Knowing the Price of Success

Rui Cao

University of British Columbia, Vancouver, BC, Canada

Pavel Naumov

Tulane University, New Orleans, Louisiana, USA

Abstract

If an agent, or a coalition of agents, knows that it has a strategy to achieve a certain outcome, it does not mean that the agent knows what the strategy is. Even if the agent knows what the strategy is, she might not know the price of executing this strategy.

The article considers modality "the coalition not only knows the strategy, but also knows an upper bound on the price of executing the strategy". The main technical result is a sound and complete bimodal logical system that describes the interplay between this modality and the distributed knowledge modality.

1. Introduction

In this article we study situations when an agent or a coalition of agents might have a strategy to achieve a certain goal at a price possibly unknown to the agent. An example of such a setting is an auction with *the right of first refusal*. In such an auction, an agent, who is not bidding at the auction, has a right to buy the item at the price offered by the winner of the auction. United States National Park Service previously used such auctions to award concession contracts with the right of first refusal given to the current concessioner [1]. Taiwanese government uses them to sell land [2]. In general, provisions with the right of first refusal are common in real estate, securities, and employment contracts [3].

Suppose that Alice and Bob are two bidders at first-price sealed-bid auction where Donald has the right of first refusal. In other words, Alice and Bob submit sealed bids and Donald has a right to match the highest bid after the bids are unsealed. We assume that all agents are rational and will not bid higher than their reservation prices.

Email addresses: `rui.cao@alumni.ubc.ca` (Rui Cao), `pgn2@cornell.edu` (Pavel Naumov)

If Alice's reservation price is lower than Donald's, then, no matter what Bob's reservation price is, she does not have a winning strategy because Donald might always decide to match the highest bid. We write this as

$$\neg\mathsf{S}_a(\text{“Alice wins the auction.”}),$$

where, in general, modal formula $\mathsf{S}_C\varphi$ stands for “coalition of agents C has a strategy to achieve φ”. The coalition power modality S_C was first introduced by Marc Pauly [4, 5]. His approach has been widely studied in the literature. Goranko studied relation between Marc Pauly's coalition logic and alternating-time temporal logics [6]. van der Hoek and Wooldridge introduced a related Coalition Logic of Propositional Control [7] in which an agent's actions are restricted to assigning Boolean values to variables. Borgo suggested a translation from the coalition logic to a fragment of the action logic [8]. Ågotnes, Balbiani, van Ditmarsch and Seban introduced a group announcement logic that combines the coalition logic and the logic of public announcements [9]. Ågotnes, van der Hoek, and Wooldridge added preference comparison to the coalition logic [10]. Goranko and Enqvist gave a complete axiomatization for modality “coalition can achieve a goal while leaving a possibility for the other agents to achieve another goal”. They call it “socially friendly” coalition power modality [11].

Suppose now that Alice's reservation price is the highest among all three of them. Then, she has a strategy to win the auction by submitting a bid equal to her reservation price:

$$\mathsf{S}_a(\text{“Alice wins the auction.”}).$$

However, since she does not know that her reservation price is the highest, she does not know that she has a winning strategy:

$$\neg\mathsf{K}_a\mathsf{S}_a(\text{“Alice wins the auction.”}).$$

In general, by formula $\mathsf{K}_C\varphi$ we denote that coalition C has a distributed knowledge of φ. A complete logical system for modalities S_C and K_C has been proposed by Ågotnes and Alechina [12].

Let us now consider the case when Alice has no reservation price and she is ready to submit a bid as high as needed to win over Bob and Donald, who still have reservation prices. In this setting Alice has a strategy to win the auction by bidding higher than Bob's and Donald's reservation prices and she knows this:

$$\mathsf{K}_a\mathsf{S}_a(\text{“Alice wins the auction.”}).$$

However, if she does not know their reservation prices, she does not know how much she should bid to win the auction. In other words, she does not *know how* to win the action. We write this as

$$\neg\mathsf{H}_a(\text{“Alice wins the auction.”}).$$

In general, modal formula $\mathsf{H}_C\varphi$ stands for “coalition of agents C *knows how* to achieve φ”. The distinction between an agent “having a strategy”, “knowing that the strategy exists", and "knowing what the strategy is" has been studied before.

Different terms have been used for what we call "know-how". While Jamroga and Ågotnes talked about "knowledge to identify and execute a strategy" [13], Jamroga and van der Hoek discussed "difference between an agent knowing that he has a suitable strategy and knowing the strategy itself" [14]. Van Benthem called such strategies "uniform" [15]. Broersen talked about "knowingly doing" [16], while Broersen, Herzig, and Troquard discussed modality "know they can do" [17]. Naumov and Tao called such strategies "executable" [18]. Wang talked about "knowing how" [19, 20].

Multiple logical systems for capturing properties of know-how have been proposed in the literature. Wang gave a complete axiomatization of "knowing how" as a binary modality [19, 20]. Li and Wang introduced a logical system for a ternary know-how modality with intermediate constraints [21]. These systems are for a single agent and they do not include a knowledge modality. Ågotnes and Alechina [22] proposed a complete axiomatization of an interplay between a single-agent knowledge modality and a coalition know-how modality to achieve a goal in one step. A complete modal logic that combines the distributed knowledge modality with the coalition know-how modality to *maintain* a goal indefinitely was given by Naumov and Tao [18]. A sound and complete logical system in a single-agent setting for know-how strategies to *achieve* a goal in multiple steps was developed by Fervari, Herzig, Li, and Wang [23]. A complete logical system that describes an interplay between modalities K_C, S_C, and H_C in one-shot was introduced by Naumov and Tao [24, 25]. They also proposed a logical system that describes the interplay between modalities K_C and H_C in the perfect recall setting [26] and another logical system for the second-order know-how [27].

Going back to our example, let us now assume that Alice and Bob have reservation prices and Donald does not. In this case, Donald has a winning strategy (use the right of first refusal to match the highest bid), he knows that he has a strategy, and he knows what the strategy is:

$$\mathsf{H}_d(\text{“Donald wins the auction”}). \tag{1}$$

Note that, before the auction starts, he does not *know the price* that he will have to pay. Not only he does not know the exact price, he cannot put any limit on how much he might end up paying if he commits to that strategy. He only knows that the price is a finite number. For this reason, we will write formula (1) as

$$\mathsf{H}_d^{\infty}(\text{“Donald wins the auction”}).$$

However, if we suppose that Donald knows reservation prices r_a and r_b of Alice and Bob respectively, then he would know that the strategy to win the auction will cost him at most $\max\{r_a, r_b\}$. We denote this by

$$\mathsf{H}_d^{\max\{r_a,r_b\}}(\text{“Donald wins the auction”}).$$

In general, by $\mathsf{H}^p_C\varphi$ we denote the fact that (a) coalition C has a strategy to achieve φ, (b) it knows that it has such strategy, (c) it knows what the strategy is, and (d) it knows that the total price of using this strategy for the whole coalition is at most p.

A complete logical system for modality S^p_C, "coalition C has a strategy with the price of at most p", has been proposed by Alechina, Logan, Nguyen, and Rakib [28] under name Resource-Bounded Coalition Logic (RBCL). Model checking for RBCL and resource-bounded Alternating-time Temporal Logic have been widely studied [29, 30, 31, 32, 33, 34, 35]. We proposed a complete logical system for modality "an agent a can achieve φ with a profit p on a budget b" [36].

Examples of real-life caps on cost of strategies are insurance deductibles and umbrella insurances. In computer science it is common to analyze worst-case algorithm execution time. At the same time, just like in our modality $\mathsf{H}^\infty_C\varphi$, upper limit on the cost or the time might not be known. For instance, a car driver might have a strategy to get from point A to point B, but the time required to do this might vary drastically in rush hours. A software development team might know how to finish a project, but it might not be able to estimate the time it would take. In many cases, an AI system acting in such situations would not only need to know that a limit to the potential cost of pursuing a strategy exists, but also to know what this limit is.

The contribution of this article is two-fold. First, we introduce a new epistemic modality "knowing the price of achieving a goal" H^p_C. This modality cannot be expressed as $\mathsf{K}_C\mathsf{S}^p_C$ or as some other combination of modalities S^p_C, K_C, and H_C, that have been studied before. Additionally, we propose a sound and complete logical system that captures the interplay between modalities K_C and H^p_C, where parameter p is either a non-negative real number or the infinity.

This work is closely related to two previous articles on know-how logics in this journal [25, 37]. These three works, however, study different aspects of know-how and use different constructions for the proof of completeness. In [25], Naumov and Tao described the interplay between coalition modalities K_C, S_C, and H_C. To prove the completeness, they proposed a novel "harmony" technique. In [37], they studied single-agent modalities K^r_a and H^r_a whose semantics is defined using metric spaces instead of indistinguishability relation. Formula $\mathsf{K}^r_a\varphi$ means that statement φ is true in a ball of radius r around the current state using the metric of agent a. Formula $\mathsf{H}^r_a\varphi$ means that single agent a has a strategy to achieve φ that works at each point of the same ball. To prove the completeness, they used "twin node" construction in which each state of the model has a twin state. Exactly the same formulae are satisfied in both twin states. In the current work we study coalition modalities K_C and H^p_C. Formula $\mathsf{K}_C\varphi$ means, just like in [25], that statement φ is true in the set of all states indistinguishable from the current state by coalition C. Formula H^p_C means that coalition C has a strategy that works in all points of this set and costs no more than p to the coalition. The semantics of this modality is a natural generalization of the semantics of modality H_C [25] and is very different from the metric space based semantics of H^p_C [37]. The proof of the completeness in the current work does not use harmony construction from [25]. Instead, it generalizes twin node construction of [37] to infinitely many twin nodes with unbounded costs, see Definition 4. This generalization was not needed in the proof of the completeness for Resource-Bounded Coalition Logic [28] because this logic includes neither knowledge nor infinite costs.

The rest of the article is structured as follows. In the next section we introduce the syntax and the formal semantics of our logical system and illustrate them using the introductory example. In Section 3 we list the axioms and the inference rules of our logical system, compare them to similar axioms in the literature, and give examples of formal proofs in our logical system. The soundness and the completeness of our system are shown in Section 4 and Section 5 respectively. Section 6 concludes.

2. Syntax and Semantics

Throughout the article we assume a fixed set of agents $\mathcal{A}$ and a fixed nonempty set of propositional variables. By a coalition we mean any finite subset of $\mathcal{A}$. In this section we define the syntax and the formal semantics of our logical system.

Definition 1. *Let Φ be the minimal set of formula such that*

1. $v \in \Phi$ *for each propositional variable* v,
2. $\neg\varphi, \varphi \to \psi, \mathsf{K}_C\varphi, \mathsf{H}^p_C\varphi \in \Phi$ *for any formulae* $\varphi, \psi \in \Phi$, *and any coalition* C, *where* p *is either a non-negative real number or* ∞.

In other words, language Φ is defined by grammar

$$\varphi := v \mid \neg\varphi \mid \varphi \to \varphi \mid \mathsf{K}_C\varphi \mid \mathsf{H}^p_C\varphi.$$

Boolean constant $\bot$ can be defined in our language in the standard way. For any sets X and Y, by X^Y we denote the set of all functions from Y to X.

Definition 2. *A tuple $(W, \{\sim_a\}_{a\in\mathcal{A}}, \Delta, \|\cdot\|, \varepsilon, M, \pi)$ is called a game, where*

1. W *is a set of states,*
2. $\sim_a$ *is an indistinguishability equivalence relation on the set of states* W *for each agent* $a \in \mathcal{A}$,
3. Δ *is a nonempty set called the* ***domain of actions****. A function that assigns an action to each agent in set* $\mathcal{A}$ *is called a* ***complete action profile****.*
4. *price* $\|d\|^a_w$ *of an action* $d \in \Delta$ *to an agent* $a \in \mathcal{A}$ *in a state* $w \in W$ *is a non-negative real number,*
5. $\varepsilon \in \Delta$ *is an zero-price action such that* $\|\varepsilon\|^a_w = 0$ *for each agent* $a \in \mathcal{A}$ *and each state* $w \in W$,
6. $M \subseteq W \times \Delta^{\mathcal{A}} \times W$ *is a relation, called* ***mechanism****, that satisfies the following* ***nontermination*** *property: for any state* $w \in W$ *and any complete action profile* $\delta \in \Delta^{\mathcal{A}}$ *there is at least one state* $u \in W$ *such that* $(w, \delta, u) \in M$,

7. *function π maps propositional variables into subsets of W.*

Next, we illustrate how our first example from the introduction, when all three agents have finite reservation prices, can be formalized as a game from Definition 2.

The introduction describes the game as an *extensive-form game* in which Alice and Bob submit simultaneous bids and Donald decides to match or not to match the highest bid only after the bids are unsealed. Here we would like to re-phrase this game as a *normal-form game* between the three players. To do this, we need Donald not to make the decision after the bids are unsealed, but to commit to a certain strategy before the auction starts. Such strategy could be described as setting the highest bid value that Donald is willing to match using the right of first refusal[1]. This value could be any number between zero and Donald's reservation price. Hence, when viewed as a normal-form game, the action consists in Alice and Bob choosing their bid values and Donald choosing the highest bid value that he is willing to match.

Definition 2 is general enough to allow system to transition between states multiple number of times as, for example, in repetitive normal-form games. In our original example the game is played only once. Thus, the states of the game naturally split into initial and final states. The initial states of the game are all possible triples of non-negative numbers (r_a, r_b, r_d) that represent Alice's, Bob's, and Donald's reservation prices. There are three final states that correspond to Alice, Bob, and Donald winning the auction.

Alice can not distinguish two initial states if she has the same reservation price r_a in both states. She can distinguish final states between themselves and final states from the initial states. The same is true for Bob and Donald with respect to their reservation prices r_b and r_d respectively.

So far, we said that Alice and Bob are choosing their bid values and Donald is choosing the highest bid value that he is willing to match. However, all three of these values cannot be higher than their respective reservation prices. Thus, the domain of action of each agent is different and it also changes from state to state. To reconcile this with Definition 2, we assume that the action consists in choosing the value as a fraction of the reservation price. For example, if Alice's reservation price is 70 and her bid is 35, then the action is 0.5. With this adjustment, the domain of actions for each agent in each state is the interval $[0, 1]$.

In our example, the price $\|d\|_w^a$ of an action $d \in [0, 1]$ for an agent a is equal to $d \cdot r_a$, where r_a is the reservation price of agent a in state w. In Definition 2, we assume that the price of an action depends on the agent invoking the action and the state in which the action is invoked. One might consider alternative definitions where the price of the action depends on the complete action profile, similarly to a utility function in game theory. In addition to the complete action profile, price also might depend on the outcome of the game. Although more

[1]Thus, we further restrict Donald's strategies to monotonic one: if Donald is willing to match a price, we assume that he is also willing to match any lower price.

general, these alternative versions of Definition 2 would not change the results in this article.

In case of the auction with the right of the first refusal, action $0 \in [0,1]$ is the zero-price action. The existence of a zero-price action, *known to each agent to be zero-price*, is an important assumption. H-Necessitation inference rule of our logical system is not valid without this assumption. The assumption of the existence of a zero-price action allows us to capture the informal meaning of the price to achieve a result. Indeed, consider a situation when an agent has only two actions priced at 2 and 5. Suppose that the action with price 5 leads to the desired outcome and the action with price 2 does not. We believe that the price of success in this case is 3, not 5. One might choose to define the price of success as the difference between the price of the needed action and the minimal price of the available actions. However, such minimum might not exist if the number of available actions is infinite. Thus, we have chosen to have an explicit zero-price action available to each agent. Alechina, Logan, Nguyen, and Rakib made the same choice in their Resource-Bounded Coalition Logic [28].

When describing our introductory example, we did not discuss what is the outcome of the auction if Alice and Bob submit equal bids and Donald decides not to exercise the right of the first refusal. Let us assume that in this situation the winner of the auction is chosen randomly from Alice and Bob. This means that the outcome of the game is not completely determined by the actions of the agents. To capture such situations we define mechanism M as a relation between the initial state, the complete action profile, and the outcome state. This approach allows multiple outcome states for the same initial state and the same complete action profile.

For any coalition C, we write $w \sim_C w'$ if $w \sim_a w'$ for each agent $a \in C$. Similarly, for any two functions f and g, we write $f =_C g$ if $f(a) = g(a)$ for each $a \in C$. By an *action profile of a coalition* C we mean any function from C to Δ.

Next is the key definition of this article. Its part (5) defines modality H^p_C. Informally, it says that a coalition C knows how to achieve φ at a price at most p if there is an action profile of the coalition C such that (a) total price of the action profile is at most p in any state w' indistinguishable to coalition C from the given state and (b) coalition C will achieve φ by using this action profile in any such state w'.

Definition 3. *For any formula $\varphi \in \Phi$ and any state $w \in W$ of a game $(W, \{\sim_a\}_{a \in \mathcal{A}}, \Delta, M, \pi)$, let satisfiability relation $w \Vdash \varphi$ be defined as follows*

1. $w \Vdash v$ *if* $w \in \pi(v)$ *where v is a propositional variable,*
2. $w \Vdash \neg\varphi$ *if* $w \nVdash \varphi$,
3. $w \Vdash \varphi \to \psi$ *if* $w \nVdash \varphi$ *or* $w \Vdash \psi$,
4. $w \Vdash \mathsf{K}_C\varphi$ *if* $w' \Vdash \varphi$ *for each* $w' \in W$ *such that* $w \sim_C w'$,
5. $w \Vdash \mathsf{H}^p_C\varphi$ *if there is an action profile* $s \in \Delta^C$ *such that*
   (a) *for each* $w' \in W$, *if* $w \sim_C w'$, *then* $\sum_{a \in C} \|s(a)\|^a_{w'} \le p$,
   (b) *for any two states* $w', u \in W$ *and any complete action profile* $\delta \in \Delta^{\mathcal{A}}$, *if* $w \sim_C w'$, $s =_C \delta$, *and* $(w', \delta, u) \in M$, *then* $u \Vdash \varphi$.

In particular, $w \Vdash \mathsf{K}_\varnothing\varphi$ means that formula φ is true in each state of the game and $w \Vdash \mathsf{H}^p_\varnothing\varphi$ means that φ will *unavoidably* happen in the next state.

If item (a) is removed from item 5 above, then the result will be the definition of the know-how modality previously used in [22, 18, 23, 24, 25, 26, 27]. Alternatively, if the condition $w \sim_C w'$ in items (a) and (b) is replaced with the condition $w = w'$, then the result is essentially equivalent to the definition of the modality S^p_C in Resource-Bounded Coalition Logic [28]. Note also that in item 5 the existential quantifier over action profile $s \in \Delta^C$ precedes the universal quantifiers over state $w' \in W$. If the order of quantifiers is changed, the definition will capture semantics of $\mathsf{K}_C\mathsf{S}^p_C\varphi$ instead of $\mathsf{H}^p_C\varphi$. In spite of the semantics in Definition 3 being so close to the know-how modality semantics and the resource bounded power modality semantics, as discussed in the introduction, the proof of the completeness of our logical system requires a significant modification to the earlier constructions.

3. Axioms and Inference Rules

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

1. Truth: $\mathsf{K}_C\varphi \to \varphi$,
2. Negative Introspection: $\neg\mathsf{K}_C\varphi \to \mathsf{K}_C\neg\mathsf{K}_C\varphi$,
3. Distributivity: $\mathsf{K}_C(\varphi \to \psi) \to (\mathsf{K}_C\varphi \to \mathsf{K}_C\psi)$,
4. Coalition Monotonicity: $\mathsf{K}_C\varphi \to \mathsf{K}_D\varphi$, where $C \subseteq D$,
5. Price Monotonicity: $\mathsf{H}^p_C\varphi \to \mathsf{H}^q_C\varphi$, where $p \leq q$,
6. Strategic Introspection: $\mathsf{H}^p_C\varphi \to \mathsf{K}_C\mathsf{H}^p_C\varphi$,
7. Cooperation: $\mathsf{H}^p_C(\varphi \to \psi) \to (\mathsf{H}^q_D\varphi \to \mathsf{H}^{p+q}_{C\cup D}\psi)$, where $C \cap D = \varnothing$,
8. Know-How of Empty Coalition: $\mathsf{H}^p_\varnothing\varphi \to \mathsf{H}^0_\varnothing\varphi$,
9. Knowledge of Empty Coalition: $\mathsf{K}_\varnothing\varphi \to \mathsf{H}^0_\varnothing\varphi$,
10. Unachievability of Falsehood: $\neg\mathsf{H}^\infty_C\bot$.

We write $\vdash \varphi$ if there is a finite sequence of formulae that ends with formula φ and each formula in the sequence is either one of the axioms or is obtained from the preceding formulae using the Modus Ponens, the K-Necessitation, or H-Necessitation

$$\frac{\varphi, \varphi \to \psi}{\psi}, \qquad \frac{\varphi}{\mathsf{K}_C\varphi}, \qquad \frac{\varphi}{\mathsf{H}^0_C\varphi}.$$

inference rule. If $\vdash \varphi$ then we say that φ is a theorem of our logical system.

We write $X \vdash \varphi$ if formula φ is provable from the theorems of our logical system and a set of additional formulae X using only the Modus Ponens inference rule. Then, $\varnothing \vdash \varphi$ is equivalent to $\vdash \varphi$.

The Truth, the Negative Introspection, the Distributivity, and the Monotonicity axioms are the standard axioms of the epistemic logic of distributed knowledge [38]. The Price Monotonicity axiom states that if a coalition knows how to achieve a goal at a price no more than p and $p \leq q$, then the same coalition also knows how to achieve the goal at a price no more than q. This axiom for modality S^p_C first appeared in [28]. The monotonicity principle for the *subscript* of the modality H^p_C is also true, but it is provable from the rest of the axioms of our logical system, see Lemma 3. The Strategic Introspection axiom states that if a coalition knows how to achieve a goal at a price no more than p, then the coalition knows that it knows how. This axiom first appeared in [22] and later was used in [18, 23, 24, 25, 26, 27].

The Cooperation axiom shows how the powers of two disjoint coalitions could be combined to achieve a common goal. This axiom for coalition in the form $\mathsf{S}_C(\varphi \to \psi) \to (\mathsf{S}_D\varphi \to \mathsf{S}_{C\cup D}\psi)$ was first proposed by Marc Pauly [4, 5]. The same axiom for modality S^p_C in the form $\mathsf{S}^p_C(\varphi \to \psi) \to (\mathsf{S}^q_D\varphi \to \mathsf{S}^{p+q}_{C\cup D}\psi)$ was given in [28], where they assume that p and q are not numbers, but vectors of "resources". The Cooperation axiom for know-how modality $\mathsf{H}_C(\varphi \to \psi) \to (\mathsf{H}_D\varphi \to \mathsf{H}_{C\cup D}\psi)$ first appeared in [22] and later was used in [18, 24, 25, 26, 27].

The Know-How of Empty Coalition axiom captures the fact that price of any action profile of the empty coalition is zero. A similar axiom for modality S^p_C comes from [28]. Finally, the Knowledge of Empty Coalition axiom states that if a statement is true in all states of the system, then the empty coalition can achieve it at zero price. This axiom for modality H_C first appeared in [18] and is also used in [24, 25, 26, 27].

The Unachievability of Falsehood axiom states that no coalition has a strategy to achieve a contradiction. This axiom captures nontermination assumption of Definition 2, item 6. The axiom first appeared in [26].

In the rest of the section we prove five statements about our logical system that will be used later in the proof of the completeness. The third of these statements, Lemma 3, is the monotonicity principle on subscript of modality H^p_C mentioned earlier in this section.

Lemma 1 (deduction). *If $X, \varphi \vdash \psi$, then $X \vdash \varphi \to \psi$.*

Proof. Since $X, \varphi \vdash \psi$ refers to the provability without the use of the K-Necessitation and H-Necessitation inference rules, the standard proof of deduction lemma for propositional logic [39, Proposition 1.9] applies to our system as well. ⊠

Lemma 2 (Lindenbaum). *Any consistent set of formulae can be extended to a maximal consistent set of formulae.*

Proof. The standard proof of Lindenbaum's lemma applies here [39, Proposition 2.14]. However, since the formulae in our logical system use real numbers in superscript, the set of formulae is uncountable. Thus, the proof of Lindenbaum's lemma in our case relies on the Axiom of Choice. ⊠

Lemma 3. $\vdash \mathsf{H}^p_C\varphi \to \mathsf{H}^p_D\varphi$, *where $C \subseteq D$.*

Proof. Note that $\varphi \to \varphi$ is a propositional tautology. Hence, $\vdash \mathsf{H}^0_{D\setminus C}(\varphi \to \varphi)$ by the Necessitation inference rule.

At the same time, $\vdash \mathsf{H}^0_{D\setminus C}(\varphi \to \varphi) \to (\mathsf{H}^p_C\varphi \to \mathsf{H}^p_D\varphi)$ by the Cooperation axiom because of the assumption $C \subseteq D$. Therefore, $\vdash \mathsf{H}^p_C\varphi \to \mathsf{H}^p_D\varphi$ by the Modus Ponens inference rule. ⊠

The next lemma is a generalization of the Distributivity and the Cooperation axioms.

Lemma 4. *If $\varphi_1,\dots,\varphi_n \vdash \psi$, then*

1. $\mathsf{K}_C\varphi_1,\dots,\mathsf{K}_C\varphi_n \vdash \mathsf{K}_C\psi$,
2. $\mathsf{H}^{p_1}_{C_1}\varphi_1,\dots,\mathsf{H}^{p_n}_{C_n}\varphi_n \vdash \mathsf{H}^{p_1+\dots+p_n}_{C_1\cup\dots\cup C_n}\psi$, *where sets C_1,...,C_n are pairwise disjoint.*

Proof. To prove statement (2), apply n times Lemma 1 to the assumption $\varphi_1,\dots,\varphi_n \vdash \psi$. Then, $\vdash \varphi_1 \to (\dots \to (\varphi_n \to \psi))$. Thus,

$$\vdash \mathsf{H}^0_{\varnothing}(\varphi_1 \to (\dots \to (\varphi_n \to \psi))),$$

by the H-Necessitation inference rule. Hence,

$$\vdash \mathsf{H}^{p_1}_{C_1}\varphi_1 \to \mathsf{H}^{p_1}_{C_1}(\varphi_2 \dots \to (\varphi_n \to \psi))$$

by the Cooperation axiom and the Modus Ponens inference rule. Then,

$$\mathsf{H}^{p_1}_{C_1}\varphi_1 \vdash \mathsf{H}^{p_1}_{C_1}(\varphi_2 \dots \to (\varphi_n \to \psi))$$

by the Modus Ponens inference rule. Thus, again by the Cooperation axiom and Modus Ponens,

$$\mathsf{H}^{p_1}_{C_1}\varphi_1 \vdash \mathsf{H}^{p_2}_{C_2}\varphi_2 \to \mathsf{H}^{p_1+p_2}_{C_1\cup C_2}(\varphi_3 \dots \to (\varphi_n \to \psi)).$$

Therefore, $\mathsf{H}^{p_1}_{C_1}\varphi_1,\dots,\mathsf{H}^{p_n}_{C_n}\varphi_n \vdash \mathsf{H}^{p_1+\dots+p_n}_{C_1\cup\dots\cup C_n}\psi$, by repeating the last two steps $n-2$ times.

The proof of the first statement is similar, but it uses the K-Necessitation inference rule and the Distributivity axiom instead of the H-Necessitation inference rule and the Cooperation axiom. ⊠

Our final statement is the well-known positive introspection principle for the distributed knowledge. We reproduce its proof here to keep the article self-contained.

Lemma 5. $\vdash \mathsf{K}_C\varphi \to \mathsf{K}_C\mathsf{K}_C\varphi$.

Proof. Formula $\mathsf{K}_C\neg\mathsf{K}_C\varphi \to \neg\mathsf{K}_C\varphi$ is an instance of the Truth axiom. Thus, $\vdash \mathsf{K}_C\varphi \to \neg\mathsf{K}_C\neg\mathsf{K}_C\varphi$ by contraposition. Hence, taking into account the following instance of the Negative Introspection axiom: $\neg\mathsf{K}_C\neg\mathsf{K}_C\varphi \to \mathsf{K}_C\neg\mathsf{K}_C\neg\mathsf{K}_C\varphi$, we have

$$\vdash \mathsf{K}_C\varphi \to \mathsf{K}_C\neg\mathsf{K}_C\neg\mathsf{K}_C\varphi. \tag{2}$$

At the same time, $\neg\mathsf{K}_C\varphi \to \mathsf{K}_C\neg\mathsf{K}_C\varphi$ is an instance of the Negative Introspection axiom. Thus, $\vdash \neg\mathsf{K}_C\neg\mathsf{K}_C\varphi \to \mathsf{K}_C\varphi$ by the law of contrapositive in the propositional logic. Hence, by the K-Necessitation inference rule, $\vdash \mathsf{K}_C(\neg\mathsf{K}_C\neg\mathsf{K}_C\varphi \to \mathsf{K}_C\varphi)$. Thus, by the Distributivity axiom and the Modus Ponens inference rule, $\vdash \mathsf{K}_C\neg\mathsf{K}_C\neg\mathsf{K}_C\varphi \to \mathsf{K}_C\mathsf{K}_C\varphi$. The latter, together with statement (2), implies the statement of the lemma by propositional reasoning. $\boxtimes$

4. Soundness

The proof of soundness of the Truth, the Negative Introspection, the Distributivity, and the Coalition Monotonicity axioms and the K-Necessitation inference rule is identical to the proof of their soundness in the epistemic logic of the distributed knowledge [38]. The soundness of the H-Necessitation inference rule follows from the existence of a zero-price action for each member of any coalition C. Below we prove the soundness of each of the remaining axioms as a separate lemma. We state the soundness of our logical system as Theorem 1 in the end of this section.

Lemma 6. *If $w \Vdash \mathsf{H}^p_C\varphi$ and $p \leq q$, then $w \Vdash \mathsf{H}^q_C\varphi$.*

Proof. By Definition 3, assumption $w \Vdash \mathsf{H}^p_C\varphi$ implies that there is an action profile $s \in \Delta^C$ of coalition C such that

1. for each $w' \in W$, if $w \sim_C w'$, then $\sum_{a\in C}\|s(a)\|^a_{w'} \leq p$,
2. for any two states $w', u \in W$ and any complete action profile $\delta \in \Delta^{\mathcal{A}}$, if $w \sim_C w'$, $s =_C \delta$, and $(w', \delta, u) \in M$, then $u \Vdash \varphi$.

Hence, by the assumption $p \leq q$ and the first statement above, for each $w' \in W$, if $w \sim_C w'$, then $\sum_{a\in C}\|s(a)\|^a_{w'} \leq p \leq q$. Therefore, $w \Vdash \mathsf{H}^q_C\varphi$ by Definition 3. $\boxtimes$

Lemma 7. *If $w \Vdash \mathsf{H}^p_C\varphi$, then $w \Vdash \mathsf{K}_C\mathsf{H}^p_C\varphi$.*

Proof. By Definition 3, assumption $w \Vdash \mathsf{H}^p_C\varphi$ implies that there is an action profile $s \in \Delta^C$ of coalition C such that

(a) for each $w' \in W$, if $w \sim_C w'$, then $\sum_{a\in C}\|s(a)\|^a_{w'} \leq p$,
(b) for any two states $w', u \in W$ and any complete action profile $\delta \in \Delta^{\mathcal{A}}$, if $w \sim_C w'$, $s =_C \delta$, and $(w', \delta, u) \in M$, then $u \Vdash \varphi$.

Consider any state $v \in W$ such that $w \sim_C v$. By Definition 3, it suffices to prove that $v \Vdash \mathsf{H}^p_C\varphi$. Again by Definition 3, it suffices to show that

1. for each $v' \in W$, if $v \sim_C v'$, then $\sum_{a\in C}\|s(a)\|^a_{v'} \leq p$,
2. for any two states $v', u \in W$ and any complete action profile $\delta \in \Delta^{\mathcal{A}}$, if $v \sim_C v'$, $s =_C \delta$, and $(v', \delta, u) \in M$, then $u \Vdash \varphi$.

Since $w \sim_C v \sim_C v'$, the first statement follows from statement (a). Similarly, the second statement follows from statement (b). $\boxtimes$

Lemma 8. *If $w \Vdash \mathsf{H}^p_C(\varphi \to \psi)$ and $w \Vdash \mathsf{H}^q_D\varphi$, then $w \Vdash \mathsf{H}^{p+q}_{C\cup D}\psi$, where $C \cap D = \varnothing$.*

PROOF. By Definition 3, assumption $w \Vdash \mathsf{H}^p_C(\varphi \to \psi)$ implies that there is an action profile $s_1 \in \Delta^C$ of coalition C such that

(a) for each $w' \in W$, if $w \sim_C w'$, then $\sum_{a\in C} \|s_1(a)\|^a_{w'} \le p$,
(b) for any two states $w', u \in W$ and any $\delta \in \Delta^{\mathcal{A}}$, if $w \sim_C w'$, $s_1 =_C \delta$, and $(w', \delta, u) \in M$, then $u \Vdash \varphi \to \psi$.

Similarly, assumption $w \Vdash \mathsf{H}^q_D\varphi$ implies that there is an action profile $s_2 \in \Delta^D$ of coalition D such that

(c) for each $w' \in W$, if $w \sim_D w'$, then $\sum_{a\in D} \|s_2(a)\|^a_{w'} \le q$,
(d) for any two states $w', u \in W$ and any $\delta \in \Delta^{\mathcal{A}}$, if $w \sim_D w'$, $s_2 =_D \delta$, and $(w', \delta, u) \in M$, then $u \Vdash \varphi$.

Consider the following action profile $s \in \Delta^{C\cup D}$ of coalition $C \cup D$:

$$s(a) = \begin{cases} s_1(a), & \text{if } a \in C,\\ s_2(a), & \text{if } a \in D.\end{cases}$$

The action profile s is well-defined due to the assumption of the lemma that sets C and D are disjoint. Note that $s =_C s_1$ and $s =_D s_2$. Thus, statements (a), (b), (c), and (d) by Definition 3 imply that

1. for each $w' \in W$, if $w \sim_{C\cup D} w'$, then
$$\sum_{a\in C\cup D} \|s(a)\|^a_{w'} \le \sum_{a\in C} \|s_1(a)\|^a_{w'} + \sum_{a\in D} \|s_2(a)\|^a_{w'} \le p + q.$$
2. for any two states $w', u \in W$ and any $\delta \in \Delta^{\mathcal{A}}$, if $w \sim_{C\cup D} w'$, $s =_{C\cup D} \delta$, and $(w', \delta, u) \in M$, then $u \Vdash \psi$.

Therefore, $w \Vdash \mathsf{H}^{p+q}_{C\cup D}\psi$ again by Definition 3. $\boxtimes$

Lemma 9. *If $w \Vdash \mathsf{H}^p_\varnothing\varphi$, then $w \Vdash \mathsf{H}^0_\varnothing\varphi$.*

PROOF. By Definition 3, assumption $w \Vdash \mathsf{H}^p_\varnothing\varphi$ implies that there is an action profile $s \in \Delta^\varnothing$ of the empty coalition such that, *in particular*,

1. for any two states $w', u \in W$ and any complete action profile $\delta \in \Delta^{\mathcal{A}}$, if $w \sim_\varnothing w'$, $s =_\varnothing \delta$, and $(w', \delta, u) \in M$, then $u \Vdash \varphi$.

At the same time, $\sum_{a\in\varnothing} \|s(a)\|^a_{w'} = 0$, for each $w' \in W$. Therefore, $w \Vdash \mathsf{H}^0_\varnothing\varphi$ by Definition 3. $\boxtimes$

Lemma 10. *If $w \Vdash \mathsf{K}_\varnothing\varphi$, then $w \Vdash \mathsf{H}^0_\varnothing\varphi$.*

PROOF. By Definition 3, assumption $w \Vdash \mathsf{K}_\varnothing\varphi$ implies that $u \Vdash \varphi$ for each state $u \in W$. Let $s \in \Delta^\varnothing$ be the unique action profile of the empty coalition. Then,

1. $\sum_{a\in\varnothing} \|s(a)\|^a_{w'} = 0$, for each $w' \in W$,
2. for any two states $w', u \in W$ and any $\delta \in \Delta^{\mathcal{A}}$, if $w \sim_\varnothing w'$, $s =_\varnothing \delta$, and $(w', \delta, u) \in M$, then $u \Vdash \varphi$.

Therefore, $w \Vdash \mathsf{H}^0_\varnothing\varphi$ by Definition 3. ⊠

Lemma 11. $w \nVdash \mathsf{H}^\infty_C\bot$.

PROOF. Suppose that $w \Vdash \mathsf{H}^\infty_C\bot$. Thus, by Definition 3, there is an action profile $s \in \Delta^C$ such that

1. for any two states $w', u \in W$ and any complete action profile $\delta \in \Delta^{\mathcal{A}}$, if $w \sim_C w'$, $s =_C \delta$, and $(w', \delta, u) \in M$, then $u \Vdash \bot$.

By item 2 of Definition 2, the domain of actions Δ is not empty. Let d be any action in set Δ. Consider action profile

$$\delta(a) = \begin{cases} s(a), & \text{if } a \in C, \\ d, & \text{otherwise.} \end{cases}$$

Note that $s =_C \delta$. Choose w' to be state w. By the nontermination condition of Definition 2, item 6, there is a state $u \in W$ such that $(w, \delta, u) \in M$. Therefore, $u \Vdash \bot$ by item 1 above, which is a contradiction. ⊠

Theorem 1 (soundness). *If $Y \vdash \varphi$, then for each state w of each game, if $w \Vdash \chi$ for each formula $\chi \in Y$, then $w \Vdash \varphi$.* ⊠

5. Completeness

In this section we prove the completeness of our logical system. We start by defining, for any maximal consistent set of formulae X_0, the canonical game $G(X_0) = (W, \{\sim_a\}_{a\in\mathcal{A}}, \Delta, \|\cdot\|, \varepsilon, M, \pi)$. Since the language of our logical system contains distributed knowledge modality K_C, as a part of our construction of the canonical game $G(X_0)$, we indirectly construct a canonical model for the epistemic logic of distributed knowledge. Constructing such a model is a significantly more complicated task than constructing a canonical model for the epistemic logic of individual knowledge. In the case of the individual knowledge, states of the canonical model are the maximal consistent sets of formulae. States w and w' are in the relation $\sim_a$ if these two sets of formulae have the same K_a-formulae. This construction needs to be modified significantly in order to be used for the distributed knowledge. Indeed, if $w \sim_{\{a,b\}} w'$, then the construction only guarantees that sets w and w' have the same K_a- and K_b-formulae. It does not guarantee that sets w and w' have the same $\mathsf{K}_{\{a,b\}}$-formulae. To guarantee that w and w' also have the same $\mathsf{K}_{\{a,b\}}$-formulae, we use a tree construction described below. Lemma 13 shows that the tree construction achieves the intended goal.

Since Resource-Bounded Coalition Logic semantics does not incorporate the distributed knowledge in any form, the proof of the completeness for RBCL does not use the tree construction [28]. The tree construction has been previously used in other works on know-how logics with distributed knowledge [18, 24, 25, 26, 27]. The proof of completeness below significantly extends the tree construction from these previous works by an introduction of an extra label r_i on each node of the tree, see Definition 4 below. Presence of this new parameter creates infinitely many otherwise indistinguishable nodes in the same tree. Such infinite sets of indistinguishable nodes are needed to model know-how strategies with unbounded costs.

Definition 4. *The set of states W consists of all finite sequences $(X_0, r_0), C_1, (X_1, r_1), \dots, C_n, (X_n, r_n)$ such that*

1. $n \geq 0$,
2. *X_i is a maximal consistent set of formulae for each $i \geq 1$,*
3. *$r_0 = 0$ and $r_i \in [0, +\infty)$ for each $i \geq 1$,*
4. *C_i is a finite subset of $\mathcal{A}$ for each $i \geq 1$,*
5. *$\{\varphi \in \Phi \mid \mathsf{K}_{C_{i-1}}\varphi \in X_{i-1}\} \subseteq X_i$ for each $i \geq 1$.*

If w is state $(X_0, r_0), C_1, (X_1, r_1), \dots, C_n, (X_n, r_n)$, then by $X(w)$ and $r(w)$ we mean set X_n and real number r_n respectively. Also, for any sequence $x = x_1, \dots, x_n$ and an element y, by $x :: y$ we mean the sequence $x_1, \dots, x_n, y$.

We define an undirected labeled tree structure on the set of states W by saying that node

$$(X_0, r_0), C_1, (X_1, r_1), \dots, C_{n-1}, (X_{n-1}, r_{n-1})$$

and node

$$(X_0, r_0), C_1, (X_1, r_1), \dots, C_{n-1}, (X_{n-1}, r_{n-1}), C_n, (X_n, r_n)$$

are adjacent. The edge between these two nodes is labeled with each agent in coalition C_n, see Figure 1.

Definition 5. *For any states $w_1, w_2 \in W$, let $w_1 \sim_a w_2$ if all edges along the unique simple path between nodes w_1 and w_2 are labeled with agent a.*

Lemma 12. *Relation $\sim_a$ is an equivalence relation on set W.* ⊠

Many previous works on the logics of coalition power [18, 24, 25, 26, 27, 40] use different modifications of the same basic idea behind the mechanism construction: the domain of actions is the set of all formulae, and $(w, \delta, w') \in M$ iff for each formula $\mathsf{H}_C\varphi \in X(w)$ if $\delta(a) = \varphi$ for each $a \in C$, then $\varphi \in X(w')$. Below we adopt this idea for our setting.

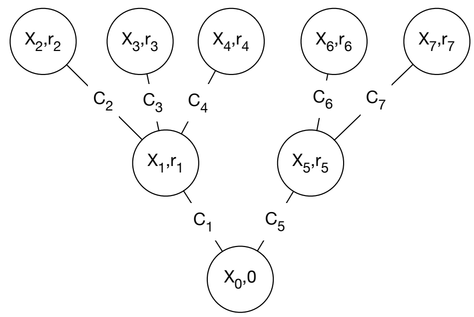

Figure 1: A fragment of the tree formed by states.

Definition 6. *Set Δ consists of a fixed zero-price action ε in addition to all triples (φ, C, p) such that $\varphi \in \Phi$ is a formula, C is a nonempty coalition, and p is either a non-negative real number or positive infinity ∞.*

Next, we define the price function in the canonical model. There are several cases to consider. The price of the zero-price action is zero. If $p < \infty$, then the price of action (φ, C, p) is set to $\frac{p}{|C|}$ so that if each member of coalition C chooses this action, then the total price to the whole coalition C is p. Finally, if $p = \infty$, then we want the price of the action to be finite but to have no upper bound in $\sim_C$ equivalence class of the current state. In other words, for any real number r we want there to be a state indistinguishable to coalition C from the current state, where the price of action (φ, C, p) to each agent is higher than r.

Definition 7. *For each action $d \in \Delta$, each agent $a \in \mathcal{A}$, and each state $w \in W$, the price function $\|d\|^a_w$ is defined as follows:*

$$\begin{aligned}
\|\varepsilon\|^a_w &= 0,\\
\|(\varphi, C, p)\|^a_w &= \begin{cases} \dfrac{p}{|C|}, & \text{if } p \in [0, +\infty),\\ r(w), & \text{if } p = \infty. \end{cases}
\end{aligned}$$

Note that according to the above definition, price $\|d\|^a_w$ in the canonical game $G(X_0)$ does not depend on the agent a. Thus, in the rest of this section we write $\|d\|^a_w$ simply as $\|d\|_w$.

Definition 8. *Mechanism M is the set of all triples $(w, \delta, u) \in W \times \Delta^{\mathcal{A}} \times W$ such that for each $\mathsf{H}^p_C\varphi \in X(w)$, if $\delta(a) = (\varphi, C, p)$ for each $a \in C$, then $\varphi \in X(u)$.*

Definition 9. $\pi(v) = \{w \in W \mid v \in X(w)\}$.

This concludes the definition of the canonical game $G(X_0)$. In Lemma 16 we prove the nontermination condition of Definition 2, item 6 is satisfied for this game. Next, we show that the tree construction works as intended, see introductory paragraph of Section 5.

Lemma 13. *If $w \sim_C w'$, then $\mathsf{K}_C\varphi \in X(w)$ iff $\mathsf{K}_C\varphi \in X(w')$.*

PROOF. By Definition 5, assumption $w \sim_C w'$ implies that each edge along the unique path between nodes w and w' is labeled with all agents in coalition C. Thus, it suffices to show that $\mathsf{K}_C\varphi \in X(w)$ iff $\mathsf{K}_C\varphi \in X(w')$ for any two *adjacent* nodes w and w' along this path. Indeed, without loss of generality, let

$$\begin{aligned} w &= (X_0, r_0), C_1, (X_1, r_1), \ldots, C_{n-1}, (X_{n-1}, r_{n-1}) \\ w' &= (X_0, r_0), C_1, (X_1, r_1), \ldots, C_{n-1}, (X_{n-1}, r_{n-1}), C_n, (X_n, r_n). \end{aligned}$$

The assumption that the edge between nodes w and w' is labeled with all agents in coalition C implies that $C \subseteq C_n$. We show next that $\mathsf{K}_C\varphi \in X(w)$ iff $\mathsf{K}_C\varphi \in X(w')$.

$(\Rightarrow)$: Suppose that $\mathsf{K}_C\varphi \in X(w) = X_{n-1}$. Thus, $X_{n-1} \vdash \mathsf{K}_C\mathsf{K}_C\varphi$ by Lemma 5. Hence, $X_{n-1} \vdash \mathsf{K}_{C_n}\mathsf{K}_C\varphi$ by the Monotonicity axiom and because $C \subseteq C_n$. Hence, $\mathsf{K}_{C_n}\mathsf{K}_C\varphi \in X_{n-1}$ because set X_{n-1} is maximal. Then, $\mathsf{K}_C\varphi \in X_n = X(w')$ by Definition 4.

$(\Leftarrow)$: Suppose that $\mathsf{K}_C\varphi \notin X(w) = X_{n-1}$. Thus, $\neg\mathsf{K}_C\varphi \in X_{n-1}$ because set X_{n-1} is maximal. Hence, $X_{n-1} \vdash \mathsf{K}_C\neg\mathsf{K}_C\varphi$ by the Negative Introspection axiom. Hence, $X_{n-1} \vdash \mathsf{K}_{C_n}\neg\mathsf{K}_C\varphi$ by the Monotonicity axiom and because $C_n \subseteq C$. Hence, $\mathsf{K}_{C_n}\neg\mathsf{K}_C\varphi \in X_{n-1}$ because set X_{n-1} is maximal. Then, $\neg\mathsf{K}_C\varphi \in X_n$ by Definition 4. Therefore, $\mathsf{K}_C\varphi \notin X_n = X(w')$ because set X_n is consistent. $\boxtimes$

The next important milestone in the proof of the completeness is Lemma 17, which sometimes is referred to as the "truth lemma". This lemma connects the syntax and the semantics of our logical system. However, to prove Lemma 17, we need the following two auxiliary results.

Lemma 14. *For any real number $p \geq 0$, any formula $\mathsf{K}_C\varphi \in \Phi$, and any state $w \in W$, if $\neg\mathsf{K}_C\varphi \in X(w)$, then there is a state $w' \in W$ such that $w \sim_C w'$, $r(w') > p$, and $\neg\varphi \in X(w')$.*

PROOF. Consider set $Y = \{\neg\varphi\} \cup \{\tau \in \Phi \mid \mathsf{K}_C\tau \in X(w)\}$. First, we prove that set Y is consistent. Suppose the opposite. Thus, there are formulae $\mathsf{K}_C\psi_1, \ldots, \mathsf{K}_C\psi_n \in X(w)$ such that $\psi_1, \ldots, \psi_n \vdash \varphi$. Hence, $\mathsf{K}_C\psi_1, \ldots, \mathsf{K}_C\psi_n \vdash \mathsf{K}_C\varphi$ by Lemma 4. Thus, $X(w) \vdash \mathsf{K}_C\varphi$ by the choice of formulae $\mathsf{K}_C\psi_1, \ldots, \mathsf{K}_C\psi_n$. Then, $\neg\mathsf{K}_C\varphi \notin X(w)$ because set $X(w)$ is consistent, which contradicts the assumption of the lemma. Therefore, set Y is consistent.

By Lemma 2, there is a maximal consistent extension $\hat{Y}$ of set Y. Let w' be the sequence $w :: C :: (\hat{Y}, p+1)$. Then, $w' \in W$ by Definition 4. Also, $w \sim_C w'$ by Definition 5 and $r(w') = p + 1 > p$. Finally, $\neg\varphi \in Y \subseteq \hat{Y} = X(w')$ by the choice of sets $Y, \hat{Y}$ and the choice of sequence w'. $\boxtimes$

Lemma 15. *If $\neg\mathsf{H}^p_C\varphi \in X(w)$, then for each action profile $s \in \Delta^C$ of coalition C there is a complete action profile $\delta \in \Delta^{\mathcal{A}}$ such that $s =_C \delta$ and at least one of the following is true*

1. *there is a state* $w' \in W$ *such that* $\sum_{a\in C} \|\delta(a)\|_{w'} > p$ *and* $w \sim_C w'$, *or*
2. *there is* $u \in W$ *such that* $(w,\delta,u) \in M$ *and* $\neg\varphi \in X(u)$.

PROOF. Let action profile $\delta \in \Delta^{\mathcal{A}}$ be defined as

$$\delta(a) = \begin{cases} s(a), & \text{if } a \in C, \\ \varepsilon, & \text{otherwise.} \end{cases} \tag{3}$$

Case I: $p \neq \infty$ and there is a formula $\mathsf{H}^{\infty}_{D}\psi \in X(w)$ and an agent $a_0 \in D$, such that $D \subseteq C$ and $\forall a \in D(s(a) = (\varphi, D, \infty))$. Thus,

$$s(a_0) = (\varphi, D, \infty). \tag{4}$$

Formula $\mathsf{K}_C\bot \to \bot$ is an instance of the Truth axiom. Thus, $\vdash \neg\mathsf{K}_C\bot$ by the laws of propositional reasoning. Hence, $\neg\mathsf{K}_C\bot \in X(w)$ because set $X(w)$ is maximal. Then, by Lemma 14, there is a state $w' \in W$ such that $w \sim_C w'$ and $r(w') > p$. Hence, by equation (3), equation (4), and Definition 7,

$$\begin{aligned} \sum_{a\in C} \|\delta(a)\|_{w'} &\geq \|\delta(a_0)\|_{w'} = \|s(a_0)\|_{w'} \\ &= \|(\varphi, D, \infty))\|_{w'} = r(w') > p, \end{aligned}$$

which proves item 1 of the lemma.
Case II: $p = \infty$ or there is no formula $\mathsf{H}^{\infty}_{D}\psi \in X(w)$ such that D is a *nonempty* subset of C and $\forall a \in D(s(a) = (\varphi, D, \infty))$. Consider set

$$\begin{aligned} Y \;=\; & \{\neg\varphi\} \cup \{\psi \in \Phi \mid \mathsf{H}^{q}_{D}\psi \in X(w), \varnothing \neq D \subseteq C, \\ & \qquad\qquad \forall a \in D(s(a) = (\psi, D, q))\} \\ & \cup \{\chi \in \Phi \mid \mathsf{H}^{t}_{\varnothing}\chi \in X(w)\} \cup \{\tau \in \Phi \mid \mathsf{K}_{\varnothing}\tau \in X(w)\}. \end{aligned}$$

Suppose set Y is not consistent. Thus, there are formulae

$$\mathsf{H}^{q_1}_{D_1}\psi_1, \dots, \mathsf{H}^{q_n}_{D_n}\psi_n \in X(w), \tag{5}$$

$$\mathsf{H}^{t_1}_{\varnothing}\chi_1, \dots, \mathsf{H}^{t_m}_{\varnothing}\chi_m \in X(w), \tag{6}$$

and

$$\mathsf{K}_{\varnothing}\tau_1, \dots, \mathsf{K}_{\varnothing}\tau_k \in X(w) \tag{7}$$

such that

$$\begin{aligned} & q_i \in [0, +\infty] \text{ for all } i \geq 1, \\ & t_i \in [0, +\infty] \text{ for all } i \geq 1, \\ & \varnothing \neq D_i \subseteq C \text{ for all } i \geq 1, && (8) \\ & s(a) = (\psi_i, D_i, q_i) \text{ for all } i \geq 1, a \in D_i, && (9) \end{aligned}$$

and

$$\psi_1, \dots, \psi_n, \chi_1, \dots, \chi_m, \tau_1, \dots, \tau_k \vdash \varphi. \tag{10}$$

Without loss of generality we can assume that formulae $\psi_1,\dots,\psi_n$ are distinct. Thus, sets $D_1,\dots,D_n$ are disjoint due to the statement (9).

Statement (10), by the Lemma 4, implies that

$$\mathsf{H}^{q_1}_{D_1}\psi_1,\dots,\mathsf{H}^{q_n}_{D_n}\psi_n,\mathsf{H}^{0}_{\varnothing}\chi_1,\dots,\mathsf{H}^{0}_{\varnothing}\chi_m,\mathsf{H}^{0}_{\varnothing}\tau_1,\dots,\mathsf{H}^{0}_{\varnothing}\tau_k\vdash\mathsf{H}^{q_1+\dots+q_n}_{D_1\cup\dots\cup D_n}\varphi.$$

Thus, by statement (5), statement (6), the Know-How of Empty Coalition axiom, statement (7), the Knowledge of Empty Coalition axiom, and the Modus Ponens inference rule,

$$X(w)\vdash\mathsf{H}^{q_1+\dots+q_n}_{D_1\cup\dots\cup D_n}\varphi.$$

Then, by statement (8) and Lemma 3,

$$X(w)\vdash\mathsf{H}^{q_1+\dots+q_n}_{C}\varphi. \tag{11}$$

Next, we observe that

$$q_1+\dots+q_n>p. \tag{12}$$

Indeed, suppose the opposite. Thus, $q_1+\dots+q_n\le p$. Then, statement (11) implies that $X(w)\vdash\mathsf{H}^{p}_{C}\varphi$ by the Price Monotonicity axiom. Thus, $\neg\mathsf{H}^{p}_{C}\varphi\notin X(w)$ due to the consistency of the set $X(w)$, which contradicts the assumption of the lemma.

Recall that, by the assumption of the case, either $p=\infty$ or there is no formula $\mathsf{H}^{\infty}_{D}\psi\in X(w)$ such that D is a nonempty subset of C and $\forall a\in D(s(a)=(\varphi,D,\infty))$. We consider these two sub-cases separately:

Case IIa: $p=\infty$. In this case $q_1+\dots+q_n>\infty$ by inequality (12), which is a contradiction.

Case IIb: there is no formula $\mathsf{H}^{\infty}_{D}\psi\in X(w)$ such that D is a *nonempty* subset of C and $\forall a\in D(s(a)=(\psi,D,\infty))$. Thus, it follows from statements (5), (8), and (9) that

$$q_i\neq\infty \text{ for all } i\le n. \tag{13}$$

Let w' be w. Since sets $D_1,\dots,D_n\subseteq C$ are pairwise disjoint and using Definition 7 together with equality (3), statement (9), and inequality (13), we conclude that

$$\begin{aligned}\sum_{a\in C}\|\delta(a)\|_{w'} &\ge \sum_{a\in D_1}\|\delta(a)\|_{w'}+\dots+\sum_{a\in D_n}\|\delta(a)\|_{w'} \\ &= \sum_{a\in D_1}\frac{q_1}{|D_1|}+\dots+\sum_{a\in D_n}\frac{q_n}{|D_n|} \\ &= q_1+\dots+q_n>p,\end{aligned} \tag{14}$$

which proves item 1 of the lemma.

We now can assume that set Y is consistent. Then, by Lemma 2, this set has a maximal consistent extension $\hat{Y}$. Let u be the sequence $w::\varnothing::(\hat{Y},0)$. Note that $u\in W$ by Definition 4 and $\neg\varphi\in Y\subseteq\hat{Y}=X(w')$ by the choice of sets Y and $\hat{Y}$ and the choice of the sequence u. To prove item 2 of the lemma it is suffices to show that $(w,\delta,u)\in M$.

Indeed, consider any formula $\mathsf{H}^q_D\psi \in X(w)$ such that $\delta(a) = (\psi, D, q)$ for each $a \in D$. Then, $\delta(a) \neq \varepsilon$ for each $a \in D$. Thus, $D \subseteq C$ due to equality (3). Hence, $\psi \in Y$ by the choice of set Y (note here that definition of set Y handles cases of nonempty and empty set D separately). Then, $\psi \in Y \subseteq \hat{Y} = X(u)$. Therefore, $(w, \delta, u) \in M$ by Definition 8. $\boxtimes$

Next, we show that canonical game $G(X_0)$ satisfies nontermination condition of Definition 2, item 6.

Lemma 16. *For any state $w \in W$ and any complete action profile $\delta \in \Delta^{\mathcal{A}}$ there is at least one state $u \in W$ such that $(w, \delta, u) \in M$.*

Proof. Recall that $\vdash \neg\mathsf{H}^\infty_{\mathcal{A}}\bot$ by the Unachievability of Falsehood axiom. Thus, $\neg\mathsf{H}^\infty_{\mathcal{A}}\bot \in X(w)$ because set $X(w)$ is consistent. Hence, by Lemma 15, there is a complete action profile $\delta' \in \Delta^{\mathcal{A}}$ such that $\delta =_{\mathcal{A}} \delta'$ and at least one of the following is true

1. there is a state $w' \in W$ such that $\sum_{a\in C} \|\delta'(a)\|_{w'} > \infty$ and $w \sim_C w'$, or
2. there is $u \in W$ such that $(w, \delta', u) \in M$ and $\neg\bot \in X(u)$.

The first out of these two statements is false because the sum cannot be greater than ∞. Thus, the second statement is true. Then, there is $u \in W$ such that $(w, \delta', u) \in M$. Therefore, $(w, \delta, u) \in M$ because $\delta =_{\mathcal{A}} \delta'$. $\boxtimes$

We are now ready to state and to prove the "truth lemma".

Lemma 17. *$w \Vdash \varphi$ iff $\varphi \in X(w)$.*

Proof. We prove the lemma by induction on the structural complexity of formula φ. If φ is a propositional variable, then the statement of the lemma follows from Definition 3 and Definition 9. The case when formula φ is a negation or an implication follows from Definition 3 and the maximality and the consistency of the set $X(w)$ in the standard way.

Suppose that formula φ has the form $\mathsf{K}_C\psi$.

$(\Rightarrow)$: Assume that $\mathsf{K}_C\psi \notin X(w)$. Thus, $\neg\mathsf{K}_C\psi \in X(w)$ because set $X(w)$ is maximal. Thus, by Lemma 14, there is a state $w' \in W$ such that $w \sim_C w'$ and $\neg\varphi \in X(w')$. Then, $\varphi \notin X(w')$ because set $X(w')$ is consistent. Hence, $w' \nVdash \varphi$ by the induction hypothesis. Therefore, $w \nVdash \mathsf{K}_C\varphi$ by Definition 3.

$(\Leftarrow)$: Assume that $\mathsf{K}_C\psi \in X(w)$. Consider any state $w' \in W$ such that $w \sim_C w'$. By Definition 3, it suffices to show that $w \Vdash \psi$. Indeed, by Lemma 13, assumptions $\mathsf{K}_C\psi \in X(w)$ and $w \sim_C w'$ imply that $\mathsf{K}_C\psi \in X(w')$. Thus, $X(w') \vdash \psi$ by the Truth axiom. Hence, $\psi \in X(w')$ by the maximality of set $X(w')$. Therefore, $w \Vdash \psi$ by the induction hypothesis.

Suppose that formula φ has the form $\mathsf{H}^p_C\psi$.

$(\Rightarrow)$: Assume that $\mathsf{H}^p_C\psi \notin X(w)$. Thus, $\neg\mathsf{H}^p_C\psi \in X(w)$ because set $X(w)$ is maximal. Consider any action profile $s \in \Delta^C$ of coalition C. By Lemma 15, there is a complete action profile $\delta \in \Delta^{\mathcal{A}}$ such that $s =_C \delta$ and at least one of the following is true

(a) there is a state $w' \in W$ such that $\sum_{a\in C} \|\delta(a)\|_{w'} > p$ and $w \sim_C w'$, or

(b) there is $u \in W$ such that $(w,\delta,u) \in M$ and $\neg\psi \in X(u)$.

In the first case, $w \nVdash \mathsf{H}^p_C\psi$ by part (5a) of Definition 3. In the second case, $\psi \notin X(u)$ because set $X(u)$ is consistent. Thus, $u \nVdash \psi$ by the induction hypothesis. Let $w' = w$. Then, $w \sim_C w'$, $s =_C \delta$, $(w',\delta,u) \in M$, and $u \nVdash \psi$. Therefore, $w \nVdash \mathsf{H}^p_C\psi$ by part (5b) of Definition 3.
$(\Leftarrow)$: Assume that $\mathsf{H}^p_C\psi \in X(w)$. Let strategy profile $s \in \Delta^C$ of coalition C be such that $s(a) = (\psi, C, p)$ for each agent $a \in C$. By Definition 3, it suffices to show that

(a) for each $w' \in W$, if $w \sim_C w'$, then $\sum_{a\in C} \|s(a)\|_{w'} \le p$,
(b) for any two states $w', u \in W$ and any complete action profile $\delta \in \Delta^{\mathcal{A}}$, if $w \sim_C w'$, $s =_C \delta$, and $(w',\delta,u) \in M$, then $u \Vdash \psi$.

We show these two statements separately. To show statement (a), consider any $w' \in W$. If $p = \infty$, then $\sum_{a\in C} \|s(a)\|_{w'} \le \infty = p$. If $C = \varnothing$, then, $\sum_{a\in C} \|s(a)\|_{w'} = 0 \le p$ because $p \ge 0$ by Definition 1. Suppose now that $p < \infty$ and $|C| > 0$. Thus, by Definition 7,

$$\sum_{a\in C} \|s(a)\|_{w'} = \sum_{a\in C} \frac{p}{|C|} = |C| \times \frac{p}{|C|} = p.$$

To show statement (b), consider any two states $w', u \in W$ and any complete action profile $\delta \in \Delta^{\mathcal{A}}$ such that $w \sim_C w'$, $s =_C \delta$, and $(w',\delta,u) \in M$. It suffices to show that $u \Vdash \psi$. Indeed, the assumption of the case $\mathsf{H}^p_C\psi \in X(w)$ implies that $X(w) \vdash \mathsf{K}_C\mathsf{H}^p_C\psi$ by the Strategic Introspection axiom and the Modus Ponens inference rule. Thus, $\mathsf{K}_C\mathsf{H}^p_C\psi \in X(w)$ by the maximality of the set $X(w)$. Hence, $\mathsf{K}_C\mathsf{H}^p_C\psi \in X(w')$ by Lemma 13 and the assumption $w \sim_C w'$. Then, $X(w') \vdash \mathsf{H}^p_C\psi$ by the Truth axiom and the Modus Ponens inference rule. Thus, $\mathsf{H}^p_C\psi \in X(w')$ because set $X(w')$ is maximal. Note also that $\delta(a) = s(a) = (\psi, C, p)$ by the assumption $s =_C \delta$ and the choice of action profile s. Thus, assumption $(w',\delta,u) \in M$ implies that $\psi \in X(u)$ by Definition 8. Therefore, $u \Vdash \psi$ by Lemma 17. ⊠

Finally, we state and prove the strong completeness theorem for our logical system.

Theorem 2. *If $Y \nvdash \varphi$, then there is a state w of a game such that $w \Vdash \chi$ for each $\chi \in Y$ and $w \nVdash \varphi$.*

Proof. Assumption $Y \nvdash \varphi$ implies that set $Y \cup \{\neg\varphi\}$ is consistent. Let set of formulae X_0 be any maximal consistent extension of set $Y \cup \{\neg\varphi\}$. Such extension exists by Lemma 2. Consider canonical game $G(X_0)$. By Definition 4, single-element sequence $(X_0, 0)$ is a state of game $G(X_0)$. We denote this state by w_0. Note that $\chi \in X_0 = X(w_0)$ for each $\chi \in Y$ by the choice of set X_0 and the choice of sequence w_0. Similarly, $\neg\varphi \in X(w_0)$. Then, $\varphi \notin X(w_0)$ because set $X(w_0)$ is consistent. Therefore, $w \Vdash \chi$ for each $\chi \in Y$ and $w \nVdash \varphi$ by the induction hypothesis. ⊠

6. Conclusion

In this article we introduced modality $\mathsf{H}_C^p\varphi$ that stands for "not only coalition C has a strategy to achieve φ, but the coalition also knows what this strategy is and the coalition knows that it will cost at most p to execute the strategy". The main technical result is a sound and complete logical system that describes the interplay between this modality and the distributed knowledge modality.

In this article we have chosen (a) to restrict the language to modalities with non-negative superscripts and (b) to restrict the models to those that have non-negative prices of actions. Neither of these restrictions is important for the proof of the soundness. In the presence of restriction (a), our canonical model construction produces a model with non-negative prices which satisfies the constraint (b). If restriction (a) is removed, then our proof of completeness will no longer work. Indeed, without this restriction, Definition 7 will assign negative prices to some of the actions. As a result, inequality (14) might no longer hold. More importantly, negative prices will deviate from the intended meaning of the *price of success* that we have discussed in Section 2. For example, suppose an agent has three actions priced at -3, 0, and 5. The negative price represents a profit from the action. Suppose also that only the action priced at 5 guarantees the success. In this situation the price of the success must account for the missed opportunity cost. Hence, the price of success should be 8, not 5. Note that this is not consistent with Definition 3. Thus, a different setting is needed to talk about negative prices, perhaps a setting that separates budget and profit as we did in [36].

A possible extension of the current work is to study modality H_C^p for strategies to achieve in multiple number of steps and for strategies to maintain. In these cases, prices of actions on different steps could be added using a *discount factor*. Know-how modality without cost parameter for these two settings has been studied in [23, 19, 20, 21] and [18] respectively. Note that the former works only deal with single-agent strategies because the Cooperation axiom does not hold for coalition strategies to achieve in multiple number of steps.

It is often possible to show decidability of a modal logic by proving completeness with respect to a class of finite models. Such completeness is usually established using filtration technique where maximal consistent sets of formulae are replaced with maximal consistent sets of subformulae of a given formula. In our case, this technique does not produce finite models because it does not impose any limit on the length of the finite sequences in Definition 4. To put a limit on the length of these sequences one might require that set X_i in this definition only contains subformulae of set X_{i-1}. However, with this modification Lemma 13 no longer holds. The decidability of the logical system proposed in this article remains an open question.

References

[1] H. H. Chouinard, Auctions with and without the right of first refusal and national park service concession contracts, American journal of agricultural

economics 87 (4) (2005) 1083–1088.

[2] Y.-M. Chiang, J. Sa-Aadu, Right of first refusal (rofr) effects in auctions with reserve price: Empirical evidence from taiwanese government land auctionsAvailable at SSRN: https://ssrn.com/abstract=2409599 or http://dx.doi.org/10.2139/ssrn.2409599.

[3] D. I. Walker, Rethinking rights of first refusal, Stan. JL Bus. & Fin. 5.

[4] M. Pauly, Logic for social software, Ph.D. thesis, Institute for Logic, Language, and Computation (2001).

[5] M. Pauly, A modal logic for coalitional power in games, Journal of Logic and Computation 12 (1) (2002) 149–166. doi:10.1093/logcom/12.1.149.

[6] V. Goranko, Coalition games and alternating temporal logics, in: Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge, Morgan Kaufmann Publishers Inc., 2001, pp. 259–272.

[7] W. van der Hoek, M. Wooldridge, On the logic of cooperation and propositional control, Artificial Intelligence 164 (1) (2005) 81 – 119.

[8] S. Borgo, Coalitions in action logic, in: 20th International Joint Conference on Artificial Intelligence, 2007, pp. 1822–1827.

[9] T. Ågotnes, P. Balbiani, H. van Ditmarsch, P. Seban, Group announcement logic, Journal of Applied Logic 8 (1) (2010) 62 – 81. doi:10.1016/j.jal.2008.12.002.

[10] T. Ågotnes, W. van der Hoek, M. Wooldridge, Reasoning about coalitional games, Artificial Intelligence 173 (1) (2009) 45 – 79. doi:10.1016/j.artint.2008.08.004.

[11] V. Goranko, S. Enqvist, Socially friendly and group protecting coalition logics, in: Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems, International Foundation for Autonomous Agents and Multiagent Systems, 2018, pp. 372–380.

[12] T. Ågotnes, N. Alechina, Epistemic coalition logic: completeness and complexity, in: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2 (AAMAS), 2012, pp. 1099–1106.

[13] W. Jamroga, T. Ågotnes, Constructive knowledge: what agents can achieve under imperfect information, Journal of Applied Non-Classical Logics 17 (4) (2007) 423–475. doi:10.3166/jancl.17.423-475.

[14] W. Jamroga, W. van der Hoek, Agents that know how to play, Fundamenta Informaticae 63 (2-3) (2004) 185–219.

[15] J. van Benthem, Games in dynamic-epistemic logic, Bulletin of Economic Research 53 (4) (2001) 219–248. doi:10.1111/1467-8586.00133.

[16] J. Broersen, A logical analysis of the interaction between ‘obligation-to-do’ and ‘knowingly doing’, in: International Conference on Deontic Logic in Computer Science, Springer, 2008, pp. 140–154.

[17] J. Broersen, A. Herzig, N. Troquard, What groups do, can do, and know they can do: an analysis in normal modal logics, Journal of Applied Non-Classical Logics 19 (3) (2009) 261–289. doi:10.3166/jancl.19.261-289.

[18] P. Naumov, J. Tao, Coalition power in epistemic transition systems, in: Proceedings of the 2017 International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2017, pp. 723–731.

[19] Y. Wang, A logic of knowing how, in: Logic, Rationality, and Interaction, Springer, 2015, pp. 392–405.

[20] Y. Wang, A logic of goal-directed knowing how, Synthese 195 (10) (2018) 4419–4439.

[21] Y. Li, Y. Wang, Achieving while maintaining: A logic of knowing how with intermediate constraints, in: Logic and Its Applications: Proceedings of 7th Indian Conference, ICLA 2017, Kanpur, India, January 5-7, 2017, Springer, 2017, pp. 154–167.

[22] T. Ågotnes, N. Alechina, Coalition logic with individual, distributed and common knowledge, Journal of Logic and Computation 29 (2019) 1041–1069. doi:10.1093/logcom/exv085.

[23] R. Fervari, A. Herzig, Y. Li, Y. Wang, Strategically knowing how, in: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17, 2017, pp. 1031–1038.

[24] P. Naumov, J. Tao, Together we know how to achieve: An epistemic logic of know-how, in: 16th conference on Theoretical Aspects of Rationality and Knowledge (TARK), July 24-26, 2017, EPTCS 251, 2017, pp. 441–453.

[25] P. Naumov, J. Tao, Together we know how to achieve: An epistemic logic of know-how, Artificial Intelligence 262 (2018) 279 – 300. doi:https://doi.org/10.1016/j.artint.2018.06.007.

[26] P. Naumov, J. Tao, Strategic coalitions with perfect recall, in: Proceedings of Thirty-Second AAAI Conference on Artificial Intelligence, 2018.

[27] P. Naumov, J. Tao, Second-order know-how strategies, in: Proceedings of the 2018 International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2018, pp. 390–398.

[28] N. Alechina, B. Logan, H. N. Nguyen, A. Rakib, Logic for coalitions with bounded resources, Journal of Logic and Computation 21 (6) (2011) 907–937.

[29] N. Alechina, B. Logan, H. N. Nguyen, F. Raimondi, Model-checking for resource-bounded ATL with production and consumption of resources, Journal of Computer and System Sciences 88 (2017) 126–144.

[30] N. Bulling, B. Farwer, On the (un-) decidability of model checking resource-bounded agents., in: ECAI, Vol. 10, 2010, pp. 567–572.

[31] N. Bulling, B. Farwer, Expressing properties of resource-bounded systems: The logics RTL* and RTL, in: Computational Logic in Multi-Agent Systems, Springer, 2010, pp. 22–45.

[32] N. Bulling, J. Dix, W. Jamroga, Model checking logics of strategic ability: Complexity, Specification and Verification of Multi-Agent Systems (2010) 125–159.

[33] D. Della Monica, M. Napoli, M. Parente, On a logic for coalitional games with priced-resource agents, Electronic Notes in Theoretical Computer Science 278 (2011) 215–228.

[34] N. Alechina, B. Logan, H. N. Nguyen, F. Raimondi, L. Mostarda, Symbolic model-checking for resource-bounded ATL, in: Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, International Foundation for Autonomous Agents and Multiagent Systems, 2015, pp. 1809–1810.

[35] N. Alechina, B. Logan, Resource logics with a diminishing resource, in: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, International Foundation for Autonomous Agents and Multiagent Systems, 2018, pp. 1847–1849.

[36] R. Cao, P. Naumov, Budget-constrained dynamics in multiagent systems, in: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017, 2017, pp. 915–921. doi:10.24963/ijcai.2017/127.

[37] P. Naumov, J. Tao, Knowing-how under uncertainty, Artificial Intelligence 276 (2019) 41 – 56.

[38] R. Fagin, J. Y. Halpern, Y. Moses, M. Y. Vardi, Reasoning about knowledge, MIT Press, Cambridge, MA, 1995.

[39] E. Mendelson, Introduction to mathematical logic, CRC press, 2009.

[40] P. Naumov, K. Ros, Strategic coalitions in systems with catastrophic failures (extended abstract), in: Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning, 2018, pp. 659–660.