The Probabilistic Logic of Communication and Change

Andreea Achimescu[*], Alexandru Baltag[†], and Joshua Sack[‡]

ILLC, University of Amsterdam
PO Box 94242, 1090 GE Amsterdam, The Netherlands

Abstract

This paper introduces a *Probabilistic Logic of Communication and Change*, which captures in a unified framework subjective probability, arbitrary levels of mutual knowledge and a mechanism for multi-agent Bayesian updates that can model complex social-epistemic scenarios, such as informational cascades. We show soundness, completeness and decidability of our logic, and apply it to a concrete example of cascade.

1 Introduction

In the analysis of many games, as well as of other social phenomena, it is important to be able to represent, not only the agents' probabilistic beliefs and their knowledge, but also higher levels of mutual knowledge, including common knowledge, relativized (i.e. conditional) common knowledge etc. Equally important is to have a rational mechanism for changing both probabilistic beliefs and the levels of knowledge, in a way that can accurately model the epistemic effects of social-dynamic scenarios involving complex multi-agent interactions.

In this paper we propose a unified framework, that combines a variant of the Logic of Communication and Change (LCC) from [4] and a variant of Dynamic-Epistemic Probabilistic Logic (DEPL) from [3]. Our Probabilistic Logic of Communication and Change (PLCC) inherits LCC's ability to express common knowledge as well is DEPL's ability to express probabilistic epistemic dynamics. Moreover, it does this in a way that is completely justifiable from Bayesian first principles (while DEPL seems to go beyond these principles). In Section 4, we provide a sound and complete proof system for PLCC.

While we think that our logic has great potential for applications to various issues in Game Theory and Social Epistemology, we only give here one such application, namely to an *informational cascade* [2]. Such phenomena involve

[*]a.c.achimescu@gmail.com
[†]thealexandrubaltag@gmail.com
[‡]joshua.sack@gmail.com

mechanisms that "rationally" explain conformity in certain social situations. Informational cascades are worth studying because they show how our intuition about social knowledge fails. The expectation is that groups are able to track the truth better than the individuals composing them, by virtue of *communication*, whereby the individual pieces of information each member of the group possesses are pooled together and analysed in a rational manner. However, informational cascades show how sequential communication can impair a group's ability to track the truth.

In this paper, we focus on a classical example of informational cascade, the Urn Example. Such examples have been analysed using logic before in [1] and [9]. But these frameworks could not fully capture syntactically informational cascades, as they lack common knowledge in their language. In contrast, as shown in Section 3, our logic can give a full syntactic encoding of the Urn Example.

2 Probabilistic logic of communication and change

In this section, we introduce our *Probabilistic Logic of Communication and Change* (PLCC), which captures in a unified framework subjective probability, arbitrary levels of mutual knowledge (including common knowledge) and a mechanism for multi-agent Bayesian updates that can model complex social-epistemic scenarios. Essentially, this framework combines an S5 variant of the logic LCC from [4] and a variant of the logic DEPL from [3]. In contrast to LCC [4], our version takes the first level of knowledge (i.e. individual knowledge modalities) to be factive and fully introspective, thus satisfying the standard epistemic system S5. As for DEPDL, the key difference between our event models in Definition 2.3 and the standard update models of [3, p. 77] is that the update models in [3] involve both objective *occurrence* probabilities (of events occurring given certain preconditions) and subjective *observation* probabilities (about what action agents think actually occurred). In our setting, we merge these two types of probabilities to form *subjective occurrence* probabilities, given by the functions pre_a. This gives us a purely Bayesian account, in which all probabilities are subjective. But moreover, the power of approach comes precisely from the *combination* of probabilistic features and higher levels of mutual knowledge.

Language of PLCC

Let At be a set of atomic propositions and Ag a set of agents. We also assume a set of informational events that we will explain later. The language of PLCC, denoted $\mathcal{L}_{\text{PLCC}}$, is given by the following Backus Naur form:

$$\phi ::= \mathit{true} \mid p \mid \neg\phi \mid \phi \wedge \psi \mid [\pi]\phi \mid [e]\phi \mid \alpha_1 \cdot P_a(\phi_1) + \cdots + \alpha_n \cdot P_a(\phi_n) \geq \beta$$
$$\pi ::= a \mid \pi;\pi \mid \pi \cup \pi \mid \pi^* \mid \phi?$$

where $p \in At$ is an atomic proposition, $\alpha_1, \ldots, \alpha_n, \beta$ are rational numbers, $a \in Ag$ is an agent, and e is an event from a given list E of symbols called "events" (which will form the domain of a probabilistic event model A).

The sort ϕ are called *formulas*, and the sort π are called *complex agent terms*. As usual, when a model is given, the formulas denote *propositions* (sets of states), while terms denote binary relations on states interpreted as complex *levels of mutual knowledge*. The formulas $[\pi]\phi$ are *epistemic formulas* inherited from epistemic PDL [4]. The formula $[a]\phi$ is read "agent a knows ϕ", $[a;b]\phi$ is read "a knows that b knows ϕ", $[a \cup b]\phi$ as "It is known to both a and b that ϕ" and $[(a \cup b)^*]\phi$ as "ϕ is common knowledge among a and b". The formula $[\psi?]\phi$ is equivalent to $\psi \to \phi$, but when involved in more complex terms, we can express such notions as "relativized common knowledge" (see [4]) $[(\psi?;(a \cup b))^*]\phi$ read as "ϕ is common knowledge among a and b conditional on ψ".

Formulas $[e]\phi$ are *dynamic formulas*, read as "ϕ holds after any successful informational event e". The "static" sublanguage of PLCC, that does not include any dynamic modalities, is called *probabilistic epistemic PDL* or PE-PDL.

Formulas $\alpha_1 P_a(\phi_1) + \cdots + \alpha_n \cdot P_a(\phi_n) \geq \beta$ are called *a-probability formulas* (for agent a), and express the fact that a linear combination of agent a's probabilities (of various propositions) is at least β.

We abbreviate: *false* by $\neg true$, $\neg[\pi]\neg\phi$ by $\langle\pi\rangle\phi$, $\neg[e]\neg\phi$ by $\langle e\rangle\phi$, $\neg(\neg\phi_1 \wedge \neg\phi_2)$ by $(\phi_1 \vee \phi_2)$, $\neg(\phi_1 \wedge \phi_2)$ by $(\phi_1 \to \phi_2)$, $(\phi_1 \to \phi_2) \wedge (\phi_2 \to \phi_1)$ by $\phi_1 \leftrightarrow \phi_2$. Furthermore we abbreviate the expression $\alpha_1 P_a(\phi_1) + \cdots + \alpha_n \cdot P_a(\phi_n)$ by $\sum_{i=1}^{n} \alpha_i P_a(\phi_i)$, and we may denote it by t if the details of the expression are irrelevant. We also use $c\sum_{j=1}^{n} \alpha_j P_a(\phi_j)$ for $\sum_{j=1}^{n} c\alpha_j P_(\phi_j)$. We then abbreviate

$$\begin{array}{lcl} t \leq \beta & \equiv & -t \geq -\beta \\ t_1 \geq t_2 & \equiv & t_1 - t_2 \geq 0 \\ t_1 > t_2 & \equiv & \neg(t_1 \leq t_2) \\ t_1 = t_2 & \equiv & (t_1 \geq t_2) \wedge (t_2 \geq t_1) \end{array}$$

Semantics of PLCC

Definition 2.1 (Bayesian Kripke models)**.** Given sets Ag and At, a *Bayesian Kripke model* is a quadruple $M = (S, \sim, \mu, V)$ where:

- S is a non-empty set of states.
- $\sim$ is a family of equivalence relations $\sim_a$ on S, one for each agent $a \in Ag$.
- μ is a family of functions $\mu_a : S \to (S \to [0,1])$, one for each agent $a \in Ag$, whose values are denoted by $\mu_a^s(s')$ and satisfy the conditions:
  - (SDP): if $s \sim_a t$ then $\mu_a^s(s') = \mu_a^t(s')$, for all $s' \in S$ (from [8]);
  - (CONS): $\mu_a^s(t) = 0$ if $s \not\sim_a t$ (from [8]);
  - (CAUT): $s \not\sim_a t$ if $\mu_a^s(t) = 0$ (from [6]);[1]

[1]In the submitted abstract, we grouped CONS and CAUT together and called it CONS, but here we think it is better to separate them, as CONS is the name used in [8] for just the one implication we have by that name in the current version.

– (PROB): for every $s \in S$, $\sum_{t \in S} \mu_a^s(t) = 1$.

- $V : At \to \mathcal{P}(S)$ is a *valuation function*.

The relation $\sim_a$ is interpreted as agent a's *epistemic indistinguishability relation*, which induces a's *information partition*, thus modeling a's *knowledge*, in Aumann's style. The function μ_a^s gives agent a's *subjective probability distribution* in state s. The condition (PROB) which is called "probability" ensures that μ_a^s is indeed a probability distribution; (SDP) which is called "state-determined probability" expresses *introspection* of probabilistic beliefs: agents know their own probabilities; (CONS) which is called "consistent" expresses *consistence of beliefs with knowledge*: agents assign probability 0 to propositions they know to be false; finally, (CAUT) which is called "cautious" expresses our assumption that rational agents are *cautious*: they assign probability 0 *only* to propositions they know to be false.[2] In satisfying these conditions, our "Bayesian models" are fundamentally very similar to the models used in [6]. As a pleasant consequence, this notion of "knowledge" is equivalent to "subjective probability = 1". This interesting, but rather unusual, equivalence has been previously adopted in [6] as one of its main principles. This means that we could in principle completely eliminate the epistemic relations $\sim_a$ from our models, by defining them in terms of probabilities: $s \sim_a t$ iff $\mu_a^s(t) \neq 0$.

We now describe an update mechanism that combines multi-agent Bayesian conditioning (for belief change) with epistemic update (for knowledge change) and fact changes (i.e. ontic changes induced by "real" events). Following [4], to model fact changes we employ substitution functions, which will reset the propositional valuation of the initial epistemic model.

Definition 2.2 (Substitutions [4])**.** A *substitution* is a function $\sigma\colon At \to \mathcal{L}_{\mathrm{PLCC}}$ that maps all but a finite number of propositional atoms into themselves. Let $\mathrm{dom}(\sigma) \stackrel{\mathrm{def}}{=} \{p \in At \mid \sigma(p) \neq p\}$ be the the *domain* of σ. Let $sub_{\mathcal{L}_{\mathrm{PLCC}}}$ denote the set of all such possible substitution functions and ϵ the identity substitution.

Definition 2.3 (Event Models)**.** An *event model* over $\mathcal{L}_{\mathrm{PLCC}}$ is the quintuple $A = (E, \sim, \mathrm{PRE}, \mathrm{pre}, \mathrm{sub})$ where:

- E is a finite non-empty set of *events*.

- $\sim$ is a set of equivalence relations $\sim_a$ for each agent $a \in Ag$.

- $\mathrm{PRE} : E \to \mathcal{P}(\mathcal{L}_{\mathrm{PLCC}})$ is a map, such that $\Phi \stackrel{\mathrm{def}}{=} \bigcup_{e \in E} \mathrm{PRE}(e)$ is finite set of pairwise inconsistent formulas.

- pre is a family of functions $\mathrm{pre}_a : \Phi \to (E \to [0,1])$ for each $a \in Ag$ assigning to each precondition $\phi \in \Phi$ a *subjective occurrence probability* distribution over E (i.e. $\sum_{e \in E} \mathrm{pre}_a(\phi)(e) = 1$), such that $\mathrm{pre}_a(\phi)(e) = 0$ iff $\phi \notin \mathrm{PRE}(e)$.

[2]This is a technical assumption, that allows us to do belief revision without having to conditionalize on propositions of probability 0. Another solution would be to move to a non-standard probabilistic setting, e.g. lexicographic probabilities, Popper functions or probabilities with values in a non-standard model of analysis.

- sub : $E \rightarrow sub_{\mathcal{L}_{\mathrm{PLCC}}}$ assigns a substitution function to each event in E.

We abbreviate $\mathrm{pre}_a(\phi)(e)$ by $\mathrm{pre}_a(e \mid \phi)$ in order to improve legibility. We denote by $\mathrm{pre}_a(e \mid s)$ the value of $\mathrm{pre}_a(e \mid \phi_s)$, where ϕ_s is the element of Φ that is satisfied in (M, s). If no such ϕ exists then $\mathrm{pre}_a(e \mid s) = 0$. Finally, we also use the abbreviation $\mathrm{pre}(e) \overset{\mathrm{def}}{=} \bigvee \mathrm{PRE}(e)$.

The equivalence relations $\sim_a$ capture agent a's knowledge about the current event: if $e \sim_a f$, then events e and f are indistinguishable to a at the moment when either of them is happening; $\mathrm{pre}_a(e \mid \phi)$ captures agent a's *(prior) conditional belief about event e given precondition ϕ*. In particular, $\mathrm{pre}_a(e \mid s)$ represents the (prior) conditional probability assigned by agent a to event e happening in state s.

Definition 2.4 (Product Update)**.** The update product of a static Bayesian Kripke model $M = (S, \sim, \mu, V)$ with an event model $A = (E, \sim, \Phi, \mathrm{pre}, \mathrm{sub})$ is the weighted epistemic model $M \otimes A = (S \otimes E, \sim, \mu, V)$ where:

- $S \otimes E \overset{\mathrm{def}}{=} \{(s, e) \mid s \in S, e \in E, (M, s) \models \bigvee \mathrm{PRE}(e)\}$.
- $(s, e) \sim_a (s', e')$ iff $s \sim_a s'$ and $e \sim_a e'$.
- Let $D \overset{\mathrm{def}}{=} \sum_{(s',e') \sim_a (w,g)} (\mu_a^w(s') \cdot \mathrm{pre}_a(e' \mid s'))$, and put:

$$\mu_a^{(w,g)}(s, e) \overset{\mathrm{def}}{=} \begin{cases} \frac{\mu_a^w(s) \cdot \mathrm{pre}_a(e|s)}{D} & \text{if } (s, e) \sim_a (w, g) \\ 0 & \text{otherwise} \end{cases}$$

  (Note that $D \neq 0$ for $(w, g) \in S \otimes E$.)
- $V(p) = \{(s, e) \mid M, s \models \mathrm{sub}(e)(p)\}$

It is easy to check that, if $S \otimes E \neq \emptyset$, then $(M \otimes A)$ is still a Bayesian Kripke model.

Definition 2.5 (Semantics of PLCC)**.** The semantics for $\mathcal{L}_{\mathrm{PLCC}}$ is given by a relation $\models$ between pointed models (M, s), with $M = (S, \sim, \mu, V)$ and $s \in S$, and formulas ϕ, such that

$M, s \models true$	iff	always
$M, s \models p$	iff	$s \in V(p)$
$M, s \models \neg\phi$	iff	$M, s \not\models \phi$
$M, s \models \phi \wedge \psi$	iff	$M, s \models \phi$ and $M, s \models \psi$
$M, s \models [a]\phi$	iff	for all $t \in S$: if $s \sim_a, t$ then $M, t \models \phi$
$M, s \models [e]\phi$	iff	$M, s \models \bigvee \mathrm{PRE}(e)$ then $M \times A, (s, e) \models \phi$, where e is an event in action model A
$M, s \models [\pi]\phi$	iff	for all $t \in S$: if $sR_\pi t$ then $M, t \models \phi$,
$M, s \models \sum_{j=1}^n \alpha_j P_a(\phi_j)$	iff	$\sum_{j=1}^n \alpha_j \cdot \mu_a^s(\phi_j) \geq \beta$

where $\mu_a^s(\phi_j)$ is an abbreviation for $\sum_{s'\in S, s'\models\phi_j} \mu_a^s(s')$, and R_π is a binary relation given by

$sR_a t$	iff	$s \sim_a t$
$sR_{\pi_1\cup\pi_2} t$	iff	$sR_{\pi_1} \cup R_{\pi_2} t$
$sR_{\pi_1;\pi_2} t$	iff	$sR_{\pi_1}; R_{\pi_2} t$ (there is w, such that $sR_{\pi_1} w$ and $wR_{\pi_2} t$)
$sR_{\pi^*} t$	iff	$s(R_\pi)^* t$ (where $(R_\pi)^*$ is the reflexive transitive closure of R_π)
$sR_{\phi?} t$	iff	$s = t$ and $s \models \phi$

We write $\models \varphi$ if $M, s \models \varphi$ for every pointed Bayesian Kripke model M, s.

3 Urn Example

A canonical example of informational cascade is the Urn Example. The narrative behind this example is close to [1]: each individual in a group tries to correctly identify the proportion of black and white balls contained in an urn, that was placed in a room by "Nature". It is common knowledge that the urn can either contain a mix denoted *MW* ("majority white"), consisting of 2 white balls and 1 black ball, or a mix denoted *MB* ("majority black"), consisting of 1 white ball and 2 black balls. It is commonly believed that one of the two mixes was chosen randomly from a uniform distribution (say, by using a fair coin). The agents enter the room one at a time. Upon entering, each agent randomly draws a ball from the urn, looks at it, and returns it to the urm. Then, after exiting, he publicly announces his guess as to which mix he thinks is more probable, *MW* or *MB*, so that all the agents can hear it. We assume that the agents are sincere: they truthfully announce the mix that they really believe to be more probable. In case an agent considers *MB* and *MW* equally likely, he will just guess the color of the ball he drew.

We model this example as follows. Let $Ag \stackrel{\text{def}}{=} \{1, \ldots, n\}$ be the set of agents. The set of atomic propositions is $At \stackrel{\text{def}}{=} \{\mathsf{MW}, \mathsf{MB}\} \cup \{\mathsf{DW}_i, \mathsf{DB}_i, \mathsf{W}_i, \mathsf{B}_i\}_{i\in Ag}$, where MW asserts that the urn has a majority-white mix, DW_i that agent i drew a white ball, W_i that agent i announced a majority-white guess, and similarly for the black. Let $At_{\geq i} \stackrel{\text{def}}{=} \{\mathsf{DW}_j, \mathsf{DB}_j, \mathsf{W}_j, \mathsf{B}_j\}_{i\leq j\leq n}$, and $At_{>i} \stackrel{\text{def}}{=} \{\mathsf{DW}_j, \mathsf{DB}_j, \mathsf{W}_j, \mathsf{B}_j\}_{i<j\leq n}$. Then for $0 < i \leq n$, let

$$\chi_i \stackrel{\text{def}}{=} (\mathsf{MW} \vee \mathsf{MB}) \wedge \bigwedge_{j<i}(\mathsf{DW}_j \vee \mathsf{DB}_j) \wedge \bigwedge_{j<i}(\mathsf{W}_j \vee \mathsf{B}_j)) \wedge \bigwedge_{p\in At_{\geq i}} \neg p$$

$$\chi_i^D \stackrel{\text{def}}{=} (\mathsf{MW} \vee \mathsf{MB}) \wedge \bigwedge_{j\leq i}(\mathsf{DW}_j \vee \mathsf{DB}_j) \wedge \bigwedge_{j<i}(\mathsf{W}_j \vee \mathsf{B}_j) \wedge \neg(\mathsf{W}_i \vee \mathsf{B}_i) \wedge \bigwedge_{p\in At_{>i}} \neg p$$

We then define an event model $A = (E, \sim, \mathrm{PRE}, \mathrm{pre}, \mathrm{sub})$ by

- $E \stackrel{\text{def}}{=} \{\mathsf{dw}_i, \mathsf{db}_i, \mathsf{w}_i, \mathsf{b}_i\}_{i\in Ag}$, where dw_i is the action by which i draws a white ball, w_i is i's public announcement of his guess that the urn is majority-white, and similarly for black.
- for each agent $j = 1, \ldots, n$, $\sim_j$ is the smallest equivalence relation on E satisfying

$$\mathsf{dw}_i \sim_j \mathsf{db}_i \text{ for every } i \neq j.$$

This means that agents cannot see the color of the balls drawn by other agents, but they can all hear every agent's guess when it is publicly announced (since $\mathsf{w}_i \not\sim_j \mathsf{b}_i$ for any i, j).

- $\mathrm{PRE}(e) \stackrel{\text{def}}{=} \begin{cases} \{\psi_i^W, \psi_i^B\} & \text{if } e \in \{\mathsf{dw}_i, \mathsf{db}_i\} \\ \{\phi_i^W, \phi_i^B\} & \text{if } e \in \{\mathsf{w}_i, \mathsf{b}_i\} \end{cases},$

  where

$$\begin{array}{lcl} \psi_i^W & \stackrel{\text{def}}{=} & \mathsf{MW} \wedge \chi_i \\ \psi_i^B & \stackrel{\text{def}}{=} & \mathsf{MB} \wedge \chi_i \\ \phi_i^W & \stackrel{\text{def}}{=} & P_i(\mathsf{MW}) > P_i(\mathsf{MB}) \vee (\mathsf{DW}_i \wedge P_i(\mathsf{MW}) = P_i(\mathsf{MB})) \wedge \chi_i^D \\ \phi_i^B & \stackrel{\text{def}}{=} & P_i(\mathsf{MB}) > P_i(\mathsf{MW}) \vee (\mathsf{DB}_i \wedge P_i(\mathsf{MW}) = P_i(\mathsf{MB})) \wedge \chi_i^D \end{array}$$

- $\mathrm{pre}(e \mid \varphi) = \begin{cases} 2/3 & \exists i : (\varphi = \psi_i^W \ \&\ e = \mathsf{dw}_i) \text{ or } (\varphi = \psi_i^B \ \&\ e = \mathsf{db}_i) \\ 1/3 & \exists i : (\varphi = \psi_i^B \ \&\ e = \mathsf{dw}_i) \text{ or } (\varphi = \psi_i^W \ \&\ e = \mathsf{db}_i) \\ 1 & \exists i : (\varphi = \phi_i^W \ \&\ e = \mathsf{w}_i) \text{ or } (\varphi = \phi_i^B \ \&\ e = \mathsf{b}_i) \\ 0 & \text{otherwise} \end{cases}$

- $\mathrm{sub}(e, p) = \begin{cases} \chi_i & (e = \mathsf{dw}_i \ \&\ p = \mathsf{DW}_i) \text{ or } (e = \mathsf{db}_i \ \&\ p = \mathsf{DB}_i) \\ \chi_i^D & (e = \mathsf{w}_i \ \&\ p = \mathsf{W}_i) \text{ or } (e = \mathsf{b}_i \ \&\ p = \mathsf{B}_i) \\ p & \text{otherwise} \end{cases}$

We depict many features of this action model in Figure 1.

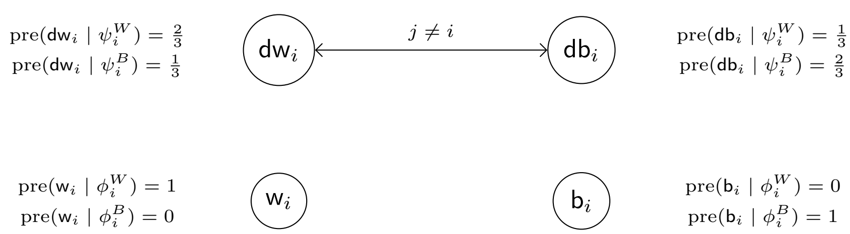

Figure 1: Our event model A for the Urn Example. Although only four events appear in the diagram above, the actual event model has $4n$ events.

Our example assumes that we start with a Bayesian Kripke model in which: there is an urn having one of the two mixes MW or MB, nobody knows which of the two mixes is in and that everybody believes the mix was chosen at random using a fair coin (i.e. everybody assigns odds $1 : 1$ to the mixes). These assumptions are encoded in the following formula:

$$\chi \stackrel{\text{def}}{=} (\mathsf{MW} \vee \mathsf{MB}) \wedge \neg(\mathsf{MW} \wedge \mathsf{MB}) \wedge \textstyle\bigwedge_{i \in Ag} \big(P_i(\mathsf{MW}) = P_i(\mathsf{MB})\big) \wedge \bigwedge_{p \in At_{\geq 1}} \neg p$$

Moreover, in our initial Bayesian model, *all the above assumptions are common knowledge*: this can be encoded in the epistemic formula

$$[(\cup_{i \in Ag} i)^*]\chi.$$

The simplest model M_0 satisfying these assumptions has two states $S = \{s, t\}$, with $\sim_i = S \times S$, $\mu_i^s(s) = \mu_i^s(t) = \mu_i^t(s) = \mu_i^t(t) = 1/2$, $V(\mathsf{MW}) = \{s\}$, $V(\mathsf{MB}) = \{t\}$ and $V(p) = \emptyset$ for all other atoms. Given such a model, assume agent 1 draws a white ball from the urn (event dw_1). By applying Bayesian reasoning, agent 1 thinks it is more likely that urn 1 is majority-white, and hence w_1 can be performed while b_1 cannot (agent 1 publicly guesses "white"). Suppose that after this agent 2 also draws a white ball, and so he similarly guesses that the urn is majority-white. All agents now know that two white balls have been drawn and gives odds of $4:1$ that the urn is majority-white (rather than majority-black). At this point a cascade begins. Each new agent to enter the room will draw either a black ball or a white ball, but either way they will still consider MW more likely. Since the other agents anticipate this, all subsequent public guesses are uninformative, and can be simply ignored.

But the first two guesses of white were still enough to ensure that all the agents will forever consider MW more likely than MB, regardless of what color ball they may draw. Thus a cascade has begun, where agents base their choices on the inferences of what agents probably believe given their actions, rather than base their choices on their own observations from nature.

We can simulate the above scenario by performing successive updates of the initial model M_0 with the event model A, obtaining successive models $((M_0 \otimes A) \otimes A) \cdots \otimes A$. But in fact there is nothing special about our model M_0: *any* model satisfying the above assumptions will lead to the same cascade, as shown in the following proposition:

Proposition 3.1. *Let* $3 \le i \le n$. *For all* $1 \le j \le i$, *let* $f_j \in \{\mathsf{dw}_j, \mathsf{db}_j\}$ *and* $g_j \in \{\mathsf{w}_j, \mathsf{b}_j\}$. *Then*

$$[(\cup_{i \in Ag} i)^*]\chi \Rightarrow [\mathsf{dw}_1][g_1][\mathsf{dw}_2][g_2][f_3][g_3] \dots [f_i][g_i]\big(P_k(\mathsf{MW}) > P_k(\mathsf{MB})\big)$$

is a valid formula, for all $1 \le k \le n$.

Proof. We focus our proof on a model M_0 (soon to be defined), such that any pointed model satisfying $[(\cup_{i \in Ag} i)^*]\chi$ is "bisimilar" one of its states, and in which $M_0 \models [\mathsf{dw}_1][g_1][\mathsf{dw}_2][g_2][f_3][g_3] \dots [f_i][g_i]\big(P_k(\mathsf{MW}) > P_k(\mathsf{MB})\big)$. Bisimulation is preserved by actions and bisimilar pointed models satisfy the same formulas. Our desired result then follows.

A *bisimulation* between Bayesian Kripke models $(M, x) = ((S, \sim^M, \mu, V^M), x)$ and $(N, y) = ((T, \sim^N, nu, V^N), x)$ is a relation Z satisfying the the following.

- For every $X \subseteq S$, $\mu(X) \le \nu(\{y \mid xZy\})$,
- For every $Y \subseteq T$, $\nu(Y) \le \mu(\{x \mid xZy\})$,
- Whenever xZy, the following hold
  - $x \in V^M(p)$ iff $y \in V^N(p)$ for every $p \in At$,
  - if $x \sim x'$, then there exists y', such that $y \sim y'$ and $x'Zy'$,
  - if $y \sim y'$, then there exists x', such that $x \sim x'$ and $x'Zy'$.

We say two pointed models are bisimilar if there is a bisimulation between them. The two lemmas are can be proved using standard methods.

Lemma 3.2. *If (M, x) and (N, y) are bisimilar, and (A, e) is a pointed action model, then $M \otimes A, (s, e)$ and $N \otimes A, (t, e)$ are bisimilar.*

Lemma 3.3. *If (M, x) and (N, y), then $(M, x) \models \varphi$ if and only if $(N, y) \models \varphi$ for any $\varphi \in \mathcal{L}_{PLCC}$.*

The model M_0 consists of two states, $\mathbb{MW}$ and $\mathbb{MB}$, where $\mathbb{MW} \sim_k \mathbb{MB}$ for all agents k, and each agent gives equal probability to both $\mathbb{MW}$ and $\mathbb{MB}$. The atomic proposition MW is true only at $\mathbb{MW}$ and MB is true only at $\mathbb{MB}$.

We graphically represent Bayesian Kripke models as follows: the worlds are depicted by rectangles, which contain the atoms true at that world. These atoms, put together in a sequence, label the world they're in. Worlds are connected by arrows, which represent equivalence relations. Every arrow is labelled by the agents who cannot distinguish between the worlds connected by the arrow. We do not denote reflexive arrows in our models, since they are always assumed to be there for any model represented. Given the SDP condition, we have that within every agent's information set, his probability assignments at each world are the same. Therefore, we represent only one probability assignment per world per agent, next to the name of the world. In model M_0 of Figure 2, the actual state $\mathbb{MB}$ is represented by the bold-font rectangle, whereas the label on the arrows designates the agents that cannot distinguish between the two states of the world, $\mathbb{MW}$ and $\mathbb{MB}$. The probabilities that each player assigns to the worlds are represented on the side of each rectangle, preceded by the players that hold these beliefs.

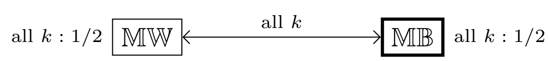

Figure 2: The initial state model M_0, after Nature picks the state of the world $\mathbb{MB}$. This action is not observable by any of the agents $k \in Ag$. All players give each world equal chances of being true.

It is clear that $M_0 \models [(\cup_{i \in Ag} i)^*]\chi$ (satisfied in both states), and it is not hard to see that any pointed model satisfying $[(\cup_{i \in Ag} i)^*]\chi$ will be bisimilar to a state in M_0.

Agent 1 draws a white ball The result of updating M_0 with event dw_1 is the new state model M_1 represented in Figure 3. Observe that agent 1 knows she has drawn a white marble dw_1, while not being able to discern the true urn $\mathbb{MB}$. All the other players in the game remain ignorant with regard to player 1's private draw, and can therefore exclude no world.

Given the new state model M_1, we can calculate the probabilities that each agent gives to the new states, using the probability update formula introduced in Section 2. For example, we compute the revised probability assignment of

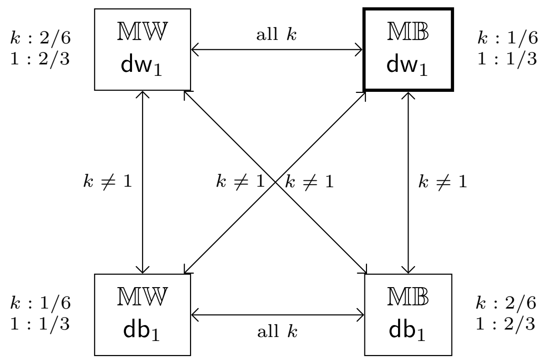

Figure 3: The product update model M_1 resulting after the private announcement of agent 1's signal.

player 1, since any other players' informational state does not change as a result of event dw_1. We drop the world indexation in the product update rule in order to improve legibility, since by the SDP condition, the probability assignment of an agent is the same at every world within the same information set. For example, applying the product update rule

$$\mu_i(s,e) = \frac{\mu_i(s) \cdot \mathrm{pre}_i(e \mid s)}{\sum_{(s',e') \in S_1} \mu_i(s') \cdot \mathrm{pre}_i(e' \mid s')}$$

we have that

$$\mu_1(\mathsf{MW}, \mathsf{dw}_1) = \frac{\frac{1}{2} \cdot \frac{1}{3}}{\frac{1}{2}} = \frac{2}{3}$$

This represents, intuitively, the probability player 1 assigns to the state of the world being $\mathbb{MW}$ (satisfying proposition MW), given he received the private signal dw_1. The other probabilities are computed in the same way and included in Figure 3.

Agent 1's action The pre function encodes the common knowledge of rationality and tie-breaking rule assumptions: agent 1 only announces w_1 if he either believes MW to be more likely than MB or he believes them to be equally probable but his private signal was dw_1. In the previous paragraph, we computed agent 1's subjective probability that MW is true, $\mu_1(\mathsf{MW}, \mathsf{dw}_1) = \frac{2}{3}$. Therefore player 1 will choose w_1. This event gives rise to the update model M_2. The product update model will still have four worlds, but the two worlds where w_1 is true will be unrelated (for all players) to the two worlds in which db_1 is true: it is common knowledge in which of these two zones the players are. Assuming the real world is $\mathbb{MB}\mathsf{dw}_1\mathsf{w}_1$, we can thus disregard the b_1 worlds as irrelevant (inaccessible, impossible). More precisely, the 4-world model with actual world $\mathbb{MB}\mathsf{dw}_1\mathsf{w}_1$ is *bisimilar* to the the 2-world model (having only the w_1 worlds) with the same actual world $\mathbb{MB}\mathsf{dw}_1\mathsf{w}_1$. So we can just delete the b_1 worlds, obtaining the model in Figure 4. From now on, we abuse notation and refer to the bisimilar model as being the product update model itself.

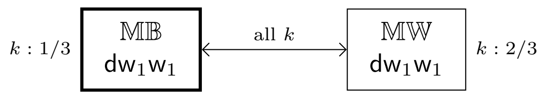

Figure 4: The product update model M_2 after agent 1's announcement. The consequence of imposing common knowledge of rationality on the current model is the deletion of the worlds at which b_1 is true. All agents, knowing that player 1 is rational, are able to deduce that agent 1 saw a white ball.

Agent 2 draws a white ball Again, as above, we model the effect of this event using a product update model $M_3 = M_2 \otimes A$ in Figure 5.

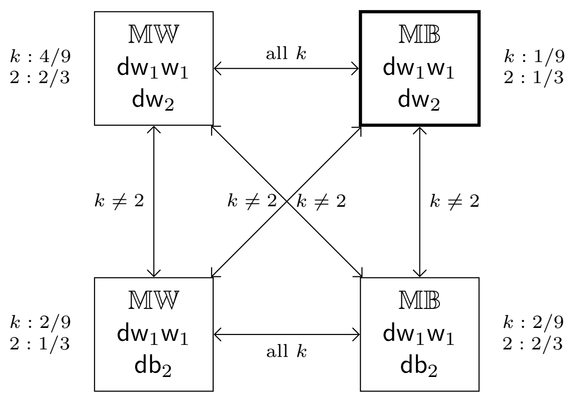

Figure 5: The product update model M_3 resulting after the private announcement of agent 2's signal.

Agent 2's action The preconditions are designed to make player 2 choose action w_2, given that the subjective belief he attaches to the world satisfying MW is given by $\mu_2(\mathsf{MW}) = \frac{4}{5} > \frac{1}{2}$. This public announcement will give rise to the product update model M_4. Graphically this is represented in Figure 6:

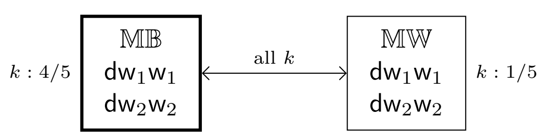

Figure 6: The product update model M_4 after agent 2's public announcement. The consequence of imposing common knowledge of rationality on the current model is the deletion of the worlds at which db_2 is true.

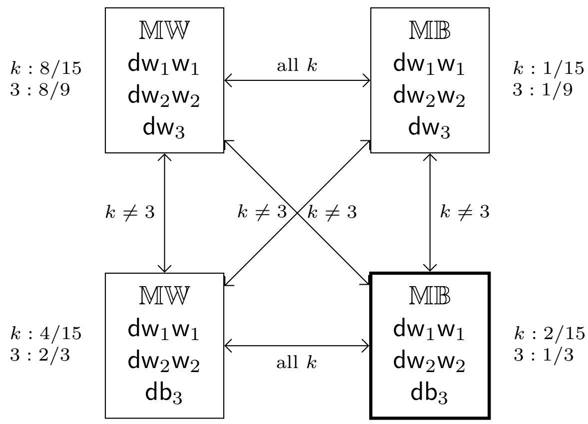

Figure 7: The product update model M_5 resulting after the private announcement of player 3's signal.

Agent 3 draws a ball Here we consider what happens when agent 3 draws a ball (either black or white). We continue our example as though agent 3 draws a black ball, but the updated model also contains information about what would have happened were he to draw white. The update model that follows this event db_3, denoted by M_5, is given by Figure 7. In the pictorial representation, you can find the subjective probability player 3 assigns to the world satisfying MW, given the observed actions of previous players and his own private signal.

Agent 3's action As we have gotten used to by now, we are going to construct a product update model M_6, as a result of the public communication of action w_3 (the only action that satisfies the preconditions in M_5. In this new model, no worlds will be deleted, since no agent except agent 3 can distinguish between dw_3 and db_3, based solely on the assumption of common knowledge of rationality. For agent 3, as we have argued, it is consistent with rationality to choose action w_3, both in the case that he receives a dw_3 and db_3. This means this agent has entered into a false cascade, and others cannot infer with his private signal though his choice.

The same reasoning can be applied to any subsequent agent, who will rationally choose w_i, regardless of his private signal. This is due to the failure of extracting any extra information for any players, except player 1, player 2 and himself. Thus, as argued before, agents will enter a cascade, in which every player $i > 2$ imitates his predecessor.

We set to prove by induction on $n \geq 3$ that the model $M_{2n-1} \models P_j(\mathsf{MW}) \geq 2 \cdot P_j(\mathsf{MB})$ for all $j \leq n$ and $M_{2n-1} \models P_j(\mathsf{MW}) \geq 4 \cdot P_j(\mathsf{MB})$ for all $j > n$. The discussion in the previous paragraphs proves the result for $n = 3$.

Inductive Hypothesis Assume the proposition holds for all $i \leq n-1$ and try to prove it holds for n. In particular the inductive hypothesis holds for

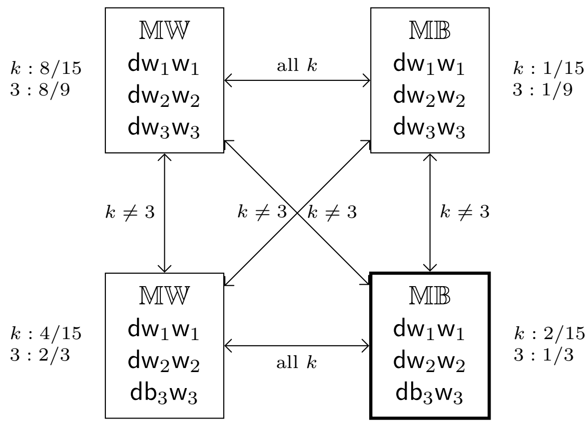

Figure 8: The product update model M_6 resulting after the private announcement of player 3's signal.

$n-1$, and therefore get state model M_{2n-3}, which we will represent partially, by lumping together all the W-worlds and, respectively, B-worlds as presented in Figure 9

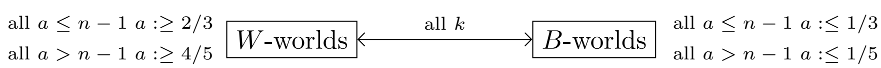

Figure 9: The state model M_{2n-3}, representing the beliefs of players after player $n-1$ has seen his signal. The probabilities express the sentence in the proposition, in terms of actual probability assignments. For example, $P_j(\mathsf{MW}) \geq 2 \cdot P_j(\mathsf{MB})$ is equivalent to saying that $\mu_j(\mathsf{MW}) \geq \frac{2}{3}$

.

Next, player $n-1$ will publicly announce w_{n-1}, as demanded by his beliefs. This announcement will not change the informational state of any agent, since no one except $n-1$ can infer anything about the private signal of player $n-1$. Therefore, the new model M_{2n-2} will be identical to M_{2n-3} in terms of beliefs of players. This is so because at every world in the model M_{2n-3}, the sentence $P_{n-1}(\mathsf{MW}) > \frac{1}{2}$ is true, and therefore common knowledge. The next event is either dw_n or db_n. The new product update model that results from M_{2n-2} and A is presented graphically in Figure 10.

Applying the technique of lumping together W-worlds and respectively B-worlds, we end up with a model of the form:

Therefore, model M_{2n-1} satisfies

$$P_j(\mathsf{MW}) \geq 2 \cdot P_j(\mathsf{MB}) \text{ for all } j \leq n, \text{ and}$$

$$P_j(\mathsf{MW}) \geq 4 \cdot P_j(\mathsf{MB}) \text{ for all } j > n$$

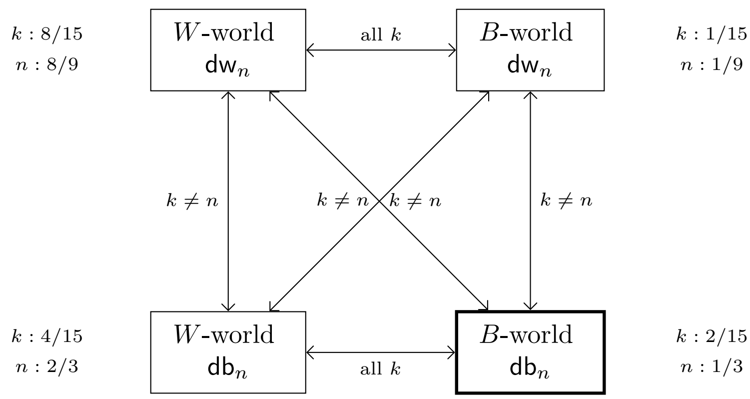

Figure 10: The product update model M_{2n-1} resulting after the private announcement of player n's signal.

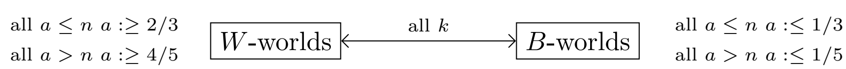

Hence we proved the induction step for n.

□

4 Proof System of $\mathcal{L}_{\mathbf{PLCC}}$

For each finite event model E, we will give a proof system for the logic PLCC having only dynamic modalities for events $e \in E$. To state our axioms, we need to introduce a "program transformer" notation $T_{ef}(\pi)$ for all programs π and events $e, f \in E$.

Definition 4.1 (T_{ij} program transformers and $K_{ijk}(\pi)$ path transformers)**.** To define $T_{ef}\pi$, let m be the number of events in E, and let $(e_1, e_2, \ldots, e_m)$ be an enumeration of all the events in E without repetitions. For all $1 \leq i, j, k \leq m$ and all programs π, we first define program transformers $T_{ij}(\pi)$ and path transformers $K_{ijk}(\pi)$. The definition is by (double) recursion on the complexity of π and on the number k.

$$T_{ij}(a) = \begin{cases} ?\mathrm{pre}(e_i); a & \text{if } e_i R(a) e_j \\ ?\perp & \text{otherwise} \end{cases}$$

$$T_{ij}(?\phi) = \begin{cases} ?(\mathrm{pre}(e_i) \wedge [e_i]\phi) & \text{if } i = j \\ ?\perp & \text{otherwise} \end{cases}$$

$$T_{ij}(\pi_1; \pi_2) = \bigcup_{k=0}^{n-1} (T_{ik}(\pi_1); T_{kj}(\pi_2))$$

$$T_{ij}(\pi_1 \cup \pi_2) = T_{ij}(\pi_1) \cup T_{ij}(\pi_2)$$

$$T_{ij}(\pi^*) = K_{ijm}(\pi)$$

and

$$K_{ij0}(\pi) = \begin{cases} ?\mathit{true} \cup T_{ij}(\pi) & \text{if } i = j \\ T_{ij}(\pi) & \text{otherwise} \end{cases}$$

$$K_{ij(k+1)}(\pi) =$$

$$\begin{cases} (K_{kkk}(\pi))^* & \text{if } i = k = j \\ (K_{kkk}(\pi))^*; K_{kjk}(\pi) & \text{if } i = k \neq j \\ K_{ikk}(\pi); (K_{kkk}(\pi))^* & \text{if } i \neq k = j \\ K_{ijk}(\pi) \cup \big(K_{ikk}(\pi); (K_{kkk}(\pi))^*; K_{kjk}(\pi)\big) & \text{otherwise } (i \neq k \neq j) \end{cases}$$

Finally, for two events $e = e_i, f = e_j \in E$, we put

$$T_{ef}\pi := T_{ij}(\pi).$$

The Proof System $\mathcal{T}_{\mathbf{PLCC}}$. Our *proof system* $\mathcal{T}_{\mathrm{PLCC}}$ for the logic $\mathcal{L}_{\mathrm{PLCC}}$ is given in Figure 11. We write $\vdash \varphi$ to indicate that φ can be proved in $\mathcal{T}_{\mathrm{PLCC}}$. Axioms K1-3 ensure that the R_a (for atomic $a \in Ag$) are all equivalence relations (though not all R_π are equivalence relations). Axioms A1-6 are basic axioms concerning complex agents. Axioms I1-6 ensure that the inequalities in probability formulas behave appropriately. Axioms W1-7 ensures that probability terms behave appropriately. Axioms D1-D5 are schema concerning dynamic operators, and depend on the "program transformer" of Definition 4.1.

The proof system $\mathcal{T}_{\text{PE-PDL}}$ of PE-PDL consists of all of the schema in Figure 11 through W7 including R1 and R2. This extends the E-PDL of [4] by incorporating the schema K1-3, I1-6, and W1-7 from [8]. We adapt a translation similar to one in [4] from LCC to E-PDL into a truth preserving translation of PLCC into PE-PDL. Although this shows that PLCC and PE-PDL are equally expressive, PLCC is much more succinct at expressing dynamic phenomena.

Theorem 4.2 (Soundness)**.** *The proof system $\mathcal{L}_{PLCC}$ is sound with respect to Bayesian Kripke structures if and only if, for $\phi \in \mathcal{L}_{PLCC}$:*

$$\vdash \phi \textit{ implies } \models \phi$$

Axiom Schemata	
PL.	All propositional tautologies
K1.	$[a]\phi \to \phi$
K2.	$[a]\phi \to [a][a]\phi$
K3.	$\neg[a]\phi \to [a]\neg[a]\phi$
A1.	$[\pi](\phi \to \psi) \to ([\pi]\phi \to [\pi]\psi)$
A2.	$[\pi_1;\pi_2]\phi \leftrightarrow [\pi_1][\pi_2]\phi$
A3.	$[\pi_1 \cup \pi_2]\phi \leftrightarrow [\pi_1]\phi \wedge [\pi_2]\phi$
A4.	$[\pi^*]\phi \leftrightarrow (\phi \wedge [\pi][\pi^*]\phi)$
A5.	$[\pi^*](\phi \to [\pi]\phi) \to (\phi \to [\pi^*]\phi)$
A6.	$[\phi?]\psi \leftrightarrow (\phi \to \psi)$
I1.	$t \geq \beta \leftrightarrow t + 0P_a(\phi) \geq \beta$
I2.	$\sum_{k=1}^n \alpha_k P_a(\phi_k) \geq \beta \to \sum_{k=1}^n \alpha_{j_k} P_a(\phi_{j_k}) \geq q\beta$ for any permutation $j_1, \ldots, j_n$ of $1, \ldots, n$
I3.	$\sum_{k=1}^n \alpha_k P_a(\phi_k) \geq \beta \wedge \sum_{k=1}^n \alpha'_k P_a(\phi_k) \geq \beta' \to \sum_{k=1}^n (\alpha_k + \alpha'_k) P_a(\phi_k) \geq (\beta + \beta')$
I4.	$t \geq \beta \leftrightarrow dt \geq d\beta$ if $d > 0$
I5.	$t \geq \beta \vee t \leq \beta$
I6.	$t \geq \beta \to t \geq \gamma$ if $\beta > \gamma$
W1.	$P_a(\phi) \geq 0$
W2.	$P_a(\mathit{true}) = 1$
W3.	$P(\phi \wedge \psi) + P(\phi \wedge \neg\psi) = P_a(\phi)$
W4.	$P_a(\phi) = P_a(\psi)$ if $\phi \leftrightarrow \psi$ is a propositional tautology.
W5.	$P_a(\mathit{false}) = 0$
W6.	$[a]\phi \leftrightarrow P_a(\phi) \geq 1$
W7.	$w \to [a]w$, for any w an a-probability formulas.
D1.	$[e]p \leftrightarrow (\mathrm{pre}(e) \to \mathrm{sub}(e)(p))$
D2.	$[e]\neg\phi \leftrightarrow (\mathrm{pre}(e) \to \neg[e]\phi)$
D3.	$[e](\phi \wedge \psi) \leftrightarrow ([e]\phi \wedge [e]\psi)$
D4.	$[e][\pi]\phi \leftrightarrow \bigwedge_{f \in E} [T_{ef}\pi][f]\phi$ (where T_{ef} is given in Definition 4.1)
D5.	$[e]\left(\sum_{k=1}^n \alpha_k \cdot P_a(\psi_k) \geq \beta\right) \leftrightarrow (\mathrm{pre}(e) \to C \geq D)$, where $C = \sum_{\phi\in\Phi} \sum_{f \sim_a e} \sum_{k=1}^n \alpha_k \cdot \mathrm{pre}_a(f \mid \phi) \cdot P_a(\phi \wedge [f]\psi_k)$, $D = \sum_{\phi\in\Phi} \sum_{f \sim_a e} \beta \cdot \mathrm{pre}_a(f \mid \phi) \cdot P_a(\phi)$.

Rules		
R1. $\dfrac{\phi \to \psi \quad \phi}{\psi}$	R2. $\dfrac{\phi}{[\pi]\phi}$	R3. $\dfrac{\phi}{[e]\phi}$

Figure 11: The Proof System $\mathcal{T}_{\mathrm{PLCC}}$

Induction on the length of the proof. It is sufficient to prove that every axiom is sound and each inference rule preserves truth. This is a routine proof, so we will only check the soundness of the most difficult reduction axiom:

$$[e]\left(\sum_{h=1}^{k} \alpha_h \cdot P_a(\psi_h) \geq \beta\right) \leftrightarrow (\mathrm{pre}(e) \to C \geq D)$$

where

$$\begin{array}{rcl} C & = & \sum_{\phi\in\Phi} \sum_{f\sim_a e} \sum_{h=1}^{k} \alpha_h \cdot \mathrm{pre}_a(f \mid \phi) \cdot P_a(\phi \wedge [f]\psi_h), \\ D & = & \sum_{\phi\in\Phi} \sum_{f\sim_a e} \beta \cdot \mathrm{pre}_a(f \mid \phi) \cdot P_a(\phi). \end{array}$$

Take an arbitrary Bayesian Kripke model M and a state s, such that:

$$M, s \models [e]\left(\sum_{1\leq h\leq k} \alpha_h \cdot P_a(\psi_h) \geq \beta\right)$$

Assume

$$M, s \models \mathrm{pre}(e).$$

Then by definition of the semantics.

$$\sum_{1\leq h\leq k} \alpha_h \sum_{\substack{(s',e')\sim_a(s,e) \\ M\otimes A,(s',e')\models\psi_h}} \mu_a^{(s,e)}(s',e') \geq \beta$$

$$\overbrace{\Leftrightarrow}^{\text{by the product update rule}}$$

$$\sum_{1\leq h\leq k} \alpha_h \frac{\sum_{\substack{(s',e')\sim_a(s,e) \\ M\otimes A,(s',e')\models\psi_h}} \mu_a^s(s')\,\mathrm{pre}(e' \mid s')}{\sum_{(w,f)\sim_a(s,e)} \mu_a^s(w)\,\mathrm{pre}(f \mid w)} \geq \beta$$

Now

$$\sum_{1\leq h\leq k} \alpha_h \sum_{\substack{(s',e')\sim_a(s,e) \\ M\otimes A,(s',e')\models\psi_h}} \mu_a^{(s,e)}(s',e') \geq \beta$$

$$\overbrace{\Leftrightarrow}^{\text{by the product update rule}}$$

$$\sum_{1\leq h\leq k} \alpha_h \frac{\sum_{\substack{(s',e')\sim_a(s,e) \\ M\otimes A,(s',e')\models\psi_h}} \mu_a^s(s')\,\mathrm{pre}(e' \mid s')}{\sum_{(w,f)\sim_a(s,e)} \mu_a^s(w)\,\mathrm{pre}(f \mid w)} \geq \beta$$

$$\overbrace{\Leftrightarrow}^{\text{re-aranging the terms}}$$

$$\sum_{1\leq h\leq k} \alpha_h \sum_{\substack{(s',e')\sim_a(s,e) \\ M\otimes A,(s',e')\models\psi_h}} \mu_a^s(s')\,\mathrm{pre}(e' \mid s') \geq \beta \sum_{(w,f)\sim_a(s,e)} \mu_a^s(w)\,\mathrm{pre}(f \mid w)$$

$$\overset{\text{by grouping worlds according to the preconditions they satisfy, for every } f \sim_a e}{\widehat{\Leftrightarrow}}$$

$$\sum_{1 \leq h \leq k} \alpha_h \sum_{\substack{f \sim_a e \\ \phi_i \in \Phi}} pre_a(f|\phi_i) \sum_{\substack{s' \sim_a s \\ s' \models [f]\psi_h \\ s' \models \phi_i}} \mu_a^s(s') \geq \beta \sum_{\substack{f \sim_a e \\ \phi_i \in \Phi}} \mathrm{pre}_a(f \mid \phi_i) \sum_{w \models \phi_i} \mu_a^s(w)$$

$$\overset{\text{by semantic definition}}{\widehat{\Leftrightarrow}}$$

$$M, s \models \sum_{1 \leq h \leq k} \alpha_h \sum_{\substack{f \sim_a e \\ \phi_i \in \Phi}} \mathrm{pre}_a(f \mid \phi_i) \cdot P_a(\phi_i \wedge [f]\psi_h) \geq \beta \sum_{\substack{f \sim_a e \\ \phi_i \in \Phi}} \mathrm{pre}_a(f \mid \phi_i) P_a(\phi_i)$$

$$\overset{\text{re-grouping the sums}}{\widehat{\Leftrightarrow}}$$

$$M, s \models \sum_{\substack{1 \leq h \leq k \\ f \sim_a e \\ \phi_i \in \Phi}} \alpha_h \, \mathrm{pre}_a(f \mid \phi_i) \cdot P_a(\phi_i \wedge [f]\psi_h) \geq \sum_{\substack{f \sim_a e \\ \phi_i \in \Phi}} \beta \, \mathrm{pre}_a(f \mid \phi_i) P_a(\phi_i)$$

Reducing the notation for the terms we have then that

$$M, s \models C \geq D. \qquad \square$$

4.1 Completeness of PLCC

We first prove the completeness of the static language PE-PDL, which we call PE-PDL, and then argue, via the reduction axioms, that every $\mathcal{L}_{\mathrm{PLCC}}$-formula can be translated into a $\mathcal{L}_{\mathrm{PLCC}}$-formula.

4.1.1 Completeness of the static language PE-PDL

The following definitions have been adapted from [5] to include the test operator.

Definition 4.3 (Fischer-Ladner closure)**.** Let X be a set of formulas. Then X is *Fischer-Ladner closed* if it is closed under subformulas and satisfies the following additional constraints:

(i) If $[\pi_1; \pi_2]\phi \in X$ then $[\pi_1][\pi_2]\phi \in X$

(ii) If $[\pi_1 \cup \pi_2]\phi \in X$ then $[\pi_1]\phi \wedge [\pi_2]\phi \in X$

(iii) If $[\pi^*]\phi \in X$ then $[\pi][\pi^*]\phi \in X$

(iv) If $[\phi?]\psi \in X$ then $\phi \to \psi \in X$.

(v) If $w \in X$ then $[a]w \in X$, for w an a-probability formula.

If Σ is any set of formulas then FL(Σ) (the *Fischer-Ladner closure* of Σ) is the smallest set of formulas containing Σ that is Fischer-Ladner closed.

Given a formula ϕ, we define $\sim\phi$ as the following formula

$$\sim\phi = \begin{cases} \psi & \text{if } \phi \text{ is of the form } \neg\psi \\ \neg\phi & \text{otherwise.} \end{cases}$$

A set of formulas X is *closed under single negations* if $\sim\phi$ belongs to X whenever $\phi \in X$.

We define $\neg$FL(Σ), the *closure of* Σ, as the smallest set containing Σ which is Fischer-Ladner closed and closed under single negations.

Definition 4.4 (Atoms)**.** Let Σ be a set of formulas. A set of formulas A is an *atom* over Σ if it is a maximal consistent subset of $\neg$FL(Σ). That is, A is an atom over Σ if $A \subseteq \neg\text{FL}(\Sigma)$, if A is consistent, and if $A \subset B \subseteq \neg\text{FL}(\Sigma)$ then B is inconsistent. $\mathsf{At}(\Sigma)$ is the set of all atoms over Σ.

The following lemma (except for item 6) closely follows [5, lemma 4.81].

Lemma 4.5. *Let Σ be any set of formulas, and A any element of $\mathsf{At}(\Sigma)$. Then*

1. *For all $\phi \in \neg FL(\Sigma)$: exactly one of ϕ and $\sim\phi$ is in A*
2. *For all $\phi \vee \psi \in \neg FL(\Sigma)$: $\phi \vee \psi \in A$ iff $\phi \in A$ or $\psi \in A$*
3. *For all $\langle\pi_1;\pi_2\rangle\phi \in \neg FL(\Sigma)$: $\langle\pi_1;\pi_2\rangle\phi \in A$ iff $\langle\pi_1\rangle\langle\pi_2\rangle\phi \in A$*
4. *For all $\langle\pi_1 \cup \pi_2\rangle\phi \in \neg FL(\Sigma)$: $\langle\pi_1 \cup \pi_2\rangle\phi \in A$ iff $\langle\pi_1\rangle\phi \in A$ or $\langle\pi_2\rangle\phi \in A$*
5. *For all $\langle\pi^*\rangle\phi \in \neg FL(\Sigma)$: $\langle\pi^*\rangle\phi \in A$ iff $\phi \in A$ or $\langle\pi\rangle\langle\pi^*\rangle\phi \in A$*
6. *For all $[\phi?]\psi \in \neg FL(\Sigma) : [\phi?]\psi \in A$ iff $\phi \to \psi$.*

Now it is time to define the canonical model over Σ.

Definition 4.6 (Canonical model over Σ)**.** Let Σ be a finite set of formulas. The canonical model over Σ is the triple $(\mathsf{At}(\Sigma), \{S^{\Sigma}_{\pi}\}_{\pi\in\Pi}, V^{\Sigma})$ where for all the propositional variables p, $V^{\Sigma}(p) = \{A \in \mathsf{At}(\Sigma) \mid p \in A\}$ and for all atoms $A, B \in \mathsf{At}(\Sigma)$ and all programs π,

$$AS_{\pi}B \text{ if } \phi_A \wedge \langle\pi\rangle\phi_B \text{ is consistent.}$$

where ϕ_A is defined as the conjunction of all formulas that belong to A.

Definition 4.7 (Regular model over Σ)**.** Let Σ be a finite set of formulas. For all basic programs a, define R^{Σ}_a as:

$$AR_aB \text{ if } \forall\phi \in \neg FL(\Sigma), [\pi]\phi \in A \text{ iff } [\pi]\phi \in B$$

For the complex programs, inductively define the PDL relations such that $\forall \pi, \pi_1, \pi_2$, we have:

$$\begin{aligned} R_{\pi_1 \cup \pi_2} &= R_{\pi_1} \cup R_{\pi_2} \\ R_{\pi_1 ; \pi_2} &= R_{\pi_1} ; R_{\pi_2} \\ R_{\pi^*} &= (R_\pi)^* \\ A R_{\phi ?} B &\Leftrightarrow A = B \text{ and } \phi \in A \end{aligned}$$

Finally, define $\mathcal{R}$, the regular model over Σ, to be

$$\mathcal{R} = (\mathsf{At}(\Sigma), \{R_\pi^\Sigma\}_{\pi \in \Pi}, V^\Sigma)$$

where V^Σ is the canonical valuation.

Given the regular model $\mathcal{R} = (\mathsf{At}(\Sigma), \{R_\pi^\Sigma\}_{\pi \in \Pi}, V^\Sigma)$, our goal is to define a probability assignment

$$\mu^\Sigma \colon Ag \to \Big(\mathsf{At}(\Sigma) \to (\mathsf{At}(\Sigma) \to [0,1])\Big)$$

s.t. if we consider the Bayesian Kripke structure

$$\mathcal{M} = (\mathsf{At}(\Sigma), \{R_\pi^\Sigma\}_{\pi \in \Pi}, \mu^\Sigma, V^\Sigma)$$

then for every state $A \in \mathsf{At}(\Sigma)$ and every $\psi \in \neg\mathrm{FL}(\Sigma)$ we have $(\mathcal{M}, A) \models \psi$ iff $\psi \in A$.

Lemma 4.8. *For any $a \in Ag$ and any atom A, there exists a probability function $\mu_a : \mathsf{At}(\Sigma) \to [0,1]$ that can realize all a-probability formulas $w \in A$ together.*

Proof. Using only propositional reasoning, we can show that:

$$\vdash \psi \leftrightarrow \bigvee_{\{A \in \mathsf{At}(\Sigma) \mid \psi \in A\}} \phi_A, \text{ for all } \psi \in \neg\mathrm{FL}(\Sigma) \tag{1}$$

$$\vdash \phi_A \to \neg\phi_B, \text{ for any } A, B \in \mathsf{At}(\Sigma), A \neq B \tag{2}$$

Using these observations and Axioms W1-W5, we can show that

$$P_a(\psi) = \sum_{\{A \in \mathsf{At}(\Sigma) \mid \psi \in A\}} P_a(\phi_A)$$

is provable in PE-PDL. Using this fact, together with I1 and I3, we can show that an a-probability formula $\psi \in \neg\mathrm{FL}(\Sigma)$ is provably equivalent to a formula of the form

$$\sum_{A \in \mathsf{At}(\Sigma)} c_A P_a(\phi_A) \geq b$$

for some appropriate coefficients c_A. Let $\mathbb{P}(A)$ be the set of atoms B, such that $\nvdash \phi_A \to P_a(\phi_B) = 0$.

Now, fix an agent a and a state $A \in \mathsf{At}(\Sigma)$. We describe a set of linear equalities and inequalities corresponding to a and s, over variables of the form $x_{aAA'}$, for $A' \in \mathsf{At}(\Sigma)$. We can think of $x_{aAA'}$ as representing $\mu_a^A(A')$, that is the probability of state A' under agent a's probability distribution at state A. We have one inequality corresponding to every a-probability formula $\psi \in A$. Assume that ψ is equivalent to

$$\sum_{A' \in \mathsf{At}(\Sigma)} c_{A'} P_a(\phi_{A'}) \geq b$$

Notice that exactly one of ψ and $\neg\psi$ is in A. If $\psi \in A$, then the corresponding inequality is

$$\sum_{A' \in \mathsf{At}(\Sigma)} c_{A'} x_{aAA'} \geq b$$

If $\neg\psi \in A$, then the corresponding inequality is

$$\sum_{A' \in \mathsf{At}(\Sigma)} c_{A'} x_{aAA'} < b$$

Further, due to W6, we have the following equalities:

$$x_{aAA'} = 0$$

for $A' \notin \mathbb{P}(A)$, and

$$x_{aAA'} > 0$$

for $A' \in A - \mathbb{P}(A)$. Finally, we have the equality

$$\sum_{A' \in At(\Sigma)} x_{aAA'} = 1$$

As shown in the proofs of [7, Thm. 2.2] and [8, Thm. 4.1], since ϕ_A is consistent, this set of linear equalities and inequalities has a solution $x^*_{a,A,A'}$, for $A' \in \mathsf{At}(\Sigma)$. Set $\mu_{a,A}(A) = x^*_{a,A,A'}$. This is the probability assignment μ that we are looking for. Before we proceed to the truth lemma, we only need to make sure that our model $\mathcal{M}$, thus constructed, satisfies the SDP condition, corresponding to the introduction of Axiom W7 in the logic. This can be easily be checked by inspecting the definition of R_a. Given this, we can assume, without loss of generality, that if AR_aA' then $\mu_{a,A} = \mu_{a,A'}$, since we have that the definition of $\mu_{a,A}$ depends only on the i-probability formulas and their negations at state A. $\square$

Before we prove the truth lemma, we need to establish two important results: an existence lemma for S_π and a theorem which states that $S_\pi \subseteq R_\pi$.

Lemma 4.9 (The Existence Lemma for S_π). *Let A be an atom and let $\langle\pi\rangle\phi$ be a formula in $\neg FL(\Sigma)$. Then $\langle\pi\rangle\phi \in A$ iff there is a B such that $AS_\pi B$ and $\phi \in B$.*

Proof. Following the strategy laid out in [5, p. 244], we set out to construct an appropriate atom B by forcing choices. We begin by enumerating the formulas in the finite set $FL(\Sigma)$ as $\sigma_1, \dots, \sigma_m$ and define B_0 to be $\{\phi\}$. Suppose as an inductive hypothesis that B_n is defined such that $\phi_A \wedge \langle\pi\rangle\phi_{B_n}$ is consistent (where $1 \leq n \leq m$). We get that

$$\vdash \langle\pi\rangle\phi_B \leftrightarrow \langle\pi\rangle\left((\phi_B \wedge \sigma_{n+1}) \vee (\phi_B \wedge \neg\sigma_{n+1})\right)$$

and thus

$$\vdash \langle\pi\rangle\phi_B \leftrightarrow (\langle\pi\rangle(\phi_B \wedge \sigma_{n+1})) \vee (\langle\pi\rangle(\phi_B \wedge \neg\sigma_{n+1}))$$

Therefore, either for $B' = B \cup \{\sigma_{n+1}\}$ or for $B' = B \cup \{\neg\sigma_{n+1}\}$, we have that $\phi_A \wedge \langle\phi_{B'}\rangle$ is consistent. Choose B_{n+1} to be this consistent expansion, and let B_m be B. Then B is the atom we want. □

Lemma 4.10 (Lemma for basic programs)**.** *For all programs $a \in Ag$, $S_a \subseteq R_a$.*

Proof. We need to show that, if AS_aB, then AR_aB, for all $A, B \in At(\Sigma)$. We begin by noting that, since $\phi_A \wedge \langle a\rangle\phi_B$ is consistent, then there exists a maximally consistent set (MCS) Γ such that $\phi_A \wedge \langle a\rangle\phi_B \in \Gamma$. Note that A is the maximal consistent subset of $\neg FL(\Sigma)$ that extends to Γ: $A = \Gamma \cap \neg FL(\Sigma)$. Since $\phi_A \wedge \langle a\rangle\phi_B \in \Gamma$ then $\langle a\rangle\phi_B \in \Gamma$ too. So, there exists a Δ, a maximally consistent set, such that $\Gamma \sim_a \Delta$, where $\sim_a$ is the canonical relation, defined by $A \sim_a B$ iff for all formulas ϕ, $\phi \in B$ implies $\langle a\rangle\phi \in A$. Let $B = \Delta \cap \neg FL(\Sigma)$. Then, we have that $A \sim_a B$. We can show that, by the standard results on canonical models, we have that if the logic includes the S5 axioms, then $\sim_a$ is an equivalence relation.

We prove the following claim:

$$T \sim_a U \leftrightarrow \forall\phi,\ \langle a\rangle\phi \in T \text{ iff } \langle a\rangle\phi \in U$$

Proof.

"⇒" Suppose $T \sim_a U$ and $\langle a\rangle\phi \in T$. Then, by the definition of $\sim_a$, we have that $\langle a\rangle\langle a\rangle\phi \in U$. By Axiom $K4$, we have that $\langle a\rangle\phi \in U$. The other direction follows from the symmetry of $\sim_a$.

"⇐" Suppose that $\forall\phi, \langle a\rangle\phi \in T$ iff $\langle a\rangle\phi \in U$. Let $\psi \in U$. We need to show that $\langle a\rangle\psi \in T$. From $\psi \in U$ and Axiom K3, we then have that $\langle a\rangle\phi \in T$.

□

Therefore, we proved then that if AS_aB then AR_aB. □

Lemma 4.11. *If Σ is finite, then $\neg FL(\Sigma)$ is finite.*

Proof. We skip the proof of this theorem, as it is a straightforward proof by induction. □

Lemma 4.12. *For all programs π, we have that $S_{\pi^*} \subseteq (S_\pi)^*$.*

Proof. Identical to the proof of Lemma 4.87 in [5, p. 244]. □

Theorem 4.13. *For all programs π, $S_\pi \subseteq R_\pi$.*

Proof by induction on the complexity of π.

Base Case: $\pi = a$ is given by the Lemma for basic programs.

Inductive Hypothesis: Assume the claim holds for all programs of complexity lower than π. Now we try to show it for π

Case 1: π is of the form $\pi_1;\pi_2$. Suppose $AS_{\pi_1;\pi_2}B$, that is, $\phi_A \wedge \langle \pi_1;\pi_2\rangle\phi_B$ is consistent. It follows, by Axiom 3, that $\phi_A \wedge \langle\pi_1\rangle\langle\pi_2\rangle\phi_B$ is consistent. By the IH, we get that $AR_{\pi_1}C$ and $CR_{\pi_2}B$. It follows immediately that $AR_{\pi_1;\pi_2}B$.

Case 2: π is of the form $\pi_1 \cup \pi_2$. Similar to Case 1. Omitted here.

Case 3: π is of the form π^*. Suppose $AS_{\pi^*}B$, that is, $\phi_A \wedge \langle\pi^*\rangle\phi_B$ is consistent. Since $S_{\pi^*} \subseteq (S_\pi)^*$, we get that there exists a chain $A = C_0 S_\pi C_1 \ldots C_k = B$, such that, for every pair C_iC_{i+1}, by the IH, if $C_i S_\pi C_{i+1}$ then $C_i R_\pi C_{i+1}$. But then $AR_{\pi^*}B$.

Case 4: π is of the form $\phi?$. Assume $\phi \in \neg FL(\Sigma)$. Suppose $AS_{\phi?}B$. Then $\phi_A \wedge [\phi?]\phi_B$ is consistent. From Axiom 7, using propositional reasoning, we get that $\langle\phi?\rangle\psi \leftrightarrow (\phi \wedge \psi)$. It follows that

$$\phi_A \wedge (\phi \wedge \psi_B) \text{ is consistent}$$

However, it's easy to notice that for any two atoms are mutually exclusive, therefore $\vdash \phi_A \to \neg\phi_B \forall A \neq B$. We can conclude then that $A = B$. Finally, since $\phi_A \wedge \phi$ is consistent and $\phi \in \neg FL(\Sigma)$, we conclude that $\phi \in A$.

□

Before we prove the truth lemma, we need to establish an existence lemma as follows:

Lemma 4.14 (Existence Lemma)**.** *Let A and B be atoms in $At(\Sigma)$ and let $[\pi]\phi \in \neg FL(\Sigma)$. Then if $[\pi]\psi \in A$ and $AR_\pi B$ then $\phi \in B$.*

Proof: induction on the complexity of π.

Base Case: π is a basic program a

We need to show that if AR_aB and $[a]\phi \in A$, then $[a]\phi \in B$. It immediately follows that if $[a]\phi \in A$ and AR_aB then, by the definition of R_a we get $[a]\phi \in B$. By this and the transitivity axiom $[a]\phi \to \phi$, it follows that $\phi \in B$.

Inductive step: Assume the claim holds for all π of a certain complexity and lower.

Case 1: ϕ is of the form $\pi_1 \cup \pi_2$. We have that $AR_{\pi_1 \cup \pi_2}B$ and $[\pi_1 \cup \pi_2]\phi \in A$. By Ax. (iv), we have that $[\pi_1]\phi, [\pi_2]\phi \in A$. Since $AR_{\pi_1 \cup \pi_2}B$ then we have that either $AR_{\pi_1}B$ or $AR_{\pi_2}B$. Applying the IH, we get that in either case $\phi \in B$.

Case 2: π is of the form $\pi_1; \pi_2$. Similar to Case 1. Will omit here.

Case 3: π is of the form $\phi?$. We have that $[\phi?]\psi \in A$ and $AR_{\phi?}B$ and we need to show that $\psi \in B$. By Axiom 7, then $\phi \to \psi \in A$. Further, from $AR_{\phi?}B$, we get that $A = B$ and $\phi \in A$. By an application of modus ponens, we get that $\psi \in A$.

Case 4: π is of the form $[\pi^*]$. In order to prove this, it suffices to show that:

Theorem 4.15. *$\forall \phi$ such that $[\pi^*] \in \neg FL(\Sigma)$, if $[\pi^*]\phi \in A$ and $AR_{\pi^*}B$ then $[\pi^*]\phi \in B$.*

Proof by induction on the length of the path k from A to B: $A = C_0 R_\pi C_1 \dots R_\pi C_k = B$.

Proof.

Base Case: the length of the path $k = 1$. We know that $AR_\pi B$ and $[\pi^*]\phi \in B$. By Axiom 5 and clause 5 of the FL closure, we get that $[\pi][\pi^*]\phi \in A$. Applying the $IH_{\exists Lemma}$, we get that $[\pi^*]\phi \in B$.

Inductive Hypothesis: Assume the claim holds for all lengths lower than k and try to prove it for k.
We have that $[\pi^*]\phi \in A$ and

$$A = C_0 R_\pi C_1 \dots C_{k-1} R_\pi C_k = B$$

By the IH, given that $AR_{\pi^*}C_{k-1}$, we have that $[\pi^*]\phi \in C_{k-1}$. By Axiom 5 and clause 5 of the FL closure, $[\pi][\pi^*]\phi \in C_{k-1}$. Since we also have that $C_{k-1}R_\pi B$, by the $IH_{\exists Lemma}$, we get have $[\pi^*]\phi \in B$.

□

□

Theorem 4.16 (Truth Lemma)**.** *Let $\mathcal{R}$ be a regular PE-PDL model over Σ. For all atoms A and all $\phi \in \neg FL(\Sigma)$, $\mathcal{R}, A \models \phi$ iff $\phi \in A$.*

Proof: Induction on the number of connectives.

Base Case Follows immediately from the definition of the canonical valuation over Σ.

Inductive step

Case 1: The Boolean case.

It follows immediately from Lemma 4.5 above.

Case 2: ϕ is of the form $[a]\psi$

"$\Rightarrow$" $\mathcal{R}, A \models [a]\psi$ then $[a]\psi \in A$.

Proof. Define $\Delta = \{[a]\chi | [a]\chi \in A\} \cup \{\neg[a]\chi | \neg[a]\chi \in A\}$. Then $\Delta \cup \{\neg\psi\}$ is inconsistent. For suppose otherwise. Then $\Delta \cup \{\neg\psi\}$ could be expanded to a MCS $B \in At(\Sigma)$. We have that by construction AR_aB. If $\neg\psi \in B$ then by the IH I get that $B \models \neg\psi$. Since $\mathcal{R}, A \models [a]\psi$ and $AR_aB \Rightarrow B \models \psi$. Contradiction. Then $\vdash \phi_\Delta \to \psi$. By $R_2 \Rightarrow \vdash [a](\phi_\Delta \to \psi)$ and by Ax. 4 and Ax. 5 we get that $\phi_\Delta \to [a]\phi_\Delta$. By $K_2 \Rightarrow \vdash \phi_\Delta \to [a]\psi$. This, together with the fact that $[a]\psi \in \neg\mathrm{FL}(\Sigma)$ and the fact that for $\forall\phi \in \mathrm{FL}(\Sigma)$, either ϕ or its negation is in $A \Rightarrow [a]\psi \in A$. □

"$\Leftarrow$" if $[a]\psi \in A$ then $\mathcal{R}, A \models [a]\psi$.

Proof. Consider $B \in At(\Sigma)$ s.t. AR_aB. Then $[a]\phi \in A \Leftrightarrow [a]\phi \in B$. This together with the assumption imply that $[a]\psi \in B$. By Ax. K3 we know that $[a]\psi \to \psi$ and since $[a]\psi \in \neg\mathrm{FL}(\Sigma)$ then $\psi \in \neg\mathrm{FL}(\Sigma)$, we get that $\psi \in B$. By the IH $\Rightarrow B \models \psi$. This holds for any B s.t. AR_aB. $\Rightarrow A \models [a]\psi$. □

Case 3: $\phi = [\pi_1 \cup \pi_2]\psi$

"$\Rightarrow$" if $\mathcal{R}, A \models [\pi_1 \cup \pi_2]\psi$ then $[\pi_1 \cup \pi_2]\psi \in A$.

Proof. $A \models [\pi_1 \cup \pi_2]\psi \Rightarrow \forall B$ s.t. $AR_{\pi_1}B$ or $AR_{\pi_2}B$ then $B \models \psi \Rightarrow A \models [\pi_1]\psi$ and $A \models [\pi_2]\psi$. By IH we get that $[\pi_1]\psi, [\pi_2]\psi \in A$. By Ax. 4, we have that $[\pi_1]\psi \wedge [\pi_2]\psi \Leftrightarrow [\pi_1 \cup \pi_2]\psi$. Since $[\pi_1 \cup \pi_2]\psi \in \neg\mathrm{FL}(\Sigma) \Rightarrow [\pi_1 \cup \pi_2]\psi \in A$. □

"$\Leftarrow$" if $[\pi_1 \cup \pi_2]\psi \in A$ then $A \models [\pi_1 \cup \pi_2]\psi$.

Proof. Consider $B \in At(\Sigma)$ s.t. $AR_{\pi_1 \cup \pi_2}B$. This means that $AR_{\pi_1}B$ or $AR_{\pi_2}B$. Now $[\pi_1 \cup \pi_2]\psi \in A$. By Ax. 4 $[\pi_1 \cup \pi_2]\psi \leftrightarrow [\pi_1]\psi \wedge [\pi_2]\psi \Rightarrow [\pi_1]\psi, [\pi_2]\psi \in A \Rightarrow$ by IH $A \models [\pi_1]\psi \wedge [\pi_2]\psi$. Since B is s.t. $AR_{\pi_1}B$ or $AR_{\pi_2}B$ then $B \models \phi \Rightarrow A \models [\pi_1 \cup \pi_2]\psi$ □

Case 4: ϕ is of the form $[\pi_1; \pi_2]\psi$

"$\Rightarrow$" $\mathcal{R}, A \models [\pi_1; \pi_2]\psi$ then $[\pi_1; \pi_2]\psi \in A$

Proof. From $\mathcal{R}, A \models [\pi_1; \pi_2]\psi$ we get that for $\forall C, \forall B \in At(\Sigma)$ s.t. $AR_{\pi_1}C$ and $CR_{\pi_2}B$, then $B \models \psi$. It follows that $C \models [\pi_2]\psi$ for any C s.t. $AR_{\pi_1}C \Rightarrow A \models [\pi_1][\pi_2]\psi$. By IH $\Rightarrow [\pi_1][\pi_2]\psi \in A$. By Ax. 3 $\Rightarrow [\pi_1; \pi_2]\psi \in A$, since $[\pi_1; \pi_2]\psi \in \neg\mathrm{FL}(\Sigma)$. □

"$\Leftarrow$" if $[\pi_1; \pi_2]\psi \in A \Rightarrow A \models [\pi_1; \pi_2]\psi$

Proof. We need to show that if $[\pi_1;\pi_2]\psi \in A$ then for $\forall C, B \in At(\Sigma)$ s.t. $AR_{\pi_1}C$ and $CR_{\pi_2}B$ then $B \models \psi$. By Ax. 3 and by clause 1, if $[\pi_1;\pi_2]\psi \in A \Rightarrow [\pi_1][\pi_2]\psi \in A$. By IH $\Rightarrow A \models [\pi_1][\pi_2]\psi \Rightarrow$ for $\forall C$ s.t. $AR_{\pi_1}C$ we have $\mathcal{R}, C \models [\pi_2]\psi \Rightarrow \forall B$ s.t. $CR_{\pi_2}B \Rightarrow B \models \psi \Rightarrow A \models [\pi_1;\pi_2]\psi$ □

Case 5: ϕ is of the form $\phi?$

"$\Rightarrow$" Assume $A \models [\phi?]\psi$. We need to show that $[\phi?]\psi \in A$. By Axiom 6, $A \models \phi \to \psi$. By the IH, we get that $\phi \to \psi \in A$. By Axiom 7, $[\phi?]\psi \in A$.

"$\Leftarrow$"Assume $[\phi?]\psi \in A$. Show that $A \models [\phi?]\psi$. By Axiom 7, we get that $\phi \to \psi \in A$. By the IH, we have that $A \models \phi \to \psi$. By Axiom 7,we get that $A \models [\phi?]\psi$.

Case 6: ϕ is of the form $[\pi^*]\psi$

"$\Leftarrow$" if $[\pi^*]\psi \in A \Rightarrow A \models [\pi^*]\psi$.

From the Existence Lemma and an application of Axiom 5, it follows immediately that $\forall \phi s.t.[\pi]\phi \in \neg FL(\Sigma)$, if $AR_{\pi^*}B$ and $[\pi^*]\phi \in A$ then $\phi \in B$.

"$\Rightarrow$" if $A \models [\pi^*]\psi \Rightarrow [\pi^*]\psi \in A$.

Proof. By contraposition, we need to prove that if $\neg[\pi^*]\phi \in A$ then $A \not\models [\pi^*]\psi$. By the Existence Lemma for S_{i^*}, we have that if $\neg[\pi^*]\phi \in A$ then $\exists B \in At(\Sigma)$ such that $AS_{\pi^*}B$, then $\neg\phi \in B$. But we have shown that $S_{\pi^*} \subseteq R_{\pi^*}$, therefore we have that $AR_{\pi^*}B$. By the $IH_{\text{Truth Lemma}}$ and the fact that $\neg\phi \in B$ we have that $B \models \neg\phi$. Therefore $A \not\models [\pi^*]\phi$. □

Case 7: ϕ is of the form $\sum_{j=1}^{k} \alpha_j P_a(\phi_j) \geq \beta$. By Lemma 4.8, $\sum_{j=1}^{k} \alpha_j P_a(\phi_j) \geq \beta \in A$ if and only if $\sum_{j=1}^{k} \sum_{\{B \in \mathsf{At}, \vdash B \to \phi_j\}} \alpha_j \mu_{a,A}(B) \geq \beta$ if and only if $\mathcal{R}, A \models \sum_{j=1}^{k} \sum_{\{B \in \mathsf{At}, \vdash B \to \phi_j\}} \alpha_j \mu a, A(B) \geq \beta$ if and only if $\mathcal{R}, A \models \sum_{j=1}^{k} \alpha_j P_a(\phi_j) \geq \beta$. □

Lemma 4.17. *If $\mathcal{R} = (\mathsf{At}(\Sigma), \{R_\pi^\Sigma\}_{\pi \in \Pi}, \mu^\Sigma, V^\Sigma)$ is a regular PE-PDL structure over Σ then $\mathcal{B} = (\mathsf{At}(\Sigma), \{R_a^\Sigma\}_{a \in Ag}, \mu^\Sigma, V^\Sigma)$ is a Bayesian Kripke model.*

Proof. We must show that SDP, CONS, CAUT, and PROB are satisfied by $\mathcal{R}$ and hence $\mathcal{B}$. That PROB is satisfied follows immediately from Lemma 4.8. The case for SDP was addressed in the proof of Lemma 4.8. For CONS and CAUT, we note that by the construction of the μ_i^A, we ensured that the support of μ_i^A was $\mathbb{P}(A)$ consisting of all atoms B, such that $\nvdash \phi_A \to P_i(\phi_B) = 0$. We thus wish to show that

$$AR_iB \quad \Leftrightarrow \quad \nvdash \phi_A \to (P_i(\phi_B) = 0)\,.$$

First observe that by W6 (and propositional reasoning),

$$\nvdash \phi_A \to [i](\neg\phi_B) \quad \Leftrightarrow \quad \nvdash \phi_A \to (P_i(\phi_B) = 0)\,.$$

It remains to show that

$$AR_iB \quad \Leftrightarrow \quad \nvdash \phi_A \to [i](\neg\phi_B).$$

The right to left direction follows almost directly from Lemma 4.10, since it is immediate from the definition of S_i that AS_iB is equivalent to $\nvdash \phi_A \to [i](\neg\phi_B)$. But here is another proof. Suppose that it is not the case that AR_iB. Then there exists a $[i]\psi$ that is in exactly one of A or B. If B, we have that $\vdash \phi_A \to \neg[i]\psi$. Then by K3 and modus ponens, we have $\vdash \phi_A \to [i]\neg[i]\psi$. As $\vdash \neg[i]\psi \to \phi_B$ (by propositional reasoning), we have that $\vdash \psi_A \to [i]\neg\phi_B$. If $[i]\psi \in A$, then we have that $\vdash \psi_A \to [i]\psi$, and by K2 and modus ponens, $\vdash \psi_A \to [i][i]\psi$. As $\vdash \neg\neg[i]\psi \in\to \phi_B$, we have that $\vdash \psi_A \to [i]\neg\phi_B$.

The left-to-right direction follows from the Truth Lemma (Lemma 4.16). Suppose that AR_iB. Then by the truth lemma (Lemma 4.16), $\mathcal{R}, B \models \psi$ for every $\psi \in B$. Then $\mathcal{R}, B \models \phi_B$ and hence $\mathcal{R}, A \models \phi_A \wedge \langle i\rangle\phi_B$. As $\phi_A \wedge \langle i\rangle\phi_B$ is satisfiable, it must be consistent. Hence $\nvdash \phi_A \to [i]\neg\phi_B$. □

Theorem 4.18 (Weak completeness of PE-PDL)**.** *PE-PDL is weakly complete with respect to the class of all Bayesian Kripke frames.*

The reduction axioms presented in Figure 11 determine a translation procedure, for reducing the $\mathcal{L}_{\text{PLCC}}$-formulas into $\mathcal{L}_{\text{PE-PDL}}$ formulas.

Definition 4.19 (Translation)**.** The function t takes a formula from the language of $\mathcal{L}_{\text{PLCC}}$ and yields a formula in the language of $\mathcal{L}_{\text{PE-PDL}}$.

$$\begin{aligned}
t(T) &= T & r(a) &= a\\
t(p) &= p & r(B) &= B\\
t(\neg\phi) &= \neg t(\phi) & r(?\phi) &= ?t(\phi)\\
t(\phi_1 \wedge \phi_2) &= t(\phi_1) \wedge t(\phi_2) & r(\pi_1;\pi_2) &= r(\pi_1); r(\pi_2)\\
t([\pi]\phi) &= [r(\pi)]t(\phi) & r(\pi_1 \cup \pi_2) &= r(\pi_1) \cup r(\pi_2)\\
t([e]T) &= T & r(\pi^*) &= (r(\pi))^*\\
t([e]p) &= t(\text{pre}(e)) \to t(\text{sub}(e)(p)) & &\\
t([e]\neg\phi) &= t(\text{pre}(e)) \to \neg t([e]\phi) & &\\
t([e](\phi_1 \wedge \phi_2)) &= t([e]\phi_1) \wedge t([e]\phi_2) & &\\
t([e_i][\pi]\phi) &= \bigwedge_{j=0}^{m-1} [T_{ij}(r(\pi))]t([e_j]\phi) & &\\
t([e][e']\phi) &= t([e]t([e']\phi)) & &\\
t\left(\sum_{1\le h\le k} \alpha_h \cdot P_a(\psi_h) \ge \beta\right) &= \sum_{1\le h\le k} \alpha_h \cdot P_a(t(\psi_h)) \ge \beta & &\\
t\left([e]\sum_{1\le h\le k} \alpha_h \cdot P_a(\psi_h) \ge \beta\right) &= t(\text{pre}(e)) \to t(C > D) & &
\end{aligned}$$

where the letters in the last line stand for

$$C = \sum_{\substack{1 \leq h \leq k \\ \phi_i \in \Phi \\ f \sim_a e}} \alpha_h \cdot \mathrm{pre}_a(f \mid \phi_i) \cdot P_a(t(\phi_i \wedge [f]\psi_h))$$

$$D = \sum_{\substack{\phi_i \in \Phi \\ f \sim_a e}} \beta \cdot \mathrm{pre}_a(f \mid \phi_i) \cdot P_a(t(\phi_i))$$

Finally, we have that:

Theorem 4.20 (Completeness of PLCC)**.** *For any ϕ, a formula of the language $\mathcal{L}_{PLCC}$, we have that:*

$$\models \phi \text{ iff } \vdash \phi$$

Proof. Given the completeness of the static language PE-PDL $\mathcal{L}_{\text{PE-PDL}}$, and the translation procedure above, which ensures every formula in the language of PLCC is equivalent to a formula in the language of PE-PDL, the result follows immediately. □

Theorem 4.21 (Decidability and Strong Finite Model Property)**.** *The satisfiability problem for PLCC is decidable. Moreover, there exists a computable function f from formulas to natural numbers, such that every satisfiable formula ϕ has a model of size at most $f(\phi)$.*

Proof. Given a consistent formula φ in PLCC, one can translate φ to in a provably equivalent way to a formula $t(\varphi)$ in PE-PDL. The length of the resulting formula is bounded by the size of φ. if Σ is the set of sub formulas of $t(\varphi)$, then $\neg FL(\Sigma)$ can be computed, and its size is bounded by the length of φ. The size of the canonical model is thus bounded by 2^n where n is the size of $\neg FL(\Sigma)$. □

Remark 4.22 (Expressivity and Succinctness)**.** PLCC and PE-PDL are equally expressive (although PLCC is much more succinct).

5 Conclusions

The Probabilistic Logic of Communication and Change (PLCC) introduced in this paper provides a unified framework for reasoning about subjective probabilities, levels of mutual knowledge and complex interactive scenarios involving changes affecting both the facts of the world and the agents' information states. We applied this logic to fully capture the higher-level reasoning involved in an informational cascade: the Urn Example.

Conceptually, the importance of our analysis of cascades comes from the fact that PLCC incorporates, not only the standard Bayesian rationality assumptions, but also meta-rationality and higher levels of reflection (via the availability of arbitrary levels of mutual knowledge, including knowledge of the long-term epistemic protocol, as encoded in our event model). So our epistemic analysis of this example shows that, contrary to some authors' opinions, *reflection and higher-order reasoning cannot in general prevent informational cascades*: even if the agents are aware of the cascade, it is still rational for them to continue engaging in it.

But our analysis of the Urn scenario is just one example of the interesting applications of our logic. Beyond this particular example, and even beyond the issue of informational cascades, we think that our logic has broader potential for applications in Game Theory and Social Epistemology, and can be used to spot hidden assumptions behind ordinary economic or multi-agent reasoning, while the axioms and inference rules can also help to analyse such assumptions.

References

[1] Baltag, A., Christoff, Z., Hansen, J. U. and Smets, S.: 2013, *Logical Models of Informational Cascades*, Logic Across the University: Foundations and Applications, College Publications, pp. 405–432.

[2] Banerjee, A. V.: 1992, A Simple Model of Herd Behavior, *The Quarterly Journal of Economics* **107**(3), 797–817.

[3] van Benthem, J., Gerbrandy, J. and Kooi, B. P.: 2009, Dynamic update with probabilities, *Studia Logica* **93**(1), 67–96.

[4] van Benthem, J., van Eijck, J. and Kooi, B.: 2006, Logics of communication and change, *Information and Computation* **204**(11), 1620–1662.

[5] Blackburn, P., de Rijke, M. and Venema, Y.: 2001, *Modal Logic*, Cambridge University Press, New York, NY, USA.

[6] van Eijck, J. and Schwarzentruber, F.: 2014, *Epistemic Probability Logic Simplified*, Manuscript.

[7] Fagin, R., Halpern, J. Y. and Megiddo, N.: 1990, A logic for reasoning about probabilities, *Inf. Comput.* **87**(1-2), 78–128.

[8] Fagin, R. and Halpern, J. Y.: 1994, Reasoning about knowledge and probability, *Journal of the ACM* **41**(2), 340–367.

[9] Rendsvig, R. K.: 2013, Aggregated beliefs and informational cascades, *in* D. Grossi, O. Roy and H. Huang (eds), *Logic, Rationality, and Interaction*, Vol. 8196 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, pp. 337–341.