically, in the proof just given, the step $p(A \cap (B \cap C)) = 0$ because $p(B | C) = 0$ fails with signed probabilities.

Next, we provide formal definitions of knowledge, common knowledge, certainty, and common certainty.

Definition 3.1. Alice knows an event E at state ω if $\mathcal{P}_A(\omega) \subseteq E$.

At state ω , Alice's information is that the true state lies in $\mathcal{P}_A(\omega)$. It follows that the true state therefore lies in any superset of $\mathcal{P}_A(\omega)$, i.e., that Alice knows all such events obtain. This is the standard definition of knowledge in the interactive epistemology literature. Some notation: The meet (finest common coarsening) of Alice's and Bob's partitions is written $\mathcal{P}_A \wedge \mathcal{P}_B$. The member of the meet that contains state ω is written $(\mathcal{P}_A \wedge \mathcal{P}_B)(\omega)$.

Definition 3.2. An event E is **common knowledge** between Alice and Bob at a state ω if $(\mathcal{P}_A \wedge \mathcal{P}_B)(\omega) \subseteq E$.

This definition of common knowledge is easily shown to be equivalent to the recursive definition (Alice knows E occurs, Bob knows E occurs, Alice knows Bob knows E occurs, etc.). Aumann (1976) proves this fact.

Definition 3.3. Alice is **certain of** an event E at state ω if $p(E | \mathcal{P}_A(\omega)) = 1$.

At state ω , Alice's information is that the true state lies in $\mathcal{P}_A(\omega)$. She is certain of E if she assigns probability 1 to E , conditional on this information. This is the standard epistemic definition of certainty.

Next, fix an event E and probabilities q_A and q_B . We define the event that it is common certainty that Alice assigns probability q_A to E and Bob assigns probability q_B to E . To do so,

let

$$A_0 = \{\omega \in \Omega : p(E | \mathcal{P}_A(\omega)) = q_A\},$$

$$B_0 = \{\omega \in \Omega : p(E | \mathcal{P}_B(\omega)) = q_B\},$$

and, in addition, let

$$A_{n+1} = A_n \cap \{\omega \in \Omega : p(B_n | \mathcal{P}_A(\omega)) = 1\},$$

$$B_{n+1} = B_n \cap \{\omega \in \Omega : p(A_n | \mathcal{P}_B(\omega)) = 1\},$$

for $n \geq 0$. The set A_0 contains all the states where Alice assigns probability q_A to E . The set A_1 contains all the states where the previous statement for Alice is true and, in addition, Alice is certain “Bob assigns probability q_B to E .” The set A_2 contains all the states where the previous statement for Alice is true and, in addition, Alice is certain “Bob assigns probability q_B to E and he is certain she assigns probability q_A to E .” And so on. In this way, the set A_n contains all the states where Alice has n th-order certainty. Likewise for Bob and the sets B_n , for all n .

Definition 3.4. 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

$$\omega^* \in \bigcap_{n=0}^{\infty} A_n \cap \bigcap_{n=0}^{\infty} B_n.$$

A special case is where $q_A = q_B = 1$. Then we can simply say that the event E is **common certainty** between Alice and Bob at ω^* .

It is clear that 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: We just gave the first step. Next, 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 can be continued to all higher levels.) But certainty is a strictly weaker modality than knowledge. (Also, common certainty is strictly weaker than common knowledge, as we will see in Theorem 4.2 below.) Figure 2 demonstrates this claim in two different instances – the first classical and the second non-classical. In both instances, Alice is certain of E but she does not know E .




Figure 2: Classical and non-classical knowledge-certainty distinction

4. Agreement and Disagreement

We can now state and prove a knowledge-based Agreement Theorem for both the classical and non-classical domains.

Theorem 4.1. Fix a common prior (which may be a signed probability measure) 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. The hypothesis of the theorem is that

$$(\mathcal{P}_A \wedge \mathcal{P}_B)(\omega^*) \subseteq A_0 \cap B_0.$$

Now, we can write $(\mathcal{P}_A \wedge \mathcal{P}_B)(\omega^*) = \bigcup_{i \in I} \pi_i$ where each $\pi_i \in \mathcal{P}_A$ and I is a (finite) index set. Since $(\mathcal{P}_A \wedge \mathcal{P}_B)(\omega^*) \subseteq A_0$, we have $p(E | \pi_i) = q_A$ for all $i \in I$. We also have

$$p(E | (\mathcal{P}_A \wedge \mathcal{P}_B)(\omega^*)) = \sum_{j \in I} p(\pi_j | \bigcup_{i \in I} \pi_i) \times p(E | \pi_j),$$

so that $p(E | (\mathcal{P}_A \wedge \mathcal{P}_B)(\omega^*))$ is an affine combination of q_A 's and is therefore equal to q_A . We can run exactly the same argument with B in place of A to conclude that $p(E | (\mathcal{P}_A \wedge \mathcal{P}_B)(\omega^*)) = q_B$. It follows that $q_A = q_B$. \square

It is of note that we did not need to impose any observability conditions in this theorem, making it fully general. (If we did add the condition that members of \mathcal{P}_A and \mathcal{P}_B receive strictly positive – as opposed to non-zero – probability, then the affine combination in the proof would become a convex combination and the proof would be exactly that in Aumann, 1976.)

Theorem 4.1 is similar to results in Leifer and Duarte (2022) on the impossibility of common knowledge of disagreement, established in the setting of generalized probability theory or GPT (Barrett, 2007). GPT is a multi-purpose operational framework for describing physical theories, including quantum mechanics.

In the classical domain, with a non-negative prior, there is also an Agreement Theorem for the certainty modality: If two agents have common certainty of each other's probabilities of E , then these probabilities must be equal, just as with common knowledge. We do not give a direct proof here, since the result will be a corollary to our Theorem 5.1 later. Taken together, Theorem 4.1 (for the classical case of a non-negative prior) and the analog for common certainty indicate that the distinction between the knowledge and certainty modalities is "small" – at least, for current purposes – in the classical domain. But the distinction is very significant in the non-classical domain, because the Agreement Theorem for certainty no longer holds there, as we are about to see.

We suggest that the certainty modality in epistemics is at least as interesting as the knowledge modality. Certainty is subjective in that an agent can be certain of an event E , but E need not happen. Knowledge is objective and satisfies the truth axiom: If an agent knows E , then E must occur. The subjective modality seems more in line with the idea that Alice and Bob are Bayesian agents forming their personalistic beliefs, beliefs about beliefs, and so on, about some event. For Alice to know Bob's beliefs (or knowledge) requires that she have direct information about his epistemic state. This introduces an ex post element to the analysis in the sense that Bob's epistemic state would need to be observed by Alice (via some information flow). The certainty modality allows an ex ante analysis where agents form prospective beliefs about events, just as in Bayesian decision theory (Savage, 1954). In any case, we think the point is made that the certainty modality is important to study and, as we now show, it is very different from the knowledge modality in the non-classical world.

Theorem 4.2. *There is a structure $(\Omega, p, \mathcal{P}_A, \mathcal{P}_B)$, where p is a signed prior, and there is an event E and a state ω such that it is common certainty at ω that Alice and Bob hold different probabilities of E .*




Figure 3: Common certainty of disagreement in a non-classical world

Proof. The state space and prior are depicted in Figure 3. As in our earlier example, Alice's partition comprises the red sets and Bob's partition comprises the blue sets. The event $E = \{\omega_2, \omega_4, \omega_5, \omega_6\}$ and the true state is ω_5 . The numbers ϵ and η are small and positive with $\epsilon \neq \eta$. Set

$$A_0 = \{\omega \in \Omega : p(E | \mathcal{P}_A(\omega)) = 1 - 2\epsilon\} = \{\omega_1, \omega_2, \omega_5\},$$

$$B_0 = \{\omega \in \Omega : p(E | \mathcal{P}_B(\omega)) = 1 - 2\eta\} = \{\omega_3, \omega_4, \omega_5\},$$

so that

$$A_1 = A_0 \cap \{\omega \in \Omega : p(B_0 | \mathcal{P}_A(\omega)) = 1\} = \{\omega_1, \omega_2, \omega_5\},$$

$$B_1 = B_0 \cap \{\omega \in \Omega : p(A_0 | \mathcal{P}_B(\omega)) = 1\} = \{\omega_3, \omega_4, \omega_5\},$$

from which $A_{n+1} = A_n$ and $B_{n+1} = B_n$ for all $n \geq 1$. It follows that $\omega_5 \in \bigcap_{n=0}^{\infty} A_n \cap \bigcap_{n=0}^{\infty} B_n$. At state ω_5 , it is common certainty between Alice and Bob that she assigns probability $1 - 2\epsilon$ to E while he assigns probability $1 - 2\eta$ to E , which proves the theorem. \square

Note that, by Theorem 4.1, the agents' probabilities of E cannot be common knowledge at ω_5 (because then the probabilities must be the same). Alternatively, this can be checked directly via the definition of common knowledge in terms of the join of $\mathcal{P}_A \wedge \mathcal{P}_B$ (which is the whole space). So, this example also serves to establish the claim that common certainty is strictly weaker than common knowledge. Note also that the example exhibits a high degree of observability: All members of \mathcal{P}_A and \mathcal{P}_B get strictly positive probability $(+1/2)$. The event of interest E gets probability $1 - \epsilon - \eta > 0$.

5. Communication

Common knowledge and common certainty are different from communication between agents. If Alice announces the probability she assigns to an event of interest E , then this communicates information to Bob and he can update his partition \mathcal{P}_B to incorporate this information. Vice versa if Bob communicates to Alice, who can then announce new probabilities. This process could continue. The communication of probabilities this way was first studied by Geanakoplos and Polemarchakis (1982).

Let us try to go down this road in the non-classical environment of this paper. Go back to the structure in Figure 3. The event of interest is $E = \{\omega_2, \omega_4, \omega_5, \omega_6\}$ as before, but now the true state is ω_1 . Alice begins by announcing her probability $1 - 2\epsilon$ of E to Bob. When Bob hears this announcement, he tries to make an inference about what information Alice has – specifically, whether she has the information $\{\omega_1, \omega_2, \omega_5\}$ (as, in fact, she does) or the information $\{\omega_3, \omega_4, \omega_6\}$ (which she does not), or whether he cannot tell which piece of information she has. Start with the second case. Bob can reason that, in this case, Alice would have announced a probability $1 - 2\eta$ of

E . Since she did not, he can infer that she did not observe $\{\omega_3, \omega_4, \omega_6\}$. Next is the case that Alice has the information $\{\omega_1, \omega_2, \omega_5\}$. Bob can reason that, in this case, Alice would have announced a probability $1 - 2\epsilon$ of E . Since she did just this, he can infer that she indeed observed $\{\omega_1, \omega_2, \omega_5\}$. Summing up, Bob will update his own information to

$$\{\omega_1, \omega_2, \omega_6\} \cap \{\omega_1, \omega_2, \omega_5\} = \{\omega_1, \omega_2\}.$$

Now ask what probability of E will Bob announce. We can try to calculate this as

$$\frac{p(E \cap \{\omega_1, \omega_2\})}{p(\{\omega_1, \omega_2\})} = \frac{-\epsilon}{0},$$

which is obviously ill-defined.

What has gone “wrong” in this process is that Bob is unable to process the information contained in Alice’s announcement in a meaningful way. There is the negative number $-\epsilon$ in the numerator and the zero in the denominator. We argue that both features pose a conceptual problem. In systems where the agents are able to communicate about an event of interest, we propose that 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 impose a condition on our epistemic structures that ensures all communication regarding an event of interest is classical. In doing this, we depart from the Geanakoplos and Polemarchakis (1982) protocol in a key way. We focus on the initial announcements that Alice and Bob can make, i.e., announcements relative to their initial partitions \mathcal{P}_A and \mathcal{P}_B , respectively. But we allow that these announcements might be of any order. That is, Alice might announce her probability of E , or she might announce her certainty (or not) of what Bob’s probability of E is, or she might announce her certainty (or not) about Bob’s certainty, and so on. Likewise for Bob. This is different from the Geanakoplos-Polemarchakis protocol, which permits announcements only of (conditional) probabilities of E , but allows these announcements to continue, as Alice and Bob successively update their information. With our approach, we are able to (re-)establish an agreement theorem even in the non-classical domain. We leave it as open whether or not such a result is possible if an appropriate notion of classicality is imposed on the Geanakoplos-Polemarchakis protocol adapted to our setting.

Formally, for all $n \geq 0$, let

$$\begin{aligned}\mathcal{M}_A^{(n)} &= \{A_n, A_n^c\}, \\ \mathcal{M}_B^{(n)} &= \{B_n, B_n^c\},\end{aligned}$$

where A_n and B_n are the sets defined back in Section 3.

Definition 5.1. For any $\pi, E \subseteq \Omega$, say π is **regular with respect to E** if $p(\pi) \geq 0$ and $p(\pi \cap E)$ lies in $[0, p(\pi)]$.

Definition 5.2. A structure $(\Omega, p, \mathcal{P}_A, \mathcal{P}_B)$ is **communication-enabled with respect to E** if, for each $n \geq 0$, all $\pi \in \mathcal{P}_A \vee \mathcal{M}_B^{(n)}$ and all $\pi' \in \mathcal{P}_B \vee \mathcal{M}_A^{(n)}$ are regular with respect to E .

These definitions capture our requirement that if an agent were to communicate their initial certainty at any level, the calculation that the other agent would then make is classical.

Remark 5.1. Fix $\pi, \pi' \subseteq \Omega$ with $\pi \cap \pi' = \emptyset$. Then if π and π' are both regular with respect to E , so is $\pi \cup \pi'$.

We have assumed throughout that all members of \mathcal{P}_A and \mathcal{P}_B receive non-zero probability. Definition 5.2 assumes nonnegative probabilities. So, at this point, we are assuming that all

members of \mathcal{P}_A and \mathcal{P}_B receive strictly positive probability. This implies that the agents are able to observe and condition on their own information using the rules of ordinary probability.

Theorem 5.1. Fix a structure $(\Omega, p, \mathcal{P}_A, \mathcal{P}_B)$ that is communication-enabled with respect to E and suppose that 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 $q_A = q_B$.

Proof. Begin by defining A_n and B_n , for $n \geq 0$, as before. Since Ω is finite, there is an N (finite) such that for all $n \geq N$, $A_{n+1} = A_n$ and $B_{n+1} = B_n$. We have

$$A_{N+1} = A_N \cap \{\omega \in \Omega : p(B_N | \mathcal{P}_A(\omega)) = 1\} = A_N,$$

from which $p(B_N | \mathcal{P}_A(\omega)) = 1$ for all $\omega \in A_N$.

Now $A_N = \bigcup_{i \in I} \pi_i$ where each $\pi_i \in \mathcal{P}_A$ and I is a (finite) index set. We just saw that $p(B_N | \pi_i) = 1$ for all such π_i . But $p(B_N | A_N)$ is a convex combination of the $p(B_N | \pi_i)$'s, so $p(B_N | A_N) = 1$. It follows that $p(A_N \setminus B_N) = 0$, which we will use shortly.

Observe that $A_N \subseteq A_0$ and so $p(E | \pi_i) = q_A$ for these same π_i . By a second convex combination argument, $p(E | A_N) = q_A$.

Next, observe that $\{A_N, A_N^c\}$ is a coarsening of \mathcal{P}_A (by definition of the A_n 's). From this and $\mathcal{M}_B^{(N)} = \{B_N, B_N^c\}$, it follows that $\{A_N \setminus B_N, (A_N \setminus B_N)^c\}$ is a coarsening of $\mathcal{P}_A \vee \mathcal{M}_B^{(N)}$. By the hypothesis of the theorem and Remark 5.1, it follows that $A_N \setminus B_N$ is regular with respect to E . Using $p(A_N \setminus B_N) = 0$, it follows that $p((A_N \setminus B_N) \cap E) = 0$, and so $p(E \cap A_N \cap B_N) = p(E \cap A_N)$. Again using $p(A_N \setminus B_N) = 0$, we get $p(A_N \cap B_N) = p(A_N) > 0$ (the set A_N is a union of members of \mathcal{P}_A). We conclude that $p(E | A_N \cap B_N) = p(E | A_N) = q_A$. We can run exactly the same argument with B in place of A to conclude that $p(E | A_N \cap B_N) = q_B$. It follows that $q_A = q_B$. \square

We are not committed to Theorem 5.1 over Theorem 4.2. There is no formal or even obvious conceptual inconsistency in the set-up of Theorem 4.2. Still, it is interesting to discover from Theorem 5.1 that if we impose the requirement that each agent be able to process classically an announcement by the other agent of their certainty at any level, then the non-classical phenomenon of common certainty of disagreement disappears. There is a subtle point here. No actual communication needs to take place. Rather, we can think of our requirement as saying that it would be possible for the two agents to confirm their disagreement, not just have common certainty of their disagreement, if they wanted to.

A corollary to Theorem 5.1 is that common certainty of disagreement is impossible in the classical world, as we mentioned earlier. This follows because the condition of being communication-enabled is automatically satisfied in the case of non-negative probabilities.

Consider another communication scenario: There is a third agent, Charlie, who starts out with no information about the true state. Alice and Bob are able to communicate with Charlie, but not with each other. (They do not necessarily undertake the communication.) We can ask if this scenario, too, rules out common certainty of disagreement. Here is the appropriate analog to Definition 6.

Definition 5.3. A structure $(\Omega, p, \mathcal{P}_A, \mathcal{P}_B)$ is **third-party communication-enabled with respect to E** if, for each $n \geq 1$, each $\pi \in \mathcal{M}_A^{(n)} \vee \mathcal{M}_B^{(n)}$ is regular with respect to E .

The idea is that the third party, Charlie, starts with the trivial partition $\{\Omega, \emptyset\}$ and is then able to make classical calculations with the information which announcements by Alice and Bob might give him. Alice and Bob do not communicate with each other.

Theorem 5.2. Fix a structure $(\Omega, p, \mathcal{P}_A, \mathcal{P}_B)$ that is third-party communication-enabled with respect to E and suppose that 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 $q_A = q_B$.

Proof. From $\mathcal{M}_A^{(N)} = \{A_N, A_N^c\}$ and $\mathcal{M}_B^{(N)} = \{B_N, B_N^c\}$ it follows that $\{A_N \setminus B_N, (A_N \setminus B_N)^c\}$ is a coarsening of $\mathcal{M}_A^{(N)} \vee \mathcal{M}_B^{(N)}$. Using the hypothesis of the theorem and Remark 2, we conclude that $A_N \setminus B_N$ is regular with respect to E . The rest of the proof follows exactly the proof of Theorem 5.1. \square

Finally, in this section, we note that a complete treatment of updating of probabilities in the non-classical domain would require the development of some new probability theory. The issue is that, as can happen in a classical setting, an agent may come to learn that an event to which they had assigned probability 0 actually obtains. We saw this in the example above where we got a zero in the denominator of Bob's updated probability. In the classical domain, the answer to the probability-0 problem is to move to the concept of a conditional probability system (Rényi, 1955). This is a family of probability measures – one measure for each event an agent might learn, including events to which the agent assigns probability 0. What would be needed in our non-classical domain is an extension of the concept of a conditional probability system to signed probabilities, which would be an exercise in pure probability theory that, to the best of our knowledge, has not been undertaken.

6. Conclusion

We end with some comments on the realizability of common certainty of disagreement (CCD) as in Theorem 4.2. In the physical domain, it can be shown that CCD is impossible when observing quantum systems, but possible when observing superquantum (no-signaling) systems (Contreras-Tejada et al., 2021). In the language of this paper, we can say that quantum mechanics somehow controls the “extent” of negativity in phase-space probability representations so that CCD cannot arise. This finding suggests there may be promise in proposing the impossibility of CCD (of “agreeing to disagree”) as an axiom in the program to derive quantum mechanics from underlying physical principles. (See Contreras-Tejada et al., 2021 for further discussion and references to the axiomatization program.)

It would be interesting to connect this paper to the study of contextuality scenarios in quantum and superquantum systems. In particular, Cabello (2013) extends the usual contextuality scenario involving a single observer to allow for a copy of the system with a second observer. He identifies a principle he calls “global exclusivity” that exactly identifies the maximum quantum violation of certain non-contextuality inequalities. Combining formal epistemics as in the current paper with this physical principle could be a promising direction.

In the setting of decision theory – more precisely, multi-person decision theory – Theorem 4.2 indicates that if we equip agents with signed probability measures, we can get highly non-classical behavior, such as betting between risk-neutral agents. With this in mind, we wonder whether it might be interesting to elevate the impossibility of CCD to an (epistemic) multi-person decision-theoretic principle. This might offer a disciplined departure from classical behavior and appears to be an open direction.

A preference basis for a decision theory with signed probabilities (perhaps, building on Perea, 2022 or Ke and Zhao, 2022) would be of interest in its own right and might also lead to a preference basis for certainty in a non-classical environment. There is a preference basis for certainty in classical decision theory. An agent is certain of an event E if and only if the complementary event E^c is Savage-null; that is, if all acts conditional on E^c are deemed indifferent. An open question is what would be the analogous definition in the signed case. (We are grateful to Miklós Pintér for raising this question.)

A different non-classical examination of the Agreement Theorem is undertaken by Khrennikov and Basieva (2014) and Khrennikov (2015). They consider quantum-like observers of a quantum system who employ either the knowledge or certainty modality. Their approach does not deliver an Agreement Theorem even for quantum systems.

Summing up, our Theorem 4.2 establish a new kind of non-classical strangeness in the form of the possibility of CCD. At the same time, we also prove that common knowledge of disagreement

and CCD in communication-enabled structures remain impossible (Theorems 4.1, 5.1, and 5.2). We believe these results open the door to further investigation of epistemics in non-classical worlds.

Conflict of interests. We declare no competing interests.

Funding. We acknowledge financial support from NYU Stern School of Business, NYU Shanghai, and J.P. Valles (Brandenburger); MTM2017-88385-P funded by MCIN/AEI/10.13039/501100011033 and by “ERDF A way of making Europe” and SEV-2015-0554-16-3 funded by MCIN/AEI/10.13039/501100011033 (Contreras-Tejada); Deutsche Bundesbank (La Mura); MTM2014- 54240-P funded by MCIN/AEI/10.13039/501100011033 and by “ERDF A way of making Europe” and ICMAT Severo Ochoa project SEV-2015-0554 funded by MCIN/AEI/10.13039/501100011033 (Scarpa); NIH grant R01DA038063 (Steverson); and QUITEMAD-CM P2018/TCS4342 funded by CAM/FEDER and PID2020-113523GB-I00 funded by MCIN/AEI/10.13039/501100011033 and by “ERDF A way of making Europe” (Contreras-Tejada and Scarpa).

Acknowledgements. We are grateful for important input from audiences at the 2023 Inaugural Science of Decision Making (SDM) Conference, Hong Kong University; the Baptist-Kyoto-NTU-Osaka-Sinica Workshop on Economic Theory; the 5th Workshop on Quantum Contextuality in Quantum Mechanics and Beyond, Czech Association of Scientific and Technological Societies; the 30th European Workshop on Economic Theory, SGH Warsaw School of Economics; One World Mathematical Game Theory Seminar; Seminar in Economic Theory, University of Warwick; Encapsulated Agents in Quantum Theory: Re-examining Wigner’s Friend, UMass Boston; and LFDSN 2019, Zhejiang University. Emiliano Catonini, Kim-Sau Chung, Dagomir Kaszlikowski, Shaowei Ke, Paweł Kurzyński, Andrés Perea, Miklós Pintér, Takuro Yamashita, Noson Yanofsky, and Chen Zhao provided valuable comments. We thank two referees and the editor for proposing substantial improvements to the paper. Eduardo Zambrano provided initial inspiration for this project. Our results were first announced in an extended abstract published in *Proceedings of the 2019 Workshop on Logics for the Formation and Dynamics of Social Norm (LFDSN 2019)*, ed. by Beishui Liao, Fenrong Liu, and Huimin Dong, at <https://www.xilogic.org/events/wp-content/uploads/2022/05/lfdn2019-proceedings.pdf>.

References

1. Abramsky S, Brandenburger A. 2011. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, **13**.
2. Aumann R. 1976. Agreeing to disagree. *Annals of Statistics*, **4**.
3. Aumann R, Brandenburger A. 1995. Epistemic conditions for Nash equilibrium. *Econometrica*, **63**.
4. Bach A, Blank H, Francke H. 1985. Bose-Einstein statistics derived from the statistics of classical particles. *Lettere al Nuovo Cimento*, **43**.
5. Barrett J. 2007. Information processing in generalized probabilistic theories. *Physical Review A*, **75**.
6. Bell J. 1964. On the Einstein Podolsky Rosen paradox. *Physics*, **1**.
7. Brandenburger A, La Mura P, Zoble S. 2022. Rényi entropy, signed probabilities, and the qubit. *Entropy*, **24**.
8. Cabello A. 2013. Simple explanation of the quantum violation of a fundamental inequality. *Physical Review Letters*, **110**.
9. Contreras-Tejada P, Scarpa G, Kubicki A, Brandenburger A, La Mura P. 2021. Observers of quantum systems cannot agree to disagree. *Nature Communications*, **12**.
10. De Finetti B. 1931. Funzione caratteristica di un fenomeno allattorio. *Atti della R. Accademia Nazionale dei Lincii Ser. 6, Memorie, Classe di Scienze, Fisiche, Matematiche e Naturali*, **4**.
11. Dellacherie C, Meyer PA. 1982. *Probabilities and Potential B*. Amsterdam: North-Holland.
12. Dirac P. 1942. The physical interpretation of quantum mechanics. *Proceedings of the Royal Society of London (Series A: Mathematical and Physical Sciences)*, **180**.
13. Feynman R. 1987. Negative probability. In Hiley B, Peat F (eds.), *Quantum Implications: Essays in Honour of David Bohm*. London: Routledge and Kegan Paul.
14. Geanakoplos J, Polemarchakis H. 1982. We can’t disagree forever. *Journal of Economic Theory*, **28**.
15. Gilboa I, Schmeidler D. 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, **18**.

16. Janson S, Konstantopoulos T, Yuan L. 2016. On a representation theorem for finitely exchangeable random vectors. *Journal of Mathematical Analysis and Applications*, **442**.
17. Jaynes E. 1986. Some applications and extensions of the de Finetti representation theorem. In Goal P, Zellner A (eds.), *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*. Amsterdam: North-Holland.
18. Kaszlikowski D, Kurzyński P. 2023. Alternative quantum mechanics? In Plotnitsky A, Haven E (eds.), *The Quantum-Like Revolution: A Festschrift for Andrei Khrennikov*. Berlin: Springer.
19. Ke S, Zhao C. 2022. From local utility to neural networks. See <https://sites.google.com/site/shaoweike/research>.
20. Kearns G, Székely G. 2006. De Finetti's theorem for abstract finite exchangeable sequences. *Journal of Theoretical Probability*, **19**.
21. Kenfack A, Życzkowski K. 2004. Negativity of the Wigner function as an indicator of non-classicality. *Journal of Optics B: Quantum and Semiclassical Optics*, **6**.
22. Khrennikov A. 2015. Quantum version of Aumann's approach to common knowledge: sufficient conditions of impossibility to agree on disagree. *Journal of Mathematical Economics*, **60**.
23. Khrennikov A, Basieva I. 2014. Possibility to agree on disagree from quantum information and decision making. *Journal of Mathematical Psychology*, **62**.
24. Leifer M, Duarte C. 2022. Generalizing Aumann's agreement theorem. See <https://arxiv.org/abs/2202.02156>.
25. Milgrom P, Stokey N. 1982. Information, trade, and common knowledge. *Journal of Economic Theory*, **26**.
26. Onggadinata K, Kurzyński P, Kaszlikowski D. 2023. Qubit from the classical collision entropy. *Physical Review A*, **107**.
27. Perea A. 2022. Expected utility as an expression of linear preference intensity. Working paper. Maastricht University.
28. Popescu S, Rohrlich D. 1994. Quantum nonlocality as an axiom. *Foundations of Physics*, **24**.
29. Rényi A. 1955. On a new axiomatic theory of probability. *Acta Mathematica Academiae Scientiarum Hungarica*, **6**.
30. Savage L. 1954. *The Foundations of Statistics*. New York: Wiley.
31. Sebenius J, Geanakoplos J. 1983. Don't bet on it: contingent agreements with asymmetric information. *Journal of the American Statistical Association*, **78**.
32. Wigner E. 1932 On the quantum correction for thermodynamic equilibrium. *Physical Review*, **40**.