NASA/TM–20260004926

A New Semantics for Belief Revision in Simplicial Complexes

Philip Sink
Department of Philosophy
Carnegie Mellon University
Pittsburgh, Pennsylvania

Alwyn E. Goodloe
NASA Langley Research Center
Hampton, Virginia

June 2026

NASA STI Program Report Series

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI Program provides access to the NTRS Registered and its public interface, the NASA Technical Report Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.
- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI Program, see the following:

- Access the NASA STI program home page at `http://www.sti.nasa.gov`

NASA/TM–20260004926

A New Semantics for Belief Revision in Simplicial Complexes

Philip Sink
Department of Philosophy
Carnegie Mellon University
Pittsburgh, Pennsylvania

Alwyn E. Goodloe
NASA Langley Research Center
Hampton, Virginia

National Aeronautics and
Space Administration

NASA Langley Research Center
Hampton, Virginia 23681-2199

June 2026

Acknowledgments

We thank the branch and directorate reviewers for their feedback on this work.

Available from:

NASA STI Program / Mail Stop 050
NASA Langley Research Center
Hampton, VA 23681-2199

Abstract

In this paper we set out to develop a novel method for handling transient faults in gossip protocols. In particular, we are interested in settings where a signal "P" can come across the network, followed by a signal "$\neg P$" at a later time. These signals could even come from the same agent, and any agent is capable of failing as such. The main idea to handle such a dramatically faulty setting is to use a technique from the philosophical literature called "belief revision". [31] This allows agents to update on information contrary to their beliefs without resorting to probabilities. However, belief revision is often difficult to implement in a practical setting. Therefore, we turn towards using simplicial semantics for epistemic modal logic to represent each single time-slice of a protocol. Such models already have a history of applications in distributed computing. [25] [36] [19] [26] [24] [22] [23] [29] In order to model the transitions between time slices, we will use action models. These too already have a history in distributed computing. [26] [19] [3] [34] [39] [15] We will modify the application of action models to simplicial complexes in order to incorporate belief revision. In so doing, we will be able to define a novel gossip protocol for handling the prescribed fault scenario.

Contents

1 Introduction

Imagine you are a disaster relief worker and you receive a message over the radio from a colleague Anne—this message tells you that a particular road is clear, one you will need to traverse to get to your next objective. Fifteen minutes later, Anne calls you again, but this time she says that the road is in fact not clear. She had gotten a new vantage point on the situation, and a fallen tree which was blocked from view previously was now clearly visible to her. Our goal in this paper is to use examples like this to motivate a new model for protocols used in distributed computing. There are a few key features in the above example that are of interest to us in this paper. The first is to note that you at one time learn a fact, call it P, and then at a later time learn precisely the negation of this fact, call it $\neg P$. Moreover, you receive this contradictory information from the same source in each case. This is not unusual in distributed computing systems comprised of many agents interacting with real world. In our example, an agent, Anne, learned new information over time that negated a fact she had previously reported in messages to colleagues. On the other hand, the same behavior can be manifested in agents that are suffering from transient faults. We will assume any agent in a distributed system is subject to such behavior. This precludes the possibility of trying to identify faulty agents by assuming most of the network is not faulty. At any given time, any node in the network might contradict what you already believe, or even what they have already told you.

Under this highly restricted setting, we aim to investigate whether a notion from philosophy, belief revision, will be of any use to us. This is a theory of learning that is intended to handle information which is contrary to an agent's beliefs: that is, it is intended to allow agents who believe some fact P to learn $\neg P$ without something catastrophic occurring, such as the agents believing all facts to be true, including inherently contradictory ones. The details of how this system works will be left to the next section. However, what is needed is a sufficiently rich mathematical setting in which we can describe protocols, and compare and contrast some traditional protocols, in particular gossip protocols, from protocols which are otherwise identical but employ belief revision.

To meet these ends, this paper will live in the intersection of four distinct areas of research: Simplicial Semantics for Epistemic Logic, Action Models, Belief Revision, and Distributed Computing, in particular gossip protocols. In general, epistemic modal logic, algebraic topology, and distributed computing are areas of research that do not often overlap. We do not expect our readers to have a background in such a broad range of topics. So the paper will largely set out to provide the necessary background. Section 2 will introduce the gossip protocol we wish to model using our formalism. Sections 3 and 4 describe background information, and can be skipped and returned to later if the reader is unfamiliar with simplicial semantics. Section 5 describes a novel simplicial semantics, and introduces a notion of belief revision which overcomes theoretical hurdles which have made belief revision difficult to implement in practical settings historically. Section 6 will then define action models, in particular define them with respect to our new simplicial models, and then explain how action models can be used to represent protocols. Section 7 will then familiarize ourselves with this setting through a few examples of action models representing some traditional protocols. Finally, in section 8 we will return to the disaster relief example, and give a gossip protocol which is descriptive of this example, and explain how through simulations of this protocol we can test the efficacy of belief revision in this setting.

2 Gossip Protocol: Overview

Picture this: You're at a conference with one thousand attendees, and you need to announce an urgent room change. You could use a PA system (centralized broadcast), but what if it fails? Instead, you whisper to three people nearby, who each tell three others, who tell three more. Within minutes, everyone knows. This is gossip in action—and it's exactly how some of the world's most resilient distributed systems stay synchronized [18].

Gossip protocols already have a storied history of applications. [35] [14] Our goal is, in part, to extend this area of applications to disaster relief networks. The notion of gossip protocol we are drawing on is Ken Birman's white paper "The Promise, and Limitations, of Gossip Protocols". [10] There he identifies six features of gossip protocols:

1. "The core of the protocol involves periodic, pairwise, inter-process interactions.
2. The information exchanged during these interactions is of (small) bounded size.
3. When nodes interact, the state of one or both changes in a way that reflects the state of the other. For example, if A pings B just to measure the round-trip-time between them, it isn't a gossip interaction.
4. Reliable communication is not assumed.
5. The frequency of the interactions is low compared to typical message latencies, so that the protocol costs are negligible
6. There is some form of randomness in the peer selection. Peer selection might occur within the full node set, or might be performed in a smaller set of neighbors." [10]

The above are informal criteria, and thus one cannot prove, specifically, whether a specified protocol satisfies them one way or the other. Moreover, when describing protocols using action models, some of the above do not apply. Though this should not stop us. Take, for example, the protocol given in "Knowledge and Gossip". [3] It's fairly easy to convince oneself that 1-4 and 6 are satisfied by the sequence of actions given therein. And this is good, because our protocol will be very, very similar.

Our protocol will function as follows. The first simplifying assumption we shall make is that the information being sent over the network can be modeled as a vector privately held bit values corresponding to each agent. Recall that, in the example described at the beginning of Section 1, Anne is sending over the network information about whether or not the road is clear. For each agent, say a, we will interpret the proposition "The road is clear" as a logical atom A, which is the private bit value held by Anne. This means that agents other than Anne are not able, at least at first, to comment on the value of A. Modifying our protocol to do away with this last fact should not be difficult, but we are starting with the simplest possible case. Without loss of generality, we assume that A being "true" is the true bit value corresponding to the real state of the world.

The second feature of the protocol is that we assume, at each time interval, an agent is randomly selected to send their vector of bit values over the network, and another agent is randomly selected to be the unique receiver. Suppose a is selected as the sender and b as receiver. Then, we assume there is a random chance that either A or $\neg A$ is the value sent over the network. In line with Anne sending that the road is not clear, when in fact it is, we assume that if $\neg A$ is sent, that a *intended* to send $\neg A$. From a's point of view, if $\neg A$ is selected, and a is observing A, no message is sent (a never intends to lie). Moreover, we assume that if $\neg A$ is sent *and received*, then b receives it correctly. However, there is also a chance that b received nothing at all and is unaware that a attempted to contact them. So, in summary, a and b are randomly selected from the pool of agents. There is a chance that a intentionally sends A or that a intentionally sends $\neg A$. Furthermore, there is a chance that b receives the intended message, or that b receives nothing and has no idea a message was intended to be sent. All other agents are not aware of what was attempted.

In order to clarify the behavior of the protocol, we shall express the protocol in pseudocode (Figure 1). Lower case roman letters denote agents and upper case roman letters denote the message being sent. Node x can send either the message X or ¬X in addition to what they've learned so far. We will assume there are n agents in total. Each agent maintains an associative store V indexed by agent name containing the messages received from other agents. Think of this as a vector of length n indexed by the agent name. So $\mathrm{V}(\mathrm{x})(X)$ holds the value of X associated with agent x. The infix update operator $\mathrm{U} \otimes \mathrm{W}$ takes two vectors and returns a new vector. Messages are assumed to have the following format:

$$\langle \mathrm{src}, \mathrm{dst}, \mathrm{message} \rangle,$$

where src and dst are the source and destination agents and message is the content being sent. The protocol makes use of an oracle $\mathrm{choose}(A)$ that randomly selects an element from the set A.

Our objective is to draw together action models, simplicial semantics, and belief revision in order to define and reason about an idealized gossip protocol for handling the kinds of faulty agents motivated in the introduction. The notion of "Belief Revision" we incorporate in our logical formalism allows us to represent the way agents update information over time. For readers unfamiliar, we will describe this idea in more detail in Section 5.3. However, it suffices here to say that revision allows an agent to learn φ, and at a later time $\neg\varphi$, without either trivializing their beliefs, or just defaulting to what was learned most recently. Naturally, there are *many* idealizing assumptions baked into this core example. Future work, will seek to do away these so as to provide a framework for reasoning about actual industrial strength gossip protocols.

3 Related Works

The reader should note that, if they are unfamiliar with simplicial semantics or belief revision, this section should be skipped and returnd to later. Various work already exists in some of these intersections. Seeing as the impetus for Simplicial Semantics is the use of combinatorial topology in distributed computing, most of that literature acknowledges and is indebted to this connection already. [25] [36] [19] [26] [24] [22] [23] [29] In particular, the papers "Knowledge and Simplicial Complexes" and "Simplicial Belief" are already attempts to give a semantics for belief, or non-factive (defeasible) knowledge. We will start with "Knowledge and Simplicial Complexes". There, they start with a typical simplicial

Gossip Protocol for Agent x

```
              𝒜 ≝ {a, b, c ...}
              ℳ ≝ {A, B, C, ...}
              var V, rcvd : 𝒜 ⟶ (ℳ ⟶ 2)
              var x, y ∈ 𝒜
              var c ∈ ℳ

              skip
     □
              c := choose{0, 1}.
              ; if c = V(x)(X) then
                y := choose 𝒜 − x
                ; send⟨x, y, V⟩
     □
              receive⟨y, x, rcvd⟩
              ; V := V ⊗ rcvd
     where
              var U, W : 𝒜 ⟶ ℳ
              U ⊗ W ≝
               forall i ∈ 𝒜
                (U ⊗ W)(i) =
                  if i ≠ x then
                    if U(i) = W(i) = ∅ then ∅
                    if U(i) ≠ ∅ ∧ W(i) = ∅ then U(i)
                    if U(i) = ∅ ∧ W(i) ≠ ∅ then W(i)
                    if U(i) ≠ ∅ ∧ W(i) ≠ ∅ then W(i)
                  if i = x then U(i)
```

Figure 1.

model of knowledge, in our notation (N,V,S,L) (see Section 5). In particular, their models satisfy the UCF condition. Then each agent $a \in Ag$ is given an idempotent function f_a on the a colored vertices. That is, if $N_a := \{n \in N | V(n) = a\}$, $f_a : N_a \to N_a$ such that $V(n) = V(f_a(n))$. The belief condition is given, then, by the following for $X \in \mathscr{F}(S)$:

$$\mathscr{M}, X \models B_a \varphi \text{ iff } \forall Y \in \mathscr{F}(S) \text{ if } f_a(\pi_a(X)) = \pi_a(Y) \text{ then } \mathscr{M}, Y \models \varphi$$

Because each f_a is idempotent, it is easy to check that B_a satisfies the **KD45** axioms. However, this definition of belief has some unsatisfying properties. First and foremost, this notion of belief does not satisfy the intuitive axiom "knowledge implies belief". Specifically, if one gave the usual intersection semantics for knowledge,

$$\mathscr{M}, X \models K_a \phi \text{ iff } \forall Y \in \mathscr{F}(S) \text{ if } \pi_a(Y) = \pi_a(X) \text{ then } \mathscr{M}, Y \models \phi,$$

Then $K_a \varphi \to B_a \varphi$ is not sound, which is undesirable. However, a natural solution to this might be to strengthen the semantics such that this notion is sound. Suppose $K_a \varphi \to B_a \varphi$ were sound. Then the following must be true for any $X \in \mathscr{F}(S)$.:

$$\{Y \in \mathscr{F}(S) | f_a(\pi_a(X)) = \pi_a(Y)\} \subseteq \{Y \in \mathscr{F}(S) | \pi_a(X) = \pi_a(Y)\}$$

Under the UCF assumptions, this equality holds if and only if f_a is the identity function. It follows that belief and knowledge must coincide. This is strongly undesirable. If we take the sense of belief from "Knowledge and Simplicial complexes", we cannot have both UCF, and the validity of "knowledge implies belief", without trivializing belief to coincide with knowledge.

The second approach to belief is given in "Simplicial Belief". [13] The author's approach here depends on breaking the UCF assumption. In particular, if $X \in \mathscr{F}(S)$, there may be $x, y \in X$ such that $x \neq y$ and $V(x) = V(y)$. For $X \in \mathscr{F}(S)$, they define $m_a(X) := |\{n \in X | V(n) = a\}|$. From this, they define the following relation:. For any $X, Y \in \mathscr{F}(S)$:

$$X \leq_a Y \text{ iff } m_a(X) \leq m_a(Y)$$

They also define the following usual intersection relation on facets. That is, in our notation, $X \sim_a Y$ if and only if $\exists n \in X \cap Y$ such that $V(n) = a$. Then $\unlhd_a = \leq_a \cap \sim_a$, and their first condition for belief is equivalent to the following:

$$\mathscr{M}, X \models B_a \varphi \text{ iff } \forall Y \in \mathscr{F}(S) \text{ if } X \rhd_a Y \text{ then } \mathscr{M}, Y \models \varphi$$

This says that φ is true at all facets 1: whose number of a-colored perspectives is capped by the number of a-colored perspectives of X, and 2: whose intersection with X contains an a colored node, then $B_a \varphi$ is true at X. They also take a second, stronger notion of belief:

$$\begin{aligned} \mathscr{M}, X \models B_a \varphi \text{ iff } \forall Y \in \mathscr{F}(S) \text{ if } X \unrhd_a Y & \\ & \text{and } \forall Z \in \mathscr{F}(S)(\text{ if } X \unrhd_a Z \text{ then } m_a(Z) \geq m_a(Y)) \\ & \text{then } \mathscr{M}, Y \models \varphi \end{aligned}$$

This second definition says that if φ is true at all facets sharing an a-intersection with X with a minimal number of perspectives over this collection, then $B_a\varphi$ is true at X. Both of these are motivated by the idea that if a world contains more a-perspectives, then that world is less *plausible*. [13] The reason is that agent is seemingly less certain of what they believe in these worlds. If X contains more a-perspectives than Y, then a is less certain of their internal configuration in X than in Y.

Each of these definitions does, clearly, have the property that $K_a\varphi \to B_a\varphi$, which is desirable. However, as mentioned earlier, their model throws out the UCF restriction. What if one were to assume that S satisfied UCF? The immediate consequence would be that for any $X \in \mathscr{F}(S)$, $m_a(X) = 1$. Hence, $\leq_a$ would be the total relation on $\mathscr{F}(S)$, and so $\trianglelefteq_a = \sim_a$. Since $\sim_a$ is the usual relation for definition knowledge, this would, as before, trivialize belief such that it coincided with knowledge. So, there is no way in this paper, either, to have both "knowledge implies belief" and UCF without trivializing knowledge. The primary advantage of our paper is that we are able to have all three considerations. We will have that knowledge implies belief, our models satisfy UCF, and belief is not trivial.

The research reported in [2] and [37] connect epistemic logic more broadly to distributed computing by using epistemic information to track which messages should be followed up on in a gossip protocol, while [3] brings in action models to describe gossip protocols themselves. The protocol developed in [3] will be of great use to us. In this protocol, each agent has a secret, which one can think of as a privately held bit value. At each time slice, a call between two agents is given. This call can take various forms: For example, it may be common knowledge who the call is between, but the information shared over the call is only known to the callers. By contrast, it may only be common knowledge that some call was made, or agents not in the call may consider it possible that no call was made. We will be interested in particular in this last case when we create our gossip protocol in Section 8. The initial configuration of a gossip protocol in [3] (which they call the "Initial Gossip Model") Is a model whose unique world is itself a model (called the "Initial Gossip Model") whose worlds are all possible configurations of the truth values of the secrets, and where (in our notation) each R_a is the equivalence relation preserving the truth values of a's secret. Translated into the simplicial setting, this will be precisely the initial configuration we will use in our gossip protocol. Similarly, the terminal configuration, or "Terminal Gossip Situation", is the model where only reflexive edges exist. This, too, defines the goal of our gossip protocol. For each of the three calls described above, this paper describes how to "update" a Gossip model using an action model corresponding to each call. These actions use a language and notation similar to what we will describe in Section 6.1.

In general, the procedures for using action models to describe protocols are understood but new, showing up in [26], [19], [3], [34], [39], and [15]. In Section 7, we will translate many of the examples of protocols in [29] into action models, broadly following the trends of these papers. Modifying action models where agents learn via belief revision is also an explored area, with work such as [7] and [8]. However, none of these models will use the specific model of revision we will motivate in Section 5.

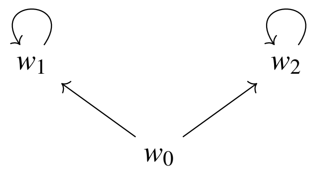

Figure 2.

4 Kripke Semantics

Many models of uncertainty in distributed computing have used the formal tool of "possible worlds" or "states" to facilitate their model. [28] [33] [38] [30] [20] [37] [2] The basic picture is thus. We start with a set of worlds/states W and for each agent $a \in Ag$ a binary relation $\mathscr{R}_a \subseteq W \times W$. Furthermore, at each world, different things are true. Let $\ell : \mathfrak{P} \to 2^W$ be an "assignment" function which assigns to every logical atom of our language, the set of atoms being $\mathfrak{P}$, to a set of worlds. Intuitively, these worlds are those where the logical atom is true. Given a world $w \in W$, we say that at an agent "knows" or "believes" a formula φ when, for all worlds u such that $(w,u) \in \mathscr{R}_a$, φ is true at u. This conception of knowledge underlies the applications of these models to distributed computing: one simple application is using this conception to model consensus as "Everybody knows that everybody knows that everybody knows that..." We say that a model $\langle W, \{\mathscr{R}_a\}_{a \in Ag}, \ell \rangle$ is a "Kripke" or "frame" model.

Let's consider a few examples. We will focus on the idea of modeling belief. Suppose there are three worlds, w_i for $i \in \{0,1,2\}$. We'll say that $w_0 \mathscr{R}_a w_1$ and $w_0 \mathscr{R}_a w_2$. Moreover, $w_1 \mathscr{R}_a w_1$ and $w_2 \mathscr{R}_a w_2$. How we draw this can be seen in Figure 2.

Suppose $\ell(P) = \{w_0, w_1, w_2\}$. Then, at world w_0, we can say that a believes P, as P is true at w_1 and w_2. Suppose by contrast $\ell(R) = \{w_1, w_2\}$. Then, as before, at world w_0, we can say that a believes R, as R is true at w_1 and w_2. However, in this scenario, by contrast, we say that a has a false belief. While a believes R at w_0, R is also false at w_0. One can observe that is is because R is irreflexive at w_0. Indeed, many epistemic properties have a tight correspondence with frame properties. See [20] for more on this.

The most important concept, however, is multi-agent belief. We will, for the duration of this paper, let the color red correspond with agent a and blue with b. So, for Kripke frames, we will draw $\mathscr{R}_a$ in red and R_b in blue.

Consider Figure 3. At w_2, we have that b believes Q, because Q is true at w_3 and w_4. Similarly, at w_1, b believes Q, since Q is true at w_1. So, we say that B_bQ is true at w_1 and w_2. So, a believes B_bQ at w_0. We would write this as B_aB_bQ, and say "a believes that b believes that Q. Note, not only is Q false at w_0, but B_bQ is false as well. Similarly, Q is false at w_2. So, we say that "a believes that $\neg Q$ is possible". Since Q is false at w_2, and B_bQ is true, then a believes it's possible that b has false beliefs.

Generally, in the previous literature, it is assumed that the $\mathscr{R}_a$ are all equivalence relations. [20] [37] [2] This is related to the fact that the modality, or quantification over worlds via the relations, is called "knowledge". Because the $\mathscr{R}_a$ are equivalence relations, indeed, because they are reflexive, then at any world w, agent a considers w possible. Thus, if a knows φ at w, then φ must be true at w. This is called "factivity". However, for our pur-

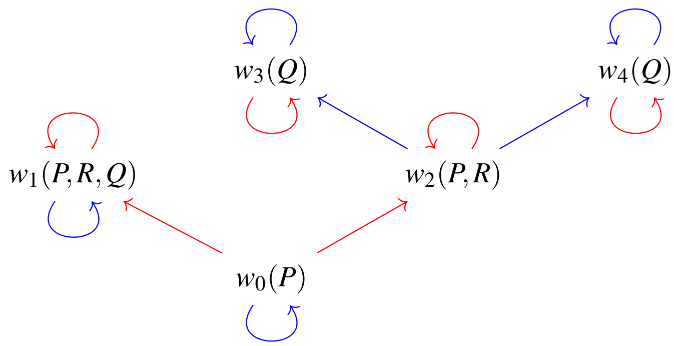

Figure 3.

poses, this assumption is too strong. Indeed, to model disaster relief, it is very important that a has some sort of affirmative attitude towards φ, despite φ being false. For our purposes, we will refer to this attitude as "belief". For those familiar with the literature connecting modal logic interpreted as knowledge to distributed computing, it suffices to interpret our use of "belief" as "knowledge without the assumption of factivity". Other domains, such as Bayesian statistics, which interpret belief as having "degrees" or "strength", are not relevant for the discussion of this paper.

5 Simplicial Semantics

5.1 Semantics for Belief

This paper will work in a similar but distinct setting to Kripke models. The basic epistemic unit that we will start with we will call a "perspective" or "point of view," very similar to the idea outlined in [22]. Given a particular agent, let's call her Alice, the perspective of that agent consists abstractly of (some of) the information available to that agent. For example, if the agent is sitting on a park bench, and sees a blue sky and ducks in the lake, all of this information is what we might consider part of the agent's perspective. This is indeed an abstract notion, but crucial to us is that the information constituting a perspective can be broken down into two pieces. The first is the *hard* information, or the information that is in some sense indefeasible. If the agent on the park bench sees ducks, for instance, this is reasonably considered hard information. By contrast, suppose Alice has a friend, Barbara, who regularly goes to the park with her. Each time the pair has been to the park, they see ducks; however today, her friend is home sick. Alice still goes to the park and sees ducks, but Barbara of course does not. Both, however, believe that there are ducks in the lake. However, for Alice, this is hard information. By contrast, Barbara believes there are ducks in the lake in a weaker sense. This is an example of *weak* information.

The line between hard and weak information is only intuitive. In fact, the distinction only matters as a modeling convention. In certain modeling contexts, it will be useful to represent some information as indefeasible, and other information as defeasible. That is all the distinction between *hard* and *weak* is meant to cache out. It rests solely on the intuition that an agent's point of view is made up of facts that are immediately obvious to

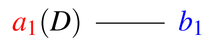

Figure 4.

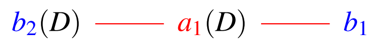

Figure 5.

the agent and cannot be contradicted by further information, and other facts that are less so. As a technical and simplifying assumption, we will assume that the hard information of a perspective is given by a possibly empty set of literals associated with that perspective.

Since we are taking perspectives as our foundational epistemic object, how can we use them to model uncertainty? One again imagines our agent at the pond with the ducks. Call her agent a. Her friend will be agent b. If D is the logical atom that means "The ducks are in the pond at the park," then we can say that D is assigned to a's perspective, call it a_1. b's perspective does not, in this language, consist of any hard information, and so we associate b_1 with the empty set. Since these perspectives are true simultaneously, we can group them together as the set $\{a_1, b_1\}$, which we draw as Figure 4.

When we formally define simplicial semantics later, it will be clear that this object is a (very basic) simplicial complex. In order to read the above, for now, it suffices to recognize that this is equivalent to a Kripke model, with one world where D is true, and two reflexive edges, one for each agent. The set with the two perspectives is the possible world, and each agent recognizes that world in their perspective, hence the reflexive edges.

Now let's imagine a slightly more complicated example. Both friends are going to the park, but do not see each other yet. We will consider this from a's point of view. She sees the ducks in the pond, but is uncertain whether or not b sees the ducks yet. She cannot imagine that b sees there are not ducks in the pond, as that would be inconsistent with her own experience. But, she cannot tell whether or not b's perspective is one where they are, in fact, seeing ducks, or they have not seen the ducks yet. The latter is the same as b_1 as above, but the former, which we call b_2, will have D associated with it, as seen in Figure 5.

This is equivalent to a Kripke model with two worlds, one for each set a, b_1 and a, b_2. The worlds are indistinguishable to agent a, and all reflexive edges for both agents are present. In the case of two agents, we will refer to sets containing two perspectives as *facets*. The facets in the above examples are red because they are associated with agent a.

We have just enough informal background to motivate the main idea behind our model of belief revision. Consider Figure 6. This model has two facets, namely $X := \{a_0, b_0\}$, and $Y := \{a_1, b_1\}$. At any facet Z, we have that $\mathscr{M}, Z \vDash B_b Q$.

However, by contrast, consider Figure 7.

Note that this model differs only in the inclusion of $\neg R$ in the vertices b_1 and a_1. When we formally specify truth in a model, we will be able to give a strong sense in which these two models are equivalent. Ultimately, at each facet, the same formulas in a basic modal language will be true. However, there is another sense in which they are not equivalent. In Figure 6, the set $\{a_0, b_1\}$ could have been included, as the vertices are consistent. The same is not true in Figure 7. If, say, $\neg Q$ were announced, agent a in perspective a_0 could

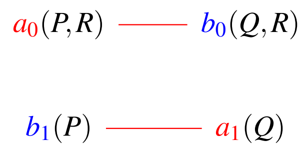

Figure 6.

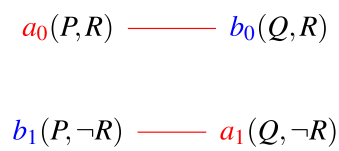

Figure 7.

"shift" the facet $\{a_0, b_0\}$ to $\{a_0, b_1\}$ in Figure 6, but not in Figure 7. We shall see that this is an example of belief revision via the process of "imaging." [31] [21] What simplicial complexes afford us is a way to read off the "nearness" between possible worlds (here modeled as facets) without needing to specify a "nearness" function as an additional component of the model.

Now we proceed to give the details of these ideas formally. This paper will work in a very slightly modified version of simplicial semantics, similar to the work in [25] [36] [19] [26] [24] [22] [23]. Let $\mathfrak{P}$ be a countable set of propositional atoms, and Ag a finite set of agents. Then, for any $P \in \mathfrak{P}$ and $a \in Ag$, we have the following grammar:

$$\varphi ::= P|\bot|\varphi \to \varphi|B_a\varphi$$

Note that, in contrast with much of the previous work, we are interpreting our modal operator B_a as belief, as opposed to knowledge [25] [36] [19] [26] [24] [22] [23]. The reason is, of course, that ultimately we will want to use our simplicial semantics to model belief revision. In light of that, we should define the models for our language.

First, we must define a simplicial complex. Informally, a simplex is a triangle. That is, it is a collection of vertices where each vertex is connected to each other. So, if one has 3 vertices, one is left with the usual triangle. 4 gives a tetrahedron, and in general, n many vertices is an $(n-1)$-dimensional triangle. Triangles of arbitrary dimensions are called simplices, and any subset of the vertices of the simplex is called a *face* of the simplex. The dimension of a face will be the number of vertices. Any face not strictly contained in some larger face is called a *facet*, following convention in the previous literature [19]. So, we could have a triangle, as below:

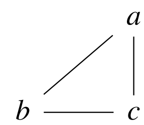

But simplicial complexes are more general than simplices. We could also have two triangles, disconnected, as a simplicial complex:

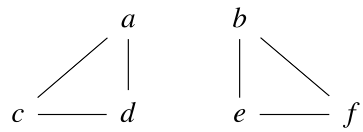

Or, those two triangles could be connected at an edge:

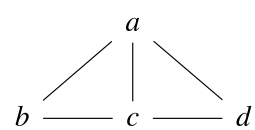

Or connected at a single vertex:

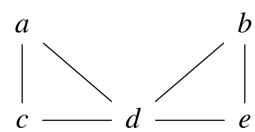

Furthermore, they do not have to be both triangles. Simplices of different sizes could be connected, such as a 2-simplex (triangle) connected to a 1-simplex (edge):

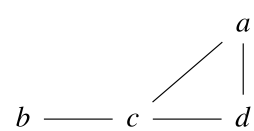

A simplex can be thought of as the collection of its component pieces. A 0-simplex is just a vertex. A 1-simplex, an edge, consists of its two vertices. A 2-simplex, a triangle, consists of 3 edges, which each themselves consist of two vertices, though those vertices will be shared among the edges. In general, simplices consist of the sub-simplexes that compose them, though simplices may have nontrivial intersection. This clues us in to the actual formal definition of a simplicial complex. Formally, a simplicial complex is a set S closed under subsets. That is, if $X \subseteq Y \in S$, then $X \in S$. Any element of $X \in S$ is a *face*, and any element $X \in S$ such that $X \subseteq Y \in S$ implies $X = Y$ is a *facet*.

In our motivating example at the beginning of this section with two agents, we used the vertices of an edge to model perspectives, and treated the edges themselves as akin to possible worlds. In this way, simplicial complexes give us the necessary machinery to model such a conception of uncertainty. We will treat the vertices of a simplicial complex as perspectives, and accordingly each vertex will be assigned to a unique agent. As per our story above, these perspectives will have logical atoms assigned to them. Uncertainty is cached out in terms of facets. To say that a_1 is uncertain between b_1 and b_2 is to say that both $\{a_1, b_1\}$ and $\{a_1, b_2\}$ are faces in the simplicial complex. This would mean that we would model uncertainty by specifying sets of perspectives for each agent. That is, each agent will have their own simplicial complex on the set of perspectives.

We will make two assumptions about simplicial complex models in this paper. The first is that for every agent's complex there can be no face containing two perspectives of the same agent. This assumption means that agents are certain of the hard information in their own perspective. The reason for this assumption is that without it, it is less clear how to interpret what the atoms assigned to a perspective really mean. We have been treating it so far as information that is readily and undeniably available to the agent given that perspective, and this is inconsistent with the idea that the agent could be uncertain about this information. The second assumption is that facets will have at least one (and so, exactly one) perspective for each agent. Facets will be our stand-in for possible worlds, and so it is important for our purposes that each agent have a perspective in each possible world. In future work, this is an assumption worth weakening, and some of the previous literature has indeed explored this in related simplicial semantics. [36] [25]

So, given this informal discussion, we shall turn towards giving the above formally. What we need to define these semantics is a set N of nodes, a function $V : N \to A$ which assigns each node to an agent, called the "coloring" function, and a function $L : N \to 3^{\mathfrak{P}}$ which assigns each node to a set of literals. The interpretation is that $L(n)(P) = 1$ if and only if P is associated with n, $L(n)(P) = 0$ if and only if $\neg P$ is associated with n, and $L(n)(P) = 2$ if and only if neither is associated with n.

Our first key idea is that we can use N, V, and L to create a kind of **maximal** simplicial complex. Like much of the previous literature, we will assume our simplicial complexes are uniquely colored. Specifically, each facet of our complexes has dimension $|Ag|$, and no two nodes are associated to the same agent. We call this condition *UCF* for "Uniquely Colored Facets".[1] Thus, the **Maximal Complex** is the set of subsets of N such that the associated logical content with that subset is consistent, and it satisfies the *UCF* condition. That is, it's the subsets $x \subseteq N$ such that $|x| = |A|$, $\bigcup_{u \in x} L(u)$ is consistent, and for all $u, v \in x$, $V(u) \neq V(v)$. We refer to $\mathfrak{M}(N,V,L)$ as the **Maximal Complex** of N, V, and L.[2] When the context is clear, we will refer to it simply as the maximal complex. The maximal complex can be defined set theoretically as follows:

$$\begin{aligned}\mathfrak{M}(N,V,L) := \{y \in 2^N \mid \exists x \in 2^N (y \subseteq x \\ &\wedge \neg(\exists P \in \mathfrak{P} \wedge (\exists u, v \in x (L(u)(P) = 1 \wedge L(v)(P) = 0))) \\ &\wedge (|x| = |A|) \\ &\wedge (\forall u, v \in x (V(u) \neq V(v))))\}\end{aligned}$$

We are taking all sets of nodes y which are a subset of some set of nodes x satisfying the following conditions: x is consistent, insofar as there is no atom P such that P is true at a

[1] Under the UCF condition, our formal use of the word "facet" aligns with our informal use at the beginning of this section.

[2] We could just as easily say that the function L is a map to consistent sets of literals, as the maximal complex selects faces on the basis of consistency. Any inconsistent vertices would simply not be able to be incorporated into the maximal complex. It is simpler to do it this way, however, as our discussion of revision will reference inconsistent perspectives.

node u in x, and P is false at a node v in x; x contains $|Ag|$-many nodes; finally, each node in x is assigned a different color. Because the subset of a consistent set remains consistent, it's clear that the maximal complex is a simplicial complex. All complexes we consider in this paper will be *UCF* subcomplexes of the maximal complex. In general, a **Simplicial Model** is a tuple $\mathscr{M} = \langle N, V, L, W, \{S_a\}_{a \in Ag}, \rangle$ where N is a set of nodes, $V : N \to Ag$ is a coloring function, $L : N \to 3^{\mathfrak{P}}$, W is a *UCF* subcomplex of $\mathfrak{M}(N,V,L)$, and for each $a \in Ag$, S_a is a *UCF* subcomplex of W.[3]

Under the *UCF* condition, there is a useful function one can define. For any facet X and agent Ag, we say that $\pi_a(X)$ is the unique vertex of X which is a-colored. Also, let $\mathscr{F}(S)$ denote the facets of S for any simplicial complex S. With this in hand, the satisfaction conditions are as follows, supposing $X \in W$:

$$\mathscr{M}, X \vDash P \text{ iff } \exists a \in A(L(\pi_a(X))(P) = 1)$$
$$\mathscr{M}, X \nvDash \bot$$
$$\mathscr{M}, X \vDash \varphi \to \psi \text{ iff, if } \mathscr{M}, X \vDash \varphi \text{ then } \mathscr{M}, X \vDash \psi$$
$$\mathscr{M}, X \vDash B_a\varphi \text{ iff } \forall Y \in \mathscr{F}(S_a) \text{ if } \pi_a(X) = \pi_a(Y) \text{ then } \mathscr{M}, Y \vDash \varphi$$

An important note at this point is that the axiom **D**, or consistency, is not true of this definition of belief. This is because it is possible to have a facet $X \in W$ such that there is no $Y \in \mathscr{F}(S_a)$ where $\pi_a(X) = \pi_a(Y)$. That is, there may be facets where agents consider no worlds possible. This would be simple enough to fix. All we would have to do is ensure that for every $n \in N$ such that $V(n) = a$, there is $X \in \mathscr{F}(S_a)$ such that $n \in X$. That is, every perspective is included in some facet. However, such restrictions will not be easy to maintain as we incorporate action models and belief revision. All of the models we consider, for the most part, will exhibit consistent beliefs. In particular, all of our initial models will validate **D**, and all action updates we define will also validate **D**. But for the ease of some proofs, we do not force this restriction in our definitions.

The previous literature does not directly define models on the basis of consistent unions of perspectives. [19] [36] [25] [26] [24] [22] [23] The closest comparison is [26], where unions of agent perspectives are assumed to be maximal consistent sets by fiat. That is, at each vertex, either P or $\neg P$ is assigned. Here, instead, consistency is the process by which we build the simplicial complex. Other literature assigns propositions directly to the facets themselves, and achieves the same effect [25]. Our semantics attempts to thread the needle between these two choices. However, we also do not assume that the assignment of atoms to vertices necessarily fixes the truth of each literal explicitly on each face. That is, we do not assume that for each literal P and each face X, there is a vertex $x \in X$ containing either P or $\neg P$. Instead, each vertex is given its own assignment of literals, determined by the desires of the modeler. Our definition of truth then determines that, on any face which

[3] As the name of this set W suggests, this is analogous to a set of worlds. Indeed, being a single, non-agent-differentiated complex, it would be very easy to define a knowledge modality using this set, standard with the rest of the literature. Such a modality would have nice properties, such as $K_a\varphi \to B_a\varphi$. Forthcoming work with Adam Bjorndahl explores a system defined as such. However, for the purposes of this memo, having both knowledge and belief present is a distraction. W is introduced here solely as a technical device which will help us prove completeness later. Moreover, without W, we must assume something like if $X \in \bigcup_{a \in Ag} \mathscr{F}(S_a)$, then it becomes valid that at every world *some* agent has true beliefs. This is undesirable. That said, sometimes we will not specify W. In this case, we assume that $W = \mathfrak{M}(N,V,L)$.

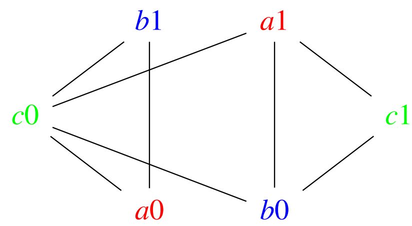

Figure 8.

does not have a vertex containing P or $\neg P$, P is assumed to be false. We call this "truth minimality." [4]

By contrast, our system will make true the following for any $P \in \mathfrak{P}$:

$$\mathbf{NU}: P \to \bigvee_{a \in A} B_a P$$

We call this axiom **NU** for "No Uncertainties," and it says that if an atom is true, then there is some agent who believes it. We verify its soundness below.

Lemma 5.1. *Validity of **NU** for arbitrary $\mathscr{M}$*

Proof. Suppose that $\mathscr{M}, X \vDash P$. Then there is some agent $a \in A$ such that $P \in L(\pi_a(X))$. Then for all Y such that $\pi_a(X) = \pi_a(Y)$, $P \in L(\pi_a(Y))$. The desired result follows. □

NU should be interpreted as saying that worlds (facets) are purely constituted by the perspectives of which they consist. That is, there are no atomic facts that are true except by way of being introduced by some agent's perspective. In general, this is an idealized assumption, like many in the previous literature, and we adopt it here for convenience.

We should pause here to discuss why we need distinct simplicial complexes for each agent. This is directly connected to our shift from interpreting the epistemic modality as belief rather than knowledge. The previous literature does not make this assumption, and it is arguably the biggest departure from that literature in our formalism. [25] [36] [19] [26] [24] [22] [23] So, consider the simplicial model drawn in Figure 8. Here, the black edges should be interpreted as shared among all agents, so in this case, all agents i have the exact same simplicial complex S_i associated with them.

Since $S_a = S_b$ for all $a, b \in Ag$, we can see that this is similar to supposing that there is a single simplicial complex S, rather than a distinct one for each agent. Suppose $a1$ is announced. The facet $X = \{a_0, b_1, c_0\}$ is inconsistent with this, and thus disappears.

[4]Truth minimality is such that our models do not make sound the "Locality" axiom present in previous literature [19]. In this literature, A partition is made on $\mathfrak{P}$ over Ag, so that each $a \in A$ corresponds with a distinct subset of $\mathfrak{P}$. The set of atoms assigned to a given agent are called the "local variables" for that agent. [19] Given this, the following is sound when P is local to a:

$$\mathbf{LOC}: B_a P \vee B_a \neg P$$

The previous literature would interpret the modality as knowledge, instead of belief, but the axiom itself is the same. Note that, in the setting where we have local variables and the axioms of **K**, **LOC** clearly implies **NU**, but the reverse does not hold.

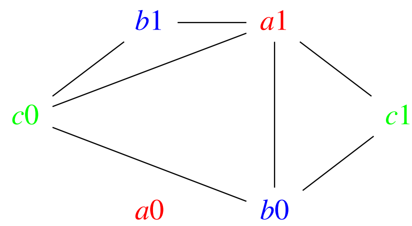

Figure 9.

However, there is a nearby[5] facet, $Y = \{a_1, b_1, c_0\}$, which is consistent with $a1$. Indeed, Y is the facet satisfying $a1$ whose intersection with X is largest. So, the agents update their belief complex as drawn in Figure 9.

However, this has some very undesirable consequences. Indeed, imagine being agent a with perspective $a1$. What is announced is your hard information. You already believe this in a very strong sense, so intuitively, you should learn only that $a1$ is common knowledge. In this case, a's set of beliefs should strictly expand. However, this is not what happened. Notably, agent a in perspective $a1$ is now uncertain about b's bit value! This of course contradicts the good intuition that when $a1$ is announced, a's set of beliefs should strictly expand. It is easy to construct examples like this if there is not a distinct belief complex for each agent; hence, we give each agent their own simplicial complex.

We conclude with an exploration of why the **D** rule is not sound in our system. Suppose that there is an a-perspective x_a and no facet X in S_a such that $\pi_a(X) = x_a$. That is, a considers no worlds possible when she is in perspective x_a. It is easy to see, then, that $B_a\bot$ is true at X. Moreover, if $B_a\bot$ is true at X, then there cannot be any Y in S_a such that $\pi_a(X) = \pi_a(Y)$. This means that when $B_a\bot$ is true at a facet X, we know that a has eliminated any facet containing $\pi_a(X)$. This interpretation seems distinct from previous literature exploring making the axiom **D** false in simplicial models [25] [36]. This literature interprets $B_a\bot$ to mean that a has had a kind of signal failure—they are no longer an active participant in the signaling game. Notably, this literature falsifies **D** by weakening *UCF*. We do so by, in contrast, restricting the maximal complex in distinct ways for each agent. Given this distinction, the different interpretations for the same formulas seem plausible.

5.2 Belief Revision in Simplicial Complexes: Informally

Belief Revision is a well trodden idea in formal philosophy. [27] [31] The main underlying intuition is the following: If an agent learns that world w is false, they may not merely remove w from their model. They may additionally incorporate new worlds, not present in the model before, which are sufficiently "nearby" to w but not contradictory with what they just learned. This is very important when dealing with modal concepts of belief. If I believe that φ, but this belief is false, then when I learn $\neg\varphi$, under normal circumstances I would cease to believe any worlds are possible. Obviously this is quite extreme!

Unfortunately, this notion of "nearer" is not easy to discern in practice. In [31], any two worlds may be as near or far from one another as the modeler desires. Simplicial

[5] "Nearby" in the sense of Section 5.3.

Semantics, we will argue, offers a particular restriction on the concept of revision which is both already present in the formalism, and independently conceptually motivated. We will begin with that conceptual motivation. We, importantly, do not want to overstate our case here. We do not claim that the intuition we motivate is the only plausible intuition one could have about what constitutes nearness. All we are hoping to demonstrate to the reader is a particular intuition which seems reasonable enough that it is worth following up on in a formal setting.

Consider the following case: You are agent a. You believe your friend b is in Schenley Park, and your other friend c is in Frick park. However, you learn from a trusted source, importantly one that you trust more than your previous beliefs, that b is in fact not in Schenley park. The first intuition is the following: If you only learn that b is not in Schenley park, why would you stop believing that c is in Frick park? When I believe that b is in Schenley and c is in Frick, every world I consider possible has to be a world where both of these facts are true. Hence, when I learn that b is not in Schenley, I have to rule out all such worlds. Intuitively, there are two kinds of candidate replacements worlds I could re-incorporate into belief set: those where b is not in Schenley, and c is in Frick, and those where b is not in Schenley and c is not in Frick. It seems natural that the former is a much better intuition.

But why are worlds where c is still in Frick better candidates for "nearer" worlds to the ones we started with? There are surely many possible explanations to this question, corresponding with various intuitions, though I will focus on the following intuition: worlds are "nearer" when they preserve more of the perspectives of your fellow agents. To see this, suppose, as before, that you believe your friend b is in Schenley Park, and your friend c is in Frick park. However, additionally, you believe that your friend d is in Highland park. Imagine in particular that each friend has reported to you the information about their location. Then, you learn from an even more trusted source the following fact: Either b and c are both indoors (i.e., not at any park), or d is indoors. Intuitively, there are three kinds of replacement worlds that you could appeal to:

1. d is not in Highland
2. b is not in Schenley and c is not in Frick
3. b is not in Schenley, c is not in Frick, and d is not in Highland

I believe that each collection of worlds is less appealing as a potential revision set than the last. The reason is that, in the first collection, you are only throwing out d's testimony. In the second, you are throwing out both b and c's testimony, and in the third collection, you believe nobody's testimony. So, in an attempt to be charitable towards as many friends as possible, you revise to the first collection. This has the following unusual property. Before learning everything, you believe that b is in Schenley, c is in Frick, and d is in Highland. Then, after learning that either b and c are both indoors, or d is indoors, you believe that d is not in Highland. This is in an effort to be charitable to b and c, who in their conjunction seem to "outrank" d. Hence, you reject the information about d.

This intuition depends on a few assumptions. First, and maybe most obviously, one must "rank" all of their fellow agents equally. If d were a far more trustworthy source of information than b or c, that would seem to force the second collection of worlds above to be the best collection to revise to. Relatedly, it seems that we need to learn information in perspectives by testimony of those perspectives. For example, we learn that b is in

Schenley because b reported it to us. In the examples below, this will be cached out in that the information in perspectives that we learn will consist of information about privately held bit values associated with each agent. And, finally, it can't be that these bits of information are dramatically different in the weight we give to them. It could be that, maybe, while we trust our friends equally, we place a lot more weight on the fact that d is in Highland than on the fact that b is in Schenley and c is in Frick. Again, we will cache out this assumption in our examples by assuming the relevant facts we learn concern information about privately held bit values. However, in a more general setting, it is good to know that in "The promise, and limitations, of gossip protocols", Birman restricts gossip protocols to be those that "The information exchanged during these interactions is of (small) bounded size." [10] This seems to line up with the idea that we can't significantly rank the information of certain signals over others. Thus, the overriding information needs to be the number of friends which we are forced to change our mind about.

There are two final notes I want to make about these informal examples. Firstly, all of these examples can be done in the language of local variables. [19]. So, if one prefers that setting to ours, modifying the below work to live in that setting seems entirely well motivated. Secondly, one might notice that in all of these examples, the information we learn rules out every world we consider possible. Perhaps, then, one might argue that this is the only setting in which belief revision is warranted. The authors would agree, in fact, that this idea is very reasonable. We will address this briefly below, though we do not formally explore this idea in this memo.

5.3 Belief Revision in Simplicial Complexes: Formally

With this intuition on hand, we shall return to describing the concept formally. Now that we've set up our semantics for belief, we should explore some examples and explain how we will draw them for the duration of this paper. In general, we will color code vertices to correspond with the agent. Vertices will be denoted by the agent and a subscript if necessary, with the literals assigned to that vertex written afterwards. Sometimes, we will denote a vertex in, where i is an agent and n is a nonsubscripted numeral 0 or 1. Here, the vertex $i0$ should be thought of as abbreviating a vertex $i(P_i)$, where P_i is a literal associated with agent i, and $i1$ abbreviates $i(\neg P_i)$. We will treat P_i as a bit value associated with agent i, which can either be 1 or 0, corresponding with P_i being false or true respectively. So, $i0$ is an i-vertex where i has bit value 0, and $i1$ an i-vertex where i has bit value 1. When we are dealing with such a setting, we will sometimes say that $i0$ is announced instead of P_i, and similarly that $i1$ is announced instead of $\neg P_i$, since there is a 1-1 correspondence in this case between atoms and perspectives. In such a setting, we can refer to P_i as the singular, unique literal at the node $i0$ and $\neg P_i$ as the unique literal at $i1$. One should note that the gossip protocol motivated in Section 2 is one such setting where this notation is useful. There, the only relevant facts are private agent bit values, and these constitute all of the perspectives.

Consider the model drawn in Figure 10. This depicts the simplicial complex S_b. There are two facets, namely $\{a_0, b\}$ and $\{a_2, b\}$. The former makes true Q and P, where as the latter makes true $\neg Q$ and P, due to truth minimality. Note, then, that if a_2 were replaced with $a_2(\neg Q)$, the truth values of all formulas in the model would be unaltered. Facts such as this will be of paramount importance in the next section. Also take note of $a_1(\neg Q, \neg P, R)$ as an unincorporated vertex. This does not contradict *UCF*, as indeed, every facet in the

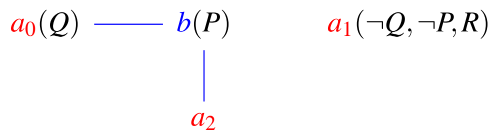

Figure 10.

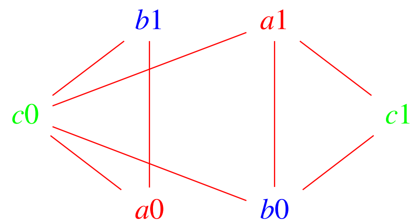

Figure 11.

model is uniquely colored. Note further that, at either of the two facets, $\neg R$ is true.

Also consider the model drawn in Figure 11. Note that this example uses the abbreviated language outlined above, where $i0$ corresponds with $i(P_i)$ and $i1$ with $i(\neg P_i)$ for all $i \in Ag$. Since it is drawn red, Figure 11 depicts the simplicial complex S_a. There are three facets: $\{a1, b0, c1\}$, $\{a1, b0, c0\}$, and $\{a0, b1, c0\}$. The literals true at each of these facets, as explained above, correspond with the vertices themselves. So at $\{a1, b0, c1\}$, b_0 is true, while a_0 and c_0 are false (equivalently, a_1 and c_1 are true).

Now to turn to how we will define belief revision in simplicial models. Let's motivate the main idea with by looking again at Figure 6. Note that S_b is depicted. At any facet X, we have that $\mathscr{M}, X \vDash \neg\hat{B}_b(\neg Q)$.[6] However, in this model, nothing prevented the possibility of the modeler selecting $Z := \{a_0, b_1\}$ as an element of S_b. If they had, of course, $\mathscr{M}, Z \vDash \hat{B}_b(\neg Q)$.

However, by contrast, recall Figure 7. Note that S_b is depicted. Note that this model differs only in the presence $\neg R$ in the vertices b_1 and a_1. Because of truth minimality, this means that the facets satisfy the exact same formulas in this example as in the previous. So in this way, the two models are equivalent. However, in this situation, the modeler could not have included the facet $\{a_0, b_1\}$ as an element of S_b. Nor could the modeler include $\{a_1, b_0\}$. Note that these are the two uniquely colored facets in the maximal complex which satisfy $\neg Q$. So, in a sense, in this model, it is not the case that b could come to consider $\neg Q$ possible. Note that this lines up with our informal examples; we are treating worlds as nearer when they preserve more perspectives. The goal of this paper is to utilize this intuition to give a semantics for Belief Revision.

Consider Figure 6 again. Imagine that the formula $\neg(P \wedge Q)$ was publicly announced, and consider what this might seem like from b's point of view. After $\neg(P \wedge Q)$ is announced, b knows the facets $\{a_0, b_0\}$ and $\{a_1, b_1\}$ are no longer tenable, as each of these facets makes true $P \wedge Q$. In a usual update, b would delete these facets and be done. However, we want to take some intuition from the literature on *imaging* [31] [21]. Imaging describes

[6] $\hat{B_i}\varphi$ is an abbreviation for $\neg B_i \neg\varphi$. It can be read as "agent i considers φ possible"

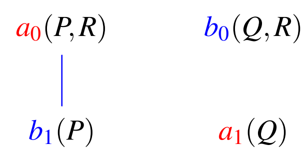

Figure 12.

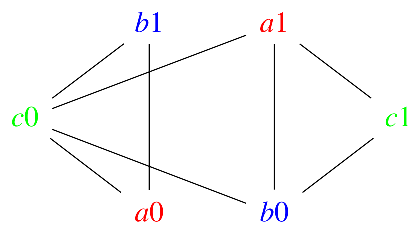

Figure 13.

an update procedure whereby worlds eliminated after learning a piece of information may be replaced by a "selection function." Suppose w is ruled out after φ is learned. Then rather than just erasing w, the agent might use the selection function to come up with a set of worlds, all of which satisfying φ, which will replace w in the model. Often this selection function takes the form of a "preference relation," as is the case in the AGM Belief Revision Literature. [27] Here, we imagine the relation as signifying "degrees of preference" or "degrees of nearness" between worlds. So, when w is ruled out, one replaces it with the nearest or most preferred set of worlds still satisfying φ. We believe simplicial complexes afford us the possibility of doing something similar.

Consider the model from the perspective b_0. The facet $\{a_0, b_0\}$ has been deleted, as this facet makes true $P \wedge Q$. Is there a compelling replacement facet? A good candidate replacement would have to contain b_0, and also make true $\neg(P \wedge Q)$. Indeed, such a replacement does exist. $\{a_1, b_0\}$ fits the bill quite nicely! Similarly, viewing the from perspective b_1, $\{a_0, b_1\}$ is a good replacement for the removed facet $\{a_1, b_1\}$. We draw this as Figure 9.

This example is a bit crude, in that there really is in every case only one candidate replacement for every world which disappears, and that replacement suffices. Consider instead, then, the model in Figure 13.

For the sake of simplicity, interpret the above model where S is the simplicial complex where every uniquely-colored triangle is a facet, and $S_a = S_b = S_c$ is the above sub-simplicial complex. Suppose a_1 is announced. We want to view what happens from the point of view of each perspective. Consider first agent a. If agent a is in perspective a_1, of course, nothing changes. In this perspective, a considers two worlds possible, and both of these are still consistent with the announced information. By contrast, if agent a is in perspective a_0, the facet $\{a_0, b_1, c_0\}$ is not consistent with the announced information. However, a_0 is welcome to try to replace that world/facet with another that would be consistent with the available information. However, because a_0 is inconsistent with a_1, the announced

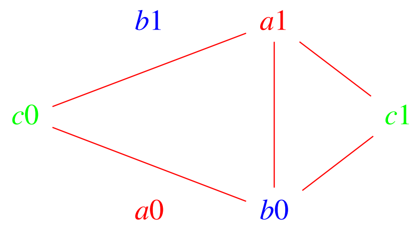

Figure 14.

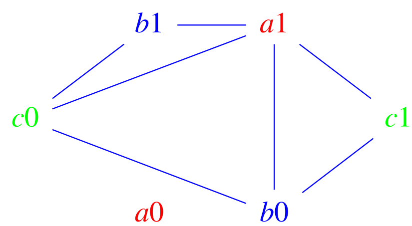

Figure 15.

information, of course no such world/facet exists. Therefore, a's updated simplicial complex is as drawn in Figure 14.

Now consider agent b. If b is in perspective b_0, they in the same situation as a_1. They consider two faces possible, neither of which are removed by the announcement of a_1. Now consider perspective b_1. They consider one facet possible, namely $\{a_0, b_1, c_0\}$. Of course, this facet is inconsistent with the announcement a_1. So b could either choose to replace this facet or erase it. As before, one plausible condition for a replacement facet is that it should contain b_1. Another is that it should make true a_1. Another necessary condition is that it should come from the maximal complex. This leaves two candidate facets for replacement, $\{a_1, b_1, c_0\}$ and $\{a_1, b_1, c_1\}$. Note that the former has a larger intersection with the removed facet, $\{a_0, b_1, c_0\}$. Because it shares more perspectives in common, we want to argue that this world is "nearer" to the removed world. The intuition is that worlds that share more perspectives have more in common, and this shared information is the sense in which they are nearer. If we use this intuition, then b's updated simplicial complex is as drawn in Figure 15.

Finally, consider agent c. At $c1$, things are easy to imagine. $c1$ considers a single facet possible, and this facet already makes true $a1$. At $c0$, however, things are different. The facet $\{a_0, b_1, c_0\}$ is obviously no longer tenable. $c0$ then seeks to replace it. Two candidates are $\{a_1, b_0, c_0\}$ and $\{a_1, b_1, c_0\}$. Both of these contain $c0$ and make true $a1$. However, why should the latter be "nearer"? Applying our intuition from before, note that the intersection between $\{a_1, b_1, c_0\}$ and $\{a_0, b_1, c_0\}$ is larger than the intersection between $\{a_1, b_0, c_0\}$ and $\{a_0, b_1, c_0\}$. In this sense, the former is "nearer" to the removed complex, and thus a better counterfactual replacement. We draw this in Figure 16.

There is a different intuition one could appeal to here. Consider Figure 16 from b's point of view. Suppose their initial complex is the same. By appealing to intersection size, their update complex is also the same as above. However, the situation from $b1$'s perspective is

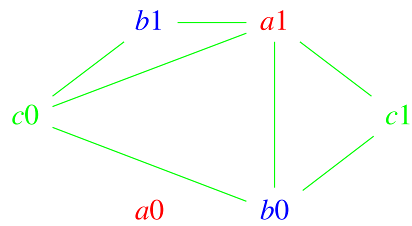

Figure 16.

distinct. The only facet $b1$ considers possible is eliminated in the update. For this reason, switching that facet to $\{a1, b1, c0\}$ may make the choice to switch seem more "forced" onto b, and therefore more reasonable. Arguably, this would mean that $c0$ should not replace the lost facet, as $c0$ still considers $\{a0, b1, c0\}$ possible after eliminating those where $a0$ is true. This intuition would capture the idea that agents should update purely by removing facets unless they remove every facet they consider possible, and only in that extreme case should they "image." This would be an intuition halfway between the classical, Bayesian update (only removing facets) and imaging. While we believe this is worth exploring, we will set such a possibility aside for the remainder of the paper.

We now move to formally define our new update procedure, where replacement worlds are those with maximal intersection with removed worlds. Call the updated model $\mathscr{M}[\varphi]$:

Definition 5.1. Consider a *UCF* simplicial model $\mathscr{M} = \langle N, V, L, \{S_a\}_{a\in A}, W\rangle$. Let $N_\varphi := N$, $V_\varphi := V$, and $L_\varphi := L$.

Let $\mathbf{F}_\varphi$ be the subset of W which satisfy φ. We need to define a function $\mathscr{R}_a : \mathscr{F}(S_a) \to W$. We call $\mathscr{R}_a$ the "replacement" function.

$$\mathscr{R}_a(F) = W \cap \{G \in \mathbf{F}_\varphi \mid (\pi_a(F) = \pi_a(G)) \\ \wedge \forall X \in \mathbf{F}_\varphi((\pi_a(X) = \pi_a(G)) \to (|F \cap G| \geq |X \cap G|))\}$$

Note that, if $G \in \mathbf{F}_\varphi$, then $\mathscr{R}_a(G) = \{G\}$. Let $S_{a,\varphi}$ be the simplicial complex whose facets are elements of the set $\bigcup_{F\in\mathscr{F}(S_a)} \mathscr{R}_a(F)$. It is easy to check that this is a *UCF* sub-simplicial complex of the maximal complex of N_φ, V_φ, and L_φ. W is unaltered. We say that $\mathscr{M}[\varphi] := \langle N_\varphi, V_\varphi, L_\varphi, \{S_{a,\varphi}\}_{a\in A}\rangle$ is the update of $\mathscr{M}$ after imaging on φ.

The next obvious question to ask would be: What axioms are these models sound and complete for? The simpler question—which for this memo is the only one we shall address—is soundness and completeness for the language without public announcement formulas. To answer that, we will use techniques already known from the existing literature [25] [36] [19] [26] [24]. There, completeness is proven by a "categorical equivalence" between a category of simplicial models and a category of Kripke frame models. The equivalence consists of two Functors, one which translates a simplicial model into a Kripke model, and another which goes the other direction. If the class of Kripke models is known to be complete for a certain collection of formulas, and the functors preserve the truth of logical formulas, then one can say that the class of simplicial models is complete with respect to the same collection of formulas. We use this technique to conjecture the following:

Theorem 5.2. *The class of UCF simplicial models is sound and complete with respect to propositional logic plus* $\mathbf{K45}+\mathbf{NU}$*, modus ponens and necessitation for each belief modality.*

The proof is given in the appendix, and is similar to those in the existing literature, beginning with a relevant categorical equivalence between a category of frame models and a category of simplicial models. [19] [26] There are some interesting technical concerns that make lifting this categorical equivalence to the completeness proof slightly unusual, though similar to concerns mentioned in [25].

6 Action Models

Action models started in the philosophical logic and epistemic logic literature as a means to model "updating" an epistemic model with what could be called an "epistemic scenario". [6] If the latter seems broad, that is because it is. An "epistemic scenario" will be a separate epistemic model, specifying its own worlds where agents can tell apart some worlds but not others. Therefore, the output of matching some input model to this new scenario will be a restriction of the Cartesian product of the input and the scenario. New worlds will be ordered pairs, whose first entry is from the input and whose second is from the scenario. This intuitively makes sense - we are matching input worlds with potential changes, as delineated by the scenario. However, the point is that the scenario will rule out some ordered pairs. If a world in the scenario is somehow incompatible with a world in the input, that world can't be added in the new update model.

The preceding discussion is deliberately vague: Action models are capable of describing a great many things. However, for our purposes, it will suffice to say that action models describe *signals*. In the previous section, our capacity to model something "dynamic" was restricted to the landscape of public announcements. However, public announcements are very unrealistic as signals in any kind of applied setting. More typically, we are talking about signals like "a sends a message to b and c is aware that a message is sent, who sent it, and who received it, but is unaware of the content of the message." Action models are ideally suited for representing such signals, and for the rest of the paper, we will explore actions purely in terms of how they can model single signals and sequences of signals such as this one.

Anyone familiar with the realm of distributed computing will recognize the possibility of a connection with action models immediately. Indeed, any distributed algorithm is a description of a sequence of signals similar to if not identical to the one described above. Such a connection is already well understood in the literature as well. [34] [24] [19] [3] In particular, the paper [3] already uses action models to model gossip protocols, which will be akin to our end goal in this paper.

6.1 Kripke Action Models

We will first explore action models as described in [6], over Kripke frames, before developing our own definition of actions similar to [19] and [24], though modified for our new particular simplicial semantics. Consider the following signal:

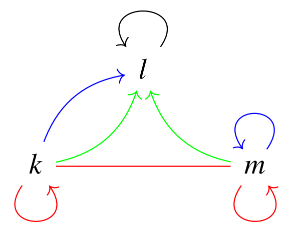

Figure 17.

"a sends a message to b whose content is P_a. However, c is unaware that any message was sent at all, and furthermore, there is a chance the message will not be received, in which case b will think no message was sent."

To model this signal as an action model, we need to represent it using something very similar to a Kripke frame. Therefore the first step is to describe the worlds. As it stands, there are three. The first world, which we shall call k, is the world where a message is sent, but not received. The second world, which we shall call m, is the world where a message is sent, but is received. Naturally, a cannot tell these worlds apart. However, b can. At m, b considers m possible. However, at k, b should not consider k possible. After all, if the message is not received, b thinks that no message was sent. This tells us that we need a third world, l, which is the world where nothing happens. At world k, b falsely believes the world is l. Similarly, at both k and m, c falsely believes the worlds is l. The only curious touch is that, at world l, a should think the world is l. After all, this is the world where no message was sent - if a doesn't send a message at all, they think they are in world l. This description gives us the Kripke frame in Figure 17

As before, red corresponds to $\mathscr{R}_a$, blue to $\mathscr{R}_b$, and green to $\mathscr{R}_c$. Black edges are those edges present in all three relations. However, the above frame is insufficient to describe the action. After all, these worlds are not compatible with every possible input. Imagine an input world w where $\neg P_a$ were true. Presumably, since P_a is a fact that a is aware of, i.e., a secret pertaining to a or a bit value which a is aware of, it is not possible that a sends out P_a at w (we are assuming that a is not lying and the signal always goes through correctly if it goes through). Hence, w is incompatible with k and m. Actions use the notion of "preconditions" to bear this out. Basically, in order for world k to obtain, it must be the case, and hence, it is a precondition, that P_a is true. We will mark preconditions with parenthesis, like we do for truth conditions in typical models, and let context differentiate. However, preconditions, at least for actions over Kripke frames, are distinct from truth conditions, as a world can have a precondition which is not an atomic formula. Writing in the preconditions gives us the the model depicted in Figure 18

l is compatible with any input, however, k and m require that P_a to be true in order for them to obtain. Nicely, this is precisely the action given on page 22 of [5], just with the precondition $K_a\varphi$ swapped out for P_a: these preconditions are equivalent so long as we only ever assign P_a to a-colored perspectives in the input, and hence, P_a will be true if and only if B_aP_a is true.

Now that we understand what actions look like, let's imagine what updating an input with an action model might yield. Consider the input model as drawn in Figure 19. This

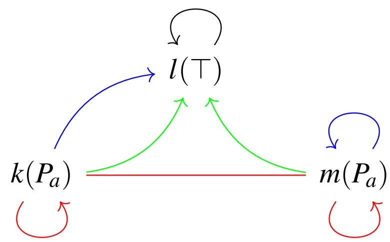

Figure 18.

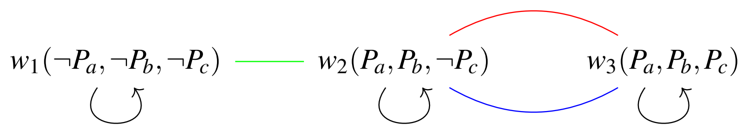

Figure 19.

input is easy enough to interpret - it is the Kripke equivalent of the simplicial example 13, with the truth values of the atomic formulae specified explicitly. Following our heuristic description of how to update some input with an action, we know that worlds in the update are ordered pairs, whose first term is from the input, and whose second term is from the action. However, not all ordered pairs are allowed - if at a world w, the precondition for the action world n is false, then the pair (w,n) cannot be allowed in the action update. Worlds in the update model are typically given as ordered pairs, with a world from the input in the first place and a world from the output in the second place. The allowable pairs are those where the first world satisfies the preconditions of the second world. For our example, using figures 18 and 19, this gives us the collection of worlds drawn in Figure 20.

What should the truth conditions at these worlds be? Well, if (w_2,k) is the world w_2, but where the message is sent but not received, then this should take on the same truth values as w_2. Indeed, actions don't change what's true at a world. Rather, they create copies of worlds which can be differently interpreted. This gives us the model as drawn in Figure 21

All that is left to determine is each R_i for $i \in Ag$. When should $(w,n)R_i(w',n')$? Well, (w,n) is world w where n has occurred. So, if w and w' are differentiable, so should (w,n) and (w',n'). The action does not change what is true at worlds, and so does not make

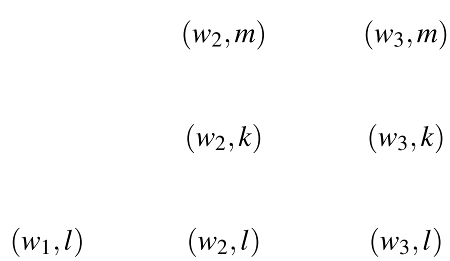

Figure 20.

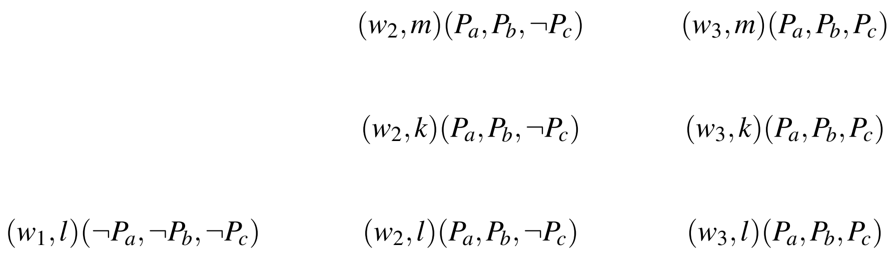

Figure 21.

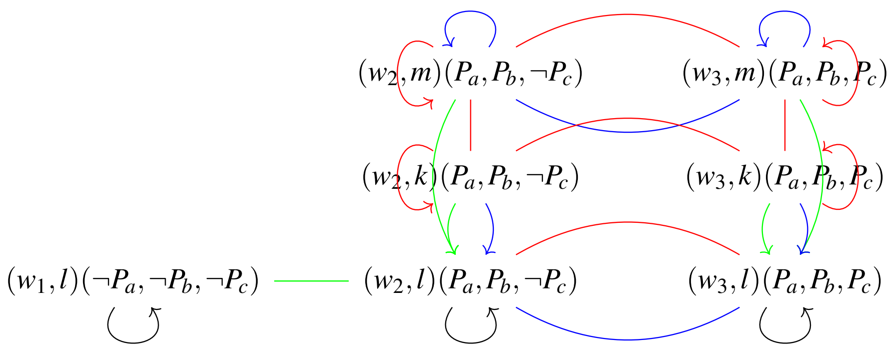

Figure 22.

worlds which were differentiable suddenly undifferentiable. However, if n and n' are differentiable, then (w,n) and (w',n') should be differentiable, even if w and w' are undifferentiable. After all, this is the main point - these were the same world, or at least undifferentiable worlds, but something different has occurred in each. This gives us a very natural definition: $(w,n)R_i(w',n')$ if and only if wR_iw' and nR_in'. This gives us the output model as drawn in Figure 22

Using our heuristic understanding, it's not difficult to imagine what each world in the above model means. At any world (w_j,n), c should only consider possible those worlds (w_h,l) such that $w_j\mathscr{R}_cw_h$, because c always falsely believes no signal was sent. Similarly, at world (w_2,k), all four of (w_2,k), (w_2,m), (w_3,k), and (w_3,m) are indistinguishable for a, as a cannot tell apart w_2 and w_3, nor can a tell apart k from m. In general, one can think of an action update as creating (restricted) "copies" of the input model, and these "copies" are indistinguishable in terms of the relations in the action. Hence, the l "copy" of the input model contains the only worlds which c considers possible, etc.

This gives us the following formal definition of an action:

Definition 6.1. Formally, an action model A is a tuple $\langle W_A,\{R_{A,i}\}_{i\in Ag},pre\rangle$ such that pre is a function from worlds $w\in W$ to sets of formulas.[7] Given a Kripke model $\mathscr{M}:=\langle W,\{R_i\}_{i\in Ag},V\rangle$, and an action model A, we define the action update model $\mathscr{M}[A]:=$

[7] W_A is a set of worlds, and each $R_{A,i}$ is a binary relation on those worlds.

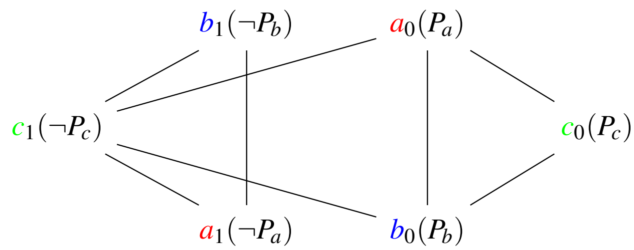

Figure 23.

$\langle W[A], \{R_i[A]\}_{i \in Ag}, V[A] \rangle$ as follows:[8]

$$W[A] := \{(x,y) \in W \times W_A | \mathscr{M}, x \vDash pre(y)\}$$
$$R_i[A](x,y) := \{(x',y') \in W[A] | x' \in R_i(x) \wedge y' \in R_{A,i}(y)\}$$
$$V[A](P) := \{(x,y) \in W[A] | x \in V(P)\}$$

6.2 Simplicial Action Models

With this in hand, we turn our attention to simplicial actions. We will work through the example just described but using simplicial, as opposed to frame, based models. The simplicial translation of the input, per the previous section, is Figure 13. We explicitly label the truth values of the atoms in Figure 23.

We now turn towards defining the simplicial version of the action model. While we could simply translate the frame version above into a simplicial complex, it will be informative to go through the exercise of directly interpreting the intuition into a simplicial complex. First we will need to itemize perspectives. The simplest agent to start with is agent c. Note that, no matter how the world looks, agent c believes that nothing has happened. So, there is a single c perspective, which we shall label "c". More complex is the situation for agent a. Whether or not the message is received by b, agent a has the same perspective: The message was sent. But there is another relevant perspective of agent a's we must consider. c, naturally, never considers that a is in the perspective where the message is sent - they believe that a has not sent a message, no matter what. So there are two perspectives: a_s, where a has sent the message, and $a_{\neg s}$, where a has not. The story is much the same for b. So long as the message is not received, b believes nothing has happened: we will call this $b_{\neg r}$. If the message is received, of course, b observes this, and so their perspective is b_r. This gives us the set of perspectives drawn in Figure 24

The next piece of information to identify are the facets. There are three - corresponding to the worlds where nothing happens, where the message is sent and not received, and where it is received. Importantly, all three agents consider the world where nothing happens possible. This would be the facet $\{a_{\neg s}, b_{\neg r}, c\}$. To see why, imagine what worlds a considers possible at $a_{\neg s}$. Similarly, $\{a_s, b_{\neg r}, c\}$ is a facet - this is the world where the message is sent, but not received. This is a world that a considers possible, but neither b nor c considers it possible. At this world, both b and c believe, falsely, that nothing has happened. And

[8] $W[A]$ is a subset of $W \times W_A$, Each $R_i[A]$ a binary relation on $W[A]$, and $V[A] : \mathfrak{P} \to 2^{W[A]}$, just as in any Kripke model.

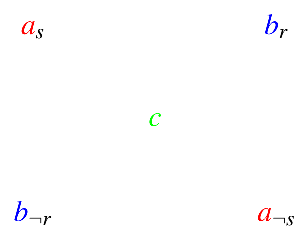

Figure 24.

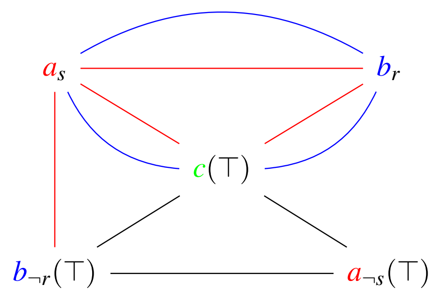

Figure 25.

finally, $\{a_s, b_r, c\}$ is a facet. This world is one that both a and b consider possible, as only c has false beliefs in this world. This gives us the simplicial complex drawn in Figure 25

The last thing to handle is the representation of the preconditions. Note that, to be in perspective a_s, there is a sense in which it must be the case that P_a is assigned to that perspective. This is because a wouldn't announce P_a unless P_a were true, and P_a is true when it is assigned to a's perspective. The other interesting perspective is b_r. It doesn't make sense to say that P_a has to be true at b_r to be in the action-perspective b_r. However, if b has perspective b_r, they observe that P_a is true by way of the received message. We will therefore use what we will call a "postcondition", similar to what is given in [19]. The perspective b will have P_a added to it after the action product is computed. We denote postconditions with brackets. This gives us the action model drawn in Figure 26

When computing the action product, we will want to take orders pairs of perspectives,

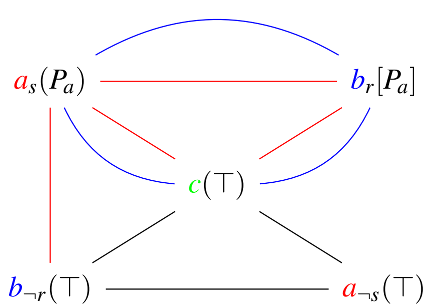

Figure 26.

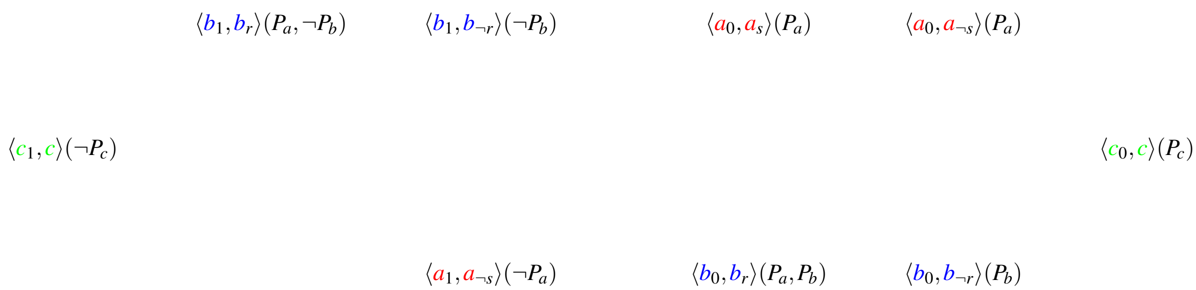

Figure 27.

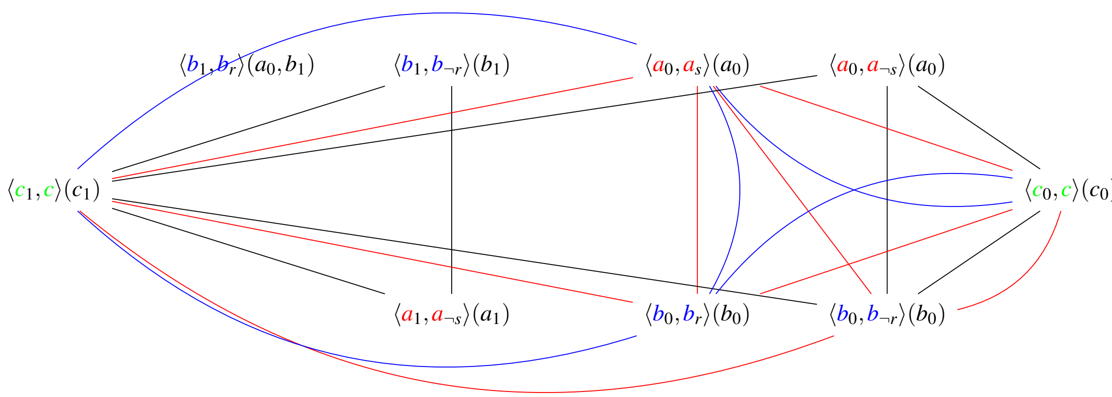

Figure 28.

akin to [19]. If x is a world in the input, and y is a world in the action model, (x,y) will be a perspective in the product if and only if the formulas assigned to x are consistent with those at y. Also, x and y need to be assigned the same color. Of note, this is not necessarily the only way to set this up, and we shall explore that when we give the definition of actions in simplicial complexes formally. Additionally, perspectives should take on the truth values of the input, except when overwritten by a postcondition. As described, this gives us the set of perspectives drawn in Figure 27

The question then is when should a set of perspectives form a facet? If F is a set of perspectives in the product, such that for each $i \in Ag$ there is a unique element of F whose color is i, then F is a facet if and only if the set of x such that there is a y such that $(x,y) \in F$ is a facet in the input, the set of y such that there is an x such that $(x,y) \in F$ is a facet in the action model, and F is consistent. By this last, we mean that there is no contradictory assignment of literals to the vertices in F. This gives us the simplicial complex drawn in Figure 28. Significantly, this model is precisely the translation of the output in the frame setting into a simplicial model.

We can now give the definitions formally. In general, we will use definitions for actions similar to those in [25] and [19]. Actions will have both pre and post conditions. We will denote these first with parenthesis, and the latter with brackets. Formally, an action A is a tuple $\langle N_A, V_A, L_A, \{S_{A,a}\}_{a \in Ag}, Post \rangle$, where $Post : N_a \to 3^{\mathfrak{P}}$.[9] Given a model $\mathscr{M} = \langle N, V, L, \{S_a\}_{a \in Ag} \rangle$, we define the update model

$$\mathscr{M}[A] := \langle N[A], V[A], L[A], \{S_a[A]\}_{a \in Ag} \rangle^{1011}$$

as follows:

$$\begin{aligned}
N[A] &:= \{(x,y) \in N \times N_A | V(x) = V_A(y) \\
&\quad \wedge \forall P \in \mathfrak{P}((L(x)(P) = 1 \to L_A(y)(P) \neq 0) \\
&\quad \wedge (L(x)(P) = 0 \to L_A(y)(P) \neq 1))\} \\
V[A](x,y) &:= V(x) \\
L[A](x,y)(P) &:= \begin{cases} Post(y)(P) \text{ if } Post(y)(P) \neq 2 \\ L_A(y)(P) \text{ if } L(x)(P) = 2 \\ L(x) \text{ otherwise} \end{cases} \\
\mathscr{F}(S_a[A]) &:= \left\{ F \in 2^{N[A]} \;\middle|\; \begin{array}{c} \pi_0(F) \in \mathscr{F}(S_a) \\ \wedge \pi_1(F) \in \mathscr{F}(S_{a,A}) \\ \wedge F \text{ is consistent w.r.t. } L[A] \end{array} \right\}
\end{aligned}$$

We say that a facet X is consistent with respect to a function L if, for all $x, y \in X$ and $P \in \mathfrak{P}$, $L(x)(P) = 1$ implies $L(y)(P) \neq 0$, and $L(x)(P) = 0$ implies $L(y)(P) \neq 1$. Let us show that this simplicial model satisfies the UCF property.

Lemma 6.1. *For any input model $\mathscr{M}$ with the UCF property, and action model A, the output model $\mathscr{M}[A]$ also satisfies the UCF property.*

Proof. Fix $b \in Ag$. If $\pi_0(F) \in \mathscr{F}(S_a)$, then there is $x \in \pi_0(F) \in \mathscr{F}(S_a)$ such that $V(x) = b$, as S_a satisfies the UCF property. So, there is $y \in V_A$ such that $(x,y) \in F$, and $V[A](x,y) = b$, as desired. Suppose $(x,y), (x',y') \in F$ and $V[A](x,y) = V[A](x',y') = b$. So $V(x) = V(y) = V_A(x') = V_A(y') = b$. The condition that $\pi_0(F) \in \mathscr{F}(S_a)$ and $\pi_1(F) \in \mathscr{F}(S_{A,a})$, and the fact that both S_A and $S_{A,a}$ are UCF, means that there is a unique $w \in N$ such that $(w,y) \in F$ and $V(w) = b$, and similarly a unique z such that $(x,z) \in F$ and $V_A(z) = b$. So, $x = x'$ and $y = y'$, it follows that $(x,y) = (x',y')$, as desired. □

Our definition of $N[A]$ is not the only one which could be motivated. In fact, all three of the definitions below are defensible.

- $\forall P \in \mathfrak{P}((L(x)(P) = 1 \to L_A(y)(P) = 1) \wedge (L(x)(P) = 0 \to L_A(y)(P) = 0))$

[9] N_A is a set of nodes, $V_A : N_A \to Ag$, $L_A : N_A \to 3^{\mathfrak{P}}$, and each $S_{A,a}$ is a UCF subcomplex of $\mathfrak{M}(N_A, V_A, L_A)$, just as in a regular simplicial model.

[10] $N[A]$ is a subset of $N \times N_A$, $V[A] : N[A] \to Ag$, $L[A] : N[A] \to 3^{\mathfrak{P}}$, and each $S_a[A]$ is a UCF subcomplex of $\mathfrak{M}(N[A], V[A], L[A])$, just as in a regular simplicial model.

[11] Since we have not been specifying W in any of these examples, we should recall that in such a case we assume W, or the set of possible facets, is always the relevant maximal complex.

- $\forall P \in \mathfrak{P}((L_A(y)(P) = 1 \rightarrow L(x)(P) = 1) \wedge (L_A(y)(P) = 0 \rightarrow L_A(x)(P) = 0))$
- $\forall P \in \mathfrak{P}((L(x)(P) = 1 \rightarrow L_A(y)(P) \neq 0) \wedge (L(x)(P) = 0 \rightarrow L_A(y)(P) \neq 1))$

We will refer to these as modes 0-2. Indeed, mode 2 is the one introduced in our definition above. For our purposes in this paper, it will suffice.

6.3 Incorporating Belief Revision

We would like to say that this notion of an action is a protocol. However, the notion of protocol standard to this literature, given in [29], is insufficient for our purposes.

Definition 6.2 (Simplicial Protocol)**.** A **Simplicial Protocol** is a triple $(\mathscr{I}, \mathscr{P}, \Xi)$ where $\mathscr{I}$ is a simplicial model for knowledge, called the **"Input Complex"**, $\mathscr{P}$ is a simplicial model for knowledge called the **"Protocol Complex"** (better thought of in our context as the output complex) and $\Xi : \mathscr{I} \rightarrow 2^{\mathscr{P}}$ is called the **"Execution Map"**, which takes faces in $\mathscr{I}$ to sets of faces in $\mathscr{P}$ such that if $\sigma \subseteq \tau$, $\Xi(\sigma) \subseteq \Xi(\tau)$, and $\Xi(\sigma \cap \sigma') = \Xi(\sigma) \cap \Xi(\sigma')$. Moreover, for any face σ in $\mathscr{I}$, $\{a \in Ag \mid \exists x \in \sigma(V_{\mathscr{I}}(x) = a)\} = \{a \in Ag \mid \exists X \in \Xi(\sigma), x \in X(V_{\mathscr{P}}(x) = a)\}$. When this is true we say that Ξ is **Chromatic**. And, additionally, $\mathscr{P} = \bigcup_{\sigma \in \mathscr{I}} \Xi(\sigma)$.

In general, Ξ specifies which input perspectives go to which output perspectives. The restrictions on Ξ bear further analysis. Consider in particular the intersection property: $\Xi(\sigma \cap \sigma') = \Xi(\sigma) \cap \Xi(\sigma')$. The textbook justifies it by saying no agent "can "tell" whether the execution started with inputs from σ or from σ'". [29] Here, "tell" refers to the usual understanding of knowledge in simplicial complexes, where the accessibility of two facets is given by their nonempty intersection.[12] Suppose that a node x is in the image of σ and σ'. That is, $x \in \Xi(\sigma) \cap \Xi(\sigma')$. Then, it follows from our intersection rule that $x \in \Xi(\sigma \cap \sigma')$. By our intersection semantics, these are precisely the nodes which are in the image of those that cannot distinguish between σ and σ'. So, nodes which cannot tell whether they are in the image of σ or σ' are precisely those whose preimage cannot tell whether they are in σ or σ'.

Intersection preservation is meant to ensure that no agent "knows" whether they came from face X or Y. Since we have shifted over to belief, that is no longer available to us. We need the corresponding version of the intersection property which uses the definition of truth for modalities in the belief setting. In particular, we will need an execution map for each agent. But, it will only be relevant that the a-colored execution map preserve the intersections of a-colored vertices.[13] This motivates the following definition of a **Belief Protocol**:

Definition 6.3 (Belief Protocol)**.** A **Belief Protocol** is a triple $(\mathscr{I}, \mathscr{P}, \{\Xi_a\}_{a \in Ag})$ where $\mathscr{I}$ is a simplicial model for belief, called the **"Input Complex"**, $\mathscr{P}$ is a simplicial model for belief called the **"Protocol Complex"** (better thought of in our context as the output complex). Let $\mathscr{I}_a$ denote the a-complex for $\mathscr{I}$ and similarly for $\mathscr{P}_a$. For each $a \in Ag$ a function $\Xi_a : \mathscr{I} \rightarrow 2^{\mathscr{P}}$ is called the **"Execution Map"**, which takes faces in $\mathscr{I}_a$ to sets of faces in $\mathscr{P}_a$

[12]This is one place where the textbook gestures at intuitions afforded by the simplicial semantics.
[13]This is especially relevant in the revision setting.

such that if $\sigma \subseteq \tau$, $\Xi_a(\sigma) \subseteq \Xi_a(\tau)$, and $\Xi_a(\pi_a(\sigma \cap \sigma')) = \pi_a(\Xi_a(\sigma)) \cap \pi_a(\Xi_a(\sigma'))$. Moreover, for any face σ in $\mathscr{I}_a$, $\{a \in Ag \mid \exists x \in \sigma(V_{\mathscr{I}}(x) = a)\} = \{a \in Ag \mid \exists X \in \Xi_a(\sigma), x \in X(V_{\mathscr{P}}(x) = a)\}$. When this is true we say that Ξ_a is **Chromatic**. And, additionally, for each $a \in Ag$, $\mathscr{P}_a = \bigcup_{\sigma \in \mathscr{I}} \Xi_a(\sigma)$.

We now show the sense in which action models for belief are protocols. Let I be a simplicial model for belief whose a complex is denoted I_a, and A a simplicial action for belief. As before, intuitively, our protocol complex should be give by $I[A]$. What we need to define is for each $a \in Ag$ an execution map $\Xi_a : I_a \to 2^{I_a[A]}$. As it turns out, the correct map is given by the analogous map to the knowledge case:

$$\Xi_a(F) := \{X \in I_a[A] \mid \pi_0(X) \subseteq F\}$$

Where $X \in I_a[A]$ is shorthand for "X is a face in the simplicial complex of $I_a[A]$." This gives us the following:

Theorem 6.2. *Let I be a simplicial model for belief and A an action model. The tuple $(I, I[A], \{\Xi_a\}_{a \in Ag})$ is a Belief Protocol.*

Proof. See 10.3. □

Morally, this is true because of the following. Chromaticity follows roughly because $V[A]((x,y)) = V(x)$. This then entails the intersection property roughly as follows. If F and G share an a-node, call it x, Ξ_a will send that a-node to a collection of a-nodes in the $I[A]$. In particular, $\Xi_a(\{x\})$ is the set of all singletons $\{(x,y)\}$ in $I_a[A]$. Now consider $\Xi_a(F)$. By definition, this is the set of all faces X in $I_a[A]$ such that $\pi_0(X) \subseteq F$. Since $\pi_0(X) = \{z \mid (z,y) \in X\}$, if $X \in \Xi_a(F)$, then $\pi_a(\pi_0(X)) = \{x\}$. Hence, $\Xi_a(F)$ contains the set of all faces X in $I_a[A]_b$ whose a-node is a tuple (x,y). Hence, $\pi_a(\Xi_a(F))$ is the set of all singletons $\{(x,y)\}$ in $I_a[A]$. The same reasoning applies to G as to F. So, we can conclude that

$$\Xi_a(\pi_a(F \cap G)) = \pi_a(\Xi_a(F)) = \pi_a(\Xi_a(G))$$

as desired.

Now we need to modify this definition of an action to use the notion of revision above. The key idea is to modify the facets from $\mathscr{F}(S_a[A])$.

Definition 6.4 (Simplicial Action for Belief with Revision)**.** Given a simplicial model for belief $\mathscr{M} = \langle N, V, L, S \rangle$, And a simplicial action for belief $A = \langle N_A, V_A, L_A, \{S_{a,A}\}_{a \in Ag}, Post \rangle$, we define the update model **with revision**

$$\mathscr{M}[A]_{BR} := \langle N[A]_{BR}, V[A]_{BR}, L[A]_{BR}, \{S_a[A]_{BR}\}_{a \in Ag} \rangle^{14}$$

as follows:

[14] $N[A]$ is a subset of $N \times N_A$, $V[A]_{BR} : N[A]_{BR} \to Ag$, $L[A]_{BR} : N[A]_{BR} \to 3^{\mathfrak{P}}$, and $S_a[A]_{BR}$ is a UCF subcomplex of $\mathfrak{M}(N[A]_{BR}, V[A]_{BR}, L[A]_{BR})$, just as in a regular simplicial model.

$$
\begin{aligned}
N[A]_{BR} := \{(x,y) \in N \times N_A | V(x) = V_A(y)\} \\
\wedge \forall P \in \mathfrak{P}((L(x)(P) = 1 \to L_A(y)(P) \neq 0) \\
\wedge (L(x)(P) = 0 \to L_A(y)(P) \neq 1))\} \\
V[A]_{BR}((x,y)) := V(x) \\
L[A]_{BR}((x,y))(P) := \begin{cases} Post(y)(P) \text{ if } Post(y)(P) \neq 2 \\ L_A(y)(P) \text{ if } L(x)(P) = 2 \\ L(x) \text{ otherwise} \end{cases}
\end{aligned}
$$

Defining $S_a[A]_{BR}$ will take some more work. Fix $a \in Ag$ and let X be a UCF facet in $2^{N \times N_A}$ such that $\pi_0(X) \in \mathscr{F}(S_a)$ and $\pi_1(X) \in \mathscr{F}(S_{a,A})$. Call the set of such X $\mathscr{F}^a(N \times N_A)$. Let X be a UCF facet in $2^{N[A]_{BR}}$ such that X is consistent, and $\pi_1(X) \in \mathscr{F}(S_{a,A})$. Call the set of such X $\mathscr{F}_a(2^{N[A]_{BR}})$.

Then we can define a replacement function as follows for $X \in \mathscr{F}^a(N \times N_A)$:

$$
\begin{aligned}
R_a(X) := \{Y \in \mathscr{F}_a(2^{N[A]_{BR}}) | \pi_a(Y) = \pi_a(X) \\
\wedge (\forall Z \in \mathscr{F}_a(2^{N[A]_{BR}})(\pi_a(Z) = \pi_a(X) \to |Y \cap X| \geq |Z \cap X|))\}
\end{aligned}
$$

We say that $S_a[A]_{BR}$ is defined as the simplicial complex whose facets come from $\bigcup_{X \in \mathscr{F}^a(N \times N_A)} R_a(X)$.

Theorem 6.3. *Let $\mathscr{M}$ be a simplicial model for belief and A a simplicial action model for belief. Then the output model $\mathscr{M}[A]_{BR}$ also satisfies the UCF property.*

Proof. See 10.4. □

Now we need to show that simplicial actions for belief with revision are belief protocols. To do this, we need the following notion. Note that if $Y \in \mathscr{F}(S_a[A]_{BR})$, then $R_a^{-1}(Y)$ is the preimage. We need to define this preimage for all faces in $S_a[A]_{BR}$. Given a face $X \in S_a[A]_{BR}$, say that $R_a^{-1}(X)$ is the set of faces which subset of all facets

$$
\begin{aligned}
R_a^{-1}(X) := \{X' \mid \exists Y \in \mathscr{F}(S_a[A]_{BR})(X \subseteq Y \wedge \exists Y' \in R_a^{-1}(Y)(X' \subseteq Y')) \\
\wedge \forall a \in Ag(\exists x \in X(V[A]_{BR}(x) = a) \leftrightarrow \exists x' \in X'(V[A]_{BR}(x') = a))\}
\end{aligned}
$$

Theorem 6.4. *Let I be a simplicial model for belief and A an action model. Define $\Xi_{a,BR}$ similarly to Ξ_a:*

$$
\Xi_{a,BR}(F) := \{X \in I_a[A]_{BR} \mid \forall X' \in R_a^{-1}(X)(\pi_0(X') \subseteq F)\}
$$

Then, the tuple $(I, I[A]_{BR}, \{\Xi_{a,BR}\}_{a \in Ag})$ is a Belief Protocol.

Proof. See 10.5. □

Morally, this is true because the replacement function takes facets to sets of facets that, among other things, share a-colored nodes. Hence, the same moral argument we gave above still applies to show the intersection property.

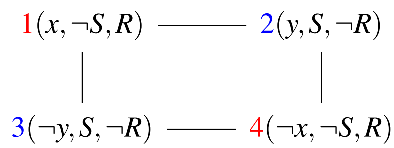

Figure 29.

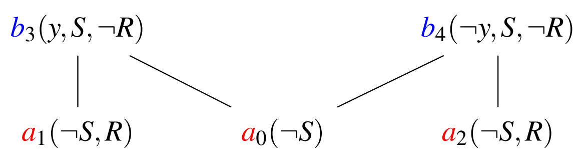

Figure 30.

7 Examples

As we have already mentioned, there is existing work which establishes the connection between action models and protocols. [26] [19] [3] [34] [39] [15] In particular, [19] even gives some examples. We will do much the same, using examples from the text "Distributed Computing through Combinatorial Topology" as our baseline. [29] These examples are nice because they are already given in the framework of simplicial complexes. All we need to do is prove that it is possible to represent these existing protocols using action models.

Consider the Alternating Message Passing Protocol on page 30 of [29]. In this protocol, there are two agents, call them a and b. As usual, these will be the red and blue agents respectively. The initial arrangement is that each agent has a personal bit-value, and they will be sharing that bit value with each other. We will call these bit values x for a's bit value and y for b's bit value. One agent is designated the first sender, and we shall pick b. The input complex will have four perspectives, two for each agent corresponding to the possible bit values, and total uncertainty. Since b is the first sender, the two b perspectives will be marked with the literals S and $\neg R$, and since a is the first received, they will be marked with the literals $\neg S$ and R. This gives us the input as drawn in Figure 29.

Now we must describe the protocol. The main idea of the alternating message protocol is that agents are sending messages back and forth, and there is always a chance a message does not go through. If an agent doesn't receive a message, they don't send. Let's imagine what happens when b sends the first message. There are two b perspectives, corresponding to where b sends y and where b sends $\neg y$. These should have preconditions S, $\neg R$, and y or $\neg y$ as appropriate. We will call these b_3 and b_4. By contrast, there are three a perspectives. These correspond with a learning y, a learning $\neg y$, and a learning nothing. We will call these a_1, a_2, and a_0 respectively. The preconditions on a_1 and a_2 are $\neg S$ and R, while the precondition on a_0 is merely $\neg S$. This last will be important for repeated applications of the protocol - any a perspective which is not a sender can receive nothing, possibly after having received nothing for many turns. This partially defined action model is drawn in Figure 30.

We now must draw the postconditions. Consider perspective b_3. Having just sent, they turn into a receiver. So the postconditions here should be $\neg S$ and R. Same for b_4. Similarly,

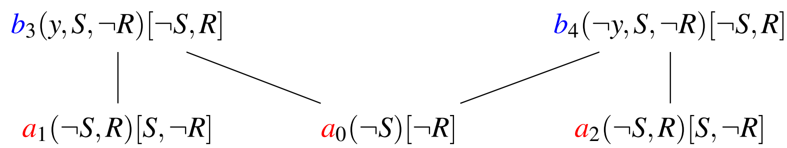

Figure 31.

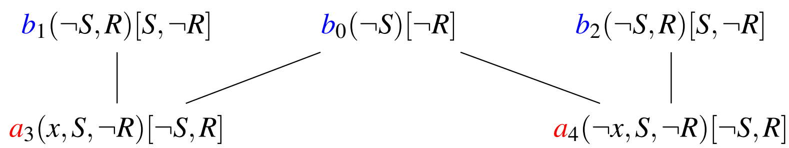

Figure 32.

receivers become senders, so the postconditions on a_1 and a_2 should be S and $\neg R$. Lastly, the postcondition on a_0 should be $\neg R$. If a receives nothing, they continue to not be a sender, but now are also not a receiver. This action model is draw in Figure 31.

The last step in the protocol is to create the mirror image of this for when a is the sender and attach it to the existing drawing. This is seen in Figure 32.

The last thing to acknowledge is that an edge must be drawn between a_0 and b_0 - faces between two edges where nothing is received continue to be faces after any step of the protocol, as these are the worlds where nothing happens. This gives us the complete action model, which is drawn in Figure 33.

Now let's compute the output of applying this action model to the specified input. The output vertices are $(1,a_0)$, $(1,a_1)$, $(1,a_2)$, $(2,b_3)$, $(3,b_4)$, $(4,a_0)$, $(4,a_1)$, $(4,a_2)$. This gives the output as drawn in Figure 34. Excitingly, this is precisely the output specified on page 30 of [29]. There are new perspectives created for when a learns the truth value of y, but also the original a perspectives are present, now corresponding to a failed signal. Let's apply the action again. The output vertices are $((1,a_0),a_0)$, $((4,a_0),a_0)$, $((2,b_3),b_0)$, $((2,b_3),b_1)$, $((2,b_3),b_2)$, $((3,b_4),b_0)$, $((3,b_4),b_1)$, $((3,b_4),b_2)$, $((4,a_1)a_4)$,

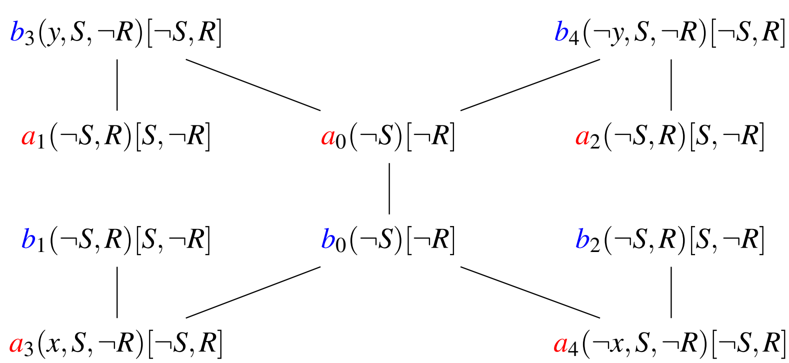

Figure 33.

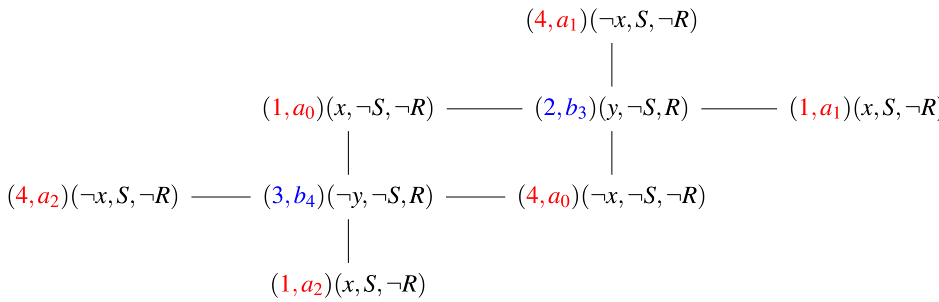

Figure 34.

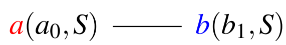

Figure 35.

$((1,a_1)a_3)$, $((4,a_2)a_4)$, $((1,a_2)a_3)$. This gives the correct result, as seen in Figure 34.

An important observation - when creating the action, we have to think through copies of perspectives. Even though creating a "base" copy of the 2 world in the input (i.e., pairing it with b_0) would seem to make sense, it doesn't. We need to transform this perspective from a sender to a non-sender. This is the role b_3 and b_4 play in the action. The base (a_0 and b_0) duplicate everything up to the sender perspective. Then, b_3 duplicates the sender, while a_1 duplicates the receiver (already duplicated by a_0, hence an additional copy).

Let us prove by induction that this example works no matter how many times we apply the protocol.

Proof. The above examples are our base case. Suppose Output(n) is a square of four nodes, with two strings off each blue node, alternating in color, of length n (so four strings total). Everything is $\neg S, \neg R$, except the last two nodes on each string, which are $\neg S, R$ just before the tip and $S, \neg R$ at the tip. We will show that applying the action increases the length of these strings by 1. Suppose, without loss of generality, that the tips of the strings are red. Then a_0 and b_0 copy everything up until the tip of each string. Each tip is copied by either a_3 or a_4, and the nodes just before the tip are additionally copied by b_1 and b_2. Because the tip of each node is either copied by a_3 or a_4, the additional copies from b_1 uniquely attach to their respective a_3 tipped string, and the b_2 copies their respective a_4 tipped string, as desired. □

We will now turn towards our next example, the "layered message passing protocol". Here, the two agents simultaneously write their bit value to a ledger at each time interval (so the protocol is synchronous). However, at each interval, there is a chance that at most one agent will fail. Upon observing a failure, the agents will stop. We will consider the situation where a has bit value 0 and b has bit value 1 first. Each of course starts as a sender, giving us the input as drawn in Figure 35.

The protocol itself is trickier. The action of course has four main perspectives. These

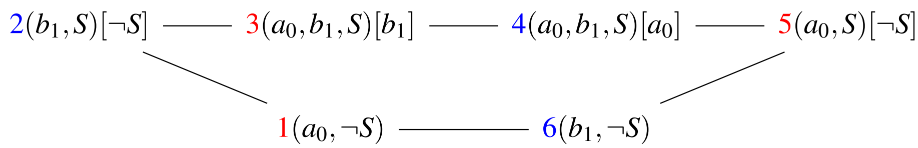

Figure 36.

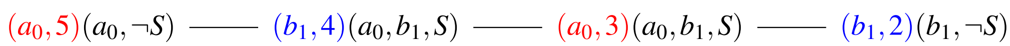

Figure 37.

being where a receives the message, a does not, and the same for b. But there are two other perspectives, one for each agent, which are less obvious because they are not relevant in the first step of the protocol. That is the perspective of a where at some point in the past, they failed to receive a message, and therefore on this particular step of the protocol did not send any message. All told, this gives the action model drawn in Figure 36. For example, perspective 2 is one where b sent, but failed to receive. 4 is where b sent and received, and 6 is where b did not send. Same for the a perspectives. The most nonideal thing about this setup is the need for the a_0 pre and post-condition in 4, and the similar setup in 3. These will be necessary to prevent unwanted duplications, corresponding to receiving a_0 and then a_1, in the setup where the bit values are not fixed in the input, as we shall see. But as presented, they do seem overly restricted. Other workarounds would be ideal.

If we apply the action model to the input a single time, we get the drawing in Figure 37, which is as given on page 31 of [29]. If we apply it again, we get the output as drawn in Figure 38, which is again as desired.

Now we want to apply this intuition to a more complicated setting, where the bit values are not predetermined in the input. This input is drawn in Figure 39.

Delightfully, the action model is easy to describe given the work we've already done. We simply need four copies of the action model previously, corresponding with the four combinations of pairs of bit values. We then stitch these together we get Figure 40.

Applying the action model once gives us the model in Figure 41, exactly as we would want per page 31 of [29]. Applying it a second time gives us the model in Figure 42. The values of variables are left off for ease of reading. It is here that it is important we include the non-ideal additional pre and post conditions. Otherwise, we could have perspectives

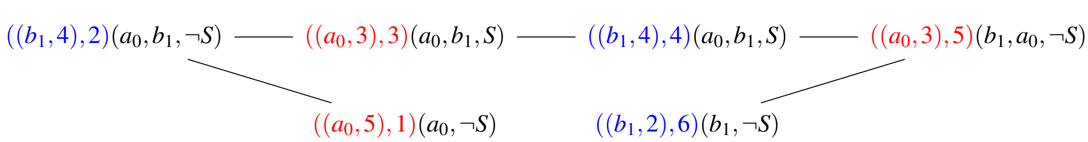

Figure 38.

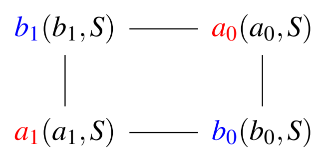

Figure 39.

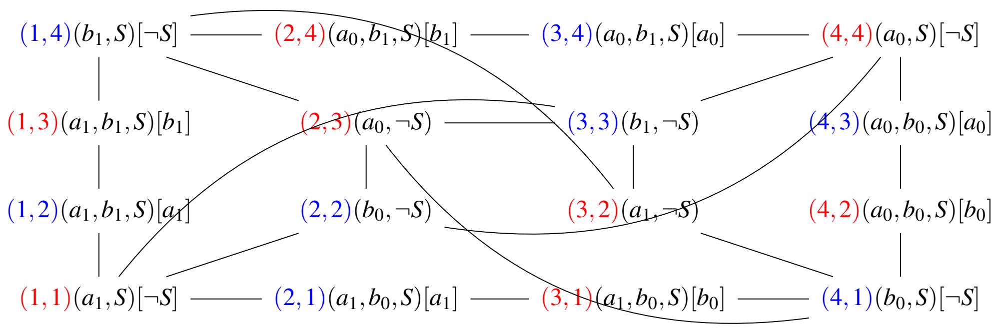

Figure 40.

in the above like $((a_0,(2,4),(4,2))$, which would be where a_0 first learns b_1 and then b_0. Nevertheless, we can show that the protocol applied n times is a square with side length $2n-1$:

Proof. The above examples are our base case. Suppose Output(n) is a square of side length $2(n+1)$. The center two nodes on each edge are labeled S, all other nodes $\neg S$, and every node but the corners contains a unique a_i,b_j pair. Applying the action duplicates every node but the centers of the edges of the input. These are duplicated by the center four nodes of the action, namely $(2,3)$, $(3,3)$, $(2,2)$, and $(3,2)$. The two center nodes of each edge are first copied by their respective a_i,b_j pair. These copies will be the center two nodes of the

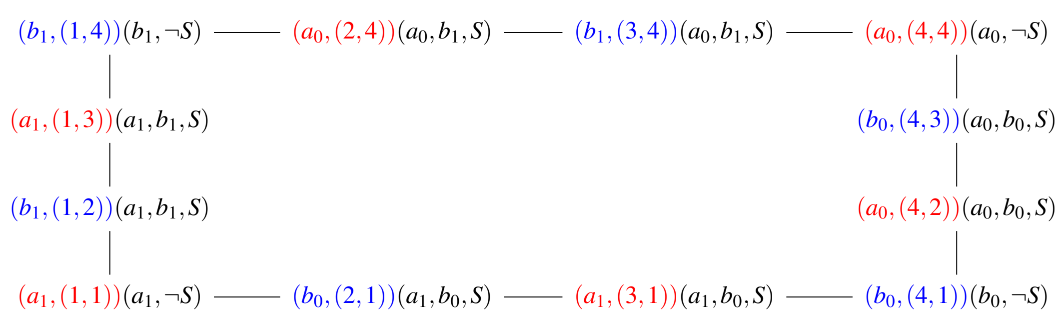

Figure 41.

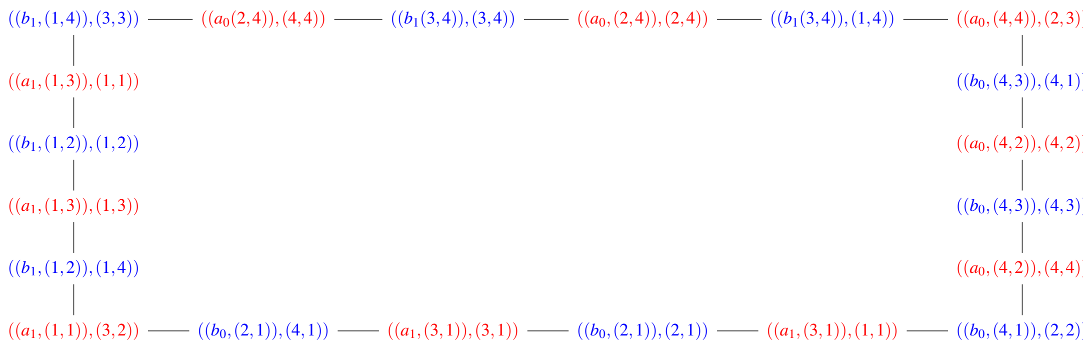

Figure 42.

string of four nodes that will replace each center pair. In this case, the a_0 b_1 edge are copied by $(2,4)$, $(3,4)$. Then the a is copied additionally by the a_i corner, and the b node is copied additionally by the b_j corner. These are the outer two nodes of that string of four nodes. No other copies are possible, as these nodes specify both the a_i and b_j value, and are S nodes. So, as desired, we have a single copy of every node except the center of each edge, which gets two copies each, as desired. □

8 Gossip Protocol

Recall that our desired setting of application was disaster relief. In this example, we are considering faults like the following. At time t, agent a announces to agent b that a certain road is clear. Call this fact P. Later, agent a gets a new vantage point over the geography of the region, and sends a new message to b. In fact, P is false, as they can see that a fallen tree is blocking it.

This example poses a few significant challenges. The first, and most obvious, is that agents have to learn information contrary to their beliefs. On the usual conception of action models (and therefore of protocols), after a announces to b that P, b believes that P is true, because the simplicial complex associates to b, namely S_b, will no longer have facets where P is true. This means that when a announces $\neg P$, S_b becomes empty, and b either believes everything or nothing, depending on how you interpret this. This is obviously undesirable, and inconsistent with what happens in practice in disaster relief scenarios. Luckily, belief revision is designed precisely to handle such problems. [31] [21]

Gossip protocols which handle faults like this or similar to this do already exist and are quite common. [4] [1] [32] [17] [9] [16] However, these tend to assume at least one of the following: there is some cap on the number or percentage of faulty agents, or there is some cap on the number of faulty signals. Neither of these, ultimately, seems tenable in the disaster relief example. Of course, at any time, any agent can send a faulty signal, either because they have bad information, and therefore bad beliefs, or because the message itself became garbled. Similarly, there is no cap on the number of faulty signals. There is no point in time at which we can imagine guarantees that signals will cease to be faulty.[15] As

[15]This second characterization may be worth weakening in the future, as it is fairly strong.

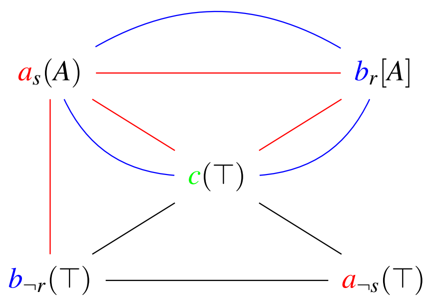

Figure 43.

such, the model will assume that faults are *transient* in the sense of [12] and faults "appear for a short time" (P is announced) and then "disappear" ($\neg P$ is announced). We assume additionally that faults are *persistent*; there is no time beyond which faulty messages cannot be sent.

So, in this section, we hope to define a gossip protocol which is descriptive of the disaster relief scenario. This protocol will be similar to the one given in [3], and the pseudocode for this protocol is given in Figure 1. In [3], the action model for a step of a protocol describes two agents sharing all of their knowledge with each other during a call. Other agents' degree of awareness is not assumed, and various possibilities are entertained. However, before we build a simplicial action model representing that protocol, we will start with a much simpler protocol. In this simpler protocol, as in the final gossip protocol, agents each have a secret, which we shall label I for each $i \in Ag$. The set of nodes N will be the set of $\{i_k | i \in Ag, k \in 2\}$, and the validation will be given by $V(I)(i_k) = k$ (otherwise, $V(I)(j_k) = 3$). The input complex will be the set of all UCF facets. More specifically, $S_i = \{X \subseteq N | h_k, j_l \in X \to h \neq j\}$. The main idea will be that the agents are trying to discover the true value of every other agent's private bit value. We will assume, without loss of generality, that the true bit value for every agent is 0. (Definition 3 of [3]) So, the goal of the task will be the complex $\{i_0 | i \in Ag\}$. This, along with the morphism mapping every facet in the input to the unique facet in the output, is very obviously a "task" as defined in [29].

The protocol itself will be very crude and simple. Unlike [3], we will assume agents only ever send some other agent *their* secret, not the sum total of secrets they have learned so far. Moreover, we will assume there is always a chance that this message fails to go through, and that agents other than the message and receiver are totally unaware any message is sent. We are already familiar with what this action model is where a sends b, as it is more or less drawn in Figure 26. It is redrawn as Figure 43 where the signal itself is labeled A. The pseudocode for this procedure is given in Figure 44.

At each time slice, a and b will be randomly selected, and the action model will be combined with the previous step's output. However, there is always a chance of course that $\neg A$ is sent instead. There are a few ways we could interpret this as an action model. However, we will use the action model given in Figure 45 to represent this. Baked into this representation is the idea that when a sends b the "faulty" signal, they are doing so intentionally, and the message, if it goes through, goes through clearly. Obviously other

Simple Protocol for Agent x

```
𝒜 ≝ {a, b, c ...}
ℳ ≝ {A, B, C, ...}
var V, rcvd : 𝒜 ⟶ (ℳ ⟶ 2)
var x, y ∈ 𝒜
var c ∈ ℳ

  skip
▯
  c := choose{0, 1}.
  ; if c = V(x)(X) then
    ; y := choose 𝒜 − x
    ; send⟨x, y, V⟩
▯
  receive⟨y, x, rcvd⟩
  ; V := rcvd
```

Figure 44.

kinds of fault are possible. *a* could "intend" to say that the road is clear, and *b* could mishear, for instance. But this way of modeling the fault most closely aligns with our initial example. So, with each step of the protocol, two agents are selected at random, call them *a* and *b*. Fix a probability $p > .5$. With probability p, the signal A is sent, and with probability $1-p$, the signal $\neg A$ is sent. This is our overly simplistic protocol. Again, the pseudocode is given in Figure 44.

However, this protocol is too simple to be called a gossip protocol in the usual sense. So, we now turn to modifying this simple protocol into a gossip protocol, and then describing that protocol as an action model. Recall the protocol we described in Section 2. Suppose *a* is randomly selected as the sender and *b* as the receiver. *a* then sends to *b* every single literal

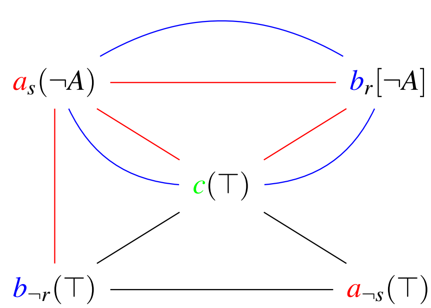

Figure 45.

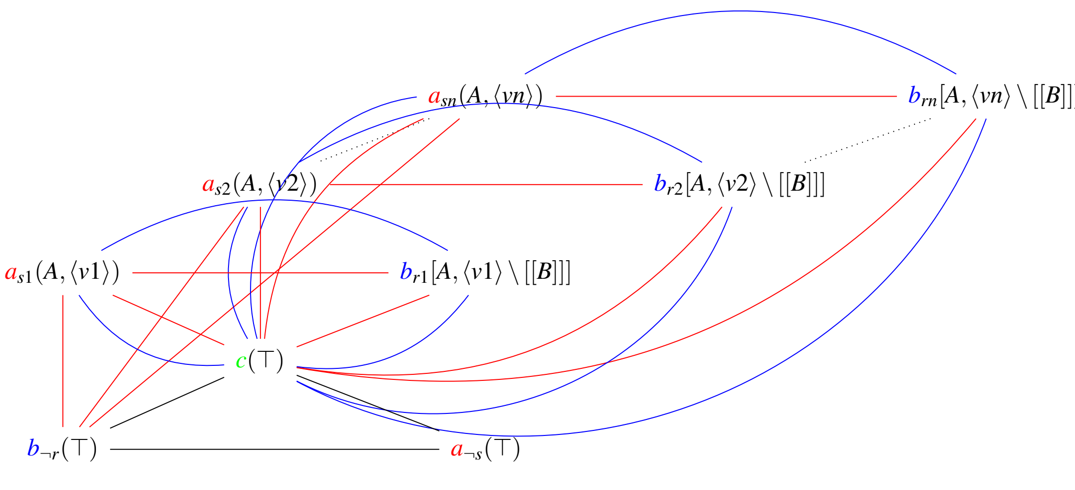

Figure 46.

they are aware of. At this step of execution, there are still only finitely many a-perspectives in the model. Each of these will have a string of literals associated with them. For each a-perspective, we imagine this string, in addition to either A or $\neg A$, chosen at random, is sent. In the event that b receives this packet, b updates on all of the information *except* the value of B, preserving this regardless of what was received. So, for example, if A sends the literals A, B, and $\neg C$, a b-perspective where B is true will update on the entire packet, while a b perspective where B is false updates merely on A and $\neg C$. The pseudocode describing this procedure is given in Figure 1.

Now we must turn this procedure into an action model. Intuitively, all we have to do to the above example is duplicate the a_s perspectives, depending on the packet of information that a perspective would send. In parallel, we duplicate the b_r perspectives similarly. Then, with some random chance, a chooses whether to share A or $\neg A$. As mentioned above, there are finitely many a-perspectives in the model at any stage of execution. Suppose without loss of generality that there are n-many. Let each $v1, \ldots, vn$ represent the tuple of literals true at each corresponding a perspective, sans either A or $\neg A$. Then $A, \langle v1 \rangle$ represents the set literals in $v1$ appended with A. Similarly, $\langle v1 \rangle[[B]]$ represents the set of literals in $v1$ sans either B or $\neg B$ if either appears in $v1$. So, the precondition for an a perspective where A is along with the rest of the packet is that that perspective honestly believes the packet, and hence is given by $A, \langle v1 \rangle$. The postcondition for b receiving a packet is to overwrite all literals except the value of B, and hence is given by $A, \langle v1 \rangle[[B]]$. This is drawn in Figure 46, for the case where A is sent. The action for when $\neg A$ is sent is obvious.

Figure 46 looks messy but it is simple to parse. In order to make it more readable, we have also drawn the protocol broken up into its constituent simplicial complexes S_i for all $i \in Ag$. S_a is given in Figure 47, S_b in Figure 48, and S_c in Figure 49. All we have done is duplicate the triangles $\{a_{s1}, b_{\neg r}, c\}$ and $\{a_{s1}, b_{r1}, c\}$ for each possible packet of information that a might send, and b may or may not receive. The preconditions on the a_{si} ensure

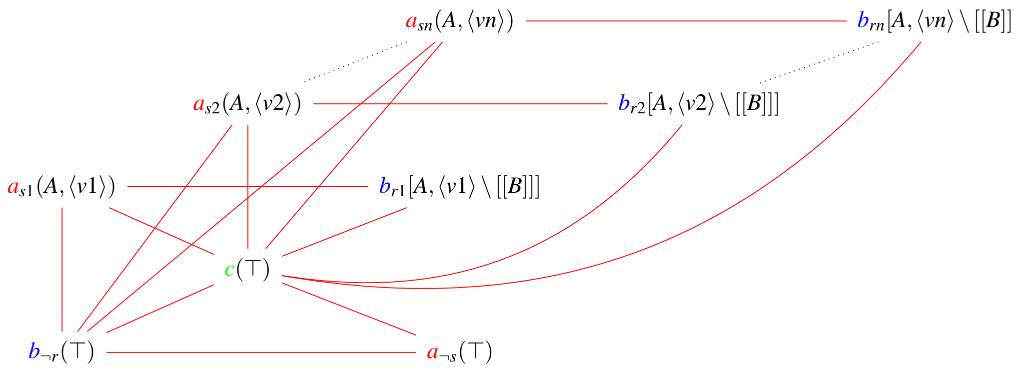

Figure 47.

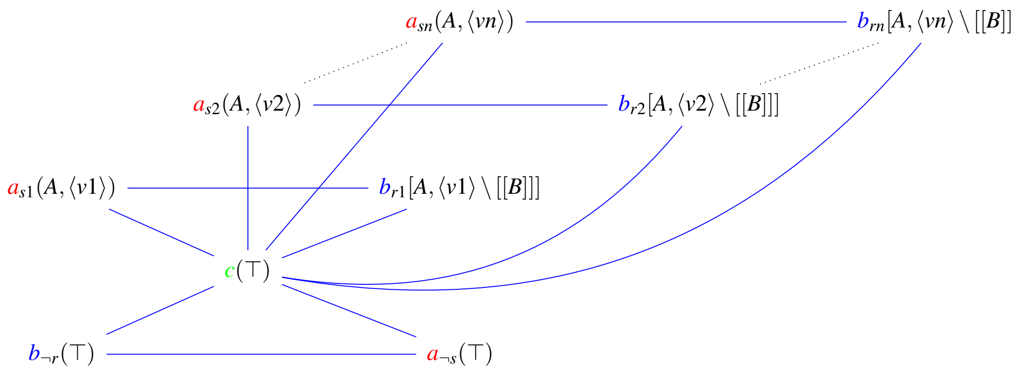

Figure 48.

that the only perspectives where the given vector of values is true will send. Similarly, the postconditions ensure b updates, here $\langle vi\rangle\setminus[[B]]$ is shorthand for the vector of literals vi, sans either B or $\neg B$ if it is present. That only means that b will not change her opinion on B if a shares it, but will still update their beliefs about what a believes about the truth of B.

One might be surprised, at this point, at the lack of fanfare. It seems that specifying the protocol was, at the end of the day, decently simple. Is this really meant to be the culmination of this entire paper? The contention of the authors is that the ease with which protocols can be articulated, as seen here, is in fact the *point* of expressing protocols using action models. Articulating protocols and modifying them is both simple and intuitive. Once the background work for defining actions using revision is given, it's also simple to incorporate this potential tool for fault tolerance into the protocols as well.

This ease of presentation will also let us make simple modifications to the protocol. Imagine instead that the receiver updates far less cautiously. That is, if b is the receiver, even the value of B may be overwritten if it is received in a packet. The pseudocode for this protocol is given in Figure 50. It is also simple to imagine how one changes the simplicial representation of such a protocol. One only has to update how the postcondition is evaluated. If $A,\langle v1\rangle$ is received, instead of updating with $A,\langle v1\rangle\setminus[[B]]$, instead b updates with merely $A,\langle v1\rangle$. This is drawn in Figures 51, 52, and 53.

By contrast, we can also imagine a much more cautious receiver. Perhaps instead,

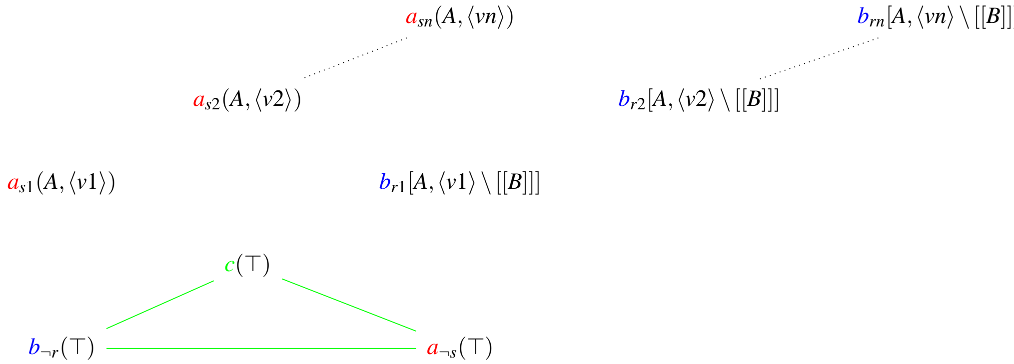

Figure 49.

Gossip Protocol 1 for Agent x

$$\mathscr{A} \stackrel{\text{def}}{=} \{\mathrm{a,b,c}\ldots\}$$
$$\mathscr{M} \stackrel{\text{def}}{=} \{\mathrm{A,B,C},\ldots\}$$
$$\texttt{var}\ \mathrm{V,rcvd} : \mathscr{A} \longrightarrow (\mathscr{M} \longrightarrow 2)$$
$$\texttt{var}\ \mathrm{x,y} \in \mathscr{A}$$
$$\texttt{var}\ \mathrm{c} \in \mathscr{M}$$

```
        skip
□
        c := choose{0,1}.
        ;if c = V(x)(X) then
          y := choose 𝒜 − x
          ;send⟨x,y,V⟩
□
        receive⟨y,x,rcvd⟩
        ;V := V ⊕ rcvd
where
        var U,W : 𝒜 ⟶ ℳ
        U ⊕ W ≝
         forall i ∈ 𝒜
          (U ⊕ W)(i) =
            if U(i) = W(i) = ∅ then ∅
            if U(i) ≠ ∅ ∧ W(i) = ∅ then U(i)
            if U(i) = ∅ ∧ W(i) ≠ ∅ then W(i)
            if U(i) ≠ ∅ ∧ W(i) ≠ ∅ then W(i)
```

Figure 50.

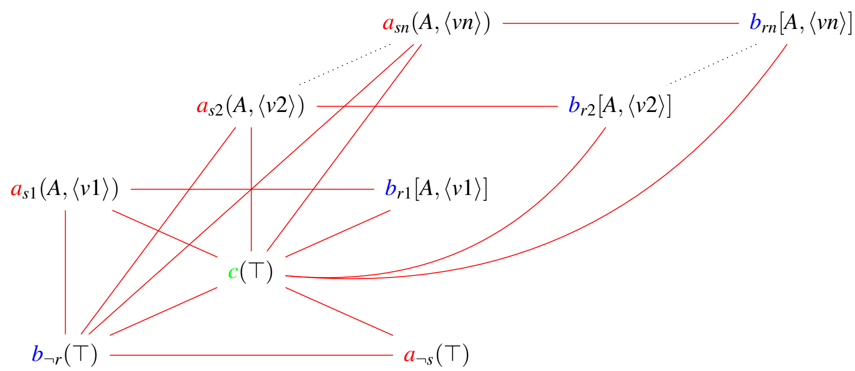

Figure 51.

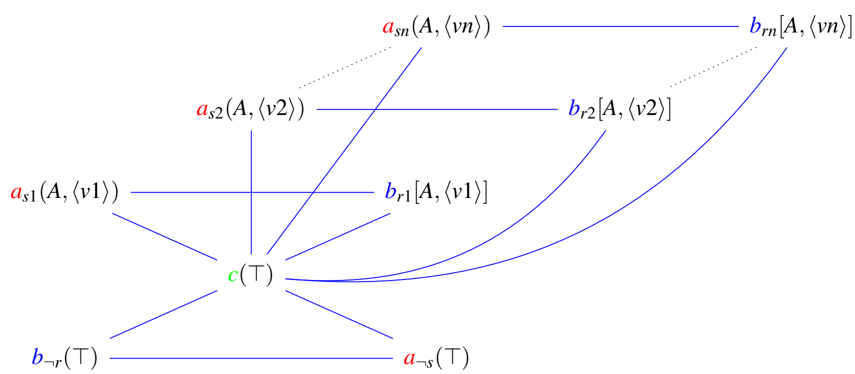

Figure 52.

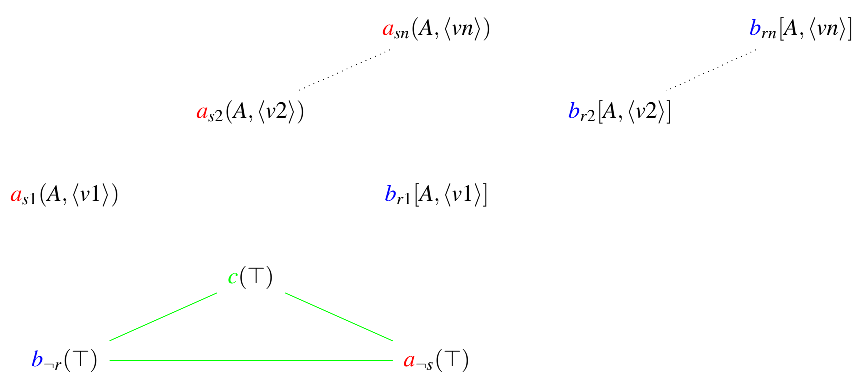

Figure 53.

whether or not they update on a received packet should be conditional on whether or not that packet agrees with their own perception of their bit value. In this case, updating is itself conditional. The receiver won't update in response to a packet of information if that packet contains a truth value for the receiver's literal contradictory to the one assigned to the perspective. The pseudocode for this is given in Figure 54.

Modifying our action model to represent this is straightforward. A precondition for b updating of course corresponds to a precondition. If the packet received is $A, \langle v1 \rangle$ the precondition that must be met is the value of B if it is present in $\langle v1 \rangle$. We denote this $\langle v1_B \rangle$. If this precondition is met, the postcondition is then the whole packet.[16] This is drawn in Figures 55, 56, and 57.

[16] If the precondition is met, it follows that updating by the whole packet won't change the value of B at the perspective.

Gossip Protocol 3 for Agent x

$\mathscr{A} \stackrel{\text{def}}{=} \{\mathrm{a,b,c}\ldots\}$
$\mathscr{M} \stackrel{\text{def}}{=} \{\mathrm{A,B,C,}\ldots\}$
$\texttt{var}\ \mathrm{V,rcvd} : \mathscr{A} \longrightarrow (\mathscr{M} \longrightarrow 2)$
$\texttt{var}\ \mathrm{x,y} \in \mathscr{A}$
$\texttt{var}\ \mathrm{c} \in \mathscr{M}$

$\texttt{skip}$

□

$\mathrm{c} := \texttt{choose}\{0,1\}.$
$;\texttt{if}\ \mathrm{c} = \mathrm{V(x)(X)}\ \texttt{then}$
$\quad \mathrm{y} := \texttt{choose}\ \mathscr{A} - \mathrm{x}$
$\quad ;\texttt{send}\langle \mathrm{x,y,V} \rangle$

□

$\texttt{receive}\langle \mathrm{y,x}, rcvd \rangle$
$;\mathrm{V} := \mathrm{V} \odot \mathrm{rcvd}$

$\texttt{where}$

$\texttt{var}\ \mathrm{U,W} : \mathscr{A} \longrightarrow \mathscr{M}$
$\mathrm{U} \odot \mathrm{W} \stackrel{\text{def}}{=}$
$\quad \texttt{forall}\ i \in \mathscr{A}$
$\quad (\mathrm{U} \odot \mathrm{W})(i) =$
$\qquad \texttt{if}\ U(x) = W(x)\ \texttt{then}$
$\qquad\quad \texttt{if}\ \mathrm{U}(i) = \mathrm{W}(i) = \emptyset\ \texttt{then}\ \emptyset$
$\qquad\quad \texttt{if}\ \mathrm{U}(i) \neq \emptyset \wedge \mathrm{W}(i) = \emptyset\ \texttt{then}\ \mathrm{U}(i)$
$\qquad\quad \texttt{if}\ \mathrm{U}(i) = \emptyset \wedge \mathrm{W}(i) \neq \emptyset\ \texttt{then}\ \mathrm{W}(i)$
$\qquad\quad \texttt{if}\ \mathrm{U}(i) \neq \emptyset \wedge \mathrm{W}(i) \neq \emptyset\ \texttt{then}\ \mathrm{W}(i)$
$\qquad \texttt{if}\ U(x) \neq W(x)\ \texttt{then}\ \mathrm{U}(i)$

Figure 54.

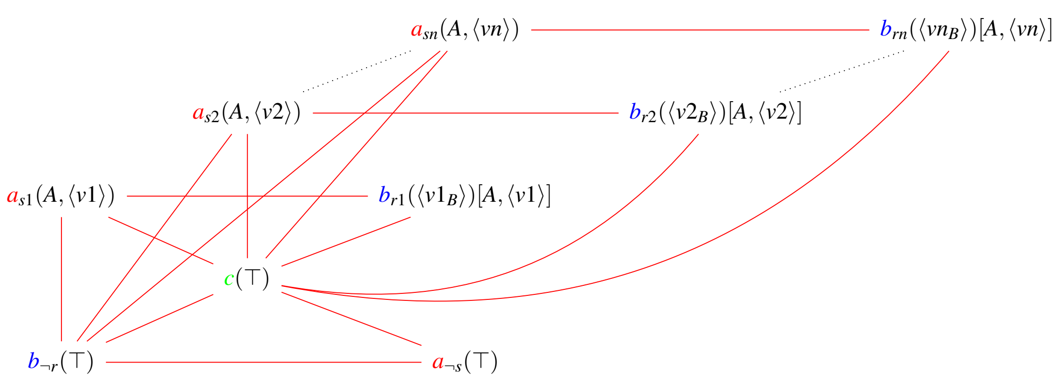

Figure 55.

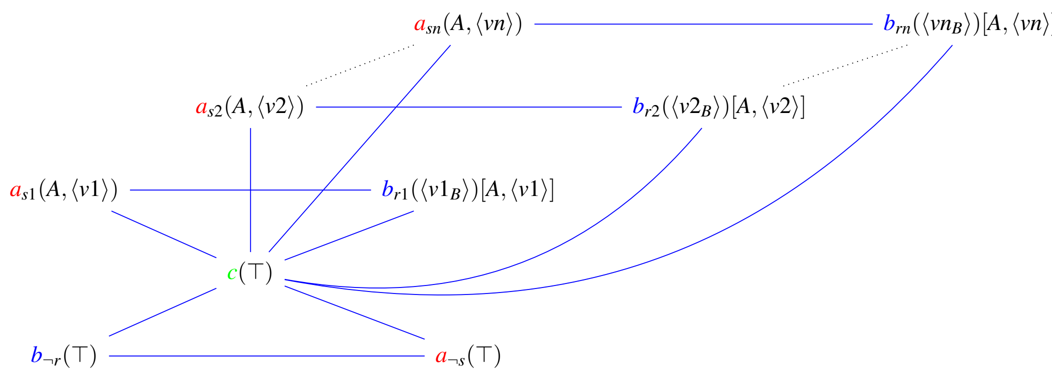

Figure 56.

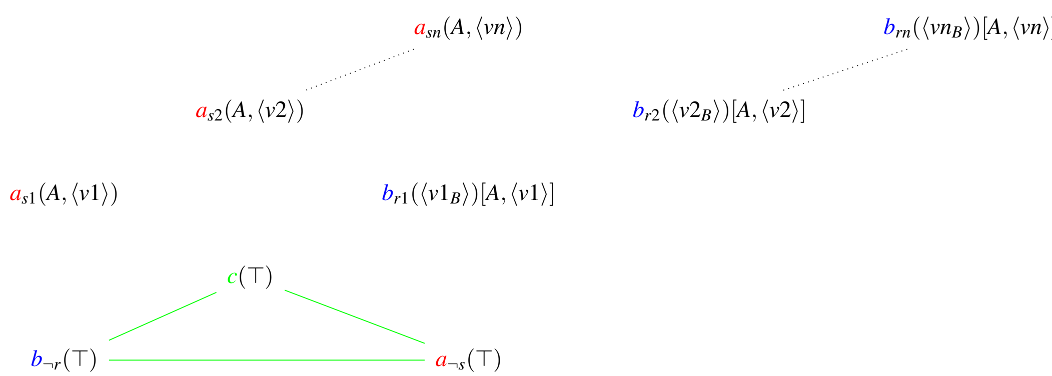

Figure 57.

9 Conclusion

In this paper we have given a novel semantics for epistemic modal logic in simplicial complexes. The goal of this semantics was to unify two ideas. The first of these ideas is already established, the second is novel to this memo. Firstly, simplicial complexes are an excellent and up-and-coming tool for modeling distributed protocols. [29] Secondly, we demonstrated here that simplicial complexes have a very natural tool built in for modeling a notion of "nearness" between worlds. This notion allows for the modeling of a distinct kind of learning, called "belief revision". Belief revision is particularly desirable when agents are likely to learn facts which they believe to be false, that is, when agents are likely to learn φ despite φ being false at every world/state the agent considers possible.

Taking these two ideas in hand, we were able to give a new definition of a protocol, in the spirit of the existing literature on the applications of simplicial complexes to distributed computing, where agents learn via "belief revision" as opposed to learning in a more conventional way. This is inspired by potential applications to disaster relief, where the possibility of false signals, or contradictory signals, is well documented. We showed that our new semantics can express existing protocols, and we defined a novel protocol meant to capture the disaster relief scenario, albeit in a crude sense.

This leads us to directions of future work. The first major direction for future work would be to implement action models in simplicial complexes in a simulation. In particular, a program which could compute action updates would be very helpful. Even small examples get large very quickly, and being able to work with more examples would help us to answer many questions, some of which we have already posed.

Firstly, does belief revision actually help with fault tolerance? If so, how? In what contexts? Answering this would require comparing and contrasting many protocols, ideally being able to compare protocols who only differ in whether or not they use the revision updating procedure. Getting clear on exactly how revision affects a protocol, especially after many stages, would go a long way towards addressing the question we started with: would belief revision help in a disaster relief scenario?

Secondly, what are the limitations of representing protocols using actions? A question this paper has not addressed is the actual, provable, scope of using action models to represent protocols. Are there protocols which cannot be represented using action models? If so, which ones? Why? A place to start would seem to be to take the definition of protocol from [29] and try to see if one can prove that all such protocols could be defined as actions.

Moreover, there are multiple nuances that can be added to the given gossip protocol. Restrictions on the number of contradictory signals as well as restrictions on who can send a contradictory signal are worth exploring. More significantly, the current protocol assumes agents only share their own private bit value with each signal. Expanding the model to share the totality of what they've learned (including, importantly, potentially false information) seems crucial.

10 Appendix

In this section we will formally prove a completeness theorem for our simplicial semantics. As is standard in the literature we will begin with a categorical equivalence. [19] [24]

Consider the following:

Definition 10.1 ($\mathfrak{S}$)**.** Let S_1 and S_2 be *UCF* simplicial models. We say that $f : S_1 \to S_2$ is a morphism iff, treating f as a function on vertices, X is a simplex in S_a iff $f[X]$ is a simplex in S_a, f preserves color, and for each vertex v, $L(v) = L(f(v))$. Note that the preservation of color and simplexes guarantees that f preserves maximality of simplexes. We say that the category whose objects are *UCF* simplicial models with these morphisms is $\mathfrak{S}$.

A quick definition:

Definition 10.2. A Kripke Model $K = \langle W, \{\mathscr{R}_a\}_{a\in A}, C\rangle$ is *proper* if and only if for any world $w \in W$ there is no world u such that $w\mathscr{R}_a u$ for all $a \in A$. We borrow this definition from [19].

And the following is the "natural" category of proper Kripke models.

Definition 10.3 ($\mathfrak{K}$)**.** Let K_1 and K_2 be transitive and Euclidean proper Kripke models such that for any world w and atom P true at that world, there exists an agent $a \in A$ such that if $w\mathscr{R}_a u$ and $w \in \ell(P)$, then $u \in \ell(P)$. It's easy to see that this makes **NU** sound. We call such Kripke models *Ideal*. We say that $g : K_1 \to K_2$ is a morphism iff, treating g as a function on worlds, $wR_{a,1}u$ iff $g(w)R_{a,2}g(u)$, and $w \in \ell_1(P)$ iff $g(w) \in \ell_2(P)$. We say that the category whose objects are ideal Kripke models with these morphisms is $\mathfrak{K}$.

One would like to say that $\mathfrak{S}$ and $\mathfrak{K}$ are equivalent categories. That is, one would like to say that there exists a pair of full and faithful functors from one category to the other which are inverses of each other. However, in this instance, this fails. The reason can be observed in the discussion of Figures 6 and 7. Intuitively, one would like to say that these two simplicial models are distinct - that is, they are not treated as equivalent in $\mathfrak{S}$. However, this cannot be the case if $\mathfrak{S}$ and $\mathfrak{K}$ are to be treated as equivalent, as the natural functor from $\mathfrak{S}$ to $\mathfrak{K}$ would map these to the exact same Kripke frame. The reason is that these models have corresponding facets which satisfy the same theories. To get around this, we will modify the category of Simplicial models.

Definition 10.4 ($\mathfrak{S}'$)**.** Let S_1 and S_2 be *UCF* simplicial models. We say that $f : S_1 \to S_2$ is a morphism iff, treating f as a function on $\mathscr{F}(W_1)$, we have that for any facets $X, Y \in W_1$ and $X \in S_{1,b}$, $Y \in S_{1,a}$ and $\pi_a(X) = \pi_a(Y)$ iff $f(Y) \in S_{2,a}$ and $\pi_a(f(X)) = \pi_a(f(Y))$, and furthermore, $\bigcup_{a\in A} L(\pi_a(X)) = \bigcup_{a\in A} L(\pi_a(f(X)))$. We say that the category whose objects are *UCF* simplicial models with these morphisms is $\mathfrak{S}'$.

Indeed, this modification works.

Theorem 10.1. *$\mathfrak{S}'$ and $\mathfrak{K}$ are equivalent categories.*

Proof. First we need to define two functors. For the first functor, define $F : \mathfrak{S}' \to \mathfrak{K}$ such that for a *UCF* model S, $F(S)$ is a Kripke model such that for each maximal facet X in some S_a, there is a unique world w_X in $F(S)$, and $w_X\mathscr{R}_a w_Y$ iff $\pi_a(X) = \pi_a(Y)$ and $Y \in \mathscr{F}(S_a)$. By design, worlds and facets are in bijection. Furthermore, $w_X \in \ell(P)$ iff $P \in \bigcup_{x\in X} L(x)$. We need to show that this is an ideal Kripke model. It's obvious that $\mathscr{R}_a$ is transitive and Euclidean for each $a \in A$. Now consider any atom P and world w_X. Suppose that $w_X \in \ell(P)$. It follows that $P \in \bigcup_{x\in X} L(x)$. Fix $x \in X$ such that $P \in L(x)$. Since S is *UCF*, fix a such that $V(x) = a$. Suppose further that $w_X \mathscr{R}_a w_Y$. Then $x = \pi_a(X) = \pi_a(Y)$, and so $P \in L(\pi_a(Y)) \subseteq \bigcup_{x \in Y} L(Y)$. It follows that $w_Y \in \ell(P)$, as desired. Finally, if $w_X \mathscr{R}_a w_Y$ for all $a \in A$, then $\pi_a(X) = \pi_a(Y)$ for all $a \in A$. This means that $X = Y$, and so $w_X = w_Y$, which tells us that $F(S)$ is a proper Kripke model. This shows that we have an ideal Kripke model.

Now we must show that F as given above is a functor. Let $f : S_1 \to S_2$ be a $\mathfrak{S}'$ morphism. Then $F(f) : F(S_1) \to F(S_2)$ is defined as follows: $F(f)(w_X) = w_{f(X)}$. We need to show that this is a $\mathfrak{K}$ morphism. Indeed, by design, $w_{f(X))}$ is a world in $F(S_2)$. Let R_1 be the relation in $F(S_1)$ and R_2 be the relation in $F(S_2)$. Suppose $w_X R_1^a w_Y$. This is true if and only if $\pi_a(X) = \pi_a(Y)$ and $Y \in \mathscr{F}(S_{1,a})$. Call this shared vertex v. This follows if and only if $f(v) \in f(X) \cap f(Y)$, and since $f(Y)$ is in $S_{2,a}$, this is true if and only if $w_{f(X)} R_2^a w_{f(Y)}$, as desired. Suppose $w_X \in \ell_1(P)$. This is true iff $P \in \bigcup_{x \in X} L(x)$. Because $\bigcup_{x \in X} L(x) = \bigcup_{x \in f(X)} L(x)$, this is true if and only if $P \in \bigcup_{x \in f(X)} L(x)$. This is true iff $w_{f(X)} \in \ell_2(P)$, as desired. This shows that F is well defined.

It's obvious that F preserves the identity. We must show that it preserves composition to establish its functoriality. Let $f_1 : S_1 \to S_2$ and $f_2 : S_2 \to S_3$ be morphisms. Then the following holds:

$$\begin{aligned} F(f_2 \circ f_1)(w_X) &= w_{f_2 \circ f_1(X)} \\ &= w_{f_2(f_1(X))} \\ &= F(f_2)(w_{f_1(X)}) \\ &= F(f_2)(F(f_1)(w_X)) \\ &= (F(f_2) \circ F(f_1))(w_X) \end{aligned}$$

From this it follows that $F(f_2 \circ f_1) = F(f_2) \circ F(f_1)$, as desired. So, we have given our first functor, and shown that it is a functor.

Now we must define the second functor. Define $G : \mathfrak{K} \to \mathfrak{S}'$ such that for an ideal model K, $G(K)$ is a simplicial model where the facets of W are of the form X_w for each world w in K, and the facets of each S_a are of the form X_w fr each world w in K such that $w \mathscr{R}_a w$. Because K is ideal, worlds and facets are in bijection. We assume each facet has a unique vertex for each agent. The only further stipulation is that $\pi_a(X_w) = \pi_a(X_u)$ and $X_u \in \mathscr{F}(S_a)$ if and only if $w \mathscr{R}_a u$.[17] For each atom P true at w, we say that $P \in L(\pi_a(X_w))$ iff P is true at all u such that $w \mathscr{R}_a u$. By the fact that K is ideal, we know such an a exists. It is easy to see that this is a *UCF* model. Indeed, the disjoint collection of facets X_w is vacuously *UCF*, and each association of two vertices preserves the *UCF* property.

Now we must show that G as given above is a functor. Let $g : K_1 \to K_2$ be a $\mathfrak{K}$ morphism. Then $G(g) : G(K_1) \to G(K_2)$ is defined as follows: $G(g)(X_w) = X_{g(w)}$. We need to show this is a $\mathfrak{S}'$ morphism. It suffices to show that G preserves facets and the logical information. Since g preserves reflexive edges, $G(g)$ therefore preserves facets in each $S_{1,a}$ in $G(K_1)$. Let X be a facet in W_1 in $G(K_1)$. Fix w in the worlds of K_1 such that $X = X_w$. Then $G(g)(X_w) = X_{g(w)}$ which is by definition a facet in $G(K_2)$, since $g(w)$ is a world in K_2. Now suppose $P \in L(\pi_a(X_w))$. This is true iff $u \in \ell(P)$ for all u such that $w \mathscr{R}_a u$. This is true iff $g(u) \in \ell(P)$ for all u such that $w \mathscr{R}_a g(u)$. This is true iff $P \in L(\pi_a(g(w)))$. This is true iff $P \in L(G(g)(\pi_a(X_w)))$, as desired. This shows that G is well defined.

[17] By the Euclidean axiom, we have that, if $w \mathscr{R}_a u$, then $u \mathscr{R}_a u$, so $X_u \in \mathscr{F}(S_a)$ is actually redundant as a condition.

It's obvious that G preserves the identity. We must show that it preserves composition to establish its functoriality. Let $g_1 : K_1 \to K_2$ and $g_2 : K_2 \to K_3$ be morphisms. Then the following holds:

$$\begin{aligned} G(g_2 \circ g_1)(X_w) &= X_{g_2 \circ g_1(w)} \\ &= X_{g_2(g_1(w))} \\ &= G(g_2)(X_{g_1(w)}) \\ &= G(g_2)(G(g_1)(X_w)) \\ &= (G(g_2) \circ G(g_1))\,(X_w) \end{aligned}$$

From this it follows that $G(g_2 \circ g_1) = G(g_2) \circ G(g_1)$, as desired.

What remains to be shown is that for any Kripke model K, $FG(K)$ is isomorphic to K and for any Simplicial model S, $GF(S)$ is isomorphic to S. For the first case, note that the worlds in $FG(K)$ are of the form w_{X_w} where w is a world in K. The above demonstrated this identification is a bijection. Moreover, for any worlds w and u in K, $w\mathscr{R}_a u$ in K iff X_u is in S_a and $\pi_a(X_w) = \pi_a(X_u)$. This is true iff $w_{X_w}\mathscr{R}_a u_{X_u}$, as desired. Moreover, $w \in \ell(P)$, iff there is some $a \in Ag$ such that $P \in L(\pi_a(X_w))$, which is true iff $w_{X_w} \in \ell(P)$. This shows that K and $FG(K)$ are isomorphic.

Suppose X and Y are facets in W, $\pi_a(X) = \pi_a(Y)$ and $Y \in S_a$. This is true iff $w_X\mathscr{R}_a w_Y$, which, because $Y \in S_a$, and so $w_Y\mathscr{R}_a w_Y$, is true iff $\pi_a(X_{w_X}) = \pi_a(Y_{w_Y})$ and $Y_{w_Y} \in \mathscr{F}(S_a)$ for $GF(S)$. Suppose instead that $\bigcup_{a\in A} L(\pi_a(X)) = \bigcup_{a\in A} L(\pi_a(Y))$ This is true iff $w_X, w_Y \in \ell(P)$. Because $F(S)$ is ideal, for each P, there is an agent $a \in Ag$ such that if $w \in \ell(P)$, then for all u such that $w\mathscr{R}_a u$, $u \in \ell(P)$. By definition of G, this means that $\bigcup_{a\in A} L(\pi_a(X_{w_X})) = \bigcup_{a\in A} L(\pi_a(Y_{w_Y}))$. This shows that S and $GF(S)$ are isomorphic. □

Theorem 10.2. *The class of UCF simplicial models is sound and complete with respect to propositional logic plus* **K45** + **NU***, modus ponens and necessitation for each belief modality in the language without public announcement.*

Proof. First, we need to show that F and G are logic preserving functors. This is a simple structural induction.

Suppose P is an atomic formula, K is a Kripke model, and S is a simplicial model. Suppose further that $K, w \vDash P$ and $S, X \vDash P$. Then, $w \in \ell(P)$ and there is $a \in Ag$ such that $P \in L(\pi_a(X))$. So, $w_X \in \ell(P)$, and $P \in L(\pi_a(X_w))$, showing that F and G preserve atomic formulas.

Assume the obvious inductive hypothesis for formulas up to depth n. The only interesting case is the modal one. Let φ be a formula of depth n. Suppose that $K, w \vDash B_a\varphi$ and $S, X \vDash B_a\varphi$. We will handle the Kripke case first. This means that, for all u such that $w\mathscr{R}_a u$, $K, u \vDash \varphi$. By the IH, and the definition of G, for all facets $Y \in S_a$ such that $\pi_a(Y) = \pi_a(X_u)$, we have that $G(K), X_u \vDash \varphi$. The result follows. Because $S, X \vDash B_a\varphi$, then for all $Y \in S_a$ such that $\pi_a(X) = \pi_a(Y)$, then $S, Y \vDash \varphi$. By the IH, and the definition of F, $F(S), w_Y \vDash \varphi$. The result follows.

The next step is to construct a canonical **K45+NU** frame model, using the usual "unboxing" method. That is, construct a model K_C whose set of worlds W_C is the set of all maximal consistent **K45+NU** sets of formulas, $\ell(P) = \{w \in W_C | P \in w\}$, and $w\mathscr{R}_a u$ if and only if all formulas φ, if $B_a\varphi \in w$, then $\varphi \in u$. It is easy to show that this model is such that $K_C, w \vDash \varphi$ if and only if $\varphi \in w$. In general, this model is not proper, but we can apply the translation from [11] to generate a bisimilar model which is. Then, applying G to this proper model gives us a canonical simplicial model. Completeness follows. □

Theorem 10.3. *Let I be a simplicial model for belief and A an action model. The tuple $(I, I[A], \{\Xi_a\}_{a\in Ag})$ is a Belief Protocol.*

Proof. Suppose F and G are faces of S_a and $F \subseteq G$. Then suppose $X \in \Xi_a(F)$. Then $\pi_0(X) \subseteq F$. Then $\pi_0(X) \subseteq G$, so $X \in \Xi_a(G)$, as desired.

In order to show that each Ξ_a is chromatic we will first establish the result for singletons. Let $n \in I$ be a node and fix $a \in Ag$ such that $V(n) = a$. Then in particular, if $X \in \Xi_a(\{n\})$, then $\pi_0(X) = n$. It follows from the UCF property that $|X| = 1$, and fixing $(n,y) \in X$, $V[A]((n,y)) = V(n) = a$.

Now we generalize this to larger faces. For each $n \in F$ where F is a face in I, we know that $V[A](\Xi_a(\{n\})) = V(n)$. Suppose $X \in \Xi_a(F)$. Then, for each $a \in Ag$ such that there's an $(x,y) \in X$ such that $V[A]((x,y)) = V(x) = a$, it because $\pi_0(X) \subseteq F$, $x \in F$ and thus there is an a-perspective in F. This suffices to show chromaticity.

Let F and G be two faces of S_a. To verify the intersection property, we can see that the only interesting case is when $\pi_a(F \cap G) \neq \emptyset$. If $\pi_a(F \cap G) = \emptyset$. Then, $\Xi_a(\pi_a(F \cap G)) = \pi_a(\Xi_a(F)) \cap \pi_a(\Xi_a(G)) = \emptyset$. Suppose $X \in \Xi_a(\pi_a(F \cap G))$. Then $\pi_0(X) \subseteq \pi_a(F \cap G)$. Note that this implies that $\pi_0(X)$ is a singleton containing the unique a-colored node in both F and G. So, $\pi_0(X) \subseteq F$ and $\pi_0(X) \subseteq G$, and we get that $X \in \pi_a(\Xi_a(F)) \cap \pi_a(\Xi_a(G))$, as desired. Suppose now that $X \in \pi_a(\Xi_a(F)) \cap \pi_a(\Xi_a(G))$. Then $X \in \pi_a(\Xi_a(F))$. So, by the UCF property, X is an a-colored singleton in $I_a[A]$. By the definition of Ξ_a, there is X' such that $X \subseteq X'$ and $\pi_0(X') \subseteq F$. Since X contains the unique a-colored node of X', we can conclude that $\pi_0(X) \subseteq F$ is the singleton set containing the unique a-colored node of F. The same argument will show that $\pi_0(X)$ is the singleton set containing the unique a-colored node of G. So, $\pi_0(X)$ is the singleton set containing the unique a-colored node of $F \cap G$. Hence, $\pi_0(X) = \pi_a(F \cap G)$. By definition, then, $X \in \Xi_a(\pi_a(F \cap G))$. This suffices to show the intersection property.

By construction, $\bigcup_{\sigma\in\mathscr{I}} \Xi_a(\sigma) \subseteq I[A]$. Suppose that X is a face in $I[A]$. Take $\sigma = \pi_0(X)$. By the definition of $I[A]$, σ is a face in I. Fix Y any facet containing X. This suffices to show $\bigcup_{\sigma\in\mathscr{I}} \Xi_a(\sigma) = I[A]$. □

Theorem 10.4. *Let $\mathscr{M}$ be a simplicial model for belief and A a simplicial action model for belief. Then the output model $\mathscr{M}[A]_{BR}$ also satisfies the UCF property.*

Proof. This follows immediately from the fact that, by assumption, each facet in $S_a[A]_{BR}$ is given by $\mathscr{R}_a(X)$ for some $X \in \mathscr{F}^a(\mathscr{M}[A])$. Because $\mathscr{R}_a(X) \subseteq \mathscr{F}_a(2^{N[A]_{BR}})$ for all $X \in \mathscr{F}^a(\mathscr{M}[A])$, and the facets in $\mathscr{F}_a(2^{N[A]_{BR}})$ are assumed to be UCF, the result follows. □

Theorem 10.5. *Let I be a simplicial model for belief and A an action model. Define $\Xi_{a,BR}$ similarly to Ξ_a:*

$$\Xi_{a,BR}(F) := \{X \in I_a[A]_{BR} \mid \forall X' \in R_a^{-1}(X)(\pi_0(X') \subseteq F)\}$$

Then, the tuple $(I, I[A]_{BR}, \{\Xi_{a,BR}\}_{a\in Ag})$ is a Belief Protocol.

Proof. Suppose F and G are faces of S_a and $F \subseteq G$. Then suppose $X \in \Xi_{a,BR}(F)$. Then for all $X' \in R_a^{-1}(X)$ $\pi_0(R_a^{-1}(X')) \subseteq F$. Then $\pi_0(R_a^{-1}(X')) \subseteq G$, so $X \in \Xi_{a,BR}(G)$, as desired.

In order to show that each $\Xi_{a,BR}$ is chromatic we will first establish the result for singletons. Let $n \in I$ be a node and fix $a \in Ag$ such that $V(n) = a$. Then in particular, if $X \in \Xi_{a,BR}(\{n\})$, then for all $X' \in R_a^{-1}(X)$, $\pi_0(X') = n$. It follows from the UCF property and the fact that R_a^{-1} preserves colors that $|X| = |X'| = 1$. Fix y such that $(n,y) \in X'$ and x and z such that $(x,z) \in X$. By the definition of R_a^{-1}, $V[A]_{BR}((x,z)) = V[A]_{BR}((n,y)) = V(n) = a$.

Now we generalize this to larger faces. For each $n \in F$ where F is a face in I, we know that $V[A]_{BR}(\Xi_{a,BR}(\{n\})) = V(n)$. Suppose $X \in \Xi_{a,BR}(F)$. Consider each $a \in Ag$ such that there's an $(x,y) \in X$ such that $V[A]_{BR}((x,y)) = V(x) = a$. For all $X' \in R_a^{-1}(X)$, we have that $\pi_0(X') \subseteq F$. Fix $(z,w) \in X'$. Then $z \in F$ and because $a = V(x) = V[A]_{BR}(x,y) = V[A]_{BR}(z,w) = V(z)$, there is an a-perspective in F. This suffices to show chromaticity.

Let F and G be two faces of S_a. To verify the intersection property, we can see that the only interesting case is when $\pi_a(F \cap G) \neq \emptyset$. If $\pi_a(F \cap G) = \emptyset$. Then, by the singleton case of chromaticity, it follows that $\Xi_{a,BR}(\pi_a(F \cap G)) = \pi_a(\Xi_{a,BR}(F)) \cap \pi_a(\Xi_{a,BR}(G)) = \emptyset$.

It will be useful for the reasoning below to realize the following. If X is a singleton, a-colored node, then $R_a^{-1}(X) = X$. Suppose $X' \in R_a^{-1}(X)$. Then by definition there is $Y \in \mathscr{F}(S_a[A]_{BR})$ such that $X \subseteq Y$ and there exists $Y' \in R_a^{-1}(Y)$ where $X' \subseteq Y'$. By the definition of R_a, $\pi_a(Y) = \pi_a(Y')$. So, by UCF, and the fact that R_a^{-1} preserves colors, $X' = \pi_a(Y')$ and $X = \pi_a(X')$. = Suppose $X \in \Xi_{a,BR}(\pi_a(F \cap G))$. Then for all $X' \in R_a^{-1}(X)$, $\pi_0(X') \subseteq \pi_a(F \cap G)$. Note that this implies that $\pi_0(X')$ is a singleton containing the unique a-colored node in both F and G. So, $\pi_0(X') \subseteq F$ and $\pi_0(X') \subseteq G$, and we get that $X \in \pi_a(\Xi_{a,BR}(F)) \cap \pi_a(\Xi_{a,BR}(G))$, as desired. Suppose now that $X \in \pi_a(\Xi_{a,BR}(F)) \cap \pi_a(\Xi_a(G))$. Then $X \in \pi_a(\Xi_{a,BR}(F))$. So, by the UCF property, X is an a-colored singleton. By the fact that R_a^{-1} preserves color, for all $X' \in R_a^{-1}(X)$, X' is an a-colored singleton. By the definition of $\Xi_{a,BR}$, $\pi_0(X') \subseteq F$. Since X' contains the unique a-colored node of X, we can conclude that $\pi_0(X') \subseteq F$ is the singleton set containing the unique a-colored node of F. The same argument will show that $\pi_0(X')$ is the singleton set containing the unique a-colored node of G. So, $\pi_0(X')$ is the singleton set containing the unique a-colored node of $F \cap G$. Hence, $\pi_0(X') = \pi_a(F \cap G)$. By definition, then, $X \in \Xi_{a,BR}(\pi_a(F \cap G))$. This suffices to show the intersection property.

By construction, $\bigcup_{\sigma \in \mathscr{I}_{\dashv}} \Xi_{a,BR}(\sigma) \subseteq I_a[A]_{BR}$. Suppose that X is a face in $I[A]_{BR}$. We first need to show that $R_a^{-1}(X)$ is nonempty. Fix $Y \in \mathscr{F}(S_a[A]_{BR})$ such that $X \subseteq Y$. Then by definition, there is $Y' \in R_a^{-1}(Y)$. Fix X' to be the subset of Y' which shares the same colors as X. Then $X' \in R_a^{-1}(X)$, as desired. Because $Y' \in \mathscr{F}^a(N \times N_A)$, $\pi_0(Y') \in \mathscr{F}(I_a)$. So, because $X' \subseteq Y'$, $\pi_0(X') \in I_a$. It follows that $X \in \Xi_{a,BR}(\pi_0(X'))$. This suffices to show $\bigcup_{\sigma \in \mathscr{I}} \Xi_{a,BR}(\sigma) = I[A]_{BR}$. □

Bibliography

References

1. L. Alvisi, J. Doumen, R. Guerraoui, B. Koldehofe, H. Li, R. van Renesse, and G. Tredan. How robust are gossip-based communication protocols? *SIGOPS Oper. Syst. Rev.*, 41(5):14–18, Oct. 2007.

2. K. R. Apt, D. Grossi, and W. van der Hoek. Epistemic protocols for distributed gossiping. *Electronic Proceedings in Theoretical Computer Science*, 215:51–66, June 2016.

3. M. Attamah, H. Ditmarsch, D. Grossi, and W. Hoek. Knowledge and gossip. *Frontiers in Artificial Intelligence and Applications*, 263:21–26, 01 2014.

4. R. Bakhshi, F. Bonnet, W. Fokkink, and B. Haverkort. Formal analysis techniques for gossiping protocols. *SIGOPS Oper. Syst. Rev.*, 41(5):28–36, Oct. 2007.

5. A. Baltag. A logic for suspicious players: Epistemic actions and belief–updates in games. *Bulletin of Economic Research*, 54(1):1–45, 2002.

6. A. Baltag, L. S. Moss, and S. Solecki. *The Logic of Public Announcements, Common Knowledge, and Private Suspicions*, pages 773–812. Springer International Publishing, Cham, 2016.

7. A. Baltag and S. Smets. Conditional doxastic models: A qualitative approach to dynamic belief revision. *Electronic Notes in Theoretical Computer Science*, 165:5–21, 2006. Proceedings of the 13th Workshop on Logic, Language, Information and Computation (WoLLIC 2006).

8. A. Baltag and S. Smets. The logic of conditional doxastic actions: A theory of dynamic multi-agent belief revision. 01 2008.

9. Y. Belal, M. Maouche, S. B. Mokhtar, and A. Simonet-Boulogne. Granite : a byzantine-resilient dynamic gossip learning framework, 2025.

10. K. Birman. The promise, and limitations, of gossip protocols. *SIGOPS Oper. Syst. Rev.*, 41(5):8–13, Oct. 2007.

11. A. Bjorndahl and P. Sink. A note on proper relational structures, 2025.

12. R. Butler. A criteria standard for conflict resolution: A vision for guaranteeing the safety of self-separation in NextGen. Technical Memorandum NASA/TM-2010-215108, NASA, Langley Research Center, Hampton VA 23681-2199, USA, February 2008.

13. C. Cachin, D. Lehnherr, and T. Studer. Simplicial belief. In U. Schmid and R. Kuznets, editors, *Structural Information and Communication Complexity*, pages 176–193, Cham, 2025. Springer Nature Switzerland.

14. A. Cassandra. Apache cassandra. *Website. Available online at http://aws.amazon.com/keyspaces/what-is-cassandra/*, 13:11, 2014.

15. A. Castañeda, H. van Ditmarsch, D. A. Rosenblueth, and D. A. Velázquez. Comparing the update expressivity of communication patterns and action models. *Electronic Proceedings in Theoretical Computer Science*, 379:157–172, July 2023.

16. R. Cohen, J. Loss, and T. Moran. Efficient agreement over byzantine gossip. Cryptology ePrint Archive, Paper 2023/1507, 2023.

17. S. Coretti, A. Kiayias, C. Moore, and A. Russell. The generals' scuttlebutt: Byzantine-resilient gossip protocols. Cryptology ePrint Archive, Paper 2022/541, 2022.

18. A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing*, PODC '87, page 1–12, New York, NY, USA, 1987. Association for Computing Machinery.

19. H. v. Ditmarsch, É. Goubault, J. Ledent, and S. Rajsbaum. Knowledge and simplicial complexes. In B. Lundgren and N. A. Nuñez Hernández, editors, *Philosophy of Computing*, pages 1–50, Cham, 2022. Springer International Publishing.

20. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, Cambridge, MA, USA, 2003.

21. P. Gardenfors. Imaging and conditionalization. *The Journal of Philosophy*, 79(12):747–760, 1982.

22. E. Goubault, R. Kniazev, and J. Ledent. A many-sorted epistemic logic for chromatic hypergraphs, 2023.

23. E. Goubault, R. Kniazev, J. Ledent, and S. Rajsbaum. Simplicial models for the epistemic logic of faulty agents, 2023.

24. É . Goubault, J. Ledent, and S. Rajsbaum. A simplicial complex model for dynamic epistemic logic to study distributed task computability. *Electronic Proceedings in Theoretical Computer Science*, 277:73–87, sep 2018.

25. É. Goubault, J. Ledent, and S. Rajsbaum. A simplicial model for kb4n : epistemic logic with agents that may die. In P. Berenbrink and B. Monmege, editors, *39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022)*, volume 219 of *Leibniz International Proceedings in Informatics, LIPIcs*, pages 33:1—33:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, FRA, March 2022.

26. E. Goubault and S. Rajsbaum. A simplicial complex model of dynamic epistemic logic for fault-tolerant distributed computing. *CoRR*, abs/1703.11005, 2017.

27. A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(2):157–170, 1988.

28. L. Hella, M. Järvisalo, A. Kuusisto, J. Laurinharju, T. Lempiäinen, K. Luosto, J. Suomela, and J. Virtema. Weak models of distributed computing, with connections to modal logic. In *Proceedings of the 2012 ACM Symposium on Principles of distributed computing*, pages 185–194, 2012.

29. M. Herlihy, D. Kozlov, and S. Rajsbaum. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edition, 2013.

30. A. Kuusisto. Modal logic and distributed message passing automata. In *Computer Science Logic 2013 (CSL 2013)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2013.

31. D. Lewis. Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85(3):297–315, 1976.

32. D. Malkhi, Y. Mansour, and M. Reiter. On diffusing updates in a byzantine environment. In *Proceedings of the 18th IEEE Symposium on Reliable Distributed Systems*, pages 134–143, 1999.

33. J. Moody. Modal logic as a basis for distributed computation. Technical report, Technical Report CMU-CS-03-194, Carnegie Mellon University, 2003.

34. D. Pfleger and U. Schmid. On knowledge and communication complexity in distributed systems. In Z. Lotker and B. Patt-Shamir, editors, *Structural Information and Communication Complexity*, pages 312–330, Cham, 2018. Springer International Publishing.

35. S. Sivasubramanian. Amazon dynamodb: a seamlessly scalable non-relational database service. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, SIGMOD '12, page 729–730, New York, NY, USA, 2012. Association for Computing Machinery.

36. H. van Ditmarsch. Wanted dead or alive: Epistemic logic for impure simplicial complexes. In A. Silva, R. Wassermann, and R. de Queiroz, editors, *Logic, Language, Information, and Computation*, pages 31–46, Cham, 2021. Springer International Publishing.

37. H. van Ditmarsch, J. van Eijck, P. Pardo, R. Ramezanian, and F. Schwarzentruber. Epistemic protocols for dynamic gossip. *Journal of Applied Logic*, 20:1–31, 2017.

38. M. Y. Vardi. Why is modal logic so robustly decidable? Technical report, 1997.

39. D. A. Velázquez, A. Castañeda, and D. A. Rosenblueth. Communication pattern models: An extension of action models for dynamic-network distributed systems. *Electronic Proceedings in Theoretical Computer Science*, 335:307–321, June 2021.