

Conditional Belief, Knowledge and Probability

Jan van Eijck
CWI and ILLC
Amsterdam, Netherlands
jan.van.eijck@cwi.nl

Kai Li*
Peking University
Beijing, China
CWI
Amsterdam, Netherlands
likaiedemon@gmail.com

A natural way to represent beliefs and the process of updating beliefs is presented by Bayesian probability theory, where belief of an agent a in P can be interpreted as a considering that P is more probable than not P . This paper attempts to get at the core logical notion underlying this.

The paper presents a sound and complete neighbourhood logic for conditional belief and knowledge, and traces the connections with probabilistic logics of belief and knowledge. The key notion in this paper is that of an agent a believing P conditionally on having information Q , where it is assumed that Q is compatible with what a knows.

Conditional neighbourhood logic can be viewed as a core system for reasoning about subjective plausibility that is not yet committed to an interpretation in terms of numerical probability. Indeed, every weighted Kripke model gives rise to a conditional neighbourhood model, but not vice versa. We show that our calculus for conditional neighbourhood logic is sound but not complete for weighted Kripke models. Next, we show how to extend the calculus to get completeness for the class of weighted Kripke models.

Neighbourhood models for conditional belief are closed under model restriction (public announcement update), while earlier neighbourhood models for belief as ‘willingness to bet’ were not. Therefore the logic we present improves on earlier neighbourhood logics for belief and knowledge. We present complete calculi for public announcement and for publicly revealing the truth value of propositions using reduction axioms. The reductions show that adding these announcement operators to the language does not increase expressive power.

1 Introduction

This paper aims at isolating a core logic of rational belief and belief update that is compatible with the Bayesian picture of rational inference [11], but that is more general, in the sense that it does not force epistemic weight models (that is, models with fixed subjective probability measures, for each agent) on us.

Epistemic neighbourhood models, as defined in [7], represent belief as truth in a neighbourhood, where the neighbourhoods for an agent are subsets of the current knowledge cell of that agent. Intuitively, a neighbourhood lists those propositions compatible with what the agent knows that the agent considers as more likely than their complements. Since we intend to use neighbourhood semantics to model a certain kind of belief, it is natural to study belief updates. However, there is an annoying obstacle for updates even in a very simple case: public announcement.

*Kai Li thanks the China Scholarship Council (CSC) for a visiting grant to CWI, and CWI for hospitality during academic year 2016/2017. We wish to thank Dick de Jongh for useful advice, and Malvin Gattinger for his help and advice when we were organizing ideas of this paper. We are also grateful to 3 anonymous reviewers for their comments on an earlier version of the paper.

Intuitively, a public announcement would make an agent restrict his/her belief to the announced case. A natural way to implement this is by restricting every belief-neighbourhood to ϕ -worlds after announcing ϕ . The following example shows that this does not work, because this kind of update does not preserve reasonable neighbourhood conditions. Suppose Alice, somewhat irrationally, believes that ticket t she has bought is the winning ticket in a lottery that she knows has 10,000 tickets. Let n represent the world where ticket n is winning (we assume that this is a single winner lottery, and that Alice knows this). Then Alice's belief is given by a neighbourhood model with

$$N_a(n) = \{X \subseteq \{0000, \dots, 9999\} \mid t \in X\},$$

for all $n \in \{0000, \dots, 9999\}$. Note that Alice's belief does not depend on the world she is in: if n, m are different worlds, then $N_a(n) = N_a(m)$.

Assume Alice gets the information that some ticket v different from the ticket t that she bought is the winning ticket. Let p be such that $V(p) = \{v\}$. Then updating with p leads to an updated model with world set $\{v\}$, and with $N'_a(v) = \{X \cap \{v\} \mid X \in N_a(v)\} = \{\emptyset, \{v\}\}$. However, $N'_a(v)$ is not a neighbourhood function, for it contradicts the condition that beliefs should be consistent (different from \emptyset).

The rest of the paper is structured as follows. Section 2 introduces conditional neighbourhood semantics as an enrichment of epistemic neighbourhood semantics, and presents a complete calculus for it. In Section 3 we show that this calculus is sound but not complete for epistemic weight models, and next, that our language is expressive enough to allow an extension to a complete system for weight models. Section 4 shows that conditional neighborhood models are an excellent starting point for an extension with public announcement update operators. Section 5 traces connections with the literature, lists further work, and concludes.

2 Conditional Neighbourhood Semantics

Epistemic neighbourhood models are defined in [7] as epistemic models $\mathfrak{M} = (W, \sim, V)$ with a neighbourhood function N added to them. The neighbourhood function assigns to each agent a and each world $w \in W$ a neighbourhood $N_a(w)$ that consists of the set of propositions that agent a believes in w .

Intuitively each element in neighbourhood $N_a(w)$ represents a belief agent a holds. Usually belief is bolder than knowledge. Indeed, most people believe many things of which they are not sure. If we equate certainty with knowledge, then this means that any belief of agent a should be a subset of agent a 's current *knowledge cell*, i.e., the set $[w]_a = \{u \in W \mid w \sim_a u\}$. It follows that each proposition in neighbourhood $N_a(w)$ is a subset of $[w]_a$. Thus it is natural (in the framework of epistemic modal logic) to assume the following neighbourhood conditions:

Monotonicity If an agent believes X , and knows that X entails Y , then the agent believes Y .

No-inconsistency An agent does not hold an inconsistent belief.

Determinacy An agent do not believe both a proposition and its complement.

However as is illustrated by Alice's Lottery example from the introductory section, public announcement update do not preserve the **No-inconsistency** condition. In order to overcome this problem, we propose to enrich neighbourhood functions N with an extra parameter for propositions. In other words, instead of focusing on what agents believe, we turn our attention to what agents would believe under some assumption. Following this intuition, for each proposition X , $N_a^w(X)$ is a set of propositions such that each of these propositions represents a belief agent a holds at state w when supposing X . In this

paper, we are interested in beliefs as ‘willingness to bet’, i.e., an agent believes a proposition Y supposing X if the agent considers $Y \cap X$ more likely to be true than its complement conditioned by X , namely $\neg Y \cap X$. We also assume the following postulate:

Equivalence of Conditions If an agent knows that two conditions are equivalent, then the agent’s beliefs are the same under both conditions.

For non-equivalent conditions, on the other hand, conditional beliefs may vary.

Assume p ranges over a set of proposition letters P , and a over a finite set of agents A . The language for conditional neighbourhood logic \mathcal{L}_{CN} is given by the following BNF definition:

$$\phi ::= \top \mid p \mid \neg\phi \mid (\phi \wedge \phi) \mid B_a(\phi, \phi)$$

$B_a(\phi, \psi)$ can be read as “assuming ϕ , agent a believes (is willing to bet) ψ ”. $\perp, \vee, \rightarrow, \leftrightarrow$ are defined as usual. Somewhat arbitrarily, we assume that conditioning with information that contradicts what the agent knows (is certain of) will cause an agent to believe nothing anymore. This means we can define knowledge in terms of conditional belief, as follows. Use $K_a\phi$ for $\neg B_a(\neg\phi, \top)$ (which can be read as “ $\neg\phi$ contradicts what agent a is certain of”) and $\check{K}_a\phi$ for $\neg K_a\neg\phi$.

Consider Alice’s lottery situation again. Alice knows there are 10,000 lottery tickets numbered 0000 through 9999. Alice believes ticket t is winning (and buys it). Let n represent the world where ticket n is winning. Then Alice’s belief is given by a conditional neighbourhood model with $N_a^w(X) = \{Y \subseteq X \mid t \in Y\}$ if $t \in X$, and $N_a^w(X) = \{Y \subseteq X \mid |Y| > \frac{1}{2}|X|\}$ if $t \notin X$. Now Alice receives the information that $v \neq t$ is the winning ticket. Then $v = w$, the updated model has universe $\{v\}$, and Alice updates her belief to N' with $N'_a(\{v\}) = \{\{v\}\}$. In the updated model, Alice knows that v is the winning ticket.

Definition 1. Let Ag be a finite set of agents. A conditional neighbourhood model \mathfrak{M} is a tuple (W, N, V) where:

- W is a non-empty set of worlds;
- $N : Ag \times W \times \mathcal{P}W \rightarrow \mathcal{P}\mathcal{P}W$ is a function that assigns to every agent $a \in Ag$, every world $w \in W$ and set of worlds $X \subseteq W$ a collection $N_a^w(X)$ of sets of worlds—each such set called a neighbourhood of X —subject to the following conditions, where

$$[w]_a = \{v \in W \mid \forall X \subseteq W, N_a^w(X) = N_a^v(X)\} :$$

- (c) $\forall Y \in N_a^w(X) : Y \subseteq X \cap [w]_a$.
- (ec) $\forall Y \subseteq W : \text{if } X \cap [w]_a = Y \cap [w]_a, \text{ then } N_a^w(X) = N_a^w(Y)$.
- (d) $\forall Y \in N_a^w(X), X \cap [w]_a - Y \notin N_a^w(X)$.
- (sc) $\forall Y, Z \subseteq X \cap [w]_a : \text{if } X \cap [w]_a - Y \notin N_a^w(X) \text{ and } Y \subsetneq Z, \text{ then } Z \in N_a^w(X)$.
- V is a valuation.

We call N a neighbourhood function; a neighbourhood $N_a^w(X)$ for agent a in w , conditioned by X is a set of propositions each of which agent a believes more likely to be true than its complement.

Property (c) expresses that what is believed is also known; (ec) expresses **equivalence of conditions**; (d) expresses **determinacy**; (sc) expresses a form of “strong commitment”: if the agent does not believe the complement of Y then she must believe any weaker Z implied by Y . It can be proved (see Appendix A, Lemma 11) that these conditions together imply that any conditional neighbourhood model $\mathfrak{M} = (W, N, V)$ also satisfies the following, for any $a \in A, w \in W, X \subseteq W$:

- (m) $\forall Y \subseteq Z \subseteq X \cap [w]_a : \text{if } Y \in N_a^w(X), \text{ then } Z \in N_a^w(X)$;

(ni) $\emptyset \notin N_a^w(X)$;

(n)* if $X \cap [w]_a \neq \emptyset$, then $X \cap [w]_a \in N_a^w(X)$;

(\emptyset) if $X \cap [w]_a = \emptyset$, then $N_a^w(X) = \emptyset$;

where (m) and (ni) expresses **monotonicity** and **no-inconsistency** respectively. (\emptyset) expresses that conditioning with information that contradicts what the agent knows will cause an agent to believe nothing anymore. Note that (\emptyset) reflects our definition for K -operators $K_a\phi ::= \neg B_a(\neg\phi, \top)$.

Let $\mathfrak{M} = (W, N, V)$ be a conditional neighbourhood model, let $w \in W$. Then the key clause of the truth definition is given by:

$$\mathfrak{M}, w \models B_a(\phi, \psi) \quad \text{iff} \quad \text{for some } Y \in N_a^w(\llbracket \phi \rrbracket_{\mathfrak{M}}), Y \subseteq \llbracket \psi \rrbracket_{\mathfrak{M}}$$

where $\llbracket \phi \rrbracket_{\mathfrak{M}} = \{w \in W \mid \mathfrak{M}, w \models \phi\}$. Because of (m), we can prove that

$$\mathfrak{M}, w \models B_a(\phi, \psi) \quad \text{iff} \quad \{v \in \llbracket \phi \rrbracket_{\mathfrak{M}} \cap [w]_a \mid \mathfrak{M}, v \models \psi\} \in N_a^w(\llbracket \phi \rrbracket_{\mathfrak{M}}).$$

It is worth noting that by (ni), $B_a(\phi, \perp)$ will always be invalid for any agent a and any formula ϕ .

Note that conditional neighborhood models do not have epistemic relations \sim_a . However, such relations can be introduced by the neighbourhood function as follows: for each $a \in Ag$, $\sim_a \subseteq W \times W$ is a relation such that

$$\forall w, v \in W, w \sim_a v \quad \text{iff} \quad \forall X \subseteq W, N_a^w(X) = N_a^v(X).$$

Then $\mathfrak{M}, w \models K_a\phi$ iff for each $v \sim_a w$, $\mathfrak{M}, v \models \phi$.

It can be proved that the version of conditional neighbourhood models with epistemic relations \sim_a is equivalent to the version without (see Appendix A). Such equivalence is guaranteed by properties (n)*, (\emptyset) and another property which can be found in [1] (however they use neighbourhoods instead of conditional neighbourhoods for beliefs):

(a) $\forall v \in [w]_a : N_a^w(X) = N_a^v(X)$,

which states that if a agent cannot distinguish two worlds, then the agent holds the same beliefs on either of these two worlds. Thus we do not differentiate conditional neighbourhood models with or without such relations.

In Figure 1, axiom (D) guarantees the truth of neighbourhood condition (d), (EC) would correspond to (ec), (M) to (ec), (M) to (m), (C) to (c) and (SC) to (sc).

Theorem 2. *The CN calculus for Conditional Neighbourhood logic given in Figure 1 is sound and complete for conditional neighbourhood models.*

Proof. See Appendix B. □

Note that the calculus does not have 4 and 5 for K ; this is because these principles are derivable from (5B) and (4B).

For an example of an interesting principle that can already be proved in the CN calculus, consider the following. Suppose we have a biased coin with unknown bias, and we want to use it to simulate a fair coin. Then we can use a recipe first proposed by John von Neumann [13]: toss the coin twice. If the two outcomes are not the same, use the first result; if not, forget the outcomes and repeat the procedure. Why does this work? Because we can assume that two tosses of the same coin have the same likelihood of showing heads, even if the coin is biased. We can express subjective likelihood comparison in our language. See Figure 2, where we use α for “the first toss comes up with heads”, and β for “the second toss comes up with heads”. This hinges on the following principle:

- (Taut) All instances of propositional tautologies
 (Dist-K) $K_a(\phi \rightarrow \psi) \rightarrow K_a\phi \rightarrow K_a\psi$
 (T) $K_a\phi \rightarrow \phi$
 (5B) $B_a(\phi, \psi) \rightarrow K_aB_a(\phi, \psi)$
 (4B) $\neg B_a(\phi, \psi) \rightarrow K_a\neg B_a(\phi, \psi)$
 (D) $B_a(\phi, \psi) \rightarrow \neg B_a(\phi, \neg\psi)$
 (EC) $K_a(\phi \leftrightarrow \psi) \rightarrow B_a(\phi, \chi) \rightarrow B_a(\psi, \chi)$
 (M) $K_a(\phi \rightarrow \psi) \rightarrow B_a(\chi, \phi) \rightarrow B_a(\chi, \psi)$
 (C) $B_a(\phi, \psi) \rightarrow B_a(\phi, \phi \wedge \psi)$
 (SC) $\check{B}_a\phi \wedge \check{K}_a(\neg\phi \wedge \psi) \rightarrow B_a(\phi \vee \psi)$

Rules:

$$\frac{\phi \rightarrow \psi \quad \phi}{\psi} \text{ (MP)} \quad \frac{\phi}{K_a\phi} \text{ (Nec-K)}$$

Figure 1: The CN Calculus for Conditional Neighbourhood Logic

If α and β have the same likelihood, then $\alpha \wedge \neg\beta$ and $\neg\alpha \wedge \beta$ should also have the same likelihood, and vice versa.

Notice that $B_a(\alpha \leftrightarrow \neg\beta, \alpha)$ expresses that agent a considers $\alpha \wedge \neg\beta$ more likely than $\neg\alpha \wedge \beta$. And it follows from the comparison principle in Figure 2 that this is equivalent to: a considers α more likely than β . From now on, let $\alpha \succ_a \beta$ abbreviate $B_a(\alpha \leftrightarrow \neg\beta, \alpha)$.

By axioms M and C , $B_a(\alpha \leftrightarrow \neg\beta, \beta) \leftrightarrow B_a(\alpha \leftrightarrow \neg\beta, \neg\alpha)$, and by axiom D , $B_a(\alpha \leftrightarrow \neg\beta, \alpha) \rightarrow \neg B_a(\alpha \leftrightarrow \neg\beta, \neg\alpha)$. Therefore, $B_a(\alpha \leftrightarrow \neg\beta, \alpha) \rightarrow \neg B_a(\alpha \leftrightarrow \neg\beta, \beta)$ is provable in the CN calculus. Using the abbreviation: $\alpha \succ_a \beta \rightarrow \neg\beta \succ_a \alpha$. We therefore have three mutually exclusive cases:

- $\alpha \succ_a \beta$.
- $\beta \succ_a \alpha$.
- $\neg\alpha \succ_a \beta \wedge \neg\beta \succ_a \alpha$.

Agreeing to abbreviate the third case as $\alpha \approx_a \beta$, we get the following totality principle.

Totality $(\alpha \succ_a \beta) \vee (\beta \succ_a \alpha) \vee (\alpha \approx_a \beta)$.

Next, we define $\alpha \succsim_a \beta$ by $\neg\beta \succ_a \alpha$. This abbreviation gives:

Refl Totality $(\alpha \succsim_a \beta) \vee (\beta \succsim_a \alpha)$.

We can also connect to logic languages concerning probability that do not have knowledge operators K_a but instead use $\succsim_a \top$ (for instance [9] and [10]), by deriving that $(\alpha \succsim_a \top) \leftrightarrow K_a\alpha$.

3 Incompleteness and Completeness for Epistemic Weight Models

In this Section we interpret \mathcal{L}_{CN} in epistemic weight models, and give an incompleteness and a completeness result.




Figure 2: **Comparison principle:** $\alpha \wedge \neg\beta \succ \beta \wedge \neg\alpha$ iff $\alpha \succ \beta$

Definition 3. An *Epistemic Weight Model* for agents Ag and basic propositions P is a tuple $\mathfrak{M} = (W, R, L, V)$ where W is a non-empty countable set of worlds, R assigns to every agent $a \in Ag$ an equivalence relation \sim_a on W , L assigns to every $a \in Ag$ a function \mathbb{L}_a from W to \mathbb{Q}^+ (the positive rationals), subject to the following boundedness condition (*).

$$\forall a \in Ag \forall w \in W \sum_{u \in [w]_a} \mathbb{L}_a(u) < \infty.$$

where $[w]_a$ is the cell of w in the partition induced by \sim_a . V assigns to every $w \in W$ a subset of P .

We can interpret conditional belief sentences in these models. Let \mathfrak{M} be a weight model, and let $\llbracket \phi \rrbracket_{\mathfrak{M}} = \{w \in W \mid \mathfrak{M}, w \models \phi\}$. Let w be a world of \mathfrak{M} . Then

$$\mathfrak{M}, w \models B_a(\phi, \psi) \text{ iff } \mathbb{L}_a([w]_a \cap \llbracket \phi \wedge \psi \rrbracket_{\mathfrak{M}}) > \mathbb{L}_a([w]_a \cap \llbracket \phi \wedge \neg\psi \rrbracket_{\mathfrak{M}}),$$

and $\mathfrak{M}, w \models K_a\phi$ iff for all $v \in [w]_a$, $\mathfrak{M}, v \models \phi$. One easily checks that the axioms of the CN calculus are true for this interpretation, so we have:

Theorem 4. *The CN calculus is sound for epistemic weight models.*

To see that we do not have completeness, observe that we can express Savage's *Sure Thing Principle* in our language. In Savage's example this is about action. If an agent would do a thing if he would know ϕ , and would do the same thing if he would know $\neg\phi$, then he should do the thing in any case:

A businessman contemplates buying a certain piece of property. He considers the outcome of the next presidential election relevant. So, to clarify the matter to himself, he asks whether he would buy if he knew that the Democratic candidate were going to win, and decides that he would. Similarly, he considers whether he would buy if he knew that the Republican candidate were going to win, and again finds that he would. Seeing that he would buy in either event, he decides that he should buy, even though he does not know which event obtains, or will obtain, as we would ordinarily say. Savage, [14, p 21].

The following formula expresses this Sure Thing Principle, not about action but about belief:

$$B_a(\phi, \psi) \wedge B_a(\neg\phi, \psi) \rightarrow B_a(\top, \psi).$$

It is not hard to see that this principle is valid for weight models: if ψ has greater weight than $\neg\psi$ within the ϕ area, and also within the $\neg\phi$ area, then it is a matter of adding these weights to see that ψ has greater weight than $\neg\psi$ in the whole domain. But the Sure Thing Principle is not a validity for neighbourhood models. Let us consider the following urn example modified from Ellsberg's paradox [8].

Example 5. An urn contains 120 balls: 30 red balls, 30 green balls, and 60 yellow or blue balls (in some unknown proportion). A ball x will be pulled out of this urn, and there are 3 pairs of gambles where Alice has to pick her choice:

$G_r : x \text{ is red}$	$G_y : x \text{ is yellow}$
$G_g : x \text{ is green}$	$G_b : x \text{ is blue}$
$G_{rg} : x \text{ is either red or green}$	$G_{yb} : x \text{ is either yellow or blue}$

Alice knows that the likelihood of G_{rg} ($1/2$) is the same as G_{yb} ($1/2$), but is uncertain of the likelihood of G_y and G_b . Alice is ambiguity averse in her beliefs, which means that she is willing to bet G_r against G_y , and G_g against G_b .

To model this example, let $W = \{\text{red, green, yellow, blue}\}$, and let the neighbourhood function be the same for any $w \in W$, with

$$N^w(\{\text{red, yellow}\}) = \{\{\text{red}\}, \{\text{red, yellow}\}\},$$

$$N^w(\{\text{green, blue}\}) = \{\{\text{green}\}, \{\text{green, blue}\}\},$$

$N^w(W)$ contains all subsets with at least 3 worlds. The reason that such model exists is that all our neighbourhood properties do not express anything about connections between different neighbourhoods. Thus, in this model we have $B(G_r \vee G_y, G_r)$, $B(G_g \vee G_b, G_g)$ and $B(\top, \neg(G_r \vee G_g))$ all true in every world, in contradiction with the Sure Thing principle.

Theorem 6. *The CN calculus is incomplete for epistemic weight models.*

Proof. As we have seen, $B(G_r \vee G_y, G_r \vee G_g) \wedge B(G_g \vee G_b, G_r \vee G_g) \rightarrow B(\top, G_r \vee G_g)$ is false in the above neighbourhood counterexample to Sure Thing. So by Theorem 2, Sure Thing cannot be proved in the CN calculus. But Sure Thing is valid in the class of epistemic weighted models. \square

A Complete Calculus for Epistemic Weight Models

[16] and [9] proposed a complete logic QP for epistemic weight models based on a language \mathcal{L}_{QP} given by the following BNF definition¹:

$$\phi ::= \top \mid p \mid \neg\phi \mid (\phi \wedge \phi) \mid \phi \succ_a \phi$$

\succ_a and \approx_a are given as usual. $\succ_a \top$ functions as K_a in epistemic modal logic. The complex formula $\alpha_0 \dots \alpha_m E_a \beta_0 \dots \beta_m$ first proposed by Segerberg in [16] is an abbreviation of the formula expressing that for all worlds w in the evaluated knowledge cell of agent a , the number of α_i among $\alpha_0 \dots \alpha_m$ true in w is the same as those of β_j among $\beta_0 \dots \beta_m$ true in w . QP logic does not assume that every world in an agent's knowledge cell has the same likelihood for all propositions; and is with the following core axioms:

$$(A4) \quad \alpha_0 \alpha_1 \dots \alpha_m E_a \beta_0 \beta_1 \dots \beta_m \wedge (\alpha_0 \succ_a \beta_0) \wedge \dots \wedge (\alpha_{m-1} \succ_a \beta_{m-1}) \rightarrow (\beta_m \succ_a \alpha_m) \text{ for all } m \geq 1.$$

Nevertheless, not only we can express every notion in the probabilistic language \mathcal{L}_{QP} , but using the results of [9] and [15] we can prove the following completeness theorem as well (the proof is a simple adaptation of a proof found in [9], just to observe that every other axioms in [9] except (A4) is provable in CN calculus).

Theorem 7. $CN \oplus (A4)$ is complete for epistemic weight models.

¹We replaced “0” and “1” in [9] with the equivalent notation “ \perp ” and “ \top ” respectively, and extend this language to multi-agent case.

Comparing Expressive Power

In this subsection we compare the expressive power between \mathcal{L}_{CN} and \mathcal{L}_{QP} , restricting our attention to the single-agent case. As is shown in Section 2, we can translate \mathcal{L}_{CN} into \mathcal{L}_{QP} by defining $Tr_1 : \mathcal{L}_{CN} \rightarrow \mathcal{L}_{QP}$ with key case:

$$Tr_1(B(\alpha, \beta)) = Tr_1(\alpha \wedge \beta) \succ Tr_1(\alpha \wedge \neg\beta),$$

which express that the agent considers $\alpha \wedge \beta$ is more likely than $\alpha \wedge \neg\beta$. Likewise we can define the translation Tr_2 from \mathcal{L}_{QP} to \mathcal{L}_{CN} by key case

$$Tr_2(\alpha \succ \beta) = \neg B(Tr_2(\alpha \leftrightarrow \neg\beta), Tr_2(\beta)).$$

It is easy to prove that both translations preserve truth on weight models. As for conditional neighbourhood models, consider the following truth definition for \mathcal{L}_{QP} :

$$\mathfrak{M}, w \models_{qp} \alpha \succ \beta \text{ iff } \llbracket \neg\alpha \wedge \beta \rrbracket_{\mathfrak{M}}^{qp} \notin N^w(\llbracket \alpha \leftrightarrow \neg\beta \rrbracket_{\mathfrak{M}}^{qp}),$$

which is equivalent to (by condition (d) and $\alpha \succ \beta ::= \neg\beta \succ \alpha$),

$$\mathfrak{M}, w \models_{qp} \alpha \succ \beta \text{ iff } \llbracket \alpha \wedge \neg\beta \rrbracket_{\mathfrak{M}}^{qp} \in N^w(\llbracket \alpha \leftrightarrow \neg\beta \rrbracket_{\mathfrak{M}}^{qp}).$$

Such truth condition parallels with our translation Tr_2 . Furthermore, to show that Tr_1 preserves truth value on conditional neighbourhood models, we can prove by induction on construction of \mathcal{L}_{CN} -formulas, and establish the following equivalences:

$$\begin{aligned} \mathfrak{M}, w \models_{qp} Tr_1(B(\alpha, \beta)) & \text{ iff } \mathfrak{M}, w \models_{qp} \alpha' \wedge \beta' \succ \alpha' \wedge \neg\beta' \\ & \text{ iff } \llbracket (\alpha' \wedge \beta') \wedge \neg(\alpha' \wedge \neg\beta') \rrbracket_{\mathfrak{M}}^{qp} \in N^w(\llbracket (\alpha' \wedge \beta') \leftrightarrow \neg(\alpha' \wedge \neg\beta') \rrbracket_{\mathfrak{M}}^{qp}) \\ & \text{ iff } \llbracket \alpha' \wedge \beta' \rrbracket_{\mathfrak{M}}^{qp} \in N^w(\llbracket \alpha' \rrbracket_{\mathfrak{M}}^{qp}) \\ & \text{ iff } \llbracket \beta' \rrbracket_{\mathfrak{M}}^{qp} \in N^w(\llbracket \alpha' \rrbracket_{\mathfrak{M}}^{qp}) \text{ (by condition (n)*)} \\ & \text{ iff } \llbracket \beta \rrbracket_{\mathfrak{M}} \in N^w(\llbracket \alpha \rrbracket_{\mathfrak{M}}) \text{ (by induction hypothesis)} \\ & \text{ iff } \mathfrak{M}, w \models B(\alpha, \beta), \end{aligned}$$

where $\alpha' = Tr_1(\alpha)$ and $\beta' = Tr_1(\beta)$. Therefore we can prove that both Tr_1 and Tr_2 preserve truth value on conditional neighbourhood models.

However we can design models violating the comparison principle of Figure 2. \mathcal{L}_{QP} can differentiate models where the principle holds from those where it does not, while \mathcal{L}_{CN} cannot. A *comparison model* \mathfrak{N} is a triple (W, \succeq, V) where W is a non-empty set of worlds, $\succeq \subseteq \mathcal{P}W \times \mathcal{P}W$ is a relation between propositions, and V is a valuation. Truth definition for \mathcal{L}_{QP} is with the following key clause:

$$\mathfrak{N}, w \models_2 \alpha \succ \beta \text{ iff } \llbracket \alpha \rrbracket_{\mathfrak{N}} \succeq \llbracket \beta \rrbracket_{\mathfrak{N}}.$$

Furthermore the key clause in the truth condition for \mathcal{L}_{CN} is given by:

$$\mathfrak{N}, w \models_1 B(\alpha, \beta) \text{ iff } \llbracket \alpha \wedge \beta \rrbracket_{\mathfrak{N}} \succeq \llbracket \alpha \wedge \neg\beta \rrbracket_{\mathfrak{N}},$$

which parallels with translation Tr_1 at the semantic level.

Let $N_1 = (W, \succeq_1, V)$ and $N_2 = (W, \succeq_2, V)$ be comparison models such that:

1. $W = \{\{p, q\}, \{p\}, \{q\}, \emptyset\}$
2. $\succeq_1 = \{(\{\{p\}, \{p, q\}\}, \{\{q\}, \{p, q\}\}), (\{p\}, \{q\})\}$,

3. $\succeq_2 = \{(\{p\}, \{q\})\}$,
4. $w \in V(r)$ iff $r \in w$.

The only difference between \mathfrak{N}_1 and \mathfrak{N}_2 is that $\llbracket p \rrbracket_{\mathfrak{N}_1} \succeq_1 \llbracket q \rrbracket_{\mathfrak{N}_1}$ but not $\llbracket p \rrbracket_{\mathfrak{N}_2} \succeq_2 \llbracket q \rrbracket_{\mathfrak{N}_2}$. Thus \mathfrak{N}_2 violates the **comparison principle**, and $\neg(p \succcurlyeq q) \wedge (p \wedge \neg q \succcurlyeq \neg p \wedge q)$ is valid on \mathfrak{N}_2 but not on \mathfrak{N}_1 . However we can prove that \mathfrak{N}_1 and \mathfrak{N}_2 satisfy the same set of \mathcal{L}_{CN} -formulas. The crucial fact is to observe that we only use the comparison relation for disjoint propositions for \models_1 . For instance $B(p \leftrightarrow \neg q, p) = Tr_2(p \succcurlyeq q) = Tr_2(p \wedge \neg q \succcurlyeq p \wedge \neg q)$ is valid on both \mathfrak{N}_1 and \mathfrak{N}_2 , because for each $i \in \{1, 2\}$,

$$\begin{aligned} \mathfrak{N}_i, w \models_1 B(p \leftrightarrow \neg q, p) & \quad \text{iff} \quad \llbracket (p \leftrightarrow \neg q) \wedge p \rrbracket_{\mathfrak{N}_i} \succeq_i \\ & \quad \llbracket (p \leftrightarrow \neg q) \wedge \neg p \rrbracket_{\mathfrak{N}_i} \\ & \quad \text{iff} \quad \llbracket p \wedge \neg q \rrbracket_{\mathfrak{N}_i} \succeq_i \llbracket \neg p \wedge q \rrbracket_{\mathfrak{N}_i} \\ & \quad \text{iff} \quad \{p\} \succeq_i \{q\}, \text{ which holds for} \\ & \quad \text{either } i \in \{1, 2\}. \end{aligned}$$

Therefore \mathcal{L}_{QP} is more expressive than \mathcal{L}_{CN} . We conclude \mathcal{L}_{CN} as a core logic for conditional belief as willingness to bet.

4 Public Announcement for Conditional Neighborhood Models

Public announcement update for weight models parallels Bayesian update in probability theory. Public announcement update for probabilistic logic was first treated in [12], and more complicated probabilistic updates were discussed in [2] and [5]. As was mentioned in the introduction, public announcement updates may destroy reasonable neighbourhood conditions. The good news is that conditional neighbourhood models are more well behaved. We propose two ways of public announcement updating: deleting points and cutting links; and show reduction axioms for either of them. These shows that our neighbourhood conditions are some core principles that preserved by Bayesian update.

4.1 Deleting Points

The first approach is the usual one for public announcement update, which is restricting the domain to ϕ -worlds after announcing ϕ . It assumes that only facts (true propositions at the current world) can be publicly announced. Public announcements create common knowledge, but it need not be the case that a fact that gets announced becomes true after the update; Moore sentences are a well-known exception.

The language \mathcal{L}_{PC} is the result of extending our base language \mathcal{L}_{CN} with a public announcement operator:

$$\phi ::= \top \mid p \mid \neg\phi \mid (\phi \wedge \phi) \mid B_a(\phi, \phi) \mid [\phi]\phi$$

If $\mathfrak{M} = (W, N, V)$ is a conditional neighborhood model, \sim is the induced epistemic relation, and ϕ is a formula of the \mathcal{L}_{PC} language, then $\mathfrak{M}^\phi = (W^\phi, {}^\phi N, V^\phi)$ is given by:

- $W^\phi = \{w \in W \mid \mathfrak{M}, w \models \phi\}$
- $w \sim_a^\phi u$ iff $\mathfrak{M}, w \models \phi$, $\mathfrak{M}, u \models \phi$ and $w \sim_a u$
- ${}^\phi N_a^w(X) = \begin{cases} N_a^w(X) & \text{if } X \subseteq W^\phi \text{ and } w \in W^\phi \\ \text{undefined} & \text{otherwise} \end{cases}$
- $V^\phi(p) = V(p) \cap W^\phi$

Example 8. As an example, consider Alice's lottery situation again. Alice knows there are 10,000 lottery tickets numbered 0000 through 9999. Alice believes ticket t is winning (and buys it). Let n represent the world where ticket n is winning. Then Alice's belief is given by a conditional neighborhood model with $N_a^w(X) = \{Y \subseteq X \mid t \in Y\}$ if $t \in X$, and $N_a^w(X) = \{Y \subseteq X \mid |Y| > \frac{1}{2}|X|\}$ if $t \notin X$. Now Alice receives the information that $v \neq t$ is the winning ticket. Then $v = w$, the updated model has universe $\{v\}$, and Alice updates her belief to N' with $N_a^{v'}(\{v\}) = \{\{v\}\}$. The updated model satisfies the conditions for a conditional neighborhood model.

Definition 9. Semantics for \models_{PC} : let $\mathfrak{M} = (W, N, V)$ be a conditional neighborhood model, let $w \in W$.

$$\mathfrak{M}, w \models_{PC} [\phi]\psi \text{ iff } \mathfrak{M}, w \models_{PC} \phi \text{ implies } \mathfrak{M}^\phi, w \models_{PC} \psi.$$

A complete calculus for \models_{PC} consists of the calculus for CN, plus the usual Reduction Axioms of public announcement update for boolean cases and the following key Reduction Axiom (call this system PC):

- $[\phi]B_a(\psi, \chi) \leftrightarrow (\phi \rightarrow B_a(\phi \wedge [\phi]\psi, [\phi]\chi))$

In Appendix C we prove that the PC Calculus is sound and complete w.r.t. \models_{PC} .

4.2 Cutting Links

We can generalize public announcement of facts to public announcement of truth values. In announcing the value of ϕ , it depends on the truth value of ϕ in the actual world whether ϕ or $\neg\phi$ gets announced.

The language $\mathcal{L}_{PC\pm}$ for this kind of update is given by the following BNF:

$$\phi ::= \top \mid p \mid \neg\phi \mid (\phi \wedge \phi) \mid B_a(\phi, \phi) \mid [\pm\phi]\phi$$

To capture the intuition of such updates for conditional neighbourhoods, we use the following mechanism that cuts epistemic relations between ϕ -worlds and $\neg\phi$ -worlds after announcing ϕ .

If $\mathfrak{M} = (W, N, V)$ is a conditional neighborhood model, \sim is the induced epistemic relation, and ϕ is a formula of the $\mathcal{L}_{PC\pm}$ language, then $\mathfrak{M}^{\pm\phi} = (W^{\pm\phi}, {}^{\pm\phi}N, V^{\pm\phi})$ is given by:

- $W^{\pm\phi} = W$
- $\sim^{\pm\phi} = \{(w, v) \in W^2 \mid w \sim_a v \text{ and } \mathfrak{M}, w \models \phi \text{ iff } \mathfrak{M}, v \models \phi\}$
- ${}^{\pm\phi}N_a^w(X) = N_a^w(X \cap [w]_a^{\pm\phi})$
- $V^{\pm\phi} = V$

Definition 10. Semantics for $\models_{PC\pm}$: let $\mathfrak{M} = (W, N, V)$ be a conditional neighborhood model, let $w \in W$.

$$\mathfrak{M}, w \models_{PC\pm} [\phi]\psi \text{ iff } \mathfrak{M}^{\pm\phi}, w \models_{PC\pm} \psi.$$

The system PC \pm for $\models_{PC\pm}$ consists of the calculus for CN, plus the following Reduction Axioms:

- $[\pm\phi]p \leftrightarrow p$
- $[\pm\phi]\neg\psi \leftrightarrow \neg[\pm\phi]\psi$
- $[\pm\phi](\psi \wedge \chi) \leftrightarrow [\pm\phi]\psi \wedge [\pm\phi]\chi$
- $[\pm\phi]B_a(\psi, \chi) \leftrightarrow (\phi \rightarrow B_a(\phi \wedge [\pm\phi]\psi, [\pm\phi]\chi)) \wedge (\neg\phi \rightarrow B_a(\neg\phi \wedge [\pm\phi]\psi, [\pm\phi]\chi))$

Also in Appendix C we prove that the PC \pm Calculus is sound and complete w.r.t. $\models_{PC\pm}$.

5 Conclusions and Future Work

In Section 1, we illustrated that public announcement update may not preserve reasonable neighbourhood conditions. To overcome this problem, we introduced conditional neighbourhood semantics in Section 2. We gave an alternative interpretation for this system in Section 3, and then in Section 4 we gave two flavours of public announcement update for conditional neighborhood semantics. Because public announcement update for epistemic weight models is basically Bayesian update in a logical setting, and because the complete calculus CN for conditional neighbourhood models is a subsystem of a complete probabilistic logic $CN \oplus (A4)$ (as shown in Section 3), our investigations show that CN can be viewed as a core logic of rational belief and belief update that is compatible with the Bayesian picture of inference, but more general. Conditional neighborhood models generalize epistemic weight models, and this generalization creates room for modelling ambiguity aversion in belief as willingness to bet.

In Section 3 we have shown that \mathcal{L}_{CN} has weaker expressive power than \mathcal{L}_{QP} . Our conditional neighbourhood semantics for \mathcal{L}_{CN} allows us to develop a reasoning system CN that is not yet committed to a probabilistic numerical interpretation of belief. This might be a convenient starting point for further investigation of counterfactual reasoning. In Section 2 we have assumed that conditioning with information that contradicts an agent's knowledge will cause the agent to refrain believing anything, but in future work we may relax this constraint, by allowing to visit knowledge cells other than the current one when a neighbourhood function is conditioned with propositions that conflict with current knowledge. A naive way to do so is to incorporate the selection function f for counterfactuals proposed by Stalnaker in [17] in our framework. When a proposition X is disjoint with agent a 's current knowledge cell $[w]_a$, the selection function f would guide us to an X -world $u = f_a(X, w)$, and then let the neighbourhood $N_a^w(X)$ be $N_a^u(X)$, which is the neighbourhood conditioned by X at $[u]_a$.

While subjective conditional beliefs given by neighbourhood functions suggest how agents' beliefs would change by public announcements, further updates like public lies and recovery from lies may allow us to represent further details of agents' beliefs in an objective way. Here, recovery is to free agents from the influence of lies that were accepted as true in the past. These two kinds of updates give us powerful tools to test what a agent would believe after providing each possible piece of information, which in turn would inform us the agent's conditional beliefs or subjective probability. It is future work to compare and combine these two approaches: subjective conditional beliefs informative for belief update and objective belief changes to conditional beliefs.

Neighbourhood structures have also been used to describe the pieces of evidence that agents accept ([6],[3],[4]). In this approach, each proposition in a neighbourhood is interpreted as a piece of evidence, instead of a belief; and because evidences are usually acquired in various situations, such evidence neighbourhoods may have contradictory propositions. Furthermore, beliefs are generated from certain consistent subsets of the evidence neighbourhood. Even though evidence models and conditional neighbourhood models provide different perspectives on belief, we may be able to combine them in future. In one direction, conditional beliefs or even subjective probability could be generated from certain evidence models, while the way evidence is involved in belief formation may provide information about the strengths of the resulting beliefs. In the other direction, our conditional beliefs might serve as prior knowledge for specifying the credence of evidence.

References

- [1] Philippe Balbiani, David Fernández-Duque & Emiliano Lorini (2016): *A Logical Theory of Belief Dynamics for Resource-Bounded Agents*. In: *Proceedings of the 2016 International Con-*

- ference on Autonomous Agents & Multiagent Systems, AAMAS '16, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, pp. 644–652. Available at <http://dl.acm.org/citation.cfm?id=2936924.2937020>.
- [2] Johan van Benthem (2003): *Conditional Probability Meets Update Logic*. *Journal of Logic, Language and Information* 12(4), pp. 409–421, doi:10.1023/A:1025002917675.
- [3] Johan van Benthem, David Fernández-Duque & Eric Pacuit (2012): *Evidence Logic: A New Look at Neighborhood Structures*. In: *Advances in Modal Logic 9, papers from the ninth conference on "Advances in Modal Logic," held in Copenhagen, Denmark, 22-25 August 2012*, pp. 97–118. Available at <http://www.aiml.net/volumes/volume9/Benthem-Fernandez-Duque-Pacuit.pdf>.
- [4] Johan van Benthem, David Fernández-Duque & Eric Pacuit (2014): *Evidence and plausibility in neighborhood structures*. *Ann. Pure Appl. Logic* 165(1), pp. 106–133, doi:10.1016/j.apal.2013.07.007.
- [5] Johan van Benthem, Jelle Gerbrandy & Barteld P. Kooi (2009): *Dynamic Update with Probabilities*. *Studia Logica* 93(1), pp. 67–96, doi:10.1007/s11225-009-9209-y.
- [6] Johan van Benthem & Eric Pacuit (2011): *Dynamic Logics of Evidence-Based Beliefs*. *Studia Logica* 99(1-3), pp. 61–92, doi:10.1007/s11225-011-9347-x.
- [7] Jan van Eijck & Bryan Renne (2016): *Update, Probability, Knowledge and Belief*. In: *Advances in Modal Logic 11, proceedings of the 11th conference on "Advances in Modal Logic," held in Budapest, Hungary, August 30 - September 2, 2016*, pp. 551–570. Available at <http://www.aiml.net/volumes/volume11/vanEijck-Renne.pdf>.
- [8] Daniel Ellsberg (1961): *Risk, Ambiguity, and the Savage Axioms**. *The Quarterly Journal of Economics* 75(4), p. 643, doi:10.2307/1884324.
- [9] Peter Gärdenfors (1975): *Qualitative probability as an intensional logic*. *J. Philosophical Logic* 4(2), pp. 171–185, doi:10.1007/BF00693272.
- [10] Sujata Ghosh & Dick de Jongh (2013): *Comparing strengths of beliefs explicitly*. *Logic Journal of the IGPL* 21(3), pp. 488–514, doi:10.1093/jigpal/jzs050.
- [11] E. T. Jaynes (2003): *Probability Theory: The Logic of Science*. Cambridge University Press, doi:10.1017/CBO9780511790423.
- [12] Barteld P. Kooi (2003): *Knowledge, Chance, and Change*. Ph.D. thesis, Groningen University.
- [13] John von Neumann (1951): *Various techniques used in connection with random digits*. Technical Report 12: 36, National Bureau of Standards Applied Math Series.
- [14] L.J. Savage (1972): *The Foundations of Statistics*. Dover Books on Mathematics Series, Dover Publications. Available at https://books.google.nl/books?id=y3_1DAAAQBAJ.
- [15] Dana Scott (1964): *Measurement structures and linear inequalities*. *Journal of Mathematical Psychology* 1(2), pp. 233 – 247, doi:10.1016/0022-2496(64)90002-1.
- [16] Krister Segerberg (1971): *Qualitative Probability in a Modal Setting*. *Studies in Logic and the Foundations of Mathematics* 63, pp. 341 – 352, doi:10.1016/S0049-237X(08)70852-8. Proceedings of the Second Scandinavian Logic Symposium.
- [17] Robert C. Stalnaker (1981): *A Theory of Conditionals*, pp. 41–55. Springer Netherlands, Dordrecht, doi:10.1007/978-94-009-9117-0_2.

A Alternative Definition of Conditional Neighborhood Models

Lemma 11. *Let $\mathfrak{M} = (W, N, V)$ be a conditional neighborhood model. Then \mathfrak{M} satisfies the following conditions for any $a \in Ag$, $x \in W$ and $X \subseteq W$:*

(m) $\forall Y \subseteq Z \subseteq X \cap [w]_a$: if $Y \in N_a^w(X)$, then $Z \in N_a^w(X)$;

(no-emptyset) $\emptyset \notin N_a^w(X)$;

(n)* if $X \cap [w]_a \neq \emptyset$, then $X \cap [w]_a \in N_a^w(X)$;

(\emptyset) if $X \cap [w]_a = \emptyset$, then $N_a^w(X) = \emptyset$.

Proof. Let $a \in Ag$, $x \in W$, $X \subseteq W$ and $X' = X \cap [w]_a$.

First consider (m). Let $Y \subseteq Z \subseteq X'$ and $Y \in N_a^w(X)$. Suppose for contradiction $Z \notin N_a^w(X)$. Then $Y \neq Z$, and hence $Y \subsetneq Z$, which implies $X' - Z \subsetneq X' - Y$. It follows, from (sc), that $X' - Y \in N_a^w(X)$, contrary to $Y \in N_a^w(X)$ and (d).

Second consider (no-emptyset). Suppose for contradiction that $\emptyset \in N_a^w(X)$. If $X' = \emptyset$, then by (c) $N_a^w(X) = \{\emptyset\}$; but by (d), since $\emptyset \in N_a^w(X)$,

$$\emptyset = X' - \emptyset \notin N_a^w(X).$$

Contradiction! If $X' \neq \emptyset$, then since $\emptyset \subseteq X'$, by (m) we have $X' \in N_a^w(X)$; but because $\emptyset \in N_a^w(X)$, by (d) $X' = X' - \emptyset$ should not in $N_a^w(X)$, contradiction.

Third for (n)*. Suppose $X' \neq \emptyset$ and for contradiction that $X' \notin N_a^w(X)$. Then $X' - \emptyset \notin N_a^w(X)$, and by (sc)

$$\emptyset \subsetneq X' \in N_a^w(X),$$

a contradiction!

Last for (\emptyset). Suppose $X' = \emptyset$. Then by (c), for all $Y \in N_a^w(X)$, $Y = \emptyset$. Because of (no-emptyset), we have $\emptyset \notin N_a^w(X)$. Therefore $N_a^w(X) = \emptyset$. \square

Note that in Definition 1 the equivalence relation for knowledge is derived from the conditional neighborhood function. Here we will show that there is an equivalent definition that takes epistemic equivalences as primary.

Definition 12. Let Ag be a finite set of agents. A conditional neighborhood model* \mathfrak{M} is a tuple (W, \sim, N, V) where:

- W is a non-empty set of worlds;
- \sim assigns to each $a \in Ag$ an equivalence relation \sim_a on W , and we use $[w]_a$ for the \sim_a class of w ;
- N assigns to each $a \in Ag$ a function N_a that assigns to every world $w \in W$ and set of worlds $X \subseteq W$ a collection $N_a^w(X)$ of sets of worlds—each such set called a neighborhood of X —subject to the following conditions:
 - (c) $\forall Y \in N_a^w(X) : Y \subseteq X \cap [w]_a$;
 - (a) $\forall v \in [w]_a : N_a^w(X) = N_a^v(X)$;
 - (d) $\forall Y \in N_a^w(X), X \cap [w]_a - Y \notin N_a^w(X)$;
 - (sc) $\forall Y, Z \subseteq X \cap [w]_a$: if $X - Y \notin N_a^w(X)$ and $Y \subsetneq Z$, then $Z \in N_a^w(X)$;
 - (ec) $\forall Y \subseteq W$: if $X \cap [w]_a = Y \cap [w]_a$, then $N_a^w(X) = N_a^w(Y)$;
- V is a valuation.

Note that in this definition we have another condition (a) on neighborhood functions. This contrasts with Definition 1, where we already make sure that (a) holds by the way we define $[w]_a$.

In Definition 12, however, $[w]_a$ is defined in terms of \sim_a , which is simply an equivalence relation that does not come with a guarantee for (a).

The following proposition assures us that the two approaches are equivalent.

Proposition 13. *Let $\mathfrak{M} = (W, \sim, N, V)$ be a conditional neighborhood model*, let $a \in Ag$ and let $R_a \subseteq W \times W$ be defined as follows:*

- $\forall w, v \in W, wR_av$ iff $\forall X \subseteq W, N_a^w(X) = N_a^v(X)$.

Then $\sim_a = R_a$, and (W, N, V) is a conditional neighborhood model.

Proof. Let $w, v \in W$. Suppose $w \sim_a v$. Then by (a), we know that for each $X \subseteq W$, $N_a^w(X) = N_a^v(X)$, which implies $(w, v) \in R_a$.

Suppose it is not the case that $w \sim_a v$. Then $[w]_a \cap [v]_a = \emptyset$. Similar to the proofs in Lemma 11, we can prove (n)* and (\emptyset) for \mathfrak{M} , and hence we have $[w]_a \in N_a^w([w]_a)$, and by (\emptyset), $[w]_a \notin N_a^v([w]_a)$. It follows that $N_a^w([w]_a) \neq N_a^v([w]_a)$, which implies $(w, v) \notin R_a$.

Therefore $[w]_a = \{v \in W \mid \forall X \subseteq W, N_a^w(X) = N_a^v(X)\}$. We can check that (W, N, V) satisfies all the conditions in Definition 1, which implies (W, N, V) is a conditional neighborhood model. \square

B Completeness of CN

In this section we prove Theorem 2. As a first step in the creation of a canonical model, we define formula closures.

Definition 14. Let $\alpha \in \mathcal{L}_{CN}$ be any CN-consistent formula, i.e., $\not\vdash_{CN} \neg\alpha$. The basic closure of α , denoted as $\Phi(\alpha)$, is the smallest set of formula Γ such that:

- if ϕ is a sub-formula of $\alpha \wedge \top$, then $\phi \in \Gamma$;
- if $\phi \in \Gamma$ and ϕ is not a negation, then $\neg\phi \in \Gamma$;
- if $\phi, \psi \in \Gamma$, then $\phi \wedge \psi \in \Gamma$.

Let $\Phi^B(\alpha)$ be the smallest extension of $\Phi(\alpha)$ such that

- if $\phi, \psi \in \Phi(\alpha)$, then $B_a(\phi, \psi) \in \Phi^B(\alpha)$ for each $a \in Ag$;
- if $\phi \in \Phi^B(\alpha)$ and ϕ is not a negation, then $\neg\phi \in \Phi^B(\alpha)$;
- if $\phi, \psi \in \Phi^B(\alpha)$, then $\phi \wedge \psi \in \Phi^B(\alpha)$.

Now we define the canonical model of α . Note that we use maximal consistent subsets of $\Phi^B(\alpha)$ instead of maximal consistent sets because we want to make sure our model is in some sense differentiable. Furthermore, for each maximal consistent subset of $\Phi^B(\alpha)$, we duplicate it Ω number of times, i.e., each possible world is a maximal consistent subset \mathbf{w} of $\Phi^B(\alpha)$ indexed by a number $i \in \Omega$, namely \mathbf{w}_i . In this way, we can define an equivalence relation \sim_a with the right properties in our canonical model.

Definition 15. A canonical conditional neighborhood model \mathfrak{M}_α of α is a tuple (W, \sim, N, V) where:

- $W = \{\mathbf{w} \subseteq \Phi^B(\alpha) \mid \mathbf{w} \text{ is a maximal consistent subset of } \Phi^B(\alpha)\} \times \Omega$.
- for each $a \in Ag$, \sim_a is an equivalence relation on W such that $\forall \mathbf{w}_i, \mathbf{v}_j \in W$:
 1. if $\mathbf{w}_i \sim_a \mathbf{v}_j$, then $\forall \phi, \psi \in \Phi(\alpha), \mathbf{w} \vdash_{CN} B_a(\phi, \psi)$ iff $\mathbf{v} \vdash_{CN} B_a(\phi, \psi)$;

2. if $\mathbf{w}_i \sim_a \mathbf{v}_j$ and $\mathbf{w} \cap \Phi(\alpha) = \mathbf{v} \cap \Phi(\alpha)$, then $\mathbf{w}_i = \mathbf{v}_j$;
 3. if $\forall \phi, \psi \in \Phi(\alpha), \mathbf{w} \vdash_{CN} B_a(\phi, \psi)$ iff $\mathbf{v} \vdash_{CN} B_a(\phi, \psi)$, then there is a $\mathbf{u}_l \in W$ such that $\mathbf{v} \cap \Phi(\alpha) = \mathbf{u}_l \cap \Phi(\alpha)$ and $\mathbf{w}_i \sim_a \mathbf{u}_l$.
- $N_a^W(X) := \{\{\mathbf{v}_j \in X \cap [\mathbf{w}_i]_a \mid \mathbf{v} \vdash_{CN} \psi\} \mid \phi \in \Phi(\alpha), \mathbf{w}_i \vdash_{CN} B_a(\|\mathbf{X}\|_a^{\mathbf{w}_i}, \psi)\}$, where $[\mathbf{w}_i]_a = \{\mathbf{v}_j \in W \mid \mathbf{w}_i \sim_a \mathbf{v}_j\}$ and $\|\mathbf{X}\|_a^{\mathbf{w}_i}$ is the characteristic formula for X w.r.t. $[\mathbf{w}_i]_a$, i.e., $\forall \mathbf{w}_i \in [\mathbf{w}]_a, \|\mathbf{X}\|_a^{\mathbf{w}_i} \in \mathbf{w}$ iff $\mathbf{w}_i \in X$.
 - $\mathbf{w}_i \in V(p)$ iff $p \in \mathbf{w}$.

Note that canonical conditional neighborhood models have equivalence relations \sim_a , unlike conditional neighborhood models. However, this is not a problem, because by Proposition 13, conditional neighborhood models with and without \sim_a relations are essentially the same.

Also note that $\Phi(\alpha)$ is finite up to logical equivalence, and because Ag is finite, $\Phi^B(\alpha)$ is finite up to logical equivalence as well. It follows that by Condition (2) for \sim_a each $[\mathbf{w}_i]_a$ is finite.

Because Condition (1) for \sim_a , $[\mathbf{w}_i]_a$ is differentiable w.r.t. $\Phi(\alpha)$ in the sense that for each subset $X \subseteq [\mathbf{w}_i]_a$, there is a characteristic formula $\phi \in \Phi(\alpha)$ such that $\forall \mathbf{v}_j \in [\mathbf{w}_i]_a, \phi \in \mathbf{v}_j$ iff $\mathbf{v}_j \in X$, and we use $\|\mathbf{X}\|_a^{\mathbf{v}_j}$ for such characteristic formula.

Lemma 16. *A canonical conditional neighborhood model \mathfrak{M}_α of α exists, given that α is CN-consistent.*

Proof. We only need to prove \sim_a exists for each $a \in Ag$. Let $a \in Ag$, let MCS be the set of maximal consistent subsets of $\Phi^B(\alpha)$, let $W = MCS \times \Omega$ and let $\approx \subseteq MCS \times MCS$ be the relation such that for all $\mathbf{w}, \mathbf{v} \in MCS, \mathbf{w} \approx \mathbf{v}$ iff:

- $\mathbf{w} \cap \Phi(\alpha) = \mathbf{v} \cap \Phi(\alpha)$,
- $\forall \phi, \psi \in \Phi(\alpha), \mathbf{w} \vdash_{CN} B_a(\phi, \psi)$ iff $\mathbf{v} \vdash_{CN} B_a(\phi, \psi)$.

It is easy to check that \approx is an equivalence relation, and because MCS is finite, $[\mathbf{w}] = \{\mathbf{v} \mid \mathbf{w} \approx \mathbf{v}\}$ is also finite. It follows that $[\mathbf{w}] \times \Omega$ is enumerable. Notice that $\{[\mathbf{v}] \times \Omega \mid \mathbf{v} \in MCS\}$ is a partition of W . For each $[\mathbf{w}] \times \Omega \in \{[\mathbf{v}] \times \Omega \mid \mathbf{v} \in MCS\}$, let w_0, w_1, \dots be an enumeration of $[\mathbf{w}] \times \Omega$.

Now we define \sim_a . Let $\sim_a \subseteq W \times W$ be the relation such that for all $\mathbf{w}_i, \mathbf{v}_j \in W, \mathbf{w}_i \sim_a \mathbf{v}_j$ iff

- $\forall \phi, \psi \in \Phi(\alpha), \mathbf{w} \vdash_{CN} B_a(\phi, \psi)$ iff $\mathbf{v} \vdash_{CN} B_a(\phi, \psi)$,
- \mathbf{w}_i and \mathbf{v}_j are both the n -th element in enumerations of $[\mathbf{w}] \times \Omega$ and $[\mathbf{v}] \times \Omega$ respectively, i.e., $\mathbf{w}_i = w_n$ and $\mathbf{v}_j = v_n$ for some $n \in \Omega$.

To check that such \sim_a is a desired equivalence relation, first it is easy to verify that it satisfies Condition (1) in Definition 15.

Now consider Condition (2) in Definition 15. Suppose $\mathbf{w}_i \sim_a \mathbf{v}_j$ and $\mathbf{w} \cap \Phi(\alpha) = \mathbf{v} \cap \Phi(\alpha)$. Then we have $[\mathbf{w}] = [\mathbf{v}]$, which implies $[\mathbf{w}] \times \Omega = [\mathbf{v}] \times \Omega$. Furthermore, since \mathbf{w}_i and \mathbf{v}_j are both the n -th element in enumerations of $[\mathbf{w}] \times \Omega$ and $[\mathbf{v}] \times \Omega$ respectively for some $n \in \Omega$, we obtain $\mathbf{w}_i = \mathbf{v}_j$.

Lastly consider Condition (2) in Definition 15. Let $\mathbf{w}_i, \mathbf{v}_j \in W$ such that

$$\forall \phi, \psi \in \Phi(\alpha), \mathbf{w} \vdash_{CN} B_a(\phi, \psi) \text{ iff } \mathbf{v} \vdash_{CN} B_a(\phi, \psi).$$

Suppose \mathbf{w}_i is the n -th element in the enumeration of $[\mathbf{w}] \times \Omega$. Then there is a $\mathbf{u}_l \in [\mathbf{v}] \times \Omega$ such that \mathbf{u}_l is the n -th element in the enumeration of $[\mathbf{v}] \times \Omega$. By the definition of \approx we have $\mathbf{v} \cap \Phi(\alpha) = \mathbf{u}_l \cap \Phi(\alpha)$ and

$$\forall \phi, \psi \in \Phi(\alpha), \mathbf{v} \vdash_{CN} B_a(\phi, \psi) \text{ iff } \mathbf{u}_l \vdash_{CN} B_a(\phi, \psi),$$

the latter of which implies

$$\forall \phi, \psi \in \Phi(\alpha), \mathbf{w} \vdash_{CN} B_a(\phi, \psi) \text{ iff } \mathbf{u} \vdash_{CN} B_a(\phi, \psi).$$

Therefore by the definition of \sim_a we have $\mathbf{w}_i \sim_a \mathbf{u}_l$ and $\mathbf{v} \cap \Phi(\alpha) = \mathbf{u} \cap \Phi(\alpha)$. \square

Lemma 17. *Let $\mathfrak{M}_\alpha = (W, \sim, N, V)$ be a canonical conditional neighborhood model of a CN-consistent formula α , let $a \in Ag$, $\mathbf{w} \in W$, and let $\phi \in \Phi(\alpha)$ such that $\mathbf{v} \vdash_{CN} \phi$ for all $\mathbf{v} \in [\mathbf{w}]_a$. Then $K_a\phi \in \mathbf{w}$.*

Proof. Because $\phi \in \Phi(\alpha)$, we know that either $K_a\phi \in \mathbf{w}$ or $\neg K_a\phi \in \mathbf{w}$. Suppose for contradiction that $\neg K_a\phi \in \mathbf{w}$. Then $\mathbf{w} \vdash_{CN} \check{K}_a\neg\phi$. Let

$$\bullet \mathbf{v}^- = \{\neg\phi\} \cup \{B_a(\psi, \chi) \in \mathbf{w} \mid \psi, \chi \in \Phi(\alpha)\} \cup \{\neg B_a(\psi, \chi) \in \mathbf{w} \mid \psi, \chi \in \Phi(\alpha)\}.$$

\mathbf{v}^- should be consistent, for otherwise there are

$$B_a(\alpha_1, \beta_1), \dots, B_a(\alpha_k, \beta_k), \neg B_a(\gamma_1, \delta_1), \dots, \neg B_a(\gamma_l, \delta_l) \in \mathbf{w}$$

such that

$$\vdash_{CN} \bigwedge_{i=1}^k B_a(\alpha_i, \beta_i) \wedge \bigwedge_{i=1}^l \neg B_a(\gamma_i, \delta_i) \rightarrow \phi,$$

which implies, by (Nec-K) and (Dist-K),

$$\vdash_{CN} \bigwedge_{i=1}^k K_a B_a(\alpha_i, \beta_i) \wedge \bigwedge_{i=1}^l K_a \neg B_a(\gamma_i, \delta_i) \rightarrow K_a\phi,$$

and then using (5B) and (4B) we have $\mathbf{w} \vdash_{CN} K_a\phi$, contrary to $\mathbf{w} \vdash_{CN} \check{K}_a\neg\phi$. It follows that \mathbf{v}^- is consistent, and then by Condition (3) in Definition 15, there is a $\mathbf{v} \in [\mathbf{w}]_a$ such that $\neg\phi \in \mathbf{v}$, contrary to our assumption that for all $\mathbf{v} \in [\mathbf{w}]_a$, $\mathbf{v} \vdash_{CN} \phi$.

Therefore $\mathbf{w} \vdash_{CN} K_a\phi$, i.e., $K_a\phi \in \mathbf{w}$. \square

Theorem 18. *Every CN-consistent formula α has a conditional neighborhood model \mathfrak{M}_α .*

Proof. Suppose $\not\vdash \neg\alpha$. Let $\mathfrak{M}_\alpha = (W, \sim, N, V)$ be a canonical conditional neighborhood model of α , and for all $a \in Ag$, $\mathbf{w} \in W$ and $X \subseteq W$, let $\|X\|_a^{\mathbf{w}}$ be the characteristic formula for X w.r.t. $[\mathbf{w}]_a$, i.e., $\forall \mathbf{v} \in [\mathbf{w}]_a$, $\|X\|_a^{\mathbf{w}} \in \mathbf{v}$ iff $\mathbf{v} \in X$. It suffices to show that (W, \sim, N, V) is a conditional neighborhood model* (see Definition 12), and then by Proposition 13 we can obtain that (W, N, V) is a conditional neighborhood model.

Clearly \sim_a are equivalence relations. It follows that for each $\mathbf{v} \in [\mathbf{w}]_a$, $[\mathbf{v}]_a = [\mathbf{w}]_a$, which implies

$$\forall \mathbf{w}, \mathbf{v} \in W \forall X \subseteq W, \mathbf{w} \sim_a \mathbf{v} \text{ only if } \|X\|_a^{\mathbf{w}} = \|X\|_a^{\mathbf{v}}. \quad (1)$$

It remains to show that N satisfies (c), (a), (d), (sc) and (ec). Let $a \in Ag$, $\mathbf{w} \in W$ and $X \subseteq W$.

First we consider (c), but it is straightforward by the definition of $N_a^{\mathbf{w}}(X)$.

Second for (a). Consider all $\mathbf{u} \in [\mathbf{w}]_a$, $Y, Z \subseteq W$. $Z \in N_a^{\mathbf{w}}(Y)$ iff

$$\exists \psi \in \Phi(\alpha), \mathbf{w} \vdash_{CN} B_a(\|Y\|_a^{\mathbf{w}}, \psi) \text{ and } Z = \{\mathbf{v} \in Y \cap [\mathbf{w}]_a \mid \mathbf{v} \vdash_{CN} \psi\}$$

iff by $\|X\|_a^{\mathbf{w}}$, $\psi \in \Phi(\alpha)$, Definition 15(1) and (1)

$$\exists \psi \in \Phi(\alpha), \mathbf{u} \vdash_{CN} B_a(\|Y\|_a^{\mathbf{u}}, \psi) \text{ and } Z = \{\mathbf{v} \in Y \cap [\mathbf{u}]_a \mid \mathbf{v} \vdash_{CN} \psi\}$$

iff $Z \in N_a^w(Y)$.

Third for (d), where we use axioms (D), (N) and (M). Consider any $Y \in N_a^w(X)$. By the definition of $N_a^w(X)$, we have that there is a $\phi \in \Phi(\alpha)$ such that $Y = \{\mathbf{v} \in X \cap [\mathbf{w}]_a \mid \mathbf{v} \vdash_{CN} \phi\}$ and $\mathbf{w} \vdash_{CN} B_a(\|X\|_a^w, \phi)$. Using (D) we can derive that

$$\mathbf{w} \vdash_{CN} \neg B_a(\|X\|_a^w, \neg\phi). \quad (2)$$

Now suppose by contradiction that there is a $\psi \in \Phi$ such that $X - Y = \{\mathbf{v} \in X \mid \mathbf{v} \vdash_{CN} \psi\}$ and $\mathbf{w} \vdash_{CN} B_a(\|X\|_a^w, \psi)$. Then by (N) we have $\mathbf{w} \vdash_{CN} B_a(\|X\|_a^w, \|X\|_a^w \wedge \psi)$. Note that for each $\mathbf{v} \in [\mathbf{w}]_a$, $\mathbf{v} \vdash_{CN} \|X\|_a^w \wedge \psi$ only if $\mathbf{v} \in X - Y$ only if $\mathbf{v} \vdash_{CN} \neg\phi$, i.e., $\mathbf{v} \vdash_{CN} \|X\|_a^w \wedge \psi \rightarrow \neg\phi$. Thus by Lemma 17 and $\|X\|_a^w, \phi, \psi \in \Phi(\alpha)$ we can get that $K_a(\|X\|_a^w \wedge \psi \rightarrow \neg\phi) \in \mathbf{w}$, and then using (M) we obtain $\mathbf{w} \vdash_{CN} B_a(\|X\|_a^w, \neg\phi)$, contrary to (2).

Then for (sc), we use (T) and (SC). Consider any $Y, Z \subseteq X \cap [\mathbf{w}]_a$ such that $X \cap [\mathbf{w}]_a - Y \notin N_a^w(X)$. Because $X \cap [\mathbf{w}]_a - Y \notin N_a^w(X)$, we have $\mathbf{w} \not\vdash_{CN} B_a(\|X\|_a^w, \neg\|Y\|_a^w)$. Recall that $\|X\|_a^w, \neg\|Y\|_a^w \in \Phi(\alpha)$, we know that either $B(\|X\|_a^w, \neg\|Y\|_a^w) \in \mathbf{w}$ or $\neg B(\|X\|_a^w, \neg\|Y\|_a^w) \in \mathbf{w}$. It follows that

$$\mathbf{w} \vdash_{CN} \neg B(\|X\|_a^w, \neg\|Y\|_a^w). \quad (3)$$

If $X \cap [\mathbf{w}]_a - Y = \emptyset$, then $X \cap [\mathbf{w}]_a = Y$, which implies there is no such Z with $Y \subsetneq Z$; thus (sc) vacuously holds. Suppose $\mathbf{v} \in X \cap [\mathbf{w}]_a - Y$ and $\mathbf{v} \in Z \supsetneq Y$. Then $\mathbf{v} \vdash_{CN} \|X\|_a^w \wedge \neg\|Y\|_a^w \wedge \|Z\|_a^w$. Using (T) we can obtain that $\mathbf{v} \vdash_{CN} \check{K}_a(\|X\|_a^w \wedge \neg\|Y\|_a^w \wedge \|Z\|_a^w)$. By $\mathbf{v} \in [\mathbf{w}]_a$ and Condition (1) in Definition 15, $\mathbf{w} \vdash_{CN} \check{K}_a(\|X\|_a^w \wedge \neg\|Y\|_a^w \wedge \|Z\|_a^w)$. It follows that, using (3) and (SC), $\mathbf{w} \vdash_{CN} B_a(\|X\|_a^w, \|Y\|_a^w \vee \|Z\|_a^w)$. Because $Y \subsetneq Z$ and thus $Z = \{\mathbf{v} \in X \cap [\mathbf{w}]_a \mid \mathbf{v} \vdash_{CN} \|Y\|_a^w \vee \|Z\|_a^w\}$, we have $Z \in N_a^w(X)$.

Last for (ec). Consider any $Y \subseteq W$ such that $X \cap [\mathbf{w}]_a = Y \cap [\mathbf{w}]_a$. Clearly $X \cap [\mathbf{w}]_a$ and $Y \cap [\mathbf{w}]_a$ have the same characteristic formula w.r.t. $[\mathbf{w}]_a$, i.e., $\|X\|_a^w = \|Y\|_a^w$. Let $Z \subseteq W$. $Z \in N_a^w(X)$ iff

$$\exists \psi \in \Phi(\alpha), \mathbf{w} \vdash_{CN} B_a(\|X\|_a^w, \psi) \text{ and } Z = \{\mathbf{v} \in X \cap [\mathbf{w}]_a \mid \mathbf{v} \vdash_{CN} \psi\}$$

iff since $\|X\|_a^w = \|Y\|_a^w$ and $X \cap [\mathbf{w}]_a = Y \cap [\mathbf{w}]_a$,

$$\exists \psi \in \Phi(\alpha), \mathbf{w} \vdash_{CN} B_a(\|Y\|_a^w, \psi) \text{ and } Z = \{\mathbf{v} \in Y \cap [\mathbf{w}]_a \mid \mathbf{v} \vdash_{CN} \psi\}$$

iff $Z \in N_a^w(Y)$. □

Lemma 19. (Truth Lemma) Let $\alpha \in \mathcal{L}_{CN}$ be a CN-consistent formula and let $\mathfrak{M}_\alpha = (W, N, V)$ be a canonical conditional neighborhood model of α removing equivalence relation \sim . Then for all formulas $\phi \in \Phi(\alpha)$ and $\mathbf{w} \in W$, $\mathfrak{M}_\alpha, \mathbf{w} \models_{CN} \phi$ iff $\phi \in \mathbf{w}$.

Proof. We prove by induction on ϕ . The cases of \top, p and the Boolean combinations are straightforward. For the case of $B_a(\psi, \chi)$. Let X be any set such that $\{\mathbf{v} \in [\mathbf{w}]_a \mid \psi \in \mathbf{v}\} \subseteq X$. Note that because $B_a(\psi, \chi) \in \Phi(\alpha)$, we have $\psi, \chi \in \Phi(\alpha)$.

Note that for each $\mathbf{v} \in [\mathbf{w}]_a$, $\psi \in \mathbf{v}$ iff $\mathbf{v} \in X$ iff $\|X\|_a^w \in \mathbf{v}$. Thus for each $\mathbf{v} \in [\mathbf{w}]_a$, $\mathbf{v} \vdash_{CN} \psi \leftrightarrow \|X\|_a^w$. It follows, by Lemma 17

$$K_a(\phi \leftrightarrow \|X\|_a^w) \in \mathbf{w}. \quad (4)$$

Also note that $\mathfrak{M}_\alpha, \mathbf{w} \models B_a(\psi, \chi)$ iff $\{\mathbf{v} \in \llbracket \psi \rrbracket_{\mathfrak{M}_\alpha} \cap [\mathbf{w}]_a \mid \mathfrak{M}_\alpha, \mathbf{v} \models \chi\} \in N_a^w(\llbracket \psi \rrbracket_{\mathfrak{M}_\alpha})$ iff (induction hypothesis) $X = \llbracket \psi \rrbracket_{\mathfrak{M}_\alpha}$ and $\{\mathbf{v} \in X \cap [\mathbf{w}]_a \mid \chi \in \mathbf{v}\} \in N_a^w(X)$. Thus

$$\mathfrak{M}_\alpha, \mathbf{w} \models B_a(\psi, \chi) \text{ iff } \{\mathbf{v} \in X \cap [\mathbf{w}]_a \mid \chi \in \mathbf{v}\} \in N_a^w(X). \quad (5)$$

Suppose $B_a(\psi, \chi) \in \mathbf{w}$. Then we have $\mathbf{w} \vdash B_a(\psi, \chi)$, which implies by (4) and (ec) $\mathbf{w} \vdash B_a(\|X\|_a^{\mathbf{w}}, \chi)$. It follows that $\{\mathbf{v} \in X \cap [\mathbf{w}]_a \mid \chi \in \mathbf{v}\} \in N_a^{\mathbf{w}}(X)$, and hence, using (5), $\mathfrak{M}_\alpha, \mathbf{w} \models B_a(\psi, \chi)$.

Suppose $\mathfrak{M}_\alpha, \mathbf{w} \models B_a(\psi, \chi)$, and hence by (5) $\{\mathbf{v} \in X \cap [\mathbf{w}]_a \mid \chi \in \mathbf{v}\} \in N_a^{\mathbf{w}}(X)$. Then there is a $\chi' \in \Phi(\alpha)$ such that $\{\mathbf{v} \in X \cap [\mathbf{w}]_a \mid \chi' \in \mathbf{v}\} = \{\mathbf{v} \in X \cap [\mathbf{w}]_a \mid \chi \in \mathbf{v}\}$ and $\mathbf{w} \vdash B_a(\|X\|_a^{\mathbf{w}}, \chi')$. It follows, by Lemma 17 that $K_a(\chi' \leftrightarrow \chi) \in \mathbf{w}$. Using (4), (ec) and (M) we have $\mathbf{w} \vdash B_a(\psi, \chi)$. Recall that $B_a(\psi, \chi) \in \Phi(\alpha)$, we can obtain $B_a(\psi, \chi) \in \mathbf{w}$. \square

C Completeness of PC and PC \pm

Theorem 20. (Soundness) *PC is sound w.r.t. \models_{PC} .*

Proof. To illustrate that $[\phi]B_a(\psi, \chi) \leftrightarrow (\phi \rightarrow B_a(\phi \wedge [\phi]\psi, \phi \wedge [\phi]\chi))$ is sound. Consider any conditional neighborhood model $\mathfrak{M} = (W, N, V)$ and any $w \in W$.

$\mathfrak{M}, w \models [\phi]B_a(\psi, \chi)$, iff $\mathfrak{M}, w \models \phi$ only if $\mathfrak{M}^\phi, w \models B_a(\psi, \chi)$, iff $\mathfrak{M}, w \models \phi$ only if $\{v \in \llbracket \psi \rrbracket_{\mathfrak{M}^\phi} \cap [w]_a^\phi \mid \mathfrak{M}^\phi, w \models \chi\} \in M_a^w(\llbracket \psi \rrbracket_{\mathfrak{M}^\phi})$.

$\mathfrak{M}, w \models \phi \rightarrow B_a(\phi \wedge [\phi]\psi, \phi \wedge [\phi]\chi)$ iff $\mathfrak{M}, w \models \phi$ only if $\mathfrak{M}, w \models B_a(\phi \wedge [\phi]\psi, \phi \wedge [\phi]\chi)$, iff $\mathfrak{M}, w \models \phi$ only if $\{v \in \llbracket \phi \wedge [\phi]\psi \rrbracket_{\mathfrak{M}} \cap [w]_a \mid \mathfrak{M}, v \models \phi \wedge [\phi]\chi\} \in N_a^w(\llbracket \phi \wedge [\phi]\psi \rrbracket_{\mathfrak{M}})$.

Let $v \in W$.

1. $\mathfrak{M}, v \models \phi \wedge [\phi]\chi$ iff $\mathfrak{M}, v \models \phi$ and $\mathfrak{M}, v \models \phi$ implies $\mathfrak{M}^\phi, v \models \chi$ iff $\mathfrak{M}^\phi, v \models \chi$.
2. $\mathfrak{M}, v \models \phi \wedge [\phi]\psi$ iff $\mathfrak{M}, v \models \phi$ and $\mathfrak{M}, v \models \psi$ implies $\mathfrak{M}^\phi, v \models \psi$ iff $\mathfrak{M}^\phi, v \models \psi$.
3. $\mathfrak{M}, w \models \phi$ implies $v \in \llbracket \phi \wedge [\phi]\psi \rrbracket_{\mathfrak{M}} \cap [w]_a$ iff $\mathfrak{M}, w \models \phi$ implies $\mathfrak{M}, v \models \phi \wedge [\phi]\psi$ and $w \sim_a v$ iff $\mathfrak{M}^\phi, v \models \psi$ and $w \sim_a^\phi v$ iff $v \in \llbracket \psi \rrbracket_{\mathfrak{M}^\phi} \cap [w]_a^\phi$.

It follows that $\mathfrak{M}, w \models \phi \rightarrow B_a(\phi \wedge [\phi]\psi, \phi \wedge [\phi]\chi)$ iff $\mathfrak{M}, w \models \phi$ implies $\{v \in \llbracket \psi \rrbracket_{\mathfrak{M}^\phi} \cap [w]_a^\phi \mid \mathfrak{M}^\phi, v \models \chi\} \in M_a^w(\llbracket \psi \rrbracket_{\mathfrak{M}^\phi})$, iff $\mathfrak{M}, w \models [\phi]B_a(\psi, \chi)$, and this completes our proof. \square

Theorem 21. *The PC Calculus is complete.*

Proof. This follows directly from the completeness of CN, plus the fact that the axioms for public announcement update are reduction axioms: we can compile out the update operators to reduce PC to CN. \square

Theorem 22. (Soundness) *PC \pm is sound w.r.t. $\models_{PC\pm}$.*

Proof. To illustrate that $[\pm\phi]B_a(\psi, \chi) \leftrightarrow (\phi \rightarrow B_a(\phi \wedge [\pm\phi]\psi, [\pm\phi]\chi)) \wedge (\neg\phi \rightarrow B_a(\neg\phi \wedge [\pm\phi]\psi, [\pm\phi]\chi))$ is sound. Consider any conditional neighborhood model $\mathfrak{M} = (W, N, V)$ and any $w \in W$. We consider two cases: $\mathfrak{M}, w \models \phi$ or $\mathfrak{M}, w \models \neg\phi$.

First suppose $\mathfrak{M}, w \models \phi$. Then $\mathfrak{M}, w \models [\pm\phi]B_a(\psi, \chi)$, iff $\mathfrak{M}^{\pm\phi}, w \models B_a(\psi, \chi)$, iff $\{v \in \llbracket \psi \rrbracket_{\mathfrak{M}^{\pm\phi}} \cap [w]_a^{\pm\phi} \mid \mathfrak{M}^{\pm\phi}, w \models \chi\} \in M_a^w(\llbracket \psi \rrbracket_{\mathfrak{M}^{\pm\phi}})$.

Furthermore, $\mathfrak{M}, w \models (\phi \rightarrow B_a(\phi \wedge [\pm\phi]\psi, [\pm\phi]\chi)) \wedge (\neg\phi \rightarrow B_a(\neg\phi \wedge [\pm\phi]\psi, [\pm\phi]\chi))$ iff (because $\mathfrak{M}, w \models \phi$), $\mathfrak{M}, w \models B_a(\phi \wedge [\pm\phi]\psi, [\pm\phi]\chi)$ iff

$$\{v \in \llbracket \phi \wedge [\pm\phi]\psi \rrbracket_{\mathfrak{M}} \cap [w]_a \mid \mathfrak{M}, v \models \phi \wedge [\pm\phi]\chi\} \in N_a^w(\llbracket \phi \wedge [\pm\phi]\psi \rrbracket_{\mathfrak{M}}).$$

Let $v \in W$.

1. $\mathfrak{M}, v \models \phi \wedge [\pm\phi]\chi$ iff $\mathfrak{M}, v \models \phi$ and $\mathfrak{M}, v \models \phi$ implies $\mathfrak{M}^{\pm\phi}, v \models \chi$ iff $\mathfrak{M}^{\pm\phi}, v \models \chi$.

2. $\mathfrak{M}, v \models \phi \wedge [\pm\phi]\psi$ iff $\mathfrak{M}, v \models \phi$ and $\mathfrak{M}, v \models \psi$ implies $\mathfrak{M}^{\pm\phi}, v \models \psi$ iff $\mathfrak{M}^{\pm\phi}, v \models \psi$.
3. $\mathfrak{M}, w \models \phi$ implies $v \in \llbracket \phi \wedge [\pm\phi]\psi \rrbracket_{\mathfrak{M}} \cap [w]_a$ iff $\mathfrak{M}, w \models \phi$ implies $\mathfrak{M}, v \models \phi \wedge [\pm\phi]\psi$ and $w \sim_a v$ iff $\mathfrak{M}^{\pm\phi}, v \models \psi$ and $w \sim_a^{\pm\phi} v$ iff $v \in \llbracket \psi \rrbracket_{\mathfrak{M}^{\pm\phi}} \cap [w]_a^{\pm\phi}$.

It follows that $\mathfrak{M}, w \models (\phi \rightarrow B_a(\phi \wedge [\pm\phi]\psi, [\pm\phi]\chi)) \wedge (\neg\phi \rightarrow B_a(\neg\phi \wedge [\pm\phi]\psi, [\pm\phi]\chi))$ iff $\{v \in \llbracket \psi \rrbracket_{\mathfrak{M}^{\pm\phi}} \cap [w]_a^{\pm\phi} \mid \mathfrak{M}^{\pm\phi}, v \models \chi\} \in M_a^w(\llbracket \psi \rrbracket_{\mathfrak{M}^{\pm\phi}})$, iff $\mathfrak{M}, w \models [\pm\phi]B_a(\psi, \chi)$.

The proof of the case that $\mathfrak{M}, w \models \neg\phi$ is similar to that of the first case, and we omit the proof. \square

Theorem 23. *The PC_{\pm} calculus is complete.*

Proof. Again, this follows directly from the completeness of CN, plus the fact that the axioms for public announcement update are reduction axioms: we can compile out the update operators to reduce PC_{\pm} to CN. \square