Simplicial Structures for Epistemic Reasoning in Multi-agent Systems

Inaugural dissertation
of the Faculty of Science,
University of Bern

Presented by

David Benjamin Lehnherr
from Wädenswil, Switzerland

Supervisors of the Doctoral Thesis:

Prof. Dr. Christian Cachin
University of Bern, Switzerland

Prof. Dr. Thomas Studer
University of Bern, Switzerland

Simplicial Structures for Epistemic Reasoning in Multi-agent Systems

Inaugural dissertation
of the Faculty of Science,
University of Bern

Presented by

David Benjamin Lehnherr
from Wädenswil, Switzerland

Supervisors of the Doctoral Thesis:

Prof. Dr. Christian Cachin
University of Bern, Switzerland

Prof. Dr. Thomas Studer
University of Bern, Switzerland

Accepted by the Faculty of Science.

Bern,
8th December 2025

The Dean
Prof. Dr. Jean-Louis Reymond

Acknowledgments

"Every triangle is a love triangle if you love it..."

– Unknown

Good things come in pairs! I want to thank my advisors:

Dear Christian, I will be forever thankful for your generosity and letting me wander on my own path.
Dear Thomas, walking one's own path is always easier with company. I want to thank you for your guidance.

I want to thank Hans for our lovely interactions during visits. To my beloved coworkers: thank you for the (occasionally intellectual) lunchtime sparring sessions. I'm pretty sure I won them all. Special thanks goes to Annalisa, Djanira, and Juan for proofreading this thesis.

I thank Valentin for many stimulating discussions, and François-Xavier for many joyful conversations about cryptography and privacy.

Special thanks goes to Mike and Vanessa for making me feel at home in a place I no longer belong. On another note, I'd also like to thank Lorenz and Ben for sticking with me through life and all its surprises.

Throughout my PhD, I have experienced more physical pain than ever before. I would like to thank the Rugby Club Bern for providing this balance.

And of course, what would those acknowledgments be without thanking both my wonderful parents, who were always there for me. To my mother:

Dear mom, without you, none of this would have been possible – you know.

The final lines of these acknowledgments are dedicated in remembrance to my grandmothers. Dear Elena, no one else's meals ever tasted as much like home as yours did. Dear Frieda, your influence on my life has been immeasurable, and my heart is filled with gratitude having been able to accompany you at the end of your journey – just as you did when mine started. With all my love.

David

Abstract

Over the past decades, using modal logic to represent the knowledge of processes has proven to be a powerful framework for formally reasoning about distributed systems. Another story of success is modeling the views of processes with the help of (chromatic) simplicial complexes. Due to the rich toolbox of combinatorial topology, this approach has yielded many fundamental insights. However, the connection between these two models, namely simplicial semantics for modal logic, has been discovered only recently.

This thesis continues the structural study of simplicial semantics by changing the standard properties of chromatic simplicial complexes. Originally, a colored vertex is interpreted as a possible local state of the agent corresponding to that color. Global states are certain faces of the simplicial complex, and indistinguishability is based on the containment of vertices. We investigate three different simplicial structures by altering simplicial complexes in the following separate ways: allowing adjacent vertices to share the same color (polychromatic complexes), introducing directed faces (directed complexes), and permitting parallel faces (semi-simplicial sets).

We explore polychromatic complexes and directed simplicial complexes as models for belief. Different to knowledge, beliefs may be false. Models for belief are of great importance when using simplicial models to reason about distributed systems with malicious agents. In such a setting, an adversarial agent can manipulate honest agents. Thus, trustworthy agents may need to make decisions based on their beliefs.

Polychromatic complexes allow us to doxastically interpret the multiplicity of a color within a face. Different to the original model, vertices correspond to doxastic states of an agent. The fewer doxastic alternatives an agent has in a face, the more plausible it becomes. If there is exactly one doxastic state, belief becomes knowledge. We analyze the notion of most plausible belief, i.e., what an agent believes when only considering adjacent faces with the lowest multiplicity of its color. The notion of most plausible belief is often hidden in cryptographic properties of the form: algorithm $\mathcal{A}$ guarantees property P with overwhelming probability. This statement could be read as: "it is most plausible that, when using $\mathcal{A}$, the property P holds". Lastly, we explore how an alternative interpretation of polychromatic complexes can model quorum systems.

Directed complexes provide an intuitive framework for modeling belief, because

they preserve the interpretation of standard (undirected) simplicial complexes, i.e., vertices still correspond to "physical" local states. Directions are assigned to faces from the perspective of a vertex and can be either 1 or 0. We interpret 1 as being possible and 0 as being impossible. For evaluating belief, an agent only considers adjacent facets with direction 1. Notably, our models allow for merely introspective beliefs, as the accessibility relation is not required to be serial. This is achieved by assigning direction 0 to all adjacent facets. Moreover, we introduce a logic of belief and prove it to be sound and complete with respect to models based on directed complexes. To conclude, we position our directed models within the literature by showing them to be as expressive as one of the original variants of simplicial models.

A novelty of simplicial semantics is that the indistinguishability relation is based on connectivity, which implies that it inherits properties of the underlying structure. A generalization of simplicial complexes are semi-simplicial sets, which may contain parallel faces. This indicates that there might be indistinguishability relations based on higher-order connectivity, and not just on the containment of vertices. Unlike usual notions of group knowledge, the rich structure of semi-simplicial sets allows us to unfold relations (or synergies) among members of a group. We refer to this construct as an agent pattern. Moreover, we introduce an explicit representation of semi-simplicial sets and identify an indistinguishability relation for agent patterns. Our indistinguishability relation induces a new modality, which we term synergistic knowledge. We present the logic of synergistic knowledge and show its soundness and completeness with respect to models based on semi-simplicial sets. Furthermore, we illustrate applications of our models for various settings of distributed computing. Examples include the hierarchy of consensus and the dining cryptographers problem. Finally, we discuss a different reading of agent patterns, which is still based on our indistinguishability relation, inducing an additional modality that enables reasoning about group knowledge of processes with respect to the underlying network topology.

Contents

1 Introduction

In 1993, three independent groups (cf. Herlihy and Shavit [37], Borowsky and Gafni [12], and Saks and Zaharoglou [47]) confirmed Chaudhuri's [18] conjecture on the impossibility of the *asynchronous k-set agreement* (cf. Herlihy and Rajsbaum [36]). Each team proved the result using distinct models of distributed computation. In this thesis, we look at the approach of Herlihy and Shavit [37], who pioneered the use of simplicial complexes in distributed systems, through the lens of modal logic.

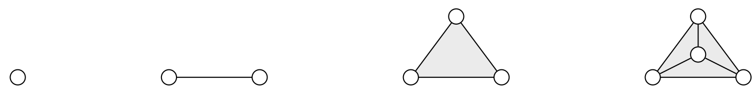

Figure 1.1: Examples of simplices.

A *simplex* is a generalization of a triangle. Figure 1.1 shows all simplices containing at most four vertices. The composition of simplices is called a *simplicial complex*. Formally, given a non-empty set of vertices $\mathcal{V}$, a simplicial complex is a pair $(\mathcal{V}, \mathcal{S})$, where $\mathcal{S}$ is a set of non-empty subsets of $\mathcal{V}$ closed under set inclusion. The elements of $\mathcal{S}$ are called *faces*. The simplicial complex in Figure 1.2 is obtained by gluing two solid triangles along one edge. In total, it has 11 faces: four vertices, five edges, and two triangles.

Herlihy and Shavit [37] interpret vertices as local states of processes. Each vertex is labeled with an identifier of a process, referred to as its *color*, indicating to which agent the local state belongs. Vertices sharing a face differ in colors, and their corresponding local states are mutually compatible. In Figure 1.2, the local states p, q, and r are compatible. However, the local states r and s are not, because they are contained in different faces. In this context, a global state is defined as a maximal set of mutually compatible local states, or, in simplicial terms, a *facet* of the simplicial complex. The simplicial complex in Figure 1.2 contains two facets, namely the two triangles.

Defining global states in terms of facets highlights the importance of processes' local states rather than the system's global state. Under this interpretation, a process cannot distinguish between two global states if and only if its local state is included in both. In Figure 1.2, a process inhabiting p or q cannot distinguish between the global state composed of the vertices p, q, and r and the one containing p, q, and s. As a result, this topological approach provides the necessary machinery to carry out indistinguishability proofs in distributed systems (cf. Attiya and Rajsbaum [5]). Over the years, this model of computation has proven to be powerful due to having access to strong results from combinatorial topology (cf. Herlihy et al. [35]).

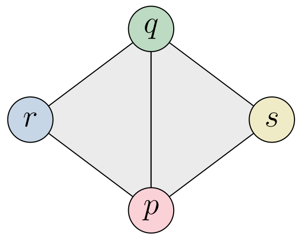

Figure 1.2: A simplicial complex that consists of two solid triangles sharing the pq-edge.

Prior to the rise of topological models, applying modal logic to describe the knowledge of agents (or processes) throughout a computation was the gold standard of formal reasoning in distributed systems. One of the most prominent results is the one-to-one correspondence between solving consensus and obtaining common knowledge (cf. Halpern and Moses [33]). In general, linking the solvability of a problem to the knowledge of the participating agents has yielded many foundational results (cf. Fagin et al. [20] and van Ditmarsch et al. [53]).

The standard semantics of modal logic is based on labeled graphs and is referred to as *relational semantics*. The nodes of the graph are called *worlds* and represent global states. An undirected edge with label a between worlds w and v means that agent a cannot distinguish the two worlds.

Apart from connectivity, worlds have no individual meaning yet. A *valuation* is a function that characterizes worlds by assigning them to a set of propositional variables. For example, a propositional variable p_a may describe a part of the local state of agent a. Thus, if p_a is assigned to a world w, agent a's local state satisfies the property associated with p_a. We say that, at a world w, a formula ϕ is known by agent a if and only if ϕ is true in each world that a cannot distinguish from w.

In 2018, Goubault et al. [28] discovered the link between the two approaches[1], laying the foundation for a new field of research: *simplicial semantics for modal logic*. Instead of reasoning about knowledge on labeled graphs, Goubault et al. [28] model knowledge on simplicial complexes. Each propositional variable describes a local attribute of an agent and may only be assigned to vertices with its color.

[1]The subsequent journal publication is available in Goubault et al. [29].

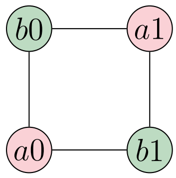

Figure 1.3: A simplicial complex that shows all possible input configurations of a two-party binary consensus protocol.

The proposition p holds in a facet (global state) if and only if p was assigned to a vertex of that facet. Consider an agent a in local state v belonging to a facet X. Agent a knows a formula ϕ if and only if ϕ is true in all facets containing v. Based on the groundwork of Herlihy and Shavit [37], an epistemic interpretation of simplicial complexes in the context of distributed systems is possible.

At first, only simplicial complexes whose facets all contain as many vertices as there are agents were studied. Such complexes correspond to settings in which processes do not crash because every agent is present in each global state. Goubault et al. [28] proposed a modal logic and showed its soundness and completeness with respect to models based on this class of complexes. In addition to the standard S5 axioms, their logic contains the *locality axiom*, stating that agents always know the truth value of their own local propositions. Although processes are often subjected to crashes, these restricted simplicial complexes can already be used to model interesting scenarios.

Figure 1.3 shows a simplicial complex representing all input configurations of a binary consensus protocol between two agents, a and b. Each agent proposes a value that is either 0 or 1. After completing the protocol, the agents must agree on one value. A vertex ai is of color a and represents that a's input is i, and an edge $\{ai, bj\}$ is the global state where a has input i and b has input j. When agent a has input 0, it cannot distinguish the global states $\{a0, b1\}$ and $\{a0, b0\}$ because their intersection contains the vertex $a0$. Hence, in the facet $\{a0, b1\}$, agent a considers it possible that b's input is 1 or 0. In other words, agent a does not know agent b's input.

Goubault et al. [28] also interpreted dynamic epistemic logic (cf. van Ditmarsch et al. [53]) on the previously mentioned type of simplicial complexes. This allowed them to model task solvability as done originally by Herlihy and Shavit [37]. In the original simplicial complex framework, the usual approach for proving impossibility involves constructing a topological obstruction. Goubault et al. [28] proposed the so-called *logical approach*, which establishes a logical obstruction by analyzing the knowledge agents must acquire in order to solve the task. It is important to point out that both approaches rely on the connectivity of the simplicial complex. Subsequently, van Ditmarsch et al. [51] used the logical approach to analyze additional distributed tasks. The logical approach was further extended by Hoshino [38], Nishimura [42], and Yagi and Nishimura [54] in various ways.

In general, simplicial complexes may contain facets of arbitrary size. A facet with fewer vertices than there are agents represents that some agents are not present (or have crashed). Therefore, simplicial complexes are compelling structures that are well suited for studying distributed systems that are subjected to crashes. The simplicial complex shown in Figure 1.4 contains two facets, X and Y, each containing a different number of vertices. The facet Y is interpreted as a global state in which agent c is not present, because it does not contain a c-vertex.

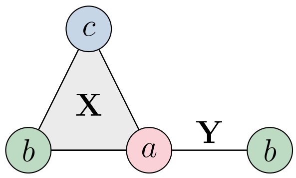

Figure 1.4: A simplicial complex whose facets may be of different size.

From the epistemic standpoint, such complexes are harder to handle. If an agent's vertex is missing in a facet, there is no vertex to assign propositional variables to. Consequently, formulas concerning local attributes of crashed agents might be undefined. Randrianomentsoa et al. [46] identified this problem and provided a three-valued logic, with the third value being *undefined*, that is sound and complete with respect to their simplicial models. Besides facets containing different amounts of vertices, the authors also allow arbitrary faces of the complex to represent global states. Moreover, Bílková et al. [11] define a notion of bisimulation to compare arbitrary simplicial complexes with respect to the proposed three-valued logic. Lastly, the recent approach of Yang [55] avoids a three-valued logic by using assignment operators that syntactically indicate if agents are present.

To avoid the problem of missing vertices, Goubault et al. [27, 30] suggested assigning propositional variables to arbitrary faces instead of vertices. Using this method, statements about crashed agents are always defined, and every face of a complex can be interpreted as a global state. The authors achieve further expressivity by extending the language with a modal operator describing the distributed knowledge of a group (cf. Halpern and Moses [33]). They also prove their logic to be sound and complete. A formula ϕ is distributed knowledge among a non-empty group of agents if and only if ϕ is true in all global states considered possible by each of its members. The logical approach was adapted to general simplicial

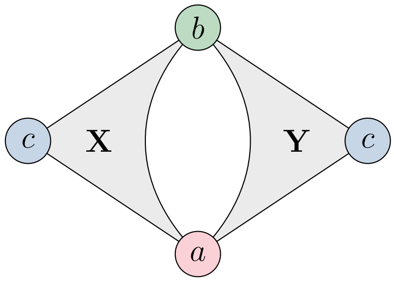

Figure 1.5: A semi-simplicial set containing parallel edges.

complexes by Nakai et al. [41].

Both methods involve trade-offs. Assigning propositional variables to vertices aligns closely with the original idea of emphasizing local states. However, using a three-valued logic or assignment operators makes the logic more difficult to apply. Conversely, assigning propositional variables to faces eliminates the problem, but also shifts the focus from a local to a more global perspective. The global approach does not require a locality axiom, but instead needs an axiom concerning the distributed knowledge of all alive agents. Randrianomentsoa et al. [45] and, earlier, Bílková et al. [11] propose an alternative approach incorporating both points of view. However, propositional variables associated with global states solely concern the liveness of agents.

On a different note, Loreti and Quadrini [39] analyzed simplicial models similar to the ones suggested by Goubault et al. [30] with respect to spatial logic instead of epistemic logic. In their model, vertices represent agents instead of local states of agents, and faces correspond to interactions or physical space. Consequently, vertices are not colored.

Until now, the underlying structure of a simplicial model has been a standard simplicial complex. However, another line of research that emerged in parallel is altering the properties of simplicial complexes and analyzing the resulting semantics. An example are semi-simplicial sets, which generalize simplicial complexes. Informally, a semi-simplicial set is a simplicial complex that may contain parallel faces. That is, if two simplices share a set of vertices, they need not share a face containing those vertices. Figure 1.5 shows such a complex where two solid triangles share two vertices, but not an edge. Goubault et al. [24, 26] and Randrianomentsoa [45] studied these structures as models for (distributed) knowledge. Another generalization of simplicial complexes are hypergraphs, which provide a more succinct representation of simplicial complexes, and were studied by Goubault et al. [25].

Overview of the Thesis

This thesis continues the structural study of simplicial semantics and can be divided into two parts. The first part is about modeling belief on simplicial structures. As for knowledge, an agent believes ϕ if and only if ϕ is true in all possible worlds. However, unlike knowledge, belief may be false. Modeling belief on simplicial structures is one of the open questions posed in the report from a Dagstuhl seminar on epistemic and topological reasoning in distributed systems by Castañeda et al. [17]. Having models for belief enables us to represent settings in which agents may act based on their beliefs.

The second part of this thesis examines semi-simplicial sets from the global point of view. Figure 1.5 can be interpreted as agents a and b individually not being able to distinguish the two facets. Together, however, they can distinguish them, because the facets do not share an ab-edge. This is different from distributed knowledge, which dictates that the agents cannot distinguish the two states, because they cannot do so on an individual basis.

Structure of the Thesis

After a preliminary chapter, we analyze simplicial complexes where adjacent vertices can share a color in Chapter 3. Such complexes enable us to define more intricate indistinguishability relations, which can be used to model different notions of belief. This chapter is based on the work:

- Christian Cachin, David Lehnherr, and Thomas Studer. Simplicial belief. In Ulrich Schmid and Roman Kuznets, editors, *Structural Information and Communication Complexity (SIROCCO)*, pages 176–193. Springer Nature Switzerland, 2025.

Such complexes allow us to model certain types of belief. However, their philosophical interpretation remains open. In Chapter 4, we analyze directed simplicial complexes as more transparent models for belief. This chapter describes a preliminary version of a result that was later generalized and extended in:

- Hans van Ditmarsch, Djanira Gomes, David Lehnherr, Valentin Müller, and Thomas Studer. Hypergraph semantics for doxastic logics. *Manuscript*, 2025.

In Chapter 5, we identify a new type of knowledge, called *synergistic knowledge*, which allows a group of agents to know more than just the consequences of their pooled knowledge. We analyze synergistic knowledge in detail and illustrate its applications in practice. This chapter is based on the two works:

- Christian Cachin, David Lehnherr, and Thomas Studer. Synergistic knowledge. In Shlomi Dolev and Baruch Schieber, editors, *Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 552–567. Springer, 2023.
- Christian Cachin, David Lehnherr, and Thomas Studer. Synergistic knowledge. *Theor. Comput. Sci.*, 1023:114902, 2025.

2 Preliminaries

Chapter Organization. We begin this chapter by defining our logical language. As outlined in the introduction, simplicial semantics can be approached from two perspectives: global and local. For each point of view, we present one relational model and one simplicial model. Section 2.1 and Section 2.2 introduce the principles of relational and simplicial models. They also provide formal definitions of the selected models. Although more nuanced simplicial models exist, we focus on simple models that are better suited for an introduction. Whenever relevant, we point out how the simple models differ from the richer ones. In Section 2.3, we identify the conditions under which the chosen relational and simplicial models can be transformed into one another.

Let Ag be a finite set of n agents, and let Prop be a countably infinite set of propositional variables. For $p \in \mathsf{Prop}$ and $a \in \mathsf{Ag}$, the language $\mathcal{L}$ is inductively defined by the following grammar:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid \Box_a \phi$$

We write $\Diamond_a \phi$ for $\neg\Box_a\neg\phi$, and define the remaining Boolean connectives as usual. In particular, we set $\phi \vee \psi = \neg(\neg\phi \wedge \neg\psi)$ and $\bot = p \wedge \neg p$, for some fixed $p \in \mathsf{Prop}$. We write $\mathsf{alive}(a)$ for $\neg\Box_a\bot$, and $\mathsf{dead}(a)$ for $\Box_a\bot$. Throughout this work, we assume a fixed a partition of Prop that consists of pairwise disjoint sets $\mathsf{P}_1, \dots, \mathsf{P}_n$, and assign to each agent $a \in \mathsf{Ag}$ its own set of propositional variables P_i, which we denote with P_a.

2.1 Relational Models

This section reviews standard concepts of relational semantics and introduces specific notions used to model both the global and local approach.

Definition 1 (Frame). *A* frame *is a pair $F = (W, R)$ such that:*

1. *W is a set of possible worlds;*
2. *R is a function assigning a relation R_a on W to each agent $a \in \mathsf{Ag}$.*

Let P be a relational property. A frame F has property P if and only if the relation R_a satisfies P for all $a \in \mathsf{Ag}$. For example, if R_a is reflexive for all $a \in \mathsf{Ag}$, then F is reflexive. Moreover, a function that assigns each world $w \in W$ to a set of propositional variables $A \in \mathsf{Pow}(\mathsf{Prop})$, where $\mathsf{Pow}(\cdot)$ denotes the powerset, is called a *valuation*. Lastly, a *relational model* is a frame F equipped with a valuation V.

Definition 2 (Relational Model)**.** *A relational model is a pair $\mathcal{M} = (F, V)$ such that:*

1. *$F = (W, R)$ is a frame;*
2. *$V : W \to \mathsf{Pow}(\mathsf{Prop})$ is a valuation.*

Definition 3 ($\Vdash$)**.** *Let $F = (W, R)$ be a frame, and let $\mathcal{M} = (F, V)$ be a relational model. For all worlds $w \in W$, we define the relation $\mathcal{M}, w \Vdash \phi$ by induction on $\phi \in \mathcal{L}$:*

$$\begin{array}{lll} \mathcal{M}, w \Vdash p & \textit{iff} & p \in V(w) \\ \mathcal{M}, w \Vdash \neg\phi & \textit{iff} & \mathcal{M}, w \nVdash \phi \\ \mathcal{M}, w \Vdash \phi \wedge \psi & \textit{iff} & \mathcal{M}, w \Vdash \phi \textit{ and } \mathcal{M}, w \Vdash \psi \\ \mathcal{M}, w \Vdash \Box_a\phi & \textit{iff} & wR_av \textit{ implies } \mathcal{M}, v \Vdash \phi, \textit{ for all } v \in W. \end{array}$$

Let $F = (W, R)$ be a frame, and let $\mathcal{M} = (F, V)$ be a relational model. If $\mathcal{M}, w \Vdash \phi$ holds, then ϕ is said to be *satisfied* at the world w. A formula $\phi \in \mathcal{L}$ is *valid in* $\mathcal{M}$, denoted $\mathcal{M} \Vdash \phi$, if it is satisfied at all worlds. Lastly, ϕ is said to be *valid*, written $\Vdash \phi$, if ϕ is valid in every relational model.

Relational models that reason about the knowledge of agents are often based on frames where each R_a is a partial equivalence relation (symmetric and transitive). Whenever a frame is symmetric and transitive, we write $F = (W, \sim)$ instead of $F = (W, R)$ to emphasize these properties. In a local model, each agent has reliable access to its own information, meaning that it always knows the truth value of its own local propositions. This property is called *locality*. It is not determined by the structure of the frame itself, but rather by how the valuation of the model is defined.

Definition 4 (Local Valuation)**.** *Let $F = (W, \sim)$ be a symmetric and transitive frame. We call a valuation $V : W \to \mathsf{Pow}(\mathsf{Prop})$ local if and only if, for all worlds $s, t \in W$ with $s \sim_a t$, and all $p \in \mathsf{P}_a$, it holds that:*

$$p \in V(s) \textit{ if and only if } p \in V(t).$$

A *local* relational model is a pair $\mathcal{M} = (F, V)$, where F is symmetric and transitive, and V is local. If $\mathcal{M}$ is a local relational model, we write $\mathcal{M}, w \Vdash^{\mathsf{loc}} \phi$ instead of $\mathcal{M}, w \Vdash \phi$ for all $\phi \in \mathcal{L}$. Validity is defined in the same way as for arbitrary relational models. It immediately follows from Definition 4 that such models validate locality, i.e., for all $a \in \mathsf{Ag}$ and $p \in \mathsf{P}_a$:

$$\Vdash^{\mathsf{loc}} \Box_a p \vee \Box_a \neg p.$$

Lastly, for clarity and consistency, we refer to relational models that are based on symmetric and transitive frames, but may not be local, as *global* relational models. The same notational conventions used for local models apply to global models. If $\mathcal{M}$ is a global model, we omit a superscript and simply write $\mathcal{M}, w \Vdash \phi$. The models presented in this thesis either resemble local or global models, and we will use the same notation for them.

2.2 Simplicial Models

We now introduce the standard definitions of simplicial semantics and provide examples of the selected models.

Definition 5 (Simplicial Complex)**.** *Let $\mathcal{V}$ be a set of vertices. The pair $\mathbb{C} = (\mathcal{V}, \mathcal{S})$ with $\mathcal{S} \subseteq \mathsf{Pow}(\mathcal{V}) \setminus \{\emptyset\}$ is called a* simplicial complex *if:*

$$\textit{for each } X \in \mathcal{S} \textit{ and each } \emptyset \neq Y \subseteq X, \textit{ we have } Y \in \mathcal{S}.$$

Remark 1. *Herlihy et al. [35] originally defined a simplicial complex to contain the singleton set $\{v\}$ for each vertex $v \in \mathcal{V}$. Not enforcing this is a purely cosmetic choice.*

The elements of $\mathcal{S}$ are called *faces*. A face of $\mathbb{C}$ that is maximal under set inclusion is a *facet*, and we denote the set of $\mathbb{C}$'s facets by $\mathcal{F}(\mathbb{C})$. If all facets contain the same number of elements, the complex is *pure*, otherwise, it is *impure*. The *dimension* of a face $X \in \mathcal{S}$ is given by $|X| - 1$, where $|X|$ denotes the cardinality of the set X. The dimension of a simplicial complex is the dimension of its largest facet. Throughout this thesis, we reserve the term "pure" for simplicial complexes with dimension $|\mathsf{Ag}| - 1$. If $X \subseteq Y$ for two faces X and Y, then X is called a *subface* of Y. A *coloring* is a mapping $\chi : \mathcal{V} \to \mathsf{Ag}$. It is called *proper* if it assigns a different value to each vertex within a face, and *improper* otherwise. A *chromatic simplicial complex* is a triple $(\mathcal{V}, \mathcal{S}, \chi)$ where $(\mathcal{V}, \mathcal{S})$ is a simplicial complex, and χ is a proper coloring on $\mathcal{V}$. Whenever the context is clear, we use the terms chromatic simplicial complex, simplicial complex, and complex interchangeably. Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$ be a complex. We say that an agent a cannot distinguish between two facets $X, Y \in \mathcal{F}(\mathbb{C})$ if and only if $a \in \chi(X \cap Y)$. The concept of knowledge can be extended to arbitrary faces and non-empty groups of agents: a group $G \subseteq \mathsf{Ag}$ cannot distinguish two faces X and Y if and only if $G \subseteq \chi(X \cap Y)$, where $\chi(X \cap Y) = \{\chi(u) \mid u \in X \cap Y\}$. If a properly colored facet X contains an a-vertex, then we denote that vertex as X_a.

A function assigning either vertices or facets to sets of propositional variables is called a *labeling*. A *simplicial model* $\mathcal{C} = (\mathbb{C}, L)$ is a chromatic simplicial complex $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$ equipped with a labeling L.

Definition 6 (Simplicial Model)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$ be a chromatic simplicial complex. A simplicial model is a pair $\mathcal{C} = (\mathbb{C}, L)$, where L is a labeling.*

If the labeling L assigns facets to sets of proportionals variables, then $\mathcal{C}$ is called a *global simplicial model*.

Definition 7 ($\Vdash_\sigma$)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$ be a chromatic simplicial complex and let $\mathcal{C} = (\mathbb{C}, L)$ be a global simplicial model. For all facets $X \in \mathcal{F}(\mathbb{C})$, we define the relation $\mathcal{C}, X \Vdash_\sigma \phi$ by induction on $\phi \in \mathcal{L}$:*

$$
\begin{array}{lll}
\mathcal{C}, X \Vdash_\sigma p & \textit{iff} & p \in L(X) \\
\mathcal{C}, X \Vdash_\sigma \neg\phi & \textit{iff} & \mathcal{C}, X \nVdash_\sigma \phi \\
\mathcal{C}, X \Vdash_\sigma \phi \wedge \psi & \textit{iff} & \mathcal{C}, X \Vdash_\sigma \phi \textit{ and } \mathcal{C}, X \Vdash_\sigma \psi \\
\mathcal{C}, X \Vdash_\sigma \Box_a\phi & \textit{iff} & a \in \chi(X \cap Y) \textit{ implies } \mathcal{C}, Y \Vdash_\sigma \phi, \textit{ for all } Y \in \mathcal{F}(\mathbb{C}).
\end{array}
$$

Remark 2. *A generalized version of global simplicial models has been introduced by Goubault et al. [27], where the labeling, and hence the definition of the relation $\Vdash_\sigma$, is extended to arbitrary faces.*

Global simplicial models allow us to reason about dead agents because the underlying simplicial complex need not be pure (see Example 1). An agent a is considered dead in a facet X if and only if $\mathcal{C}, X \Vdash \Box_a\bot$.

Example 1. *Figure 2.1 shows an impure chromatic complex $\mathbb{C}$ over a set of vertices $\mathcal{V} = \{a, b, b', c\}$ that contains exactly two facets:*

$$X = \mathsf{Pow}(\{a, b, c\}) \setminus \{\emptyset\} \quad \textit{and} \quad Y = \mathsf{Pow}(\{a, b'\}) \setminus \{\emptyset\}.$$

The coloring of the vertices χ is determined by their names, i.e., $\chi(a) = a$ and $\chi(b') = b$. Further, let $\mathsf{Prop} = \{p\}$, and consider the global simplicial model $\mathcal{C} = (\mathbb{C}, L)$, where the labeling L is such that $p \in L(X)$ and $p \notin L(Y)$.

The agent a cannot distinguish the facets X and Y because its local state is contained in both, i.e., $a \in \chi(X \cap Y)$. Moreover, agent c is dead in Y because Y does not contain a vertex with color c. Thus, the following holds:

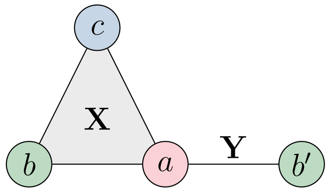

Figure 2.1: The impure chromatic simplicial complex discussed in Example 1.

- $\mathcal{C}, X \not\Vdash_\sigma \Box_a p$*;*
- $\mathcal{C}, Y \Vdash_\sigma \mathsf{dead}(c)$*, and* $\mathcal{C}, X \not\Vdash_\sigma \Box_a \mathsf{alive}(c)$*.*

We now introduce *local simplicial models.* Intuitively, every propositional variable describes a component of the state of a single agent. Consequently, vertices are assigned to sets of local propositional variables. Definition 8 formalizes this intuition: a labeling is *local* if and only if it assigns each local state of an agent to a set of propositional variables that concern only that agent. A simplicial model $\mathcal{C} = (\mathbb{C}, L)$, where $\mathbb{C}$ is pure with dimension $|\mathsf{Ag}| - 1$ and L is a local labeling, is called a local simplicial model.

Definition 8 (Local Labeling)**.** *Let* $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$ *be a chromatic simplicial complex. A labeling* $L : \mathcal{V} \to \mathsf{Pow}(\mathsf{Prop})$ *is called a* local *labeling if and only if, for all vertices* $v \in \mathcal{V}$*:*

$$\chi(v) = a \text{ implies } L(v) \subseteq \mathsf{P}_a.$$

Assuming that the underlying simplicial complex is pure with dimension $|\mathsf{Ag}| - 1$ is done for simplicity. Hence, all agents are always alive. Since every facet contains a vertex of each agent's color, we can define truth in a facet similarly to global simplicial models (see Definition 9). Allowing the underlying simplicial complex of a local simplicial model to be impure would make statements about dead agents undefined. For example, what is the truth value of a local proposition $p \in \mathsf{P}_a$ at a facet that does not contain a vertex of color a? An extension of local simplicial models that evaluates satisfiability at arbitrary faces and whose underlying complex may be impure was proposed by Randrianomentsoa et al. [46]. The issue of undefined propositions is addressed by using a three-valued logic, with the third value being *undefined.*

Definition 9 ($\Vdash_\sigma^{\mathsf{loc}}$)**.** *Let* $\mathcal{C} = (\mathbb{C}, L)$ *be a local simplicial model. For all facets*

$X \in \mathcal{F}(\mathbb{C})$, we define the relation $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \phi$ by induction on $\phi \in \mathcal{L}$:

$$\begin{array}{lll} \mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} p & \textit{iff} & p \in L(X_a), \textit{ and } p \in \mathsf{P}_a \\ \mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \neg\phi & \textit{iff} & \mathcal{C}, X \not\Vdash^{\mathsf{loc}}_{\sigma} \phi \\ \mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \phi \wedge \psi & \textit{iff} & \mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \phi \textit{ and } \mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \psi \\ \mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a\phi & \textit{iff} & a \in \chi(X \cap Y) \textit{ implies } \mathcal{C}, Y \Vdash^{\mathsf{loc}}_{\sigma} \phi, \textit{ for all } Y \in W. \end{array}$$

Let $\mathcal{C} = (\mathbb{C}, L)$ be a local simplicial model. If $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \phi$ is true, we say that ϕ is *satisfied* at X in $\mathcal{C}$. Further, ϕ is *σ-valid in $\mathcal{C}$*, denoted by $\mathcal{C} \Vdash^{\mathsf{loc}}_{\sigma} \phi$, if it is satisfied at all $X \in \mathcal{F}(\mathbb{C})$. Moreover, ϕ is *σ-valid*, written $\Vdash^{\mathsf{loc}}_{\sigma} \phi$, if it is valid in all global simplicial models. The same notational approach is used for global simplicial models.

Example 2 shows a local simplicial model. It is straightforward to show that local simplicial models validate locality (see Example 2), i.e., for all $a \in \mathsf{Ag}$ and $p \in \mathsf{P}_a$, it holds that:

$$\Vdash^{\mathsf{loc}}_{\sigma} \Box_a p \vee \Box_a \neg p.$$

Example 2. *Figure 2.2 shows a pure chromatic complex $\mathbb{C}$ over a set of vertices $\mathcal{V} = \{1_a, 1_b, 1_c, 0_b\}$ that contains exactly the two facets:*

$$X = \mathsf{Pow}(\{1_a, 1_b, 1_c\}) \setminus \{\emptyset\} \quad \textit{and} \quad Y = \mathsf{Pow}(\{1_a, 1_b, 0_c\}) \setminus \{\emptyset\}.$$

The coloring of the vertices χ is determined by their names. Let $\mathsf{P}_a = \{p_a\}$, $\mathsf{P}_b = \{p_b\}$, and $\mathsf{P}_c = \{p_c\}$ be the sets of local propositional variables. Consider the local simplicial model $\mathcal{C} = (\mathbb{C}, L)$, where the labeling L is as indicated by the names of vertices, i.e., $L(1_c) = \{p_c\}$ and $L(0_c) = \emptyset$.

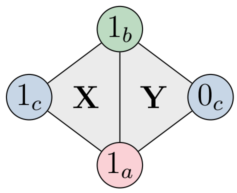

Figure 2.2: The pure simplicial complex discussed in Example 2.

The facets X and Y are indistinguishable to agents a and b. Therefore, neither agent knows the truth value of p_c. However, the model satisfies locality because $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a p_a \wedge \Box_b p_b \wedge \Box_c p_c$.

2.3 Correspondence

This section establishes the correspondence between global (local) relational models and global (local) simplicial models. The presented definitions and results for the global approach were originally established by Goubault et al. [27]. In the local case, we refer to Goubault et al. [28]. Let $F = (W, R)$ be a frame. For all worlds $w \in W$ and agents $a \in \mathsf{Ag}$, we define:

$$\mathsf{Alive}(w) = \{a \in \mathsf{Ag} \mid \exists v \in W. wR_a v\}.$$

Moreover, if F is symmetric and transitive, we define:

$$[w]_a = \{v \mid w \sim_a v\}, X_w = \{([w]_a, a) \mid a \in \mathsf{Alive}(w)\}, \text{ and } S_w = \mathsf{Pow}(X_w) \setminus \{\emptyset\}.$$

We begin by transforming a frame into a simplicial complex. In general, this cannot be done in a structure-preserving way unless the frame satisfies certain properties.

In our construction, the pairs $([w]_a, a)$ with $[w]_a \neq \emptyset$ represent the vertices, which reflects the idea that an agent has the same local state in two worlds if and only if it cannot distinguish between them. Definition 10 ensures that all worlds of the frame correspond to a simplex.

Definition 10 (Non-empty Frame)**.** *A frame $F = (W, R)$ is non-empty if and only if, for all worlds $w \in W$, there exists an agent $a \in \mathsf{Alive}(w)$.*

To avoid duplicate faces, every two distinct worlds must differ in a local state. Frames that meet this condition are called *proper* (Definition 11). Notice that if we considered only local simplicial models, the condition that $a \in \mathsf{Alive}(w) \cap \mathsf{Alive}(v)$ would not be needed, because all agents are assumed to be alive.

Definition 11 (Properness)**.** *Let $F = (W, \sim)$ be a symmetric and transitive frame. We say that F is* proper *if and only if, for all $w, v \in W$ where $w \neq v$, there exists an agent $a \in \mathsf{Alive}(w) \cap \mathsf{Alive}(v)$ such that $w \not\sim_a v$.*

Additionally, indistinguishability and connectivity must align in both structures. Since connectivity on simplicial complexes is a partial equivalence relation (symmetric and transitive), the frame must be symmetric and transitive (Definition 12). Definition 13 additionally imposes reflexivity on the frame. This requirement is needed only for local models.

Definition 12 (Partial Epistemic Frame)**.** *A frame $F = (W, \sim)$ is a partial epistemic frame if and only if each $\sim_a$ is a partial equivalence relation.*

Definition 13 (Epistemic Frame)**.** *A frame $F = (W, \sim)$ is an epistemic frame if and only if each $\sim_a$ is an equivalence relation.*

Definition 14 describes how to transform a partial epistemic frame into a chromatic complex. Example 3 illustrates the construction.

Definition 14 ($\mathbb{C}^F$)**.** *Let $F = (W, \sim)$ be a partial epistemic frame. We define the chromatic simplicial complex $\mathbb{C}^F = (\mathcal{V}, \mathcal{S}, \chi)$:*

- $\mathcal{V} = \{([w]_a, a) \mid a \in \mathsf{Ag}, w \in W, \text{ and } [w]_a \neq \emptyset\}$*;*
- $\mathcal{S} = \bigcup_{w \in W} S_w$*;*
- $\chi(([w]_a, a)) = a$*, for all* $w \in W$ *and* $a \in \mathsf{Ag}$*.*

Example 3. *Consider the set of worlds $W = \{w, v\}$ and the set of agents $\mathsf{Ag} = \{a, b, c\}$. Further, let $F = (W, \sim)$ be the partial epistemic frame as depicted in Figure 2.3 on the left.*

Figure 2.3: A frame F (left) and the corresponding chromatic simplicial complex $\mathbb{C}^F$ (right). We write $[w]_a$ instead of $([w]_a, a)$ for readability.

We will now apply our construction to F. The set of vertices of the corresponding simplicial complex is $\mathcal{V} = \{([w]_a, a), ([w]_b, b), ([w]_c, c), ([v]_b, b)\}$, where

$$[w]_a = [v]_a = \{w, v\},\ [w]_b = [w]_c = \{w\},\ \text{and } [v]_b = \{v\}.$$

Notice that since $[v]_c = \emptyset$, there is no vertex $([v]_c, c)$ in the set $\mathcal{V}$. The set $\mathcal{S}$ is obtained by building the union of the two simplices S_w and S_v. Lastly, the coloring is as indicated by the names of vertices, i.e., $\chi(([u]_a, a)) = a$, for all $a \in \mathsf{Ag}$ and $u \in W$ with $a \in \mathsf{alive}(u)$. The corresponding chromatic simplicial complex $\mathbb{C}^F$ is depicted in Figure 2.3. It is the same as shown in Figure 2.1.

If the partial epistemic frame F is non-empty, then $\mathcal{V} \neq \emptyset$ and $\mathbb{C}^F$ is a chromatic simplicial complex. This follows because $\mathcal{S}$ is the union of powersets that exclude the empty set, and $\chi(\cdot)$ is well-defined since $[w]_a \neq \emptyset$ for all vertices

$([w]_a, a) \in \mathcal{V}$. Non-emptiness is thus required for all subsequent definitions and lemmas. Nonetheless, the complex $\mathbb{C}^F$ may not fully reflect the structure of a non-empty frame.

We now show that if the partial epistemic non-empty frame F is proper in addition, then $\mathbb{C}^F$ preserves the structure of F. Lemma 1 states that an agent a cannot distinguish between two worlds w and v if and only if X_w and X_v both contain its local state.

Lemma 1. *Let $F = (W, \sim)$ be a non-empty partial epistemic frame, and consider the chromatic simplicial complex $\mathbb{C}^F = (\mathcal{V}, \mathcal{S}, \chi)$ as defined in Definition 14. For all $w, v \in W$ with $([w]_a, a), ([v]_a, a) \in \mathcal{V}$, it holds that:*

$$w \sim_a v \text{ iff } a \in \chi(X_w \cap X_v).$$

Proof. The direction from left to right is straightforward. The other direction, follows because $\sim_a$ is a partial equivalence relation. Indeed, by assumption we have that $[w]_a = [v]_a$. Further, since $([w]_a, a), ([v]_a, a) \in \mathcal{V}$ it holds that $[w]_a \neq \emptyset$ and $[v]_a \neq \emptyset$. Thus, there exists $u \in [w]_a \cap [v]_a$. By definition of $\mathcal{V}$, it holds that $w \sim_a u$ and $v \sim_a u$. Finally, by symmetry and transitivity of $\sim_a$, we find that $w \sim_a v$. □

Lemma 2 demonstrates that the sets X_w are exactly the facets of $\mathbb{C}^F$. Finally, Lemma 3 states that there are as many facets as worlds.

Lemma 2. *Let $F = (W, \sim)$ be a proper and non-empty partial epistemic frame, and let $\mathbb{C}^F = (\mathcal{V}, \mathcal{S}, \chi)$ be according to Definition 14. It holds that:*

$$\mathcal{F}(\mathbb{C}^F) = \{X_w \mid w \in W\}.$$

Proof. Let $X = \{X_w \mid w \in W\}$. We first observe that by construction, for any face $Y \in \mathcal{S}$ there exists $Z \in X$ with $Y \subseteq Z$. Thus, we immediately obtain $\mathcal{F}(\mathbb{C}^F) \subseteq X$. We proceed to show that $X \subseteq \mathcal{F}(\mathbb{C}^F)$. Towards a contradiction, assume that there exists $X_v \in X$ such that X_v is not a facet. Thus, there exists a face $Y \in \mathcal{S}$ and $X_v \subsetneq Y$. By our observation, there exists $w \in W$ such that $Y \subseteq X_w$. This implies $X_v \subsetneq X_w$, which contradicts the properness of F because $v \neq w$. □

Lemma 3. *Let $F = (W, \sim)$ be a proper and non-empty partial epistemic frame. The function:*

$$b(w) = X_w$$

is a bijection.

Proof. By Lemma 2, it holds that $\mathcal{F}(\mathbb{C}^F) = \{X_w \mid w \in W\}$ and surjectivity follows immediately from the definition of b. Regarding injectivity, we will prove that

$b(w) = b(v)$ implies $w = v$. Assume towards a contradiction that $b(w) = X_w = X_v = b(v)$ for two different worlds $w, v \in W$. Thus, for all $a \in \mathsf{Alive}(w) = \mathsf{Alive}(v)$, we find that $[w]_a = [v]_a$, and $w \sim_a v$ (by Lemma 1). Since we assumed that $w \neq v$, the properness of F implies that there exists an agent $b \in \mathsf{Alive}(w)$ with $w \not\sim_b v$. This is a contradiction. Hence, we conclude that $w = v$, and that b is injective. □

Remark 3. *Given the role of properness in the previous proofs, a slightly different formulation is: $F = (W, \sim)$ is proper if and only if, for all distinct worlds $w, v \in W$, there exists an agent $a \in \mathsf{Alive}(w) \cap \mathsf{Alive}(v)$ for which $[w]_a \neq [v]_a$. Although non-standard, this definition places greater emphasis on local states, i.e., no two different facets of $\mathbb{C}^F$ have the same set of vertices.*

Definition 15 states how a global relational model based on a proper and non-empty partial epistemic frame can be transformed into an equivalent global simplicial model.

Definition 15. *Let $F = (W, \sim)$ be a proper and non-empty partial epistemic frame, and let $\mathcal{M} = (W, V)$ be a global relational model. We construct the global simplicial model $\mathcal{C}^{\mathcal{M}}_{\mathsf{glob}} = (\mathbb{C}^F, L)$ as follows:*

- *$\mathbb{C}^F = (S, \mathcal{V}, \chi)$ is defined as in Definition 14;*
- *the labeling L is defined such that for all worlds $w \in W$ and propositional variables $p \in \mathsf{Prop}$:*
$$p \in L(X_w) \text{ iff } p \in V(w).$$

Lemma 4 establishes the *pointwise equivalence*[1] of every global relational model $\mathcal{M}$ and its corresponding global simplicial model $\mathcal{C}^{\mathcal{M}}_{\mathsf{glob}}$. This is achieved only if our construction induces a bijection between the points of both models, as ensured by Lemma 3. If there were no bijection, two distinct worlds w and v would map to the same facet X. Consequently, if a propositional variable p is satisfied at w but not at v, then X cannot satisfy both, which is a contradiction.

Lemma 4. *Let $F = (W, \sim)$ be a proper and non-empty partial epistemic frame, and let $\mathcal{M} = (F, V)$ be a global relational model. It holds that for all formulas $\varphi \in \mathcal{L}$:*
$$\mathcal{M}, w \Vdash \varphi \quad \text{iff} \quad \mathcal{C}^{\mathcal{M}}_{\mathsf{glob}}, X_w \Vdash_\sigma \varphi.$$

Proof. Follows by induction on the length of formulas and using the Lemmas 1, 2, and 3. □

[1] A global relational model and its corresponding global simplicial model are pointwise equivalent if and only if the same formulas are satisfied at every world of the relational model and its corresponding facet.

Transforming a local relational model follows the same principle, except that propositional variables are assigned to vertices instead of facets.

Definition 16. *Let $F = (W, \sim)$ be a proper and non-empty epistemic frame, and let $\mathcal{M} = (F, V)$ be a local relational model. We construct the local simplicial model $\mathcal{C}^{\mathcal{M}}_{\mathsf{loc}} = (\mathbb{C}^F, L)$ as follows:*

- *$\mathbb{C}^F = (\mathcal{V}, \mathcal{S}, \chi)$ is defined as in Definition 14;*
- *the labeling L is defined such that for all worlds $w \in W$, agents $a \in \mathsf{Ag}$, and propositional variables $p \in \mathsf{P}_a$:*

$$p \in L(([w]_a, a)) \textit{ iff } p \in V(w).$$

As previously noted, the frame F must be epistemic to ensure that the labeling L is well-defined. For example, let $\mathsf{Ag} = \{a, b\}$, and consider a proper and non-empty partial epistemic frame $F = (\{w\}, \sim)$, where $w \sim_a w$ but $w \not\sim_b w$. Further, let $\mathcal{M} = (F, V)$ be a local relational model such that $V(w) = \{p\}$ for some $p \in \mathsf{P}_b$. Since $[w]_b = \emptyset$, there is no b-vertex in $\mathbb{C}^F$, and the expression $L(([w]_b, b)) = p$ is undefined. If F is an epistemic frame, then $w \sim_b w$ is ensured, and $L(([w]_b, b)) = p$ becomes well-defined.

Lemma 5. *Let $F = (W, \sim)$ be a proper and non-empty partial epistemic frame, and let $\mathcal{M} = (F, V)$ be a local relational model. It holds that for all formulas $\varphi \in \mathcal{L}$:*

$$\mathcal{M}, w \Vdash^{\mathsf{loc}} \varphi \quad \textit{iff} \quad \mathcal{C}^{\mathcal{M}}_{\mathsf{loc}}, X_w \Vdash^{\mathsf{loc}}_{\sigma} \varphi.$$

Proof. Follows by induction on the length of formulas and using the reflexivity of F as well as the Lemmas 1, 2, and 3. □

Transforming chromatic simplicial complexes to frames, and consequently simplicial models to relational models, is less difficult. We only sketch the transformation and provide detailed constructions for similar models in subsequent chapters. Given a chromatic simplicial complex $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$, the corresponding frame $F = (W, \sim)$ is constructed as follows:

- for each facet $X \in \mathcal{F}(\mathbb{C})$, we define a world w_X;
- the set of worlds is given by $W = \{w_X \mid X \in \mathcal{F}(\mathbb{C})\}$;
- $w_X \sim_a w_Y$ if and only if $a \in \chi(X \cap Y)$.

It is straightforward to verify that F is a proper and non-empty partial epistemic frame. Moreover, if $\mathbb{C}$ is a pure chromatic complex with dimension $|\mathsf{Ag}| - 1$, then F is an epistemic frame. If $\mathcal{C} = (\mathbb{C}, L)$ is a global simplicial model, we define the valuation V such that:

$$p \in V(w_X) \text{ iff } p \in L(X).$$

Given a local simplicial model $\mathcal{C} = (\mathbb{C}, L)$, we define the valuation V such that for all agents $a \in \mathsf{Ag}$ and $p \in \mathsf{P}_a$, it holds that:

$$p \in V(w_X) \text{ iff } p \in L(X_a).$$

Since $\mathbb{C}$ is assumed to be pure with dimension $|\mathsf{Ag}| - 1$, it holds for every facet $X \in \mathcal{F}(\mathbb{C})$ that $\chi(X) = \mathsf{Ag}$, which makes the expression X_a well-defined. It can be shown that $\mathcal{C} = (\mathbb{C}, L)$ and $\mathcal{M} = (F, V)$, constructed as described, are pointwise equivalent. We refer to the work of Goubault et al. [28] for a detailed proof.

3 Polychromatic Simplicial Models

Chapter Organization. This chapter begins with an introduction to generalized simplicial models in Section 3.1. Section 3.2 presents polychromatic models and defines different notions of belief on them. Next, Section 3.3 analyzes knowledge and belief gain in polychromatic models. An alternative interpretation of polychromatic complexes, inducing a semantics for a somebody-knows modality, is discussed in Section 3.4. Lastly, Section 3.5 concludes the chapter by outlining some possible directions for future work.

The central focus of this thesis is to alter the structural properties of simplicial complexes, and examine the resulting structures epistemically. To motivate this approach, we shift our focus from knowledge to *belief*. Unlike knowledge, belief is an epistemic notion that is non-factive, i.e., an agent can believe ϕ, despite ϕ being false in the actual state. False beliefs arise frequently when reasoning about Byzantine-fault-tolerant systems, i.e., systems in which agents may deviate arbitrarily from the protocol (cf. Cachin et al. [13]). For instance, a malicious agent may deceive an honest one by lying about its own input for a distributed task.

Reasoning about belief on simplicial structures poses a challenge for simplicial semantics (cf. Castañeda et al. [17, Section 4.3]). In relational models, the semantics of belief are the same as for knowledge, i.e., an agent believes ϕ if and only if ϕ is true in all worlds it considers possible. False beliefs can then be addressed by omitting the reflexivity condition on the accessibility relation (cf. Fagin et al. [20]). Under simplicial semantics, however, an agent's local state must be contained in the actual global state. Thus, for any facet X in which an agent a is alive, it is always the case that a considers X possible, i.e., $a \in \chi(X)$. Therefore, if ϕ is true in all global states that a considers possible, then ϕ is true at X. As a result, a belief modality solely based on the inclusion of vertices is inherently factive.

One way of overcoming this problem is to introduce *belief functions* (cf. van Ditmarsch et al. [52]). An agent's belief function f_a maps a facet X, in which a is alive, to another facet Y, in which a is present as well. An agent a believes a formula ϕ in X if and only if it knows ϕ in $f_a(X)$. Hence, belief becomes knowledge if $f_a(X) = X$. Figure 3.1 depicts a simplicial complex with two disconnected facets X and Y, together with a belief function f_a that maps X to Y, i.e., a believes that it is in Y although it is in X. Consequently, a falsely believes that c is not present in the current global state.

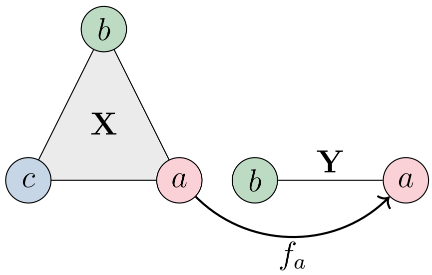

Figure 3.1: Despite being in X, agent a thinks that the actual state is Y.

In this chapter, we present and interpret polychromatic simplicial complexes, i.e., complexes in which adjacent vertices may be of the same color. Such models admit a notion of belief that satisfies the two previously mentioned conditions. The belief studied here is based on the plausibility of states rather than on their possibility alone. That is, an agent believes a formula ϕ if and only if ϕ is true in all states that it deems plausible enough. Since the actual global state need not be among them, the agent's beliefs might be wrong.

We start by defining an agent's plausibility relation between global states, based on the multiplicity of its color within a state. If the color of an agent a has a lower or equal multiplicity in a possible global state X than in another possible state Y, then a considers X to be at least as plausible as Y. If vertices are interpreted as doxastic states, a possible reading of our relation is that an agent a considers worlds with fewer doxastic alternatives more plausible. Since this relation is a wellfounded preorder, we can use the machinery of plausibility models (cf. Baltag and Smets [7, 8]) to define various notions of belief, such as *safe belief* and *most plausible belief*.

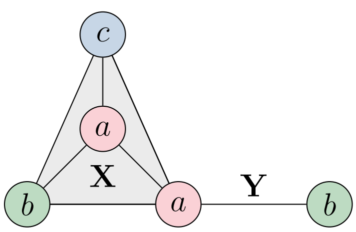

Figure 3.2: A polychromatic simplicial complex that models the same situation as the simplicial complex in Figure 3.1, but without a belief function.

An agent a most plausibly believes ϕ if and only if ϕ is true in the worlds that it considers to be the most plausible ones. This kind of belief is often used in reliable distributed computing, where agents may act based on guarantees that hold true with overwhelming probability, i.e., the states in which those guarantees hold are the most plausible ones. For example, when communicating over authenticated links, it is the most plausible case that if Alice receives a message m from Bob over such a link, then m was actually sent by Bob and not by an impostor. Our simplicial models can represent this kind of belief while taking the topology of the model into account. The complex in Figure 3.2 depicts the same situation as the complex in Figure 3.1, but without a belief function. When in X, a considers Y more plausible than X because its multiplicity in X is 2 and only 1 in Y. In this case, a considers Y to be the most plausible world and falsely believes that c crashed. Hence, polychromatic simplicial complexes are an interesting generalization of chromatic simplicial complexes that allow us to define more intricate accessibility relations. The subsequent sections will explore this structure in detail.

3.1 Simplicial Knowledge

Definition 17 provides a generalization of the representation of a chromatic simplicial complex and was first introduced by Goubault et al. [27]. It augments a simplicial complex $\mathbb{C} = (\mathcal{V}, \mathcal{S})$ with a set of worlds $\mathcal{F}(\mathbb{C}) \subseteq W \subseteq \mathcal{S}$, determining which faces are global states. Consequently, non-empty subsets of facets can now be valid global states as well. This enables modeling scenarios where some agents may have crashed, even when the underlying complex is pure. It is important to note, however, that W is required to contain all facets.

Consider, for example, the augmented chromatic simplicial complex shown in Figure 3.3. Its underlying pure simplicial complex contains only one facet, $X = \{a, b, c\}$, and is augmented with the set $Y = \{a, b\}$. Therefore, the non-facet Y is a global state, and since Y does not contain a c-vertex, agent c is considered dead in Y. As a result, agent a does not know if c is alive because $a \in \chi(X \cap Y)$.

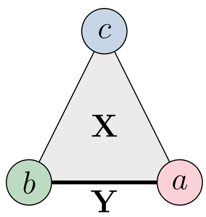

Figure 3.3: An augmented chromatic simplicial complex. The thick ab-edge Y indicates that it is considered a global state.

Definition 17 (Augmented Chromatic Simplicial Complex)**.** *The quadruple* $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ *is an* augmented chromatic simplicial complex *if and only if the following holds:*

1. $\mathbb{S} = (\mathcal{V}, \mathcal{S}, \chi)$ *is a chromatic simplicial complex;*

2. $\mathcal{F}(\mathbb{S}) \subseteq W \subseteq \mathcal{S}$.

The notation used for chromatic simplicial complexes extends naturally to augmented chromatic simplicial complexes. For instance, if $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ is an augmented chromatic simplicial complex, then $\mathcal{F}(\mathbb{C})$ denotes the set of maximal elements of $\mathcal{S}$. Similarly, stating that $\mathbb{C}$ is pure with dimension n means that its underlying complex is pure and of dimension n. From this point onward, we refer to augmented chromatic simplicial complexes as *augmented complexes*.

The notion of a *generalized global simplicial model* in Definition 18 was proposed by Goubault et al. [27]. Our language of distributed knowledge $\mathcal{L}_{\mathcal{D}}$ extends $\mathcal{L}$ by the modal operator $\Box_G$ for each non-empty set of agents G. It is inductively defined as follows:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid \Box_G \phi$$

where $p \in \mathsf{Prop}$ and $\emptyset \neq G \subseteq \mathsf{Ag}$. Moreover, we set $\mathsf{alive}(G) \equiv \Box_G \top$ and $\mathsf{dead}_G = \Box_G \bot$, and explicitly define an epistemic indistinguishability relation in Definition 19 for clarity. Since we adopt the global perspective, we overload notation and use $\Vdash_\sigma$ for the satisfaction relation (Definition 20), as well as the term *σ-validity*.

Definition 18 (Generalized Global Simplicial Model)**.** *A* generalized global simplicial model *is a pair* $\mathcal{C} = (\mathbb{C}, L)$*, where:*

1. $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ *is an augmented complex;*

2. $L : W \to \mathsf{Pow}(\mathsf{Prop})$ *is a labeling.*

Definition 19 ($\sim_G^{\mathbb{C}}$)**.** *Let* $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ *be an augmented complex. For each non-empty set of agents* G*, we define:*

$$\sim_G^{\mathbb{C}} = \{(X, Y) \mid X, Y \in \mathcal{S} \text{ and } G \subseteq \chi(X \cap Y)\}.$$

We write $X \sim_G^{\mathbb{C}} Y$ *if and only if* $(X, Y) \in \sim_G^{\mathbb{C}}$.

If $G = \{a\}$, we write $X \sim_a^{\mathbb{C}} Y$ and $\Box_a$ instead of $X \sim_{\{a\}}^{\mathbb{C}} Y$ and $\Box_{\{a\}}$ respectively. Note that $\Box_G$ is the usual notion of the distributed knowledge of a group of agents G, which is semantically given by the intersection of the individual indistinguishability relations of group members (cf. Halpern and Moses [33]).

Definition 20 ($\Vdash_\sigma$)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be an augmented complex, and let $\mathcal{C} = (\mathbb{C}, L)$ be a generalized global simplicial model. For every face $X \in W$, we define the relation $\mathcal{C}, X \Vdash_\sigma \phi$ by induction on $\phi \in \mathcal{L}_\mathcal{D}$:*

$$
\begin{array}{lll}
\mathcal{C}, X \Vdash_\sigma p & \textit{iff} & p \in L(X) \\
\mathcal{C}, X \Vdash_\sigma \neg\phi & \textit{iff} & \mathcal{C}, X \nVdash_\sigma \phi \\
\mathcal{C}, X \Vdash_\sigma \phi \wedge \psi & \textit{iff} & \mathcal{C}, X \Vdash_\sigma \phi \textit{ and } \mathcal{C}, X \Vdash_\sigma \psi \\
\mathcal{C}, X \Vdash_\sigma \Box_G\phi & \textit{iff} & X \sim_G^{\mathbb{C}} Y \textit{ implies } \mathcal{C}, Y \Vdash_\sigma \phi, \textit{ for all } Y \in W.
\end{array}
$$

Lemma 6 states the standard result of $\sim_G^{\mathbb{C}}$ being a partial equivalence relation (cf. Goubault et al. [27]).

Lemma 6. *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be an augmented complex. For each group of agents G, the relation $\sim_G^{\mathbb{C}}$ is an equivalence relation on $\{X \in W \mid G \subseteq \chi(X)\}$ and empty otherwise.*

3.2 Simplicial Belief

We now drop the assumption that chromatic simplicial complexes are properly colored. A *polychromatic complex* is an augmented complex whose coloring need not be proper. We will define a well-founded preorder on the states of a polychromatic model, which will serve as a plausibility relation (cf. Baltag and Smets [7, 8]). This makes it possible to interpret various notions of belief on simplicial models based on polychromatic complexes.

Lemma 6 does not hold for polychromatic complexes because $\sim_a^{\mathbb{C}}$ need not be transitive. Indeed, let $\mathcal{V} = \{0, 1, 2, 3\}$, and consider the augmented complex $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$, where:

- $\mathcal{S} = \{\{0\}, \{1\}, \{2\}, \{3\}, \{0, 1\}, \{1, 2\}, \{2, 3\}\}$;
- $W = \{X, Y, Z\}$, where $X = \{0, 1\}$, $Y = \{1, 2\}$, and $Z = \{2, 3\}$;
- $\chi(v) = a$ for all $v \in \mathcal{V}$.

We find that $X \sim_a^{\mathbb{C}} Y$ and $Y \sim_a^{\mathbb{C}} Z$, but not $X \sim_a^{\mathbb{C}} Z$. Hence, the relation $\sim_G^{\mathbb{C}}$ is not transitive. To re-establish transitivity of $\sim_G^{\mathbb{C}}$, we must require that for any three worlds $X, Y, Z \in W$ and every non-empty group of agents G:

$$G \subseteq \chi(X \cap Y) \text{ and } G \subseteq \chi(Y \cap Z) \text{ implies } G \subseteq \chi(X \cap Z). \qquad (\star)$$

Polychromatic models (Definition 21) are based on polychromatic complexes satisfying ($\star$). Due to their global nature, we overload notation and use $\Vdash_\sigma$ for the satisfaction relation as well (Definition 23).

Definition 21 (Polychromatic Model)**.** *A* polychromatic model $\mathcal{C} = (\mathbb{C}, L)$ *is a pair where:*

- $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ *be a polychromatic complex satisfying* $(\star)$*;*
- $L : W \to \mathsf{Pow}(\mathsf{Prop})$ *is a labeling.*

Requiring condition $(\star)$ is similar to requiring a transitive accessibility relation in certain relational models.

The multiplicity of a color within a face (Definition 22) induces for each agent a a wellfounded relation $\leq_a$ on worlds. We call this the (a priori) plausibility relation.

Definition 22 (Multiplicity)**.** *Let* $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ *be a polychromatic complex. We define the* multiplicity *of* $a \in \mathsf{Ag}$ *in a face* $X \in \mathcal{S}$ *as:*

$$m_a(X) = |\{v \in X \mid \chi(v) = a\}|.$$

For $X, Y \in W$ *and* $a \in \mathsf{Ag}$*, we write:*

$$X \leq_a Y \quad \textit{iff} \quad m_a(X) \leq m_a(Y).$$

Next, we introduce a local plausibility relation:

$$\trianglelefteq_a \; = \leq_a \cap \sim_a^{\mathbb{C}},$$

which captures the agent's plausibility relation at a given global state. Furthermore, we write

$$X \geq_a Y \quad \text{iff} \quad m_a(X) \geq m_a(Y),$$

and we use $\trianglerighteq_a$ and $\triangleleft_a$ in the obvious way. The following lemma shows that $\sim_G^{\mathbb{C}}$ can be given in terms of the local plausibility relation.

Lemma 7. $\sim_a^{\mathbb{C}} = \trianglelefteq_a \cup \trianglerighteq_a$.

Proof. Observing that $\leq_a$ is strongly connected and unfolding the definition yields

$$\sim_a^{\mathbb{C}} = (\leq_a \cup \geq_a) \cap \sim_a^{\mathbb{C}} = (\leq_a \cap \sim_a^{\mathbb{C}}) \cup (\geq_a \cap \sim_a^{\mathbb{C}}) \; = \; \trianglelefteq_a \cup \trianglerighteq_a. \qquad \square$$

From the relation $\trianglerighteq_a$, we get a corresponding modal operator $[\trianglerighteq]_a$, which is referred to in the literature as *safe belief* (cf. Baltag and Smets [8]). This notion of belief is sometimes also called *defeasible knowledge* (cf. Baltag and Renne [6]). Our language of distributed knowledge and belief $\mathcal{L}_{\mathcal{BD}}$ extends $\mathcal{L}_{\mathcal{D}}$ by the modal operator $[\trianglerighteq]_a$ for each agent $a \in \mathsf{Ag}$. It is inductively defined by:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid \Box_G \phi \mid [\trianglerighteq]_a \phi$$

where $p \in \mathsf{Prop}$, $a \in \mathsf{Ag}$, and $\emptyset \neq G \subseteq \mathsf{Ag}$. As usual, the dual of safe belief is defined as $\langle \trianglerighteq \rangle_a \varphi \equiv \neg [\trianglerighteq]_a \neg \varphi$.

Definition 23 ($\Vdash_\sigma$)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be a polychromatic complex such that $\mathcal{C} = (\mathbb{C}, L)$ is a polychromatic model. For every face $X \in W$, we define the relation $\mathcal{C}, X \Vdash_\sigma \phi$ by induction on $\phi \in \mathcal{L}_{\mathcal{BD}}$:*

$\mathcal{C}, X \Vdash_\sigma p$	*iff*	$p \in L(X)$
$\mathcal{C}, X \Vdash_\sigma \neg\phi$	*iff*	$\mathcal{C}, X \nVdash_\sigma \phi$
$\mathcal{C}, X \Vdash_\sigma \phi \wedge \psi$	*iff*	$\mathcal{C}, X \Vdash_\sigma \phi$ *and* $\mathcal{C}, X \Vdash_\sigma \psi$
$\mathcal{C}, X \Vdash_\sigma \Box_G\phi$	*iff*	$X \sim_G^{\mathbb{C}} Y$ *implies* $\mathcal{C}, Y \Vdash_\sigma \phi$, *for all* $Y \in W$
$\mathcal{C}, X \Vdash_\sigma [\trianglerighteq]_a\phi$	*iff*	$X \trianglerighteq_a Y$ *implies* $\mathcal{C}, Y \Vdash_\sigma \phi$, *for all* $Y \in W$.

The $[\trianglerighteq]_a$-modality satisfies the $\mathsf{S4.2}$ principles for alive agents.

Lemma 8. *The following formulas are valid:*

1. $[\trianglerighteq]_a(\phi \to \psi) \to ([\trianglerighteq]_a\phi \to [\trianglerighteq]_a\psi)$;
2. $\mathsf{alive}(a) \to ([\trianglerighteq]_a\phi \to \phi)$;
3. $[\trianglerighteq]_a\phi \to [\trianglerighteq]_a[\trianglerighteq]_a\phi$;
4. $\langle\trianglerighteq\rangle_a[\trianglerighteq]_a\phi \to [\trianglerighteq]_a\langle\trianglerighteq\rangle_a\phi$.

Proof. Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be a polychromatic complex satisfying ($\star$), and consider the polychromatic model $\mathcal{C} = (\mathbb{C}, L)$. We will only show the last claim. Assume

$$\mathcal{C}, X \Vdash \langle\trianglerighteq\rangle_a[\trianglerighteq]_a\phi.$$

There exists Z with $X \trianglerighteq_a Z$ and

$$\mathcal{C}, Z \Vdash [\trianglerighteq]_a\phi. \tag{3.1}$$

From $X \trianglerighteq_a Z$ we get that a is alive in Z and thus $Z \trianglerighteq_a Z$. Therefore, by (3.1) it follows that:

$$\mathcal{C}, Z \Vdash \phi. \tag{3.2}$$

Let Y be arbitrary with $X \trianglerighteq_a Y$. By ($\star$), we find $Y \sim_a^{\mathbb{C}} Z$. Thus we get by Lemma 7, that $Z \trianglerighteq_a Y$ or $Y \trianglerighteq_a Z$. In the first case, we use (3.1) to obtain $\mathcal{C}, Y \Vdash \phi$. Further we get $Y \trianglerighteq_a Y$ from $X \trianglerighteq_a Y$, and thus

$$\mathcal{C}, Y \Vdash \langle\trianglerighteq\rangle_a\phi. \tag{3.3}$$

In the second case, (3.3) follows immediately from (3.2). Since Y was arbitrary with $X \trianglerighteq_a Y$, (3.3) implies $\mathcal{C}, X \Vdash [\trianglerighteq]_a\langle\trianglerighteq\rangle_a\phi$. □

Plausibility models are not limited to safe belief, and can represent other notions of belief. We start by defining the set of most plausible worlds[1].

Definition 24 ($\mathsf{Min}_{\trianglelefteq_a}$)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be a polychromatic complex. For $X \in W$ and $a \in \mathsf{Ag}$, we define:*

$$\mathsf{Min}_{\trianglelefteq_a}(X) = \{Y \in W \mid Y \sim_a^{\mathbb{C}} X \text{ and } \nexists Z \in W. Z \lhd_a Y\}.$$

Observe that, since $\leq_a$ is wellfounded, we have $\mathsf{Min}_{\trianglelefteq_a}(X) \neq \emptyset$ if agent a is alive in the world X.

The polychromatic complex $\mathbb{C}$, where $W = \mathcal{F}(\mathbb{C})$, given in Figure 3.4 shows a situation with two minimal worlds. Namely, we have $Y, Z \in \mathsf{Min}_{\trianglelefteq_a}(X)$. Further, the following relations hold: $Y \lhd_a X$, $Z \lhd_a X$, $Z \trianglelefteq_a Y$, and $Y \trianglelefteq_a Z$.

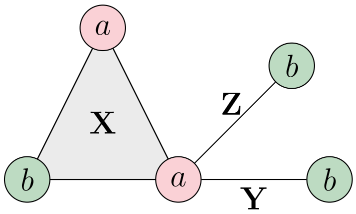

Figure 3.4: A polychromatic complex with two minimal worlds.

The following lemma ensures that the later defined notion of most plausible belief (see Definition 25), behaves as expected (Theorem 1).

Lemma 9. *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be a polychromatic complex satisfying ($\star$). Further, let $X \in W$ and $Y \trianglelefteq_a X$. For each $Z \in \mathsf{Min}_{\trianglelefteq_a}(X)$, we have $Z \trianglelefteq_a Y$.*

Proof. From $Y \trianglelefteq_a X$ we get $X \sim_a^{\mathbb{C}} Y$. Let $Z \in \mathsf{Min}_{\trianglelefteq_a}(X)$. By definition, this means:

$$\nexists V \in W. V \lhd_a Z \tag{3.4}$$

and $X \sim_a^{\mathbb{C}} Z$. Since $\sim_a^{\mathbb{C}}$ is a partial equivalence relation, we get $Z \sim_a^{\mathbb{C}} Y$. Suppose towards a contradiction that $Z \trianglelefteq_a Y$ does not hold. Since $Z \sim_a^{\mathbb{C}} Y$ holds, we must have $Z \not\leq_a Y$ and thus $Y <_a Z$. Therefore, $Y \lhd_a Z$. This contradicts (3.4), and we conclude $Z \trianglelefteq_a Y$. □

If we let Y be X in Lemma 9, then we obtain the following instance:

$$Z \in \mathsf{Min}_{\trianglelefteq_a}(X) \quad \text{implies} \quad Z \trianglelefteq_a X. \tag{3.5}$$

We now include a new modality $\mathcal{B}_a$ for each agent a in our language $\mathcal{L}_{\mathcal{BD}}$. The resulting language is $\mathcal{L}_{\mathcal{BD}}^+$. Definition 25 states the semantics of $\mathcal{B}_a$.

[1]The term "worlds" instead of "faces" emphasizes that we are only referring to faces belonging to W.

Definition 25 ($\mathcal{B}_a$)**.** *Let $\mathcal{C} = (\mathbb{C}, L)$ be a polychromatic model based on the polychromatic complex $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$. For every agent $a \in \mathsf{Ag}$, face $X \in W$, and formula $\phi \in \mathcal{L}^{+}_{\mathcal{BD}}$, we define:*

$$\mathcal{C}, X \Vdash_\sigma \mathcal{B}_a\phi \text{ iff } \quad Y \in \mathsf{Min}_{\trianglelefteq_a}(X) \text{ implies } \mathcal{C}, Y \Vdash_\sigma \phi, \text{ for all } Y \in W.$$

The modality $\mathcal{B}_a$ models agent a's most plausible belief. It is well-known that $\mathcal{B}_a$ can be expressed in terms of the $[\trianglerighteq]_a$-modality (cf. Baltag and Smets [8] and Stalnaker [48]). Theorem 1 states that the same is true for polychromatic models.

Theorem 1. *Let $\mathcal{C} = (\mathbb{C}, L)$ be a polychromatic model based on the polychromatic complex $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$. For all $a \in \mathsf{Ag}$ and $X \in W$ with $a \in \chi(X)$, we find that:*

$$\mathcal{C}, X \Vdash_\sigma \mathcal{B}_a\varphi \quad \text{if and only if} \quad \mathcal{C}, X \Vdash_\sigma \langle\trianglerighteq\rangle_a[\trianglerighteq]_a\varphi.$$

Proof. For the direction from right to left, we assume $\mathcal{C}, X \Vdash_\sigma \langle\trianglerighteq\rangle_a[\trianglerighteq]_a\varphi$. Thus, there exists Y with $X \trianglerighteq_a Y$ and $\mathcal{C}, Y \Vdash_\sigma [\trianglerighteq]_a\varphi$. Now consider an arbitrary $Z \in \mathsf{Min}_{\trianglelefteq_a}(X)$. By Lemma 9 we find $Y \trianglerighteq_a Z$ and, therefore, $\mathcal{C}, Z \Vdash_\sigma \varphi$. This yields $\mathcal{C}, X \Vdash_\sigma \mathcal{B}_a\varphi$.

For the direction from left to right, we have $X \sim_a^{\mathbb{C}} X$ since agent a is alive. Thus $\mathsf{Min}_{\trianglelefteq_a}(X)$ is non-empty and we let $Y \in \mathsf{Min}_{\trianglelefteq_a}(X)$. By (3.5), we obtain $Y \trianglelefteq_a X$. It remains to show $\mathcal{C}, Y \Vdash_\sigma [\trianglerighteq]_a\varphi$. Let $Z \in W$ be such that $Z \trianglelefteq_a Y$. Then $Z \sim_a^{\mathbb{C}} X$ by transitivity of $\sim_a^{\mathbb{C}}$. Now we find $Z \in \mathsf{Min}_{\trianglelefteq_a}(X)$, for otherwise, we would find $V \in W$ with $V \lhd_a Z$, which yields $V \lhd_a Y$ and thus contradicts $Y \in \mathsf{Min}_{\trianglelefteq_a}(X)$. From $Z \in \mathsf{Min}_{\trianglelefteq_a}(X)$ and the assumption $\mathcal{C}, X \Vdash_\sigma \mathcal{B}_a\varphi$, we get $\mathcal{C}, Z \Vdash_\sigma \varphi$, which concludes the proof. □

Remark 4. *A consequence of Theorem 1 is that the properties of the $\mathcal{B}_a$-modality follow from properties of $[\trianglerighteq]_a$ such as the ones given in Lemma 8. For instance, we find that the following is σ-valid:*

$$\mathcal{B}_a\phi \wedge \mathcal{B}_a\psi \to \mathcal{B}_a(\phi \wedge \psi).$$

Our model satisfies the knowledge yields belief principle. In particular, we have the following lemma.

Lemma 10. *Let $\mathcal{C} = (\mathbb{C}, L)$ be a polychromatic model based on the polychromatic complex $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$. For all $X \in W$, we have:*

$$\mathcal{C}, X \Vdash_\sigma \Box_a\varphi \to [\trianglerighteq]_a\varphi \quad \text{and} \quad \mathcal{C}, X \Vdash_\sigma [\trianglerighteq]_a\varphi \to \mathcal{B}_a\varphi.$$

Proof. For the first claim, assume

$$\mathcal{C}, X \Vdash_\sigma \Box_a \varphi. \tag{3.6}$$

Let $X \trianglerighteq_a Y$, i.e., $X \sim_a^{\mathbb{C}} Y$ and $X \geq_a Y$. By (3.6) we immediately get $\mathcal{C}, Y \Vdash_\sigma \varphi$ and hence, $\mathcal{C}, X \Vdash_\sigma [\trianglerighteq]_a \varphi$.

For the second claim, assume

$$\mathcal{C}, X \Vdash_\sigma [\trianglerighteq]_a \varphi. \tag{3.7}$$

Let $Y \in \mathsf{Min}_{\trianglelefteq_a}(X)$ be arbitrary. Using (3.5), we obtain $Y \trianglelefteq_a X$. Now $\mathcal{C}, Y \Vdash_\sigma \varphi$ follows immediately from (3.7), which yields $\mathcal{C}, X \Vdash_\sigma \mathcal{B}_a \varphi$. ☐

Example 4 illustrates that, as typical for preference-based semantics, our models are non-monotone. That is, agents may drop their beliefs when learning new information.

Example 4. *Consider the set $\mathcal{V} = \{1, 2, 3, 4\}$ and a coloring χ such that:*

1. $\chi(1) = \chi(3) = a$;
2. $\chi(2) = \chi(4) = b$.

Moreover, let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ and $\mathbb{C}^\psi = (\mathcal{V}, \mathcal{S}^\psi, \chi, W^\psi)$ be two polychromatic complexes given by:

$$\mathcal{S} = (\mathsf{Pow}(\{1, 2, 3\}) \cup \mathsf{Pow}(\{3, 4\})) \setminus \{\emptyset\} \text{ and } \mathcal{S}^\psi = \mathsf{Pow}(\{1, 2, 3\}) \setminus \{\emptyset\}.$$

Figure 3.5 shows the polychromatic complexes. Further, let $X = \{1, 2, 3\}$ and $Y = \{3, 4\}$. We define $W = \{X, Y\}$, and $W^\psi = \{X\}$. Consider the two polychromatic models:

$$\mathcal{C} = (\mathbb{C}, L\}) \text{ and } \mathcal{C}^\psi = (\mathbb{C}^\psi, L^\psi).$$

We choose the labeling L such that for some propositional formulas $\psi, \phi \in \mathcal{L}^+_{\mathcal{BD}}$:

$$\mathcal{C}, X \Vdash_\sigma \neg\phi \wedge \psi \quad \text{and} \quad \mathcal{C}, Y \Vdash_\sigma \phi \wedge \neg\psi.$$

We set $L^\psi(X) = L(X)$. The model $\mathcal{C}^\psi$ represents the situation after the agents in $\mathcal{C}$ learn that ψ is true. That is, it is the same as $\mathcal{C}$ but without the worlds where ψ is false. We observe that

$$\mathcal{C}, X \Vdash_\sigma \mathcal{B}_a \phi \quad \text{and} \quad \mathcal{C}^\psi, X \nVdash_\sigma \mathcal{B}_a \phi.$$

Hence, a only believes ϕ in X before it learns ψ. This is because removing worlds from $\mathcal{C}$ can result in a new world becoming a most plausible world.

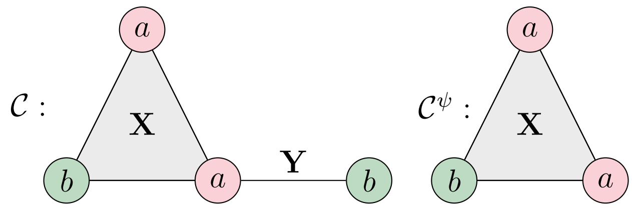

Figure 3.5: The polychromatic complex $\mathcal{C}^{\psi}$ represents the state of affairs after the agents in $\mathcal{C}$ learn that ψ is true.

In the context of properly colored augmented complexes, properness as in Definition 11 is not satisfied. An example of such an augmented complex is the one shown in Figure 3.3, where no alive agent belonging to both X and Y can distinguish them. Consequently, the definition of properness has to be adapted to the generalized setting: there exists an agent that can distinguish between two worlds if and only if those worlds contain the same set of alive agents. Under this condition, augmented complexes satisfy properness. Goubault et al. [27] show that global relational models based on non-empty partial epistemic frames satisfying the following axiom, can be transformed to meet the adapted definition of properness:

$$\mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \wedge \varphi \rightarrow \Box_G(\mathsf{dead}(G^C) \rightarrow \varphi), \tag{P}$$

where G^C denotes the complement of G. Intuitively, P states that if two global states have the same set of alive agents and no alive agent can distinguish between them, then both states must satisfy the same formulas. Generalized global simplicial models naturally satisfy this property. However, Lemma 11 shows that P is no longer true for polychromatic models.

Lemma 11. P *is not σ-valid for polychromatic models.*

Proof. Consider the following counter-example: let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be the polychromatic complex given by:

$$\mathcal{V} = \{1, 2, 3\} \quad \text{and} \quad \mathcal{S} = \mathsf{Pow}(\mathcal{V}) \setminus \{\emptyset\},$$

where W contains only the two worlds $X = \{1, 2, 3\}$ and $Y = \{1, 3\}$. Further, let $\mathsf{Ag} = \{a, b\}$, and assume that the vertices are colored as follows:

$$\chi(1) = \chi(2) = a \quad \text{and} \quad \chi(3) = b.$$

Consider a labeling such that $L(X) = \{p\}$, and $L(Y) = \emptyset$. Further, let $\mathcal{C} = (\mathbb{C}, L)$. For $G = \{a, b\}$. We find that

$$\mathcal{C}, X \Vdash_{\sigma} \mathsf{alive}(G) \wedge \mathsf{dead}(G^c) \wedge p.$$

However, we have $X \sim_G^{\mathbb{C}} Y$ and $\mathcal{C}, Y \Vdash_\sigma \mathsf{dead}(G^c) \wedge \neg p$. Hence

$$\mathcal{C}, X \nVdash_\sigma \Box_G(\mathsf{dead}(G^c) \to p).$$

Therefore, P is not valid on polychromatic models. $\square$

3.3 Knowledge Gain

An important result for simplicial models is that agents cannot gain new knowledge along morphisms. This property is essential for showing that certain distributed tasks are not solvable (cf. Goubault et al. [29]). We adapt the notion of a morphism between simplicial complexes to the setting of polychromatic models. The fact that our models are polychromatic does not matter for the definition of morphism. Given a function $f : U \to V$ and a set $W \subseteq U$, we let:

$$f(W) = \{f(x) \mid x \in W\}.$$

Definition 26 (Simplicial Map). *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S})$ and $\mathbb{C}' = (\mathcal{V}', \mathcal{S}')$ be two simplicial complexes. A* simplicial map *from $\mathbb{C}$ to $\mathbb{C}'$ is a function $f : \mathcal{V} \to \mathcal{V}'$ such that if $X \in \mathcal{S}$ then $f(X) \in \mathcal{S}'$.*

Definition 27 (Morphism). *Let $(\mathcal{V}, \mathcal{S})$ and $(\mathcal{V}', \mathcal{S}')$ be two simplicial complexes, and let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ and $\mathbb{C}' = (\mathcal{V}', \mathcal{S}', \chi', W')$ be two polychromatic complexes satisfying ($\star$). A* morphism *from a polychromatic model $\mathcal{C} = (\mathbb{C}, L)$ to a polychromatic model $\mathcal{C}' = (\mathbb{C}', L')$ is a function f such that:*

1. *f is a simplicial map from $(\mathcal{V}, \mathcal{S})$ to $(\mathcal{V}', \mathcal{S}')$;*
2. *$\chi'(f(v)) = \chi(v)$ for all $v \in \mathcal{V}$;*
3. *$f(X) \in W'$ for all $X \in W$;*
4. *$L'(f(X)) = L(X)$ for all $X \in W$.*

Morphisms respect the indistinguishability relation. We have the following lemma.

Lemma 12. *Let $\mathcal{C} = (\mathbb{C}, L)$ and $\mathcal{C}' = (\mathbb{C}', L')$ be polychromatic models based on the polychromatic complexes $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ and $\mathbb{C}' = (\mathcal{V}', \mathcal{S}', \chi', W')$. If f is a morphism from $\mathcal{C}$ to $\mathcal{C}'$, we find for all $X, Y \in W$:*

$$X \sim_G^{\mathbb{C}} Y \quad \textit{implies} \quad f(X) \sim_G^{\mathbb{C}'} f(Y).$$

Proof. Assume $G \subseteq \chi(X \cap Y)$ and let a be an element of G. There exists $v \in \mathcal{V}$ with $v \in X$, $v \in Y$, and $\chi(v) = a$. We find that $f(v) \in f(X)$, $f(v) \in f(Y)$, and $\chi'(f(v)) = a$. Hence, $a \in \chi'(f(X) \cap f(Y))$. □

The positive formulas are the formulas of $\mathcal{L}_\mathcal{D}$ where the operator $\Box_G$ occurs only in its unnegated form. Formally, we use the following definition.

Definition 28 (Positive Formulas)**.** *We consider the following grammar:*

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi$$
$$\psi ::= \phi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \Box_G \psi$$

where $p \in \mathsf{Prop}$ *and* $a \in \mathsf{Ag}$*. Formulas given by* ψ *are called* positive formulas*.*

The result about no knowledge gain is standard (cf. Goubault et al. [29]). Note that we are in a setting where agents may crash. However, since we adopt the global point of view, we can employ the usual formulation of positive formulas in the following theorem. Also, the fact that we have polychromatic models does not matter.

Theorem 2. *Let* $\mathcal{C} = (\mathbb{C}, L)$ *and* $\mathcal{C}' = (\mathbb{C}', L')$ *be polychromatic models based on the polychromatic complexes* $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ *and* $\mathbb{C}' = (\mathcal{V}', \mathcal{S}', \chi', W')$*. Further, let* $X \in W$ *and* ψ *be a positive formula. If* f *is a morphism from* $\mathcal{C}$ *to* $\mathcal{C}'$*, we find that:*

$$\mathcal{C}', f(X) \Vdash_\sigma \psi \quad \textit{implies} \quad \mathcal{C}, X \Vdash_\sigma \psi,$$

for all $X \in W$*.*

Proof. First we show that for a formula given by ϕ according to Definition 28, we have:

$$\mathcal{C}', f(X) \Vdash_\sigma \phi \quad \text{iff} \quad \mathcal{C}, X \Vdash_\sigma \phi, \tag{3.8}$$

for all $X \in W$. We proceed by induction on the structure of ϕ and distinguish:

1. ϕ is an atomic proposition p. Since f is a morphism, we have that $L'(f(X)) = L(X)$. Thus $\mathcal{C}, f(X) \Vdash_\sigma p$ iff $\mathcal{C}, X \Vdash_\sigma p$.

2. ϕ is a negation or a conjunction. The claim follows immediately by I.H.

Now we proceed by induction on the structure of ψ (according to Definition 28) and assume $\mathcal{C}', f(X) \Vdash_\sigma \psi$. We distinguish the following cases:

1. ψ is a formula ϕ (according to Definition 28). The claim follows from (3.8).

2. ψ is a conjunction or a disjunction. The claim follows immediately by the induction hypothesis.

3. ψ is of the form $\Box_G\psi'$. Since $\mathcal{C}', f(X) \Vdash_\sigma \Box_G\psi'$ holds, we obtain that $f(X) \sim_G^{\mathbb{C}'} Y$ implies $\mathcal{C}', Y \Vdash_\sigma \psi'$ for all $Y \in W'$.

   Let $Z \in W$ be such that $X \sim_G^{\mathbb{C}} Z$. Since f is a morphism, we find by Lemma 12 that
   $$f(X) \sim_G^{\mathbb{C}'} f(Z) \quad \text{and} \quad f(Z) \in W'.$$
   Thus, $\mathcal{C}', f(Z) \Vdash_\sigma \psi'$. By I.H., we find $\mathcal{C}, Z \Vdash_\sigma \psi'$. Since $Z \in W$ was arbitrary with $X \sim_G^{\mathcal{C}} Z$, we conclude $\mathcal{C}, X \Vdash_\sigma \Box_G\psi'$. □

The previous theorem only holds for knowledge but not for belief. That is, the operator $[\trianglerighteq]_a$ cannot be included in the class of positive formulas. We have the following lemma.

Lemma 13. *There exist polychromatic models $\mathcal{C} = (\mathbb{C}, L)$ and $\mathcal{C}' = (\mathbb{C}', L')$ based on $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ and $\mathbb{C}' = (\mathcal{V}', \mathcal{S}', \chi', W')$, together with a morphism f from $\mathcal{C}$ to $\mathcal{C}'$, and $X \in W$ such that for some agent $a \in \mathsf{Ag}$ and $p \in \mathsf{Prop}$:*

$$\mathcal{C}', f(X) \Vdash_\sigma [\trianglerighteq]_a p \quad \textit{but} \quad \mathcal{C}, X \nVdash_\sigma [\trianglerighteq]_a p.$$

Proof. Consider the set $\mathcal{V} = \{1, 2, 3\}$. We let a be an agent and set

$$\chi(1) = \chi(2) = \chi(3) = a.$$

We let $\mathcal{S} = \{\{1, 2\}, \{2, 3\}, \{1\}, \{2\}, \{3\}\}$ and $\mathcal{S}' = \{\{1, 2\}, \{1\}, \{2\}\}$. Further, we set $X = \{1, 2\}$, $Y = \{2, 3\}$, and $Z = \{2\}$, as well as:

$$W = \{X, Y\}, \qquad L(X) = \emptyset \quad \text{and} \quad L(Y) = \{p\}$$

and

$$W' = \{X, Z\}, \quad L'(X) = \emptyset \quad \text{and} \quad L'(Z) = \{p\}.$$

It is straightforward to that $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ and $\mathbb{C}' = (\mathcal{V}, \mathcal{S}', \chi, W')$ are polychromatic complexes such that

$$\mathcal{C} = (\mathbb{C}, L) \quad \text{and} \quad \mathcal{C}' = (\mathbb{C}', L')$$

are polychromatic models. Moreover, let f be such that $f(1) = 1$, $f(2) = 2$, and $f(3) = 2$. Obviously, f is a morphism from $\mathcal{C}$ to $\mathcal{C}'$ and we have that $f(Y) = Z$. Finally, we obtain that:

$$\mathcal{C}', f(Y) \Vdash_\sigma [\trianglerighteq]_a p \quad \text{but} \quad \mathcal{C}, Y \nVdash_\sigma [\trianglerighteq]_a p.$$ □

Besides the statement that belief gain is possible, this lemma could also be interpreted in such a way that condition ($\star$) is not strong enough or that we need a different notion of morphism.

3.4 Somebody Knows

We now change the interpretation of the set Ag. Instead of consisting of *agents themselves*, we assume that Ag is the set of *agent names*. This interpretation is rather natural in the presence of an improper coloring, because two distinct agents can share a name. Moreover, a name can represent the type of a group. For example, in distributed computing, we might be interested in the two names f and c, where f stands for "faulty" and c stands for "correct". Logics without a fixed one-to-one correspondence between names and agents have been extensively studied by Grove and Halpern [31] and Bílková et al. [10]. In what follows, we will demonstrate how polychromatic simplicial complexes provide a semantics for a *somebody-knows* modality. We enrich $\mathcal{L}$ with the modality S_a, which reads as *somebody named a knows*. This modality is similar to the *somebody in-group-G-knows* modality S_G studied by Ågotnes and Wáng [2], where each agent has a unique identifier and $G \subseteq \mathsf{Ag}$. Unlike S_a, the modality S_G can be defined in terms of the individual knowledge of the group members, i.e.,

$$\mathsf{S}_G\phi \equiv \bigvee_{a\in G} \Box_a\phi,$$

because the identity of each group member is contained in the group G. Moreover, the modality S_G satisfies $\mathsf{S}_G\phi \to \mathsf{S}_H\phi$ for $G \subseteq H$, which does not have a meaningful formulation with names since our language cannot express that a specific agent is named a.

Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be a polychromatic simplicial complex. We define the *star* of a vertex $v \in \mathcal{V}$ as:

$$\mathsf{st}(v) = \{X \in W \mid v \in X\}.$$

The *set of a-stars* of a face $X \in W$ is given by:

$$\mathcal{N}_a(X) = \{\mathsf{st}(v) \mid v \in X \text{ and } \chi(v) = a\}.$$

Since $\mathcal{F}(\mathbb{C}) \subseteq W$, the set $\mathcal{N}_a(X)$ is never empty if there is an alive agent with name a in face X. We denote our language with $\mathcal{L}_\mathsf{S}$. A *simplicial named model* is a pair $\mathcal{C} = (\mathbb{C}, L)$, where $\mathbb{C}$ is a polychromatic complex and L is a global labeling. Notice that we no longer require $(\star)$. The semantics of S_a are:

$$\mathcal{C}, X \Vdash_\sigma \mathsf{S}_a\phi \text{ iff } \exists Y \in \mathcal{N}_a(X).\forall Z \in Y.\mathcal{C}, Z \Vdash \phi. \tag{3.9}$$

Thus, if $\mathcal{C}, X \Vdash_\sigma \neg\mathsf{S}_a\phi$, then all agents in X with name a consider at least one world possible in which $\neg\phi$ holds.

We can formulate (3.9) more compact with the help of *truth sets*. For a formula $\phi \in \mathcal{L}_\mathsf{S}$ and a simplicial named model $\mathcal{C} = (\mathbb{C}, L)$, we denote the *truth-set* of ϕ as

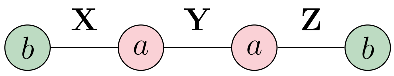

Figure 3.6: The polychromatic complex considered in the proof of Lemma 14.

$\phi^{\mathcal{C}} = \{X \in W \mid \mathcal{C}, X \Vdash_\sigma \phi\}$. Therefore, (3.9) becomes:

$$\mathcal{C}, X \Vdash_\sigma \mathsf{S}_a\phi \text{ iff } \exists N \in \mathcal{N}_a(X).N \subseteq \phi^{\mathcal{C}}.$$

The dual $\langle \mathsf{S}_a \rangle$ reads as: *everyone named a considers ϕ possible*. Semantically, it is given by:

$$\mathcal{C}, X \Vdash_\sigma \langle \mathsf{S}_a \rangle\phi \text{ iff } \forall N \in \mathcal{N}_a(X).N \not\subseteq (\neg\phi)^{\mathcal{C}}.$$

Similar to neighborhood semantics (cf. Pacuit [43]), Lemma 14 shows that the formula $\mathsf{S}_a(\phi \to \psi) \to (\mathsf{S}_a\phi \to \mathsf{S}_a\psi)$ is not σ-valid.

Lemma 14. *The formula* $\mathsf{S}_a(\phi \to \psi) \to (\mathsf{S}_a\phi \to \mathsf{S}_a\psi)$ *is not σ-valid.*

Proof. We will construct a counterexample. Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be the polychromatic complex shown in Figure 3.6 with $W = \mathcal{F}(\mathbb{C})$. The set of Y's a-stars is: $\mathcal{N}_a(Y) = \{\{X, Y\}, \{Y, Z\}\}$.

We define a simplicial named model $\mathcal{C} = (\mathbb{C}, L)$, where L is a labeling such that for some $p, q \in \mathsf{Prop}$, p is only satisfied at X, Y, and q is only satisfied at Y. By the definition of truth we have:

$\mathcal{C}, Y \Vdash_\sigma \mathsf{S}_a(p \to q)$ because $\{Y, Z\} \subseteq (p \to q)^{\mathcal{C}}$ and $\{Y, Z\} \in \mathcal{N}_a(Y)$;

$\mathcal{C}, Y \Vdash_\sigma \mathsf{S}_a p$ because $\{X, Y\} \subseteq \phi^{\mathcal{C}}$ and $\{X, Y\} \in \mathcal{N}_a(Y)$;

$\mathcal{C}, Y \nVdash_\sigma \mathsf{S}_a q$ because $q^{\mathcal{C}} = \{Y\}$ and for all $N \in \mathcal{N}_a(Y)$, we have $N \not\subseteq \{Y\}$.

Therefore, $\mathsf{S}_a(\phi \to \psi) \to (\mathsf{S}_a\phi \to \mathsf{S}_a\psi)$ is not σ-valid. □

Lemma 15 states that S_a is factive. Notice that it does not matter whether the complex is impure because $\mathcal{C}, X \Vdash_\sigma \mathsf{S}_a\phi$ implies the existence of an agent named a in X knowing ϕ. In contrast, knowledge is only truthful for alive agents in models based on impure simplicial complexes.

Lemma 15. *The formula* $\mathsf{S}_a\phi \to \phi$ *is σ-valid.*

Proof. Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be a polychromatic complex and consider an arbitrary simplicial named model $\mathcal{C} = (\mathbb{C}, L)$. Furthermore, assume that $\mathcal{C}, X \Vdash_\sigma \mathsf{S}_a\phi$. By definition, there exists an a-vertex $v \in X$ with $\mathsf{st}(v) \subseteq \phi^{\mathcal{C}}$. Moreover, by definition $X \in \mathsf{st}(v)$, and thus $X \in \phi^{\mathcal{C}}$, i.e., $\mathcal{C}, X \Vdash_\sigma \phi$. □

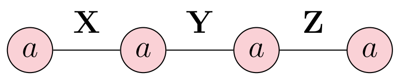

Figure 3.7: The polychromatic complex considered in the proof of Lemma 17.

Remark 5. *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be a polychromatic complex. For every two faces $X, Y \in W$ such that $v \in X \cap Y$ with $\chi(v) = a$, we find:*

$$\mathsf{st}(v) \in \mathcal{N}_a(X) \cap \mathcal{N}_a(Y).$$

Lemma 16 states that S_a satisfies positive introspection.

Lemma 16. *The formula $\mathsf{S}_a\phi \to \mathsf{S}_a\mathsf{S}_a\phi$ is σ-valid.*

Proof. Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be a polychromatic complex, and consider a simplicial named model $\mathcal{C} = (\mathbb{C}, L)$. Assume that $\mathcal{C}, X \Vdash \mathsf{S}_a\phi$. By the definition of truth, there exists $\mathsf{st}(v) \in \mathcal{N}_a(X)$ with $v \in X$ and $\chi(v) = a$ such that $\mathsf{st}(v) \subseteq \phi^{\mathcal{C}}$. Let $Y \in \mathsf{st}(v)$. By Remark 5 we have $\mathsf{st}(v) \in \mathcal{N}_a(Y)$. Since $\mathsf{st}(v) \subseteq \phi^{\mathcal{C}}$, it follows that $\mathcal{C}, Y \Vdash \mathsf{S}_a\phi$. Moreover, because Y was arbitrary, we conclude that for all $Z \in \mathsf{st}(v)$, we have that $\mathcal{C}, Z \Vdash \mathsf{S}_a\phi$. Lastly, since $\mathsf{st}(v) \in \mathcal{N}_a(X)$ and $\mathsf{st}(v) \subseteq (\mathsf{S}_a\phi)^{\mathcal{C}}$, it follows that $\mathcal{C}, X \Vdash \mathsf{S}_a\mathsf{S}_a\phi$. □

The next lemma establishes that the standard axioms $\mathsf{B} : \phi \to \mathsf{S}_a\langle\mathsf{S}_a\rangle\phi$ and $\mathsf{5} : \langle\mathsf{S}_a\rangle\phi \to \mathsf{S}_a\langle\mathsf{S}_a\rangle\phi$ are not valid.

Lemma 17. *The formulas $\phi \to \mathsf{S}_a\langle\mathsf{S}_a\rangle\phi$ and $\langle\mathsf{S}_a\rangle\phi \to \mathsf{S}_a\langle\mathsf{S}_a\rangle\phi$ are both not σ-valid.*

Proof. We construct a simplicial named model that violates both formulas. Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, W)$ be the polychromatic complex shown in Figure 3.7, with $W = \mathcal{F}(\mathbb{C})$. The sets of a-stars are:

- $\mathcal{N}_a(X) = \{\{X\}, \{X, Y\}\}$;
- $\mathcal{N}_a(Y) = \{\{X, Y\}, \{Y, Z\}\}$;
- $\mathcal{N}_a(Z) = \{\{Z\}, \{Y, Z\}\}$.

We define a model $\mathcal{C} = (\mathbb{C}, L)$ such that for some $p \in \mathsf{Prop}$, it holds that $p^{\mathcal{C}} = \{Y\}$. The following is straightforward to verify:

1. $\mathcal{C}, Y \Vdash_\sigma p$;
2. $\mathcal{C}, Y \Vdash_\sigma \neg\mathsf{S}_a\neg p$ because $(\neg p)^{\mathcal{C}} = \{X, Z\}$, and there is no $N \in \mathcal{N}_a(Y)$ such that $N \subseteq \{X, Z\}$;

3. Similarly, it follows that $\mathcal{C}, X \Vdash_\sigma \mathsf{S}_a \neg p$ and $\mathcal{C}, Z \Vdash_\sigma \mathsf{S}_a \neg p$.

Since for all $N \in \mathcal{N}_a(Y)$, it holds that $(\neg \mathsf{S}_a \neg p)^{\mathcal{C}} = \{Y\} \not\supseteq N$, we find that $\mathcal{C}, Y \Vdash_\sigma \neg \mathsf{S}_a \neg \mathsf{S}_a \neg p$, which is equivalent to $\mathcal{C}, Y \Vdash_\sigma \neg \mathsf{S}_a \langle \mathsf{S}_a \rangle p$. Thus, both formulas are violated. □

3.4.1 Modeling Quorums

A proper coloring allows us to express to which agent a local state belongs, while an improper coloring specifies the name (or type) of the agent inhabiting that local state. Combining both, i.e., assigning colors **and** names to vertices, yields more expressive models. Let Ag be a set of $n \geq 1$ agents, and let N be a set of $1 \leq k \leq n$ names. Moreover, we assume that $\mathsf{Ag} \cap \mathsf{N} = \emptyset$. In this setting, we can express that an agent a has name n. A *mixed complex* is a quadruple $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \nu)$ where:

- $(\mathcal{V}, \mathcal{S})$ is a simplicial complex;
- $\chi : \mathcal{V} \to \mathsf{Ag}$ is a proper coloring;
- $\nu : \mathcal{V} \to \mathsf{N}$ is an improper coloring.

We refrain from augmenting a mixed complex with a set W for simplicity. The star of a vertex is now defined over $\mathcal{F}(\mathbb{C})$ instead of W. A *mixed model* is a pair $\mathcal{C} = (\mathbb{C}, L)$, where $\mathbb{C}$ is a mixed complex, and L is a labeling $L : \mathcal{F}(\mathbb{C}) \to \mathsf{Pow}(\mathsf{Prop})$. Finally, the n-stars of $X \in \mathcal{F}(\mathbb{C})$ are defined in terms of ν instead of χ:

$$\mathcal{N}_n(X) = \{\mathsf{st}(v) \mid v \in X \text{ and } \nu(v) = n\}.$$

The satisfaction relation $\Vdash_\sigma$ can be defined on the facets of $\mathbb{C}$ as expected. Specifically, we define:

$$\begin{array}{ll} \mathcal{C}, X \Vdash_\sigma \Box_G \phi & \text{iff } X \sim_G^{\mathbb{C}} Y \quad \text{implies} \quad \mathcal{C}, Y \Vdash_\sigma \phi, \text{ for all } Y \in W \\ \mathcal{C}, X \Vdash_\sigma \mathsf{S}_n \phi & \text{iff } \exists N \in \mathcal{N}_n(X). N \subseteq \phi^{\mathcal{C}}. \end{array}$$

Having two modalities relying on different colorings allows us to formulate more elaborate connections between groups of agents. Consider, for example, a mixed model $\mathcal{C} = (\mathbb{C}, L)$ with a facet $X \in \mathcal{F}(\mathbb{C})$ containing an a-vertex named n. If $G \subseteq \mathsf{Ag}$ is a group of agents with $a \in G$, the following holds:

$$\mathcal{C}, X \Vdash_\sigma \bigwedge_{a \in G} \Box_a \phi \to \mathsf{S}_n \phi.$$

There are further interesting variants of the above formula, and a lot of those statements arise frequently in distributed systems or game theory. We will look at a specific scenario in distributed systems: *quorums*.

Assume that we are given N processes (or agents). Some processes are *faulty* and may deviate arbitrarily from the protocol. Processes that are not faulty are called *correct*. Given an assumption on which sets of processes may jointly fail, a *(Byzantine) quorum system* (cf. Malkhi and Reiter [40]) is a set of sets of processes (called quorums) such that they pairwise intersect in at least one correct process. Quorums are essential for the reliable transfer of information in networks, because a process cannot receive conflicting information from two quorums. Indeed, consider for example a decentralized payment system such as FastPay (cf. Baudet et al. [9]), in which clients submit their transactions to servers that maintain records of all client balances. Upon receiving a transaction, the server needs to verify whether the client's balance is sufficient. Once verified, the server signs the transaction and informs the other servers about its validity. If a server receives such a message from a quorum of servers, then it can update its state. Two transactions with the same issuer conflict if the sum of the transfers exceeds its balance. Since we use quorum systems, there cannot exist quorums for two conflicting transactions. Thus, *double-spending* is not possible. Towards a contradiction, assume that it is possible. Hence, there exist quorums for both transactions. Since those quorums intersect in a correct server, this means that there exists a correct server that deemed two conflicting transactions valid, which is a contradiction. We now formulate this idea in mixed models.

Let $\mathsf{N} = \{f, c\}$ be the set of names, where f stands for "faulty" and c stands for "correct". Let Ag be the set of processes and let $\mathbb{C}$ be a pure mixed complex. The complex $\mathbb{C}$ is pure for simplicity. Consider the two groups of processes G and H such that both groups are elements of a quorum system $\mathcal{Q}$. Consequently, they intersect in a process with name c. Thus, for every mixed model $\mathcal{C} = (\mathbb{C}, L)$ and $X \in \mathcal{F}(\mathbb{C})$, it holds that

$$\mathcal{C}, X \Vdash_\sigma \left(\bigwedge_{a \in G} \Box_a \phi \wedge \bigwedge_{a \in H} \Box_a \psi \right) \to \mathsf{S}_c(\phi \wedge \psi),$$

which captures quorum systems in an intuitive manner. Moreover, we can generalize this to arbitrary names and write $\mathcal{Q}_n$ for a quorum system such that any two elements of $\mathcal{Q}_n$ intersect in a process with name n. The following would be an axiom of our logic:

$$\left(\bigwedge_{a \in G} \Box_a \phi \wedge \bigwedge_{a \in H} \Box_a \psi \right) \to \mathsf{S}_n(\phi \wedge \psi), \text{ if } G, H \in \mathcal{Q}_n.$$

Quorum systems can be further generalized to the setting in which each agent has its own quorum system. This setting is called *asymmetric trust* because every process can have its own trust assumption (cf. Alpos et al. [4]).

3.5 Conclusion and Outlook

Polychromatic models naturally support logics of (most) plausible belief, because the multiplicity of a color within a face can be interpreted as an inverse plausibility measure, giving rise to a plausibility relation $\leq_a$ on worlds. Restricting $\leq_a$ to adjacent worlds yields a local plausibility relation $\trianglerighteq_a$, which in turn induces a modality $[\trianglerighteq]_a$. If further restricted to worlds with minimal multiplicity, we obtain the most plausible belief modality $\mathcal{B}_a$. Since a most plausible world is a possible world, it immediately follows that most plausible belief satisfies KYB (Lemma 10). At the same time, since the set of most plausible worlds need not include the actual world, an agent's most plausible beliefs may be false. Thus, our notion of belief depends entirely on the topological structure of the complex, while satisfying KYB without requiring additional machinery like belief functions. In addition to providing first definitions, we also established key differences from simplicial complexes such as the observations about non-proper models and belief gain. Finally, we demonstrate how polychromatic complexes can serve as models supporting a somebody-knows modality. This observation is of particular interest, since these models do not require ($\star$).

There are various ways for extending polychromatic complexes. In general, analyzing models in which ($\star$) does not hold, such as mixed models, and exploring which epistemic attitudes they model is certainly worthwhile. In the context of mixed models, a further study of a simplicial variant of coalition logic (cf. Pauli [44]) appears promising.

Regarding belief, a next step is to identify relational frames equivalent to polychromatic complexes, and to develop a sound and complete axiomatization of (most) plausible belief. This requires us to better understand properness with respect to polychromatic complexes. Furthermore, despite including a notion of group knowledge in our logic, we only considered individual belief. Obviously, different notions of group belief provide an interesting topic for future research. A reasonable variant of group belief is:

$$X \leq_G Y \text{ if and only if } \min\{m_a(X) \mid a \in G\} \leq \min\{m_a(Y) \mid a \in G\},$$

from which we could define the interpretation of the modalities $[\trianglerighteq]_G$ and $\mathcal{B}_G$ as in the individual case. Some principles of this notion of group belief are immediate, e.g., that group belief does not imply belief of subgroups (or individual belief). A detailed analysis of this approach, along with its relationship to existing methods, e.g., Gaudou et al. [23], should be studied.

4 Directed Simplicial Models

Chapter Organization. This chapter discusses directed pure simplicial complexes as models of belief, adopting the local perspective throughout. Section 4.1 introduces these structures and defines a semantics of belief over them. In Section 4.2, a correspondence between directed simplicial models and a particular class of local relational models is established. Next, Section 4.3 presents the logic of local beliefs (LLB) and proves its soundness and completeness with respect to our models. Section 4.4 shows that our models, despite being restricted to pure complexes, are at least as expressive as global simplicial models. Finally, Section 4.5 concludes this chapter with a discussion of potential directions for future work.

Modeling belief on simplicial structures is challenging. Chapter 3 presented polychromatic models as a potential solution, but agents having multiple doxastic local states simultaneously remains difficult to interpret. This chapter proposes an alternative approach based on *directed simplicial complexes*, which offer a natural and transparent framework for modeling belief.

Formally, a directed simplicial complex is a triple $(\mathcal{V}, \mathcal{S}, \rho)$ where the pair $\mathbb{C} = (\mathcal{V}, \mathcal{S})$ is a simplicial complex, and $\rho : \mathcal{V} \times \mathcal{S} \to \{0, 1\}$ is a function that assigns a direction to every face of $\mathbb{C}$ from a vertex's perspective. If a face is assigned 1, then it is considered possible from the perspective of that vertex, and impossible otherwise. The function ρ is subjected to two conditions that reflect the structure of the underlying simplicial complex. First, faces that are considered possible from a vertex v must belong to a facet containing v. Second, ρ must be downward closed with respect to set inclusion, i.e., whenever a face X is considered possible from v, all of its subfaces are also considered possible.

In this chapter, we adopt the local point of view and only treat facets as global states, even though directions are defined on arbitrary faces. This is to avoid the introduction of a three-valued logic, and to maintain compatibility with potential extensions to augmented complexes (see Definition 17).

Directions induce a *possibility relation* on facets. An agent a in facet X considers a facet Y possible if and only if $a \in \chi(X \cap Y)$ **and** $\rho(X_a, Y) = 1$. Since the agent may not consider the actual world possible, this relation can be used to model belief. Requiring agent a being contained in both X and Y preserves the principle of knowledge-yields-belief, because every facet that is considered for belief would also be considered when evaluating knowledge. We refrain from introducing a knowledge modality for simplicity.

As already mentioned, the possibility relation is able to model false beliefs. Indeed, consider for example the simplicial complex $\mathbb{C} = (\mathcal{V}, \mathcal{S})$ in Figure 4.1. Let ρ be a possibility function such that $\rho(a, X) = 1$ and $\rho(a, Y) = 0$. In this complex, agent a considers the facet X possible, but not Y. Thus, when in Y, agent a does not consider the actual global state and may form false beliefs. Moreover, the possibility relation is not symmetric because a considers X possible from Y, but it does not consider Y when in X. Consequently, directed simplicial complexes allow us to model false beliefs, while taking the topology of the complex into account.

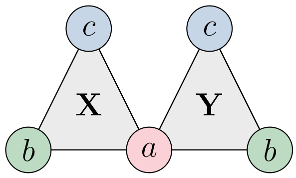

Figure 4.1: A simplicial complex to which directions can be assigned, such that the induced possibility relation is not reflexive or symmetric.

Another interesting feature of directed simplicial complexes is that, in contrast to previous variants of simplicial models, the possibility relation need not be serial. Given a directed simplicial complex $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \rho)$, we can set $\rho(v, X) = 0$, for all $v \in \mathcal{V}$ and $X \in S$. This allows agents to have inconsistent beliefs. Being able to drop reflexivity, symmetry, and seriality is intriguing, making directed simplicial complexes a structure worth studying.

4.1 Semantics

Directions take values 1 or 0, with 1 indicating possibility and 0 representing impossibility. Given a simplicial complex $\mathbb{C}$ and a vertex v, the possibility function selects a subset of the faces belonging to facets that are accessible to v via the standard indistinguishability relation.

Definition 29 (ρ)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S})$ be a simplicial complex. A* possibility function *$\rho : \mathcal{V} \times \mathcal{S} \to \{0, 1\}$ is a function such that $\rho(v, X) = 1$ implies:*

1. *there exists $Y \in \mathcal{F}(\mathbb{C})$ such that $X \subseteq Y$ and $v \in Y$;*

2. *for all $Y \subseteq X$, it holds that $\rho(v, Y) = 1$.*

A simplicial complex equipped with a possibility function is called a *directed simplicial complex.*

Definition 30 (Directed Simplicial Complex)**.** *The triple* $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \rho)$ *is called a* directed simplicial complex *if:*

- $(\mathcal{V}, \mathcal{S})$ *is a simplicial complex;*
- $\rho : \mathcal{V} \times \mathcal{S} \to \{0, 1\}$ *is a possibility function.*

If $(\mathcal{V}, \mathcal{S}, \rho)$ is a directed simplicial complex and χ is a proper coloring on $\mathcal{V}$, we call $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$ a *chromatic directed simplicial complex.* However, if clear from the context, we omit explicitly mentioning that a directed simplicial complex is chromatic. The same notational conventions and terminologies used for simplicial complexes are adopted in the directed setting. For instance, if $\mathbb{C}$ is a directed simplicial complex, then $\mathcal{F}(\mathbb{C})$ denotes the facets of its underlying simplicial complex. Similarly, a directed simplicial complex is called pure if and only if its underlying simplicial complex is pure.

The possibility function gives rise to a *possibility relation* upon which a modality for belief can be defined.

Definition 31 (Possibility Relation)**.** *Let* $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$ *be a directed simplicial complex, and let* $C = (\mathcal{V}, \mathcal{S})$*. For* $a \in \mathsf{Ag}$*, we call*

$$\mathsf{Pos}_a = \{(X, Y) \in \mathcal{F}(C) \times \mathcal{F}(C) \mid \rho(X_a, Y) = 1\},$$

the possibility relation *of the agent* a*.*

Lemma 18. *Let* $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$ *be a directed simplicial complex. For each agent* $a \in \mathsf{Ag}$*, the relation* Pos_a *is transitive.*

Proof. Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$ be a directed simplicial complex such that there exist facets $X, Y, Z \in \mathcal{F}(\mathbb{C})$ and an agent $a \in \mathsf{Ag}$ such that:

$$\rho(X_a, Y) = 1 \quad \text{and} \quad \rho(Y_a, Z) = 1.$$

By Definition 29, $\rho(X_a, Y) = 1$ implies that $X_a = Y_a$. Therefore, it holds that $\rho(X_a, Z) = 1$. Thus, the relation Pos_a is transitive. □

Lemma 19. *Let* $\mathcal{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$ *be a directed simplicial complex. For each agent* $a \in \mathsf{Ag}$*, the relation* Pos_a *is euclidean.*

Proof. Similar to the proof of Lemma 18. □

Example 5 illustrates some properties of the possibility relation.

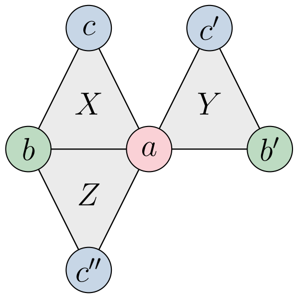

Figure 4.2: A directed simplicial complex.

Example 5. *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \rho, \chi)$ be the directed simplicial complex in Figure 4.2, where χ is as indicated by the color of vertices and ρ is such that:*

- $\rho(a, Y) = \rho(b', Y) = \rho(c', Y) = 0$;
- $\rho(a, X) = \rho(b, X) = \rho(c, X) = 1$;
- $\rho(b, Z) = \rho(c'', Z) = 1$ *and* $\rho(a, Z) = 0$.

Some example statements that are true about $\mathbb{C}$ are:

1. *when in Z, agent a rules out Z, but considers X;*
2. *when in X, agent a does not consider Z possible;*
3. *when in Y, agent c does not deem a facet to be possible.*

Thus, the relation Pos_a need not be reflexive, serial, or symmetric.

A directed simplicial model is a directed simplicial complex equipped with a local valuation (see Definition 8). Since the labeling is local, we overload notation and use $\Vdash_\sigma^{\mathsf{loc}}$ for the satisfaction relation.

Definition 32 (Directed Simplicial Model)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$ be a pure directed simplicial complex. The pair $\mathcal{C} = (\mathbb{C}, L)$, where L is a local labeling, is called a* directed simplicial model*.*

Definition 33 ($\Vdash_\sigma^{\mathsf{loc}}$)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$ be a pure directed simplicial complex. For a directed simplicial model $\mathcal{C} = (\mathbb{C}, L)$, all facets $X \in \mathcal{F}(\mathbb{C})$, we define the relation $\mathcal{C}, X \Vdash_\sigma^{\mathsf{loc}} \phi$ by induction on $\phi \in \mathcal{L}$:*

$\mathcal{C}, X \Vdash_\sigma^{\mathsf{loc}} p$	*iff*	$p \in L(X_a)$ *and* $p \in \mathsf{P}_a$
$\mathcal{C}, X \Vdash_\sigma^{\mathsf{loc}} \neg\phi$	*iff*	$\mathcal{C}, X \nVdash_\sigma^{\mathsf{loc}} \phi$
$\mathcal{C}, X \Vdash_\sigma^{\mathsf{loc}} \phi \wedge \psi$	*iff*	$\mathcal{C}, X \Vdash_\sigma^{\mathsf{loc}} \phi$ *and* $\mathcal{C}, X \Vdash_\sigma^{\mathsf{loc}} \psi$
$\mathcal{C}, X \Vdash_\sigma^{\mathsf{loc}} \Box_a\phi$	*iff*	*for all* $Y \in \mathcal{F}(\mathbb{C})$, $(X, Y) \in \mathsf{Pos}_a$ *implies* $\mathcal{C}, Y \Vdash_\sigma^{\mathsf{loc}} \phi$.

Let $\mathcal{C} = (\mathbb{C}, L)$, where $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$, be a directed simplicial model. We write $\mathcal{C} \Vdash^{\mathsf{loc}}_{\sigma} \phi$, if $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \phi$ for all $X \in \mathcal{F}(\mathbb{C})$. Moreover, $\Vdash^{\mathsf{loc}}_{\sigma} \phi$ denotes that $\mathcal{C} \Vdash^{\mathsf{loc}}_{\sigma} \phi$ for every directed simplicial model $\mathcal{C}$.

Remark 6. *Definition 33 restricts satisfiability to facets, although directions are defined on arbitrary faces. This is done to avoid a three-valued logic, and to maintain compatibility with potential extensions to augmented complexes (see Definition 17). Using directed hypergraphs offers a more succinct representation (cf. Gomes et al. [50]).*

We refrain from interpreting the formula $\Box_a \bot$ as agent a being dead, since $\Box_a \bot$ can be true in a facet containing a vertex with color a. Instead, we may view agents for which $\Box_a \bot$ holds as being infinitely slow but not crashed. This interpretation aligns well with asynchronous distributed systems, in which crashed processes and infinitely slow ones are treated the same.

Example 6 illustrates the previously introduced notions. It also shows that nested modalities behave as expected.

Example 6. *Consider the simplicial complex depicted in Figure 4.3, and let $\mathcal{C} = (\mathbb{C}, L)$, where $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$, be a directed simplicial model where:*

- $\mathcal{V} = \{a, b, c, b', c'\}$;
- $\mathcal{S} = (\mathsf{Pow}(X) \cup \mathsf{Pow}(Y) \cup \mathsf{Pow}(Z)) \setminus \{\emptyset\}$, *for* $X = \{a, b, c\}$, $Y = \{a, b', c\}$, *and* $Z = \{a, b', c'\}$;
- χ *as indicated by the name of vertices;*
- ρ *is a possibility function such that:*
  - $\rho(a, X) = \rho(a, Z) = 0$ *and* $\rho(a, Y) = 1$;
  - $\rho(b, X) = 1$, $\rho(b', Y) = 0$ *and* $\rho(b', Z) = 1$;
  - $\rho(c, X) = \rho(c, Y) = 1$ *and* $\rho(c', Z) = 1$.
- L *is such that for some* $p \in \mathsf{P}_b$ *we have* $p \in L(b')$ *and* $p \notin L(b)$.

Agent a always considers Y to be the only possible facet. In X, agent c considers both X and Y possible, and in Z, it solely considers Z. Conversely, when in X, agent b deems only X possible, and, while being in Y or Z, it only considers Z. It can be verified that $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a p$, $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a \neg \Box_c p$, and $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a \Box_b \Box_c p$. Moreover, it is also the case that a falsely believes p while in X, i.e., $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a p \wedge \neg p$. Lastly, since $\rho(a, Y) = 1$, we find that for all $\phi \in \mathcal{L}$ it holds that $\mathcal{C}, Y \Vdash^{\mathsf{loc}}_{\sigma} \Box_a \phi \to \phi$.

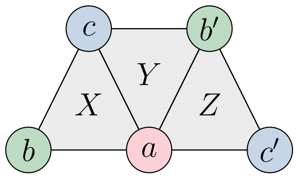

Figure 4.3: The underlying simplicial complex of the directed simplicial model in Example 6.

4.2 Correspondence to Relational Semantics

Chapter 2 introduced the notions of local valuation (see Definition 4) and properness (see Definition 11) on symmetric and transitive frames. However, Example 5 showed that Pos_a need not be symmetric. Consequently, frames corresponding to directed simplicial complexes need not be symmetric either. Thus, we must generalize these notions to frames that are not necessarily symmetric. Prior to this, we define what local and global states mean for such frames. In fact, Definition 34 extends the concept of local and global states to arbitrary frames.

Definition 34 (Local and Global States)**.** *Let $F = (W, R)$ be a frame. For every world $w \in W$ and agent $a \in \mathsf{Ag}$, we define the following:*

- $\widetilde{w}_a = \{v \in W \mid (w, v) \in (R_a \cup R_a^{-1})^+\} \cup \{w\}$, *where $(R_a \cup R_a^{-1})^+$ is the transitive closure of the symmetric closure of R_a. We call $(\widetilde{w}_a, a)$ the* local state of a in w*;*
- $X_w = \{(\widetilde{w}_a, a) \mid a \in \mathsf{Ag}\}$ *is the* global state *represented by w.*

Remark 7. *The requirement that every local state $\widetilde{w}_a$ contains the actual world w is for technical reasons that we will address later.*

The motivation for locality in the directed setting remains unchanged. A valuation is local if and only if it assigns the same local propositional variables to worlds where the agent has the same local state. Again, we refer to relational models with a local valuation as local relational models. The same notational conventions of local models in Chapter 2 apply to local models introduced here. It is straightforward to verify that locality is satisfied under Definition 35.

Definition 35 (Local Valuation)**.** *Let $F = (W, R)$ be a frame. We call $V : W \to \mathsf{Pow}(\mathsf{Prop})$ a* local valuation *if and only if for all worlds $w, v \in W$:*

$$\widetilde{w}_a \cap \widetilde{v}_a \neq \emptyset \text{ implies } p \in V(w) \text{ iff } p \in V(v), \text{ for all } a \in \mathsf{Ag} \text{ and } p \in \mathsf{P}_a.$$

Properness (Definition 36) is defined in terms of sequences of a-transitions, and is similar to the alternative characterization proposed in Remark 3: a frame F is proper if and only if every two distinct worlds differ in at least one agent's local state. The generalized definition presented here differs slightly from the standard definition of properness. Nevertheless, both definitions coincide when evaluated on epistemic frames (see Definition 13).

Definition 36 (Proper)**.** *Let $F = (W, R)$ be a frame. We say that F is* proper *if and only if, for all worlds $w, v \in W$ with $w \neq v$, there exists an agent $a \in \mathsf{Ag}$ such that $\widetilde{w}_a \cap \widetilde{v}_a = \emptyset$.*

We show the correspondence between local relational models and directed simplicial models, beginning with the more complicated direction: transforming a relational model into a directed simplicial model.

Definition 37 (Transformation from Relational Model to Simplicial Model)**.** *Let $F = (W, R)$ be a proper frame, and let $\mathcal{M} = (F, V)$ be a local relational model. We construct a directed complex $\mathbb{C} = (\mathcal{V}^T, \mathcal{S}^T, \chi^T, \rho^T)$ and a directed simplicial model $\mathcal{C}^T = (\mathbb{C}, L^T)$ as follows:*

- $\mathcal{V}^T = \{(\widetilde{w}_a, a) \mid w \in W \text{ and } a \in \mathsf{Ag}\}$;
- $\mathcal{S}^T = \bigcup_{w \in W} \mathsf{Pow}(\{X_w \mid w \in W\}) \setminus \{\emptyset\}$;
- $\chi^T((\widetilde{w}_a, a)) = a$ *for all* $w \in W$ *and* $a \in \mathsf{Ag}$;
- $\rho^T((\widetilde{w}_a, a), Y) = 1$ *if and only if there exists a facet* X_v *such that* $Y \subseteq X_v$ *and* $(w, v) \in R_a$, *and* $\rho^T((\widetilde{w}_a, a), Y) = 0$, *otherwise;*
- $p \in L^T((\widetilde{w}_a, a))$ *if and only if* $p \in V(w)$ *and* $p \in \mathsf{P}_a$.

We will now go over the construction of the directed simplicial complex proposed in Definition 37. Consider the relational structure in Figure 4.4 on the left. Although the model has three different worlds, agent a has only one local state. This is reflected in the fact that:

$$\widetilde{w_X}_a = \widetilde{w_Y}_a = \widetilde{w_Z}_a = \{w_X, w_Y, w_Z\}.$$

The relational structure is equivalent to the directed simplicial complex shown in Figure 4.4 on the right, where the possibility function is defined such that:

$$\rho(b_1, X) = 1,\ \rho(b_2, Y) = \rho(a, Y) = 1,\ \text{and } \rho(b_3, Z) = 1.$$

As indicated in Remark 7, we include the actual world in every local state. This ensures that we can transform a frame to a directed simplicial complex.

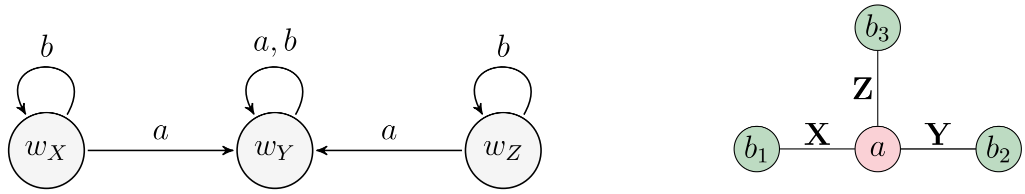

Figure 4.4: A relational structure (left) where a has only one local state, and its corresponding simplicial complex (right).

For example, consider the two-agent relational frame depicted on the left in Figure 4.5. Under any valuation, agent b will have inconsistent beliefs in both worlds. Since the accessibility relation R_b is empty, clearly $(R_b \cup R_b^{-1})^+$ is also empty. However, because all agents are assumed to be alive in all facets, agent b can actually distinguish w from v. Thus, $(\widetilde{v}_b, b)$ and $(\widetilde{w}_b, b)$ should represent distinct local states. If this were not the case, b would have only one vertex in the associated directed simplicial complex, and it would collapse into one facet. The relations for local states are:

$$\widetilde{w}_a = \widetilde{v}_a = \{w, v\};\ \widetilde{v}_b = \{v\};\ \text{and}\ \widetilde{w}_b = \{w\}.$$

Indeed, we find that $(\widetilde{v}_b, b) \neq (\widetilde{w}_b, b)$. However, if we had not included the actual global state, it would hold that $\widetilde{w}_b = \emptyset = \widetilde{v}_b$, which would mean that we cannot transform this frame to a directed complex. The directed simplicial complex corresponding to the frame is shown in Figure 4.5 on the right. The directions are:

$$\rho((\widetilde{w}_a, a), X_v) = 1 \text{ and } \rho((\widetilde{v}_b, b), X_v) = \rho((\widetilde{w}_b, b), X_w) = \rho((\widetilde{w}_a, a), X_w) = 0.$$

Lemma 20 shows that the worlds of the proper frame and its transformed complex coincide.

Lemma 20. *Let $F = (W, R)$ be a proper frame, and consider $\mathbb{C} = (\mathcal{V}^T, \mathcal{S}^T)$. The function:*

$$b : \mathcal{F}(\mathbb{C}) \to W,\ \textit{where}\ b(X_w) = w,\ \textit{for all}\ w \in W$$

is a bijection from $\mathcal{F}(\mathbb{C})$ to W.

Proof. First, we remark that by construction of $\mathcal{S}^T$, it holds that:

$$\mathcal{F}(\mathbb{C}) = \{X_w \mid w \in W\},$$

Figure 4.5: A relational frame (left) where agent b considers no world. On the right is the underlying simplicial complex of the transformed frame.

We now show that b is a bijection. The mapping b is surjective, because for every X_w, there exists $w \in W$ by definition. Regarding injectivity, suppose that for some facets $X_w, X_v \in \mathcal{F}(\mathbb{C})$, we have $b(X_w) = b(X_v)$. We show that this implies $X_w = X_v$. By Definition 34, $X_w = \{(\widetilde{w}_a, a) \mid a \in \mathsf{Ag}\}$ and $X_v = \{(\widetilde{v}_a, a) \mid a \in \mathsf{Ag}\}$. Now we claim that for all $a \in \mathsf{Ag}$, we have that $(\widetilde{w}_a, a) = (\widetilde{v}_a, a)$. By properness (Definition 36), this implies that $X_w = X_v$. To see this, let $a \in \mathsf{Ag}$ be arbitrary. By definition of $b(\cdot)$, our assumption implies that $w = v$. But then we have that:

$$\begin{aligned}\widetilde{w}_a &= \{u \in W \mid (w,u) \in (R_a \cup R_a^{-1})^+\} \cup \{w\} && \text{(Def. 34)}\\ &= \{u \in W \mid (v,u) \in (R_a \cup R_a^{-1})^+\} \cup \{v\} && (w = v)\\ &= \widetilde{v}_a\end{aligned}$$

i.e., $(\widetilde{w}_a, a) = (\widetilde{v}_a, a)$ for all $a \in \mathsf{Ag}$. □

We use the following lemma to show that the associated model satisfies the conditions of a directed simplicial model.

Lemma 21. *ρ^T is a possibility function.*

Proof. We show that ρ^T has the properties of a possibility function:

1. if $\rho^T((\widetilde{w}_a, a), Y) = 1$, then there exists a facet X_v such that $Y \subseteq X_v$ and $(w, v) \in R_a$. Since $\widetilde{w}_a \cap \widetilde{v}_a \neq \emptyset$, we find that $(\widetilde{w}_a, a) = (\widetilde{v}_a, a) \in X_v$;

2. for all $Z \subseteq Y \subseteq X_v$, we have $\rho((\widetilde{w}_a, a), Z) = 1$ by definition. □

Showing that $\mathcal{V}^T, \mathcal{S}^T$, and χ^T are well-defined is straightforward. However, the fact that L^T is well-defined is not obvious.

Lemma 22. *L^T is a well-defined local labeling.*

Proof. We first prove that L^T is well-defined, i.e., for every vertex $v \in \mathcal{V}^T$, there is at most one element $A \in \mathsf{Pow}(\mathsf{Prop})$ such that $L^T(v) = A$. Towards a contradiction, assume that this is not the case. That is, there exist two elements

$A, B \in \mathsf{Pow}(\mathsf{Prop})$ such that $A \neq B$ and $A = L^T(v) = B$. Since $A \neq B$, there exists $p \in \mathsf{P}_a$ such that $p \in A$ and $p \notin B$. The only way that this can occur is if there exist two worlds $w, x \in W$ with $\widetilde{w}_a = \widetilde{x}_a$ such that $\mathcal{M}, w \Vdash^{\mathsf{loc}} p$ and $\mathcal{M}, x \nVdash^{\mathsf{loc}} p$. If this is the case, L^T would not be well-defined because

$$p \in L^T((\widetilde{w}_a, a)) = L^T((\widetilde{x}_a, a)) \not\ni p.$$

However, by Definition 35, this cannot happen because V is a local valuation. Lastly, by definition, L^T is a local labeling because it only assigns elements of P_a to vertices of the form $(\widetilde{w}_a, a)$. ☐

Corollary 1. *$\mathcal{C}^T$ is a directed simplicial model.*

We now show that every relational model $\mathcal{M}$ and its corresponding directed simplicial model $\mathcal{C}^T$ are pointwise equivalent. The next lemma is crucial for this proof.

Lemma 23. *We have $(w, v) \in R_a$ if and only if $\rho^T((\widetilde{w}_a, a), X_v) = 1$.*

Proof. We show both directions separately.

- $\Longrightarrow$: By definition.
- $\Longleftarrow$: If $\rho^T((\widetilde{w}_a, a), X_v) = 1$, then there exists Y and $u \in W$ such that $X_v \subseteq Y = X_u$ and $(w, u) \in R_a$. Since X_v is a facet, it follows that $X_v = Y = X_u$. Moreover, due to Lemma 20, we have that $u = v$ and thus $(w, v) \in R_a$. ☐

Lemma 24. *Let $F = (W, R)$ be a proper frame, and let $\mathcal{M} = (F, V)$ be a proper relational model with a local valuation. It holds that:*

$$\mathcal{M}, w \Vdash^{\mathsf{loc}} \phi \textit{ if and only if } \mathcal{C}^T, X_w \Vdash^{\mathsf{loc}}_{\sigma} \phi.$$

Proof. The claim can be proven by simple induction on the length of ϕ and using Lemma 23. ☐

Next, we show the straightforward direction: given a directed simplicial model $\mathcal{C}$, we define a logically equivalent proper relational model $\mathcal{M}^T$ with a local valuation. We take the usual approach (cf. Goubault et al. [29]).

Definition 38 (Transformation from Simplicial Model to Relational Model)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$ be a directed pure simplicial complex, and let $\mathcal{C} = (\mathbb{C}, L)$ be a directed simplicial model. We construct a proper frame $F^T = (W^T, R^T)$ and a local relational model $\mathcal{M}^T = (F^T, V^T)$ as follows:*

- $W^T = \{w_X \mid X \in \mathcal{F}(\mathbb{C})\}$;

- $(w_X, v_Y) \in R_a^T$ *if and only if* $\rho(X_a, Y) = 1$, *for all* $a \in \mathsf{Ag}$ *and* $w_X, v_Y \in W$*;*
- $p \in V^T(w_X)$ *if and only if* $p \in \ell(v)$ *for some* $v \in X$*, for all* $w_X \in W$*.*

Clearly, there is a bijection from $\mathcal{F}(\mathbb{C})$ to W^T. Properness of F^T follows immediately because $(\mathcal{V}, \mathcal{S})$ is a simplicial complex. Showing that V^T is well-defined is straightforward. We only need to show that V^T is local.

Lemma 25. V^T *is a local valuation.*

Proof. Let $w_X, w_Y \in W$ and suppose that $\widetilde{w}_X \cap \widetilde{w}_Y \neq \emptyset$. Let $a \in \mathsf{Ag}$ and $p \in \mathsf{P}$. We show that $p \in V^T(w_X)$ if and only if $p \in V^T(w_Y)$. We use, again, the observation that $\widetilde{w}_X \cap \widetilde{w}_Y \neq \emptyset$ implies $\widetilde{w}_X = \widetilde{w}_Y$. Hence, between w_X and w_Y, there exists a finite sequence of a-transitions

$$w_X = w_{X_0}(R_a^T \cup R_a^{T-1})w_{X_1}(R_a^T \cup R_a^{T-1}) \ldots (R_a^T \cup R_a^{T-1})w_{X_n} = w_Y.$$

If $w_X = w_Y$ then we are done, thus assume otherwise. We prove by induction that for all $n \geq 1$, this implies that $X_{0a} = X_{na}$. Then, by locality of $\mathcal{C}$, $p \in V^T(w_X)$ if and only if $p \in L(X_a) = L(Y_a)$, if and only if $p \in V^T(w_Y)$.

For the base case, we have either $w_{X_0} R_a^T w_{X_1}$, or $w_{X_1} R_a^T w_{X_0}$. By definition of W^T, X_0 and X_1 are both facets, so X_{0a} and X_{1a} are defined. But then, in either case, $X_{0a} = X_{1a}$ by the definitions of R_a^T and ρ. For the induction step, the same reasoning yields $X_{na} = X_{n+1a}$, such that $X_{0a} = X_{n+1a}$ (induction hypothesis). Thus, V^T is indeed a local valuation. □

Finally, transitivity and euclideanity of each R_a follow immediately from transitivity and euclideanity of the possibility relation Pos_a (Lemmas 18 and 19). We obtain the following corollary.

Corollary 2. $\mathcal{M}^T$ *is a proper relational model with a local valuation.*

It remains to show that $\mathcal{C}$ and $\mathcal{M}^T$ are pointwise equivalent, i.e. that the two models satisfy the same formulas over $\mathcal{L}$ at each world.

Lemma 26. *Let* $\mathcal{C}$ *be a directed simplicial model. It holds that*

$$\mathcal{C}, X \Vdash_\sigma^{\mathsf{loc}} \phi \text{ if and only if } \mathcal{M}^T, w_X \Vdash^{\mathsf{loc}} \phi.$$

Proof. We show the claim by induction on ϕ. The atomic and boolean cases are immediate. We focus on the case where $\phi = \Box_a \psi$. For the left-to-right direction, suppose $\mathcal{C}, X \Vdash_\sigma^{\mathsf{loc}} \Box_a \psi$. To show that $\mathcal{M}^T, w_X \Vdash^{\mathsf{loc}} \Box_a \psi$, let $w_Y \in W^T$ and suppose $(w_X, w_Y) \in R_a^T$. By definition, this means that $Y \in \mathcal{F}(\mathbb{C})$ with $\rho(X_a, Y) = 1$. Then, $\mathcal{C}, Y \Vdash_\sigma^{\mathsf{loc}} \psi$ and, by the induction hypothesis, $\mathcal{M}^T, w_Y \Vdash^{\mathsf{loc}} \psi$. The right-to-left direction is similar. □

We combine Lemmas 24 and 26 to conclude that for each directed simplicial model, there exists a logically equivalent proper relational model with a local valuation, and vice versa.

4.3 Soundness and Completeness

The *local logic of belief* (LLB) is a modal logic consisting of the following axioms:

$$\text{all propositional tautologies} \tag{Taut}$$
$$\Box_a(\phi \to \psi) \to (\Box_a\phi \to \Box_a\psi) \tag{K}$$
$$\Box_a\phi \to \Box_a\Box_a\phi \tag{4}$$
$$\Diamond_a\phi \to \Box_a\Diamond_a\phi \tag{5}$$
$$p_a \to \Box_a p_a \tag{L+}$$
$$\neg p_a \to \Box_a\neg p_a \tag{L−}$$

and the inference rules modus ponens MP and $\Box_a$-necessitation $\Box_a$-Nec for all agents $a \in \mathsf{Ag}$. We write $\vdash^{\mathsf{LLB}} \varphi$ to denote that the formula $\varphi \in \mathcal{L}$ can be deduced in our system.

$$\frac{A \qquad A \to B}{B} \;\text{(MP)} \qquad \frac{A}{\Box_a A} \;(\Box_a\text{-Nec})$$

Lemma 27 states that locality can be derived in LLB. It is a direct consequence of the axioms L+ and L−.

Lemma 27. *For all $a \in \mathsf{Ag}$ and $p \in \mathsf{P}_a$, it holds that $\vdash^{\mathsf{LLB}} \Box_a p \vee \Box_a\neg p$.*

Proof.

1. $p \to \Box_a p$ — instance of L+
2. $\neg p \to \Box_a\neg p$ — instance of L−
3. $(p \vee \neg p) \to (\Box_a p \vee \Box_a\neg p)$ — propositional reasoning
4. $\Box_a p \vee \Box_a\neg p$ — propositional reasoning.

□

For an agent a with inconsistent beliefs, both $\Box_a p$ and $\Box_a\neg p$ hold.

Remark 8. *Axioms L+ and L− ensure that if $\Box_a\top$ is true in a world, then for each $p \in \mathsf{P}_a$, the agent a believes exclusively p or $\neg p$.*

The proof of soundness (Lemma 28) is carried out in the usual way. To show completeness, rather than converting directed simplicial models to equivalent relational models and proving their completeness, we construct the canonical directed simplicial model directly. Our approach is similar to that used by Randrianomentsoa et al. [46].

Lemma 28. *For every formula $\phi \in \mathcal{L}$ it holds that*

$$\vdash^{\mathsf{LLB}} \phi \textit{ implies } \Vdash^{\mathsf{loc}}_{\sigma} \phi.$$

Proof. If ϕ is an instance of a propositional tautology, the proof that $\Vdash^{\mathsf{loc}}_{\sigma} \phi$ is standard. We only show the statement for the axioms K,4,5, $\mathsf{L}+$, and $\mathsf{L}-$. Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi, \rho)$ be an arbitrary pure directed simplicial complex and let $\mathcal{C} = (\mathbb{C}, L)$ be a directed simplicial model.

1. K: let $X \in \mathcal{F}(\mathbb{C})$ be arbitrary and assume that $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a(\phi \to \psi)$ and $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a\phi$. Further, let $Y \in \mathcal{F}(\mathbb{C})$ be such that $\rho(X_a, Y) = 1$. By the definition of truth, it follows that $\mathcal{C}, Y \Vdash^{\mathsf{loc}}_{\sigma} \phi \to \psi$ as well as $\mathcal{C}, Y \Vdash^{\mathsf{loc}}_{\sigma} \phi$. Consequently, we have that $\mathcal{C}, Y \Vdash^{\mathsf{loc}}_{\sigma} \psi$. Because Y was arbitrary, it holds that $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a\psi$.

2. 4: let $X \in \mathcal{F}(\mathbb{C})$ be arbitrary and assume that $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a\phi$. Additionally, let $Y \in \mathcal{F}(\mathbb{C})$ be such that $\rho(X_a, Y) = 1$ and let $Z \in \mathcal{F}(\mathbb{C})$ be such that $\rho(Y_a, Z) = 1$. Since ρ preserves connectivity, and with χ being a proper coloring, we find that $X_a = Y_a = Z_a$. Therefore, $\rho(X_a, Z) = 1$ and $\mathcal{C}, Z \Vdash^{\mathsf{loc}}_{\sigma} \phi$ which yields $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a\Box_a\phi$.

3. 5: let $X \in \mathcal{F}(\mathbb{C})$ be arbitrary and assume that $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Diamond_a\phi$. Thus, there exists $Y \in \mathcal{F}(\mathbb{C})$ such that $\rho(X_a, Y) = 1$ and $\mathcal{C}, Y \Vdash^{\mathsf{loc}}_{\sigma} \phi$. By the same reasoning as earlier, we find that $X_a = Y_a$. Hence, for all $Z \in \mathcal{F}(C)$ with $\rho(X_a, Z) = 1$, it holds that $\rho(Z_a, Y) = 1$, because $X_a = Z_a$. Therefore, $\mathcal{C}, Z \Vdash^{\mathsf{loc}}_{\sigma} \Diamond_a\phi$ by definition. Since Z was arbitrary, $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a\Diamond_a\phi$.

4. $\mathsf{L}+$: let $X \in \mathcal{F}(\mathbb{C})$ be such that for some $p_a \in \mathsf{P}_a$ we have $p_a \in \ell(X_a)$. If there is no $Y \in \mathcal{F}(\mathbb{C})$ with $\rho(X_a, Y) = 1$, then $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a p_a$ is vacuously true. If there exists such a Y, then, again because ρ preserves connectivity, we find that $X_a = Y_a$ and thus $p_a \in \ell(Y_a)$ and $\mathcal{C}, Y \Vdash^{\mathsf{loc}}_{\sigma} p_a$. Since Y was arbitrary, we conclude that $\mathcal{C}, X \Vdash^{\mathsf{loc}}_{\sigma} \Box_a p_a$.

5. $\mathsf{L}-$: analogous to the case for $\mathsf{L}+$. □

We will now introduce the necessary machinery to show completeness.

Definition 39. *Let $a \in \mathsf{Ag}$ be an agent and let $\Gamma \subseteq \mathcal{L}$. We define:*

- $\Gamma \setminus \Box_a = \{\phi \mid \Box_a\phi \in \Gamma\}$;
- $\Box_a\Gamma = \{\Box_a\phi \mid \Box_a\phi \in \Gamma\}$;
- $\mathsf{P}^a_\Gamma = \{p \in \mathsf{P}_a \mid p \in \Gamma\}$.

We denote that a formula ϕ is derivable in a normal modal logic L by $\vdash^{\mathsf{L}} \phi$. If ϕ is not derivable, we write $\nvdash^{\mathsf{L}} \phi$.

Definition 40 (Maximal Consistent Set)**.** *A set $\Gamma \subseteq \mathcal{L}$ is* consistent *with respect to a normal modal logic L if and only if $\Gamma \nvdash^{\mathsf{L}} \bot$. Γ is* maximal consistent *regarding L if it is consistent and none of its proper supersets are consistent. We define:*

$$\mathbb{G}_{\mathsf{L}} = \{\Gamma \subseteq \mathcal{L} \mid \Gamma \textit{ is a maximal consistent set for } \mathsf{L}\}.$$

Whenever clear from the context, we do not mention that a set is maximal consistent explicitly for LLB. Specifically we will write $\mathbb{G}$ instead of $\mathbb{G}_{\mathsf{LLB}}$. The Lindenbaum Lemma (Lemma 29) is standard and used throughout our completeness proof.

Lemma 29 (Lindenbaum Lemma)**.** *Any consistent set $\Gamma \neq \emptyset$ of a normal modal logic can be extended to a maximal consistent set Δ such that $\Gamma \subseteq \Delta$.*

Proof. The proof is standard and is thus omitted. □

As usual, global states will be described by maximal consistent stets. Definition 41 shows how exactly global states are formed.

Definition 41 (X_Γ)**.** *Let $\Gamma \in \mathbb{G}$. We define $\Gamma_a = \Box_a\Gamma \cup \mathsf{P}^a_\Gamma$ and*

$$X_\Gamma = \{\Gamma_a \mid a \in \mathsf{Ag}\}.$$

In the canonical directed simplicial complex, the vertices will be given by the set $\{\Gamma_a \mid \Gamma \in \mathbb{G} \text{ and } a \in \mathsf{Ag}\}$, and the facets are specified by the following set $\{X_\Gamma \mid \Gamma \in \mathbb{G}\}$.

Unlike the canonical model construction by Randrianomentsoa et al. [46], our construction explicitly encodes the true local propositions in the vertices. Since beliefs about local propositions are not necessarily truthful, the truth of a local proposition $p \in \mathsf{P}_a$ is not equivalent to agent a's belief in its truth. Therefore, the truth value of p cannot be recovered from the set $\Box_a\Gamma$. Consequently, omitting the set P^a_Γ from $\Box_a\Gamma$ would result in an improper canonical directed simplicial model. In order to see this, let $\mathsf{Ag} = \{a\}$ and $\mathsf{Prop} = \{p, q\} = \mathsf{P}_a$, and define $\Delta_0, \Gamma_0 \subseteq \mathcal{L}$ as follows:

$$\Delta_0 = \{p, \neg q, \Box_a p, \Box_a q\}$$
$$\Gamma_0 = \{\neg p, q, \Box_a p, \Box_a q\}.$$

Both Δ_0 and Γ_0 are consistent (despite agent a's inconsistent beliefs). By the Lindenbaum Lemma (Lemma 29), these sets can be extended to maximal consistent sets $\Delta, \Gamma \in \mathbb{G}$. Observe that both Δ and Γ will contain $\Box_a \bot$ and, as a consequence, $\Box_a \phi$ for all $\phi \in \mathcal{L}$. But then $\Box_a \Delta = \Box_a \Gamma = \{\Box_a \phi \mid \phi \in \mathcal{L}\}$. Now, if we omit the local propositions P^a_Δ and P^a_Γ from the respective vertices in the canonical construction, then $X_\Delta = \Box_a \Delta = \Box_a \Gamma = X_\Gamma$, while clearly $\Delta \neq \Gamma$. In fact, all local states of the model in which agent a has inconsistent beliefs will be described by the same local state $\{\Box_a \phi \mid \phi \in \mathcal{L}\}$. We avoid this by defining $\Gamma_a = \Box_a \Gamma \cup \mathsf{P}^a_\Gamma$ for all $\Gamma \in \mathbb{G}$.

The next lemma ensures that there is a bijection from the set $\mathbb{G}$ to the set $\{X_\Gamma \mid \Gamma \in \mathbb{G}\}$. Although seemingly trivial, it is a key result for proving Lemma 34 later.

Lemma 30. *For all* $\Gamma, \Delta \in \mathbb{G}$ *it holds that* $X_\Gamma = X_\Delta$ *if and only if* $\Gamma = \Delta$.

Proof. We show both directions separately:

1. $\Longrightarrow$: Assume that $X_\Gamma = X_\Delta$. We proceed to show that

   $$\phi \in \Gamma \text{ if and only if } \phi \in \Delta$$

   by induction on the length of ϕ.

   i. Case $\phi \equiv p$ for $a \in \mathsf{Ag}$ and $p \in \mathsf{P_a}$: by assumption $\mathsf{P}^a_\Gamma = \mathsf{P}^a_\Delta$ and thus $p \in \Delta$ if and only if $p \in \Gamma$.

   ii. Case $\phi \equiv \neg\psi$: $\neg\psi \in \Gamma$ if and only if $\psi \notin \Gamma$ (by maximal consistency) if and only if $\psi \notin \Delta$ (by induction hypothesis) if and only if $\neg\psi \in \Delta$.

   iii. The case for $\phi \equiv \psi \wedge \psi'$ is similar.

   iv. Case $\phi \equiv \Box_a \psi$: by assumption $\Box_a \Gamma = \Box_a \Delta$ and thus $\Box_a \psi \in \Gamma$ if and only if $\Box_a \psi \in \Delta$.

2. $\Longleftarrow$: If $\Delta = \Gamma$, then $\Box_a \Gamma = \Box_a \Delta$ and $\mathsf{P}^a_\Gamma = \mathsf{P}^a_\Delta$ for all $a \in \mathsf{Ag}$, which yields the claim. □

Lemma 31 states that if a set $\Delta \in \mathbb{G}$ is consistent with the beliefs of an agent in $\Gamma \in \mathbb{G}$, then that agent has the same view in Δ as in Γ.

Lemma 31. *Let* $\Gamma \in \mathbb{G}$. *For all* $\Delta \in \mathbb{G}$ *with* $\Gamma \setminus \Box_a \subseteq \Delta$, *it holds that*

$$\Box_a \Gamma = \Box_a \Delta \text{ and } \mathsf{P}^a_\Gamma = \mathsf{P}^a_\Delta.$$

Proof. We first show that $\mathsf{P}^a_\Gamma = \mathsf{P}^a_\Delta$. Since $\Gamma \setminus \Box_a \subseteq \Delta$, we find $\Box_a \bot \notin \Box_a \Gamma$ due to the consistency of Δ. Because $\mathsf{L+}$ and $\mathsf{L-}$ are axioms of LLB and Γ is a maximal consistent set, we find that $\mathsf{P}^a_\Gamma = \mathsf{P}^a_\Delta$. Indeed, as a consequence of Γ being a maximal consistent set, and by $\mathsf{L+}$ and $\mathsf{L-}$, it holds that:

$$\Box_a p \in \Box_a \Gamma \text{ or } \Box_a \neg p \in \Box_a \Gamma \text{ for all } p \in \mathsf{P}_a. \qquad (\star)$$

Conditioned on $\Box_a \bot \notin \Box_a \Gamma$, exactly one belongs to $\Box_a \Gamma$. Thus,

$$p \in \mathsf{P}^a_\Gamma \text{ if and only if } \Box_a p \in \Box_a \Gamma.$$

As a consequence of $(\star)$, we obtain $\mathsf{P}^a_\Gamma = \mathsf{P}^a_\Delta$.

The proof that $\Box_a \Gamma = \Box_a \Delta$ is straightforward. Since 4 and 5 are axioms of LLB and $\Gamma, \Delta \in \mathbb{G}$, it holds that:

1. if $\Box_a \phi \in \Box_a \Gamma$ then $\Box_a \phi \in \Box_a \Delta$;
2. if $\Box_a \phi \notin \Box_a \Gamma$ then $\Box_a \phi \notin \Box_a \Delta$.

Therefore, we conclude that $\Box_a \Gamma = \Box_a \Delta$. □

We are now able to specify the canonical construction (Definition 42).

Definition 42 (Canonical Construction)**.** *The* canonical directed simplicial complex *for* LLB *is a quadruple*

$$\mathbb{C}^c = (\mathcal{V}^c, \mathcal{S}^c, \chi^c, \rho^c),$$

and the canonical directed simplicial model *for* LLB *is a pair* $\mathcal{C}^c = (\mathbb{C}^c, L^c)$*, such that:*

- $\mathcal{V}^c = \{\Gamma_a \mid \Gamma \in \mathbb{G} \text{ and } a \in \mathsf{Ag}\}$*;*
- $S^c = \bigcup_{\Gamma \in \mathbb{G}} \mathsf{Pow}(X_\Gamma) \setminus \{\emptyset\}$*;*
- $\chi^c(\Gamma_a) = a$*;*
- $\rho^c(\Gamma_a, Y) = \begin{cases} 1 & \text{if there exists } \Delta \in \mathbb{G} \text{ with } Y \subseteq X_\Delta \text{ and } \Gamma \setminus \Box_a \subseteq \Delta. \\ 0 & \text{otherwise.} \end{cases}$
- $p \in L^c(\Gamma_a)$ *if and only if* $p \in \Gamma_a$.

Using the following observations and lemmas, we prove that $\mathcal{C}^c$ is indeed a directed simplicial model.

Lemma 32. $(\mathcal{V}^c, \mathcal{S}^c)$ *is a simplicial complex.*

Proof. $\mathcal{S}^c$ is the union of powersets that do not contain $\{\emptyset\}$. Hence, it is closed under non-empty subsets and $(\mathcal{V}^c, \mathcal{S}^c)$ is a simplicial complex. □

The following observation follows directly from the definition of χ^c.

Observation 1. χ^c *is a proper coloring.*

Since every facet of $\mathbb{C}^c$ contains at most $|\mathsf{Ag}|$ elements, and $\chi(X_\Gamma) = |\mathsf{Ag}|$ for all $\Gamma \in \mathbb{G}$, we obtain the following corollary from Lemma 30.

Corollary 3. *$\mathcal{F}((\mathbb{C}) = \{X_\Gamma \mid \Gamma \in \mathbb{G}\}$ and $\mathbb{C}$ is pure.*

As a consequence of Lemma 31 we obtain the following corollary:

Corollary 4. *If $\rho^c(\Gamma_a, X_\Delta) = 1$, then $\Gamma_a \in X_\Delta$.*

By definition, $\rho^c(\Gamma_a, Y) = 1$ requires the existence of a facet Y containing Γ_a. Therefore, we obtain the following observation.

Observation 2. *ρ^c is a possibility function on $(\mathcal{V}^c, \mathcal{S}^c)$.*

Lemma 33. *$\mathcal{C}^c$ is a directed simplicial model.*

Proof. By Lemma 32 and Observation 1 and 2, $\mathbb{C}^c$ is a directed simplicial complex. Finally, by Definition 41, L^c is a local labeling. Thus, $\mathcal{C}^c$ is a directed simplicial model. □

The next lemma states that if an agent considers a facet possible, then all its beliefs must be true in that facet. More formally, if $\Box_a\phi \in \Gamma$ and $\rho^c(\Gamma_a, X_\Delta) = 1$, then $\phi \in \Delta$.

Lemma 34. *For all $\Delta, \Gamma \in \mathbb{G}$, $\rho^c(\Gamma_a, X_\Delta) = 1$ implies $\Gamma \setminus \Box_a \subseteq \Delta$.*

Proof. Assume that $\rho^c(\Gamma_a, X_\Delta) = 1$. By Definition 42, there exists a set $\Omega \in \mathbb{G}$ such that $X_\Delta \subseteq X_\Omega$ and $\Gamma \setminus \Box_a \subseteq \Omega$. By Corollary 3, both X_Δ and X_Ω are facets. Since facets are maximal under inclusion, it holds that $X_\Delta = X_\Omega$. Finally, by Lemma 30, we find that $\Delta = \Omega \supseteq \Gamma \setminus \Box_a$ as desired. □

The last ingredient for showing the truth lemma is proving that if an agent does not believe ϕ in X_Γ, then $\neg\phi$ is consistent with its belief in X_Γ.

Lemma 35. *Let $\Gamma \in \mathbb{G}$ such that $\Box_a\phi \notin \Gamma$ for some ϕ. The set*

$$\Gamma \setminus \Box_a \cup \{\neg\phi\}$$

is consistent.

Proof. The claim follows immediately from LLB being a normal modal logic. Indeed, assume that $\Gamma \setminus \Box_a \cup \{\neg\phi\}$ is inconsistent. Thus, there exist formulas $\phi_1, \ldots, \phi_m \in \Gamma \setminus \Box_a$ and $(\phi_1 \wedge \ldots \wedge \phi_m) \to \phi$ is derivable. We show that this implies that $\Box_a\phi \in \Box_a\Gamma$:

1. $(\phi_1 \wedge \cdots \wedge \phi_m) \to \phi$ — Derivable by assumption.

2. $\Box_a((\phi_1 \wedge \cdots \wedge \phi_m) \to \phi)$ $\Box_a$-Nec.

3. $\Box_a\phi_1 \wedge \ldots \wedge \Box_a\phi_m$ Elements of $\Box_a\Gamma$.

4. $\Box_a(\phi_1 \wedge \ldots \wedge \phi_m)$ Distribution of $\Box_a$.

5. $\Box_a\phi$ K and MP (2,4).

Therefore, if $(\phi_1, \ldots, \phi_m) \to \phi$ is derivable, then $\Box_a\phi \in \Box_a\Gamma$, whenever $\Box_a\phi_i \in \Box_a\Gamma$ for $i = 1, \ldots, m$. However, since we assumed $\Box_a\phi \notin \Box_a\Gamma$, this is a contradiction and $\Gamma \setminus \Box_a \cup \{\neg\phi\}$ is consistent. □

Given Lemmas 34 and 35, the proof of the truth lemma (Lemma 36) is straightforward.

Lemma 36 (Truth Lemma)**.** *Let $\mathcal{C}^c$ be the canonical model for* LLB*. For each $\Gamma \in \mathbb{G}$ and each formula $\phi \in \mathcal{L}$ we have*

$$\phi \in \Gamma \textit{ iff } \mathcal{C}^c, X_\Gamma \Vdash^{\mathsf{loc}}_\sigma \phi.$$

Proof. The proof of the truth lemma can be carried out in the usual way by induction over the length of ϕ. We only show the case for formulas of the form $\phi = \Box_a\psi$:

- $\Longrightarrow$: let $\Box_a\psi \in \Gamma$. By Lemma 34, it holds that $\psi \in \Delta$ for every accessible world X_Δ. By the inductive hypothesis, we obtain $\mathcal{C}^c, X_\Delta \Vdash^{\mathsf{loc}}_\sigma \psi$. Thus, by the definition of truth, it holds that $\mathcal{C}^c, X_\Gamma \Vdash^{\mathsf{loc}}_\sigma \Box_a\psi$.

- $\Longleftarrow$: we show the contrapositive. Assume that $\Box_a\psi \notin \Gamma$. We need to show that there exists a world X_Δ such that:
  - $\rho^c(\Gamma_a, X_\Delta) = 1$;
  - $\mathcal{C}^c, X_\Delta \nVdash^{\mathsf{loc}}_\sigma \psi$.

  By Lemma 35, the set $\Gamma \setminus \Box_a \cup \{\neg\psi\}$ is consistent. By Lemma 29, we can extend this set to a maximal consistent set Δ. Further, by the definition of ρ^c, it holds that $\rho^c(\Gamma_a, X_\Delta) = 1$. Finally, by the induction hypothesis, we obtain $\mathcal{C}^c, X_\Delta \nVdash^{\mathsf{loc}}_\sigma \psi$ and thus, $\mathcal{C}^c, X_\Gamma \nVdash^{\mathsf{loc}}_\sigma \Box_a\psi$. □

We can finally conclude the statement of completeness of LLB with respect to directed simplicial models.

Theorem 3. *For all $\phi \in \mathcal{L}$, it holds that*

$$\Vdash^{\mathsf{loc}}_\sigma \phi \textit{ implies } \vdash^{\mathsf{LLB}} \phi.$$

Given the back-and-forth transformation to proper relational models with a local valuation, defined in Section 4.2, LLB is also sound and complete with respect to the latter.

Corollary 5. *The logic* LLB *is sound and complete with respect to the class of local relational models based on proper, transitive, and euclidean frames.*

4.4 Relation to Impure Simplicial Complexes

We now demonstrate that restricting ourselves to pure complexes does not reduce expressivity when compared to global simplicial models (see Definition 6). The *pure extension* $\mathbb{C}^p = (\mathcal{V}^p, \mathcal{S}^p)$ of a possibly impure simplicial complex $\mathbb{C} = (\mathcal{V}, \mathcal{S})$ is defined in Definition 43. It is obtained by adding vertices to facets such that each facet has dimension $|\mathsf{Ag}| - 1$. Thus, it holds that $\mathcal{S} \subseteq \mathcal{S}^p$. Next, Definition 44 shows how to color the added vertices, and introduces a possibility function ρ^p on $\mathbb{C}^p$. The directions are chosen in such a way that the indistinguishability relation of agents in $\mathbb{C}$ remain unchanged. If v is a vertex that was added to a facet X while extending $\mathbb{C}$, then $\rho(v, Y) = 0$, for all faces $Y \in \mathcal{S}^p$. This reflects that the agent is present, but has inconsistent beliefs. Based on this method, we can construct a *corresponding directed simplicial model* (see Definition 45) for every global simplicial model. This is achieved by extending the underlying complex, coloring it accordingly, adding directions as described, and choosing a local labeling that mirrors the labeling of the global simplicial model. Since the extended simplicial complex is pure, we do not run into any problems regarding formulas not being defined.

Definition 43 ($\mathbb{C}^p$)**.** *Let* $\mathbb{C} = (\mathcal{V}, \mathcal{S})$ *be a simplicial complex. Its* pure extension $\mathbb{C}^p = (\mathcal{V}^p, \mathcal{S}^p)$ *is the pure simplicial complex such that*

1. *for all* $X \in \mathcal{F}(\mathbb{C}^p)$ *it holds that* $|X| = |\mathsf{Ag}|$*;*
2. $\mathcal{V} \subseteq \mathcal{V}^p$ *and* $\mathcal{S} \subseteq \mathcal{S}^p$*;*
3. *there exists a bijection:*
$$b : \mathcal{F}(\mathbb{C}) \to \mathcal{F}(\mathbb{C}^p)$$
such that for all $X, Y \in \mathcal{F}(\mathbb{C})$ *and* $v \in \mathcal{V}$*:*
$$v \in X \cap Y \text{ if and only if } v \in b(X) \cap b(Y).$$

Given a simplicial complex $\mathbb{C}$, its pure extension $\mathbb{C}^p$ is constructed by adding vertices to facets with a lower dimension. Vertices that are added this way are only connected to vertices of the facet to which they were added.

Definition 44 ($\mathbb{C}^{\mathsf{dir}}$)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$ be a chromatic simplicial complex. The corresponding directed chromatic simplicial complex $\mathbb{C}^{\mathsf{dir}} = (\mathcal{V}^p, \mathcal{S}^p, \chi^p, \rho^p)$ is defined as follows:*

- *$\mathbb{C}^p = (S^p, \mathcal{V}^p)$ is the pure extension of $\mathbb{C}$;*
- *$\chi^p : \mathcal{V}^p \to \mathsf{Ag}$ is a proper coloring such that $\chi(v) = a$ implies $\chi^p(v) = a$ for all $v \in \mathcal{V}$;*
- *$\rho^p : \mathcal{V}^p \times S^p \to \{0, 1\}$ is a function such that:*
  1. *$\rho^p(v, X) = 1$ if and only if there exists $Y \in \mathcal{F}(\mathbb{C}^p)$ with $X \subseteq Y$ and $v \in Y \cap b^{-1}(Y)$;*
  2. *$\rho^p(v, X) = 1$ implies $\rho^p(v, Y) = 1$, for all $Y \subseteq X$.*

Definition 45 ($\mathcal{C}^{\mathsf{dir}}$)**.** *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$ be a chromatic simplicial complex and let $\mathcal{C} = (\mathbb{C}, L)$ be a global simplicial model. The corresponding directed simplicial model is a pair $\mathcal{C}^{\mathsf{dir}} = (\mathbb{C}^{\mathsf{dir}}, L^p)$, where L^p is a labeling such that:*

$$p_a \in L(X) \text{ if and only if } p_a \in L^p(b(X)_a) \text{ for all } a \in \mathsf{Ag} \text{ and } X \in \mathcal{F}(\mathbb{C}).$$

The definition of ρ^p implies that $\rho^p(v, v) = 0$ whenever $v \in \mathcal{V}^p \setminus \mathcal{V}$. Hence,

$$\rho^p(v, X) = 1 \text{ implies } v \in \mathcal{V}. \tag{4.1}$$

It is straightforward to verify that ρ^p is a possibility function on $\mathbb{C}^p$, and that χ^p as well as L^p are well-defined. Thus, the corresponding directed simplicial model $\mathcal{C}^{\mathsf{dir}}$ is indeed a directed simplicial model. Example 7 shows the construction.

Example 7. *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$ be a simplicial complex, and let $\mathcal{C} = (\mathbb{C}, L)$ be a global simplicial model. Figure 4.6 depicts $\mathbb{C}$ on the left and its colored pure extension $\mathbb{C}^p$ on the right. To construct the pure extension of $\mathbb{C}$ only one vertex has to be added to $Y \in \mathcal{F}(\mathbb{C})$. The corresponding directed simplicial complex $\mathbb{C}^{\mathsf{dir}}$ is obtained by coloring $\mathbb{C}^p$ and assigning directions. Since Y is missing a vertex with color c, the vertex v is colored accordingly. Directions are as follows:*

- *$\rho^p(a, b(X)) = \rho^p(b, b(X)) = \rho^p(c, b(X)) = 1$;*
- *$\rho^p(a, b(Y)) = \rho^p(b', b(Y)) = 1$;*
- *$\rho^p(c', b(Y)) = 0$.*

The corresponding directed simplicial model $\mathcal{C}^{\mathsf{dir}} = (\mathbb{C}^{\mathsf{dir}}, L^p)$ is obtained by choosing a labeling L^p that agrees with L. That is, if $p \in \mathsf{P}_a$ is assigned to a facet $X \in \mathcal{F}(\mathbb{C})$ by L, then L^p assigns p to the corresponding a-vertex of the corresponding directed complex $\mathbb{C}^{\mathsf{dir}}$.

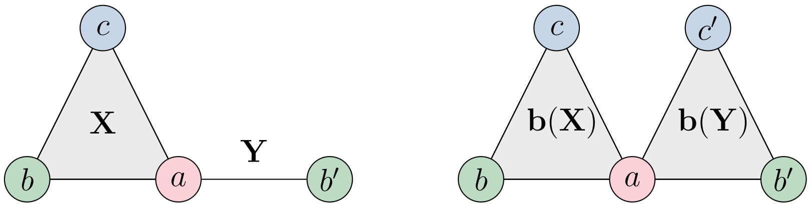

Figure 4.6: An impure simplicial complex (left) and the corresponding colored pure extension (right).

Finally, Lemma 37 shows the pointwise equivalence between global simplicial models and their corresponding directed simplicial models.

Lemma 37. *Let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \chi)$ be a simplicial complex, and let $\mathcal{C} = (\mathbb{C}, L)$ be a global simplicial model. Consider the corresponding directed simplicial complex $\mathbb{C}^{\mathsf{dir}} = (\mathcal{V}^p, \mathcal{S}^p, \chi^p, \rho^p)$, as well as the corresponding directed simplicial model $\mathcal{C}^{\mathsf{dir}} = (\mathbb{C}^{\mathsf{dir}}, L^p)$. It holds that for all $\phi \in \mathcal{L}$:*

$$\mathcal{C}, X \Vdash_\sigma \phi \text{ if and only if } \mathcal{C}^{\mathsf{dir}}, b(X) \Vdash_\sigma^{\mathsf{loc}} \phi.$$

Proof. We show the claim by induction on the length of ϕ:

1. Let $\phi \equiv p$ for some $p \in \mathsf{P}_a$, where $a \in \mathsf{Ag}$. By Definition 45:

$$\mathcal{C}, X \Vdash_\sigma p \text{ iff } p \in L(X) \text{ iff } p \in L^p(b(X)_a) \text{ iff } \mathcal{C}^{\mathsf{dir}}, b(X) \Vdash_\sigma^{\mathsf{loc}} p.$$

2. Let $\phi \equiv \neg\psi$; then the claim follows by the induction hypothesis.

3. Let $\phi \equiv \psi \wedge \psi'$; then the claim follows by the induction hypothesis.

4. Let $\phi \equiv \Box_a \psi$. We show the direction from left to right. The other direction is similar. If $\mathcal{C}, X \Vdash_\sigma \Box_a \bot$, then X does not contain an a-vertex, and we find that $\mathcal{C}^{\mathsf{dir}}, b(X) \Vdash_\sigma^{\mathsf{loc}} \Box_a \bot$, because of (4.1). Thus, we assume that X contains an a-vertex. By construction of $\mathbb{C}^{\mathsf{dir}}$, there exist facets that are considered possible by a in $b(X)$. Let $b(Y) \in \mathcal{F}(\mathbb{C}^{\mathsf{dir}})$ such that $\rho^p(b(X)_a, b(Y)) = 1$. Since ρ^p is a possibility function, it holds that $b(X)_a \in b(Y)$. By Definition 44, $b(X)_a = X_a$, and by Definition 43, we have that $X_a \in Y$ because $X_a \in b(X) \cap b(Y)$. By assumption, it holds that $\mathcal{C}, Y \Vdash_\sigma \psi$, and $\mathcal{C}^{\mathsf{dir}}, b(Y) \Vdash_\sigma^{\mathsf{loc}} \psi$ follows by the induction hypothesis. Therefore, $\mathcal{C}^{\mathsf{dir}}, b(X) \Vdash \Box_a \psi$, as desired.

□

4.5 Conclusion and Outlook

This chapter presented pure directed simplicial complexes as models for belief. We introduced the logic of local belief LLB, and proved it to be sound and complete with respect to simplicial models based on these structures. The logic LLB extends multi-agent K45 with axioms ensuring that agents believe the actual values of their own local propositional variables. We proved that directed simplicial models are equivalent to local relational models based on a proper, transitive, and euclidean frame. Additionally, we showed that impure global simplicial models can be expressed as directed simplicial models. However, this does not hold the other way around because directed simplicial complexes capture a less constrained accessibility relation. Thus, the presented directed simplicial models can be considered a generalization of global simplicial models.

A natural next step is to direct augmented simplicial complexes as mentioned in Remark 6, which would allow us to reason in a more detailed manner about the local views of agents. Indeed, let $\mathcal{V} = \{a, b, c\}$ be a set of vertices that are colored as indicated by their names, and let $\mathbb{S} = (\mathcal{V}, \mathcal{S})$ be the solid triangle spanned by those vertices. The only facet of $\mathbb{S}$ is denoted with X and we refer to its ac-edge as Y. Furthermore, let $\mathbb{C} = (\mathcal{V}, \mathcal{S}, \rho, \{X, Y\})$ be an augmented complex. We can direct $\mathbb{C}$ in many ways. Specifically, we will look at two directed versions $\mathbb{C}_1$ and $\mathbb{C}_2$ shown in Figure 4.7. The thick ac-edge indicates that this edge belongs to the set W. Both complexes are structurally equivalent. However, the difference between $\mathbb{C}_1$ and $\mathbb{C}_2$ is that agent a assigns 1 to Y in $\mathbb{C}_2$, and 0 in $\mathbb{C}_1$. In both complexes, agent a does not deem X possible. Therefore, in $\mathbb{C}_1$, agent a has inconsistent beliefs since it does not consider a world possible, whereas in $\mathbb{C}_2$ it has consistent beliefs because it considers Y possible. Thus, we can distinguish between the two structures. If we did not allow arbitrary faces to be global states, we could not distinguish between $\mathbb{C}_1$ and $\mathbb{C}_2$, because agent a has inconsistent beliefs in both.

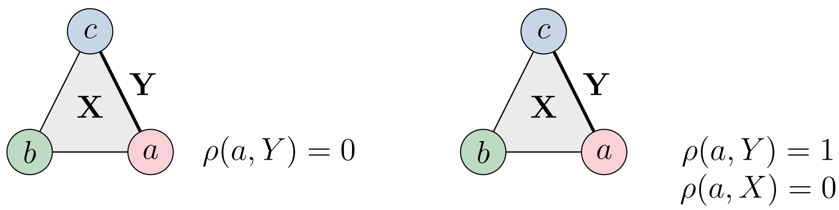

Figure 4.7: The directed augmented simplicial complexes $\mathbb{C}_1$ (left) and $\mathbb{C}_2$ (right) only differ in the directions that agent a assigns to the ac-edge.

Due to augmented complexes possibly being impure, we must adapt our framework to the three-valued semantics of Randrianomentsoa et al. [46]. It seems plausible that their results transfer to our setting, since directions only restrict the accessibility relation to a subset of the indistinguishable states. Moreover, our construction of the canonical model closely follows the one used for three-valued semantics. Thus, already having gained insights about directions on arbitrary faces is valuable.

Lastly, having an intuitive and simple solution to model beliefs allows us to investigate belief dynamics on simplicial models, thereby addressing the next open question posed by Castañeda et al. [17]: modeling agents capable of lying and deceiving others. Such agents occur in distributed systems where agents may deviate arbitrarily from the protocol (cf. Cachin et al. [13]). In such settings, agents may need to act based on their beliefs.

5 Multi-simplicial Models

Chapter Organization. This chapter examines semi-simplicial sets as models for group knowledge. The concept of group knowledge in this context extends the traditional notion of distributed knowledge and is referred to as synergistic knowledge. Instead of studying groups of agents, we examine agent patterns, which capture synergies among group members. Section 5.1 introduces an explicit representation of semi-simplicial sets, as well as an indistinguishability relation for agent patterns. In Section 5.2, the logic Syn for synergistic knowledge is presented. Section 5.3 introduces multi-simplicial models, and shows soundness of Syn with respect to them. Applications of synergistic knowledge to distributed computing are showcased in Section 5.4. Next, completeness of Syn with respect to multi-simplicial models is established in Section 5.5. An alternative interpretation of agent patterns is discussed in Section 5.6 . Finally, Section 5.7 concludes the chapter and outlines possible future work.

So far, we have only considered simplicial complexes that forbid parallel faces. Simplicial complexes that may contain repeated faces are called *semi-simplicial sets*. These differ from *simplicial sets* in that they do not admit faces containing multiple copies of the same vertex. For an introduction to simplicial sets, as well as a discussion of the differences between semi-simplicial sets and simplicial sets, we refer to Friedman [22].

Figure 5.1 shows a simplicial complex $\mathbb{C}_1$ and a semi-simplicial set $\mathbb{C}_2$. The complex $\mathbb{C}_1$ is obtained by gluing together two solid triangles along their ab-edge, whereas the semi-simplicial set $\mathbb{C}_2$ is formed by gluing them together at their a-vertex and b-vertex. As a result, they do not have the same ab-boundary in $\mathbb{C}_2$.

Interestingly, for a group of agents G, the standard notion of indistinguishability between two faces X and Y, based on

$$G \subseteq \chi(X \cap Y), \tag{5.1}$$

does not offer the necessary means to differentiate between simplicial complexes and semi-simplicial sets. Indeed, despite the different structure, every group of agents has the same indistinguishability relation in $\mathbb{C}_1$ and $\mathbb{C}_2$. For example, the group $G = \{a, b\}$ cannot distinguish between X and Y in either $\mathbb{C}_1$ or $\mathbb{C}_2$, because the two triangles intersect at the a- and b-vertex. This raises the question of

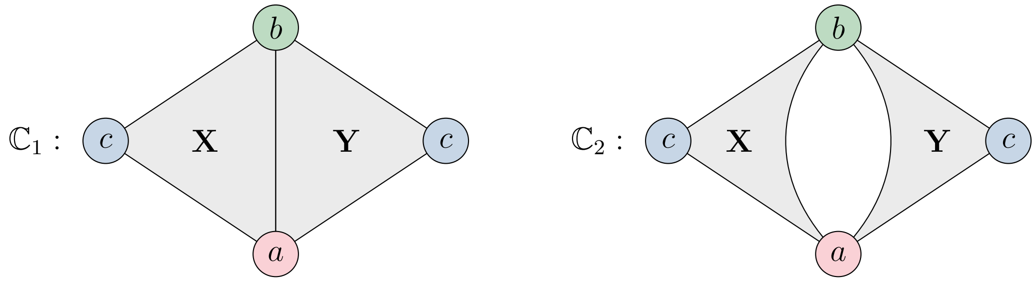

Figure 5.1: A difference between simplicial complexes and semi-simplicial sets is that the latter can contain multiple edges connecting the same vertices.

whether an indistinguishability relation based on higher order connectivity could describe semi-simplicial sets better.

The first to study semi-simplicial sets as models for group knowledge were Goubault et al. [26]. To distinguish between complexes such as $\mathbb{C}_1$ and $\mathbb{C}_2$, the authors base the notion of a group's indistinguishability relation on the respective boundary of simplices. Informally, for a group $G \subseteq \mathsf{Ag}$, the new notion of indistinguishability between simplices is equivalent to:

$$S \text{ and } T \text{ have the same } G\text{-boundary.} \tag{5.2}$$

We denote the simplices containing the facets X and Y in Figure 5.1 with S_X and S_Y. In $\mathbb{C}_1$, we see that S_X and S_Y have the same ab-boundary because the edge between the a-vertex and the b-vertex belongs to both S_X and S_Y. Thus, agents a and b together cannot distinguish S_X from S_Y according to (5.2). On the other hand, the semi-simplicial set $\mathbb{C}_2$ contains two edges between the a-vertex and the b-vertex, one belonging to S_X and the other one belonging to S_Y. Therefore, the ab-boundary of S_X is not equal to the ab-boundary of S_Y. Consequently, agents a and b together are able to distinguish S_X from S_Y based on (5.2). Condition (5.2) is equivalent to the usual notion of indistinguishability if we consider simplicial complexes without parallel faces.

Goubault et al. [26] observe that each of their semi-simplicial set models is equivalent modulo bisimulation to a model without parallel faces. For example, in Figure 5.2, the semi-simplicial set model $\mathcal{C}_3$ is bisimilar to the generalized global simplicial model $\mathcal{C}_4$, because in both models a and b together can always determine the actual state. Consequently, the group $G = \{a, b\}$ knows whether p or $\neg p$ holds at all times. Goubault et al. [26] indicate that other logics may describe semi-simplicial sets more adequately, and pose the following question: *Could we define a logic that is able to capture better the global geometry of the model?* Answering that question is the main concern of this chapter.

We propose the epistemic attitude of *synergistic knowledge*, which is able to describe scenarios where a group of agents can know more than just the consequences of their pooled knowledge. Instead of reasoning about what a group of agents knows, we unfold relations within a group. Thus, the synergistic knowledge of two seemingly equal groups need not be the same, because the agents can have different relations in both groups. We refer to this as *synergy*. Moreover, unlike the standard category-theoretic definition of semi-simplicial sets, we employ an explicit representation of semi-simplicial sets.

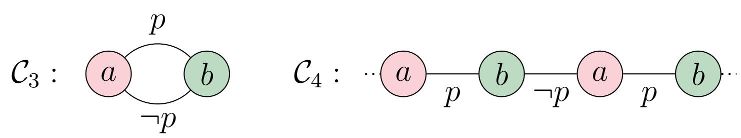

Figure 5.2: The semi-simplicial set model $\mathcal{C}_3$ and the infinite generalized global simplicial model $\mathcal{C}_4$ are logically equivalent according to the notion of indistinguishability (5.2).

Simplicial complexes with parallel faces are represented by indexing faces. For example, to describe $\mathbb{C}_2$, the indexed vertices are $(\{a\},1),(\{b\},1),(\{c\},1)$, and$(\{c\},2)$; the indexed edges are

$$(\{a,b\},1),(\{a,b\},2),(\{a,c\},1),(\{a,c\},2),(\{b,c\},1),(\{b,c\},2);$$

and the indexed triangles are $(\{a,b,c\},1),(\{a,b,c\},2)$. With this new type of faces comes the need for a different representation of simplicial complexes. This is because we cannot tell if $(\{c\},1)$ is contained in $(\{a,b,c\},1)$ or $(\{a,b,c\},2)$. In order to keep track of which indexed faces are contained in which indexed faces, we specify the simplicial structure induced by an indexed face. We call this structure the *explicit simplex*. The semi-simplicial set $\mathbb{C}_2$ can be represented as the set containing the two explicit simplices:[1]

$$S_1 = \left\{ \begin{matrix} (\{a,b,c\},1) \\ (\{a,b\},1),\ldots,(\{b,c\},1) \\ (\{a\},1),(\{b\},1),(\{c\},1) \end{matrix} \right\} \text{ and } S_2 = \left\{ \begin{matrix} (\{a,b,c\},2) \\ (\{a,b\},2),\ldots,(\{b,c\},2) \\ (\{a\},1),(\{b\},1),(\{c\},2) \end{matrix} \right\}.$$

Under this representation, it is evident that $(\{c\},1)$ is only contained in S_1. Further, it is clear that the explicit simplices S_1 and S_2 share an a-vertex and a b-vertex, but not an ab-edge. Due to the different representation of our structures, we

[1] In the previous chapters, we could also have represented a simplicial complex as a set of simplices. However, in the case of non-indexed faces, this was not necessary, because a face determines its subfaces.

refrain from calling them semi-simplicial sets and refer to them as *multi-simplicial complexes* instead.

We accomplish to differentiate between $\mathcal{C}_3$ and $\mathcal{C}_4$ by reasoning not about what a group of agents $\emptyset \neq G \subseteq \mathsf{Ag}$ can distinguish, but rather about what an *agent pattern* $\emptyset \neq G \subseteq \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}$ is able to distinguish. Elements of G are thought of as relations (or synergies) among agents. Hence, G represents a group of agents as well as their synergies. For an agent pattern G, we base indistinguishability between two explicit simplices S and T on:

$$G \subseteq (S \cap T)^{\circ}, \tag{5.3}$$

where the elements of $(S \cap T)^{\circ}$ are the colors contained in the indexed faces shared by S and T. For example, in $\mathbb{C}_1$ we have that $(S_X \cap S_Y)^{\circ}$ is the set $\{\{a, b\}, \{a\}, \{b\}\}$. Our new notion of indistinguishability (5.3) has the same structure as (5.1), which is possible due to our explicit representation of semi-simplicial sets. Based on (5.3), the synergistic knowledge of agent patterns differs in $\mathcal{C}_3$ and $\mathcal{C}_4$. Indeed, since $\mathcal{C}_3$ contains two parallel ab-edges, the agent pattern $G = \{\{a\}, \{b\}\}$ cannot determine whether p or $\neg p$ holds in $\mathcal{C}_3$. Hence, in $\mathcal{C}_3$ the agents need additional synergy to identify the actual state. In $\mathcal{C}_4$, however, the pattern G can always identify the current state, and thus knows if p or $\neg p$. This chapter introduces our simplicial structures, presents the logic for synergistic knowledge Syn, and examines its properties.

5.1 Multi-simplicial Structures

This section formally introduces multi-simplicial structures and shows some key properties. Let $\mathcal{I}$ be a set of indices. The set of all indexed faces is:

$$\mathsf{Agsi} = \{(A, i) \mid \emptyset \neq A \subseteq \mathsf{Ag} \text{ and } i \in \mathcal{I}\}.$$

A coloring is unnecessary, since an agent's identifier is embedded in an indexed face. For $a \in \mathsf{Ag}$, the pair $(\{a\}, i)$ can be thought of as agent a in local state i. Let $S \subseteq \mathsf{Agsi}$. An element $(A, i) \in S$ is *maximal in* S if and only if:

$$\forall (B, j) \in S. |A| \geq |B|.$$

Definition 46 (Explicit Simplex). *Let $\emptyset \neq S \subseteq \mathsf{Agsi}$. S is an* explicit simplex *if and only if:*

- $\mathsf{S1}$*: The maximal element is unique, i.e.,*

  $$\textit{if } (A, i) \in S \textit{ and } (B, j) \in S \textit{ are maximal in } S, \textit{ then } A = B \textit{ and } i = j.$$

  The maximal element of S is denoted as $\max(S)$.

$\mathsf{S2}$*: S is uniquely downwards closed, i.e., for $(B,i) \in S$ and $\emptyset \neq C \subseteq B$:*

$$\exists! j \in \mathcal{I}.(C,j) \in S, \text{ where } \exists! j \text{ means that there exists exactly one } j.$$

$\mathsf{S3}$*: S contains nothing else, i.e.,*

$$(B,i) \in S \text{ and } (A,j) = \max(S) \text{ implies } B \subseteq A.$$

Definition 47 (Multi-simplicial Complex)**.** *Let $\mathbb{C}$ be a set of explicit simplices. $\mathbb{C}$ is a* multi-simplicial complex *if and only if the following condition is met:*

C*: For every two explicit simplices $S, T \in \mathbb{C}$, if there exist $A \subseteq \mathsf{Ag}$ and $i \in \mathcal{I}$ with $(A,i) \in S$ and $(A,i) \in T$, then:*

$$\text{for all } \emptyset \neq B \subseteq A \text{ and all } j \quad (B,j) \in S \quad \text{iff} \quad (B,j) \in T.$$

When clear from context, we use the terms *multi-simplicial complex* and *complex* interchangeably. However, the term *simplicial complex* is reserved and always corresponds to the notion introduced in Definition 5. We will also refer to *explicit simplices* as *simplices* when no ambiguity arises.

We require condition C to ensure that a complex cannot contain two different simplices with the same maximal indexed face (Lemma 38).

Lemma 38. *Let $\mathbb{C}$ be a complex and $S, T \in \mathbb{C}$. We find that:*

$$\max(S) = \max(T) \text{ implies } S = T.$$

Proof. We only show $S \subseteq T$, because the other direction is symmetric. Let $(A,i) = \max(S)$ and assume that $(B,j) \in S$. Because of $\mathsf{S3}$, we have $B \subseteq A$. By Condition C, we conclude $(B,j) \in T$. $\square$

Whenever the context is clear, we abbreviate $(\{a_1, \ldots, a_n\}, i)$ as $a_1 \ldots a_n i$. Additionally, we assume $\mathcal{I} = \mathbb{N}$ for examples and use a row (or a mixed row-column) notation for explicit simplices. For example, the complex $\mathbb{C}_1$ in Figure 5.1 is represented as

$$\mathbb{C}_1 = \left\{ \left\{ \begin{matrix} abc0 \\ ab0, ac0, bc0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc1 \\ ab0, ac1, bc1 \\ a0, b0, c1 \end{matrix} \right\} \right\},$$

where both elements of $\mathbb{C}_1$ share the element $ab0$. The complex $\mathbb{C}_2$, on the other hand, is represented as

$$\mathbb{C}_2 = \left\{ \left\{ \begin{matrix} abc0 \\ ab0, ac0, bc0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc1 \\ ab1, ac1, bc1 \\ a0, b0, c1 \end{matrix} \right\} \right\},$$

where we can clearly distinguish between the elements $ab0$ and $ab1$. Whenever we refer to a simplex within a complex, we write $\langle a_1, \dots, a_n i\rangle$ to denote the simplex whose maximal element is $(\{a_1, \dots, a_n\}, i)$. Lemma 38 ensures that this notation is well-defined.

Definition 48 (Indistinguishability)**.** *Let $S \subseteq \mathsf{Agsi}$, we define:*

$$S^\circ = \{A \mid \exists i \in \mathcal{I}.(A, i) \in S\}.$$

An agent pattern G is a non-empty subset of $\mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\}$. The set of all agent patterns is AP. We say that G cannot distinguish between two explicit simplices S and T if and only if $G \subseteq (S \cap T)^\circ$.

Definition 49 ($\sim_G^{\mathbb{C}}$)**.** *Let $\mathbb{C}$ be a complex. For $G \in \mathsf{AP}$ we define:*

$$\sim_G^{\mathbb{C}} = \{(S, T) \mid S, T \in \mathbb{C} \textit{ and } G \subseteq (S \cap T)^\circ\},$$

and write $S \sim_G^{\mathbb{C}} T$ if and only if $(S, T) \in \sim_G^{\mathbb{C}}$.

We motivate our choice of $\sim_G^{\mathbb{C}}$ as an indistinguishability relation with Lemma 39 and Lemma 40.

Lemma 39. *For a complex $\mathbb{C}$, the relation $\sim_G^{\mathbb{C}}$ is symmetric and transitive.*

Proof. Symmetry immediately follows from the fact that set intersection is commutative. To show transitivity, let $S, T, U \in \mathbb{C}$ such that $S \sim_G^{\mathbb{C}} T$ and $T \sim_G^{\mathbb{C}} U$, i.e.,

$$G \subseteq (S \cap T)^\circ \tag{5.4}$$

$$G \subseteq (T \cap U)^\circ \tag{5.5}$$

Let $A \in G$. Because of (5.4), there exists i with:

$$(A, i) \in S \quad \text{and} \quad (A, i) \in T. \tag{5.6}$$

By (5.5), there exists j with:

$$(A, j) \in T \quad \text{and} \quad (A, j) \in U. \tag{5.7}$$

From (5.6), (5.7), and Condition S2, we obtain $i = j$. Thus, by (5.6) and (5.7), we get $A \in (S \cap U)^\circ$. Since A was arbitrary in G, we conclude that $G \subseteq (S \cap U)^\circ$, i.e., $S \sim_G^{\mathbb{C}} T$. □

Given an agent pattern G, the set $G^\star$ (see Definition 50) consists of the singleton sets of agents contained in G. If S is a simplex and $G^\star \subseteq S^\circ$, then each agent a contained in G has a corresponding vertex $(\{a\}, i)$ in S.

Definition 50 ($G^\star$)**.** *Let $G \in \mathsf{AP}$ be an agent pattern, we define:*

$$G^\star = \{\{a\} \mid \exists A \in G \text{ and } a \in A\}.$$

Lemma 40. *Let $\mathbb{C}$ be a complex and consider a maximal set of simplices $\mathcal{S}_G \subseteq \mathbb{C}$ such that for any $S \in \mathcal{S}_G$ we have $G^\star \subseteq S^\circ$. The indistinguishability relation $\sim_G^{\mathbb{C}}$ is reflexive on $\mathcal{S}_G \times \mathcal{S}_G$ and empty otherwise.*

Proof. We first show reflexivity. Let $S \in \mathcal{S}_G$. For each $B \in G$, we have to show that $B \in (S \cap S)^\circ$, i.e.,

$$\text{there exists } i \text{ with } (B, i) \in S. \tag{5.8}$$

Let $(A, i) = \max(S)$, and let $b \in B$. Because of $G^\star \subseteq S^\circ$, there exists l such that $(\{b\}, l) \in S$. By $\mathsf{S3}$ we get $b \in A$. Since b was arbitrary in B, we get $B \subseteq A$. By $\mathsf{S2}$ we conclude that (5.8) holds and reflexivity is established.

We now show that $\sim_G^{\mathbb{C}}$ is empty otherwise. Let S be a simplex such that $G^\star \nsubseteq S^\circ$ and let T be an arbitrary simplex. Then there exists a, A with $a \in A \in G$ and $\{a\} \notin S^\circ$, i.e.,

$$\text{for all } i,\ (\{a\}, i) \notin S. \tag{5.9}$$

Suppose towards a contradiction that:

$$G \subseteq (S \cap T)^\circ. \tag{5.10}$$

Because of $A \in G$, we get $A \in (S \cap T)^\circ$. Hence $A \in S^\circ$, i.e., there exists l with $(A, l) \in S$. With $\mathsf{S2}$ and $\{a\} \subseteq A$, we find that there exists j with $(\{a\}, j) \in S$. This is a contradiction to (5.9). Thus, (5.10) cannot hold. □

Corollary 6. *Let $\mathbb{C}$ be a complex and let $\mathcal{S}_G$ be as in Lemma 40. It holds that $\sim_G^{\mathbb{C}}$ is an equivalence relation on $\mathcal{S}_G \times \mathcal{S}_G$.*

Lemma 41 and Lemma 42 are needed to show that adding synergies to an agent pattern makes it stronger in the sense that it can distinguish between more simplices (Lemma 43).

Lemma 41. *Let $\mathbb{C}$ be a complex and let $S, T \in \mathbb{C}$. Further, let $A \in (S \cap T)^\circ$ and $\emptyset \neq B \subseteq A$. We find $B \in (S \cap T)^\circ$.*

Proof. From $A \in (S \cap T)^\circ$, we obtain that there exists i such that $(A, i) \in S$ and $(A, i) \in T$. From $\mathsf{S2}$ we find that there exists j such that $(B, j) \in S$. Thus, by C, we get $(B, j) \in T$ and we conclude $B \in (S \cap T)^\circ$. □

Corollary 7. *Let $\mathbb{C}$ be a complex, and let G be an agent pattern. Further let $A, B \subseteq \mathsf{Ag}$ be such that $\emptyset \neq B \subseteq A \in G$. It holds that:*

$$\sim_{G \cup \{B\}}^{\mathbb{C}} = \sim_G^{\mathbb{C}}.$$

Lemma 42 (Anti-Monotonicity)**.** *Let $\mathbb{C}$ be a complex. For all agent patterns $G, H \in \mathsf{AP}$, it holds that $G \subseteq H$ implies $\sim_H^{\mathbb{C}} \subseteq \sim_G^{\mathbb{C}}$.*

Proof. Assume $G \subseteq H$. For any two simplices S and T with $S \sim_H^{\mathbb{C}} T$, we have $G \subseteq H \subseteq (S \cap T)^\circ$ by Definition 48 and hence, $S \sim_G^{\mathbb{C}} T$. □

Lemma 43. *Let $\mathbb{C}$ be a complex, and let $H_1, H_2, \ldots, H_n \subseteq \mathsf{Ag}$ be non-empty with $n \geq 2$. We have*

$$\sim_{\{H_1 \cup H_2, \ldots, H_n\}}^{\mathbb{C}} \subseteq \sim_{\{H_1, H_2, \ldots, H_n\}}^{\mathbb{C}}.$$

Proof. From Lemma 41 and Lemma 42 we find that

$$\sim_{\{H_1 \cup H_2, \ldots, H_n\}}^{\mathbb{C}} = \sim_{\{H_1 \cup H_2, H_1, H_2, \ldots, H_n\}}^{\mathbb{C}} \subseteq \sim_{\{H_1, H_2, \ldots, H_n\}}^{\mathbb{C}} .$$ □

In traditional relational semantics, distributed knowledge of a set of agents is modeled by the indistinguishability relation that is given by the intersection of the indistinguishability relations of the individual agents. Therefore, we call the property:

$$\bigcap_{B \in G} \sim_{\{B\}}^{\mathbb{C}} = \sim_G^{\mathbb{C}} \tag{SGK}$$

standard group knowledge.

Lemma 44. *Let $\mathbb{C}$ be a complex, and let G be an agent pattern. The relation $\sim_G^{\mathbb{C}}$ satisfies* (SGK)*.*

Proof. $(S,T) \in \bigcap_{B \in G} \sim_{\{B\}}^{\mathbb{C}}$ iff for each $B \in G$, we have $B \in (S \cap T)^\circ$ iff $G \subseteq (S \cap T)^\circ$ iff $S \sim_G^{\mathbb{C}} T$. □

The following lemma captures the intuition of distributed knowledge in our framework. The intersection of the individual indistinguishability relations of all agents of a non-empty group G corresponds to the agent pattern consisting of singleton sets for each agent of G.

Lemma 45. *Let $\emptyset \neq G \subseteq \mathsf{Ag}$ and $H = \bigcup_{a \in G} \{\{a\}\}$. We have:*

$$\bigcap_{a \in G} \sim_{\{\{a\}\}}^{\mathbb{C}} = \sim_H^{\mathbb{C}} .$$

Proof. $(S,T) \in \bigcap_{a \in G} \sim_{\{\{a\}\}}^{\mathbb{C}}$ iff for each $a \in G$, we have $\{a\} \in (S \cap T)^\circ$ iff (by the definition of H) $H \subseteq (S \cap T)^\circ$ iff $S \sim_H^{\mathbb{C}} T$. □

5.2 Syntax

The logic of synergistic knowledge is a normal modal logic that includes a modality $[G]$ for each agent pattern $G \in \mathsf{AP}$. Formulas of the language of synergistic knowledge $\mathcal{L}^{\mathsf{syn}}$ are inductively defined by the following grammar:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid [G]\phi$$

where $p \in \mathsf{Prop}$ and $G \in \mathsf{AP}$ is an agent pattern. The remaining Boolean connectives are defined as usual. In particular, we set $\bot = p \wedge \neg p$ for some fixed $p \in \mathsf{Prop}$, and we write $\mathsf{alive}(G)$ for $\neg[G]\bot$. If G is an agent pattern, then G^C denotes its complement, which is defined as:

$$G^C = \{H \in \mathsf{Pow}(\mathsf{Ag}) \setminus \{\emptyset\} \mid \nexists B \in G.H \subseteq B\}.$$

Moreover, we define:

$$\mathsf{dead}(G) = \bigwedge_{B \in G} \neg\mathsf{alive}(\{B\}).$$

Notice that $\mathsf{dead}(G) \not\equiv \neg\mathsf{alive}(G)$. Indeed, $\mathsf{dead}(G)$ expresses that for each $B \in G$, the pattern $\{B\}$ is dead, whereas $\neg\mathsf{alive}(G)$ is true if some $\{B\} \subseteq G$ is dead. The axiom system Syn consists of the following axioms:

$$\text{all propositional tautologies} \qquad \text{(Taut)}$$
$$[G](\phi \to \psi) \to ([G]\phi \to [G]\psi) \qquad \text{(K)}$$
$$\phi \to [G]\neg[G]\neg\phi \qquad \text{(B)}$$
$$[G]\phi \to [G][G]\phi \qquad \text{(4)}$$
$$\mathsf{alive}(G) \to ([G]\phi \to \phi) \qquad \text{(T)}$$
$$\mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \wedge \phi \to [G](\mathsf{dead}(G^C) \to \phi) \qquad \text{(P)}$$
$$\bigvee_{G \in \mathsf{AP}} \mathsf{alive}(G) \qquad \text{(NE)}$$
$$[G]\phi \to [H]\phi \quad \text{if } G \subseteq H \qquad \text{(Mono)}$$
$$[G \cup \{B\}]\phi \to [G]\phi \quad \text{if there exists } A \in G \text{ and } \emptyset \neq B \subseteq A \qquad \text{(Equiv)}$$
$$\mathsf{alive}(G) \wedge \mathsf{alive}(H) \to \mathsf{alive}(G \cup H) \qquad \text{(Union)}$$
$$\mathsf{alive}(G) \to \mathsf{alive}(\{A \cup B\}) \quad \text{if } A, B \in G \qquad \text{(Clo)}$$

and the inference rules modus ponens MP and $[G]$-necessitation $[G]$-Nec. We write $\vdash \varphi$ to denote that $\varphi \in \mathcal{L}^{\mathsf{syn}}$ can be deduced in the system Syn.

$$\frac{A \quad A \to B}{B} \ \text{(MP)} \qquad \frac{A}{[G]A} \ \text{([}G\text{]-Nec)}$$

Axiom NE states that there is always an alive agent pattern; Mono asserts that an agent pattern cannot know more than its supersets; Equiv ensures that adding already existing synergies to an agent pattern does not strengthen its knowledge; and Union and Clo enforce standard closure properties. Lastly, the axiom P reflects Lemma 38.

5.3 Simplicial Semantics

A complex $\mathbb{C}$ equipped with a labeling $L : \mathbb{C} \to \mathsf{Pow}(\mathsf{Prop})$ is a *multi-simplicial model*.

Definition 51 (Multi-simplicial Model)**.** *A* multi-simplicial model $\mathcal{C} = (\mathbb{C}, L)$ *is a pair such that:*

1. $\mathbb{C}$ *is a complex;*
2. $L : \mathbb{C} \to \mathsf{Pow}(\mathsf{Prop})$ *is a labeling.*

Given the resemblance of multi-simplicial models to global simplicial models (see Definition 6), we overload notation and use $\Vdash_\sigma$ for satisfiability on multi-simplicial models as well. We also predominantly use the letters S and T to refer to global states, rather than X and Y as in the previous chapters, to emphasize that global states are explicit simplices.

Definition 52 ($\Vdash_\sigma$)**.** *Let* $\mathcal{C} = (\mathbb{C}, L)$ *be a multi-simplicial model, and let* $S \in \mathbb{C}$ *be an explicit simplex. We define the relation* $\mathcal{C}, S \Vdash_\sigma \phi$ *by induction on* $\phi \in \mathcal{L}^{\mathsf{syn}}$*:*

$$
\begin{array}{lll}
\mathcal{C}, S \Vdash_\sigma p & \textit{iff} & p \in L(S) \\
\mathcal{C}, S \Vdash_\sigma \neg\phi & \textit{iff} & \mathcal{C}, S \nVdash_\sigma \phi \\
\mathcal{C}, S \Vdash_\sigma \phi \wedge \psi & \textit{iff} & \mathcal{C}, S \Vdash_\sigma \phi \textit{ and } \mathcal{C}, S \Vdash_\sigma \psi \\
\mathcal{C}, S \Vdash_\sigma [G]\phi & \textit{iff} & S \sim_G^{\mathbb{C}} T \textit{ implies } \mathcal{C}, T \Vdash_\sigma \phi, \textit{ for all } T \in \mathbb{C}.
\end{array}
$$

We write $\mathcal{C} \Vdash_\sigma \phi$, if $\mathcal{C}, S \Vdash_\sigma \phi$ for all $S \in \mathbb{C}$. A formula ϕ is σ-valid, denoted by $\Vdash_\sigma \phi$, if $\mathcal{C} \Vdash_\sigma \phi$ for all models $\mathcal{C}$. Whenever it is clear from the context we omit the subscript σ.

Remark 9. *Since a multi-simplicial complex is represented as a set of explicit simplices, it can contain two simplices* S *and* T *with* $S \subsetneq T$*. Consequently, a multi-simplicial model based on such a complex contains subsimplices as global states. In Chapter 3, we employed a different modeling approach to represent that certain faces of a simplicial complex are global states. Specifically, we used the approach by Goubault et al. [27], which augments a simplicial complex* $\mathbb{C} = (\mathcal{V}, \mathcal{S})$

with a set of worlds W such that $\mathcal{F}(\mathbb{C}) \subseteq W \subseteq \mathcal{S}$. However, this is not needed for multi-simplicial complexes because the information of the set W is contained in their definition.

Corollary 8 is an immediate consequence of the previously introduced Corollary 6. It relates the formula $\mathsf{alive}(G)$ to the structure of the underlying complex in the expected way.

Corollary 8. *Let $\mathcal{C} = (\mathbb{C}, L)$ be a simplicial model. We find:*

$$\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G) \quad \textit{iff} \quad S \sim_G^{\mathbb{C}} S.$$

Soundness of Syn with respect to simplicial models follows as usual. We present the proof of completeness in Section 5.5.

Theorem 4 (Soundness)**.** $\vdash \varphi$ *implies* $\Vdash_\sigma \varphi$.

Proof. We only show T, P, NE, Union, Clo, Mono, Equiv, and $[G]$-Nec. Let $\mathcal{C} = (\mathbb{C}, L)$ be an arbitrary multi-simplicial model.

1. T: Consider a simplex $S \in \mathbb{C}$ and assume that $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G)$ and $\mathcal{C}, S \Vdash_\sigma [G]\varphi$. By Corollary 8, we find $S \sim_G^{\mathbb{C}} S$ and thus $\mathcal{C}, S \Vdash_\sigma \varphi$ because $\mathcal{C}, S \Vdash_\sigma [G]\varphi$.

2. P: Consider a simplex $S \in \mathbb{C}$ with $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G) \land \mathsf{dead}(G^C) \land \phi$. By assumption, G must have a unique maximal element. Indeed, towards a contradiction, assume that it is not the case and G has the maximal elements $B_1, B_2, \ldots, B_n$. Furthermore, let:

$$B = \bigcup_{i=1}^{n} B_i.$$

   It is straightforward to verify that $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(\{B\})$. However, since $B \in G^C$, by assumption, it holds that $\mathcal{C}, S \Vdash_\sigma \mathsf{dead}(\{B\})$, which is a contradiction. Additionally, by $\mathsf{S3}$, the maximal element of G, say A, is the set of all agents alive in S. Let $T \in \mathbb{C}$ be such that $S \sim_G^{\mathbb{C}} T$ and assume $\mathcal{C}, T \Vdash_\sigma \mathsf{dead}(G^C)$. By transitivity of $\sim_G^{\mathbb{C}}$ and Corollary 8, we have that $\mathcal{C}, T \Vdash_\sigma \mathsf{alive}(G)$. Thus, by the same reasoning as before, A is the set of all alive agents for T as well, i.e., $\max(S) = (A, i)$ and $\max(T) = (A, i)$ for some $i \in \mathcal{I}$. Finally, by Lemma 38, we find that $S = T$, and thus $\mathcal{C}, T \Vdash_\sigma \phi$.

3. NE: Follows because simplices are not empty.

4. Union: Consider a simplex $S \in \mathbb{C}$ and assume that $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G)$ and $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(H)$. By Corollary 8 it holds that $S \sim_G^{\mathbb{C}} S$ and $S \sim_H^{\mathbb{C}} S$. By Lemma 44 we have

$$\sim_{G \cup H}^{\mathbb{C}} = \bigcap_{B \in G \cup H} \sim_{\{B\}}^{\mathbb{C}} = \left(\bigcap_{B \in G} \sim_{\{B\}}^{\mathbb{C}} \right) \cap \left(\bigcap_{B \in H} \sim_{\{B\}}^{\mathbb{C}} \right),$$

and thus $S \sim_{G \cup H}^{\mathbb{C}} S$ and $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G \cup H)$.

5. Clo: Consider a world $S \in \mathbb{C}$ and assume that $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(G)$, and let $A, B \in G$. By Lemma 41 we find $S \sim_{\{A\}}^{\mathbb{C}} S$ and $S \sim_{\{B\}}^{\mathbb{C}} S$, i.e., there exist $i, j \in \mathcal{I}$ such that $(A, i) \in S$ and $(B, j) \in S$. Furthermore, let $(C, k) = \max(S)$. By $\mathsf{S}3$, we find $A \subseteq C$ as well as $B \subseteq C$, and thus $A \cup B \subseteq C$. Since S is downwards closed by $\mathsf{S}2$, there exists $k \in \mathcal{I}$ such that $(\{A \cup B\}, k) \in S$. Hence, $S \sim_{\{A \cup B\}}^{\mathbb{C}} S$ and $\mathcal{C}, S \Vdash_\sigma \mathsf{alive}(\{A \cup B\})$ by Lemma 49.

6. Mono: Follows from Lemma 42.

7. Equiv: Follows from Corollary 7.

Lastly, we show ($[G]$-Nec). Let $\phi \in \mathcal{L}^{\mathsf{syn}}$ and assume that ϕ is σ-valid, i.e., for any multi-simplicial model $\mathcal{C} = (\mathbb{C}, L)$ and $S \in \mathbb{C}$, we have $\mathcal{C}, S \Vdash_\sigma \phi$. Let $S \in \mathbb{C}$ be arbitrary. By assumption, it holds that for all T with $S \sim_G^{\mathbb{C}} T$, we have $\mathcal{C}, T \Vdash_\sigma \phi$. Thus, $\mathcal{C}, S \Vdash_\sigma [G]\phi$ by the definition of truth. Since S was arbitrary, we have $\mathcal{C} \Vdash_\sigma [G]\phi$. Lastly, since $\mathcal{C}$ was arbitrary, the formula $[G]\phi$ is σ-valid, i.e., $\Vdash_\sigma [G]\phi$. □

5.4 Examples

The purpose of this section is to illustrate the expressivity of multi-simplicial models. We first present a simple multi-simplicial model in Example 5.3, which contains a subsimplex as a world. Multi-simplicial models with subsimplices as worlds can be used to model the uncertainty about the presence of other agents. Such scenarios are common in distributed systems where processes may crash. Example 9 elevates the example of $\mathbb{C}_2$ in Figure 5.1 from two triangles to two tetrahedrons. Lastly, Example 10 and Example 11 are of special interest because they apply our logic to well-known concepts in distributed computing. In both examples, synergistic knowledge can be interpreted as the knowledge that can be obtained when being able to access the functionality provided by some service. For example, the meaning of the agent pattern $G = \{\{a, b, c\}, \{b, c\}\}$ is that i) the agents a, b, and c can use the functionality provided by the service, and ii) the agents b and c alone can do so as well. In Example 10, the service is a *consensus* *object*, which allows processes to synchronize. The pattern G represents that i) all agents can reach consensus, and ii) b and c can reach consensus without agent a. In Example 11, the service is a *shared coin*, and G represents that i) all agents together can establish a shared coin, and ii) the agents b and c can establish a shared coin. Moreover, Example 10 captures the idea that for some applications, the agent pattern must include the area of the triangle and not just its edges. Thus, Example 10 shows the difference between mutual and pairwise synergy. Example 11 demonstrates that the patterns $\{\{a,b\},\{a,c\}\}$, $\{\{a,b\},\{b,c\}\}$, and $\{\{b,c\},\{a,c\}\}$ are weaker than the pattern $\{\{a,b\},\{a,c\},\{b,c\}\}$.

Regarding notation, from now on we will omit the set parentheses for agent patterns whenever it is clear from the context and write for example $[abc, ab, ac]$ instead of $[\{\{a,b,c\},\{a,b\},\{a,c\}\}]$.

Example 8 (Two Agents)**.** *Let* $\mathsf{Ag} = \{a,b\}$*, and consider the complex:*

$$\mathbb{C} = \left\{ \left\{ \begin{matrix} ab0 \\ a0, b0 \end{matrix} \right\}, \{a0\} \right\},$$

depicted in Figure 5.3. Notice that $\mathbb{C}$ *contains the simplex* $\langle a0 \rangle$*, which is also included in the edge* $\langle ab0 \rangle$*. Furthermore, let* L *be an arbitrary labeling and let* $\mathcal{C} = (\mathbb{C}, L)$ *be a multi-simplicial model. It is straightforward to verify that* $\mathcal{C}, \langle ab0 \rangle \Vdash_\sigma \mathsf{alive}(a)$ *and* $\mathcal{C}, \langle ab0 \rangle \Vdash_\sigma \mathsf{alive}(b)$*. However, it holds that* $\mathcal{C}, \langle a0 \rangle, \nVdash_\sigma \mathsf{alive}(b)$ *because* $\{b\} \nsubseteq \langle a0 \rangle^\circ$*, and hence* $\mathcal{C}, \langle a0 \rangle \Vdash_\sigma [b]\bot$*. Moreover, agent* a *alone does not know whether* $\mathsf{alive}(b)$ *because agent* a *cannot distinguish* $\langle a0 \rangle$ *from* $\langle ab0 \rangle$ *due to* $\{\{a\}\} \subseteq (\langle a0 \rangle \cap \langle ab0 \rangle)^\circ$.

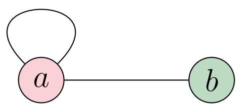

Figure 5.3: A model in which a considers it possible that it is the only agent alive. The loop indicates that the a-vertex belongs to the complex as a subsimplex. The loop is not part of the complex.

Example 9 (Two Tetrahedrons)**.** *Consider the complex:*

$$\mathbb{C} = \left\{ \left\{ \begin{matrix} abcd0 \\ abc0, bcd0, acd0, abd0, \\ ab0, bc0, ac0, ... \\ a0, b0, c0, d0 \end{matrix} \right\}, \left\{ \begin{matrix} abcd1 \\ abc1, bcd1, acd1, abd1 \\ ab1, bc1, ac1, ... \\ a0, b0, c0, d1 \end{matrix} \right\} \right\},$$

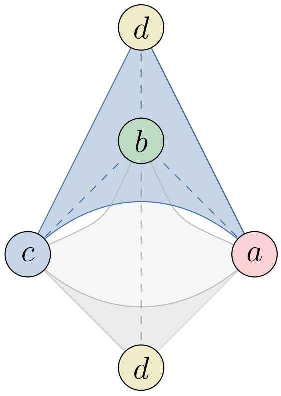

Figure 5.4: Two tetrahedron sharing only the vertices a, b, and c. Dashed lines indicate edges that are hidden from the front view. The two shades of gray highlight the visible faces of the lower tetrahedron.

depicted in Figure 5.4. Let L be a labeling such that for some $p \in \mathsf{Prop}$, we have that $p \in L(\langle abcd0\rangle)$ and $p \notin L(\langle abcd1\rangle)$. Finally, let $\mathcal{C} = (\mathbb{C}, L)$ be a multi-simplicial model. We observe that:

$$(\langle abcd0\rangle \cap \langle abcd1\rangle)^\circ = \big(\{a0, b0, c0\}\big)^\circ$$
$$= \{\{a\}, \{b\}, \{c\}\}.$$

Without synergy, the agents a, b, and c cannot know p in $\langle abcd0\rangle$ because

$$\{\{a\}, \{b\}, \{c\}\} \subseteq (\langle abcd0\rangle \cap \langle abcd1\rangle)^\circ,$$

and hence, the agents cannot distinguish $\langle abcd0\rangle$ from $\langle abcd1\rangle$, i.e.

$$\mathcal{C}, \langle abcd1\rangle \not\Vdash_\sigma [a, b, c]p.$$

However, the agent pattern $G = \{\{a, b, c\}\}$ can distinguish $\langle abcd0\rangle$ and $\langle abcd1\rangle$, and thus:

$$\mathcal{C}, \langle abcd0\rangle) \Vdash_\sigma [G]p, \text{ as well as } \mathcal{C}, \langle abcd1\rangle \Vdash_\sigma [G]\neg p.$$

Example 10 (Consensus Number)**.** *This example demonstrates how multi-simplicial models can be used to reason about the hierarchy of synchronization objects (cf. Herlihy [34]). In the asynchronous shared memory model, an n-consensus object allows n processes to reach wait-free agreement. The consensus number of an object O is the largest n for which it is possible to implement an n-consensus object by only using objects of type O and atomic registers. This introduces a hierarchy on shared memory objects because:*

1. *no combination of objects with consensus number $k < n$ can implement an object with consensus number n (cf. Herlihy [34, Thm. 1]);*

2. *objects with consensus number $n \geq k$ can implement objects with consensus number k (cf. Herlihy [34, Thm. 14]).*

The lowest elements of this hierarchy are atomic registers with consensus number 1 (cf. Herlihy [34, Thm. 2]).

Let d be a propositional variable such that processes can only know whether d or $\neg d$ is the case, if they can carry out a computation that relies on objects with a consensus number of at least k. For example, d could be computed by a smart contract with consensus number k (cf. Guerraoui et al. [32] and Alpos et al. [3]).

Our goal is to capture this scenario by a multi-simplicial model $\mathcal{C} = (\mathbb{C}, L)$. Let Ag be the set of finitely many processes. We construct $\mathcal{C}$ as follows: the complex $\mathbb{C}$ consists of two simplices that contain all processes and that only differ in their maximal element, i.e., the two simplices have maximal elements (Ag, i) and (Ag, j) with $i \neq j$. The labeling L is chosen such that it assigns different values for d to both simplices. We interpret an element $\{p_1, \ldots, p_\ell\}$ of an agent pattern G as the processes of that element having access to objects with consensus number ℓ. For three agents $a, b,$ and c, the model $\mathcal{C} = (\mathbb{C}, L)$ is given by the complex

$$\mathbb{C} = \left\{ \left\{ \begin{matrix} abc0 \\ ab0 \\ bc0 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc1 \\ ab0 \\ bc0 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\} \right\}$$

with a labeling L such that

$$\mathcal{C}, \langle abc0 \rangle \Vdash_\sigma d \quad \text{and} \quad \mathcal{C}, \langle abc1 \rangle \nVdash_\sigma d.$$

It is easy to check that $\langle abc0 \rangle \sim^{\mathbb{C}}_{ab,ac,bc} \langle abc1 \rangle$ because the simplices only differ in their maximal element. Thus, the agents cannot know whether d or $\neg d$ if they have only access to objects with consensus number 2. However,

$$\mathcal{C} \Vdash_\sigma [abc]d \vee [abc]\neg d$$

is true which shows that, when having access to an object with consensus number 3, the agents will always be able to determine whether d or $\neg d$. Furthermore, observe that the pattern $G = \{\{a\}, \{b\}, \{c\}\}$ represents traditional distributed knowledge where the processes only communicate by reading and writing to atomic registers. Lastly, since

$$\mathcal{C} \nVdash_\sigma [G]d \vee [G]\neg d,$$

$\mathcal{C}$ captures that distributed knowledge is not sufficient in this case.

Example 11 (Dining Cryptographers)**.** *The dining cryptographers problem, proposed by Chaum [19], illustrates how a shared-coin primitive can be used by three cryptographers (i.e., agents) to find out whether their employer or one of their peers paid for the dinner. However, if their employer did not pay, the payer wishes to remain anonymous. For the sake of space, we do not give a full formalization of the dining cryptographers problem. Instead, we solely focus on the ability of agreeing on a coin-flip and the resulting knowledge. In what follows, we will provide a multi-simplicial model in which the agents* a, b *and* c *can determine whether or not their employer paid if and only if they have pairwise access to a shared coin. Let the propositional variable* p *denote that their employer paid. We interpret an agent pattern* $G = \{\{a, b\}\}$ *as* a *and* b*, having access to a shared coin. Our model* $\mathcal{C} = (\mathbb{C}, L)$*, depicted in Figure 5.5, is given by the complex:*

$$\mathbb{C} = \left\{ \begin{matrix} \left\{ \begin{matrix} abc0 \\ ab0 \\ bc0 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc1 \\ ab1 \\ bc0 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc2 \\ ab0 \\ bc1 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc3 \\ ab0 \\ bc0 \\ ac1 \\ a0, b0, c0 \end{matrix} \right\}, \\ \left\{ \begin{matrix} abc4 \\ ab1 \\ bc1 \\ ac0 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc5 \\ ab1 \\ bc0 \\ ac1 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc6 \\ ab0 \\ bc1 \\ ac1 \\ a0, b0, c0 \end{matrix} \right\}, \left\{ \begin{matrix} abc7 \\ ab1 \\ bc1 \\ ac1 \\ a0, b0, c0 \end{matrix} \right\} \end{matrix} \right\}.$$

The labeling L *is chosen such that:*

$$\begin{matrix} p \in L(\langle abc0 \rangle), & p \notin L(\langle abc1 \rangle), & p \notin L(\langle abc2 \rangle), & p \notin L(\langle abc3 \rangle), \\ p \in L(\langle abc4 \rangle), & p \in L(\langle abc5 \rangle), & p \in L(\langle abc6 \rangle), & p \notin L(\langle abc7 \rangle). \end{matrix}$$

Consider the agent pattern $G = \{\{a, b\}, \{a, c\}, \{b, c\}\}$*, then:*

$$\mathcal{C} \Vdash_\sigma [G]p \vee [G]\neg p, \tag{5.11}$$

i.e., in any world, if all agents have pairwise access to shared coins, they can know the value of p*. Further, for each agent pattern* $H \subsetneq G$ *and each* $S \in \mathbb{C}$*:*

$$\mathcal{C}, S \nVdash_\sigma [H]p \vee [H]\neg p. \tag{5.12}$$

Notice that (5.12) states that there is no world, where an agent pattern $H \subsetneq G$ *can know whether* p *or* $\neg p$*, and hence, it is stronger than* $\mathcal{C} \nVdash_\sigma [H]p \vee [H]\neg p$*.*

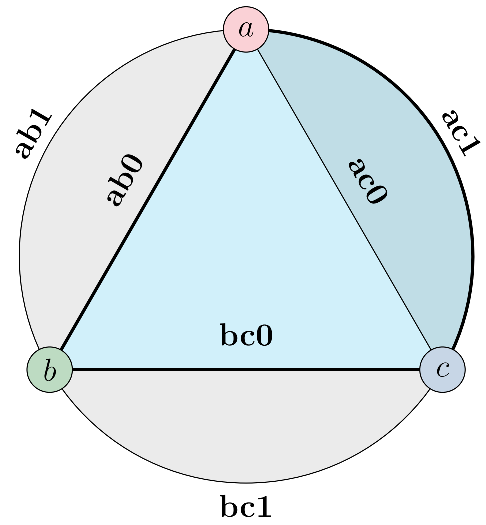

Figure 5.5: Dining cryptographers model. In total, there are eight different ways to form a solid triangle. For simplicity, the solid triangles are not labeled. The simplex $\langle abc3 \rangle$ is illustrated by the blue area, and the simplex $\langle abc7 \rangle$ is represented by the gray area.

5.5 Completeness

To establish completeness of Syn with respect to multi-simplicial models, we take a detour via relational models and focus on a fragment of Syn, denoted by Syn^{-}, which omits the axiom P. In this chapter, we adopt the global point of view. For brevity, we will omit explicitly referring to our models as global. We first introduce κ-models, which represent multi-agent relational models in which the property (SGK), i.e., standard group knowledge, need not be satisfied. We can employ the standard techniques to show that Syn^{-} is sound and complete with respect to κ-models. Next, we present δ-models, which are κ-models that additionally satisfy (SGK). The proof that Syn^{-} is complete with respect to δ-models is more involved and requires the so-called *unraveling method* (cf. Halpern et al. [21] and Van der Hoek and Meyer [49]). Once soundness and completeness of Syn^{-} with respect to those models is established, we show that Syn is sound and complete with respect to proper[2] δ-models. Lastly, we prove completeness of Syn with respect to multi-simplicial models by relating them to proper δ-models.

[2]Properness with respect to agent patterns will be defined later. The definition follows the same principles as Definition 11.

5.5.1 κ-models

This section introduces a generalization of frames. Instead of reasoning about the view of a group of agents, a *pattern frame* describes which agent patterns can distinguish which worlds. Relational models based on such frames are called *pattern models*. A special type of pattern frames are κ-frames. A pattern model whose underlying frame is a κ-frame is referred to as a κ-model. After discussing the properties of κ-models, we show that Syn^{-} is sound and complete with respect to κ-models.

Definition 53 (Pattern Frame). *A* pattern frame *is a pair* (W, R) *such that:*

1. *W is a set of possible worlds;*

2. *R is a function assigning a relation R_G to each agent pattern $G \in \mathsf{AP}$.*

The naming conventions for frames apply to pattern frames. If a frame is symmetric and transitive, we write $F = (W, \sim)$ instead of $F = (W, R)$ to indicate those properties. Throughout this chapter, if clear from context, we refer to pattern frames as *frames*.

Definition 54 (Pattern Model). *A* pattern model *is a pair $\mathcal{M} = (F, V)$ where:*

1. *$F = (W, R)$ is a frame;*

2. *$V : W \to \mathsf{Pow}(\mathsf{Prop})$ is a valuation.*

Remark 10 (Notation). *We reserve capital letters for explicit simplices and lowercase letters for worlds. For example, S, T, and U denote explicit simplices, and u, v, and w represent worlds of a pattern frame.*

We define the satisfaction relation for pattern models as expected. We overload notation and use $\Vdash$ for pattern models and global relational models.

Definition 55 ($\Vdash$). *Let $F = (W, R)$ be a frame, and let $\mathcal{M} = (F, V)$ be a pattern model. For all worlds $w \in W$, we define the relation $\mathcal{M}, w \Vdash \phi$ by induction on $\phi \in \mathcal{L}^{\mathsf{syn}}$:*

$$
\begin{array}{lll}
\mathcal{M}, w \Vdash p & \textit{iff} & p \in V(w) \\
\mathcal{M}, w \Vdash \neg\phi & \textit{iff} & \mathcal{M}, w \nVdash \phi \\
\mathcal{M}, w \Vdash \phi \wedge \psi & \textit{iff} & \mathcal{M}, w \Vdash \phi \textit{ and } \mathcal{M}, w \Vdash \psi \\
\mathcal{M}, w \Vdash [G]\phi & \textit{iff} & wR_G v \textit{ implies } \mathcal{M}, v \Vdash \phi, \textit{ for all } v \in W.
\end{array}
$$

Let $F = (W, R)$ be a frame and consider the pattern model $\mathcal{M} = (F, V)$. If $\mathcal{M}, w \Vdash \phi$ holds, then ϕ is said to be *satisfied* at the world w in $\mathcal{M}$. A formula $\phi \in \mathcal{L}^{\mathsf{syn}}$ is *valid in* $\mathcal{M}$, denoted by $\mathcal{M} \Vdash \phi$, if it is satisfied at all worlds $w \in W$. Moreover, ϕ is said to be *valid*, written $\Vdash \phi$, if ϕ is valid in every pattern model.

Definition 56 specifies what it means for an agent pattern to be alive.

Definition 56 ($\mathsf{Alive}(G)_F$)**.** *Let $F = (W, R)$ be a frame. We define:*

$$\mathsf{Alive}(G)_F = \{w \in W \mid wR_G w\}.$$

If the frame is clear from the context, we omit the subscript F and write $\mathsf{Alive}(G)$ instead of $\mathsf{Alive}(G)_F$.

Definition 57 introduces κ-frames. Frames that need not satisfy (SGK) are sometimes referred to as *pseudo-frames* (cf. Ågotnes et al. [1]).

Definition 57 (κ-frame)**.** *A symmetric and transitive frame $F = (W, \sim)$ is called κ-frame if and only if for all agent patterns G and H:*

$\mathsf{K1}$*: $\mathsf{Alive}(G)_F \cap \mathsf{Alive}(H)_F \subseteq \mathsf{Alive}(G \cup H)_F$;*

$\mathsf{K2}$*: $\mathsf{Alive}(G)_F \subseteq \mathsf{Alive}(\{A \cup B\})_F$ for $A, B \in G$;*

$\mathsf{K3}$*: $\sim_H \subseteq \sim_G$, if $G \subseteq H$;*

$\mathsf{K4}$*: $\sim_G \subseteq \sim_{G \cup \{B\}}$ if there exists $A \in G$ with $\emptyset \neq B \subseteq A$;*

NE*: for all $w \in W$, there exists an agent pattern G such that $w \sim_G w$.*

Besides reasoning about the view of agent patterns instead of the view of a group of agents, κ-frames are similar to the by Goubault et al. [26] introduced *generalized epistemic frames*. Those frames satisfy conditions $\mathsf{K1}, \mathsf{K3}$, and NE, if G and H were sets of agents. Property $\mathsf{K1}$ ensures that for each world, there exists a maximal alive agent pattern; $\mathsf{K3}$ guarantees that an agent pattern G cannot know more than its supersets; and, lastly, NE ensures that there are no empty-worlds, i.e., worlds in which no agent pattern is alive.

The condition $\mathsf{K2}$ forces $\mathsf{alive}(G)$ to be downwards closed, and $\mathsf{K4}$ states that adding subpatterns to G does not strengthen its knowledge. Nonetheless, as shown below in Example 12, κ-frames do not necessarily satisfy (SGK). The properties $\mathsf{K3}$ and $\mathsf{K4}$ together yield Lemma 46.

Lemma 46. *Let $F = (W, \sim)$ be a κ-frame, and let G be an agent pattern such that there exists $A \in G$ with $\emptyset \neq B \subseteq A$. It holds that $\sim_G \subseteq \sim_{\{B\}}$.*

Proof. By $\mathsf{K3}$ and $\mathsf{K4}$ we have $\sim_G \subseteq \sim_{G \cup \{B\}} \subseteq \sim_{\{B\}}$. □

Remark 11. *Let $F = (W, \sim)$ be a κ-*frame*, and let G be an agent pattern as in Lemma 46. By* K3*, we have $\sim_G = \sim_{G \cup \{B\}}$.*

Example 12. *Let $\mathsf{Ag} = \{a, b\}$ and consider the κ-frame F shown in Figure 5.6. The agents patterns $\{\{a\}\}$ and $\{\{b\}\}$ cannot distinguish between the worlds w_1 and w_2. However, their union $G = \{\{a\}, \{b\}\}$ can tell the two worlds apart. Thus, the κ-frame F does not satisfy property* (SGK).

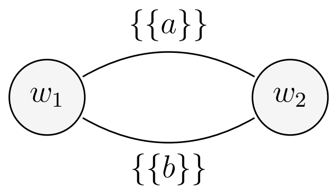

Figure 5.6: The κ-frame used in Example 12 does not satisfy standard group knowledge. Reflexive arrows are implicit.

Definition 58 (κ-model)**.** *A pattern model $\mathcal{M} = (F, V)$ is a κ-*model *if and only if F is a κ-frame.*

A formula $\phi \in \mathcal{L}^{\mathsf{syn}}$ is κ-valid, denoted $\Vdash_\kappa \phi$, if $\mathcal{M} \Vdash \phi$ for all κ-models $\mathcal{M}$. Lemma 47 shows a formula that is κ-valid.

Lemma 47. *Let $G \in \mathsf{AP}$ be an agent pattern such that there exists $A \in G$ and $\emptyset \neq B \subseteq A$. The following holds:*

$$\Vdash_\kappa \mathsf{alive}(G) \to \mathsf{alive}(\{B\}).$$

Proof. Let $F = (W, \sim)$ be a frame such that the pattern model $\mathcal{M} = (F, V)$ is a κ-model. Further, assume that $\mathcal{M}, w \Vdash \mathsf{alive}(G)$. By definition, there exists a world $v \in W$ such that $w \sim_G v$ and $\mathcal{M}, v \Vdash \top$. By symmetry and transitivity of $\sim_G$, we obtain that $w \sim_G w$. By Lemma 46, it holds that $(w, w) \in \sim_{\{B\}}$, and therefore we find that $\mathcal{M}, w \Vdash \mathsf{alive}(\{B\})$. $\square$

Lemma 48 states that if $\mathcal{M} = (F, V)$ is a model based on a symmetric and transitive frame $F = (W, \sim)$, then the formula $\mathsf{alive}(G)$ is satisfied at a world w if and only if that world has a reflexive arrow labeled with G.

Lemma 48. *Let $F = (W, \sim)$ be a symmetric and transitive frame, and let $\mathcal{M} = (F, V)$ be a pattern model. It holds that:*

$$\mathcal{M}, w \Vdash \mathsf{alive}(G) \quad \textit{iff} \quad w \in \mathsf{Alive}(G)_F.$$

Proof. We first show that

$$\mathcal{M}, w \Vdash \mathsf{alive}(G) \quad \text{implies} \quad w \in \mathsf{Alive}(G)_F.$$

Assume $\mathcal{M}, w \Vdash \mathsf{alive}(G)$, i.e., $\mathcal{M}, w \nVdash [G]\bot$. By the definition of truth, there must exist $v \in W$ with $w \sim_G v$ and $\mathcal{M}, v \Vdash \top$. By symmetry, we have $v \sim_G w$ and by transitivity we have $w \sim_G w$. Hence, $w \in \mathsf{Alive}(G)_F$. We now show that

$$w \in \mathsf{Alive}(G)_F \quad \text{implies} \quad \mathcal{M}, w \Vdash \mathsf{alive}(G).$$

Assume $w \in \mathsf{Alive}(G)_F$. By definition of $\mathsf{Alive}(G)_F$, we have that $w \sim_G w$. Therefore, $\mathcal{M}, w \nVdash [G]\bot$ by the definition of truth. □

The soundness proof is straightforward.

Theorem 5 (Soundness)**.** $\mathsf{Syn}^{\text{-}}$ *is sound with respect to κ-models.*

Proof. We only show the cases for the axioms T, NE, Union, Clo, Mono, Equiv, and $[G]$-Nec. Let $F = (W, \sim)$ be a frame such that $\mathcal{M} = (F, V)$ is a κ-model.

1. T: Consider a world $w \in W$ and assume that $\mathcal{M}, w \Vdash \mathsf{alive}(G)$ and $\mathcal{M}, w \Vdash [G]\phi$. By Lemma 48 we have $w \in \mathsf{Alive}(G)$, i.e., $w \sim_G w$. By the definition of truth we have $\mathcal{M}, w \Vdash \phi$.

2. NE: Let $w \in W$ be arbitrary. By NE, there exists an agent pattern G such that $w \sim_G w$. By Lemma 48, $\mathcal{M}, w \Vdash \mathsf{alive}(G)$ and thus

$$\mathcal{M}, w \Vdash \bigvee_{G \in \mathsf{AP}} \mathsf{alive}(G).$$

3. Union: Assume $\mathcal{M}, w \Vdash \mathsf{alive}(G)$ and $\mathcal{M}, w \Vdash \mathsf{alive}(H)$. By Lemma 48, $w \in \mathsf{Alive}(G) \cap \mathsf{Alive}(H)$. By $\mathsf{K1}$, $w \in \mathsf{Alive}(G \cup H)$ and by Lemma 48 $\mathcal{M}, w \Vdash \mathsf{alive}(G \cup H)$.

4. Clo: Assume $\mathcal{M}, w \Vdash \mathsf{alive}(G)$ and let $A, B \in G$. By Lemma 48, we have $w \in \mathsf{Alive}(G)$. By $\mathsf{K2}$, $w \in \mathsf{Alive}(\{A \cup B\})$ and by Lemma 48 we obtain $\mathcal{M}, w \Vdash \mathsf{alive}(\{A \cup B\})$.

5. Mono: Assume that $G \subseteq H$ for arbitrary G, H. Let $w \in W$ be arbitrary such that $\mathcal{M}, w \Vdash [G]\phi$. By the definition of truth, $w \sim_G v$ implies $\mathcal{M}, v \Vdash \phi$ for all $v \in W$. By $\mathsf{K3}$ we have that $\sim_H \subseteq \sim_G$, i.e., $w \sim_H v$ implies $w \sim_G v$. Thus, $\mathcal{M}, v \Vdash \phi$ whenever $w \sim_H v$. Therefore, it follows that $\mathcal{M}, w \Vdash [H]\phi$.

6. Equiv: Assume that for an arbitrary G, there exists a set $A \in G$ with $B \subseteq A$. Further, let $w \in W$ be arbitrary such that $\mathcal{M}, w \Vdash [G \cup \{B\}]\phi$. Therefore, for all $v \in W$, $w \sim_{G\cup\{B\}} v$ implies $\mathcal{M}, v \Vdash \phi$ by assumption. By $\mathsf{K4}$ we have that $\sim_G \subseteq \sim_{G\cup\{B\}}$ and thus $w \sim_G v$ implies $w \sim_{G\cup\{B\}} v$. Hence, $\mathcal{M}, v \Vdash \phi$ whenever $w \sim_G v$ and thus $\mathcal{M}, w \Vdash [G]\phi$.

Lastly, we show G-Nec. Let $\phi \in \mathcal{L}^{\mathsf{syn}}$ and assume $\Vdash_\kappa \phi$. We need to show that $[G]\phi$ is κ-valid. Let $F = (W, \sim)$ be an arbitrary κ-frame and let $\mathcal{M} = (F, V)$ be a κ-model. By assumption, $\mathcal{M}, w \Vdash \phi$ for all $w \in W$. Thus, for any $v \in W$ with $w \sim_G v$, it holds that $\mathcal{M}, v \Vdash \phi$. By the definition of truth, $\mathcal{M}, w \Vdash [G]\phi$, and since $w \in W$ was arbitrary, $\mathcal{M} \Vdash [G]\phi$. Moreover, due to $\mathcal{M}$ being arbitrary, $[G]\phi$ is κ-valid. □

We will now establish completeness of Syn^- with respect to κ-models. In what follows, we set up the usual machinery to construct the canonical model based on maximal consistent sets (see Definition 40).

Definition 59. *Let $G \in \mathsf{AP}$ and let $\Gamma \subseteq \mathcal{L}^{\mathsf{syn}}$, we define*

$$\Gamma \setminus [G] = \{\phi \mid [G]\phi \in \Gamma\}.$$

Definition 60 (Canonical Model)**.** *The canonical frame $F^c = (W^c, \sim^c)$ and the* canonical model *$\mathcal{M}^c = (F^c, V^c)$ for Syn^- are defined as:*

1. *$W^c = \{\Gamma \subseteq \mathcal{L}^{\mathsf{syn}} \mid \Gamma$ is a maximal consistent set for $\mathsf{Syn}^-\}$ is the set of possible worlds;*

2. *$\sim^c$ is a function that assigns to each agent pattern G a relation*

$$\sim^c_G = \{(\Gamma, \Delta) \in W^c \times W^c \mid \Gamma \setminus [G] \subseteq \Delta\};$$

3. *$V^c : W^c \to \mathsf{Pow}(\mathsf{Prop})$ is a function defined by*

$$V^c(\Gamma) = \{p \in \mathsf{Prop} \mid p \in \Gamma\}.$$

Lemma 49. *Let $G \in \mathsf{AP}$ and $\Gamma \in W^c$, then*

$$\Gamma \in \mathsf{Alive}(G)_{F^c} \quad \textit{iff} \quad \mathsf{alive}(G) \in \Gamma.$$

Proof. We first show that

$$\Gamma \in \mathsf{Alive}(G)_{F^c} \quad \text{implies} \quad \mathsf{alive}(G) \in \Gamma.$$

Assume $\Gamma \in \mathsf{Alive}(G)_{F^c}$. By Definition 56, $\Gamma \sim^c_G \Gamma$. Towards a contradiction, assume that $\mathsf{alive}(G) \notin \Gamma$, i.e., $[G]\bot \in \Gamma$ by the maximal consistency of Γ. Since

$\Gamma \sim_G^c \Gamma$, i.e., $\Gamma \setminus [G] \subseteq \Gamma$, this yields $\bot \in \Gamma$, which contradicts the consistency of Γ. Thus $\mathsf{alive}(G) \in \Gamma$. We now show that

$$\mathsf{alive}(G) \in \Gamma \quad \text{implies} \quad \Gamma \in \mathsf{Alive}(G)_{F^c}.$$

Assume $\mathsf{alive}(G) \in \Gamma$. We need to show $\Gamma \sim_G^c \Gamma$. Let $\phi \in \Gamma \setminus [G]$, i.e., $[G]\phi \in \Gamma$. Since

$$\mathsf{alive}(G) \to ([G]\phi \to \phi)$$

is an axiom of $\mathsf{Syn}^{\text{-}}$ and by the maximal consistency of Γ, it follows that $\phi \in \Gamma$, i.e., $\Gamma \sim_G^c \Gamma$. By Definition 56, we obtain $\Gamma \in \mathsf{Alive}(G)_{F^c}$. □

The next lemma states that F^c satisfies the properties of κ-frames.

Lemma 50. *F^c is a κ-frame.*

Proof. We show that F^c satisfies all properties of κ-frames.

1. Symmetry: Assume $\Gamma \setminus [G] \subseteq \Delta$. We need to show $\Delta \setminus [G] \subseteq \Gamma$. Let $\phi \in \Delta \setminus [G]$, i.e., $[G]\phi \in \Delta$. Since Δ is a maximal consistent set, we also have $\neg\neg\phi \in \Delta \setminus [G]$. Assume now towards a contradiction that $\phi \notin \Gamma$. Since Γ is a maximal consistent set and B is an axiom of $\mathsf{Syn}^{\text{-}}$, we have $\neg\phi \in \Gamma$ as well as $[G]\neg[G]\neg\neg\phi \in \Gamma$ by the maximal consistency of Γ. Therefore, $\neg[G]\neg\neg\phi \in \Gamma \setminus [G]$ and thus $\neg[G]\neg\neg\phi \in \Delta$. This is a contradiction because by assumption, we have $[G]\neg\neg\phi \in \Delta$. Hence, we conclude $\phi \in \Gamma$ which shows that $\Delta \setminus [G] \subseteq \Gamma$.

2. Transitivity: Assume $\Gamma \sim_G^c \Delta$ and $\Delta \sim_G^c \Phi$. We need to show $\Gamma \sim_G^c \Phi$. By assumption, we have $\Gamma \setminus [G] \subseteq \Delta$ and $\Delta \setminus [G] \subseteq \Phi$. Let $\phi \in \Gamma \setminus [G]$, i.e., $[G]\phi \in \Gamma$. Since Γ is maximally consistent and because 4 is an axiom of $\mathsf{Syn}^{\text{-}}$, we have $[G][G]\phi \in \Gamma$, and thus $[G]\phi \in \Delta$. Further, since $\Delta \sim_G^c \Phi$, we have $\phi \in \Phi$ and hence $\Gamma \setminus [G] \subseteq \Phi$.

3. $\mathsf{K1}$: Assume $\Gamma \in \mathsf{Alive}(G)$ and $\Gamma \in \mathsf{Alive}(H)$. By Lemma 49 it follows that $\mathsf{alive}(G) \in \Gamma$ and $\mathsf{alive}(H) \in \Gamma$. Since Γ is a maximal consistent set and Union is an axiom of $\mathsf{Syn}^{\text{-}}$ it follows that $\mathsf{alive}(G \cup H) \in \Gamma$. By Lemma 49 we have $\Gamma \in \mathsf{Alive}(G \cup H)$.

4. $\mathsf{K2}$: Assume $A, B \in G$ and let $\Gamma \in \mathsf{Alive}(G)$. By Lemma 49 it follows that $\mathsf{alive}(G) \in \Gamma$. Since Γ is a maximal consistent set and Clo is an axiom of $\mathsf{Syn}^{\text{-}}$ it follows that $\mathsf{alive}(\{A \cup B\}) \in \Gamma$ by the maximal consistency of Γ. By Lemma 49 we have $\Gamma \in \mathsf{Alive}(\{A \cup B\})$.

5. $\mathsf{K3}$: Assume $(\Gamma, \Delta) \in \sim_H^c$, we need to show that $(\Gamma, \Delta) \in \sim_G^c$. Let $\phi \in \Gamma \setminus [G]$, i.e., $[G]\phi \in \Gamma$. Since Γ is maximally consistent and Mono is an axiom of $\mathsf{Syn}^{\text{-}}$ we have that $[H]\phi \in \Gamma$. Since we assumed that $\Gamma \setminus [H] \subseteq \Delta$ we have $\phi \in \Delta$.

6. K4: Assume $(\Gamma, \Delta) \in \sim^c_G$, we need to show that $(\Gamma, \Delta) \in \sim^c_{G\cup\{B\}}$. Let $\phi \in \Gamma \setminus [G \cup \{B\}]$, i.e., $[G\cup\{B\}]\phi \in \Gamma$. Since Γ is maximally consistent and Equiv is an axiom of Syn⁻ we have that $[G]\phi \in \Gamma$. Since we assumed $\Gamma \setminus [G] \subseteq \Delta$ we have $\phi \in \Delta$.

7. NE: Let $\Gamma \in W^c$ be arbitrary. Since Γ is maximally consistent and NE is an axiom of Syn⁻, there exists an agent pattern G such that $\mathsf{alive}(G) \in \Gamma$. By Lemma 49, $\Gamma \sim^c_G \Gamma$. □

The truth lemma is standard and completeness follows.

Lemma 51 (Truth Lemma). *For each world $\Gamma \in W^c$ and each formula $\phi \in \mathcal{L}^{\mathsf{syn}}$, it holds that:*

$$\mathcal{M}^c, \Gamma \Vdash \phi \text{ iff } \phi \in \Gamma.$$

Theorem 6 (Completeness). Syn⁻ *is complete with respect to κ-models.*

5.5.2 δ-models

Example 12 showed that κ-models need not satisfy the property (SGK). In what follows, we will introduce δ-models, which are κ-models that satisfy (SGK). We show that the system Syn⁻ is sound and complete with respect to δ-models by applying the unraveling method to the canonical frame F^c.

Definition 61 (δ-frame). *A symmetric and transitive frame $F = (W, \sim)$ is called* δ-frame *if and only if for all agent patterns G and H:*

K2: $\mathsf{Alive}(G)_F \subseteq \mathsf{Alive}(\{A \cup B\})_F$ *for* $A, B \in G$;

K3: $\sim_H \subseteq \sim_G$, *if* $G \subseteq H$;

K4: $\sim_G \subseteq \sim_{G\cup\{B\}}$ *if there exists* $A \in G$ *with* $\emptyset \neq B \subseteq A$;

NE: *for all* $w \in W$, *there exists an agent pattern* G *such that* $w \sim_G w$;

D: $\sim_G = \bigcap_{B\in G} \sim_{\{B\}}$.

Observe that D implies K1 and thus, δ-models are κ-models.

Definition 62 (δ-model). *A pattern model $\mathcal{M} = (F, V)$ is a* δ-model *if and only if F is a δ-frame.*

A formula ϕ is δ-valid, denoted $\Vdash_\delta \phi$, if $\mathcal{M} \Vdash \phi$ for all δ-models $\mathcal{M}$. Before presenting the unraveling method, we need to introduce the notion of a history. Given a symmetric and transitive frame F, the worlds of its unraveled frame will be all its histories.

Definition 63 (History)**.** *Let $F = (W, R)$ be a frame. A history is a non-empty and finite sequence of triples (w, G, v) where:*

1. *wR_Gv and G is maximal under set inclusion. That means, there does not exist an agent pattern G' with $G \subsetneq G'$ and $wR_{G'}v$;*
2. *if (w', G', v') is the successor of (w, G, v), then $v = w'$.*

We write $\ell(h)$ to denote the last world of a history h. That is, if (w, G, v) is the last element of h, then $\ell(h) = v$. Furthermore, $h \parallel (\ell(h), G, v)$ denotes the extension of h with $(\ell(h), G, v)$. The set of all histories over a frame F is denoted by H_F. Definition 64 specifies a prefix relation $\rightarrow_G$ on histories. We can use H_F and $\rightarrow_G$ to define the unraveled frame (Definition 65).

Definition 64 ($\rightarrow_G$)**.** *Let $F = (W, R)$ be a frame and let $h, h' \in H_F$. For $G \in \mathsf{AP}$, we define $\rightarrow_G \subseteq H_F \times H_F$ as follows:*

$$h \rightarrow_G h' \quad \textit{iff} \quad h' = h \parallel (\ell(h), U, \ell(h')) \textit{ and } G \subseteq U \in \mathsf{AP}.$$

Definition 65 ($U(F)$)**.** *Let $F = (W, R)$ be a frame. We define the unraveled frame $U(F) = (H_F, \{\approx_G\}_{G \in \mathsf{AP}})$ where $\approx_G$ is the transitive closure of the symmetric closure of $\rightarrow_G$, i.e., $\approx_G = (\rightarrow_G \cup \rightarrow_G^{-1})^*$.*

An unraveled model (Definition 66) is an unraveled frame equipped with a valuation L on histories that mirrors the valuation V of the original model.

Definition 66 ($U(\mathcal{M})$)**.** *Let $F = (W, R)$ be a frame and consider the pattern model $\mathcal{M} = (F, V)$. We call $U(\mathcal{M}) = (U(F), L)$ such that:*

$$p \in L(h) \quad \textit{iff} \quad p \in V(\ell(h)),$$

the unraveled model of $\mathcal{M}$.

Lemma 52 shows that histories of κ-frames are downwards closed. This is because we are requiring that for each element (w, G, v) of a history, the set G is maximal under set inclusion (see Definition 63).

Lemma 52 (Downwards Closure)**.** *Let $F = (W, \sim)$ be a κ-frame and consider a history $h \in H_F$ and let (w, U, v) be an element of h. If there exists $A \in U$ with $\emptyset \neq B \subseteq A$, then $B \in U$.*

Proof. Towards a contradiction, suppose that $B \notin U$. Consider the agent pattern $U' = U \cup \{B\}$. Clearly $U \subsetneq U'$ and $w \sim_{U'} v$ because of $\mathsf{K4}$. This contradicts U being maximal under set inclusion and thus, $B \in U$. □

Lemma 53 establishes some natural properties of $\to_G$ on κ-frames.

Lemma 53. *Let $F = (W, \sim)$ be a κ-frame and consider $h, h' \in H_F$ with $h \to_G h'$ for some agent pattern G. The following hold:*

1. *if $H \subseteq G$, then $h \to_H h'$;*
2. *if there exists $A \in G$ and $\emptyset \neq B \subseteq A$, then $h \to_{\{B\}} h'$;*
3. *if $h \to_H h'$, then $h \to_{G \cup H} h'$.*

Proof. By assumption, $h' = h \parallel (w, U, v)$ with $G \subseteq U$. For the first case, we have $H \subseteq G \subseteq U$ and thus $h \to_H h'$ by Definition 64. For the second case, we have $\{B\} \subseteq U$ by Lemma 52 and thus $h \to_{\{B\}} h'$ by Definition 64. Lastly, assume $h \to_G h'$ and $h \to_H h'$, i.e., $G \subseteq U$ and $H \subseteq U$. Hence $G \cup H \subseteq U$ by the properties of set union, and therefore $h \to_{G \cup H} h'$ by Definition 64. □

To show that an unraveled κ-frame remains to be a κ-frame (Theorem 7), we briefly recall the standard notion of paths generated by a relation R. Corollary 9 and Remark 13 are crucial in the proof of Theorem 7.

Definition 67 (R-path)**.** *Let R be a relation on a set X. An R-*path *from x_1 to x_n is a sequence*

$$\tau = (x_1, x_2), (x_2, x_3), \ldots, (x_{n-2}, x_{n-1}), (x_{n-1}, x_n)$$

with $(x_i, x_{i+1}) \in R$ for $1 \leq i \leq n - 1$.

The composition of a relation with itself is defined as usual.

Definition 68 ($R \circ R$)**.** *Let R be a relation on a set X. We define:*

$$R \circ R = \{(x, y) \in X \times X \mid \text{ There exists } z \in X \text{ with } (x, z) \in R \text{ and } (z, y) \in R\}.$$

We abbreviate the n-fold composition of R with itself as R^n.

Remark 12. *Let R be relation on a set X, then there is an R-path of length n from a to b if and only if $(a, b) \in R^n$.*

Corollary 9. *Let $F = (W, \sim)$ be a κ-frame. The following two are equivalent*

1. *$(h, h') \in \approx_G$;*
2. *there exists a $(\to_G \cup \to_G^{-1})$-path τ from h to h'.*

For brevity we refer to $(\to_G \cup \to_G^{-1})$-paths as $\to_G$-paths.

Remark 13. *If $F = (W, \sim)$ is a κ-frame, then a $\to_G$-path from h to h' implies the existence of a $\sim_G$-path from $\ell(h)$ to $\ell(h')$. It follows by transitivity that $\ell(h) \sim_G \ell(h')$.*

The established properties of $\to_G$ enable us to show that $U(F)$ is a κ-frame, which is an important step towards completeness.

Theorem 7. *Let $F = (W, \sim)$ be a κ-frame. Then $U(F)$ is a κ-frame.*

Proof. Observe that $\approx_G = (\to_G \cup \to_G^{-1})^*$ is transitive and symmetric, because the transitive closure of a symmetric relation is transitive and symmetric. Moreover, NE follows because histories are not empty.

- $\mathsf{K1}$: Assume $(h, h) \in \approx_G$ and $(h, h) \in \approx_H$. By Remark 13, $\ell(h) \sim_G \ell(h)$ as well as $\ell(h) \sim_H \ell(h)$. Due to F satisfying $\mathsf{K1}$, it holds that $\ell(h) \sim_{G \cup H} \ell(h)$. Hence, $h^* = h \parallel (\ell(h), U, \ell(h))$ with $G \cup H \subseteq U$ is a valid history and $(h, h^*) \in \approx_{G \cup H}$. We obtain $(h, h) \in \approx_{G \cup H}$ by symmetry and transitivity.

- $\mathsf{K2}$: Assume $(h, h) \in \approx_G$ and let $A, B \in G$. By assumption and Remark 13, $\ell(h) \sim_G \ell(h)$. Due to F satisfying $\mathsf{K2}$, we have $\ell(h) \sim_{\{A \cup B\}} \ell(h)$. Thus $h^* = h \parallel (\ell(h), U, \ell(h))$ with $\{A \cup B\} \subseteq U$ is a valid history which implies that $(h, h) \in \approx_{\{A \cup B\}}$ by symmetry and transitivity.

- $\mathsf{K3}$: Assume $(h, h') \in \approx_H$ and $G \subseteq H$. By Corollary 9, there must exist a $\to_H$-path τ from h to h'. Let $(s, s') \in \tau$ be arbitrary. Since F satisfies $\mathsf{K3}$, Lemma 53 implies that $s \approx_G s'$. Therefore, τ is a $\to_G$-path and $(h, h') \in \approx_G$ by Corollary 9.

- $\mathsf{K4}$: Assume $(h, h') \in \approx_G$ and that there exists $A \in G$ with $\emptyset \neq B \subseteq A$. By Corollary 9 there exists a $\to_G$-path τ from h to h'. Let $(s, s') \in \tau$ be arbitrary. Since F satisfies $\mathsf{K4}$, Lemma 53 implies that $s \approx_{G \cup \{B\}} s'$. Therefore, τ is a $\to_{G \cup \{B\}}$-path and $(h, h') \in \approx_{G \cup \{B\}}$ by Corollary 9. □

We call an R-path from x_1 to x_n non-redundant if and only if $x_i \neq x_{i+2}$ for $1 \leq i < n - 1$. Further, we often write

$$\tau = x_1 R x_2 R x_3 \ldots x_{n-1} R x_n$$

instead of $\tau = (x_1, x_2), (x_2, x_3), \ldots, (x_{n-2}, x_{n-1}), (x_{n-1}, x_n)$.

If a relation R is the union of another relation S and its converse, i.e., $R = S \cup S^{-1}$, we will use S and S^{-1} in a path instead of R. Moreover, we say that a $\to_G$-path has a *change of direction*, if $(x_i, x_{i+1}) \in \to_G$ (or $(x_i, x_{i+1}) \in \to_G^{-1}$) and $(x_{i+1}, x_{i+2}) \in \to_G^{-1}$ (or $(x_{i+1}, x_{i+2}) \in \to_G$). Lastly, we will write $h \leftarrow_G h'$ subsequently instead of $h \to_G^{-1} h'$ for better readability.

Lemma 54 states that if two histories of a κ-frame are related under $\approx_G$, i.e., $h \approx_G h'$, then h and h' have a common prefix.

Lemma 54. *Let $F = (W, \sim)$ be a κ-frame and consider the unraveled frame $U(F)$. If $(h_1, h_n) \in \approx_G$, then h_1 and h_n have a common prefix.*

Proof. If one history is a prefix of the other, the claim follows trivially. Hence, we assume that neither of them is a prefix of the other. Corollary 9 ensures that there exists a $\rightarrow_G$-path $\tau = (h_1, h_2), \ldots, (h_{n-1}, h_n)$, which we can assume to be non-redundant. Moreover, due to our first assumption, τ must have at least one change of direction. In order to show that a common prefix of h_1 and h_n exists, we show that τ has exactly one change of direction, and that change is of the form $h_{i-1} \leftarrow_G h_i \rightarrow_G h_{i+1}$ with $i > 1$.

Let h_i be the history at which the first change of direction occurs. Notice, that this implies that $i > 1$. First, we observe that this change of direction cannot be of the form $h_{i-1} \rightarrow_G h_i \leftarrow_G h_{i+1}$, because by Definition 64, it would follow that $h_{i-1} = h_{i+1}$, which would contradict τ being non-redundant. Thus, the first change of direction must be of the form $h_{i-1} \leftarrow_G h_i \rightarrow_G h_{i+1}$. Consequently, if there was an additional change of direction, then it would be of the form $h_{i+k-1} \rightarrow_G h_{i+k} \leftarrow_G h_{i+k+1}$ for $k \geq 1$, which again contradicts τ being non-redundant. Therefore, h_i is a common prefix. □

Lemma 55 states that if two agent patterns G and H cannot distinguish between two histories, their union cannot do so as well. Together with Lemma 53 and Lemma 54, we can state Lemma 56, which says that unraveled κ-frames satisfy the property D.

Definition 69 specifies the meaning of two κ-models being logically equivalent. Lemma 57 states that κ-models satisfying this notion of equivalence are pointwise equivalent (cf. Goubault et al. [26]).

Lemma 55. *Let $F = (W, \sim)$ be a κ-frame and consider the unraveled frame $U(F)$. For any two agent patterns G and H, the following holds*

$$(h, h') \in \approx_G \text{ and } (h, h') \in \approx_H \text{ implies } (h, h') \in \approx_{G \cup H} .$$

Proof. Assume $h \approx_G h'$ and $h \approx_H h'$. By Lemma 54 the paths below exist:

- $h''_G \rightarrow_G \ldots \rightarrow_G h$,
- $h''_G \rightarrow_G \ldots \rightarrow_G h'$,
- $h''_H \rightarrow_H \ldots \rightarrow_H h$, and
- $h''_H \rightarrow_H \ldots \rightarrow_H h'$.

Observe that either $h''_G = h''_H$ or one of the histories is a proper prefix of the other. If $h''_G \neq h''_H$, let h'' be the longer history. If they are of the same length, fix either $h'' = h''_G$ or $h'' = h''_H$. We can write h and h' as:

$$h = h'' \parallel (w_1, G_1, w_2) \parallel \ldots \parallel (w_{n-1}, G_n, w_n), \text{ and}$$

$$h' = h'' \parallel (w'_1, G'_1, w'_2) \parallel \ldots \parallel (w'_{m-1}G'_m, w'_m).$$

By Definition 65 we have $G \subseteq G_i, G'_j$ and $H \subseteq G_i, G'_j$ for all $i, j \geq 0$. Thus, by Lemma 53 and Corollary 9, we have $h \approx_{G \cup H} h'$. □

Lemma 56. *Let $F = (W, \sim)$ be a κ-frame. $U(F)$ satisfies* D.

Proof. $\approx_G \subseteq \bigcap_{B \in G} \approx_{\{B\}}$ follows directly by K3. For the other direction, let $G = \{A_1, \ldots, A_n\}$ be an agent pattern and consider the sets

$$B_1 = \{A_1\} \text{ and } B_i = \{A_i\} \cup B_{i-1} \text{ for } 2 \leq i \leq n.$$

For $1 \leq i \leq n$, if $(u, v) \in \bigcap_{B \in G} \approx_{\{B\}}$, then $(u, v) \in \approx_{A_i}$. Applying Lemma 55 inductively yields $(u, v) \in \approx_{B_i}$. Since $B_n = G$, we obtain $(u, v) \in \approx_G$. □

Definition 69 (Functional Bisimulation)**.** *A function $f : W_{\mathcal{M}} \to W_{\mathcal{N}}$, where $\mathcal{M} = (F^{\mathcal{M}}, V_{\mathcal{M}})$ and $\mathcal{N} = (F^{\mathcal{N}}, V_{\mathcal{N}})$ are κ-models based on the κ-frames $F^{\mathcal{M}} = (W_{\mathcal{M}}, \sim^{\mathcal{M}})$ and $F^{\mathcal{N}} = (W_{\mathcal{N}}, \sim^{\mathcal{N}})$, is a functional bisimulation if and only if:*

1. ***Atom**: for all $w \in W_{\mathcal{M}}$, $V_{\mathcal{M}}(w) = V_{\mathcal{N}}(f(w))$;*
2. ***Forth**: for any agent pattern G, $w \sim^{\mathcal{M}}_G v$ implies $f(w) \sim^{\mathcal{N}}_G f(v)$;*
3. ***Back**: for any agent pattern G, $f(w) \sim^{\mathcal{N}}_G v'$ implies that there exists w' with $f(w') = v'$ such that $w \sim^{\mathcal{M}}_G w'$.*

Lemma 57. *If $f : W_{\mathcal{M}} \to W_{\mathcal{N}}$ is a functional bisimulation, then for any formula φ we have $\mathcal{M}, w \Vdash \varphi$ if and only if $\mathcal{N}, f(w) \Vdash \varphi$.*

Proof. Let f be a functional bisimulation. We show the claim by induction on the length of φ.

1. If $\varphi \in \mathsf{Prop}$, then the claim follows by **Atom**.
2. $\varphi \equiv \neg\phi$. Follows by the induction hypothesis.
3. $\varphi \equiv \phi \wedge \psi$. Follows by the induction hypothesis.
4. $\varphi \equiv [G]\phi$. We show equivalently that

$$\mathcal{M}, w \nVdash [G]\phi \quad \text{iff} \quad \mathcal{N}, f(w) \nVdash [G]\phi.$$

From left to right, assume $\mathcal{M}, w \nVdash [G]\phi$. Hence, there exists $v \in W_{\mathcal{M}}$ with $w \sim^{\mathcal{M}}_G v$ and $\mathcal{M}, v \nVdash \phi$. By **Forth** we have that $f(w) \sim^{\mathcal{N}}_G f(v)$ and by the induction hypothesis we have $\mathcal{N}, f(v) \nVdash \phi$. Therefore, it holds that $\mathcal{N}, f(w) \nVdash [G]\phi$. For the other direction, assume $\mathcal{N}, f(w) \nVdash [G]\phi$. Hence, there exists $v' \in W_{\mathcal{N}}$ with $f(w) \sim^{\mathcal{N}}_G v'$ and $\mathcal{N}, v' \nVdash \phi$. By **Back** there exists $w' \in W_{\mathcal{M}}$ such that $f(w') = v'$ and $w \sim^{\mathcal{M}}_G w'$. By the induction hypothesis we obtain $\mathcal{M}, w' \nVdash \phi$ and thus $\mathcal{M}, w \nVdash [G]\phi$. □

Let $F = (W, \sim)$ be a κ-frame and consider a κ-model $\mathcal{M} = (F, V)$ as well as its unraveled model $U(\mathcal{M}) = (U(F), L)$. It is straightforward to show that the mapping $\mathsf{last} : H_F \to W$ that maps each $h \in H_F$ to $\ell(h) \in W$ is a functional bisimulation. We only show the claim for **Back**. Assume that $w \in H_F$ and $v' \in W$ such that $\ell(w) \sim_G v'$. It follows that, $w' = w \parallel (\ell(w), U, v')$ with $G \subseteq U$ is a valid history with $w' \approx_G w$ and $\ell(w') = v'$. Finally, we can apply Lemma 57 to the canonical model and its unraveled model, which establishes completeness of $\mathsf{Syn}^{\text{-}}$ with respect δ-models (Theorem 8).

Theorem 8. $\mathsf{Syn}^{\text{-}}$ *is sound and complete with respect to δ-models.*

It only remains to show completeness of Syn with respect to proper δ-models. The next section formally defines properness, and shows how to make the unraveled canonical model proper.

5.5.3 Properness

This section introduces the notion of properness on pattern frames. As it turns out, the unraveled canonical model need not be proper. However, it can be made proper by constructing a quotient model. This suffices to show that Syn is sound and complete with respect to proper δ-models. Different to Chapter 4, we define properness for symmetric and transitive frames.

Definition 70 specifies the maximal alive agent pattern of a world w. Two worlds w and v are equivalent (Definition (71) if and only if they have the same maximal alive agent pattern, and that agent pattern cannot distinguish between them. A frame is *proper* (Definition 72) if and only if two worlds being equivalent implies them being equal.

Definition 70. *Let $F = (W, \sim)$ be a symmetric and transitive frame. For $w \in W$, we define*

$$\overline{w} = \{B \subseteq \mathsf{Ag} \setminus \{\emptyset\} \mid \exists G. B \in G \text{ and } w \in \mathsf{Alive}(G)_F\}.$$

Definition 71 ($\equiv$)**.** *Let $F = (W, \sim)$ be a symmetric and transitive frame. We define the relation $\equiv$ on $W \times W$ as*

$$w \equiv v \quad \text{iff} \quad \overline{w} = \overline{v} \text{ and } w \sim_{\overline{w}} v.$$

Definition 72 (Proper)**.** *A symmetric and transitive frame $F = (W, \sim)$ is called proper if and only if for all $w, v \in W$, $w \equiv v$ implies $w = v$. A pattern model $\mathcal{M} = (F, V)$ is proper if and only if F is proper.*

To see why the unraveled canonical model is not proper, let $\Gamma \in W^c$ be a world of the canonical frame in which only one agent, say a, is alive. It is straightforward to verify that such a world exists. Since F^c is reflexive for alive agents, we have that $\Gamma \sim_a^c \Gamma$. Consequently, the unraveled frame contains the two different histories $h = (\Gamma, \{\{a\}\}, \Gamma)$ and $h' = h \parallel (\Gamma, \{\{a\}\}, \Gamma)$. Agent a is the only agent alive in both histories. By definition, agent a cannot distinguish between h and h', which violates properness.

If G is an agent pattern, we denote the set of maximal elements of G with $\max(G)$. Notice that $\max(\overline{w})$ always contains exactly one element. Definition 73 and Lemma 58 show under which conditions symmetric and transitive frames can be made proper. Since κ-frames are symmetric and transitive, the results carry over to them.

Definition 73 ($\mathcal{M}^\rho$)**.** *Let $F = (W, \sim)$ be a symmetric and transitive frame, and let $\mathcal{M} = (F, V)$ be a pattern model. We define the frame $F^\rho = (W^\rho, \sim^\rho)$ and the pattern model $\mathcal{M}^\rho = (F^\rho, V^\rho)$ as:*

1. *$W^\rho = W/\equiv$ is the set of equivalence classes of $\equiv$;*
2. *$[w] \sim_G^\rho [v]$ if and only if $w \sim_G v$;*
3. *for any $p \in \mathsf{Prop}$, $p \in V^\rho([w])$ if and only if $p \in V(w)$.*

$\mathcal{M}^\rho$ is well-defined, if for any two worlds $w, v \in W$ and $p \in \mathsf{Prop}$,

$$w \equiv v \text{ implies } p \in V(w) \iff p \in V(v).$$

Lemma 58. *Let $F = (W, \sim)$ be a symmetric and transitive frame, and let $\mathcal{M} = (F, V)$ be a pattern model such that $\mathcal{M}^\rho$ is well-defined. We find that:*

1. *$\mathcal{M}^\rho$ is proper;*
2. *$\mathcal{M}, w \Vdash \phi \quad$ if and only if $\quad \mathcal{M}^\rho, [w] \Vdash \phi$.*

Proof. In order to show that $\mathcal{M}^\rho$ is proper, observe that $[w] \equiv [v]$ implies $w \equiv v$. Indeed, since $[w] \sim_G^\rho [v]$ if and only if $w \sim_G v$, it holds that $\overline{[w]} = \overline{w}$ and $w \sim_{\overline{w}} v$. Hence, w and v belong to the same equivalence class, i.e., $[w] = [v]$.

For the second claim, we show the direction from right to left by induction on the length of ϕ. The other direction is symmetric. The base case follows because $\mathcal{M}^\rho$ is well-defined. The only case left is $\phi = [G]\psi$. Assume $\mathcal{M}^\rho, [w] \Vdash [G]\psi$. We need to show $\mathcal{M}, w \Vdash [G]\psi$. Let $v \in W$ be such that $w \sim_G v$, i.e., $[w] \sim_G^\rho [v]$. By the definition of truth $\mathcal{M}^\rho, [v] \Vdash \psi$ and by the induction hypothesis $\mathcal{M}, v \Vdash \psi$, which concludes the proof. $\square$

Remark 14 lists some useful properties of κ-models needed below.

Remark 14. *Let $F = (W, \sim)$ be a frame such that $\mathcal{M} = (F, V)$ is a κ-model. The following can be shown by using the properties of κ-models and the fact that $\overline{w}$ is maximal under subsets:*

1. $\mathcal{M}, w \Vdash \mathsf{alive}(\overline{w})$ *and* $\mathcal{M}, w \Vdash \mathsf{dead}(\overline{w}^C)$ *are always the case;*
2. $\mathcal{M}, w \Vdash \mathsf{alive}(G) \wedge \mathsf{dead}(G^C)$ *if and only if* $\max(G) = \max(\overline{w})$*;*
3. *if* $\max(G) = \max(\overline{w})$*, then* $w \sim_G v$ *and* $\mathcal{M}, v \Vdash \mathsf{dead}(G^C)$ *imply* $\overline{w} = \overline{v}$*.*

Proving soundness and completeness of Syn with respect to proper κ-models requires us to show that proper κ-models validate P, and that the canonical model for Syn is proper.

Lemma 59 (Soundness)**.** Syn *is sound with respect to proper κ-models.*

Proof. We showed the cases for Syn^{-} in the proof of Theorem 5. Hence, we only need to show the case for P. Let $F = (W, \sim)$ be a proper κ-frame, and let $\mathcal{M} = (F, V)$ be a κ-model. Assume $\mathcal{M}, w \Vdash \mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \wedge \phi$. By Remark 14, we find that $\max(G) = \max(\overline{w})$. This implies that for any $v \in W$ with $w \sim_G v$ such that $\mathcal{M}, v \Vdash \mathsf{dead}(G^C)$, we have $\overline{w} = \overline{v}$, i.e., $w \equiv v$ and by the properness of F, it follows that $w = v$. Therefore $\mathcal{M}, v \Vdash \phi$. □

The canonical model $\mathcal{M}^c$ for Syn is defined as in Definition 42, but worlds are maximal consistent sets for Syn instead of Syn^{-}. For simplicity, we do not distinguish between the two by the means of notation.

Theorem 9. *$\mathcal{M}^c$ is a proper κ-model.*

Proof. We already showed the properties of κ-models in the proof of Theorem 6. Hence, it suffices to show that F^c is proper. Let $\Gamma, \Delta \in W^c$ such that $\Gamma \equiv \Delta$, i.e., $\overline{\Gamma} = \overline{\Delta} = G$ and $\Gamma \sim^c_G \Delta$. We now show $\Gamma = \Delta$, i.e., for any $\phi \in \mathcal{L}^{\mathsf{syn}}$, $\phi \in \Gamma$ if and only if $\phi \in \Delta$. We show the direction from left to right. The other direction is symmetric. Assume that $\phi \in \Gamma$. Since Γ is a maximal consistent set, it follows by Remark 14 and Lemma 49 that $\mathsf{alive}(G) \wedge \mathsf{dead}(G^C) \wedge \phi \in \Gamma$. Furthermore, by P, it follows that $[G](\mathsf{dead}(G^C) \rightarrow \phi) \in \Gamma$ and thus $\mathsf{dead}(G^C) \rightarrow \phi \in \Delta$. By assumption, $\mathsf{dead}(G^C) \in \Delta$ and thus $\phi \in \Delta$, because Δ is a maximal consistent set. □

Corollary 10. Syn *is sound and complete with respect to proper κ-models.*

Lemma 60 below shows that the construction of Definition 73 can be applied to the unraveled canonical model.

Lemma 60. *Let $U(\mathcal{M}^c) = (H, L)$ be the unraveled canonical model. For any $h, h' \in H$ with $h \equiv h'$ and $p \in \mathsf{Prop}$ it holds that*

$$p \in L(h) \quad \textit{iff} \quad p \in L(h').$$

Proof. Consider two histories h, h' of $U(\mathcal{M}^c)$ with $h \equiv h'$, i.e., $\overline{h} = \overline{h'} = G$ and $h \approx_G h'$. Let $\ell(h) = \Gamma \in W^c$ and $\ell(h') = \Delta \in W^c$. By Remark 13, $\ell(h) \sim_G \ell(h')$, i.e., $\Gamma \sim^c_G \Delta$. We now show the direction from left to right. The other direction is symmetric. Let $p \in \mathsf{Prop}$ with $p \in L(h)$, i.e., $p \in V^c(\Gamma)$. Since Γ is a maximal consistent set it follows by Remark 14 and Lemma 49 that $\mathsf{alive}(G) \wedge \mathsf{dead}(G^c) \wedge p \in \Gamma$. By P, $[G](\mathsf{dead}(G^C) \to p) \in \Gamma$. Therefore, $\mathsf{dead}(G^C) \to p \in \Delta$. Lastly, by assumption and Remark 14, we have $\mathsf{dead}(G^C) \in \Delta$ and since Δ is a maximal consistent set, it follows that $p \in \Delta$. Therefore, $p \in V^c(\ell(h'))$, and by definition $p \in L(h')$. □

Applying the construction in Definition 73 to the unraveled canonical model yields the following corollary.

Corollary 11. *$U(\mathcal{M}^c)^\rho$ is proper and*

$$U(\mathcal{M}^c)^\rho, [h] \Vdash \phi \quad \textit{iff} \quad U(\mathcal{M}), h \Vdash \phi.$$

By Corollary 11, we obtain that Syn is sound and complete with respect to proper δ-models.

Corollary 12. *For all $\phi \in \mathcal{L}^{\mathsf{Syn}}$, it holds that:*

$$\vdash \phi \textit{ if and only if } \Vdash_\delta \phi.$$

5.5.4 δ-translations

In this section, we show how every proper δ-model can be transformed to an equivalent multi-simplicial model. Completeness of Syn with respect to multi-simplicial models follows immediately. A δ-translation is a multi-simplicial model that represents a δ-model based on a proper frame.

Definition 74 (δ-translation)**.** *Let $F = (W, \sim)$ be a δ-frame and $\mathcal{M} = (F, V)$ be a δ-model. A multi-simplicial model $\mathcal{C} = (\mathbb{C}, L)$ is a* δ-translation *of $\mathcal{M}$ if and only if there exists a mapping $T : W \to \mathbb{C}$ such that for all $w, v \in W$:*

1. *$G \subseteq (T(w) \cap T(v))^\circ$ iff $w \sim_G v$;*
2. *$p \in L(T(w))$ iff $p \in V(w)$.*

We present a general construction of a δ-translation for proper δ-frames, as well as an algorithmic one for δ-models based on finite proper δ-frames.

General δ-translation

Let $F = (W, \sim)$ be a proper δ-frame, and let $\mathcal{M} = (F, V)$ be a δ-model. We provide a general translation based on equivalence classes. The construction is similar to the one presented in Chapter 2, where we associated equivalence classes $[w]_a$ with vertices. In the context of group knowledge, we can map equivalence classes of the form $[w]_{\{G\}}$, where $w \in W$ and $\emptyset \neq G \subseteq \mathsf{Ag}$, to faces of an explicit simplex. The intuition is that the set

$$[w]_{\{G\}} = \{v \mid w \sim_{\{G\}} v\} \neq \emptyset,$$

represents the G-boundary of the explicit simplex that corresponds to w. The explicit simplices that correspond to worlds $v \in [w]_{\{G\}}$ share the same G-boundary with w. The explicit simplex induced by a world $w \in W$ is of the form $\{(G, [w]_{\{G\}}) \mid [w]_{\{G\}} \neq \emptyset\}$, where G is a subset of agents alive in w. Since $[w]_{\{G\}}$ is an equivalence class and $w \in [w]_{\{G\}}$, the explicit simplices inherit closure properties of the proper δ-frame F.

Definition 75. *Let $F = (W, \sim)$ be a proper δ-frame. For $w \in W$ and $\emptyset \neq G \subseteq \mathsf{Ag}$, we define:*

$$[w]_{\{G\}} = \{v \in W \mid w \sim_{\{G\}} v\} \text{ and } T(w) = \{(G, [w]_{\{G\}}) \mid [w]_{\{G\}} \neq \emptyset\}.$$

We also write $T = \{T(w) \mid w \in W\}$.

Since F is a proper δ-frame, it is straightforward to show that each $T(w)$ is an explicit simplex (Lemma 61), and that T is a complex (Lemma 62). Moreover, Lemma 63 states that $T(\cdot)$ is a bijection.

Lemma 61. *For all $w \in W$, $T(w)$ is an explicit simplex.*

Proof. Since F satisfies NE, it holds that $T(w) \neq \emptyset$, for all $w \in W$. We show that $T(w)$ satisfies $\mathsf{S1}, \mathsf{S2}$, and $\mathsf{S3}$.

- $\mathsf{S1}$: Let $(G, [w]_{\{G\}}) \in T(w)$ and $(H, [w]_{\{H\}}) \in T(w)$ be two distinct maximal elements. This immediately implies $G \neq H$. By Definition 75, we obtain that $w \in \mathsf{Alive}(\{G\}) \cap \mathsf{Alive}(\{H\})$. Since δ-frames satisfy $\mathsf{K3}$, we obtain that $w \in \mathsf{Alive}(\{G, H\})$. Moreover, by $\mathsf{K2}$ we find that $w \in \mathsf{Alive}(\{G \cup H\})$. Therefore, it holds that $(G \cup H, [w]_{\{G \cup H\}}) \in T(w)$. Consequently, $(G, [w]_{\{G\}})$ and $(H, [w]_{\{H\}})$ cannot be maximal, which is a contradiction.

- $\mathsf{S2}$: Let $(B, [w]_{\{B\}}) \in T(w)$ and $\emptyset \neq C \subseteq B$. By Lemma 46, it holds that $\sim_{\{B\}} \subseteq \sim_{\{C\}}$. By assumption, we have that $w \sim_{\{B\}} v$, which implies that $w \sim_{\{C\}} v$, and thus $[w]_{\{C\}} \neq \emptyset$. Consequently, it holds that $(C, [w]_{\{C\}}) \in T(w)$. Uniqueness of $[w]_{\{C\}}$ follows by construction because $[w]_{\{C\}}$ is an equivalence class.

S3: Let $\max(T(w)) = (G, [w]_{\{G\}})$ and suppose that $(H, [w]_{\{H\}}) \in T(w)$ with $H \not\subseteq G$. By Definition 75, we have $w \sim_{\{G\}} w$ and $w \sim_{\{H\}} w$. Since the δ-frame satisfies K3, we also have $w \sim_{\{G,H\}} w$, and by K2 it holds that $w \sim_{\{G\cup H\}} w$. As a result $(G \cup H, [w]_{\{G\cup H\}}) \in T(w)$, but this contradicts $(G, [w]_{\{G\}})$ being the maximal element of $T(w)$. □

Lemma 62. *T is a multi-simplicial complex.*

Proof. We need to show that T satisfies C. Let $T(w), T(v) \in T$ such that $(G, [w]_{\{G\}}) \in T(w) \cap T(v)$. By construction, $w \sim_{\{G\}} v$. Therefore, by Lemma 46, we obtain $w \sim_{\{H\}} v$ for all $\emptyset \neq H \subseteq G$. As a result, the explicit simplices $T(w)$ and $T(v)$ share all elements of the form $(H, [w]_{\{H\}})$, and T satisfies condition C. □

Lemma 63. *$T(\cdot)$ is a bijection.*

Proof. Surjectivity follows by the construction of T. Regarding injectivity, assume that $T(w) = T(v)$. This implies that $w \equiv v$. By properness of F we obtain that $w = v$. □

A consequence of Lemma 63 is that the labeling introduced in Definition 76 is well-defined.

Definition 76 (δ-labeling)**.** *Let V be a valuation on a set of worlds W. We define the δ-labeling L on T as follows:*

$$p \in L(T(w)) \textit{ iff } p \in V(w).$$

The last step is to show that our construction preserves the indistinguishability relation (Lemma 64).

Lemma 64. *$T(w) \sim_G^T T(v)$ if and only if $w \sim_G v$.*

Proof. By definition, G is not empty. Since F satisfies D, we can assume without loss of generality that G contains exactly one element $A \subseteq \mathsf{Ag}$. It holds that $(A, [w]_{\{A\}}) \in T(w) \cap T(v)$ iff $[w]_{\{A\}} = [v]_{\{A\}} \neq \emptyset$ iff $w \sim_{\{A\}} v$. □

We obtain the following corollary from Lemma 64 and Definition 76.

Corollary 13. *$\mathcal{C} = (T, L)$, where L is the δ-labeling, is a δ-translation of $\mathcal{M}$.*

Theorem 10. *Let $F = (W, \sim)$ be a proper δ-frame, and consider the δ-model $\mathcal{M} = (F, V)$. Further, let $\mathcal{C} = (\mathbb{C}, L)$ be a δ-translation of $\mathcal{M}$. It holds that*

$$\mathcal{M}, w \Vdash \phi \textit{ if and only if } \mathcal{C}, T(w) \Vdash_\sigma \phi.$$

Proof. By induction on the length of formulas.

1. Let $\phi \equiv p \in \mathsf{Prop}$. We have $\mathcal{M}, w \Vdash \phi$ iff $p \in V(w)$ iff $p \in L(T(w))$ iff $\mathcal{C}, T(w) \Vdash_\sigma \phi$ (by the definition of L).

2. Let $\phi \equiv \neg\psi$. Follows by the induction hypothesis.

3. Let $\phi \equiv \psi \wedge \varphi$. Follows by the induction hypothesis.

4. Let $\phi \equiv [G]\psi$. We equivalently show

$$\mathcal{M}, w \nVdash [G]\psi \text{ iff } \mathcal{C}, T(w) \nVdash_\sigma [G]\psi.$$

It holds that $\mathcal{M}, w \nVdash [G]\psi$ iff there exists $v \in W$ with $w \sim_G v$ and $\mathcal{M}, v \nVdash \psi$ iff $G \subseteq (T(w) \cap T(v))^\circ$ (Lemma 64) and $\mathcal{C}, T(v) \nVdash_\sigma \psi$ (by hypothesis) iff $\mathcal{C}, T(w) \nVdash_\sigma [G]\psi$ by definition. □

Hence, if $\nvdash \varphi$, then there exists a proper δ-model $\mathcal{M}$ (Corollary 12) such that $\mathcal{M} \nVdash \varphi$. By Corollary 13, we can construct a δ-translation $\mathcal{C} = (T, L)$ of $\mathcal{M}$ such that $\mathcal{C} \nVdash_\sigma \varphi$. Thus, Syn is sound and complete with respect to multi-simplicial models.

Corollary 14. $\vdash \varphi$ *if and only if* $\Vdash_\sigma \varphi$.

Algorithmic Translation for Finite Frames

Given a proper δ-frame $F = (W, \sim)$ such that W finite, Construction 1 shows how we can algorithmically build a δ-translation of F.

We assume an arbitrary enumeration of worlds and write w_i for the i-th world. The simplicial image of a world w_i under a mapping T, i.e., $T(w_i)$, is denoted by S_i. The next lemmas show that there exists a multi-simplicial complex $\mathbb{C}$ that preserves the structure of the proper frame F. Let $w_i^* \subseteq \mathsf{Ag}$ be the maximum set of all agents alive in w_i. Construction 1 on input F first initializes a simplex $S_i = \{(A, i) \mid A \subseteq w_i^*\}$ for each world w_i (lines 3 to 5). At this point, no two different simplices S_i and S_j are connected. Throughout the transformation (lines 6 to 16), Construction 1 glues related simplices together according to the indistinguishability relation $\sim$ of the frame F. It iterates through all pairs (w_i, w_j) with $i < j$, which suffices by the symmetry of $\sim_G$, and checks for all G, whether $(w_i, w_j) \in \sim_G$. If so, for each $B \in G$, the pair (B, j) is replaced by the pair $(B, k) \in S_i$, where k is the smallest index such that $w_k \sim_{\{B\}} w_i$. After the replacement, the simplices S_i and S_j are connected. Example 13 shows a possible execution.

We will now prove the correctness of Construction 1. Specifically, we show that the multi-simplicial model $\mathcal{C} = (\mathbb{C}, L)$, where $\mathbb{C}$ is the complex returned by Construction 1 (line 20), and L is a labeling such that $p \in L(S_i)$ if and only if $p \in V(w_i)$, is a δ-translation of $\mathcal{M}$.

Construction 1 δ-translation

1: **Input**
2: A proper δ-frame $F = (W, \sim)$

3: **Initialization**
4: $w_i^* = \max\{A \subseteq \mathsf{Ag} \mid w_i \sim_{\{A\}} w_i\}$
5: $S_i = \{(A, i) \mid A \subseteq w_i^*\}$ for $1 \leq i \leq n$

6: **Transformation**
7: $i = 1$
8: $j = 1$
9: **while** w_i exists **do**
10: $j \leftarrow i + 1$
11: **while** w_j exists **do**
12: **for each** $G \in \mathsf{AP}$ **with** $w_i \sim_G w_j$ **do**
13: **for each** $B \in G$ **do**
14: $k \leftarrow \min\{l \mid w_l \sim_{\{B\}} w_i\}$
15: $S_j \leftarrow S_j \setminus \{(B, j)\}$
16: $S_j \leftarrow S_j \cup \{(B, k)\}$
17: $j \leftarrow j + 1$
18: $i \leftarrow i + 1$

19: **Output**
20: $\mathbb{C} = \{S_i \mid w_i \in W\}$

Example 13. *Let* $\mathsf{Ag} = \{a, b\}$ *and consider the proper* δ*-frame* $F = (W, \sim)$ *depicted in Figure 5.7. Construction 1 first initializes:*

$$S_1 = \left\{ \begin{matrix} ab1 \\ a1, b1 \end{matrix} \right\}, \; S_2 = \left\{ \begin{matrix} ab2 \\ a2, b2 \end{matrix} \right\}, \text{ and } S_3 = \left\{ \begin{matrix} ab3 \\ a3, b3 \end{matrix} \right\}.$$

During the transformation phase, when $i = 1$ *and* $j = 2$*, Construction 1 replaces* $(a, 2)$ *and* $(b, 2)$ *with* $(a, 1)$ *and* $(b, 1)$ *since it holds that* $w_1 \sim_{\{a\}} w_2$ *and* $w_1 \sim_{\{b\}} w_2$*. Moreover, it replaces* $(b, 3) \in S_3$ *with* $(b, 1)$*. Observe that if Line 14 was missing, then, for* $i = 2$ *and* $j = 3$*, the Construction would add* $(b, 2)$ *to* S_3 *as well, which would make it an ill-formed simplex. The resulting complex is:*

$$\mathbb{C} = \left\{ \left\{ \begin{matrix} ab1 \\ a1, b1 \end{matrix} \right\}, \left\{ \begin{matrix} ab2 \\ a1, b1 \end{matrix} \right\}, \left\{ \begin{matrix} ab3 \\ a3, b1 \end{matrix} \right\} \right\}.$$

The next Lemma ensures that we can safely assume the existence of a unique and non-empty S_i for each w_i throughout our proofs.

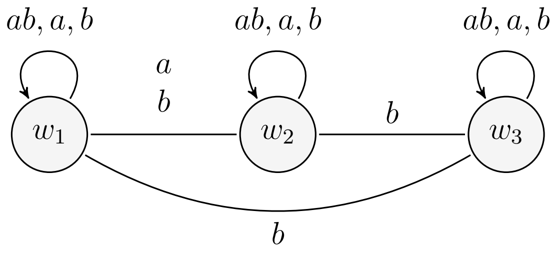

Figure 5.7: The δ-model for Example 13. Only arrows for maximal agent patterns are shown.

Lemma 65 (Uniqueness)**.** *Let $F = (W, \sim)$ be a proper δ-frame. After the initialization of Construction 1 on input F, and after each execution of Line 16, the following holds for all worlds $w_i, w_j \in W$:*

1. *$S_i \neq \emptyset$;*
2. *$S_i = S_j$ if and only if $i = j$.*

Proof. $S_i \neq \emptyset$ follows immediately because δ-frames satisfy NE and thus, each S_i is initialized to some non-empty set. Furthermore, since each element that is removed gets replaced, we conclude that $S_i \neq \emptyset$ for all $w_i \in W$ after Line 16. For the second claim, the direction from right to left follows immediately. For the other direction, observe that if $S_i = S_j$, then $\overline{w_i} = \overline{w_j}$ and $w_i \sim_{\overline{w_i}} w_j$, i.e., $w_i \equiv w_j$. By the properness of F, it follows that $w_i = w_j$, i.e., $i = j$. □

Since the smaller index gets precedence, some elements of a simplex may never be exchanged. For example, the simplex S_1 in Example 13 remains unchanged throughout Construction 1. Lemma 66 shows that such simplices are well-formed. Moreover, Lemma 67 states that replacements are final.

Lemma 66. *Let $F = (W, \sim)$ be a proper δ-frame and consider Construction 1 on input F. After the initialization (lines 3 to 5), for all $w_j \in W$, it holds that $(A, j) \in S_j$ iff $w_j \sim_{\{A\}} w_j$.*

Proof. From left to right, we have $(A, j) \in S_j$ if and only if $A \subseteq w_j^*$. Since $w_j \sim_{\{w_j^*\}} w_j$ by definition, we obtain $w_j \sim_{\{A\}} w_j$ by Lemma 46. Regarding the other direction, assume $w_j \sim_{\{A\}} w_j$. By the definition of w_j^* it follows that $A \subseteq w_j^*$ which implies that $(A, j) \in S_j$ after the initialization phase. □

Lemma 67. *Let $F = (W, \sim)$ be a proper δ-frame and consider $i, j, k \in \mathbb{N}$ such that $i \neq j, i \neq k$, and $j \neq k$. There does not exist $B \subseteq \mathsf{Ag}$ such that, while running Construction 1 on input F, (B, i) is replaced with (B, j) and (B, j) is replaced with the pair (B, k).*

Proof. Towards a contradiction, assume that there exists such a $B \subseteq \mathsf{Ag}$, i.e., for some S_i, (B, i) is replaced with (B, j) and (B, j) is replaced with (B, k). By assumption, $i < j < k$ because Construction 1 replaces (B, l) with (B, m) only if $l > m$. Since lines 15 and 16 are executed at least twice, there exist $G, G' \in \mathsf{AP}$ such that $B \in G$ and $B \in G'$ with $w_i \sim_G w_j$ and $w_j \sim_{G'} w_k$. Since $\{B\} \subseteq G$ and $\{B\} \subseteq G'$, we have by $\mathsf{K3}$ that $w_i \sim_{\{B\}} w_j$ and $w_j \sim_{\{B\}} w_k$. Further, $w_i \sim_{\{B\}} w_k$ follows by transitivity of $\sim_{\{B\}}$. But this means, that $i = k$ because (B, k) was replaced by (B, i) prior which contradicts that $i < k$. □

The next lemmas show that after the Construction 1, elements of $\mathbb{C}$ are well-formed simplices and $\mathbb{C}$ is a well-formed complex preserving the indistinguishability relation.

Lemma 68. *Let $F = (W, \sim)$ be a proper δ-frame. The output $\mathbb{C}$ of Construction 1 on input F satisfies the following two properties:*

$\mathsf{T1}$*: Let $S \in \mathbb{C}$. If $(A, j) \in S$ and $(A, k) \in S$, then $j = k$.*

$\mathsf{T2}$*: Let $w_i, w_j \in W$ and $G \in \mathsf{AP}$, then:*

$$\forall B \in G.\exists k \in \mathbb{N}.(B, k) \in S_i \wedge (B, k) \in S_j \quad \textit{iff} \quad w_i \sim_G w_j.$$

Proof. $\mathsf{T1}$ follows by construction. Regarding $\mathsf{T2}$, we start by showing the direction from left to right. Assume that:

$$\forall B \in G.\exists k \in \mathbb{N}.(B, k) \in S_i \wedge (B, k) \in S_j.$$

Without loss of generality, we fix $i \leq j$. We have that $(B, k) \in S_i \cap S_j$ only if there exists G' with $w_i \sim_{G'} w_j$ and $B \in G'$. By $\mathsf{K3}$ we get $w_i \sim_{\{B\}} w_j$. Since B is arbitrary, D implies that $w_i \sim_G w_j$.

For the other direction, let $w_i \sim_G w_j$ and let $B \in G$. By construction, $(B, j) \in S_j$ is replaced by (B, k) with $k \leq i$. Let k be the smallest index such that $w_k \sim_G w_i$, i.e., if $k < i$, then $(B, i) \in S_i$ was replaced with (B, k) before. Lemma 66 ensures that $(B, k) \in S_k$. Hence, if $k < i$, $(B, j) \in S_j$ and $(B, i) \in S_i$ are both replaced by (B, k) in the same iteration. If $k = i$, only (B, j) is replaced by (B, i), and $(B, i) \in S_i$ at the end of the Construction due to the symmetry of $\sim_G$ and k being minimal. By Lemma 67 there cannot be any further replacements and thus $(B, k) \in S_i$ and $(B, k) \in S_j$ at the end of Construction 1. □

A consequence of $\mathsf{T2}$ is that:

$$S_i \sim_G^{\mathbb{C}} S_j \text{ if and only if } w_i \sim_G w_j. \tag{5.13}$$

Lemma 69. *Let $F = (W, \sim)$ be a proper δ-frame. Let $\mathbb{C}$ be the output returned by Construction 1 on input F. For each $w_i \in W$, the explicit simplex $S_i \in \mathbb{C}$ is well-formed.*

Proof. We show that S_i satisfies $\mathsf{S1}, \mathsf{S2}$, and $\mathsf{S3}$.

- $\mathsf{S1}$: Let $(A, j) \in S_i$ and $(B, k) \in S_i$ be two distinct maximal elements. First, we note that $A \neq B$ due to $\mathsf{T1}$. By $\mathsf{T2}$, $w_i \in \mathsf{Alive}(\{A\})$ and $w_i \in \mathsf{Alive}(\{B\})$. By $\mathsf{K3}$, $w_i \in \mathsf{Alive}(\{A\} \cup \{B\})$. Further, by $\mathsf{K2}$, $w_i \in \mathsf{Alive}(\{A \cup B\})$. Hence, $(A \cup B, l) \in S_i$ for some l. But since $A \subsetneq A \cup B$, this contradicts that (A, i) is the maximal element of S_i.

- $\mathsf{S2}$: Let $(B, j) \in S_i$ and $\emptyset \neq C \subseteq B$, we need to show that there exists a unique k such that $(C, k) \in S_i$. By Lemma 46 we have $\sim_{\{B\}} \subseteq \sim_{\{C\}}$. Since we assume that $(B, j) \in S_i$, we get $w_i \sim_{\{B\}} w_i$ by $\mathsf{T2}$. Due to $\sim_{\{B\}} \subseteq \sim_{\{C\}}$, we have that $w_i \sim_{\{C\}} w_i$. Thus, by $\mathsf{T2}$, there exists $k \in \mathbb{N}$ such that $(C, k) \in S_i$. Condition $\mathsf{T1}$ ensures that k is unique.

- $\mathsf{S3}$: Let $\max(S_i) = (B, k)$ and suppose that $(A, j) \in S_i$ for $A \nsubseteq B$. We have $w_i \sim_{\{A\}} w_i$ and $w_i \sim_{\{B\}} w_i$ by $\mathsf{T2}$. By $\mathsf{K3}$, $w_i \in \mathsf{Alive}(\{A\} \cup \{B\})$ follows, and by $\mathsf{K2}$ we have that $w_i \in \mathsf{Alive}(\{A \cup B\})$. Therefore it holds that $(A \cup B, l) \in S_i$ for some l which contradicts the maximality of (B, k). □

Lemma 70. *Let $F = (W, \sim)$ be a proper δ-frame and consider Construction 1 on input F. The output set $\mathbb{C} = \{S_i \mid w_i \in W\}$ is a complex.*

Proof. In order to show that $\mathbb{C}$ is a complex, we need to show that it satisfies Condition C. Consider the simplices $S_m, S_n \in \mathbb{C}$ and assume that there exists $A \subseteq \mathsf{Ag}$ and i with $(A, i) \in S_m$ and $(A, i) \in S_n$. We need to show that for all $\emptyset \neq B \subseteq A$ and all $j \in \mathbb{N}$ it holds that

$$(B, j) \in S_m \quad \text{iff} \quad (B, j) \in S_n.$$

Since S_m and S_n are arbitrary, it is enough to show only one direction. Assume that $(B, j) \in S_m$. By $\mathsf{T2}$, we have $w_m \sim_{\{A\}} w_n$ and by Lemma 46, it holds that $w_m \sim_{\{B\}} w_n$. Since $(B, j) \in S_m$, $\mathsf{T2}$ implies $w_m \sim_{\{B\}} w_j$. Further, by symmetry and transitivity we obtain $w_j \sim_{\{B\}} w_n$. Hence, it holds that $j \leq n$. Therefore, (B, j) replaced (B, n) in S_n. By Lemma 67, no more replacements of that pair can happen during the execution of Construction 1 and we conclude $(B, j) \in S_n$. □

Finally, we obtain that Construction 1 together with an appropriate labeling is a δ-translation.

Theorem 11. *Let $F = (W, \sim)$ be a proper δ-frame, and let $\mathcal{M} = (F, V)$ be δ-model. Further, let $\mathcal{C} = (\mathbb{C}, L)$ be such that $\mathbb{C}$ is the output of Construction 1 on input F and $L : \mathbb{C} \to \mathsf{Pow}(\mathsf{Prop})$ is a labeling such that:*

$$p \in L(S_i) \quad \textit{iff} \quad p \in V(w_i). \tag{5.14}$$

It holds that $\mathcal{C}$ is a δ-translation of $\mathcal{M}$.

Proof. By Lemma 70, $\mathbb{C}$ is a complex. It follows from (5.13) and (5.14) that $\mathcal{C}$ is a δ-translation of $\mathcal{M}$. □

5.6 Communication

In this section, we will explore a different reading of agent patterns, namely as a description of the communication happening between the agents. Let G be the agent pattern $\{\{a\}, \{b, c\}\}$. We interpret this as the agents *b and c communicating with each other*, but *there is no communication between a and b or c*. A formula $[G]\phi$ will thus be interpreted as *a knows ϕ* ***and*** *the group b, c has distributed knowledge of ϕ*. We can also distinguish the patterns $\{\{a, b\}, \{b, c\}\}$ and $\{\{a, b\}, \{b, c\}, \{a, c\}\}$. In the first one, a and c can only communicate via b whereas in the second one, a and c have a direct communication channel.

Definition 77 (Connected)**.** *Let $C \subseteq \mathsf{Pow}(\mathsf{Ag})$, two elements $X, Y \in C$ are called* connected in C *if and only if there exist $Z_0, ..., Z_k \in C$ such that $Z_i \cap Z_{i+1} \neq \emptyset$ for $0 \leq i < k$ and $Z_0 = X$ and $Z_k = Y$.*

Definition 78 (Connected Component)**.** *Let $C \in \mathsf{AP}$, we call C a* connected component *if and only if for any $X, Y \in C$ with $X \neq Y$ it holds that X and Y are connected in C. Let $G \in \mathsf{AP}$. We call $H \neq \emptyset$ a* maximal connected component of G *if and only if $H \subseteq G$ and there is no connected component $H' \subseteq G$ such that H is a proper subset of H'.*

We can represent an agent pattern G as the union of its maximal connected components. Let $C_1, \ldots, C_k$ be the maximal connected components of G. We have that $G = \bigcup_{i=1}^{k} C_i$, and if $X \in C_i$ and $Y \in C_j$ with $i \neq j$, then it holds that $X \cap Y = \emptyset$.

Definition 79 states how we can define an alternative indistinguishability relation that is based on the original one in Definition 48.

Definition 79 (E_G)**.** *Let $G = \bigcup_{i=1}^{k} C_i$ be an agent pattern with k maximal connected components C_i. Two simplices $S, T \in \mathbb{C}$ cannot be distinguished componentwise by an agent pattern G, denoted by $S \,\mathsf{E}_G\, T$, if and only if:*

$$\exists 1 \leq j \leq k.\ S \sim_{C_j}^{\mathbb{C}} T.$$

We use the notation E_G since this relation is used to model statements of the form *every component of G knows that*. For this section, we adapt the truth definition as follows:

$$\mathcal{C}, w \Vdash_\sigma [G]\phi \qquad \text{iff} \qquad w \mathrel{\mathsf{E}_G} v \text{ implies } \mathcal{C}, v \Vdash_\sigma \phi \quad \text{for all } v \in \mathbb{C}.$$

Let $G = \{\{a\} \mid a \in \mathsf{Ag}\}$, we can read $[G]\phi$ as *everybody knows that* ϕ. By the symmetry of $\sim^{\mathbb{C}}$ and Lemma 40, we immediately obtain Lemma 71.

Lemma 71. *Let $G = \bigcup_{i=1}^{k} C_i$ be an agent pattern with k maximal connected components C_i. Then E_G is symmetric. Moreover, let $\mathcal{S}_G \subseteq \mathbb{C}$ be a maximal set of simplices such that for any $S \in \mathcal{S}_G$ we have $C_i^\star \subseteq S^\circ$ for some $1 \leq i \leq n$. Then the relation E_G is reflexive on $\mathcal{S}_G \times \mathcal{S}_G$.*

Note that E_G is not transitive. Indeed, consider the complex $\mathbb{C}$ in Figure 5.8 and the agent pattern $G = \{\{a\}, \{b\}\}$. Since X and Y share a b-vertex, we have that $X \mathrel{\mathsf{E}_G} Y$. By similar reasoning, we find that $Y \mathrel{\mathsf{E}_G} Z$. However, since X and Z do not intersect in any vertex, it does not hold that $X \mathrel{\mathsf{E}_G} Z$.

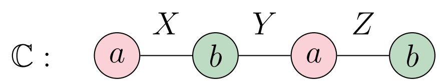

Figure 5.8: A complex that shows that componentwise indistinguishability need not be transitive.

Anti-monotonicity does also not hold in general. Consider again $\mathbb{C}$ in Figure 5.8, and let $G = \{\{a\}\}$ and $H = \{\{b\}\}$. Although $X \mathrel{\mathsf{E}_G} Y$ does not hold because X and Y do not intersect in an a-vertex, we find that $X \mathrel{\mathsf{E}_{G \cup H}} Y$ due to X and Y sharing a b-vertex, which violates anti-monotonicity. However, anti-monotonicity does hold componentwise.

Lemma 72 (Anti-monotonicity)**.** *Let $G = \bigcup_{i=1}^{k} C_i$ be an agent pattern with k maximal connected components C_i. Let C be a connected component with $C_i \subseteq C$ for some $1 \leq i \leq k$ and let $H = G \cup C$. We find that $\mathsf{E}_H \subseteq \mathsf{E}_G$.*

Proof. Let $\mathbb{C}$ be a complex, and consider the simplices $X, Y \in \mathbb{C}$. Combining G with C might create a new maximal connected component. Let $\mathcal{C}$ be the maximal connected component of H such that $C_i \subseteq C \subseteq \mathcal{C}$. Assume that $X \mathrel{\mathsf{E}_H} Y$. If $X \not\sim^{\mathbb{C}}_{\mathcal{C}} Y$, then there exists already a connected component C_j of G and H with $C_j \not\subseteq \mathcal{C}$ such that $X \mathrel{\mathsf{E}_{C_j}} Y$ and the claim follows. Therefore, assume $X \sim^{\mathbb{C}}_{\mathcal{C}} Y$. By assumption, we find that $X \sim^{\mathbb{C}}_{C_i} Y$ because $\sim^{\mathbb{C}}$ satisfies anti-monotonicity (see Lemma 42). Therefore, $X \mathrel{\mathsf{E}_G} Y$ by Definition 79. $\square$

Lemma 73 (Link). *Let $F, G, H \subseteq \mathsf{Pow}(\mathsf{Ag})$ be connected components such that $F \cup G$ is connected and $F \cup H$ is connected. The following formula is valid:*

$$[G]\phi \wedge [H]\psi \to [F \cup G \cup H](\phi \wedge \psi).$$

Proof. First, observe that $F \cup G \cup H$ is connected. Thus, by Lemma 72, $[G]\phi$ implies $[F \cup G \cup H]\phi$ and $[H]\psi$ implies $[F \cup G \cup H]\psi$. Since $[F \cup G \cup H]$ is a normal modality, we conclude $[F \cup G \cup H](\phi \wedge \psi)$. □

Example 14 (Missing Link). *Two networks G and H, each modeled as a connected component, both know that if malicious activity is detected, certain services must be stopped. Let* mact *be a propositional variable that indicates whether an intruder has been spotted and let* stop *indicates that the services are disabled. Since the procedure is known to both networks, we have*

$$[G](\mathsf{mact} \to \mathsf{stop}) \wedge [H](\mathsf{mact} \to \mathsf{stop}) \text{ as well as } [G \cup H](\mathsf{mact} \to \mathsf{stop}).$$

Suppose now that G detects malicious activity, i.e. $[G]\mathsf{mact}$. Thus, G will stop certain services, i.e. $[G]\mathsf{stop}$. If the networks cannot communicate with each other, i.e. $G \cup H$ is not connected, then H will not stop the services. Hence, G and H as a whole are not following the security protocol, i.e. $\neg[G \cup H]\mathsf{stop}$, and might leave the system in a vulnerable state. However, if a coordinating node relays messages from G to H, then H could shut down its services as well. By Lemma 73 we find that for some network F, such that $F \cup G$ as well as $F \cup H$ is connected, it holds that

$$([G \cup H](\mathsf{mact} \to \mathsf{stop}) \wedge [G]\mathsf{mact}) \to [F \cup G \cup H]\mathsf{stop}.$$

5.7 Conclusion and Outlook

This chapter presented multi-simplicial complexes, which are an explicit representation of semi-simplicial sets. We interpreted multi-simplicial complexes from the global point of view. The fact a that semi-simplicial set may contain parallel edges allowed us to define a novel indistinguishability relation, which in turn gave rise to a new modality that we call synergistic knowledge. Furthermore, we presented the logic of synergistic knowledge Syn, and showed that it is sound and complete with respect to models that are based on multi-simplicial complexes.

Unlike standard notions of group indistinguishability, our indistinguishability relation is based on an agent pattern and not on a group of agents. This added detail allows us to distinguish between the two models shown in Figure 5.2, thereby addressing an open question posed by Goubault et al. [26]. Furthermore, it allowed us to identify an alternative interpretation of an agent pattern. This led to an indistinguishability relation called componentwise indistinguishability, which is not necessarily transitive. An axiomatization of *componentwise knowledge* would exclude axiom 4 and present an interesting logic to explore in the context of multi-simplicial models.

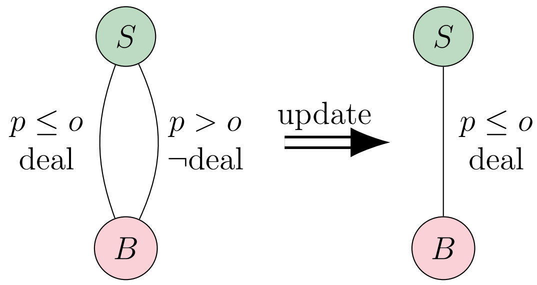

Figure 5.9: Without accessing a buyer-seller smart contract, the buyer B and the seller S do not know whether a trade can happen. We can represent the model after accessing the smart contract by removing the parallel edge that corresponds to the global state which is not the case.

Another direction for future research is exploring the notion of synergy in distributed computing, particularly for tasks where agents must commit to certain roles. For example, a smart contract that matches buyers and sellers can be seen as a synergistic primitive involving two parties, which is similar to the consensus number object in Example 10. The synergy arises from the fact that a seller cannot sell without a buyer, just as a buyer cannot buy without a seller. This task, therefore, inherently requires two different roles: buyer and seller. The multi-simplicial model depicted on the left in Figure 5.9 illustrates how our models can capture such tasks. In this scenario, the buyer B offers an amount o, while the seller S offers a product at price p. On their own, neither agent can determine whether a transaction is possible, i.e., whether $p \leq o$ (i.e., deal) or $p > o$ (i.e., no deal). Only after querying the previously mentioned smart contract, they know if an exchange can take place. Lastly, it is also interesting to analyze this from a dynamic standpoint in which we update the model by removing simplices. In this example, after querying the smart contract, we could eliminate one edge depending on the values of o and p. Figure 5.9 shows an update for the model on the left, in the case that $p \leq o$. Investigating updates that eventually transform models with parallel faces to standard simplicial complexes is a promising next step.

6 Conclusion

This thesis continued the structural study of simplicial semantics. It explored simplicial complexes that may contain adjacent vertices of the same color, directed faces, or parallel faces. While doing so, different notions of belief were studied, and a new notion of group knowledge was established. Moreover, modeling quorum systems with polychromatic complexes, as well as reasoning about the topology of a network in semi-simplicial sets, showed how generalizations of simplicial complexes can be used to formally reason about distributed systems.

Due to the infancy of the field, there are many things to be studied. As outlined at the end of each chapter, we can further refine the presented structures, or look at them from a different angle. For example, the study of polychromatic complexes and semi-simplicial sets followed the global approach. An analysis from the local point of view is much needed in order to understand these structures better. On the other hand, looking at directed complexes from the global perspective might simplify future extensions.

Besides shifting perspectives, we identify three lines of future work. First, based on our investigations of belief from the global and local point of view, a fruitful direction is to explore the open question posed by Castañeda et al. [17]: How can we reason about malicious agents on simplicial complexes? Second, a question from the Dagstuhl report (cf. Castañeda et al. [17]) that remains open is an intuitive philosophical interpretation of polychromatic complexes with respect to distributed systems. Lastly, there is no semantics based on simplicial sets that epistemically interprets degeneracy maps, i.e., that gives meaning to faces containing multiple copies of the same subface. This thesis serves as an entry point to all three lines of research.

Bibliography

[1] Thomas Ågotnes and Yì N. Wáng. Resolving distributed knowledge. *Artif. Intell.*, 252:1–21, 2017.

[2] Thomas Ågotnes and Yì N. Wáng. Somebody knows. In Meghyn Bienvenu, Gerhard Lakemeyer, and Esra Erdem, editors, *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021, Online event, November 3-12, 2021*, pages 2–11, 2021.

[3] Orestis Alpos, Christian Cachin, Giorgia Azzurra Marson, and Luca Zanolini. On the synchronization power of token smart contracts. In *41st IEEE International Conference on Distributed Computing Systems, ICDCS 2021, Washington DC, USA, July 7-10, 2021*, pages 640–651. IEEE, 2021.

[4] Orestis Alpos, Christian Cachin, Björn Tackmann, and Luca Zanolini. Asymmetric distributed trust. *Distributed Comput.*, 37(3):247–277, 2024.

[5] Hagit Attiya and Sergio Rajsbaum. Indistinguishability. *Commun. ACM*, 63(5):90–99, April 2020.

[6] Alexandru Baltag and Bryan Renne. Dynamic Epistemic Logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2016 edition, 2016.

[7] Alexandru Baltag and Sonja Smets. Dynamic belief revision over multi-agent plausibility models. *Proceedings of the 7th Conference on Logic and the Foundations of Game and Decision (LOFT)*, pages 11–24, 2006.

[8] Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT7)*, pages 13–60, 2008.

[9] Mathieu Baudet, George Danezis, and Alberto Sonnino. Fastpay: High-performance byzantine fault tolerant settlement. In Sarah Meiklejohn and Abhi Shelat, editors, *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies, AFT 2020, New York, NY, USA, October 21-23, 2020*, pages 163–177. ACM, 2020.

[10] Marta Bílková, Zoé Christoff, and Olivier Roy. Revisiting epistemic logic with names. In Joseph Y. Halpern and Andrés Perea, editors, *Proceedings Eighteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2021, Beijing, China, June 25-27, 2021*, volume 335 of *EPTCS*, pages 39–54, 2021.

[11] Marta Bílková, Hans van Ditmarsch, Roman Kuznets, and Rojo Randrianomentsoa. Bisimulation for impure simplicial complexes. In Agata Ciabattoni, David Gabelaia, and Igor Sedlár, editors, *Advances in Modal Logic, AiML 2024, Prague, Czech Republic, August 19-23, 2024*, pages 225–248. College Publications, 2024.

[12] Elizabeth Borowsky and Eli Gafni. Generalized FLP impossibility result for t-resilient asynchronous computations. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 91–100. ACM, 1993.

[13] Christian Cachin, Rachid Guerraoui, and Luís E. T. Rodrigues. *Introduction to Reliable and Secure Distributed Programming (2. ed.)*. Springer, 2011.

[14] Christian Cachin, David Lehnherr, and Thomas Studer. Synergistic knowledge. In Shlomi Dolev and Baruch Schieber, editors, *Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 552–567. Springer, 2023.

[15] Christian Cachin, David Lehnherr, and Thomas Studer. Simplicial belief. In Ulrich Schmid and Roman Kuznets, editors, *Structural Information and Communication Complexity (SIROCCO)*, pages 176–193. Springer Nature Switzerland, 2025.

[16] Christian Cachin, David Lehnherr, and Thomas Studer. Synergistic knowledge. *Theor. Comput. Sci.*, 1023:114902, 2025.

[17] Armando Castañeda, Hans van Ditmarsch, Roman Kuznets, Yoram Moses, and Ulrich Schmid. Epistemic and Topological Reasoning in Distributed Systems (Dagstuhl Seminar 23272). *Dagstuhl Reports*, 13(7):34–65, 2024.

[18] Soma Chaudhuri. Agreement is harder than consensus: set consensus problems in totally asynchronous systems. In *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing*, PODC '90, page 311–324, New York, NY, USA, 1990. Association for Computing Machinery.

[19] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *J. Cryptol.*, 1(1):65–75, 1988.

[20] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.

[21] Ronald Fagin, Joseph Y. Halpern, and Moshe Y. Vardi. What can machines know? on the epistemic properties of machines. In *Proceedings of the Fifth AAAI National Conference on Artificial Intelligence*, AAAI'86, page 428–434. AAAI Press, 1986.

[22] Greg Friedman. Survey Article: An elementary illustrated introduction to simplicial sets. *Rocky Mountain Journal of Mathematics*, 42(2):353–423, 2012.

[23] Benoit Gaudou, Andreas Herzig, Dominique Longin, and Emiliano Lorini. On modal logics of group belief. In Andreas Herzig and Emiliano Lorini, editors, *The Cognitive Foundations of Group Attitudes and Social Interaction*, pages 75–106. Springer, 2015.

[24] Éric Goubault, Bernardo Hummes Flores, Roman Kniazev, Jérémy Ledent, and Sergio Rajsbaum. A categorical and logical framework for iterated protocols. *CoRR*, abs/2505.10071, 2025.

[25] Éric Goubault, Roman Kniazev, and Jérémy Ledent. A many-sorted epistemic logic for chromatic hypergraphs. In Aniello Murano and Alexandra Silva, editors, *32nd EACSL Annual Conference on Computer Science Logic, CSL 2024, February 19-23, 2024, Naples, Italy*, volume 288 of *LIPIcs*, pages 30:1–30:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

[26] Éric Goubault, Roman Kniazev, Jérémy Ledent, and Sergio Rajsbaum. Semi-simplicial set models for distributed knowledge, 2023.

[27] Éric Goubault, Roman Kniazev, Jérémy Ledent, and Sergio Rajsbaum. Simplicial models for the epistemic logic of faulty agents. *Boletín de la Sociedad Matemática Mexicana*, 30(3):90, 2024.

[28] Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. A simplicial complex model for dynamic epistemic logic to study distributed task computability. In Andrea Orlandini and Martin Zimmermann, editors, *Proceedings Ninth International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2018, Saarbrücken, Germany, 26-28th September 2018*, volume 277 of *EPTCS*, pages 73–87, 2018.

[29] Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. A simplicial complex model for dynamic epistemic logic to study distributed task computability. *Inf. Comput.*, 278:104597, 2021.

[30] Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. A simplicial model for $KB4_n$: Epistemic logic with agents that may die. In P. Berenbrink and B. Monmege, editors, *39th International Symposium on Theoretical Aspects of Computer Science*, volume 219 of *LIPIcs*, pages 33:1–33:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[31] Adam J. Grove and Joseph Y. Halpern. Naming and identity in epistemic logics part I: the propositional case. *J. Log. Comput.*, 3(4):345–378, 1993.

[32] Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovič, and Dragos-Adrian Seredinschi. The consensus number of a cryptocurrency. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, PODC '19, page 307–316, New York, NY, USA, 2019. Association for Computing Machinery.

[33] Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. In T. Kameda, J. Misra, J. G. Peters, and N. Santoro, editors, *Proc. of the Third Annual ACM Symposium on Principles of Distributed Computing*, pages 50–61. ACM, 1984.

[34] Maurice Herlihy. Wait-free synchronization. *ACM Trans. Program. Lang. Syst.*, 13(1):124–149, 1991.

[35] Maurice Herlihy, Dmitry N. Kozlov, and Sergio Rajsbaum. *Distributed Computing Through Combinatorial Topology.* Morgan Kaufmann, 2013.

[36] Maurice Herlihy and Sergio Rajsbaum. Algebraic topology and distributed computing: A primer. In Jan van Leeuwen, editor, *Computer Science Today: Recent Trends and Developments*, volume 1000 of *Lecture Notes in Computer Science*, pages 203–217. Springer, 1995.

[37] Maurice Herlihy and Nir Shavit. The asynchronous computability theorem for t-resilient tasks. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 111–120. ACM, 1993.

[38] Sou Hoshino. Determining existence of logical obstructions to the distributed task solvability, 2022.

[39] Michele Loreti and Michela Quadrini. A spatial logic for simplicial models. *Log. Methods Comput. Sci.*, 19(3), 2023.

[40] Dahlia Malkhi and Michael K. Reiter. Byzantine quorum systems. *Distributed Comput.*, 11(4):203–213, 1998.

[41] Daisuke Nakai, Masaki Muramatsu, and Susumu Nishimura. Partial product updates for agents of detectable failure and logical obstruction to task solvability. *CoRR*, abs/2303.16437, 2023.

[42] Susumu Nishimura. Defining logical obstruction with fixpoints in epistemic logic. *J. Appl. Comput. Topol.*, 8(4):941–970, 2024.

[43] Eric Pacuit. *Neighborhood Semantics for Modal Logic.* Springer Publishing Company, Incorporated, 1st edition, 2017.

[44] Marc Pauly. A modal logic for coalitional power in games. *J. Log. Comput.*, 12(1):149–166, 2002.

[45] Rojo Randrianomentsoa. *Epistemic Logic for Distributed Systems with Crash Failures.* PhD thesis, TU Wien, 2025.

[46] Rojo Randrianomentsoa, Hans van Ditmarsch, and Roman Kuznets. Impure Simplicial Complexes: Complete Axiomatization. *Logical Methods in Computer Science*, 19(4), 2023.

[47] Michael E. Saks and Fotios Zaharoglou. Wait-free k-set agreement is impossible: the topology of public knowledge. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 101–110. ACM, 1993.

[48] Robert Stalnaker. On logics of knowledge and belief. *Philosophical Studies*, 128:169–199, 2006.

[49] W. van der Hoek and J-J.Ch. Meyer. *A Complete Epistemic Logic for Multiple Agents: Combining Distributed and Common Knowledge*, volume 1996-52. Utrecht University: Information and Computing Sciences, uu-cs edition, 1996.

[50] Hans van Ditmarsch, Djanira Gomes, David Lehnherr, Valentin Müller, and Thomas Studer. Hypergraph semantics for doxastic logics. *Manuscript*, 2025.

[51] Hans van Ditmarsch, Éric Goubault, Marijana Lazic, Jérémy Ledent, and Sergio Rajsbaum. A dynamic epistemic logic analysis of equality negation and other epistemic covering tasks. *J. Log. Algebraic Methods Program.*, 121:100662, 2021.

[52] Hans van Ditmarsch, Éric Goubault, Jérémy Ledent, and Sergio Rajsbaum. Knowledge and simplicial complexes. In Björn Lundgren and Nancy Abigail Nuñez Hernández, editors, *Philosophy of Computing*, pages 1–50, Cham, 2022. Springer International Publishing.

[53] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer Publishing Company, Incorporated, 1st edition, 2007.

[54] Koki Yagi and Susumu Nishimura. Logical obstruction to set agreement tasks for superset-closed adversaries. *CoRR*, abs/2011.13630, 2020.

[55] Yuanzhe Yang. Impure simplicial complex and term-modal logic with assignment operators. In Adam Bjorndahl, editor, *Proceedings Twentieth Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2025), Düsseldorf, Germany, July 14-16, 2025*, pages 393–410. Electronic Proceedings in Theoretical Computer Science, To appear.

Erklärung

gemäss Art. 18 PromR Phil.-nat. 2019

Name/Vorname: Lehnherr, David

Matrikelnummer: 14-930-275

Studiengang:

Bachelor ☐ Master ☐ Dissertation ☑

Titel der Arbeit: Simplicial Structures for Epistemic Reasoning in Multi-agent Systems

LeiterIn der Arbeit: Prof. Dr. Christian Cachin
Prof. Dr. Thomas Studer

Ich erkläre hiermit, dass ich diese Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen benutzt habe. Alle Stellen, die wörtlich oder sinngemäss aus Quellen entnommen wurden, habe ich als solche gekennzeichnet. Mir ist bekannt, dass andernfalls der Senat gemäss Artikel 36 Absatz 1 Buchstabe r des Gesetzes über die Universität vom 5. September 1996 und Artikel 69 des Universitätsstatuts vom 7. Juni 2011 zum Entzug des Doktortitels berechtigt ist. Für die Zwecke der Begutachtung und der Überprüfung der Einhaltung der Selbständigkeitserklärung bzw. der Reglemente betreffend Plagiate erteile ich der Universität Bern das Recht, die dazu erforderlichen Personendaten zu bearbeiten und Nutzungshandlungen vorzunehmen, insbesondere die Doktorarbeit zu vervielfältigen und dauerhaft in einer Datenbank zu speichern sowie diese zur Überprüfung von Arbeiten Dritter zu verwenden oder hierzu zur Verfügung zu stellen.

Ort/Datum Bern, 10.11.2025

Unterschrift D. Le