

Knowledge and Common Knowledge of Strategies

Borja Sierra Miranda^[0009-0002-0079-8473] and Thomas
Studer^[0000-0002-0949-3302] ✉

University of Bern, Bern, Switzerland
{borja.sierra,thomas.studer}@unibe.ch

Abstract. Most existing work on strategic reasoning simply adopts either an informed or an uninformed semantics. We propose a model in which knowledge of strategies can be specified at a fine-grained level. It is possible to distinguish first-order, higher-order, and common knowledge of strategies. We illustrate the effect of higher-order knowledge of strategies by studying the game Hanabi. Further, we show that common knowledge of strategies is necessary to solve the consensus problem. Finally, we study the decidability of the model checking problem.

Keywords: strategic reasoning · epistemic logic · consensus

1 Introduction

Strategic reasoning is a fundamental aspect of decision-making in multi-agent systems, game theory, and artificial intelligence. It involves the ability of agents to anticipate the actions of others, formulate plans, and achieve desired outcomes within competitive or cooperative settings. To formally analyze strategic interactions, researchers have developed logical frameworks such as Coalition Logic [21], Alternating-time Temporal Logic [1], and Strategy Logic [9, 20]. Many of these systems have been extended to allow for epistemic reasoning [5, 18, 13].

Most works adopt one of two semantics regarding knowledge of strategies: in the uninformed semantics, agents do not know each other’s strategies [18, 13], while in the informed one, agents know everyone’s strategies [23]. The uninformed semantics models situations where agents have neither a priori knowledge of other agents’ strategies nor any means to communicate information about their strategies. In the informed semantics the agents’ strategies are common knowledge, which makes it possible for the agents to infer additional information from observed actions.

A notable exception to the above-mentioned dichotomy is [5], which introduces a semantics for strategy logic where one can specify, for each agent a , the set of agents A_a whose strategy a is informed of. It captures the case of informed semantics, the case of uninformed semantics, and intermediate cases where, e.g., some agents may privately communicate their strategies to other agents.

The logical presentation of strategic games introduced in [2] also includes knowledge of strategies via formulas of the form $K_a s_j^l$ stating *a knows that j has chosen her strategy l*. Artemov uses this approach to study whether a game has a

definitive solution under given epistemic/rationality conditions. However, since the paper does not introduce a concrete language or deductive system, questions of expressivity, decidability or complexity are not discussed.

Contribution. We introduce a novel approach where knowledge of strategies can be specified on a fine-grained level. In particular, our formulation makes it possible to distinguish first-order and various forms of higher-order knowledge of strategies.¹ Moreover, we can express that *everybody knows* someone’s strategy but also that a strategy is *common knowledge* among a group of agents [11].

Our semantics is based on concurrent game structures. In order to specify (higher-order) knowledge of an agent about other agents’ strategies, we introduce the notion of an *information perspective* that can be seen as a generalization of the sets A_a from [5] mentioned before. An information perspective is a set I containing sequences of agents. Such sequences are interpreted as follows: if, e.g., $abc \in I$, then agent a knows that agent b knows agent c ’s strategy. The logical language we evaluate over these structures is the usual language of epistemic logic extended with a temporal *next*-operator. This is sufficient for our examples and to investigate the power and the effect of information perspectives.

We present two small examples showing the difference between having no knowledge and first-order knowledge of an agent’s strategy. These examples will be useful to investigate basic epistemic properties such as positive and negative introspection and reflection. Moreover, these initial examples also show that the concurrent game structure is general background knowledge for all agents. The same phenomenon occurs in traditional epistemic logic where the Kripke structure is common knowledge among all agents [3].

Hanabi is a cooperative game with imperfect information. We investigate the role of higher-order knowledge of strategies in Hanabi. We show that if player b knows player a ’s strategy, then b can obtain new information from observing a ’s actions. Further, if a knows that b knows a ’s strategy, then a knows that b will obtain new information from observing a ’s action. In the situation that we study, this knowledge is a prerequisite for a to safely perform the given action. This illustrates how higher-order knowledge is necessary for successful gameplay.

Since an information perspective may be infinite, it is possible to express common knowledge of strategies. We show that if all strategies are common knowledge, then the individual knowledge operators satisfy the negative introspection property. Next, we turn to the binary consensus problem [7] and show that common knowledge of strategies is required to solve this task successfully.

Finally, we investigate the problem of model checking where we restrict ourselves to formulas that do not contain common knowledge operators. The key technical lemma shows that, in this case, only information perspectives of bounded size have to be considered. This yields the decidability of the corresponding model checking problem.

¹ First-order knowledge means factual knowledge, i.e. knowledge about the objective state of the world, whereas higher-order knowledge denotes knowledge about knowledge, i.e. knowledge about the subjective epistemic state of some agent.

2 The model

We start with a countable set of atomic propositions Prop and a finite set of agents Ag . $\mathcal{P}(X)$ denotes the power set of X and X^Y denotes the set of functions from Y to X .

A *concurrent game structure* is a tuple $(\text{Ac}, \mathbf{V}, \mathbf{E}, \ell, \sim_a)$ where

1. Ac is a finite set of actions,
2. \mathbf{V} is a finite set of positions,
3. $\mathbf{E} : \mathbf{V} \times \text{Ac}^{\text{Ag}} \rightarrow \mathbf{V}$ is a transition function,
4. $\ell : \mathbf{V} \rightarrow \mathcal{P}(\text{Prop})$ a valuation function,
5. $\sim_a \subseteq (\mathbf{V} \times \mathbf{V}) \cup (\text{Ac} \times \text{Ac})$ is an equivalence relation for each agent $a \in \text{Ag}$, called a 's observation relation.

A *joint action* is a function $\alpha : \text{Ag} \rightarrow \text{Ac}$ mapping each agent to an action. A *history* is a finite sequence of positions and joint actions $\rho = v_0 \alpha_1 v_1 \dots \alpha_n v_n$ such that $\mathbf{E}(v_i, \alpha_{i+1}) = v_{i+1}$. For a history $\rho = v_0 \alpha_1 v_1 \dots \alpha_n v_n$ and $i \leq n$, we set $\rho_{\leq i} := v_0 \alpha_1 v_1 \dots \alpha_i v_i$ and $\text{last}(\rho) := v_n$. We let Hist be the set of all histories.

Two positions $v, v' \in \mathbf{V}$ are *indistinguishable* for an agent a if $v \sim_a v'$, and similarly for actions. Two joint actions α and β are indistinguishable for agent a , in symbols $\alpha \sim_a \beta$ if they are pointwise indistinguishable for a , i.e. for all $b \in \text{Ag}$, we have that $\alpha(b) \sim_a \beta(b)$.

We assume synchronous perfect recall. That is, agents remember their observations of all past positions and actions. Hence, we extend the observation relations to histories as follows. Two histories $\rho = v_0 \alpha_1 v_1 \dots \alpha_n v_n$ and $\rho' = v'_0 \alpha'_1 v'_1 \dots \alpha'_m v'_m$ are *indistinguishable* for agent a , in symbols $\rho \sim_a \rho'$, if $m = n$, $v_i \sim_a v'_i$ for all $0 \leq i \leq n$, and $\alpha_i \sim_a \alpha'_i$ for all $1 \leq i \leq n$.

A *strategy* is a function $\sigma : \text{Hist} \rightarrow \text{Ac}$. We let Str denote the set of all strategies.

An *assignment* $\chi : \text{Ag} \rightarrow \text{Str}$ is a mapping of agents to strategies. A history $\rho = v_0 \alpha_1 v_1 \dots \alpha_n v_n$ is *consistent* with an assignment χ for a set of agents X if $\alpha_{i+1}(b) = \chi(b)(\rho_{\leq i})$ for all $i < n$ and agents $b \in X$.

Given a concurrent game structure \mathcal{G} , an assignment χ , and a history $\rho = v_0 \alpha_1 v_1 \dots \alpha_n v_n$, we define the *one-step continuation of ρ*

$$X_{\mathcal{G}}^{\chi} \rho := v_0 \alpha_1 v_1 \dots \alpha_n v_n \alpha_{n+1} v_{n+1}$$

where the joint action α_{n+1} is given by $\alpha_{n+1}(b) := \chi(b)(\rho)$ for all agents b and $v_{n+1} := \mathbf{E}(v_n, \alpha_{n+1})$ is the next position.

Let X be a finite set. We use X^* to denote the set of all finite words with symbols from X without adjacent repetitions, that is if $x_n \dots x_1 \in X^*$, then $x_i \neq x_{i+1}$ for all $1 \leq i < n$.

For a word $w = x_n \dots x_1 \in X^*$ and an element $y \in X$, we write yw for the word $yx_n \dots x_1$. The *length of a word* w , $\text{len}(w)$, is the number of symbols of w . That is, for a word $w = x_n \dots x_1 \in X^*$, we have $\text{len}(w) = n$. Further, we set

$$X^{\geq n} := \{w \in X^* \mid \text{len}(w) \geq n\},$$

that is $X^{\geq n}$ is the set of those words of X^* with a length of at least n .

An information perspective stores which agents are informed about the strategies of which other agents. Formally, an *information perspective* I is a subset of $\text{Ag}^{\geq 2}$. Further, for an agent a , we set $I_a := \{b \in \text{Ag} \mid ab \in I\} \cup \{a\}$ and

$$I[a] := \{w \in \text{Ag}^{\geq 2} \mid aw \in I\} \cup \{w \in I \mid w = aw' \text{ for some } w' \in \text{Ag}^*\}.$$

Let $a, b, c \in \text{Ag}$ be agents. We interpret $ab \in I$ as *agent a knows (is informed about) agent b 's strategy*. Hence, I_a is the set of all agents whose strategies agent a is informed about. We assume that each agent knows its own strategy. Therefore, we let $a \in I_a$ by definition.

Further, $abc \in I$ means that agent a knows that agent b knows agent c 's strategy, and so on. The operation $I[a]$ changes the information perspective to agent a 's point of view. That is, $I[a]$ contains the information that a knows. We assume positive introspection, that is, agents know about their own knowledge. This is reflected by the condition that for all w with $w = aw'$ for some w' , we have $w \in I$ implies $w \in I[a]$.

An information perspective I is *truthful* if it is closed under suffixes, i.e. for any $a \in \text{Ag}$ and any $w \in \text{Ag}^{\geq 2}$ we have that

$$aw \in I \text{ implies } w \in I. \quad (1)$$

This means, for instance, if agent a knows that agent b knows agent c 's strategy, then agent b must know agent c 's strategy. Formally, (1) guarantees that $abc \in I$ implies $bc \in I$.

On a technical level, an important consequence of (1) is that for any truthful information perspective I and any agent a , we have

$$I[a] \subseteq I. \quad (2)$$

This will be essential when we establish a version of the knowledge axiom T.

An agent a cannot distinguish two assignments χ and χ' given an information perspective I if χ and χ' agree on the strategies of all agents that agent a is informed about according to I . As mentioned before, this includes the agent's own strategy. We say that two assignments χ and χ' are *indistinguishable* for agent a under the information perspective I , in symbols $\chi \sim_a^I \chi'$, if we have $\chi(b) = \chi'(b)$ for all $b \in I_a$.

A *state* in a concurrent game structure is a triple (χ, I, ρ) where

1. χ is an assignment,
2. I is an information perspective, and
3. ρ is a history.

Let a be an agent. We say that (χ, I, ρ) is *a -consistent* if ρ is consistent with χ under I_a . Further, we call a state (χ', I', ρ') *accessible from* (χ, I, ρ) *for* a , in symbols $(\chi, I, \rho) \trianglelefteq_a (\chi', I', \rho')$, if

1. $\chi \sim_a^I \chi'$,
2. $I[a] \subseteq I'$,
3. $\rho \sim_a \rho'$,
4. (χ, I, ρ) is a -consistent and
5. (χ', I', ρ') is a -consistent.

In order to deal with common knowledge of a group of agents, we need the following definitions where $G \subseteq \text{Ag}$:

$$\begin{aligned} Z \trianglelefteq_G Z' &\text{ iff } Z \trianglelefteq_a Z' \text{ for some } a \in G, \\ Z \trianglelefteq_G^* Z' &\text{ iff } Z \trianglelefteq_G Z_1 \trianglelefteq_G \cdots \trianglelefteq_G Z_n \trianglelefteq_G Z' \text{ for some states } Z_1, \dots, Z_n. \end{aligned}$$

Hence $Z \trianglelefteq_G^* Z'$ states that Z' is reachable from Z using agents from G . Its definition includes the case that $Z \trianglelefteq_G^* Z'$ if $Z \trianglelefteq_a Z'$ for some $a \in G$.

In the following, we study some basic properties of the accessibility relation. First, we observe that replacing the history in a state Z with a history that cannot be distinguished by agent a gives a state Z' that is accessible from Z for a (given that Z and Z' are a -consistent and the information perspective of Z is truthful). The following is an immediate consequence of $\chi \sim_a^I \chi$ and (2).

Lemma 1. *Let $\mathcal{G} = (\text{Ac}, \text{V}, \text{E}, \ell, \sim_a)$ be a concurrent game structure and let a be an agent. We let (χ, I, ρ) be an a -consistent state with a truthful information perspective I . For each history ρ' such that $\rho \sim_a \rho'$ and ρ' is consistent with χ under I_a , we have $(\chi, \text{I}, \rho) \trianglelefteq_a (\chi, \text{I}, \rho')$.*

The next lemma states that if $Z \trianglelefteq_a Z'$, then the history of the state Z' is consistent from the point of view of agent a in Z . This is related to the notation $\rho \sim_a^{\chi, \text{I}} \rho'$ used in [5]. A proof of this lemma is given in Appendix A.

Lemma 2. *Let $\mathcal{G} = (\text{Ac}, \text{V}, \text{E}, \ell, \sim_a)$ be a concurrent game structure and let $Z = (\chi, \text{I}, \rho)$ and $Z' = (\chi', \text{I}', \rho')$ be states with $Z \trianglelefteq_a Z'$ for some agent a . Then ρ' is consistent with χ under I_a .*

We now introduce our languages \mathcal{L} and \mathcal{L}^{C} , which includes formulas:

1. $\text{K}_a \phi$ meaning *agent a knows that ϕ* ;
2. $\text{C}_G \phi$ meaning *it is common knowledge among G that ϕ* ;
3. $\text{X}\phi$ meaning *at the next point in time ϕ holds*.

Formulas of the language \mathcal{L} are given by the following grammar

$$\phi ::= p \mid \perp \mid \phi \rightarrow \phi \mid \text{K}_a \phi \mid \text{X}\phi,$$

and formulas of the language \mathcal{L}^{C} are given by

$$\phi ::= p \mid \perp \mid \phi \rightarrow \phi \mid \text{K}_a \phi \mid \text{C}_G \phi \mid \text{X}\phi,$$

where $p \in \text{Prop}$, $a \in \text{Ag}$, and $G \subseteq \text{Ag}$. The remaining Boolean connectives are defined as usual.

The next-operator is the only temporal operator that we include in our language, as it is the only one needed for our examples. It is straightforward to include further temporal operators, see, e.g., [5] for the case of the until-operator.

Let $\mathcal{G} = (\text{Ac}, \text{V}, \text{E}, \ell, \sim_a)$ be a concurrent game structure and $Z = (\chi, \text{I}, \rho)$ be a state. *Truth* of a formula ϕ in \mathcal{G} at state Z , in symbols $\mathcal{G}, Z \Vdash \phi$, is inductively defined as follows:

1. $\mathcal{G}, Z \Vdash p$ iff $p \in \ell(\text{last}(\rho))$ where $p \in \text{Prop}$,
2. $\mathcal{G}, Z \not\vdash \perp$,
3. $\mathcal{G}, Z \Vdash \phi \rightarrow \psi$ iff $\mathcal{G}, Z \not\vdash \phi$ or $\mathcal{G}, Z \Vdash \psi$,
4. $\mathcal{G}, Z \Vdash K_a \phi$ iff $\mathcal{G}, Z' \Vdash \phi$ for all states Z' with $Z \trianglelefteq_a Z'$,
5. $\mathcal{G}, Z \Vdash C_G \phi$ iff $\mathcal{G}, Z' \Vdash \phi$ for all states Z' with $Z \trianglelefteq_G^* Z'$,
6. $\mathcal{G}, Z \Vdash X\phi$ iff $\mathcal{G}, Z' \Vdash \phi$ where $Z' = (\chi, I, X_G^x \rho)$.

We say that a formula ϕ is *valid* if $\mathcal{G}, Z \Vdash \phi$ holds for all concurrent game structures \mathcal{G} and states Z .

3 First examples and basic properties

We give two basic examples to illustrate our truth definition. These examples will also be referred to in the proof of Lemma 5.

We consider a set Ag of two agents a and b and the concurrent game structure shown in Figure 1. In position v_1 , only agent a will perform an action. This can be either action α leading to position v_2 where p is true or action β leading to position v_3 where q is true. We assume that every agent can distinguish the different positions and also the two actions. Since the other agents' actions do not matter in these examples, we will not mention them and consider α and β also as joint actions. We also omit the loops on final positions. Let the strategy σ_α be




Fig. 1. The concurrent game structure \mathcal{G}_1

such that $\sigma_\alpha(v_1) = \alpha$ and let χ_α be an assignment with $\chi_\alpha(a) = \sigma_\alpha$. Similarly, let the strategy σ_β be such that $\sigma_\beta(v_1) = \beta$ and let χ_β be an assignment with $\chi_\beta(a) = \sigma_\beta$.

For the first example, we assume that agent b is informed about agent a 's strategy.

Example 1. We set $I := \{ba\}$ and consider the state $Z = (\chi_\alpha, I, v_1)$. In this setting, we find that $\mathcal{G}_1, Z \Vdash K_b X p$ holds.

Indeed, let $Z' = (\chi', I', v_1)$ be an arbitrary state with $Z \trianglelefteq_b Z'$. This implies $\chi_\alpha \sim_b^I \chi'$, which by $a \in I_b$ yields $\chi_\alpha(a) = \chi'(a)$. Therefore, we get $X_{\mathcal{G}_1}^{\chi'} v_1 = v_1 \alpha v_2$ and thus $\mathcal{G}_1, Z' \Vdash X p$. Since Z' was arbitrary with $Z \trianglelefteq_b Z'$, we conclude $\mathcal{G}_1, Z \Vdash K_b X p$.

For the second example, we consider the case when b is not informed about a 's strategy.

Example 2. We set $I = \emptyset$, that is, in particular, $ba \notin I$, and consider the state $Z = (\chi_\alpha, I, v_1)$. We show that

$$\mathcal{G}_1, Z \not\models K_b Xp. \quad (3)$$

Indeed, we find $\chi_\alpha \sim_b^I \chi_\beta$. Hence we obtain $Z \trianglelefteq_b Z'$ for $Z' = (\chi_\beta, I, v_1)$. We get $\mathcal{G}_1, Z' \models X\neg p$ and thus conclude (3).

Similarly, we can also establish $\mathcal{G}_1, Z \not\models K_b Xq$. However, let us mention that, as usual, the concurrent game structure is general background knowledge. That means agent b knows that agent a will perform either action α or action β and thus reach a position where p holds or one where q holds. Formally, we have $\mathcal{G}_1, Z \models K_b(Xp \vee Xq)$.

Let us now examine basic properties of the individual knowledge operators. The proofs of the following three lemmas are omitted here for reasons of space; they can be found in the accompanying technical report [24]. We begin with positive introspection.

Lemma 3. *The formula $K_a\phi \rightarrow K_aK_a\phi$ is valid.*

We defined truthful information perspectives to satisfy (1). This is a necessary condition for $Z \trianglelefteq_a Z$ to hold. However, it is not sufficient as the state Z may not be a -consistent.

We have the following lemma, which states that the Truth axiom for K_a holds at a -consistent states with truthful information perspectives. This is similar to the situation in simplicial semantics [8, 12], which provides a model of epistemic logic for distributed systems, where the Truth axiom only holds for correct processes.

Lemma 4. *Let \mathcal{G} be a concurrent game structure and $Z = (\chi, I, \rho)$ be a state with a truthful information perspective. We have $\mathcal{G}, Z \models \neg K_a \perp \rightarrow (K_a\phi \rightarrow \phi)$.*

In general, agents cannot perform negative introspection.

Lemma 5. *The schema of negative introspection is not valid, i.e. there is a countermodel to $\neg K_b\phi \rightarrow K_b\neg K_b\phi$.*

4 Hanabi

We illustrate the workings of our logic and the importance of higher-order knowledge of strategies by studying the game Hanabi. Hanabi is a cooperative game of imperfect information. Because of this particular combination, Hanabi has been proposed as a challenge for machine learning techniques in multi-agent settings [4]. Early work on Hanabi was mainly concerned with evaluating fixed

strategies [10]. Later, Hanabi has been used to study, e.g., zero-shot coordination, that is, constructing agents that can coordinate with novel partners they have not seen before [14]. Most recently, Perrotin [22] provides an epistemic logic analysis of Hanabi that is based on (common) knowledge of game states. In contrast to our examples, knowledge of strategies is not relevant for that work.

The goal of the game is to play cards so as to form ordered stacks, one for each color, beginning with a card of rank 1 and ending with a card of rank 5. The game’s twist is the imperfect information, which arises from each player being unable to see their own cards (but they can see the hands of the other players). Hence, players have to rely on receiving information from the other players.

Players take turns doing one of three actions: giving a hint, playing a card from their hand, or discarding a card. The number and types of hints that are allowed are very restricted. However, each action taken in a play is observed by all players and can, therefore, convey implicit information. This is possible, in particular, when players have pre-established conventions and tactics. In our example, we assume that the players use the following convention: if a player thinks that several of her cards are possible to play, she plays the leftmost of them.

We use Hanabi as a prime example of the role of higher-order knowledge when reasoning about strategies. In particular, we investigate a standard tactic of Hanabi, the so-called *finesse* move, see [4, 26] for more details.

We consider a simplified situation, ignoring the color of the cards and the exact form of the hints. There are three players a , b , and c . It is player a ’s turn, then player b ’s, and then player c ’s. Player a sees that player b has a card of rank 1 on her leftmost position and player c has a card of rank 2. Player a could now

1. inform player b of her card of rank 1. Then player b will play that card.
2. inform player c of her card of rank 2 and assume that player b will infer that she must have a card 1. So player b will play her card 1, and then player c will play her card 2.

Of course, the second option is the better choice, as with one hint, two correct cards will be played. This is the essence of the *finesse*-move.

We model this situation with the concurrent game structure \mathcal{G}_2 given in Figure 2. Proposition p means that player b has a card 1. The positions v_1 and v_4 are two possible positions before a ’s move. At position v_1 , player b does have a card 1, while she does not at v_4 . The positions v_2 , v_3 , v_5 , and v_6 represent the different possible results of a ’s action.

Action α means that player a gives player b the information that b has a card 1. Action β means that player a gives player c the information that she has a card 2. Action γ means that player a gives player b the information that b does not have a card 1. As in our previous examples, only the action of player a matters, and we consider α , β , and γ also as joint actions.

All actions are public, and each agent can distinguish them. In particular, we have

$$\alpha \not\sim_b \beta \text{ and } \beta \not\sim_b \gamma \text{ and } \alpha \not\sim_b \gamma. \quad (4)$$




Fig. 2. The concurrent game structure \mathcal{G}_2

Player a can see player b 's cards. Thus, a can distinguish the different states, i.e.

$$v_i \not\sim_a v_j \quad \text{for } i \neq j. \quad (5)$$

Player b does not know whether she has a card 1. Thus

$$v_1 \sim_b v_4. \quad (6)$$

If player a informs player b , then she knows whether she has a card 1. Thus $v_2 \not\sim_b v_5$. However, if player a gives information to player c , then b still does not know whether she has a card 1. Hence

$$v_3 \sim_b v_6. \quad (7)$$

Player a 's strategy is as follows.

1. If player b does not have a card 1, then a gives b the information that b does not have a card 1, i.e. she performs action γ .
2. If player b has a card 1, then a informs c , i.e. she performs action β .

We denote this strategy by σ . That is $\sigma(v_1) := \beta$ and $\sigma(v_4) := \gamma$. Further, we let the assignment χ be such that $\chi(a) := \sigma$. We consider the history $\rho := v_1\beta v_3$, i.e. player b has a card 1 and player a gives information to c .

Assume that player b knows player a 's strategy. We show that in this situation, b can infer from observing action β that she must have a card of rank 1.

Example 3. We consider an information perspective I with $ba \in I$ and let Z be the state (χ, I, ρ) . We find that in this situation, player b knows that she has a card 1, although player a informed player c . Formally, we have $\mathcal{G}_2, Z \Vdash K_b p$. Indeed, let $Z' = (\chi', I', \rho')$ be any state with $Z \trianglelefteq_b Z'$. This implies $\chi \sim_b^1 \chi'$, which by $a \in I_b$ yields $\chi'(a) = \chi(a) = \sigma$.

Since ρ' must be consistent with χ' under I'_b , we get $\rho' = v_1\beta v_3$ or $\rho' = v_4\gamma v_5$. Since $\rho \sim_b \rho'$, we find by (4) that $\rho' = v_1\beta v_3$ or $\rho' = v_4\beta v_6$. Hence, we conclude $\rho' = v_1\beta v_3$. Therefore, $\mathcal{G}_2, Z' \Vdash p$ and $\mathcal{G}_2, Z \Vdash K_b p$.

Now, we assume that player b does not know player a 's strategy. In this situation, player b does not know that she has a card of rank 1 unless player a tells her so.

Example 4. Formally, we consider an information perspective I with $ba \notin I$ and let Z be the state (χ, I, ρ) . We will show $\mathcal{G}_2, Z \not\models K_b p$.

First, we observe that Z is b -consistent. Now let σ' be a strategy such that $\sigma'(v_1) = \beta$ and $\sigma'(v_4) = \beta$. Further let χ' be an assignment with $\chi'(a) = \sigma'$. Because of $ba \notin I$ we find $a \notin I_b$. Therefore $\chi \sim_b^I \chi'$. We set $\rho' = v_4 \beta v_6$ and observe by (6) and (7) that $\rho \sim_b \rho'$. Further we have that ρ' is consistent with χ' under I_b , which means that (χ', I', ρ') is b -consistent. Therefore, we obtain $(\chi, I, \rho) \sqsubseteq_b (\chi', I', \rho')$. We have $(\chi', I', \rho') \not\models p$ and thus conclude $\mathcal{G}_2, Z \not\models K_b p$.

Let us again consider Example 3 where we established that

$$\begin{aligned} \text{if } b \text{ knows } a\text{'s strategy, then } b \text{ will infer from observing action } \beta \\ \text{that she must have a card of rank 1.} \end{aligned} \quad (8)$$

However, note that this is not enough for player a to safely choose action β . What is required is that player a knows (8). Only then can a be sure that if she informs c , player b will play her card 1, and thus informing c about her card 2 is a good choice.

Hence, this is a situation where higher-order knowledge is relevant. Namely, we have that if player a knows that player b knows a 's strategy, then player a knows that if she informs player c , then player b will know that she has a card 1. We can formalize this situation as follows.

Example 5. We let I be an information perspective with $aba \in I$ and consider the state $Z = (\chi, I, v_1)$. Let $Y = (\chi', I', \rho')$ be an arbitrary state with $Z \sqsubseteq_a Y$. Thus $\rho \sim_a \rho'$. Therefore by (5), we get $\rho' = v_1$. Moreover, we have $\chi \sim_a^I \chi'$. Because of $a \in I_a$, we get $\chi'(a) = \chi(a) = \sigma$. Hence $X_{\mathcal{G}_2}^{\chi'} \rho' = v_1 \beta v_3$. So we let X be the state $(\chi', I', v_1 \beta v_3)$. We observe $ba \in I[a] \subseteq I'$. Therefore, we obtain $\mathcal{G}_2, X \models K_b p$ as in Example 3. Hence $\mathcal{G}_2, Y \models \mathbf{X}K_b p$. Since Y was arbitrary with $Z \sqsubseteq_a Y$, we conclude $\mathcal{G}_2, Z \models K_a \mathbf{X}K_b p$.

The reasoning of the previous example is only possible if agent a knows that b is informed about a 's strategy. Otherwise, a must take into account a situation where b is not informed about a 's strategy, and thus a considers the situation of Example 4 possible. Hence, even though b will know that she has a card 1, player a will not know that b knows this.

Example 6. We let I be an information perspective with $ba \in I$ but $aba \notin I$ and consider an a -consistent state $Z = (\chi, I, v_1)$. We set $Y = (\chi, I[a], v_1)$ and find $Z \sqsubseteq_a Y$ and $ba \notin I[a]$. Further, we have $X_{\mathcal{G}_2}^{\chi} v_1 = v_1 \beta v_3$ and let X be the state $(\chi, I[a], v_1 \beta v_3)$. As in Example 4, we find $\mathcal{G}_2, X \not\models K_b p$. Hence $\mathcal{G}_2, Y \not\models \mathbf{X}K_b p$ and finally $\mathcal{G}_2, Z \not\models K_a \mathbf{X}K_b p$.

5 Common knowledge

In this section, we investigate common knowledge of strategies. Agent a 's strategy is common knowledge among a group of agents G if all agents in G are

informed about it, all agents in G know that all agents in G are informed about it, all agents in G know that all agents in G know that all agents in G are informed about it, and so on. Hence, an information perspective I models that agent a 's strategy is common knowledge among G if $wa \in I$ for each $w \in G^{\geq 1}$.

We can recover negative introspection if all strategies are common knowledge. This is captured by the notion of a fully informed information perspective. Formally, we have the following definition.

Definition 1. *We say that the information perspective I is fully informed if $I = \text{Ag}^{\geq 2}$.*

A fully informed information perspective I satisfies the following important property for any agent a

$$I[a] = I = \text{Ag}^{\geq 2}. \quad (9)$$

The following lemma states that Axiom B, which in modal logic is $\phi \rightarrow \Box \Diamond \phi$, and negative introspection hold for fully informed information perspectives. We omit the proof for reasons of space.

Lemma 6. *Let \mathcal{G} be a concurrent game structure and $Z = (\chi, I, \rho)$ be a state with a fully informed information perspective. We have $\mathcal{G}, Z \Vdash \phi \rightarrow K_a \neg K_a \neg \phi$. and $\mathcal{G}, Z \Vdash \neg K_a \phi \rightarrow K_a \neg K_a \phi$.*

We illustrate the need for common knowledge of strategies with the problem of binary consensus. In this problem, we have two agents, each of whom has an input value that can be either 0 or 1, and they have to agree on a common value, which has to be one of their input values. In this context, agreeing on a value means that the value is common knowledge for the two agents.

We consider the following scenario. There are two agents, a and b , and the group G consists of these two agents. In the following, we will just say common knowledge when we mean common knowledge among G . One agent has 0 as the input value; the other agent has an input value of 1. They exchange their values using reliable communication. So, it is common knowledge that the agents have 0 and 1 as their respective input values. We represent this by the position v_1 in a concurrent game structure \mathcal{G}_3 that has the same form as the one given in Fig. 1.

Now, agent a has to choose its output value, and we only consider the question of whether her choice will be common knowledge. Action α means agent a uses the minimum function to choose its output value. Action β means agent a uses the maximum function. So p stands for a decided for 0, and q means that a decided for 1 as the output value.

We assume agent a uses the minimum function, i.e. her strategy is to perform action α . We let χ be such that $\chi(a)(v_1) = \alpha$ and χ' be such that $\chi'(a)(v_1) = \beta$. Agent a does not announce which output value she has chosen. That is in \mathcal{G}_3 , we have $v_2 \sim_b v_3$ and $\alpha \sim_b \beta$. Let $\rho := v_1 \alpha v_2$ and $\rho' := v_1 \beta v_3$. We find that $\rho \sim_b \rho'$.

We observe that common knowledge of a 's strategy is necessary for common knowledge of p . A proof of this theorem is given in Appendix B.

Theorem 1. *Let I be an information perspective such that a 's strategy is not common knowledge. We have $\mathcal{G}_3, (\chi, I, \rho) \not\models C_G p$.*

Further, if a 's strategy is common knowledge, then p is common knowledge, see Appendix C for a proof.

Theorem 2. *Let I be such that a 's strategy is common knowledge. We have $\mathcal{G}_3, (\chi, I, \rho) \models C_G p$.*

6 Model checking

Finally, we investigate the model checking problem. We will show that it is decidable for \mathcal{L} -formulas.

We begin with the following auxiliary notion. The K-depth of an \mathcal{L} -formula ϕ , in symbols $\text{kd}(\phi)$, is inductively defined by:

$$\begin{aligned} \text{kd}(p) &:= 0 & \text{kd}(\phi \rightarrow \psi) &:= \max(\text{kd}(\phi), \text{kd}(\psi)) \\ \text{kd}(\perp) &:= 0 & \text{kd}(K_a \phi) &:= \text{kd}(\phi) + 1 \\ \text{kd}(X\phi) &:= \text{kd}(\phi). \end{aligned}$$

For $X \subseteq \text{Ag}^*$ and a natural number n , we set $X \upharpoonright_n := \{w \in X \mid \text{len}(w) \leq n\}$. The following lemma is proved in Appendix D.

Lemma 7. *Let \mathcal{G} be a concurrent game structure, $Z = (\chi, I, \rho)$ be a state, and ϕ be an \mathcal{L} -formula. For any information perspective I' with $I \upharpoonright_{\text{kd}(\phi)+2} = I' \upharpoonright_{\text{kd}(\phi)+2}$, we have $\mathcal{G}, (\chi, I, \rho) \models \phi$ iff $\mathcal{G}, (\chi, I', \rho) \models \phi$.*

The *duration* of a history is the number of joint actions occurring in it, i.e. the duration of $v_0 \alpha_1 v_1 \dots \alpha_n v_n$ is n . The duration of a state (χ, I, ρ) is the duration of ρ . Let Z be a state with duration n . By the definition of truth, we find that

1. to evaluate a formula of the form $X\phi$ at Z , we have to evaluate ϕ at a state with duration $n + 1$;
2. to evaluate a formula of the form $K_a \phi$ at Z , we have to evaluate ϕ at states with duration n .

Let ϕ be an \mathcal{L} -formula with a maximal nesting depth of X -operators of m . We find that to evaluate ϕ at Z , we only have to consider states with a duration of at most $n + m$.

This, Lemma 7, and the fact that a concurrent game structure consists only of finitely many actions and finitely many positions together yield that to evaluate an \mathcal{L} -formula at a given state, we only have to consider finitely many other states. Therefore, we obtain the following corollary.

Corollary 1. *Model checking is decidable for \mathcal{L} -formulas.*

7 Conclusion

In this work, we introduced a model that can express first-order, higher-order, and common knowledge of strategies. We studied how higher-order knowledge of strategies affects agents' strategic abilities in the game Hanabi, and we showed that common knowledge of strategies is necessary to solve the consensus task. Finally, we established the decidability of the model checking problem for the language without common knowledge.

Three important topics are left for future work:

1. studying the model checking problem in the presence of common knowledge;
2. considering a language that includes more expressive temporal operators, in particular *until*;
3. extending our approach to proper strategy logic [9, 20], i.e. to a language that treats strategies as explicit first-order objects.

With respect to the last topic, first steps towards to an extension to strategy logic are presented in [16]. The second topic, i.e. the combination of epistemic fixed points (common knowledge) and temporal fixed points (unti) [17], is also highly interesting from a proof-theoretic perspective. Cyclic and non-wellfounded systems [6, 15, 19, 25] provide an promising approach to handle this.

Acknowledgments. We would like to thank the anonymous reviewers for their helpful comments.

Thomas Studer is supported by the Swiss National Science Foundation (SNSF) under grant number 200021-227544 (Epistemic Group Attitudes).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

A Proof of Lemma 2

We have

$$\alpha'_{i+1}(b) = \chi'(b)(\rho'_{\leq i}) \quad \text{for all } i < n \text{ and agents } b \in I'_a.$$

By $I[a] \subseteq I'$ we get

$$\alpha'_{i+1}(b) = \chi'(b)(\rho'_{\leq i}) \quad \text{for all } i < n \text{ and agents } b \in I_a.$$

Because of $\chi \sim_a^I \chi'$, we conclude

$$\alpha'_{i+1}(b) = \chi(b)(\rho'_{\leq i}) \quad \text{for all } i < n \text{ and agents } b \in I_a,$$

which means that ρ' is consistent with χ under I_a .

B Proof of Theorem 1

We define sequences $w_i \in \mathbf{Ag}^{\geq 2}$ for $i \geq 1$ as follows: $w_1 := ba$, $w_2 := aba$, $w_3 := baba$, and so on. For a sequence $w = x_n \cdots x_2 x_1 \in \mathbf{Ag}^{\geq 2}$, we let ϕ_w be the formula $K_{x_n} \cdots K_{x_2} p$. Further, we write ϕ_i for ϕ_{w_i} . For instance, we have $\phi_1 = K_b p$ and $\phi_2 = K_a K_b p$.

We show

$$\text{for all information perspectives } I, \quad w_i \notin I \text{ implies } \mathcal{G}_3, (\chi, I, \rho) \not\models \phi_i$$

by induction on i .

Case $i = 1$. Since $ba \notin I$, we $a \notin I_b$. Therefore, $\chi \sim_b^I \chi'$. This yields $(\chi, I, \rho) \leq_b (\chi', I[b], \rho')$. Because of $\mathcal{G}_3, (\chi', I[b], \rho') \not\models p$, we conclude

$$\mathcal{G}_3, (\chi, I, \rho) \not\models K_b p.$$

Case $i = j + 1$. Let $x \in \mathbf{Ag}$ be such that $w_i = xw_j$. We have

$$(\chi, I, \rho) \leq_x (\chi, I[x], \rho).$$

By the definition of $I[x]$ and $w_i \notin I$, we find $w_j \notin I[x]$. Thus, by I.H., we obtain

$$\mathcal{G}_3, (\chi, I[x], \rho) \not\models \phi_j.$$

Therefore, $\mathcal{G}_3, (\chi, I, \rho) \not\models K_x \phi_j$, which is

$$\mathcal{G}_3, (\chi, I, \rho) \not\models \phi_i.$$

If a 's strategy is not common knowledge, there exists i with $w_i \notin I$ and thus $\mathcal{G}_3, (\chi, I, \rho) \not\models \phi_i$. Since $C_G p \rightarrow \phi_i$ is valid, we conclude $\mathcal{G}_3, (\chi, I, \rho) \not\models C_G p$.

C Proof of Theorem 2

By induction on i we show that for all states Z_0, \dots, Z_i such that

$$(\chi, I, \rho) =: Z_0 \triangleleft_{x_1} Z_1 \triangleleft_{x_2} \cdots \triangleleft_{x_i} Z_i,$$

we have

1. $\mathcal{G}_3, Z_i \Vdash p$ and
2. for $Z_i = (\sigma^i, I^i, \rho^i)$, agent a 's strategy is common knowledge in I^i .

Case $i = 0$. We have $\text{last}(\rho) = v_2$ and hence $\mathcal{G}_3, (\chi, I, \rho) \Vdash p$. By assumption, we also have that a 's strategy is common knowledge in I .

Case $i = j + 1$. We let $Z_j = (\chi^j, I^j, \rho^j)$. By I.H. we find $\mathcal{G}_3, Z_j \Vdash p$, which means that $\text{last}(\rho^j) = v_2$, and

$$\text{agent } a \text{'s strategy is common knowledge in } I^j. \tag{10}$$

Therefore, $a \in I_{x_i}^j$. Since ρ^j is consistent with χ^j under $I_{x_i}^j$, we find $\chi^j(a)(v_1) = \alpha$. We have $\chi^j \sim_{x_i}^j \chi^i$. Therefore, $\chi^i(a)(v_1) = \chi^j(a)(v_1) = \alpha$. Since ρ^i is consistent with χ^i under $I_{x_i}^i$, we get $\text{last}(\rho^i) = v_2$. Hence, we conclude

$$\mathcal{G}_3, Z_i \Vdash p.$$

To show the second claim, we let $w = w_h \cdots w_1 \in \text{Ag}^{\geq 1}$. If $w_h \neq x_i$, then by (10) we have $x_i w a \in I^j$ and thus

$$w a \in I[x_i] \subseteq I^i.$$

If $w_h = x_i$, then by (10) we have $w a \in I^j$ and again

$$w a \in I[x_i] \subseteq I^i.$$

Hence, a 's strategy is common knowledge in I^i .

We have shown that for any state Z_i that is reachable from (χ, I, ρ) , we have $\mathcal{G}_3, Z_i \Vdash p$. Therefore $\mathcal{G}_3, (\chi, I, \rho) \Vdash C_G p$ as desired.

D Proof of Lemma 7

We show by induction on n and subinduction on ϕ : if $I \upharpoonright_{n+2} = I' \upharpoonright_{n+2}$, then for any \mathcal{L} -formula ϕ with $\text{kd}(\phi) \leq n$, we have

$$\mathcal{G}, (\chi, I, \rho) \Vdash \phi \quad \text{iff} \quad \mathcal{G}, (\chi, I', \rho) \Vdash \phi. \quad (11)$$

Case $n = 0$. The formula ϕ does not contain any \mathbf{K} -modality, and (11) trivially holds.

Case $n = k + 1$. For the subinduction, we only show the case $\phi = \mathbf{K}_a \psi$. We show the direction from left to right by contraposition. The other direction is analogous. Assume

$$\mathcal{G}, (\chi, I', \rho) \not\Vdash \mathbf{K}_a \psi.$$

Thus, there exists an information perspective I'_0 such that

$$\mathcal{G}, (\chi, I'_0, \rho) \not\Vdash \psi \quad (12)$$

and

$$(\chi, I', \rho) \leq_a (\chi, I'_0, \rho). \quad (13)$$

Hence

$$I'[a] \subseteq I'_0. \quad (14)$$

Now we set

$$I_0 := I[a] \cup I'_0. \quad (15)$$

We trivially have

$$I[a] \subseteq I_0. \quad (16)$$

Further we have

$$I_0 \upharpoonright_{n+1} = I'_0 \upharpoonright_{n+1} . \quad (17)$$

Indeed, \supseteq holds by definition (15). For \subseteq , we assume $w \in I_0 \upharpoonright_{n+1}$ and distinguish the cases according to (15): If $w \in I'_0$, then $w \in I'_0 \upharpoonright_{n+1}$ and we are done. If $w \in I[a]$, then we distinguish:

1. $aw \in I$. We have $aw \in I \upharpoonright_{n+2} = I' \upharpoonright_{n+2}$. Thus $w \in I'[a]$ and by (14), we get $w \in I'_0$. Hence $w \in I'_0 \upharpoonright_{n+1}$.
2. $w \in I$ and $w = aw'$ for some w' . Similar to the previous case.

From (17) and (12), we get by I.H. that $\mathcal{G}, (\chi, I_0, \rho) \not\models \psi$. Further we have $I \upharpoonright_2 = I' \upharpoonright_2$ by assumption and $I_0 \upharpoonright_2 = I'_0 \upharpoonright_2$ by (17). We obtain $I_a = I'_a$ and $(I_0)_a = (I'_0)_a$, respectively. Hence, using (13), we find that (χ, I, ρ) and (χ, I_0, ρ) are a -consistent. Together with (16), this yields

$$(\chi, I, \rho) \preceq_a (\chi, I_0, \rho).$$

We finally conclude $\mathcal{G}, (\chi, I, \rho) \not\models K_a \psi$.

References

1. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. *Journal of the ACM* **49**(5), 672–713 (2002). <https://doi.org/10.1145/585265.585270>
2. Artemov, S.: On definitive solutions of strategic games. In: Baltag, A., Smets, S. (eds.) *Johan van Benthem on Logic and Information Dynamics*, pp. 487–507. Springer International Publishing (2014). https://doi.org/10.1007/978-3-319-06025-5_17
3. Artemov, S.: Beyond knowledge of the model. In: Weiss, Y., Birman, R. (eds.) *Saul Kripke on Modal Logic*, pp. 23–41. Springer International Publishing (2024). https://doi.org/10.1007/978-3-031-57635-5_3
4. Bard, N., Foerster, J.N., Chandar, S., Burch, N., Lanctot, M., Song, H.F., Parisotto, E., Dumoulin, V., Moitra, S., Hughes, E., Dunning, I., Mourad, S., Larochelle, H., Bellemare, M.G., Bowling, M.: The Hanabi challenge: A new frontier for AI research. *Artificial Intelligence* **280**, 103216 (2020). <https://doi.org/https://doi.org/10.1016/j.artint.2019.103216>
5. Belardinelli, F., Knight, S., Lomuscio, A., Maubert, B., Murano, A., Rubin, S.: Reasoning about agents that may know other agents' strategies. In: Zhou, Z.H. (ed.) *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*. pp. 1787–1793. International Joint Conferences on Artificial Intelligence Organization (2021). <https://doi.org/10.24963/ijcai.2021/246>, <https://doi.org/10.24963/ijcai.2021/246>
6. Bucheli, S., Kuznets, R., Studer, T.: Two ways to common knowledge. *Electronic Notes in Theoretical Computer Science* **262**, 83–98 (2010). <https://doi.org/10.1016/j.entcs.2010.04.007>, *Proceedings of the 6th Workshop on Methods for Modalities (M4M-6 2009)*
7. Cachin, C., Guerraoui, R., Rodrigues, L.: *Introduction to Reliable and Secure Distributed Programming*. Springer (2011)

8. Cachin, C., Lehnherr, D., Studer, T.: Synergistic knowledge. *Theoretical Computer Science* **1023**, 114902 (2025). <https://doi.org/https://doi.org/10.1016/j.tcs.2024.114902>
9. Chatterjee, K., Henzinger, T.A., Piterman, N.: Strategy logic. *Information and Computation* **208**(6), 677–693 (2010). <https://doi.org/https://doi.org/10.1016/j.ic.2009.07.004>
10. Cox, C., De Silva, J., Deorsey, P., Kenter, F.H.J., Retter, T., Tobin, J.: How to make the perfect fireworks display: Two strategies for Hanabi. *Mathematics Magazine* **88**(5), 323–336 (2015). <https://doi.org/10.4169/math.mag.88.5.323>
11. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: *Reasoning about Knowledge*. MIT Press (1995)
12. Goubault, É., Kniazev, R., Ledent, J., Rajsbaum, S.: Simplicial models for the epistemic logic of faulty agents. *Boletín de la Sociedad Matemática Mexicana* **30**(3) (2024). <https://doi.org/10.1007/s40590-024-00656-x>
13. van der Hoek, W., Wooldridge, M.J.: Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica* **75**(1), 125–157 (2003). <https://doi.org/10.1023/A:1026185103185>
14. Hu, H., Lerer, A., Peysakhovich, A., Foerster, J.: "Other-play" for zero-shot coordination. In: *Proceedings of the 37th International Conference on Machine Learning*. JMLR.org (2020)
15. Kokkinis, I., Studer, T.: Cyclic proofs for linear temporal logic. In: Probst, D., Schuster, P. (eds.) *Concepts of Proof in Mathematics, Philosophy, and Computer Science*. pp. 171–192. De Gruyter, Berlin, Boston (2016). <https://doi.org/doi:10.1515/9781501502620-011>
16. Luo, J., Studer, T.: Strategic knowledge in multi-agent systems (in preparation)
17. Marinković, B., Glavan, P., Ognjanović, Z., Studer, T.: A temporal epistemic logic with a non-rigid set of agents for analyzing the blockchain protocol. *Journal of Logic and Computation* **29**(5), 803–830 (2019)
18. Maubert, B., Murano, A.: Reasoning about knowledge and strategies under hierarchical information. In: Thielscher, M., Toni, F., Wolter, F. (eds.) *Principles of Knowledge Representation and Reasoning 2018*. pp. 530–540. AAAI Press (2018)
19. Menéndez Turata, G.: Cyclic proof systems for modal fixpoint logics. Ph.D. thesis, Universiteit van Amsterdam (2024)
20. Mogavero, F., Murano, A., Perelli, G., Vardi, M.Y.: Reasoning about strategies: On the model-checking problem. *ACM Trans. Comput. Logic* **15**(4), 1–47 (2014). <https://doi.org/10.1145/2631917>
21. Pauly, M.: A modal logic for coalitional power in games. *Journal of Logic and Computation* **12**(1), 149–166 (2002). <https://doi.org/10.1093/logcom/12.1.149>
22. Perrotin, E.: A logical analysis of Hanabi. In: *Proceedings of the 39th AAAI Conference on Artificial Intelligence* (2025)
23. Saffidine, A., Schwarzenruber, F., Zanuttini, B.: Knowledge-based policies for qualitative decentralized POMDPs. *Proceedings of the AAAI Conference on Artificial Intelligence* **32**(1) (2018). <https://doi.org/10.1609/aaai.v32i1.12085>, <https://ojs.aaai.org/index.php/AAAI/article/view/12085>
24. Sierra Miranda, B., Studer, T.: Knowledge and common knowledge of strategies (2025), <https://arxiv.org/abs/2510.19298>
25. Sierra Miranda, B., Studer, T., Zenger, L.: Coalgebraic proof translations for non-wellfounded proofs. In: Ciabattini, A., Gabelaia, D., Sedlár, I. (eds.) *Advances in Modal Logic 2024*. pp. 527–548 (2024)
26. Volrath, Z.: H-group conventions (2026), <https://hanabi.github.io>, last accessed 19 January 2026