Agreement and Disagreement in a Non-Classical World*

Adam Brandenburger[†], Patricia Contreras-Tejada[‡], Pierfrancesco La Mura[§], Giannicola Scarpa[¶], and Kai Steverson[‖]

May 19, 2019
Very Preliminary

Abstract

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, and so we cannot assume that the same result holds when agents observe non-classical phenomena. This matter has been brought to the fore in a recent paper by Frauchiger and Renner (2018), which produces a quantum scenario exhibiting a kind of extreme disagreement that is impossible classically. Inspired by their use in quantum mechanics, we employ signed probability measures ("quasi-probabilities") to investigate the epistemics of the non-classical world and ask, in particular: What conditions lead to agreement or allow for disagreement when agents may use signed probabilities?

*We thank audiences at UMass Boston and Zhejiang University (Hangzhou) for important input, and Noson Yanofsky and Eduardo Zambrano for valuable comments. We acknowledge financial support from NYU Stern School of Business, NYU Shanghai, and J.P. Valles (Brandenburger); Deutsche Bundesbank (La Mura); MTM2014- 54240-P (MINECO), QUITEMAD+- CM Reference: S2013/ICE-2801 (Comunidad de Madrid), and ICMAT Severo Ochoa project SEV-2015-0554 (MINECO) (Scarpa); and NIH grant R01DA038063 (Steverson).

[†]Stern School of Business, Tandon School of Engineering, NYU Shanghai, New York University, New York, NY 10012, U.S.A., adam.brandenburger@nyu.edu, http://www.adambrandenburger.com

[‡]Institituto de Ciencias Mathemáticas, 28049 Madrid, Spain, patrcont@ucm.es

[§]HHL Leipzig Graduate School of Management, 04109 Leipzig, Germany, plamura@hhl.de

[¶]Universidad Complutense de Madrid, 28040 Madrid, Spain, giannicolascarpa@gmail.com

[‖]Department of Neuroscience, New York University, New York, NY 10003, U.S.A., ksteverson@nyu.edu, https://sites.google.com/site/ksteverson/

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.

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.

All this applies in the classical domain where classical probability theory applies. But in settings where non-classical probability theory does not apply – the canonical case is the quantum domain, of course — we cannot assume that the same facts about agreement and disagreement between Bayesian agents will hold. A recent paper by Frauchiger and Renner (2018) brings this matter to the fore. The heart of their argument is (in simplified form) a quantum scenario in 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). We show that this scenario cannot arise in a classical setting and we verify, in a very simple setting, that it can arise with non-classical probabilities.

Frauchiger and Renner (2018) take their finding further, by adopting an epistemic collapse axiom: If Alice is certain Bob is certain of an event F, then she is certain of F. Clearly, this axiom applied to the scenario of singular disagreement yields a contradiction. This conclusion has provoked considerable discussion in the quantum foundations community (see Pusey [2018] for commentary and additional references). Here, our focus is not on the physics but on the epistemics of the Frauchiger-Renner and related scenarios. (In particular, we do not adopt the epistemic collapse axiom.)

Inspired by their use in quantum mechanics (Wigner [1932], Dirac [1942], Feynman [1987], Wooters [1987]), we employ signed probability measures ("quasi-probabilities") to investigate the epistemics of the non-classical world and ask, in particular: What conditions lead to agreement or allow for disagreement when agents may use signed probabilities?[1]

We establish three results:

a. In a non-classical domain, and as in the classical domain, it cannot be common knowledge that two agents assign different probabilities to an event of interest.

b. 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.

c. 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.

2 Example

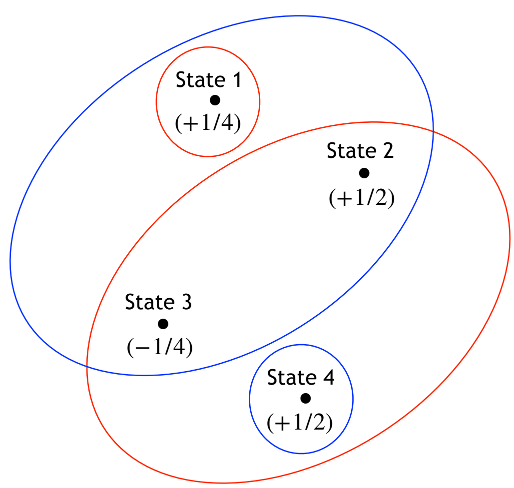

Figure 1: Singular disagreement in a non-classical world

Figure 1 depicts the state-space description of some physical system that contains non-classical components. There are four states, labeled 1 through 4. There is a common prior, and the prior probabilities of the states are given in parentheses. Notice that the (prior) probability of state 3 is negative, which cannot, of course, happen in a classical setting. (The assumption of a common prior, whether non-negative or not, is natural for the application to physical experiments where there is a commonly held — if probabilistic — theory which

[1] Abramsky and Brandenburger (2011) show that, in fact, signed probabilities characterize all no-signaling systems.

the agents employ to make their calculations.) 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 will be interested in the agents' respective (conditional) probabilities of the event $E = \{1, 3, 4\}$, when the true state of the world is 1.

Our formulation is inspired by the use of negative probabilities in the phase-space formulation of quantum mechanics (Wigner [1932], Dirac [1942], Feynman [1987], Wooters [1987]). Still, our treatment is abstract, and further work is needed to close the gap between our framework and quantum phase space.

Let's now calculate what probability Alice and Bob, respectively, assign to the event E, when the true state is 1. Letting p denote the common prior (which is a signed probability measure), we can write Alice's probability of E as

$$\frac{p(\{1,3,4\} \cap \{1\})}{p(\{1\})} = \frac{p(\{1\})}{p(\{1\})} = \frac{+\frac{1}{4}}{+\frac{1}{4}} = 1,$$

and Bob's probability of E as

$$\frac{p(\{1,3,4\} \cap \{1,2,3\})}{p(\{1,2,3\})} = \frac{p(\{1,3\})}{p(\{1,2,3\})} = \frac{+\frac{1}{4} - \frac{1}{4}}{+\frac{1}{4} + \frac{1}{2} - \frac{1}{4}} = \frac{0}{+\frac{1}{2}} = 0.$$

Thus, at state 1, Alice assigns probability 1 to E and Bob assigns probability 0 to E. Next, let's find the event, which we label G, that Bob assigns probability 0 to E. We know that $1 \in G$. Bob's probability of E is again 0 at states 2 or 3. At state 4, Bob's probability of E is

$$\frac{p(\{1,3,4\} \cap \{4\})}{p(\{4\})} = \frac{p(\{4\})}{p(\{4\})} = \frac{+\frac{1}{2}}{+\frac{1}{2}} = 1,$$

so that $G = \{1, 2, 3\}$. At state 1, Alice's probability of G is

$$\frac{p(\{1,2,3\} \cap \{1\})}{p(\{1\})} = \frac{p(\{1\})}{p(\{1\})} = \frac{+\frac{1}{4}}{+\frac{1}{4}} = 1.$$

Thus, at state 1, Alice assigns probability 1 to E while at the same time she assigns probability 1 to Bob's assigning probability 0 to E. This is an example of what we called singular disagreement in the Introduction. The example shows how 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 that they can be associated to actual outcomes and observed frequencies. In the example, all partition cells, namely $\{1\},\{2,3,4\},\{1,2,3\}$, and $\{4\}$, 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. (Less central, all events in the join of the two agents' partitions, namely, $\{1\},\{2,3\}$, and $\{4\}$, receive probability in $(0,1]$, so that, if they pooled their information, the agents could also observe and condition on the resulting events.) The event E receives probability in $(0,1]$, which means it is observable and non-trivial. We will consider observability further in the next section.

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

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

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

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

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

Remark 1 says that singular disagreement is impossible in the classical domain, verifying that this phenomenon is non-classical. The example of the previous section uses the fact that signed measures do not satisfy monotonicity. Specifically, in the proof here, the step $p(A\cap(B\cap C))=0$ because $p(B\mid C)=0$ fails with signed measures, thereby allowing singular disagreement.

Next, we establish the implication of common knowledge of the agents' probabilities. Formally, the result is an extension of the existing Agreement Theorem (Aumann [1976]), which treats the classical case of non-negative probabilities. We interpret the extension as providing a limit to the degree of disagreement possible in the non-classical world. Disagreement cannot be complete in the precise sense that there cannot be common knowledge of disagreement.

We need formal definitions of knowledge and then common knowledge. (We will provide definitions of certainty and common certainty later.)

Definition 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 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.

Proposition 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. Define

$$A_0 = \{\omega \in \Omega : p(E \mid \mathcal{P}_A(\omega)) = q_A\},$$
$$B_0 = \{\omega \in \Omega : p(E \mid \mathcal{P}_B(\omega)) = q_B\}.$$

The hypothesis of the proposition 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 an index set (finite). Since $(\mathcal{P}_A \wedge \mathcal{P}_B)(\omega^*) \subseteq A_0$, we have $p(E \mid \pi_i) = q_A$ for all $i \in I$. We also have

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

so that $p(E \mid (\mathcal{P}_A \wedge \mathcal{P}_B)(\omega^*))$ 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 $p(E \mid (\mathcal{P}_A \wedge \mathcal{P}_B)(\omega^*)) = q_B$. It follows that $q_A = q_B$. □

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].)

This result shows that the Agreement Theorem, which establishes a limit to disagreement in the classical world, extends to the non-classical world. Next, we show that, this result notwithstanding, there can be a strong type of disagreement in the non-classical case, different from the singular disagreement we have already seen.

Definition 3. Alice is *certain of* an event E at state ω if $p(E \mid \mathcal{P}_A(\omega)) = 1$.

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 definition of certainty in the interactive epistemology literature. 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, again let

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

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

and, in addition, let

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

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

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 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 is true and, in addition, Alice is certain Bob 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 nth-order certainty of E, and similarly for Bob and the sets B_n for all n.

Definition 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^*.

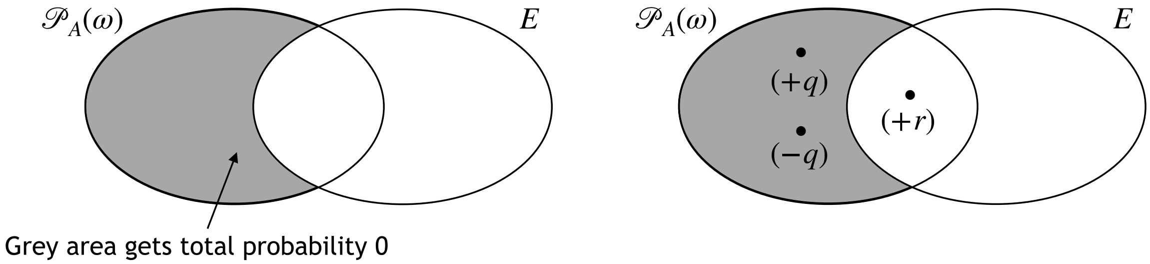

Figure 2: Classical and non-classical knowledge-certainty distinction

Going back to the definition of certainty, 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 later.) 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.

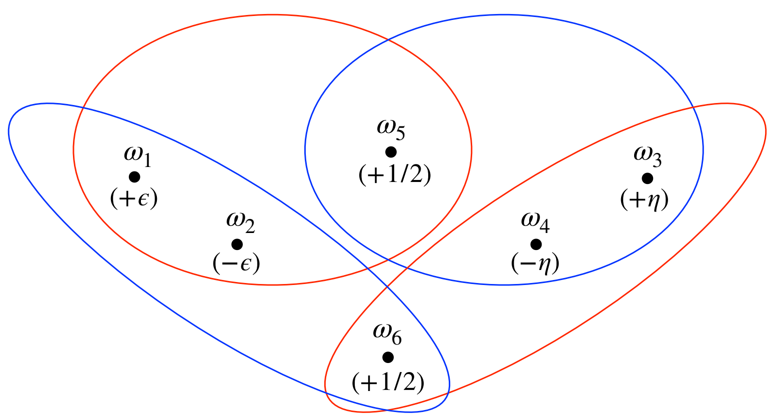

Figure 3: Common certainty of disagreement in a non-classical world

This distinction is not significant in the classical world, since there is also an Agreement Theorem for certainty: If two agents have common certainty of each other's probabilities of E, then these probabilities must be equal, just as with common knowledge. (This is a corollary to Proposition 2 below.) But the distinction is fundamental in the non-classical world. The next example shows that it is possible to have common certainty of disagreement. This is a new point about how the non-classical and classical worlds differ.

As before, Alice receives private information about the true state as represented by the red sets partitioning the state space in Figure 3, while Bob receives private information as represented by the blue sets. The underlying event of interest is $E = \{\omega_2, \omega_4, \omega_5, \omega_6\}$ and the true state of the world is ω_5. The numbers ϵ and η are small and positive with $\epsilon \neq \eta$. Set

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

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

so that

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

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

from which $A_{n+1} = A_n$ and $B_{n+1} = B_n$ for all $n \geq 1$. It follows that $\omega_5 \in \bigcap_{n=0}^{\infty} A_n \cap \bigcap_{n=0}^{\infty} B_n$. In words, 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. We conclude that in a setting with signed probabilities, it is possible for there to be common certainty of disagreement.

Note that, by Proposition 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 ($+1/2$). All members of the join (coarsest common refinement) $\mathcal{P}_A \vee \mathcal{P}_B$ get non-negative probability (0 or $+1/2$). The event of interest E gets probability $1 - \epsilon - \eta > 0$.

4 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 Ann, who can then announce new probabilities. This process could continue. (The communication of probabilities this way was first studied by Geanakoplos and Polemarchakis [1982].)

We assume that communication yields observable probabilities. If not, then communication would cease being meaningful. To see the issue, go back to the example of Figure 3 and suppose the true state is ω_1. Then both Alice and Bob assign probability $1-2\epsilon$ to $E=\{\omega_2,\omega_4,\omega_5,\omega_6\}$. Alice now communicates her probability to Bob, which tells Bob she has the information $\{\omega_1,\omega_2,\omega_5\}$. (If she had the information $\{\omega_3,\omega_4,\omega_6\}$, she would have communicated a probability of $1-2\eta$.) So, Bob's updated partition is $\mathcal{P}_B \vee \{\{\omega_1,\omega_2,\omega_5\},\{\omega_3,\omega_4,\omega_6\}\} = \{\{\omega_1,\omega_2\},\{\omega_3,\omega_4\},\{\omega_5\},\{\omega_6\}\}$ and his information is that the true state lies in $\{\omega_1,\omega_2\}$. Next, Bob calculates his new probability of E as $p(E\cap\{\omega_1,\omega_2\})/p(\{\omega_1,\omega_2\}) = -\epsilon/0$, which does not lie in $[0,1]$ (and is not even well-defined). Communication ceases being meaningful at this point. We will say that the set-up in Figure 3 is not communication-enabled.

For a formal definition, we first define a sequence of partitions for Alice, corresponding to the sequence of announcements she could make about her probability of E, her certainty or not of what Bob's probability of E is, her certainty or not about Bob's certainty, and so on. Likewise for Bob.[2] Formally, for all $n \geq 0$, let

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

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

Definition 5. 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 6. A structure $(\Omega, p, \mathcal{P}_A, \mathcal{P}_B)$ *is communication-enabled with respect to* E if, for each $n \geq 0$, each $\pi \in \mathcal{P}_A \vee \mathcal{M}_B^{(n)}$ and each $\pi \in \mathcal{P}_B \vee \mathcal{M}_A^{(n)}$ is regular with respect to E.

These definitions capture the idea that if the players communicate their probabilities at any level, the calculations they then make are meaningful. Notice that the condition does not require that the players actually communicate, only that they would be able to communicate meaningfully if they decided to do so. We come back to this point below.

Remark 2. Fix $\pi, \pi' \subseteq \Omega$ with $\pi \cap \pi' = \emptyset$. Then if π and π' are both regular with respect to E, so is $\pi \cup \pi'$.

From now on we will assume that all members of $\mathcal{P}_A$ and $\mathcal{P}_B$ get strictly positive probability. This ensures that the agents are able to observe and condition on their own information.

Proposition 2. *Fix a structure $(\Omega, p, \mathcal{P}_A, \mathcal{P}_B)$ 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 $q_A = q_B$.*

[2]Notice that all of Alice's potential announcements are made relative to her (initial) partition $\mathcal{P}_A$; and likewise for Bob. This is different from Geanakolpos and Polemarchakis (1982) where the (actual) announcements are relative to partitions which are successively updated based on previous announcements.

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

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

from which $p(B_N \mid \mathcal{P}_A(\omega)) = 1$ 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 an index set (finite). We just saw that $p(B_N \mid \pi_i) = 1$ for all such π_i. But $p(B_N \mid A_N)$ is a convex combination of the $p(B_N \mid \pi_i)$'s, so $p(B_N \mid A_N) = 1$. It follows that $p(A_N \backslash B_N) = 0$, which we will use shortly.

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

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 $\{A_N \backslash B_N, (A_N \backslash B_N)^c\}$ is a coarsening of $\mathcal{P}_A \vee \mathcal{M}_B^{(N)}$. By the hypothesis of the proposition and Remark 2, it follows that $A_N \backslash B_N$ is regular with respect to E. Using $p(A_N \backslash B_N) = 0$, it follows that $p((A_N \backslash B_N) \cap E) = 0$, and so $p(E \cap A_N \cap B_N) = p(E \cap A_N)$. Again using $p(A_N \backslash B_N) = 0$, 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 $p(E \mid A_N \cap B_N) = p(E \mid A_N) = q_A$. We can run exactly the same argument with B in place of A to conclude that $p(E \mid A_N \cap B_N) = q_B$. It follows that $q_A = q_B$. $\square$

This result says that common certainty of disagreement is possible only when the agents are unable to communicate their probabilities meaningfully. We can think of meaningful communication as necessary if Alice and Bob want to confirm their disagreement — as opposed to just being certain of disagreement (and certain they are certain, etc.). Interestingly, the mere ability to confirm their disagreement is enough to rule out common certainty of disagreement between Alice and Bob. As we said in the Introduction, a signature of non-classicality is that a system is affected by the mere availability of information – the information need not be observed.

A corollary to Proposition 2 is that common certainty of disagreement is impossible in the classical world. This follows because the condition of being communication-enabled is automatically satisfied in the case of unsigned 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 7. 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 meaningful calculations with the information which announcements by Alice and Bob would give him. (Again, Alice and Bob do not actually have to undertake this communication.) Alice and Bob do not communicate with each other.

Proposition 3. *Fix a structure* $(\Omega, p, \mathcal{P}_A, \mathcal{P}_B)$ *that is third-party 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* $q_A = q_B$*.*

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

5 Related Work

Our findings help clarify a statement in the Sudbery (2017) comment on Frauchiger and Renner (2018). Sudbery writes "[T]here is no logical reason why statements existing in different perspectives should be consistent (think of statements about the order of events in different frames of reference, in special relativity)." We show that certain types of inconsistency of statements existing in different perspectives are, in fact, impossible. First, in classical settings, if the statements concern probabilistic assessments of some underlying event in common, then what we called singular disagreement is impossible (Remark 1). Second, even in non-classical settings, both common knowledge of disagreement (Proposition 1) and common certainty of disagreement under communication (Propositions 2 and 3) are impossible, which imposes a clear limit on the way in which agents can disagree.

As discussed in the Introduction, Frauchiger and Renner (2018) promote singular disagreement to a formal contradiction. Brukner (2018) proves a no-go theorem by articulating an analogous assumption he calls "observer-independent facts" and which states that "One can jointly assign truth values to the propositions about observed outcomes (facts) of different observers ...." As in Brukner, our formalism would not justify this move except as an additional ingredient.[3]

Khrennikov and Basieva (2015) and Khrennikov (2015) study extensions of the classical Agreement Theorem to settings where the agents process information in a non-Bayesian and quantum-like way, while the substrate remains classical. This is the opposite of our approach

[3]Effectively, certainty would have to become knowledge.

(or that in Frauchiger and Renner [2018]). Our use of signed probabilities is to capture the non-classicality of the physical substrate, not a departure of the agents' information processing from Bayesianism.

Our results establish, on the one hand, a new kind of non-classical strangeness in the form of the possibility of common certainty of disagreement. On the other hand, by proving that both common knowledge of disagreement and common certainty of disagreement under communication are impossible, even in non-classical settings, we also establish a basic consistency of the non-classical world.

References

Abramsky, S., and A. Brandenburger, "The Sheaf-Theoretic Structure of Non-Locality and Contextuality," *New Journal of Physics*, 13, 2011, 113036., at https://doi.org/10.1088/1367-2630/13/11/113036.

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

Aumann, R., and A. Brandenburger, "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63, 1995, 1161-1180.

Brukner, Č., "A No-Go Theorem for Observer-Independent Facts," *Entropy*, 20, 2018, at https://doi.org/10.3390/e20050350.

Dirac, P., "The Physical Interpretation of Quantum Mechanics," *Proceedings of the Royal Society of London* (Series A: Mathematical and Physical Sciences), 180, 1942, 1-40.

Feynman. R., "Negative Probability," in B. Hiley and F. Peat (eds.), *Quantum Implications: Essays in Honour of David Bohm*, Routledge and Kegan Paul, 1987, 235-248.

Frauchiger, D., and R. Renner, "Quantum Theory Cannot Consistently Describe the Use of Itself," *Nature Communications*, 9, 2018, 3711, at https://doi. org/10.1038/s41467-018-05739-8 (2018).

Geanakoplos, J., and H. Polemarchakis, "We Can't Disagree Forever," *Journal of Economic Theory*, 28, 1982, 192-200.

Khrennikov, A., “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, at https://doi.org/10.1016/j.jmateco.2015.06.018.

Khrennikov, A., and I. Basieva, “Quantum(-Like) Decision Making: On Validity of the Aumann Theorem,” in H. Atmanspacher, C. Bergomi, T. Filk, and K. Kitto (eds.), *Quantum Interaction*, QI 2014, Lecture Notes in Computer Science, 8951, Springer, 2015, at https://doi.org/10.1007/978-3-319-15931-7_9.

Pusey, M., “An Inconsistent Friend,” *Nature Physics*, 14, 2018, 973-978, at https://doi.org/10.1038/s41567-018-0293-7.

Sudbery, A., “Single-World Theory of the Extended Wigner’s Friend Experiment,” *Foundations of Physics*, 2017, 47, 658-669, at https://doi.org/10.1007/s10701-017-0082-7.

Wigner, E., “On the Quantum Correction for Thermodynamic Equilibrium,” *Physical Review*, 40, 1932, 749-759, at https://doi:10.1103/PhysRev.40.749.

Wigner, E., “Remarks on the Mind-Body Question,” in I. Good (ed.), *The Scientist Speculates*, Heinemann, 1961.

Wooters, W., “A Wigner-Function Formulation of Finite-State Quantum Mechanics,” *Annals of Physics*, 176, 1987, 1-21, at https://doi.org/10.1016/0003-4916(87)90176-X.