Noname manuscript No.
(will be inserted by the editor)

The Ryōan-ji Axiom for Common Knowledge on Hypergraphs

Jeffrey Kane · Pavel Naumov

the date of receipt and acceptance should be inserted later

Abstract The article studies common knowledge in communication networks with a fixed topological structure. It introduces a non-trivial principle, called the Ryōan-ji axiom, which captures logical properties of common knowledge of all protocols with a given network topology. A logical system, consisting of the Ryōan-ji axiom and two additional axioms, is proven to be sound and complete.

Keywords common knowledge, hypergraph, axiomatization, completeness

1 Introduction

In this article, we investigate properties of common knowledge between groups of agents. As usual, common knowledge of φ means each agent knows φ, each agent knows that each agent knows φ, and so on ad infinitum. As an example, consider a situation where two intelligence agencies, called A and B, are operating in an island country in which people are not permitted to travel between the islands. Suppose that the island government observes all communication between the islands but does not observe intra-island communication. Assume that operatives of agencies A and B eavesdrop on some of the intra-island channels. Furthermore, suppose both of the agencies and the government know to which intra-island channels each agency has access. We are interested in the question of whether the two agencies can have common knowledge of any intelligence data unknown to the government.

If there is at least one island where both agencies have operatives eavesdrop on some of the channels, then it *might* happen that the agencies will have common knowledge of intelligence data not known to the government.

Department of Mathematics and Computer Science, McDaniel College
Westminster, Maryland 21157, the United States

For example, this is the case when operatives eavesdrop on the same intra-island channel or two different intra-island channels transmitting the same information.

The situation changes if the two agencies do not eavesdrop on channels located on the same island. Of course, an operative could encrypt data with a one-time encryption pad then transmit it over inter-island channels to another agent. We do not consider this acceptable since it relies on the two agents having a pre-existing common knowledge of the one-time encryption pad. We also exclude public key encryption since it can only delay the moment when the government learns the intelligence data. In this article we formalize the setting above and prove that in our model no common knowledge of intelligence data can exist between the two agencies without government knowing it as well. Later in this article, we prove a stronger claim about this situation: if any data is common knowledge between the agencies A and B, then the same data is also common knowledge between the two agencies and the government G. We write this property of common knowledge as

$$[A,B] \to [A,B,G]. \tag{1}$$

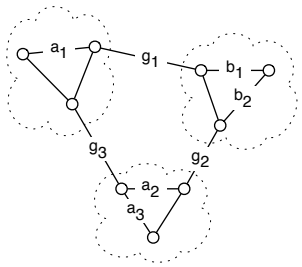

Fig. 1 $A = \{a_1, a_2, a_3\}$, $B = \{b_1, b_2\}$, and $G = \{g_1, g_2, g_3\}$. Dotted lines show the borders of the islands.

Alternatively, one can write the above principle as $[A,B]\varphi \to [A,B,G]\varphi$ to show that any common knowledge φ of A and B is also the common knowledge of A, B, and G. In this article, however, we restrict ourselves to studying common knowledge of a fixed single φ, thus we use notation $[A,B]$ instead of $[A,B,G]\varphi$. The more general setting of standard epistemic modal logic that allows for different propositions phi and for nesting (agent a knows that agent b knows that φ) will be briefly discussed in the conclusion.

In a formal model, one can represent communication channels as edges of a graph. In this setting, a vertex of the graph enforces dependencies that might exist between messages on different channels incident to the vertex. Intelligence agencies and the government are represented by the set of channels that they can eavesdrop on.

A specific example of the general principle (1) for the graph depicted on Figure 1 is

$$[\{a_1, a_2, a_3\}, \{b_1, b_2\}] \to [\{a_1, a_2, a_3\}, \{b_1, b_2\}, \{g_1, g_2, g_3\}].$$

We say that a set of edges G is a gateway between the sets of edges A and B if for every edge $a \in A$ and every edge $b \in B$, each path between a and b contains some edge in the set G. In our example, the set $\{g_1, g_2, g_3\}$ is a gateway between the sets $\{a_1, a_2, a_3\}$ and $\{b_1, b_2\}$. In general, the principle (1) is true for any gateway G between arbitrary sets of edges A and B (see Proposition 4). We refer to this principle as the *gateway principle*.

Note that while the gateway principle is valid for common knowledge, it is not valid for "mutual" knowledge of the group (knowledge by all agents in the group). Indeed, if an operative on one island knows φ and another operative on a different island knows ψ, then they mutually know $\varphi \vee \psi$, but the government does not have to know $\varphi \vee \psi$.

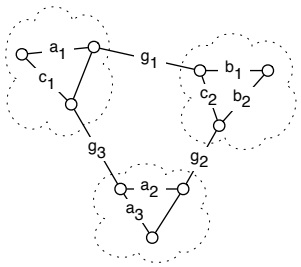

Fig. 2 $A = \{a_1, a_2, a_3\}$, $B = \{b_1, b_2\}$, $C = \{c_1, c_2\}$, and $G = \{g_1, g_2, g_3\}$.

Assume now that three intelligence agencies, called A, B, and C, operate in the same island country. We want to see if they can have common knowledge between all three of them of intelligence data not known to the government. As before, this trivially might be possible if all three agencies have operatives eavesdrop on intra-island channels of the same island. Assume however, as in Figure 2, that agencies A and C share an island and agencies B and C do as well, but there is no island where agencies A and B both operate. Furthermore, each agency operative can communicate to the agency but cannot receive a reply from the agency. In other words, operatives cannot share information through their agency headquarters[1]. In this case the three agencies cannot have common knowledge of intelligence data unknown to the government. Indeed, assume $[A, B, C]$. Thus, $[A, B]$ since common knowledge between any group of agencies implies common knowledge between any of its subgroup. Hence, by the gateway principle, $[A, B, G]$. In other words, the intelligence data must

[1] Otherwise, two operatives of agency C on different islands could easily form common knowledge amongst A, B, and C.

be a common knowledge between agencies A and B and the government G. Therefore, the data is known to the government. The above argument shows that

$$[A, B, C] \to [A, B, G], \tag{2}$$

where G is a gateway between sets of edges A and B.

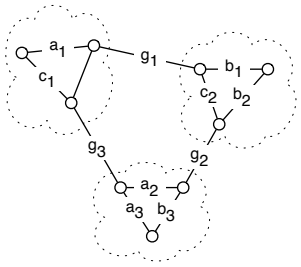

Fig. 3 $A = \{a_1, a_2, a_3\}$, $B = \{b_1, b_2, b_3\}$, $C = \{c_1, c_2\}$, and $G = \{g_1, g_2, g_3\}$.

The gateway principle is an intuitively expected statement that captures a topological property of common knowledge. However, there are topological properties of common knowledge that do not follow from this principle. Consider the example depicted in Figure 3, where agencies A and C are present on one island, agencies B and C on another island, and agencies A and B on a third island. As we show in Lemma 7, even in this setting any common knowledge between agencies A, B, and C must also be known to the government. This fact, however, does not follow from the gateway principle, since in this case G is not a gateway between any pair of agencies.

The main contribution of this article is the discovery of a topological property of common knowledge more general than the gateway principle. We call it the Ryōan-ji principle. Ryōan-ji is a famous zen garden in Kyoto, Japan that contains fifteen rocks positioned in such a way that from any point in the garden one can see at most fourteen of these rocks. The legend says that only through enlightenment one would be able to see all fifteen rocks at the same time. In the case of the example on Figure 3 the Ryōan-ji principle can be stated as

$$[A, B, C] \to [A, B, C, G].$$

Similar to the Ryōan-ji garden, from any edge in the graph depicted in Figure 3 one can not "see" at least one of the sets A, B, or C without "looking through" the set G. Of course, in the garden, the rocks themselves obstruct each other from view, while in our principle, the set G "obstructs" the sets A, B, and C. The general form of the Ryōan-ji principle is given in Section 4.

The main technical result of this article is a sound and complete axiomatization of the topological properties of common knowledge. The axiomatic

system consists of the Ryōan-ji principle in its general form, the Omniscience axiom and the Subgroup axiom that also will be stated later.

Lewis (1969) first studied the idea of common knowledge in the context of conventions. Halpern and Moses (1990); Fagin et al (1995) described modal properties of common knowledge. More recently, Studer (2009) has shown that modal logic of common knowledge does not have the Beth property. Our work is significantly different from the modal logic approach both in the language and in the content. We do not treat common knowledge as a modality that can be applied to different statements but instead focus on common knowledge of the same statement by various groups of agents connected by a communication network with a given topology. Various information flow properties of such networks have been studied before by More and Naumov (2011c,b,a); Donders et al (2011); Kane and Naumov (2013). The closest work to our current article probably is a study by Holbrook and Naumov (2012) of fault tolerance in belief formation networks, where they introduced a "shield-wall" principle capturing logical properties of such networks. However, the shield-wall principle differs significantly from the Ryōan-ji principle. In the shield-wall principle, unlike the Ryōan-ji principle, all vertices must be completely blocked by the combination of the shields.

This article is structured as follows. In Section 2, we review relevant graph theory terminology. Next, in Section 3, we introduce the formal syntax and semantics for our logical system. This is followed by sections 4 and 5 which contain axioms and examples of formal proofs in the system. In sections 6 and 7, we prove the soundness and the completeness of the system.

2 Graph Theory Terminology

In the introduction we assumed that communication between agents happens along the edges of a graph. In the rest of this work, we consider a more general setting in which messages can be shared between more than two parties. Thus, we will allow graphs in which edges might have an arbitrary number of ends. Such graphs are commonly referred to as hypergraphs. Formally, a hypergraph is defined as a pair $H = (V, E)$, where V is as an arbitrary finite set and E is an arbitrary set of subsets of V. Elements of the set V are called "vertices", and elements of the set E are called "edges" of the hypergraph H. Although it is common not to allow empty edges (Berge, 1989), this restriction is not significant in our case. An example of a hypergraph is depicted in Figure 4.

By a path in a hypergraph we mean any sequence of edges such that any two adjacent edges in the path share at least one vertex. By a path from an edge e to an edge f we mean any path whose first edge is e and whose last edge is f. For example, a, b, c is a path from edge a to edge c in the hypergraph depicted in Figure 4. By $Inc(v)$ we refer to the incident edges to v, the set of all edges containing vertex v.

The next definition is one of the key definitions of this article. It is used to state the general form of the Ryōan-ji principle.

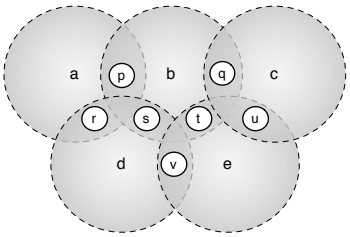

Fig. 4 Hypergraph (V, E), where $V = \{p, q, r, s, t, u, v\}$ and $E = \{a, b, c, d, e\}$.

Definition 1 For any family of sets of edges $\{A_i\}_{i\le n}$ of a hypergraph H and any set of edges O of H, we say that O *obstructs* the family of sets $\{A_i\}_{i\le n}$ if for each edge e of the hypergraph there exists an $i \le n$ such that for each edge $a \in A_i$, each path from e to a contains at least one edge from the set O.

For example, the set G obstructs the family of sets $\{A, B\}$ in the hypergraph depicted in Figure 1 and the set G obstructs the family of sets $\{A, B, C\}$ in the hypergraphs depicted in Figure 2 and Figure 3. The set $\{b\}$ does not obstruct the family of sets $\{\{a\}, \{c\}\}$ on the hypergraph depicted in Figure 4 due to existence of the path a, d, e, c. However, the set $\{b, d\}$ *does* obstruct the family of sets $\{\{a\}, \{c\}\}$ on the same hypergraph.

We conclude this section with a technical lemma involving this definition. This lemma is used in Section 5.

Lemma 1 *The family of sets of edges $\{A_i\}_{i\le n}$ of a hypergraph $H = (V, E)$ is obstructed by any set $O \subseteq E$ if A_i is empty for some i.* □

3 Syntax and Semantics

Let us introduce the following syntax that can be seen as the syntactic fragment of standard epistemic modal logic with only false, common knowledge of a fixed specific statement, and implications.

Definition 2 For any hypergraph $H = (V, E)$, let $\Phi(H)$ be the minimal set such that

1. $\bot \in \Phi(H)$,
2. $[A_1, \dots, A_n] \in \Phi(H)$, for any $n \ge 0$ and any distinctive subsets $A_1, \dots, A_n$ of set E,
3. $\varphi \to \psi \in \Phi(H)$ for each $\varphi, \psi \in \Phi(H)$.

Statement $[A, ..., A_n]$ is read "there is common knowledge of the fixed statement by the group of n agents observing subsets of edges A_1, A_2, ..., and A_n". The content inside the common knowledge brackets $[A_1, \dots, A_n]$ is treated as a set. Thus, for example, we do not distinguish $[A_1, A_2]$ from $[A_2, A_1]$. If $X = \{A_1, \dots, A_n\}$ and $Y = \{B_1, \dots, B_k\}$ are two families of subsets of E, we

will sometimes write $[A_1, \ldots, A_n, B_1, \ldots, B_k]$ to mean $[X \cup Y]$. In other words, $[A, B, B, C]$ means $[A, B, C]$.

We now define a protocol over hypergraph. With the exception of the "knowledge" component K, it is the definition previously used by More and Naumov (2011b,a). Knowledge is often represented by a proposition, but here it is more convenient to think about knowledge K as the set of "runs" of the protocol where this proposition is true.

Definition 3 A protocol over a hypergraph (V, E) is $(\{W_e\}_{e \in E}, \{L_v\}_{v \in V}, K)$ such that

1. W_e is an arbitrary set of "values" of edge $e \in E$,
2. $L_v \subseteq \prod_{e \in Inc(v)} W_e$ is a "local condition" at vertex $v \in V$,
3. $K \subseteq \prod_{e \in E} W_e$ is the "knowledge".

A value in the set W_e represents information which could be communicated over edge e. The local condition L_v captures dependencies between the values of edges in $Inc(v)$ that are "enforced" by vertex v. In the next definition we formally introduce the notion of a run of a protocol as a combination of values of all edges that satisfies local conditions at each vertex.

Definition 4 A run of a protocol $P = (\{W_e\}_{e \in E}, \{L_v\}_{v \in V}, K)$ over a hypergraph (V, E) is an arbitrary tuple $\langle r_e \rangle_{e \in E}$ such that $r_e \in W_e$ for each $e \in E$ and

$$\langle r_e \rangle_{e \in Inc(v)} \in L_v \tag{3}$$

for each $v \in V$.

We say that tuple $\langle r_e \rangle_{e \in E}$ satisfies the local conditions at vertex v if statement (3) is true. The (possibly empty) set of all runs of a protocol P is denoted $R(P)$. If $r = \langle r_e \rangle_{e \in E}$ and $r' = \langle r'_e \rangle_{e \in E}$, then by $r \equiv_A r'$ we mean that $r_a = r'_a$ for all $a \in A$. Note that each run can be viewed as an epistemic world in a Kripke model and $\equiv_A$ is the indistinguishability relation of an agent observing only the edges in set A. In this article we allow protocols that do not have runs. In Kripke model terms it means that we allow frames with empty set of worlds. Restricting protocols only to those that have at least one run will not affect any of the results in this article.

Consider the statement "any agent observing the set of channels A on the run r knows that the proposition φ is true". It can be rephrased as "φ is true on any run r^1 such that $r^1 \equiv_A r$". Similarly, the statement "any agent observing the set of channels A on run r knows that any agent observing the set of channels B on the same run r knows that φ is true" can be rephrased as "φ is true on any run r^2 such that $r^2 \equiv_B r^1 \equiv_A r$, where r^1 is an arbitrary run".

Following the same pattern, the statement "φ is common knowledge on run r between any two agents observing the sets of channels A and B, respectively" can be rephrased as "for any $X_1, X_2, \ldots, X_m \in \{A, B\}$ and any runs $r^0, r^1, \ldots, r^m$, if $r = r^0 \equiv_{X_1} r^1 \equiv_{X_2} \cdots \equiv_{X_n} r^m$, then on run r^m proposition φ is true".

Thus, the statement "$r \in K$ is common knowledge on run r between any two agents observing the sets of channels A and B, respectively" can be rephrased as "for any $X_1, X_2, \dots, X_m \in \{A, B\}$ and any runs $r^0, r^1, \dots, r^m$, if $r = r^0 \equiv_{X_1} r^1 \equiv_{X_2} \cdots \equiv_{X_n} r^m$, then $r^m \in K$". Part 2 of the next definition states this in a more general setting with multiple agents. This definition is the standard definition of $\vDash$ on Kripke models rephrased in terms of a protocol and its runs.

Definition 5 For any run r of a protocol $P = (\{W_e\}_{e\in E}, \{L_v\}_{v\in V}, K)$ over a hypergraph $H = (V, E)$, and any formula $\varphi \in \Phi(H)$, the relation $r \vDash_P \varphi$ is defined recursively on φ as follows:

1. $r \nvDash_P \bot$,
2. $r \vDash_P [A_1, \dots, A_n]$ means that for all $m \geq 0$, all $r^0, r^1, \dots, r^m \in R(P)$, and all $i_1, \dots, i_m \leq n$, if $r = r^0 \equiv_{A_{i_1}} r^1 \equiv_{A_{i_2}} \cdots \equiv_{A_{i_m}} r^m$, then $r^m \in K$,
3. $r \vDash_P \varphi \to \psi$ if $r \nvDash_P \varphi$ or $r \vDash_P \psi$.

We omit the subscript P from the expression $r \vDash_P \varphi$ when its value is clear from the context.

Note that $r \vDash [A_1]$ means that $r^1 \in K$ for each run r^1 such that $r^1 \equiv_{A_1} r$. Thus, going back to the introductory examples, $[G]$ is a statement that the government knows the intelligence data. Hence, the statement "if A and B have common knowledge of intelligence data, then this data must be known to the government G" can be expressed in our language as $[A, B] \to [G]$.

An interesting case of the common knowledge statement $[A_1, \dots, A_n]$ occurs when $n = 0$. It is easy to see from Definition 5 that $r \vDash [\ ]$ is equivalent to $r \in K$. Thus, $[\ ]$ is simply a statement that the knowledge is true. Note that if $r \in K$, then any agent observing the set of all edges E of the hypergraph H on run r will know that $r \in K$. Therefore, it is always true that $[\ ] \to [E]$. The last statement is our Omniscience axiom.

Another interesting formula is $[\varnothing]$. Since $r \equiv_\varnothing r'$ for any two runs r and r' of any protocol P, statement $r \vDash [\varnothing]$ means that $r' \in K$ is true on any run r' of the protocol P. This, of course, is different from the previously discussed statement $r \vDash [\ ]$.

Finally, note also that if the set A_1 contains several channels, then $[A_1]$ could also be interpreted as "the set of parties observing individual channels in the set A_1 have distributed knowledge of K". The statement $[A_1, \dots, A_n]$, could similarly be interpreted as a combination of common knowledge and distributed knowledge concepts.

The following proposition illustrates Definition 3 and Definition 5.

Proposition 1 *There exists a protocol over hypergraph $H_0 = (V, E)$ depicted in Figure 5 and a run r of this protocol such that*

$$r \nvDash [\{a\}, \{b\}, \{d\}] \to ([\{a\}, \{c\}, \{d\}] \to [\{b\}, \{c\}]).$$

Proof Let P be the triple $(\{W_e\}_{e\in E}, \{L_v\}_{v\in V}, K)$ such that

1. $W_a = W_d = \{0, 1\}^2$,

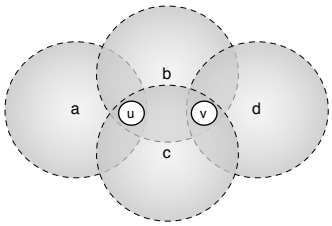

Fig. 5 Hypergraph $H_0 = (V, E)$. $V = \{u, v\}$. $E = \{a, b, c, d\}$.

2. $W_b = W_c = \{0, 1\}$,
3. $L_u = \{\langle\langle s_1, s_2\rangle, t, p\rangle \in W_a \times W_b \times W_c \mid s_1 = t \text{ and } s_2 = p\}$,
4. $L_w = \{\langle t, p, \langle q_1, q_2\rangle\rangle \in W_b \times W_c \times W_d \mid t = q_1 \text{ and } p = q_2\}$,
5. $K = \{\langle\langle s_1, s_2\rangle, t, p, \langle q_1, q_2\rangle\rangle \in W_a \times W_b \times W_c \times W_d \mid t = 0 \text{ or } p = 0\}$.

Triple P is a protocol by Definition 3. Let run r of the protocol P be the tuple $\langle\langle 0, 0\rangle, 0, 0, \langle 0, 0\rangle\rangle$.

Lemma 2 $r \vDash [\{a\}, \{b\}, \{d\}]$.

Proof Consider any $x_1, \ldots, x_m \in \{a, b, d\}$ and any runs $r^0, r^1, \ldots, r^m$ of the protocol P such that $r = r^0 \equiv_{\{x_1\}} r^1 \equiv_{\{x_2\}} \cdots \equiv_{\{x_m\}} r^m$. Note that any two runs equivalent on a are also equivalent on b due to the local condition L_u. Similarly, any two runs equivalent on d are also equivalent on b due to the local condition L_w. Thus, $r = r^0 \equiv_{\{b\}} r^1 \equiv_{\{b\}} \cdots \equiv_{\{b\}} r^m$. Hence, $r^m = \langle\langle 0, y\rangle, 0, y, \langle 0, y\rangle\rangle$ for some $y \in \{0, 1\}$. Therefore, $r^m \in K$ by the choice of K. □

Lemma 3 $r \vDash [\{a\}, \{c\}, \{d\}]$.

Proof The argument is similar to the proof of Lemma 2. □

Lemma 4 $r \nvDash [\{b\}, \{c\}]$.

Proof Notice $r = \langle\langle 0, 0\rangle, 0, 0, \langle 0, 0\rangle\rangle \equiv_{\{b\}} \langle\langle 0, 1\rangle, 0, 1, \langle 0, 1\rangle\rangle \equiv_{\{c\}} \langle\langle 1, 1\rangle, 1, 1, \langle 1, 1\rangle\rangle$. Thus $r \nvDash [\{b\}, \{c\}]$, as $\langle\langle 1, 1\rangle, 1, 1, \langle 1, 1\rangle\rangle \notin K$. □

We are ready to conclude the proof of the proposition. By Lemma 2, Lemma 3, and Lemma 4, $r \nvDash [\{a\}, \{b\}, \{d\}] \to ([\{a\}, \{c\}, \{d\}] \to [\{b\}, \{c\}])$. □

4 Axioms

For any hypergraph $H = (V, E)$, in addition to propositional tautologies in the language $\Phi(H)$ and the Modus Ponens inference rule, our logical system contains the following axioms:

1. Omniscience: $[\ ] \to [E]$,

2. Subgroup: $[A_1,\dots,A_n]\to[A_1,\dots,A_{n-1}]$, where $n\geq 1$,
3. Ryōan-ji: $[A_1,\dots,A_n]\to[A_1,\dots,A_n,O]$, where the set of edges O of the hypergraph H obstructs the family of sets of edges $\{A_i\}_{i\leq n}$.

We write $X\vdash_H\varphi$ if formula φ is provable in our system extended by an additional set of axioms X. We write $\vdash_H\varphi$ instead of $\varnothing\vdash_H\varphi$.

5 Examples

In this section, we give several examples of formal proofs in our logical system. We prove soundness and completeness of the logical system in the next two sections.

As we discussed in the introduction, any common knowledge between agencies A and B must also be known by the government G on the hypergraph H_1 shown in Figure 1. This statement, as mentioned in Section 3, can be expressed in our language as $[A,B]\to[G]$. We formally prove this statement as our first example.

Proposition 2 $\vdash_{H_1}[\{a_1,a_2,a_3\},\{b_1,b_2\}]\to[\{g_1,g_2,g_3\}]$.

Proof Recall from Section 2, $\{g_1,g_2,g_3\}$ obstructs $\{\{a_1,a_2,a_3\},\{b_1,b_2\}\}$. Thus, by the Ryōan-ji axiom,

$$\vdash_{H_1}[\{a_1,a_2,a_3\},\{b_1,b_2\}]\to[\{a_1,a_2,a_3\},\{b_1,b_2\},\{g_1,g_2,g_3\}].$$

Thus, $\vdash_{H_1}[\{a_1,a_2,a_3\},\{b_1,b_2\}]\to[\{g_1,g_2,g_3\}]$, by the Subgroup axiom and the laws of the propositional logic. □

In the introduction, we discussed that any common knowledge between agencies A, B, and C is also common knowledge between A, B, and the government G for the hypergraph H_2 depicted in Figure 2. This statement is formally captured by equation (2). We prove an even stronger claim below.

Proposition 3

$$\vdash_{H_2}[\{a_1,a_2,a_3\},\{b_1,b_2\},\{c_1,c_2\}]\to[\{a_1,a_2,a_3\},\{b_1,b_2\},\{c_1,c_2\},\{g_1,g_2,g_3\}].$$

Proof Recall $\{g_1,g_2,g_3\}$ obstructs $\{\{a_1,a_2,a_3\},\{b_1,b_2\},\{c_1,c_2\}\}$, as discussed in Section 2. Therefore,

$$\vdash_{H_2}[\{a_1,a_2,a_3\},\{b_1,b_2\},\{c_1,c_2\}]\to[\{a_1,a_2,a_3\},\{b_1,b_2\},\{c_1,c_2\},\{g_1,g_2,g_3\}],$$

by the Ryōan-ji axiom. □

We now state and prove the gateway principle discussed in the introduction.

Proposition 4 $\vdash_H[A,B]\to[A,B,G]$, *for any hypergraph H and any sets of edges A, B, and G of the hypergraph H such that for every edge $a\in A$ and every edge $b\in B$, each path between a and b contains some edge in G.*

Proof Suppose G does not obstruct the family of sets $\{A, B\}$. Thus, by Definition 1, there exists an edge e of H such that there is a path from e to an edge in set A and a path from e to an edge in set B, where neither path contains an edge in G. These two paths can be combined into a path from A to B containing no edges from G, in contradiction with the assumption of the proposition. Hence, G must obstruct $\{A, B\}$. Therefore, $\vdash_H [A, B] \to [A, B, G]$, by the Ryōan-ji axiom. □

According to Definition 2, any set inside the common knowledge brackets could be empty. As we have discussed above, $[\varnothing]$ means that the knowledge is true on every run. Thus, as shown in the next proposition, there is common knowledge of this among an arbitrary set of agents.

Proposition 5 $\vdash_H [\varnothing] \to [A_1, \ldots, A_n]$*, where* $A_1, \ldots, A_n$ *are arbitrary sets of edges of hypergraph* H*.*

Proof Due to the Subgroup axiom, it suffices to prove $\vdash_H [\varnothing] \to [\varnothing, A_1, \ldots, A_n]$ for any $n \geq 0$. We show this by induction on n. If $n = 0$, then $[\varnothing] \to [\varnothing]$ is a propositional tautology. Assume $\vdash_H [\varnothing] \to [\varnothing, A_1, \ldots, A_{n-1}]$. Then by Lemma 1, A_n obstructs the family of sets $\{\varnothing, A_1, \ldots, A_{n-1}\}$. Hence, by the Ryōan-ji axiom, $\vdash_H [\varnothing, A_1, \ldots, A_{n-1}] \to [\varnothing, A_1, \ldots, A_n]$. Therefore, $\vdash_H [\varnothing] \to [\varnothing, A_1, \ldots, A_n]$. □

The next proposition illustrates that any number of agents observing the same set of channels have common knowledge about anything known to any one of these agents. We use this proposition in the proof of completeness.

The following proposition captures a monotonicity property different from the one expressed in the Subgroup axiom.

Proposition 6 $\vdash_H [A_1, \ldots, A_{n-1}, A_n] \to [A_1, \ldots, A_{n-1}, A'_n]$*, where* $A_n \subseteq A'_n$ *and* H *is an arbitrary hypergraph.*

Proof Note that from each edge e of H, the path from e to an edge a in A_n contains at least one edge in A'_n since a is in A'_n. Thus A'_n obstructs any family of sets which contains A_n. In particular, A'_n obstructs $\{A_i\}_{i \leq n}$. Hence, by the Ryōan-ji axiom, $\vdash_H [A_1, \ldots, A_{n-1}, A_n] \to [A_1, \ldots, A_{n-1}, A_n, A'_n]$. Therefore, by the Subgroup axiom, $\vdash_H [A_1, \ldots, A_{n-1}, A_n] \to [A_1, \ldots, A_{n-1}, A'_n]$. □

6 Soundness

Theorem 1 *For any* $\varphi \in \Phi(H)$*, if* $\vdash_H \varphi$*, then* $r \vDash_P \varphi$ *for any run* r *of any protocol* P *over hypergraph* $H = (V, E)$*.*

Proof Soundness of propositional tautologies and the Modus Ponens inference rule is trivial. The next three lemmas establish soundness of the Omniscience, Subgroup, and Ryōan-ji axioms.

Lemma 5 (Omniscience) *If* $r \vDash_P [\ ]$*, then* $r \vDash_P [E]$*.*

Proof Suppose $r \nvDash_P [E]$. Thus, there are $r^0, r^1, \dots, r^m \in R(P)$ such that $r = r^0 \equiv_E r^1 \equiv_E \cdots \equiv_E r^m \notin K$. Hence, as E is the set of all edges, $r = r^0 = r^1 = \cdots = r^m \notin K$. Therefore, $r = r^0 \notin K$, which is a contradiction with the assumption $r \vDash_P [\ ]$. □

Lemma 6 (Subgroup) *If $r \vDash_P [A_1, \dots, A_n]$, then $r \vDash_P [A_1, \dots, A_{n-1}]$.*

Proof Let $r^1, \dots, r^m \in R(P)$ and $i_1, \dots, i_m \le n-1$ be such that

$$r = r^0 \equiv_{A_{i_1}} r^1 \equiv_{A_{i_2}} \cdots \equiv_{A_{i_m}} r^m.$$

Then $r^m \in K$ by the assumption of the lemma and Definition 5. □

Lemma 7 (Ryōan-ji) *For any protocol $P = (\{W_e\}_{e\in E}, \{L_v\}_{v\in V}, K)$ over a hypergraph $H = (V, E)$, if $r \vDash_P [A_1, \dots, A_n]$ and a set of edges O obstructs some family of sets of edges $\{A_i\}_{i\le n}$, then $r \vDash_P [A_1, \dots, A_n, O]$.*

Proof Due to Definition 5, it suffices to show that for any runs $r' = \langle r'_e\rangle_{e\in E}$ and $r'' = \langle r''_e\rangle_{e\in E}$, if $r' \equiv_O r''$, then there exists $m \ge 0$, runs $r^0, r^1, \dots, r^m$, and $t_1, \dots, t_m \le n$ such that $r' = r^0 \equiv_{A_{t_1}} r^1 \equiv_{A_{t_2}} \cdots \equiv_{A_{t_m}} r^m = r''$.

Consider the hypergraph $H' = (V, E \setminus O)$. Let m be the number of connected components in H' and $H_1 = (V_1, E_1), \dots, H_m = (V_m, E_m)$ be these connected components. By Definition 1, for each $j \le m$ there is $t_j \le n$ such that $E_j \cap A_{t_j} = \varnothing$. Let tuples $r^0 = \langle r^0_e\rangle_{e\in E}, \dots, r^m = \langle r^m_e\rangle_{e\in E}$ be defined as

$$r^j_e = \begin{cases} r'_e = r''_e & \text{if } e \in O,\\ r'_e & \text{if } e \in E_{j+1} \cup \cdots \cup E_m,\\ r''_e & \text{if } e \in E_1 \cup \cdots \cup E_j. \end{cases}$$

Note that $r^{j-1} \equiv_{A_{t_j}} r^j$ because $E_i \cap A_{t_j} = \varnothing$. Furthermore, $r^0 = r'$ and $r^m = r''$. Next, we show that $r^1, \dots, r^m$ are runs of the protocol P. It is sufficient to show that r^j satisfies the local conditions at each vertex $v \in V$. Indeed, take an arbitrary vertex v. Let ℓ be such that $v \in V_\ell$. Thus, $Inc(v) \subseteq E_\ell \cup O$. We handle the cases $j < \ell$ and $j \ge \ell$ separately. If $j < \ell$, then $r^j \equiv_{E_\ell \cup O} r'$. Hence, $r^j \equiv_{Inc(v)} r'$. Therefore, $\langle r'_e\rangle_{e\in Inc(v)} \in L_v$ implies $\langle r^j_e\rangle_{e\in Inc(v)} \in L_v$. If $j \ge \ell$, then $r^j \equiv_{E_\ell \cup O} r''$. Hence, $r^j \equiv_{Inc(v)} r''$. Therefore, $\langle r''_e\rangle_{e\in Inc(v)} \in L_v$ implies $\langle r^j_e\rangle_{e\in Inc(v)} \in L_v$. □

This concludes the proof of the theorem. □

7 Completeness

The completeness theorem is stated later as Theorem 2. For each formula φ not provable in our system, the proof of this theorem constructs a protocol P and a run r of this protocol such that $r \nvDash_P \varphi$. The protocol P will be defined as a composition of several atomic protocols $P(B, F)$. In this section we first define atomic protocols, followed by two different protocol composition operations. We then conclude by combining atomic protocols into the desired protocol P using these two operations.

7.1 Protocol $P(B, F)$

Definition 6 Let $H = (V, E)$ be an arbitrary hypergraph and B be an arbitrary subset of E. Let $(V_1, E_1), \ldots, (V_n, E_n)$ be the connected components of hypergraph $(V, E \backslash B)$ that contain at least one edge and F be such that $F = E_i$ for some $i \leq n$. Let protocol $P(B, F)$ be the triple $(\{W_e\}_{e \in E}, \{L_v\}_{v \in V}, K)$ where

1. $\{W_e\}_{e \in E}$ is the family of sets, such that
$$W_e = \begin{cases} \{0, 1\} & \text{if } e \in F, \\ \{0\} & \text{otherwise.} \end{cases}$$
2. $\{L_v\}_{v \in V}$ is the family of sets, such that
$$L_v = \left\{ \langle r_e \rangle_{e \in Inc(v)} \in \prod_{e \in Inc(v)} W_e \;\middle|\; r_{f_1} = r_{f_2} \text{ for each } f_1, f_2 \in Inc(v) \cap F \right\}.$$
3. K is the single-element set $\{\langle \rho^0_e \rangle_{e \in E}\}$, such that $\rho^0_e = 0$ for each $e \in E$.

For example, let $H = (V, E)$ be the hypergraph depicted in Figure 4, if $B = \{b, d\}$, then hypergraph $(V, E \setminus B)$ has two connected components: $\{a\}$ and $\{c, e\}$. Let F be the set $\{c, e\}$. The values of edges c and e under the protocol $P(B, F)$ are 0 or 1. All other edges have the value 0. Local condition L_u forces the values of edges c and e to be equal.

The unique element of the set K will be referred to as ρ^0. By ρ^1, we mean tuple $\langle \rho^1_e \rangle_{e \in E}$ where
$$\rho^1_e = \begin{cases} 1 & \text{if } e \in F, \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 8 $\rho^0, \rho^1 \in R(P(B, F))$. □

Lemma 9 $\rho^0 \neq \rho^1$.

Proof Due to Definition 6, component F contains at least one edge. □

Lemma 10 *ρ^0 and ρ^1 are the only runs of the protocol $P(B, F)$.*

Proof Suppose there is a run $r = \langle r_e \rangle_{e \in E}$ of the protocol $P(B, F)$ such that $r \neq \rho^0$ and $r \neq \rho^1$. By the choice of W_e, there must exist edges $f_0, f_1 \in F$ such that $r_{f_0} = 0$ and $r_{f_1} = 1$. As F is the set of edges of a connected component of $(V, E \setminus B)$, consider any path between f_0 and f_1 with each of its edges belonging to F. Due to the local conditions at each of the vertices shared by the adjacent edges of the path, the value of the run at each edge on this path is the same. Therefore, $r_{f_0} = r_{f_1}$, which is a contradiction with the assumptions that $r_{f_0} = 0$ and $r_{f_1} = 1$. □

Informally, the next lemma states that several intelligence agencies might have a common knowledge on run ρ^0 of the protocol $P(B, F)$ if and only if they all have operatives on island F.

Lemma 11 *If $A_1, \dots, A_n$ are any sets of edges, then $\rho^0 \vDash [A_1, \dots, A_n]$ if and only if $A_i \cap F \neq \varnothing$ for each $i \leq n$.*

Proof ($\Rightarrow$) : Assume $\rho^0 \vDash [A_1, \dots, A_n]$ and $A_i \cap F = \varnothing$ for some $i \leq n$. Note that $\rho^0 \equiv_{A_i} \rho^1$. Hence, $\rho^1 \in K$, by Definition 5 and the assumption $\rho^0 \vDash [A_1, \dots, A_n]$, which is a contradiction.
($\Leftarrow$) : Assume $\rho^0 \nvDash [A_1, \dots, A_n]$. Then by Definition 5, there is an $m \geq 0$ and $i_1, \dots, i_m \leq n$ such that $\rho^0 = r^0 \equiv_{A_{i_1}} r^1 \equiv_{A_{i_2}} \cdots \equiv_{A_{i_m}} r^m$ and $r^m \notin K$. By Lemma 10, $r^k \in \{\rho^0, \rho^1\}$ for every $k \leq m$. In particular, $r^m = \rho^1$ due to $r^m \notin K$. Thus, there must exist $1 \leq \ell \leq m$ such that $r^{\ell-1} = \rho^0$ and $r^\ell = \rho^1$. Hence, $\rho^0 \equiv_{A_{i_\ell}} \rho^1$. Therefore, $A_{i_\ell} \cap F = \varnothing$. $\square$

Building towards the proof of completeness, the next lemma shows that for any pair of atomic common knowledge statements, one of them can be made true and the other false, as long as the former does not imply the latter in our logical system.

Lemma 12 *If $\nvdash_H [A_1, \dots, A_n] \to [B_1, \dots, B_m]$, where $H = (V, E)$ is an arbitrary hypergraph, then there exists $i \leq m$, a connected component $H' = (V', E')$ of the hypergraph $(V, E \setminus B_i)$, and a run r of a protocol $P(B_i, E')$ such that $r \vDash [A_1, \dots, A_n]$ and $r \nvDash [B_1, \dots, B_m]$.*

Proof Case 1: $n = 0$. Let us suppose first that $m = 1$ and $B_1 = E$. By the Omniscience axiom, $\nvdash_H [\ ] \to [E]$. which is a contradiction to the assumption $\nvdash_H [\ ] \to [B_1]$.

We may assume now that there is $e \in E \setminus B_i$ for some $i \leq m$. Let $H' = (V', E')$ be the connected component of the hypergraph $(V, E \setminus B_i)$ such that $e \in E'$. Note that $A_k \cap E' \neq \varnothing$ for all $k \leq n$ since $n = 0$. Consider run ρ^0 of the protocol $P(B_i, E')$. By Lemma 11, $\rho^0 \vDash [\ ]$. At the same time, $B_i \cap E' = \varnothing$ by the choice of H'. Hence, by Lemma 11, $\rho^0 \nvDash [B_1, \dots, B_m]$.

Case 2: $n > 0$. First, assume that there exists an $i \leq m$ such that B_i does not obstruct the family of sets $\{A_k\}_{k \leq n}$. In other words, there exists an edge $e \in E$ such that there is a path from each of the family of sets $\{A_k\}_{k \leq n}$ to e which does not go through edges in set B_i. Let $H' = (V', E')$ be the connected component of the hypergraph $(V, E \setminus B_i)$ such that $e \in E'$. Note that $A_k \cap E' \neq \varnothing$ for all $k \leq n$. Consider run ρ^0 of the protocol $P(B_i, E')$. By Lemma 11, $\rho^0 \vDash [A_1, \dots, A_n]$. At the same time, $B_i \cap E' = \varnothing$ by the choice of H'. Hence, by Lemma 11, $\rho^0 \nvDash [B_1, \dots, B_m]$.

Now, assume that B_i obstructs the family of sets of edges $\{A_k\}_{k \leq n}$ for every $i \leq m$. Thus, by Definition 1, the set B_i obstructs any family of sets of edges which contains subfamily $\{A_k\}_{k \leq n}$. Hence, by the Ryōan-ji axiom applied m times,

$$\vdash_H [A_1, \dots, A_n] \to [A_1, \dots, A_n, B_1],$$
$$\vdash_H [A_1, \dots, A_n, B_1] \to [A_1, \dots, A_n, B_1, B_2],$$
$$\dots$$
$$\vdash_H [A_1, \dots, A_n, B_1, \dots, B_{m-1}] \to [A_1, \dots, A_n, B_1, \dots, B_m].$$

Thus, by the laws of propositional logic,

$$\vdash_H [A_1, \ldots, A_n] \to [A_1, \ldots, A_n, B_1, \ldots, B_m].$$

By the Subgroup axiom,

$$\vdash_H [A_1, \ldots, A_n] \to [B_1, \ldots, B_m],$$

which is a contradiction with the assumption of the lemma. □

7.2 Protocol Compositions

For any family of protocols over the same hypergraph, one can consider a composition of these protocols in which all protocols are executed concurrently and independently of each other. Definition 3 stipulates that each protocol includes set K, which represents knowledge. For the composed protocol, knowledge could be understood as a conjunction or a disjunction of knowledges from the individual protocols. We denote the first type of composition by $\bigotimes$ and the second by $\bigoplus$. The proof of completeness uses both of these operations.

Definition 7 If $\{P^i\}_{0<i\leq\ell} = \{(\{W_e^i\}_{e\in E}, \{L_v^i\}_{v\in V}, K^i)\}_{0<i\leq\ell}$ is any (possibly empty) family of protocols over the same hypergraph (V, E), by $\bigotimes_{i\leq\ell} P^i$ we mean the protocol $(\{W_e\}_{e\in E}, \{L_v\}_{v\in V}, K)$ such that

1. W_e is the Cartesian product $\prod_{i\leq\ell} W_e^i$,
2. L_v is such that

$$L_v = \left\{ \langle\langle t_e^i\rangle_{i\leq\ell}\rangle_{e\in Inc(v)} \in \prod_{e\in Inc(v)} \prod_{i\leq\ell} W_e^i \;\middle|\; \langle t_e^i\rangle_{e\in Inc(v)} \in L_v^i \text{ for all } i \leq \ell \right\},$$

3. K is such that

$$K = \left\{ \langle\langle t_e^i\rangle_{i\leq\ell}\rangle_{e\in E} \in \prod_{e\in E} \prod_{i\leq\ell} W_e^i \;\middle|\; \langle t_e^i\rangle_{e\in E} \in K^i \text{ for all } i \leq \ell \right\}.$$

Lemma 13 *$\langle\langle r_e^i\rangle_{i\leq\ell}\rangle_{e\in E} \in R(\bigotimes_{i\leq\ell} P^i)$ if and only if $\langle r_e^i\rangle_{e\in E} \in R(P^i)$ for all $0 < i \leq \ell$.* □

The next lemma connects common knowledge under the composed protocol with common knowledge under the individual protocols.

Lemma 14 *If $\{P^i\}_{0<i\leq\ell} = \{(\{W_e^i\}_{e\in E}, \{L_v^i\}_{v\in V}, K^i)\}_{0<i\leq\ell}$ is a family of protocols over a hypergraph (V, E), then $\langle\langle r_e^i\rangle_{i\leq\ell}\rangle_{e\in E} \vDash_{\bigotimes_{i\leq\ell} P^i} [A_1, \ldots, A_n]$ if and only if $\langle r_e^i\rangle_{e\in E} \vDash_{P^i} [A_1, \ldots, A_n]$ for every $i \leq \ell$.*

Proof ($\Rightarrow$) : Suppose that there is a sequence of runs $\langle s^0_e\rangle_{e\in E}, \langle s^1_e\rangle_{e\in E}, \dots, \langle s^T_e\rangle_{e\in E}$ of the protocol P^q and $j_1, \dots, j_T \le n$ such that

$$\langle r^q_e\rangle_{e\in E} = \langle s^0_e\rangle_{e\in E} \equiv_{A_{j_1}} \langle s^1_e\rangle_{e\in E} \equiv_{A_{j_2}} \cdots \equiv_{A_{j_T}} \langle s^T_e\rangle_{e\in E} \notin K^q. \tag{4}$$

For each $0 < i \le \ell$, $e \in E$, and $0 \le t \le T$, consider $S^{i,t}_e$ such that

$$S^{i,t}_e = \begin{cases} s^t_e & \text{if } i = q,\\ r^i_e & \text{otherwise.}\end{cases} \tag{5}$$

By Lemma 13, $\langle\langle S^{i,t}_e\rangle_{i\le\ell}\rangle_{e\in E}$ is a run of the protocol $\bigotimes_{i\le\ell} P^i$ for each $t \le T$. Hence, due to (4) and (5),

$$\langle\langle r^i_e\rangle_{i\le\ell}\rangle_{e\in E} = \langle\langle S^{i,0}_e\rangle_{i\le\ell}\rangle_{e\in E} \equiv_{A_{j_1}} \langle\langle S^{i,1}_e\rangle_{i\le\ell}\rangle_{e\in E} \equiv_{A_{j_2}} \cdots \equiv_{A_{j_T}} \langle\langle S^{i,T}_e\rangle_{i\le\ell}\rangle_{e\in E}$$

$$\langle\langle S^{i,T}_e\rangle_{i\le\ell}\rangle_{e\in E} \notin \left\{ \langle\langle S^i_e\rangle_{i\le\ell}\rangle_{e\in E} \in \prod_{e\in E}\prod_{i\le\ell} W^i_e \;\middle|\; \langle S^i_e\rangle_{e\in E} \in K^i \text{ for all } i \le \ell \right\},$$

which is a contradiction with the assumption $\langle\langle r^i_e\rangle_{i\le\ell}\rangle_{e\in E} \vDash_{\bigotimes_{i\le\ell} P^i} [A_1, \dots, A_n]$.
($\Leftarrow$) : Suppose that $\langle\langle s^{i,0}_e\rangle_{i\le\ell}\rangle_{e\in E}, \langle\langle s^{i,1}_e\rangle_{i\le\ell}\rangle_{e\in E}, \dots, \langle\langle s^{i,T}_e\rangle_{i\le\ell}\rangle_{e\in E}$ is a sequence of runs of the protocol $\bigotimes_{i\le\ell} P^i$ and $j_1, \dots, j_T \le n$ are such that

$$\langle\langle r^i_e\rangle_{i\le\ell}\rangle_{e\in E} = \langle\langle s^{i,0}_e\rangle_{i\le\ell}\rangle_{e\in E} \equiv_{A_{j_1}} \langle\langle s^{i,1}_e\rangle_{i\le\ell}\rangle_{e\in E} \equiv_{A_{j_2}} \cdots \equiv_{A_{j_T}} \langle\langle s^{i,T}_e\rangle_{i\le\ell}\rangle_{e\in E}$$

$$\langle\langle s^{i,T}_e\rangle_{i\le\ell}\rangle_{e\in E} \notin \left\{ \langle\langle t^i_e\rangle_{i\le\ell}\rangle_{e\in E} \in \prod_{e\in E}\prod_{i\le\ell} W^i_e \;\middle|\; \langle t^i_e\rangle_{e\in E} \in K^i \text{ for all } i \le \ell \right\}.$$

Thus, there exists $q \le \ell$ such that

$$\langle r^q_e\rangle_{e\in E} = \langle s^{q,0}_e\rangle_{e\in E} \equiv_{A_{j_1}} \langle s^{q,1}_e\rangle_{e\in E} \equiv_{A_{j_2}} \cdots \equiv_{A_{j_T}} \langle s^{q,T}_e\rangle_{e\in E} \notin K^q,$$

which is a contradiction with the assumption $\langle r^q_e\rangle_{e\in E} \vDash_{P^q} [A_1, \dots, A_n]$. □

Next, we define the other type of protocol composition and prove its basic properties.

Definition 8 If $\{P^i\}_{0<i\le\ell} = \{(\{W^i_e\}_{e\in E}, \{L^i_v\}_{v\in V}, K^i)\}_{0<i\le\ell}$ is any (possibly empty) family of protocols over the same hypergraph (V, E), then by $\bigoplus_{i\le\ell} P^i$ we mean the protocol $(\{W_e\}_{e\in E}, \{L_v\}_{v\in V}, K)$ such that

1. W_e is the Cartesian product $\prod_{i\le\ell} W^i_e$,
2. L_v is such that

$$L_v = \left\{ \langle\langle t^i_e\rangle_{i\le\ell}\rangle_{e\in Inc(v)} \in \prod_{e\in Inc(v)}\prod_{i\le\ell} W^i_e \;\middle|\; \langle t^i_e\rangle_{e\in Inc(v)} \in L^i_v \text{ for all } i \le \ell \right\},$$

3. K is such that

$$K = \left\{ \langle \langle t_e^i \rangle_{i \le \ell} \rangle_{e \in E} \in \prod_{e \in E} \prod_{i \le \ell} W_e^i \;\middle|\; \langle t_e^i \rangle_{e \in E} \in K^i \text{ for some } i \le \ell \right\}.$$

Note that the only difference between Definition 8 and Definition 7 is in the way the individual sets K^i are combined into the composed set K.

Lemma 15 *$\langle \langle r_e^i \rangle_{i \le \ell} \rangle_{e \in E} \in R(\bigoplus_{i \le \ell} P^i)$ if and only if $\langle r_e^i \rangle_{e \in E} \in R(P^i)$ for all $0 < i \le \ell$.* □

Lemma 16 *If $\{P^i\}_{0<i\le\ell} = \{(\{W_e^i\}_{e\in E}, \{L_v^i\}_{v\in V}, K^i)\}_{0<i\le\ell}$ is a family of protocols over a hypergraph (V, E), then $\langle \langle r_e^i \rangle_{i \le \ell} \rangle_{e \in E} \vDash_{\bigoplus_{i \le \ell} P^i} [A_1, \ldots, A_n]$ if and only if there is $0 < i \le \ell$ such that $\langle r_e^i \rangle_{e \in E} \vDash_{P^i} [A_1, \ldots, A_n]$.*

Proof ($\Rightarrow$) : Suppose that for each $0 < i \le \ell$ there is a sequence of runs $\langle s_e^{i,0} \rangle_{e \in E}, \langle s_e^{i,1} \rangle_{e \in E}, \ldots, \langle s_e^{i,T_i} \rangle_{e \in E}$ of the protocol P^i and $j_1^i, \ldots, j_{T_i}^i \le n$ such that

$$\langle r_e^i \rangle_{e \in E} = \langle s_e^{i,0} \rangle_{e \in E} \equiv_{A_{j_1^i}} \langle s_e^{i,1} \rangle_{e \in E} \equiv_{A_{j_2^i}} \cdots \equiv_{A_{j_{T_i}^i}} \langle s_e^{i,T_i} \rangle_{e \in E} \notin K^i. \tag{6}$$

Informally, this means that for each individual protocol P^i, there is a sequence of "jumps" leading outside of the knowledge K^i. We will now combine these sequences of jumps in individual protocols into a single sequence of jumps in the composite protocol $\bigoplus_{i \le \ell} P^i$. This sequence will end outside of the knowledge of the composite protocol. The jumps on different components of the composite protocol will be executed not concurrently, but one component protocol at a time.

For each $\alpha, i \in \{1, \ldots, \ell\}$, $e \in E$, and $t \in \{0, 1, \ldots, T_i\}$, consider $S_e^{\alpha,i,t}$ such that

$$S_e^{\alpha,i,t} = \begin{cases} s_e^{i,T_\beta} & \text{if } \alpha > i, \\ s_e^{i,t} & \text{if } \alpha = i, \\ r_e^i & \text{otherwise.} \end{cases} \tag{7}$$

By Lemma 15, $\langle \langle S_e^{\alpha,i,t} \rangle_{i \le \ell} \rangle_{e \in E}$ is a run of the protocol $\bigoplus_{i \le \ell} P^i$ for each $0 < \alpha \le \ell$ and each $t \le T_i$ for every $0 < i \le \ell$. Hence, due to (6) and (7),

$$\begin{aligned}
\langle \langle r_e^i \rangle_{i \le \ell} \rangle_{e \in E} &= \langle \langle S_e^{1,i,0} \rangle_{i \le \ell} \rangle_{e \in E} \equiv_{A_{j_1^1}} \langle \langle S_e^{1,i,1} \rangle_{i \le \ell} \rangle_{e \in E} \equiv_{A_{j_2^1}} \cdots \equiv_{A_{j_{T_1}^1}} \\
\langle \langle S_e^{1,i,T_i} \rangle_{i \le \ell} \rangle_{e \in E} &= \langle \langle S_e^{2,i,0} \rangle_{i \le \ell} \rangle_{e \in E} \equiv_{A_{j_1^2}} \langle \langle S_e^{2,i,1} \rangle_{i \le \ell} \rangle_{e \in E} \equiv_{A_{j_2^2}} \cdots \equiv_{A_{j_{T_2}^2}} \\
\langle \langle S_e^{2,i,T_i} \rangle_{i \le \ell} \rangle_{e \in E} &= \langle \langle S_e^{3,i,0} \rangle_{i \le \ell} \rangle_{e \in E} \equiv_{A_{j_1^3}} \langle \langle S_e^{3,i,1} \rangle_{i \le \ell} \rangle_{e \in E} \equiv_{A_{j_2^3}} \cdots \equiv_{A_{j_{T_3}^3}} \\
\langle \langle S_e^{3,i,T_i} \rangle_{i \le \ell} \rangle_{e \in E} &= \langle \langle S_e^{4,i,0} \rangle_{i \le \ell} \rangle_{e \in E} \equiv_{A_{j_1^4}} \cdots \equiv_{A_{j_{T_\ell}^\ell}} \\
\langle \langle S_e^{\ell,i,T_i} \rangle_{i \le \ell} \rangle_{e \in E} &= \langle \langle s_e^{i,T_i} \rangle_{i \le \ell} \rangle_{e \in E}.
\end{aligned}$$

We now show that

$$\langle\langle s_e^{i,T_i}\rangle_{i\le\ell}\rangle_{e\in E} \notin \left\{ \langle\langle t_e^i\rangle_{i\le\ell}\rangle_{e\in E} \in \prod_{e\in E}\prod_{i\le\ell} W_e^i \;\middle|\; \langle t_e^i\rangle_{e\in E} \in K^i \text{ for some } i\le\ell \right\}.$$

In other words,

$$\langle\langle s_e^{i,T_i}\rangle_{i\le\ell}\rangle_{e\in E} \in \left\{ \langle\langle t_e^i\rangle_{i\le\ell}\rangle_{e\in E} \in \prod_{e\in E}\prod_{i\le\ell} W_e^i \;\middle|\; \langle t_e^i\rangle_{e\in E} \notin K^i \text{ for all } i\le\ell \right\},$$

which is true due to statement (6). Therefore, $\langle\langle r_e^i\rangle_{i\le\ell}\rangle_{e\in E} \nvdash_{\bigoplus_{i\le\ell} P^i} [A_1,\dots,A_n]$.

($\Leftarrow$) : Suppose $\langle\langle s_e^{i,0}\rangle_{i\le\ell}\rangle_{e\in E}, \langle\langle s_e^{i,1}\rangle_{i\le\ell}\rangle_{e\in E}, \dots, \langle\langle s_e^{i,T}\rangle_{i\le\ell}\rangle_{e\in E}$ is a sequence of runs of the protocol $\bigoplus_{i\le\ell} P^i$ and $j_1,\dots,j_T \le n$ such that

$$\langle\langle r_e^i\rangle_{i\le\ell}\rangle_{e\in E} = \langle\langle s_e^{i,0}\rangle_{i\le\ell}\rangle_{e\in E} \equiv_{A_{j_1}} \langle\langle s_e^{i,1}\rangle_{i\le\ell}\rangle_{e\in E} \equiv_{A_{j_2}} \cdots \equiv_{A_{j_T}} \langle\langle s_e^{i,T}\rangle_{i\le\ell}\rangle_{e\in E}$$

$$\langle\langle s_e^{i,T}\rangle_{i\le\ell}\rangle_{e\in E} \notin \left\{ \langle\langle t_e^i\rangle_{i\le\ell}\rangle_{e\in E} \in \prod_{e\in E}\prod_{i\le\ell} W_e^i \;\middle|\; \langle t_e^i\rangle_{e\in E} \in K^i \text{ for some } i\le\ell \right\}.$$

Hence,

$$\langle\langle s_e^{i,T}\rangle_{i\le\ell}\rangle_{e\in E} \in \left\{ \langle\langle t_e^i\rangle_{i\le\ell}\rangle_{e\in E} \in \prod_{e\in E}\prod_{i\le\ell} W_e^i \;\middle|\; \langle t_e^i\rangle_{e\in E} \notin K^i \text{ for all } i\le\ell \right\}.$$

Thus, for every $0 < i \le \ell$

$$\langle r_e^i\rangle_{e\in E} = \langle s_e^{i,0}\rangle_{e\in E} \equiv_{A_{j_1}} \langle s_e^{i,1}\rangle_{e\in E} \equiv_{A_{j_2}} \cdots \equiv_{A_{j_T}} \langle s_e^{i,T}\rangle_{e\in E} \notin K^i,$$

which is a contradiction with the assumption that there is an $0 < i \le \ell$ such that $\langle r_e^i\rangle_{e\in E} \vDash_{P^i} [A_1,\dots,A_n]$. □

7.3 Completeness: The Final Steps

Theorem 2 *For any hypergraph $H = (V,E)$ and any formula $\varphi \in \Phi(H)$, if $r \vDash_P \varphi$ for each run r of each protocol P over hypergraph H, then $\vdash_H \varphi$.*

Proof Suppose that $\nvdash_H \varphi$. Let X be a maximal consistent set such that $\varphi \in X$. By Lemma 12, for each $\bar{A} = [A_1,\dots,A_n]$ and $\bar{B} = [B_1,\dots,B_m]$ such that $X \nvdash_H \bar{A} \to \bar{B}$, there exists a protocol $P(\bar{A},\bar{B})$ and a run $r^{\bar{A},\bar{B}} = \langle r_e^{\bar{A},\bar{B}}\rangle_{e\in E}$ such that $r^{\bar{A},\bar{B}} \vDash \bar{A}$ and $r^{\bar{A},\bar{B}} \nvDash \bar{B}$. Note that if $X \vdash_H \bar{A}$ and $X \nvdash_H \bar{B}$, then $X \nvdash_H \bar{A} \to \bar{B}$ due to the Modus Ponens inference rule. Hence, $P(\bar{A},\bar{B})$ and $r^{\bar{A},\bar{B}}$ exist for every $\bar{A}$ and $\bar{B}$ such that $X \vdash \bar{A}$ and $X \nvdash \bar{B}$.

Consider the protocol P such that

$$P = \bigotimes_{X\nvdash_H \bar{B}} \bigoplus_{X\vdash_H \bar{A}} P(\bar{A},\bar{B}).$$

Let $r = \langle r_e\rangle_{e\in E} = \langle\langle\langle r_e^{\bar{A},\bar{B}}\rangle_{X\vdash_H\bar{A}}\rangle_{X\nvdash_H\bar{B}}\rangle_{e\in E}$. By Lemma 15 and Lemma 13, tuple r is a run of the protocol P.

Lemma 17 *$X \vdash_H \bar{C}$ if and only if $r \vDash_P \bar{C}$ for any $\bar{C} = [C_1, \ldots, C_n]$ where $C_1, \ldots, C_n \subseteq E$.*

Proof ($\Rightarrow$) : Suppose $X \vdash_H \bar{C}$. Note that $r^{\bar{C},\bar{B}} \vDash_{P(\bar{C},\bar{B})} \bar{C}$ for every $\bar{B}$ such that $X \nvdash_H \bar{B}$ due to the choice of the run $r^{\bar{C},\bar{B}}$. Thus, by Lemma 16,

$$\langle \langle r_e^{\bar{C},\bar{B}} \rangle_{X \vdash \bar{C}} \rangle_{e \in E} \vDash_{(\bigoplus_{X \vdash_H \bar{C}} P(\bar{C},\bar{B}))} \bar{C}.$$

Then, by Lemma 14,

$$\langle \langle \langle r_e^{\bar{C},\bar{B}} \rangle_{X \vdash_H \bar{C}} \rangle_{X \nvdash_H \bar{B}} \rangle_{e \in E} \vDash_{(\bigotimes_{X \nvdash_H \bar{B}} \ \bigoplus_{X \vdash_H \bar{C}} P(\bar{C},\bar{B}))} \bar{C}.$$

Therefore, $r \vDash_P \bar{C}$.
($\Leftarrow$) : Suppose now that $X \nvdash_H \bar{C}$. Note that $r^{\bar{A},\bar{C}} \nvDash_{P(\bar{A},\bar{C})} \bar{C}$ for every $\bar{A}$ such that $X \vdash_H \bar{A}$, due to the choice of the run $r^{\bar{A},\bar{C}}$. Thus, by Lemma 16,

$$\langle \langle r_e^{\bar{A},\bar{C}} \rangle_{X \vdash \bar{A}} \rangle_{e \in E} \nvDash_{(\bigoplus_{X \vdash_H \bar{A}} P(\bar{A},\bar{C}))} \bar{C}.$$

Then, by Lemma 14,

$$\langle \langle \langle r_e^{\bar{A},\bar{C}} \rangle_{X \vdash_H \bar{A}} \rangle_{X \nvdash_H \bar{C}} \rangle_{e \in E} \nvDash_{(\bigotimes_{X \nvdash_H \bar{C}} \ \bigoplus_{X \vdash_H \bar{A}} P(\bar{A},\bar{C}))} \bar{C}.$$

Therefore, $r \nvDash_P \bar{C}$. □

Lemma 18 *$X \vdash_H \psi$ if and only if $r \vDash_P \psi$ for each $\psi \in \Phi(H)$.*

Proof Induction on the structural complexity of formula ψ. The base case when ψ is an atomic formula $\bar{C}$ follows from Lemma 17, the base case when ψ is $\bot$ follows from Definition 5. The induction steps follows from Definition 5 in the standard way due to maximality and consistency of the set X. □

To finish the proof of the theorem, notice that $\varphi \notin X$ due to consistency of the set X. Therefore, by Lemma 18, $r \nvDash \varphi$. □

Since completeness has been shown with respect to a class of finite protocols, this result, as we show below, implies decidability of our logical system.

Corollary 1 *Set $\{\varphi \in \Phi(H) \mid \vdash_H \varphi\}$ is decidable for any hypergraph H.*

Proof Recursive enumerability of this set follows from recursive enumerability of the axioms. Recursive enumerability of the complement of this set follows from the completeness with the respect to the class of finite protocols finiteness of hypergraphs, and decidability of binary relation $r \vDash \varphi$. □

8 Conclusion

In this article we have developed a logical system for reasoning about common knowledge on a given hypergraph. At the core of this system is the Ryōan-ji principle, which forms a sound and complete logical system together with the Omniscience and Subgroup axioms.

Our system is designed to reason about common knowledge, denoted by $[A_1, \dots, A_n]$, of given knowledge specified by the set K. A natural extension of this system is a modal logic with modal formula $[A_1, \dots, A_n]\varphi$ being interpreted as "agents observing sets of channels $A_1, \dots, A_n$ have common knowledge of φ", where formula φ also can include modality $[\dots]$. The Omniscience, Ryōan-ji, and Subgroup axioms remain true in this setting. A complete description of such a logical system remains an open question. Such extension of the current work is not straightforward, given that completeness of modal epistemic logic with modalities indexed by sets of agents is not trivial even for distributed knowledge case (see Fagin et al (1995)).

Another possible extension of this work is adding dynamics to the information propagation through announcements, belief formations, or dynamic hypergraph modifications.

References

Berge C (1989) Hypergraphs, North-Holland Mathematical Library, vol 45. North-Holland Publishing Co., Amsterdam, combinatorics of finite sets, Translated from the French

Donders MS, More SM, Naumov P (2011) Information flow on directed acyclic graphs. In: Beklemishev LD, de Queiroz R (eds) WoLLIC, Springer, Lecture Notes in Computer Science, vol 6642, pp 95–109

Fagin R, Halpern JY, Moses Y, Vardi MY (1995) Reasoning about knowledge. MIT Press, Cambridge, MA

Halpern JY, Moses Y (1990) Knowledge and common knowledge in a distributed environment. J ACM 37(3):549–587

Holbrook S, Naumov P (2012) Fault tolerance in belief formation networks. In: del Cerro LF, Herzig A, Mengin J (eds) JELIA, Springer, Lecture Notes in Computer Science, vol 7519, pp 267–280

Kane J, Naumov P (2013) Epistemic logic for communication chains. In: 14th conference on Theoretical Aspects of Rationality and Knowledge (TARK '13), January 2013, Chennai, India, pp 131–137

Lewis D (1969) Convention: a philosophical study. Harvard University Press

More SM, Naumov P (2011a) The functional dependence relation on hypergraphs of secrets. In: Leite J, Torroni P, Ågotnes T, Boella G, van der Torre L (eds) CLIMA, Springer, Lecture Notes in Computer Science, vol 6814, pp 29–40

More SM, Naumov P (2011b) Hypergraphs of multiparty secrets. Ann Math Artif Intell 62(1-2):79–101

More SM, Naumov P (2011c) Logic of secrets in collaboration networks. Ann Pure Appl Logic 162(12):959–969
Studer T (2009) Common knowledge does not have the Beth property. Inform Process Lett 109(12):611–614