

Coalition Power in Epistemic Transition Systems

Pavel Naumov
Computer Science
Vassar College
Poughkeepsie, NY 12604
pnaumov@vassar.edu

Jia Tao
Computer Science
The College of New Jersey
Ewing, NJ 08628
taoj@tcnj.edu

ABSTRACT

The paper proposes a bimodal logic that describes an interplay between coalition strategies and distributed knowledge. Unlike the existing literature, the paper assumes that a strategy must be not only executable but also verifiable. That is, the strategy of a coalition should be based only on the information distributively known by the coalition and the coalition must be able to verify the result after the strategy is executed. The main technical result of the paper is a sound and complete logical system describing all universal properties expressible in the proposed bimodal language.

CCS Concepts

•Theory of computation → Modal and temporal logics; •Computing methodologies → Reasoning about belief and knowledge; Multi-agent systems; Cooperation and coordination; *Theory of mind*;

Keywords

coalition power, knowledge, formal epistemology, strategy, game theory, social choice, transition system, multiagent system, modal logic, axiomatization, completeness

1. INTRODUCTION

In this paper we study an interplay between distributed knowledge and strategic behavior in epistemic transition systems. We start, however, with an example of a *non-epistemic* transition system T_1 depicted in Figure 1. This




Figure 1: Non-epistemic transition system T_1 .

Appears in: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AA-MAS 2017)*, S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.

Copyright © 2017, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

system has three states, denoted in the figure by u , v , and w . We assume that the system has three agents called a , b , and c . To make the system to transition from the current state into another state, each of the agents takes a certain action (or casts a vote). We will consider different mechanisms for aggregation of votes of individual agents into a group decision. Although in general such mechanisms can be different from one state to another state, for our introductory example we assume that each agent has only two votes: 0 and 1 and the system transitions into a state the majority vote calls for. In other words, the labels on the edges in Figure 1 represent the outcome of the vote rather than the individual votes of the agents. For example, if the majority of agents vote 1 in state u , then the system transitions into state w ; otherwise it transitions into state v . Upon transitioning into a new state, a new vote is taken to determine the next transition. Since any two agents in system T_1 constitute a majority, their coalition can pre-determine the outcome of any vote. In other words, any coalition of two players has a way to control the behavior of the transition system T_1 . For example, if the system is currently in state u , then coalition $\{a, b\}$ can make the system to cycle between states u and w without ever visiting state v . As a result, statement p will remain permanently true, see Figure 1. In such a situation, we say that in state u coalition $\{a, b\}$ has a *strategy* to permanently enforce condition p . We denote this by $u \Vdash S_{\{a,b\}}p$ or simply as $u \Vdash S_{a,b}p$.

Note that in system T_1 neither of the three agents alone has a strategy to enforce condition p . Thus, $u \Vdash \neg S_a p$. Also, in system T_1 coalition $\{a, b\}$ cannot enforce condition p in state v because condition p is not satisfied in state v , see Figure 1. Hence, $v \Vdash \neg S_{a,b} p$.

1.1 Single Epistemic Agent

In this paper we propose a bimodal logic for reasoning about an interplay between strategies and knowledge. To achieve this goal we introduce the notion of an *epistemic* transition system. Figure 2 depicts a simple example of an epistemic transition system T_2 with just a single agent a . In addition to available transitions between the states, the figure also specifies an epistemic indistinguishability relation between the states using a dotted line. In the case of system T_2 , states u and v are indistinguishable to agent a . Since system T_2 has only one agent, we assume a trivial vote aggregation mechanism: the transition that the system takes is completely determined by the vote of agent a . For example, if in state u agent a votes 1, then the system transitions into state w ; otherwise, it transitions into state v , see Figure 2.




Figure 2: Epistemic transition system T_2 .

Executability. At first glance, it appears that agent a has a strategy in states u and v to enforce condition p by voting 0 in state u and voting 1 in state v so that the system alternates between states u and v . However, agent a is not able to distinguish states u and v and, thus, would not be able to determine the right vote in either of these states. In other words, the hypothetical strategy described above is not *executable* based on the information available to agent a . In this paper we consider only strategies of an agent that can be executed based on the information available to the agent. Since agent a does not have an executable strategy at state u of system T_2 to enforce condition p , we write $u \Vdash \neg S_a p$.

The requirement that a strategy of an agent a at a state u should be executable implies that *the strategy could be used in all states indistinguishable from state u by agent a* . Thus, if a strategy exists in state u it also exists in all states in the equivalence class of state u with respect to the indistinguishability equivalence relation of agent a . Hence, due to the standard semantics of knowledge, if such a strategy exists, its existence must be known to agent a . In other words, $S_a \varphi \rightarrow K_a S_a \varphi$ is a universally true logical principle, where K_a is the knowledge modality. We call this principle *positive strategic introspection*. Lemma 2 shows that this principle is provable from our axioms.

Similarly, if there is no strategy in a state u for an agent a to enforce φ , then this agent cannot have such a strategy in any of the states indistinguishable from state u by the agent. Thus, *negative strategic introspection* principle $\neg S_a \varphi \rightarrow K_a \neg S_a \varphi$ is also universally true and implied by our assumption of the executability of the strategies. We prove this principle from our axioms in Lemma 3.

Verifiability. Consider a single-agent epistemic transition system T_3 depicted in Figure 3. It again appears that agent




Figure 3: Epistemic transition system T_3 .

a has a strategy in state u to enforce condition p by alternating between two states: u and w . Indeed, if under this hypothetical strategy agent a always votes 1 in states u and v and she always votes 0 in state w , then the system will never transition into state v and the condition p will be en-

forced. Although agent a is not able to distinguish states u and v , such a hypothetical strategy is executable because the same vote is cast in these two states. However, she cannot verify if the strategy is successful. In fact, in this example the agent does not even know if the condition p is satisfied at state u because she is not able to distinguish states u and v . In this paper we only consider *verifiable* strategies. In other words, we assume that the agent should be able to verify that the condition that the strategy is enforcing is satisfied in each state at which the system might transition under this strategy. This assumption implies the *verifiability* principle $S_a \varphi \rightarrow S_a K_a \varphi$. This principle is derived from the axioms of our logical systems in Lemma 4.

The verifiability assumption implies that in system T_3 the agent a also has no strategy to enforce condition p in the state w of the transition system. Indeed, even if the agent can verify condition p in the state w itself, she will not be able to do this once the system transitions to state u . One can raise an objection to this argument by claiming that since the transitions start at state w , the agent should be able to distinguish state u from v by remembering that she voted 0 to get to this state. This claim, however, is valid only under an assumption of a perfect memory by the agent. We assume that the agents do not have a perfect memory and that an epistemic state description captures whatever memories the agent has at this state. In other words, the only knowledge that an agent possesses is the knowledge captured by the indistinguishability relation on the epistemic states.

Let us now consider a single-agent epistemic transition system T_4 depicted in Figure 4. Although in this system




Figure 4: Epistemic Transition system T_4 .

agent a is not able to distinguish states u and v , she has an executable and verifiable strategy in both of these states to enforce condition p . This strategy is defined as always voting 0, which results in the system alternating between these two states. This strategy is executable because it requires the agent to vote the same way in the two indistinguishable states. Thus, it does not use any information not available to the agent. The strategy is verifiable because in both states not only is condition p true, but it is known to be true to agent a . That is, $u \Vdash S_a p$ and $v \Vdash S_a p$.

1.2 Distributed Knowledge

In the previous section we have discussed our two basic requirements on strategies in epistemic transition systems, executability and verifiability, for single-agent strategies. In a more general case of coalition strategies, the same requirements are expressed through the notion of distributed knowledge by the coalition. In other words, we assume that members of a coalition not only coordinate their actions, but also share information available to them.

The executability of a coalition strategy requires that *the*

coalition should be able to execute the strategy based on the information distributively known to the members of the coalition. That is, if a coalition cannot distinguish two epistemic states, then each member of the coalition must cast the same vote in both states.




Figure 5: Multiagent epistemic transition system T_5 .

As an example, consider epistemic transition system T_5 depicted in Figure 5. Just like in transition system T_1 discussed earlier, we assume that this system has three agents: a , b , and c . Each of these agents votes either 0 or 1 and the system transitions based on the majority vote according to the labels shown in Figure 5. Since coalition $\{a, b\}$ constitutes a majority in this system, it should be able to enforce condition p in states u , w , and v by forcing the system to remain in one of these states. Hypothetically, they can achieve this by both voting 0 in states u and v and voting 1 in state w . Such a strategy, however, is not executable based on the individual knowledge of the agents a or b . For instance, agent a cannot vote according to this strategy in states u and w because she needs to vote differently in these two states, but she cannot distinguish them herself. At the same time, both agents can follow this strategy (and verify that condition p holds) if they share information available to them. Since we assume that members' actions are based on the information distributively known to the whole coalition, $w \Vdash S_{a,b}p$.

Similarly, the verifiability requirement for coalition strategies states that a coalition must be able to verify the enforced condition, before and after each transition, using the information distributively known to the members of the coalition. Consider the epistemic transition system T_6 depicted in Figure 6. We again assume that this system has three agents:




Figure 6: Multiagent epistemic transition system T_6 .

a , b , and c . Each of these agents votes either 0 or 1 and the system transitions based on the majority vote according to the labels shown in the figure. Note that in this system coalition $\{a, b\}$ can enforce condition p in states u and v by forcing system to alternate between these two states. This could be achieved by adopting the strategy for both of them

to vote 0 in all states of the system. Such a strategy does not require any knowledge about the current state and thus it is executable. However, this strategy is not verifiable based on the individual knowledge of the coalition members. For example, agent a cannot distinguish state u from state w and thus she would not be able to verify based on her knowledge alone that condition p holds at state u . At the same time, coalition $\{a, b\}$ can distinguish all states of this system based on the information distributively known by the members of the coalition: $u \Vdash S_{a,b}p$.

1.3 Related Literature

Logics of coalition power were developed by Marc Pauly [1, 2], who also proved the completeness of the basic logic of coalition power. The setting of the current paper resembles the “goal maintenance” in Pauly’s “extended coalition logic” [1, p. 80]. Pauly’s approach has been widely studied in the literature [3, 4, 5, 6, 7, 8, 9]. An alternative, binary-modality-based, logical system was proposed by More and Naumov [10].

Alur, Henzinger, and Kupferman introduced Alternating-Time Temporal Logic (ATL) that combines temporal and coalition modalities [11]. Van der Hoek and Wooldridge proposed to combine ATL with epistemic modality to form Alternating-Time Temporal Epistemic Logic [12]. They did not prove the completeness theorem for the proposed logical system. A completeness result for a logical system that combines coalition power and epistemic modalities was presented at AAMAS’12 by Ågotnes and Alechina [13]. Unlike our approach, their system does not impose any epistemic requirements on strategies. Our system requires strategies to be executable and verifiable. As a result, the logical system proposed by Ågotnes and Alechina does not contain the principles of positive strategic introspection, negative strategic introspection, and verifiability discussed above. In fact, their system does not contain any axiom describing an interplay of the two modalities.

Executability of strategies have been studied before under different names. While Jamroga and Ågotnes talk about “knowledge to identify and execute a strategy” [14], Jamroga and van der Hoek discuss “difference between an agent knowing that he has a suitable strategy and knowing the strategy itself” [15]. Van Benthem calls such strategies “uniform” [16]. None of these works proposes a complete system to capture the properties of executable strategies.

2. SYNTAX AND SEMANTICS

Throughout the rest of this paper we fix a set of propositional variables and a set of agents \mathcal{A} .

DEFINITION 1. Φ is the minimal set of formulas such that

1. $p \in \Phi$ for each propositional variable p ,
2. $\neg\varphi, \varphi \rightarrow \psi \in \Phi$ for each $\varphi, \psi \in \Phi$,
3. $K_C\varphi, S_C\varphi \in \Phi$ for each finite $C \subseteq \mathcal{A}$ and each $\varphi \in \Phi$.

The definition of an epistemic transition system given below is an extension of a standard Kripke semantics for multiagent epistemic logic S5 by two additional components: a domain of votes for each agent and a mechanism. The domain of votes is the set of all potential actions (or votes) that can be taken by an agent in each epistemic state to

influence the next transition of the system. In all examples above the domain of votes of each agent is set $\{0, 1\}$.

A mechanism is a procedure that determines the next state of the system based on the current state and the individual votes of the agents. We do not assume that all agents are treated equally by the mechanism and we do not assume that the mechanism is deterministic. An example of a non-deterministic mechanism is choosing a transition randomly if there is a tie. Formally we represent mechanism as a ternary relation between the current state, the set of votes, and the new state.

DEFINITION 2. *An epistemic transition system is a tuple $(W, \{\sim_a\}_{a \in \mathcal{A}}, \{V_a\}_{a \in \mathcal{A}}, M, \pi)$, where*

1. W is a set of states,
2. \sim_a is an “indistinguishability” equivalence relation for each $a \in \mathcal{A}$,
3. V_a is a set called “domain of choices” of agent $a \in \mathcal{A}$,
4. $M \subseteq W \times \prod_{a \in \mathcal{A}} V_a \times W$ is an aggregation mechanism,
5. π is a function that maps propositional variables into subsets of W .

Two epistemic states are indistinguishable by a coalition $C \subseteq \mathcal{A}$ if they are indistinguishable by every member of the coalition.

DEFINITION 3. *For any epistemic worlds $w_1, w_2 \in W$ and any set $C \subseteq \mathcal{A}$, let $w_1 \sim_C w_2$ if $w_1 \sim_c w_2$ for each $c \in C$.*

COROLLARY 1. *Relation \sim_C is an equivalence relation on the set of states W for each set of agents $C \subseteq \mathcal{A}$.*

Recall from the introduction that we expect all strategies to be executable. In other words, agents must use the same strategy in all states indistinguishable by the coalition. We achieve this by specifying a strategy as a function on equivalence classes.

DEFINITION 4. *A strategy of a coalition $C \subseteq \mathcal{A}$ is a tuple $\{s_c\}_{c \in C}$ such that $s_c : W / \sim_C \rightarrow V_c$.*

The remaining two definitions in this section formally specify the meaning of modal formula $S_C \varphi$ as “a coalition C has a strategy to enforce a condition φ indefinitely”. To understand our formal semantics it would be convenient to first discuss what it means for a coalition C to have a strategy \mathbf{s} to enforce condition φ in state w after exactly one transition. In the case of a *non-epistemic* transition system it means that “condition φ is satisfied in any state u into which the system could transition from state w if coalition C is using strategy \mathbf{s} ”. The executability requirement for *epistemic* systems forces the above quoted statement to be true not only for state w , but for each state $w' \in [w]_C$, where $[w]_C$ is the equivalence class of the state w with respect to relation \sim_C . Furthermore, the verifiability requirement forces the condition φ to be satisfied not just in every state u' reachable from $w' \in [w]$ but in every state $u \in [u']_C$. This is captured by a transition-like relation $\rightarrow_{\mathbf{s}}$ between equivalence classes.

DEFINITION 5. *For any states $w, u \in W$, any strategy $\mathbf{s} = \{s_c\}_{c \in C}$ of a coalition $C \subseteq \mathcal{A}$, let $[w]_C \rightarrow_{\mathbf{s}} [u]_C$ if there are $w' \in [w]_C$, $u' \in [u]_C$, and a set of votes $\{v_a\}_{a \in \mathcal{A}}$ such that*

1. $v_c = s_c([w]_C)$ for each $c \in C$,
2. $(w', \{v_a\}_{a \in \mathcal{A}}, u') \in M$.

Let $\rightarrow_{\mathbf{s}}^*$ be a relation on set W / \sim_C defined as the transitive reflexive closure of the relation $\rightarrow_{\mathbf{s}}$. The next definition is the key definition of this paper. It specifies the meaning of the two modalities in our logical system.

DEFINITION 6. *For any state $w \in W$ of a transition system $(W, \{\sim_a\}_{a \in \mathcal{A}}, \{V_a\}_{a \in \mathcal{A}}, M, \pi)$ and any formula $\varphi \in \Phi$, let relation $w \Vdash \varphi$ be defined as follows*

1. $w \Vdash p$ if $w \in \pi(p)$ for each propositional variable p ,
2. $w \Vdash \neg \varphi$ if $w \not\Vdash \varphi$,
3. $w \Vdash \varphi \rightarrow \psi$ if $w \not\Vdash \varphi$ or $w \Vdash \psi$,
4. $w \Vdash K_C \varphi$ if $w' \Vdash \varphi$ for each w' such that $w \sim_C w'$,
5. $w \Vdash S_C \varphi$ if there is a strategy \mathbf{s} of coalition C such that $w' \Vdash \varphi$ for every $w' \in W$ such that $[w]_C \rightarrow_{\mathbf{s}}^* [w']_C$.

3. AXIOMS

In addition to propositional tautologies in language Φ , our logical system consists of the following axioms:

1. Empty Coalition: $K_{\emptyset} \varphi \rightarrow S_{\emptyset} \varphi$,
2. Truth: $K_C \varphi \rightarrow \varphi$,
3. Negative Introspection: $\neg K_C \varphi \rightarrow K_C \neg K_C \varphi$,
4. Distributivity: $K_C(\varphi \rightarrow \psi) \rightarrow (K_C \varphi \rightarrow K_C \psi)$,
5. Monotonicity: $K_C \varphi \rightarrow K_D \varphi$, if $C \subseteq D$,
6. Knowledge: $S_C \varphi \rightarrow K_C \varphi$,
7. Persistence: $S_C \varphi \rightarrow S_C S_C \varphi$,
8. Cooperation: $S_C(\varphi \rightarrow \psi) \rightarrow (S_D \varphi \rightarrow S_{C \cup D} \psi)$, if $C \cap D = \emptyset$.

We say that formula φ is a theorem in our logical system and write $\vdash \varphi$ if formula φ is derivable from the axioms of our system using Modus Ponens and Strategic Necessitation inference rules:

$$\frac{\varphi, \varphi \rightarrow \psi}{\psi} \quad \frac{\varphi}{S_C \varphi}$$

We write $X \vdash \varphi$ if formula φ is derivable from the theorems of our logical system and an additional set of axioms X using only Modus Ponens inference rule.

LEMMA 1. *Knowledge Necessitation inference rule $\frac{\varphi}{K_C \varphi}$ is admissible in our logical system for each set $C \subseteq \mathcal{A}$.*

PROOF. Suppose that $\vdash \varphi$. Thus, $\vdash S_C \varphi$ by Strategic Necessitation inference rule. At the same time, $\vdash S_C \varphi \rightarrow K_C \varphi$ by Knowledge axiom. Therefore, $\vdash K_C \varphi$ by Modus Ponens inference rule. \square

4. EXAMPLES OF DERIVATIONS

The soundness of our logical system is shown in Section 5. In this section we give several examples of formal proofs in the system. The first three of these examples are *positive strategic introspection*, *negative strategic introspection*, and *verifiability* principles mentioned in the introduction. Lemma 5 is a strategic monotonicity principle.

LEMMA 2. $\vdash \mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\mathcal{S}_C\varphi$.

PROOF. By Persistence axiom, $\vdash \mathcal{S}_C\varphi \rightarrow \mathcal{S}_C\mathcal{S}_C\varphi$. At the same time, $\vdash \mathcal{S}_C\mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\mathcal{S}_C\varphi$ by Knowledge axiom. Therefore, $\vdash \mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\mathcal{S}_C\varphi$ by the laws of propositional reasoning. \square

LEMMA 3. $\vdash \neg\mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\neg\mathcal{S}_C\varphi$.

PROOF. First, note that $\vdash \mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\mathcal{S}_C\varphi$, by Lemma 2. Thus, $\vdash \neg\mathcal{K}_C\mathcal{S}_C\varphi \rightarrow \neg\mathcal{S}_C\varphi$ by the law of contrapositive. Hence, $\vdash \mathcal{K}_C(\neg\mathcal{K}_C\mathcal{S}_C\varphi \rightarrow \neg\mathcal{S}_C\varphi)$ by Lemma 1. Then, by Distributivity axiom and Modus Ponens inference rule,

$$\vdash \mathcal{K}_C\neg\mathcal{K}_C\mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\neg\mathcal{S}_C\varphi.$$

At the same time, $\vdash \neg\mathcal{K}_C\mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\neg\mathcal{K}_C\mathcal{S}_C\varphi$ by Negative Introspection axiom. Thus, by the laws of propositional reasoning, $\vdash \neg\mathcal{K}_C\mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\neg\mathcal{S}_C\varphi$. Hence, by the law of contrapositive, $\vdash \neg\mathcal{K}_C\neg\mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\mathcal{S}_C\varphi$. Note that $\vdash \mathcal{K}_C\mathcal{S}_C\varphi \rightarrow \mathcal{S}_C\varphi$ by Truth axiom. Thus, by the laws of propositional reasoning, $\vdash \neg\mathcal{K}_C\neg\mathcal{S}_C\varphi \rightarrow \mathcal{S}_C\varphi$. Therefore, $\vdash \neg\mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\neg\mathcal{S}_C\varphi$ again by the law of contrapositive. \square

LEMMA 4. $\vdash \mathcal{S}_C\varphi \rightarrow \mathcal{S}_C\mathcal{K}_C\varphi$.

PROOF. By Knowledge axiom, $\vdash \mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\varphi$. Thus, $\vdash \mathcal{S}_\emptyset(\mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\varphi)$ by Strategic Necessitation inference rule. At the same time, by Cooperation axiom,

$$\vdash \mathcal{S}_\emptyset(\mathcal{S}_C\varphi \rightarrow \mathcal{K}_C\varphi) \rightarrow (\mathcal{S}_C\mathcal{S}_C\varphi \rightarrow \mathcal{S}_C\mathcal{K}_C\varphi).$$

Thus, $\vdash \mathcal{S}_C\mathcal{S}_C\varphi \rightarrow \mathcal{S}_C\mathcal{K}_C\varphi$ by Modus Ponens inference rule. Therefore, $\vdash \mathcal{S}_C\varphi \rightarrow \mathcal{S}_C\mathcal{K}_C\varphi$ by Persistence axiom and the laws of propositional reasoning. \square

LEMMA 5. $\vdash \mathcal{S}_C\varphi \rightarrow \mathcal{S}_D\varphi$ where $C \subseteq D$.

PROOF. Note that $\varphi \rightarrow \varphi$ is a propositional tautology. Thus, $\vdash \varphi \rightarrow \varphi$. Hence, $\vdash \mathcal{S}_{D \setminus C}(\varphi \rightarrow \varphi)$ by Strategic Necessitation rule. At the same time, by Distributivity axiom, $\vdash \mathcal{S}_{D \setminus C}(\varphi \rightarrow \varphi) \rightarrow (\mathcal{S}_C\varphi \rightarrow \mathcal{S}_D\varphi)$ due to the assumption $C \subseteq D$. Thus, $\vdash \mathcal{S}_C\varphi \rightarrow \mathcal{S}_D\varphi$ by Modus Ponens rule. \square

We conclude the section with a well-known argument showing that the positive introspection principle follows from the rest of the axioms of modal logic S5. Later we use this observation in the proof of the completeness of our system.

LEMMA 6. $\vdash \mathcal{K}_C\varphi \rightarrow \mathcal{K}_C\mathcal{K}_C\varphi$.

PROOF. Formula $\neg\mathcal{K}_C\varphi \rightarrow \mathcal{K}_C\neg\mathcal{K}_C\varphi$ is an instance of Negative Introspection axiom. Thus, $\vdash \neg\mathcal{K}_C\neg\mathcal{K}_C\varphi \rightarrow \mathcal{K}_C\varphi$ by the law of contrapositive in the propositional logic. Hence, $\vdash \mathcal{K}_C(\neg\mathcal{K}_C\neg\mathcal{K}_C\varphi \rightarrow \mathcal{K}_C\varphi)$ by Knowledge Necessitation inference rule, see Lemma 1. Thus, by Distributivity axiom and Modus Ponens inference rule,

$$\vdash \mathcal{K}_C\neg\mathcal{K}_C\neg\mathcal{K}_C\varphi \rightarrow \mathcal{K}_C\mathcal{K}_C\varphi. \quad (1)$$

At the same time, $\mathcal{K}_C\neg\mathcal{K}_C\varphi \rightarrow \neg\mathcal{K}_C\varphi$ is an instance of Truth axiom. Thus, $\vdash \mathcal{K}_C\varphi \rightarrow \neg\mathcal{K}_C\neg\mathcal{K}_C\varphi$ by contraposition. Hence, taking into account the following instance of

Negative Introspection axiom $\neg\mathcal{K}_C\neg\mathcal{K}_C\varphi \rightarrow \mathcal{K}_C\neg\mathcal{K}_C\neg\mathcal{K}_C\varphi$, one can conclude that $\vdash \mathcal{K}_C\varphi \rightarrow \mathcal{K}_C\neg\mathcal{K}_C\neg\mathcal{K}_C\varphi$. The latter, together with statement (1), implies the statement of the lemma by the laws of propositional reasoning. \square

5. SOUNDNESS

In this section we prove the soundness of our logical system. We start the proof by introducing a notion of a composition of strategies of two disjoint coalitions and by proving several properties of this operation. These properties are used to prove the soundness of Cooperation axiom.

DEFINITION 7. For any two disjoint coalitions $C, D \subseteq \mathcal{A}$ and any two strategies $\mathbf{s}^1 = \{s_c^1\}_{c \in C}$ and $\mathbf{s}^2 = \{s_d^2\}_{d \in D}$ of these coalitions, let strategy $\mathbf{s}^1 \oplus \mathbf{s}^2$ of coalition $C \cup D$ be $\{s_a\}_{a \in C \cup D}$ where

$$s_a([w]_{C \cup D}) = \begin{cases} s_a^1([w]_C), & \text{if } a \in C, \\ s_a^2([w]_D), & \text{if } a \in D. \end{cases}$$

The next lemma shows that the composition of strategies is a well-defined operation.

LEMMA 7. For any disjoint coalitions C and D , if $w \sim_{C \cup D} w'$, then $s_a([w]_{C \cup D}) = s_a([w']_{C \cup D})$ for each $a \in C \cup D$.

PROOF. Without loss of generality, assume that $a \in C$. Note that $w \sim_{C \cup D} w'$ implies that $w \sim_C w'$. Thus, $[w]_C = [w']_C$. Therefore, $s_a([w]_{C \cup D}) = s_a^1([w]_C) = s_a^1([w']_C) = s_a([w']_{C \cup D})$. \square

LEMMA 8. For any two disjoint coalitions C and D , if $[w]_{C \cup D} \rightarrow_{\mathbf{s}^1 \oplus \mathbf{s}^2} [w']_{C \cup D}$, then $[w]_C \rightarrow_{\mathbf{s}^1} [w']_C$ and $[w]_D \rightarrow_{\mathbf{s}^2} [w']_D$.

PROOF. Let $\mathbf{s}^1 = \{s_c^1\}_{c \in C}$, $\mathbf{s}^2 = \{s_d^2\}_{d \in D}$, and $\mathbf{s}^1 \oplus \mathbf{s}^2 = \{s_a\}_{a \in C \cup D}$. By Definition 5, assumption $[w]_{C \cup D} \rightarrow_{\mathbf{s}^1 \oplus \mathbf{s}^2} [w']_{C \cup D}$ implies that there are epistemic states $u \in [w]_{C \cup D}$, $u' \in [w']_{C \cup D}$, and a set of votes $\{v_a\}_{a \in \mathcal{A}}$ such that

1. $v_a = s_a([w]_{C \cup D})$ for each $a \in C \cup D$,
2. $(u, \{v_a\}_{a \in \mathcal{A}}, u') \in M$.

Hence, by Definition 7,

3. $v_c = s_c([w]_C)$ for each $c \in C$,
4. $v_d = s_d([w]_D)$ for each $d \in D$,
5. $(u, \{v_a\}_{a \in \mathcal{A}}, u') \in M$.

Note that assumptions $u \in [w]_{C \cup D}$ and $u' \in [w']_{C \cup D}$ imply that $u \sim_{C \cup D} w$ and $u' \sim_{C \cup D} w'$. Hence, $u \sim_C w$, $u' \sim_C w'$, $u \sim_D w$, $u' \sim_D w'$, by Definition 3. Thus, $u \in [w]_C$, $u' \in [w']_C$, $u \in [w]_D$, and $u' \in [w']_D$. Therefore, $[w]_C \rightarrow_{\mathbf{s}^1} [w']_C$ and $[w]_D \rightarrow_{\mathbf{s}^2} [w']_D$ by items 3, 4, and 5 above and Definition 5. \square

COROLLARY 2. If $[w]_{C \cup D} \rightarrow_{\mathbf{s}^1 \oplus \mathbf{s}^2}^* [w']_{C \cup D}$, then $[w]_C \rightarrow_{\mathbf{s}^1}^* [w']_C$ and $[w]_D \rightarrow_{\mathbf{s}^2}^* [w']_D$, for any two disjoint coalitions C and D .

In the rest of this section we prove the soundness of our system with respect to the semantics introduced in Section 2.

THEOREM 1. If $\vdash \varphi$, then $w \Vdash \varphi$ for each formula $\varphi \in \Phi$ and each state w of each epistemic transition system.

The soundness of propositional tautologies, S5 axioms (Distributivity, Truth, Negative Introspection, and Monotonicity), and Modus Ponens and Strategic Necessitation inference rules is standard. Below we show the soundness of each of the remaining axioms as a separate lemma.

LEMMA 9. *If $w \Vdash K_{\emptyset}\varphi$, then $w \Vdash S_{\emptyset}\varphi$.*

PROOF. By Definition 4, empty tuple $\mathbf{s} = \{s_a\}_{a \in \emptyset}$ is a strategy of the empty coalition. Consider any $u \in W$ such that $[w]_{\emptyset} \rightarrow_{\mathbf{s}} [u]_{\emptyset}$. By Definition 6, it suffices to show that $u \Vdash \varphi$. Indeed, $w \sim_{\emptyset} u$ by Definition 3. Hence, $u \Vdash \varphi$ by the assumption $w \Vdash K_{\emptyset}\varphi$ and Definition 6. \square

LEMMA 10. *If $w \Vdash S_C\varphi$, then $w \Vdash K_C\varphi$.*

PROOF. Suppose that $w \Vdash S_C\varphi$. Thus, by Definition 6, there is a strategy \mathbf{s} of coalition C such that $w' \Vdash \varphi$ for each $w' \in W$ such that $[w]_C \rightarrow_{\mathbf{s}}^* [w']_C$.

Consider any $u \in W$ such that $w \sim_C u$. By Definition 6, it suffices to show that $u \Vdash \varphi$. Indeed, $[w]_C \rightarrow_{\mathbf{s}}^* [w]_C$ due to $\rightarrow_{\mathbf{s}}^*$ being the transitive reflexive closure of the relation $\rightarrow_{\mathbf{s}}$. Hence, $[w]_C \rightarrow_{\mathbf{s}}^* [u]_C$ by the assumption $w \sim_C u$. Therefore, $u \Vdash \varphi$ due to the choice of strategy \mathbf{s} . \square

LEMMA 11. *If $w \Vdash S_C\varphi$, then $w \Vdash S_C S_C\varphi$.*

PROOF. Assumption $w \Vdash S_C\varphi$ by Definition 6 implies that there is a strategy \mathbf{s} of coalition C such that for each $u \in W$ if $[w]_C \rightarrow_{\mathbf{s}}^* [u]_C$, then $u \Vdash \varphi$.

Consider any $w' \in W$ such that $[w]_C \rightarrow_{\mathbf{s}}^* [w']_C$. By Definition 6, we need to show that $w' \Vdash S_C\varphi$. Let $w'' \in W$ be any epistemic world such that $[w']_C \rightarrow_{\mathbf{s}}^* [w'']_C$. Again by Definition 6, it suffices to show that $w'' \Vdash \varphi$. Note that $[w]_C \rightarrow_{\mathbf{s}}^* [w'']_C$ due to $\rightarrow_{\mathbf{s}}^*$ being the transitive reflexive closure of the relation $\rightarrow_{\mathbf{s}}$. Therefore, $w'' \Vdash \varphi$ due to the choice of strategy \mathbf{s} . \square

LEMMA 12. *If $w \Vdash S_C(\varphi \rightarrow \psi)$, $w \Vdash S_D\varphi$, and $C \cap D = \emptyset$, then $w \Vdash S_{C \cup D}\psi$.*

PROOF. Suppose that $w \Vdash S_C(\varphi \rightarrow \psi)$. Thus, by Definition 6, there is a strategy $\mathbf{s}^1 = \{s_c^1\}_{c \in C}$ of coalition C such that $u \Vdash \varphi \rightarrow \psi$ for each u where $[w]_C \rightarrow_{\mathbf{s}^1}^* [u]_C$. Similarly, assumption $w \Vdash S_D\varphi$ implies that there is a strategy $\mathbf{s}^2 = \{s_d^2\}_{d \in D}$ of coalition D such that $u \Vdash \varphi$ for each u where $[w]_D \rightarrow_{\mathbf{s}^2}^* [u]_D$.

Consider any $w' \in W$ such that $[w]_{C \cup D} \rightarrow_{\mathbf{s}^1 \oplus \mathbf{s}^2}^* [w']_{C \cup D}$. By Definition 6, it suffices to show that $w' \Vdash \psi$. Indeed, assumption $[w]_{C \cup D} \rightarrow_{\mathbf{s}^1 \oplus \mathbf{s}^2}^* [w']_{C \cup D}$, by Corollary 2, implies that $[w]_C \rightarrow_{\mathbf{s}^1}^* [w']_C$ and $[w]_D \rightarrow_{\mathbf{s}^2}^* [w']_D$. Thus, $w' \Vdash \varphi \rightarrow \psi$ and $w' \Vdash \varphi$ by the choice of strategies \mathbf{s}^1 and \mathbf{s}^2 . Therefore, $w' \Vdash \psi$ by Definition 6. \square

6. COMPLETENESS

The section is dedicated to the proof of the following completeness theorem.

THEOREM 2. *Let φ be any formula in set Φ . If $w \Vdash \varphi$ for every epistemic state $w \in W$ of every epistemic transition system $(W, \{\sim_a\}_{a \in A}, \{V_a\}_{a \in A}, M, \pi)$, then $\vdash \varphi$.*

6.1 Properties of Maximal Consistent Sets

We start the proof of the completeness by establishing two properties of maximal consistent sets that are used later. The first property, Lemma 13, is typical for a proof of completeness of a modal logic. The second property, Lemma 14 below, is specific to our construction.

LEMMA 13. *For any consistent set of formulae X and any $\neg K_C\varphi \in X$, set $\{\neg\varphi\} \cup \{\psi \mid K_C\psi \in X\}$ is consistent.*

PROOF. Assume the opposite. Thus, there must exist formulae $K_C\psi_1, \dots, K_C\psi_n \in X$ such that $\psi_1, \dots, \psi_n \vdash \varphi$. Thus, by the deduction theorem for propositional logic,

$$\vdash \psi_1 \rightarrow (\psi_2 \rightarrow \dots (\psi_n \rightarrow \varphi) \dots).$$

Hence, by Knowledge Necessitation inference rule (Lemma 1),

$$\vdash K_C(\psi_1 \rightarrow (\psi_2 \rightarrow \dots (\psi_n \rightarrow \varphi) \dots)).$$

By Distributivity axiom and Modus Ponens inference rule,

$$K_C\psi_1 \vdash K_C(\psi_2 \rightarrow \dots (\psi_n \rightarrow \varphi) \dots).$$

By repeating the last step $(n - 1)$ times,

$$K_C\psi_1, \dots, K_C\psi_n \vdash K_C\varphi.$$

Hence, $X \vdash K_C\varphi$ by the choice of formula $K_C\psi_1, \dots, K_C\psi_n$, which contradicts to the consistency of set X due to the assumption $\neg K_C\varphi \in X$. \square

LEMMA 14. *For any consistent set of formulae X , and any subsets D_1, \dots, D_n of a finite set $C \subseteq A$, any formula $\neg S_C\psi \in X$, any formulae $S_{D_1}\varphi_1, \dots, S_{D_n}\varphi_n \in X$, and any formulae $K_{\emptyset}\chi_1, \dots, K_{\emptyset}\chi_m \in X$, if $D_i \cap D_j = \emptyset$ for each $i \neq j$ where $1 \leq i, j \leq n$, then the set of formulae $\{\neg\psi, \varphi_1, \dots, \varphi_n, \chi_1, \dots, \chi_m\}$ is consistent.*

PROOF. Suppose $\varphi_1, \dots, \varphi_n, \chi_1, \dots, \chi_m \vdash \psi$. Hence, by the deduction theorem for propositional logic applied $n + m$ times, $\vdash \varphi_1 \rightarrow (\dots (\varphi_n \rightarrow (\chi_1 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots))) \dots)$. Then, by Strategic Necessitation inference rule,

$$\vdash S_{\emptyset}(\varphi_1 \rightarrow (\dots (\varphi_n \rightarrow (\chi_1 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots))) \dots)).$$

Thus, by Cooperation axiom and Modus Ponens rule,

$$\vdash S_{D_1}\varphi_1 \rightarrow S_{\emptyset \cup D_1}(\varphi_2 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots)).$$

In other words,

$$\vdash S_{D_1}\varphi_1 \rightarrow S_{D_1}(\varphi_2 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots)).$$

Then, by Modus Ponens inference rule,

$$S_{D_1}\varphi_1 \vdash S_{D_1}(\varphi_2 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots)).$$

By Cooperation axiom and Modus Ponens inference rule,

$$S_{D_1}\varphi_1 \vdash S_{D_2}\varphi_2 \rightarrow S_{D_1 \cup D_2}(\dots (\chi_m \rightarrow \psi) \dots).$$

Again, by Modus Ponens inference rule,

$$S_{D_1}\varphi_1, S_{D_2}\varphi_2 \vdash S_{D_1 \cup D_2}(\dots (\chi_m \rightarrow \psi) \dots).$$

By repeating the previous steps $n - 2$ times,

$$S_{D_1}\varphi_1, \dots, S_{D_n}\varphi_n \vdash S_{D_1 \cup \dots \cup D_n}(\chi_1 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots)).$$

Recall that $D_1 \cup \dots \cup D_n \subseteq C$ by the assumption of the lemma. Hence, by Lemma 5,

$$S_{D_1}\varphi_1, \dots, S_{D_n}\varphi_n \vdash S_C(\chi_1 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots)).$$

Note that $S_{D_1}\varphi_1, S_{D_2}\varphi_2, \dots, S_{D_n}\varphi_n \in X$ by the assumption of the lemma. Thus, $X \vdash S_C(\chi_1 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots))$. By Cooperation axiom and Modus Ponens inference rule,

$$X \vdash S_{\emptyset}\chi_1 \rightarrow S_{C \cup \emptyset}(\chi_2 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots)).$$

In other words,

$$X \vdash S_{\emptyset}\chi_1 \rightarrow S_C(\chi_2 \rightarrow (\dots (\chi_m \rightarrow \psi) \dots)).$$

By Modus Ponens inference rule,

$$X, S_{\emptyset}\chi_1 \vdash S_C(\chi_2 \rightarrow (\dots(\chi_m \rightarrow \psi)\dots)).$$

by Empty Coalition axiom and Modus Ponens inference rule,

$$X, K_{\emptyset}\chi_1 \vdash S_C(\chi_2 \rightarrow (\dots(\chi_m \rightarrow \psi)\dots)).$$

By repeating the previous steps $m - 1$ times,

$$X, K_{\emptyset}\chi_1, \dots, K_{\emptyset}\chi_m \vdash S_C\psi.$$

Recall that $K_{\emptyset}\chi_1, \dots, K_{\emptyset}\chi_m \in X$ by the assumption of the lemma. Thus, $X \vdash S_C\psi$. This contradicts the assumption $\neg S_C\psi \in X$ of the lemma due to the consistency of set X . \square

6.2 Canonical Model

In this section we construct a canonical model based on a fixed maximal consistent set of formulae X_0 . This construction is based on the ‘‘unravelling’’ technique [17]. Informally, states in this model are nodes in a tree. Nodes of the tree are labeled with maximal consistent sets of formulae and edges of the trees are labeled with coalitions. The root node of the tree is labeled with set X_0 . Formally, epistemic states are defined as sequences representing paths in such a tree.

DEFINITION 8. *The set of epistemic states W consists of all finite sequences $X_0, C_1, X_1, C_2, \dots, C_n, X_n$, such that*

1. $n \geq 0$,
2. X_i is a maximal consistent subset of Φ for each $i \geq 1$,
3. C_i is a finite subset of \mathcal{A} for each $i \geq 1$,
4. $\{\varphi \mid K_{C_i}\varphi \in X_{i-1}\} \subseteq X_i$ for each $i \geq 1$.

DEFINITION 9. *The domain of choices V_a is $\Phi \times \mathcal{P}(\mathcal{A})$ for each $a \in \mathcal{A}$.*

DEFINITION 10. *For any $w = X_0, C_1, X_1, C_2, \dots, C_n, X_n$ and any $w' = X_0, C'_1, X'_1, C'_2, \dots, C'_m, X'_m$, let $w \sim_a w'$ if there is an integer number k such that*

1. $0 \leq k \leq \min\{n, m\}$,
2. $X_i = X'_i$ for each i such that $0 < i \leq k$,
3. $C_i = C'_i$ for each i such that $0 < i \leq k$,
4. $a \in C_i$ for each i such that $k < i \leq n$,
5. $a \in C'_i$ for each i such that $k < i \leq m$.

For any state $w = X_0, C_1, X_1, C_2, \dots, C_n, X_n$, by $hd(w)$ we mean set X_n .

LEMMA 15. *For any $X_0, C_1, X_1, C_2, \dots, C_n, X_n \in W$ and any $k \leq n$, if $K_C\varphi \in X_n$ and $C \subseteq C_i$ for all i such that $k < i \leq n$, then $K_C\varphi \in X_k$.*

PROOF. Suppose that there is $k \leq n$ such that $K_C\varphi \notin X_k$. Let m be the maximal such k . Note that $m < n$ due to the assumption $K_C\varphi \in X_n$ of the lemma. Thus, $m < m+1 \leq n$.

Assumption $K_C\varphi \notin X_m$ implies $\neg K_C\varphi \in X_m$ due to the maximality of set X_m . Hence, $X_m \vdash K_C\neg K_C\varphi$ by Negative Introspection axiom. Thus, $X_m \vdash K_{C_{m+1}}\neg K_C\varphi$ by Monotonicity axiom and the assumption $C \subseteq C_{m+1}$ of the lemma (recall that $m+1 \leq n$). Then, $K_{C_{m+1}}\neg K_C\varphi \in X_m$ due to the maximality of the set X_m . Hence, $\neg K_C\varphi \in X_{m+1}$ by Definition 8. Thus, $K_C\varphi \notin X_{m+1}$ due to the consistency of set X_{m+1} , which is a contradiction with the choice of m . \square

LEMMA 16. *For any $X_0, C_1, X_1, C_2, \dots, C_n, X_n \in W$ and any $k \leq n$, if $K_C\varphi \in X_k$ and $C \subseteq C_i$ for all i such that $k < i \leq n$, then $\varphi \in X_n$.*

PROOF. We prove the lemma by induction on the distance between n and k . In the base case, $n = k$. Then the assumption $K_C\varphi \in X_n$ implies $X_n \vdash \varphi$ by Truth axiom. Therefore, $\varphi \in X_n$ due to the maximality of set X_n .

Suppose that $k < n$. Assumption $K_C\varphi \in X_k$ implies $X_k \vdash K_C K_C\varphi$ by Lemma 6. Thus, $X_k \vdash K_{C_{k+1}} K_C\varphi$ by Monotonicity axiom, the condition $k < n$ of the inductive step, and the assumption $C \subseteq C_{k+1}$ of the lemma. Then, $K_{C_{k+1}} K_C\varphi \in X_k$ by the maximality of set X_k . Hence, $K_C\varphi \in X_{k+1}$ by Definition 8. Then, $\varphi \in X_n$ by the induction hypothesis. \square

LEMMA 17. *If $K_C\varphi \in hd(w)$ and $w \sim_C w'$, then $\varphi \in hd(w')$ for any states $w, w' \in W$ and any finite $C \subseteq \mathcal{A}$.*

PROOF. The statement follows from Lemma 15, Lemma 16, and Definition 10. \square

LEMMA 18. *If $K_C\varphi \notin hd(w)$, then there is an epistemic state $w' \in W$ such that $w \sim_C w'$ and $\varphi \notin hd(w')$.*

PROOF. Due to the maximality of the set $hd(w)$, assumption $K_C\varphi \notin hd(w)$ implies that $\neg K_C\varphi \in hd(w)$. Thus, by Lemma 13, set $Y_0 = \{\neg\varphi\} \cup \{\psi \mid K_C\psi \in hd(w)\}$ is consistent. Let Y be a maximal consistent extension of set Y_0 and w' be sequence w, C, Y . In other words, sequence w' is an extension of sequence w by two additional elements: C and Y . Note that $w' \in W$ due to Definition 8 and the choice of set Y_0 . Furthermore, $w \sim_C w'$ by Definition 10. To finish the proof, we need to show that $\varphi \notin hd(w')$. Indeed, $\neg\varphi \in Y_0 \subseteq Y = hd(w')$ by the choice of Y_0 . Therefore, $\varphi \notin hd(w')$ due to the consistency of the set $hd(w')$. \square

DEFINITION 11. *M is the set of all tuples $(w, \{v_a\}_{a \in \mathcal{A}}, w')$ such that*

$$\{\varphi \mid \exists C \subseteq \mathcal{A} (S_C\varphi \in hd(w) \wedge \forall c \in C (v_c = (\varphi, C)))\} \subseteq hd(w').$$

LEMMA 19. *Let $w, w' \in W$ be epistemic state, $S_C\varphi \in \Phi$ be a formula, and $\mathbf{s} = \{s_c\}_{c \in C}$ be the strategy of coalition C such that $s_c([u]_C) = (\varphi, C)$ for each $c \in C$ and each $u \in W$. If $S_C\varphi \in hd(w)$ and $[w]_C \rightarrow_{\mathbf{s}} [w']_C$, then $S_C\varphi \in hd(w')$.*

PROOF. Suppose that $[w]_C \rightarrow_{\mathbf{s}} [w']_C$. Thus, by Definition 5, there are epistemic states $u \in [w]_C$ and $u' \in [w']_C$ and a set of votes $\{v_a\}_{a \in \mathcal{A}}$ such that

1. $v_c = s_c([u]_C) = (\varphi, C)$ for each $c \in C$,
2. $(u, \{v_a\}_{a \in \mathcal{A}}, u') \in M$.

Recall that $S_C\varphi \in hd(w)$ by the assumption of the lemma. Thus, $hd(w) \vdash S_C S_C\varphi$ by Persistence axiom. Hence, $hd(w) \vdash K_C S_C S_C\varphi$ by Lemma 2. Then, $K_C S_C S_C\varphi \in hd(w)$ due to the maximality of the set $hd(w)$. Thus, $S_C S_C\varphi \in hd(u)$ by Lemma 17 and the assumption $u \in [w]_C$. Hence, $hd(u) \vdash S_C K_C S_C\varphi$ by Lemma 4. Thus, $S_C K_C S_C\varphi \in hd(u)$ due to the maximality of the set $hd(u)$. Hence, $K_C S_C\varphi \in hd(u')$ by Definition 11 and due to item 1. and item 2. above. Therefore, $S_C\varphi \in hd(w')$ by Lemma 17 and the assumption $u' \in [w']_C$. \square

LEMMA 20. *Let $w, w' \in W$ be epistemic states, $S_C\varphi \in \Phi$ be a formula, and $\mathbf{s} = \{s_c\}_{c \in C}$ be the strategy of coalition C such that $s_c([u]_C) = (\varphi, C)$ for each $c \in C$ and each $u \in W$. If $S_C\varphi \in hd(w)$ and $[w]_C \rightarrow_{\mathbf{s}}^* [w']_C$, then $\varphi \in hd(w')$.*

PROOF. Suppose that $S_C\varphi \in hd(w)$. Then it follows from Lemma 19 and the fact relation \rightarrow_s^* being a transitive reflexive closure of the relation \rightarrow_s that $S_C\varphi \in hd(w')$. Hence, $hd(w') \vdash K_C\varphi$ by Knowledge axiom. Thus, $hd(w') \vdash \varphi$ by Truth axiom. Therefore, $\varphi \in hd(w')$ due to the maximality of the set $hd(w')$. \square

LEMMA 21. *For any epistemic state $w \in W$, any formula $\neg S_C\psi \in hd(w)$, and any strategy $\mathbf{s} = \{s_c\}_{c \in C}$ of coalition C , there is a state $w' \in W$ such that $[w]_C \rightarrow_s [w']_C$ and $\psi \notin hd(w')$.*

PROOF. Let Y_0 be the following set of formulas

$$\{\varphi \mid \exists D \subseteq C [S_D\varphi \in hd(w) \wedge \forall d \in D (s_d([w]_C) = (\varphi, D))]\} \\ \cup \{\neg\psi\} \cup \{\chi \mid K_{\emptyset}\chi \in hd(w)\}.$$

We first show that set Y_0 is consistent. Suppose the opposite. Thus, there must exist formulae $\varphi_1, \dots, \varphi_n, \chi_1, \dots, \chi_m \in Y_0$ and subsets $D_1, \dots, D_n \subseteq C$ such that (i) $S_{D_i}\varphi_i \in hd(w)$ for each integer $i \leq n$, (ii) $s_d([w]_C) = (\varphi, D_i)$ for each $i \leq n$ and each $d \in D_i$, (iii) $K_{\emptyset}\chi_i \in hd(w)$ for each integer $i \leq m$, and (iv) set $\{\neg\psi, \varphi_1, \dots, \varphi_n, \chi_1, \dots, \chi_m\}$ is inconsistent. Without loss of generality we can assume that formulae $\varphi_1, \dots, \varphi_n$ are pairwise distinct.

CLAIM 1. *Sets D_i and D_j are disjoint for each $i \neq j$.*

PROOF OF CLAIM. Assume that $d \in D_i \cap D_j$, then $s_d([w]_C) = (\varphi_i, D_i)$ and $s_d([w]_C) = (\varphi_j, D_j)$. Hence, $(\varphi_i, D_i) = (\varphi_j, D_j)$. Thus, $\varphi_i = \varphi_j$, which contradicts to the assumption that formulae $\varphi_1, \dots, \varphi_n$ are pairwise distinct. \square

By Lemma 14, it follows from Claim 1 that set Y_0 is consistent. Let Y be a maximal consistent extension of Y_0 and w' be the sequence w, \emptyset, Y . In other words, sequence w' is an extension of sequence w by two additional elements: \emptyset and Y . Note that $w' \in W$ by Definition 8.

Let \top be any propositional tautology. For example, \top could be formula $\psi \rightarrow \psi$. Finally, let votes $\{v_a\}_{a \in \mathcal{A}}$ be defined as follows

$$v_a = \begin{cases} s_a([w]_C), & \text{if } a \in C, \\ (\top, \mathcal{A}), & \text{otherwise.} \end{cases} \quad (2)$$

CLAIM 2. *The following set is a subset of $hd(w')$:*

$$\{\varphi \mid \exists D \subseteq \mathcal{A} (S_D\varphi \in hd(w) \wedge \forall d \in D (v_d = (\varphi, D)))\}.$$

NOTE. The claim might appear to be true simply by the choice of set Y_0 . However, the claim allows D to be an arbitrary set of agents, not just a subset of coalition C as in the case of Y_0 .

PROOF OF CLAIM. Consider any φ and any $D \subseteq \mathcal{A}$ such that $S_D\varphi \in hd(w)$ and $v_d = (\varphi, D)$ for each $d \in D$. We need to show that $\varphi \in hd(w')$.

Case 1: $D \subseteq C$. Thus, $s_d([w]_C) = (\varphi, D)$ for each $d \in D$ by definition (2). Thus, $\varphi \in Y_0$ by the choice of set Y_0 . Therefore, $\varphi \in hd(w')$ because $Y_0 \subseteq Y = hd(w')$.

Case 2: There is $d \in D$ such that $d \notin C$. Thus, $v_d = (\top, \mathcal{A})$ by definition (2). Note that $v_d = (\varphi, D)$ by the choice of the set D . Thus, $(\top, \mathcal{A}) = (\varphi, D)$. Hence, formula φ is the tautology \top . Therefore, $\varphi \in hd(w')$ because set $hd(w')$ is maximal. \square

By Definition 11, Claim 2 implies that $(w, \{v_a\}_{a \in \mathcal{A}}, w') \in M$. Hence, $[w]_C \rightarrow_s [w']_C$ by Definition 5. Thus, $[w]_C \rightarrow_s^* [w']_C$ because relation \rightarrow_s^* is a transitive reflexive closure

of \rightarrow_s . To finish the proof of the lemma note that $\psi \notin hd(w')$ because set $hd(w')$ is consistent and $\neg\psi \in Y_0 \subseteq Y = hd(w')$. \square

DEFINITION 12. $\pi(p) = \{w \in W \mid p \in hd(w)\}$.

LEMMA 22. $w \Vdash \varphi$ iff $\varphi \in hd(w)$ for each formula $\varphi \in \Phi$.

PROOF. We prove by induction on the structural complexity of formula φ . If formula φ is a propositional variable, then the required follows from Definition 6 and Definition 12. The cases of formula φ being a negation or an implication follow by Definition 6 from the maximality and the consistency of the set $hd(w)$ in the standard way.

Let formula φ have the form $K_C\psi$.

(\Rightarrow) Suppose that $K_C\psi \notin hd(w)$. Thus, by Lemma 18, there is $w' \in W$ such that $w \sim_C w'$ and $\psi \notin hd(w')$. Hence, $w' \not\Vdash \psi$ by the induction hypothesis. Therefore, $w \not\Vdash K_C\psi$ by Definition 6.

(\Leftarrow) Assume that $K_C\psi \in hd(w)$. Consider any $w' \in W$ such that $w \sim_C w'$. By Definition 6, it suffices to show that $w' \Vdash \psi$. Indeed, $\psi \in hd(w')$ by Lemma 17. Therefore, by the induction hypothesis, $w' \Vdash \psi$.

Let formula φ have the form $S_C\psi$.

(\Rightarrow) Suppose that $S_C\psi \notin hd(w)$. Thus, $\neg S_C\psi \in hd(w)$ due to the maximality of the set $hd(w)$. Hence, by Lemma 21, for any strategy \mathbf{s} of coalition C there is a world $w' \in W$ such that $[w]_C \rightarrow_s [w']_C$ and $\psi \notin hd(w')$. Thus, by the induction hypothesis, for any strategy \mathbf{s} of coalition C there is a world $w' \in W$ such that $[w]_C \rightarrow_s [w']_C$ and $w' \not\Vdash \psi$. Then, $w \not\Vdash S_C\psi$ by Definition 6.

(\Leftarrow) Assume that $S_C\psi \in hd(w)$. Let $\mathbf{s} = \{(s, C)\}_{s \in C}$. Then, by Lemma 20, for any epistemic world $w' \in W$, if $[w]_C \rightarrow_s [w']_C$, then $\psi \in hd(w')$. Hence, by the induction hypothesis, for any epistemic world $w' \in W$, if $[w]_C \rightarrow_s [w']_C$, then $w' \Vdash \psi$. Therefore, $w \Vdash S_C\psi$ by Definition 6. \square

We are now ready to finish the proof of Theorem 2. Suppose that $\not\Vdash \varphi$. Let X_0 be any maximal consistent subset of Φ such that $\neg\varphi \in X_0$. Consider single-element sequence $w = X_0$. Note that $w \in W$ by Definition 8. Thus, $w \Vdash \neg\varphi$ by Lemma 22. Therefore, $w \not\Vdash \varphi$ by Definition 6. \square

7. CONCLUSION

In this paper we have introduced a bimodal logical system for reasoning about coalition power in epistemic transition systems. Unlike the existing axiomatic logical system [13], the proposed approach requires strategies to be executable and verifiable. These requirements manifest themselves in new logical principles: positive and negative strategic introspections and verifiability. The main technical result of our work is the soundness and the completeness of our system.

REFERENCES

- [1] Marc Pauly. *Logic for Social Software*. PhD thesis, Institute for Logic, Language, and Computation, 2001.
- [2] Marc Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.
- [3] Valentin Goranko. Coalition games and alternating temporal logics. In *Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge*,

- pages 259–272. Morgan Kaufmann Publishers Inc., 2001.
- [4] Wiebe van der Hoek and Michael Wooldridge. On the logic of cooperation and propositional control. *Artificial Intelligence*, 164(1):81 – 119, 2005.
 - [5] Stefano Borgo. Coalitions in action logic. In *20th International Joint Conference on Artificial Intelligence*, pages 1822–1827, 2007.
 - [6] Luigi Sauro, Jelle Gerbrandy, Wiebe van der Hoek, and Michael Wooldridge. Reasoning about action and cooperation. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '06, pages 185–192, New York, NY, USA, 2006. ACM.
 - [7] Thomas Ågotnes, Philippe Balbiani, Hans van Ditmarsch, and Pablo Seban. Group announcement logic. *Journal of Applied Logic*, 8(1):62 – 81, 2010.
 - [8] Thomas Ågotnes, Wiebe van der Hoek, and Michael Wooldridge. Reasoning about coalitional games. *Artificial Intelligence*, 173(1):45 – 79, 2009.
 - [9] Francesco Belardinelli. Reasoning about knowledge and strategies: Epistemic strategy logic. In *Proceedings 2nd International Workshop on Strategic Reasoning, SR 2014, Grenoble, France, April 5-6, 2014*, volume 146 of *EPTCS*, pages 27–33, 2014.
 - [10] Sara Miner More and Pavel Naumov. Calculus of cooperation and game-based reasoning about protocol privacy. *ACM Trans. Comput. Logic*, 13(3):22:1–22:21, August 2012.
 - [11] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
 - [12] Wiebe van der Hoek and Michael Wooldridge. Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica*, 75(1):125–157, 2003.
 - [13] Thomas Ågotnes and Natasha Alechina. Epistemic coalition logic: completeness and complexity. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 1099–1106. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
 - [14] Wojciech Jamroga and Thomas Ågotnes. Constructive knowledge: what agents can achieve under imperfect information. *Journal of Applied Non-Classical Logics*, 17(4):423–475, 2007.
 - [15] Wojciech Jamroga and Wiebe van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2-3):185–219, 2004.
 - [16] Johan Van Benthem. Games in dynamic-epistemic logic. *Bulletin of Economic Research*, 53(4):219–248, 2001.
 - [17] Henrik Sahlqvist. Completeness and correspondence in the first and second order semantics for modal logic. *Studies in Logic and the Foundations of Mathematics*, 82:110–143, 1975. (Proc. of the 3rd Scandinavian Logic Symposium, Uppsala, 1973).