Bisimulation for conditional modalities

Alexandru Baltag and Giovanni Ciná

Institute for Logic, Language and Computation, University of Amsterdam

March 21, 2016

Abstract

We give a general definition of bisimulation for conditional modalities interpreted on selection functions and prove the correspondence between bisimilarity and modal equivalence. We describe how this framework specializes to unary modalities and further investigate the operators and semantics to which these results apply. First, we show that our models encompass both Grove's sphere models and Gabbay's models for non-monotonic reasoning. Second, we show how to derive a solid notion of bisimulation for conditional belief, behaving as desired both on plausibility models and on evidence models. These novel definitions of bisimulations are exploited in a series of undefinability results. Third, we treat relativized common knowledge, underlining how the same observations still hold for a different modality in a different semantics. Finally, we show the flexibility of the approach by generalizing it to multi-agent systems, encompassing the case of multi-agent plausibility models and suggesting how to cover group conditional belief.

1 Introduction

The Modal Logic literature offers a number of examples of conditional modalities, developed for a variety of reasons: conditionals from conditional logic, conditional belief, relativized common knowledge, to name a few. Yet there has been little work so far in developing model-theoretic tools to study such operators, which have been used mainly for the purpose of modelling our intuitions. The notable exception is conditional belief. The problem of finding the right notion of bisimulation for conditional belief has been the focal point of some recent publications in the field of formal epistemology [1, 2, 3, 10, 11].

In this paper we attempt to understand what is *conditional* about conditional modalities, proposing a framework that covers all the aforementioned operators. We do so by giving a general notion of bisimulation for conditional modalities, where the latter are interpreted on selection functions. Conditional logics, together with selection functions, have a long history and tradition in philosophical logic [18, 16, 9, 27]; they have been used in various applications such as non-monotonic inference, belief change and the analysis of intentions and desires. We thus tackle the problem at a high level of generality; this bird's eye perspectives enables a streamlined presentation of the main arguments, avoids repetitions and highlights the crucial assumptions.

To ensure that the notion of bisimulation is a good fit for the logic, the key result that one would like to obtain is the classical theorem 'bisimilarity iff modal equivalence, on some restricted class of models', echoing the analogous theorem for basic modal logic.[1] In other words, one wants to characterize exactly when two models are indistinguishable by means of a conditional modality. Such a result is however not the end of the story, a well behaved notion of bisimulation should also satisfy the following list of desiderata:

1. The bisimulation should be **structural**, that is, it should not make reference to formulas of the modal language besides the atomic propositions featuring in the basic condition "if w and w' are bisimilar then $\forall p$ we have $w \in V(p)$ iff $w' \in V(p)$".[2]

2. Ideally such bisimulation should be **closed under unions and relational composition**. The former ensures the existence of a largest bisimulation, while the latter guarantees that the related notion of bisimilarity is transitive.

3. The definition of such bisimulation should be in principle independent from additional parts of the structure that do not appear in the semantics of the conditional modality: two states should be indistinguishable only if they behave in the same with respect to the features that the conditional modality can "detect". This feature makes the bisimulation **modular**, allowing us to add further conditions to it in order to take care of additional operators in the language and still retain the correspondence with modal equivalence.

4. When the unconditional modality is amenable to different semantics, the bisimulation for the conditional version should **generalize** the bisimulation for the **unconditional modality uniformly across semantics**.

In this paper we provide a notion of bisimulation for conditional modalities that satisfies all the desiderata in the list and prove the correspondence between bisimilarity and modal equivalence for the semantics on selection functions. In the rest of the paper we discuss a number of examples of our general framework and reap the benefits of the general results. We begin by showing that our models encompass both the sphere models for belief revision and the selection function models for non-monotonic reasoning. We proceed by discussing how this

[1]See [8] for a standard reference in Modal Logic.

[2]For example, a non-structural notion of bisimulation for conditional belief on epistemic plausibility models was given in [10], but was regarded as problematic by the author himself for the same reason.

approach provides a solid notion of bisimulation for conditional belief, fulfilling the list of desiderata that we put forward. We then treat another conditional operator, relativized common knowledge, underlining how the same results still hold for a different modality with a different semantics. Finally, we show the flexibility of the approach by generalizing it to multi-agent systems, on one hand encompassing the case of multi-agent plausibility models and on the other hand suggesting how to cover *group* conditional belief.

The paper proceeds from the abstract to the concrete. In Section 2 we define conditional logic with semantics of selection function and give a general notion of bisimulation for conditional modalities; such notion is used to prove the main results of the paper. We pause to explain how such results translate to unconditional modalities. We then proceed to compare our conditional models with Grove's sphere models and the selection function models for non-monotonic reasoning, arguing that the latter are both examples of our definition. In Section 3 we introduce plausibility models, derive a notion of bisimulation for conditional belief on such structures and instantiate the previous result to this context. The bisimulation for conditional belief is then exploited to obtain some undefinability results. Section 4 deals with evidence models, structures introduced to describe how belief and knowledge relate to a body of available evidences. We show how they can be seen as conditional models equipped with selection functions and again obtain another version of the previous results, to which we add more observations about undefinability. In Section 5 we cover the case of relativized common knowledge interpreted on Kripke models, commenting on how to relate our framework to the original Hennessy-Milner result. Section 6 describes how our analysis extends to the multi-agent case by treating the example of multi-agent plausibility models and subsequently discussing the interesting issue of merging the beliefs of a group of agents. In Section 7 we discuss related work; we conclude in Section 8.

2 Bisimulation for conditional modalities

Consider the language

$$\phi ::= \quad p \,|\, \neg\phi \,|\, \psi \wedge \phi \,|\, \psi \rightsquigarrow \phi$$

where $p \in At$, a set of atomic propositions. The formulas $\psi \rightsquigarrow \phi$ are supposed to encode statements such as "ϕ is the case, conditional on ψ". As a semantic of this logic we consider selection functions of type $W \times \wp(W) \to \wp(W)$, along the lines of [16]. Similar considerations can be cast in the more general framework proposed by Chellas in [9], but the generality of neighborhood selection functions is not really needed here, neither to prove our results or to encompass the examples of the following sections; we thus limit ourselves to Lewis' original proposal.

Definition 2.1. A *conditional model* is a tuple $\mathcal{M} = \langle W, f, V \rangle$ with W a non-empty set of worlds, a function $f : W \times \wp(W) \to \wp(W)$, called a *selection* *function*, and $V : W \to \wp(At)$ a valuation function. The selection function is required to satisfy two conditions:

1. for all $w \in W$ we have $f(w, X) \subseteq X$;
2. if $X \subseteq Y$ then for all $w \in W$ we have that if $f(w, Y) \subseteq X$ then $f(w, Y) = f(w, X)$.

The intuition behind the selection function is that $f(w, X)$ selects the worlds in X that are more 'relevant' or 'important' from the perspective of world w. This reading immediately explains the first condition. The second requirements prescribe what the behaviour of the selection function should be with respect to stronger conditions. If X is a stronger condition than Y and all worlds v that are relevant to w within Y are still contained in X then the worlds relevant to w with respect to X coincide with the worlds relevant to w with respect to Y.

The semantics for the propositional part of the language is the usual one, while for conditionals we have the clause:

$$\mathcal{M}, w \vDash \psi \rightsquigarrow \phi \quad \text{iff} \quad f(w, [\![\psi]\!]_{\mathcal{M}}) \subseteq [\![\phi]\!]_{\mathcal{M}}$$

This encodes the intuition that "ϕ is the case, conditional on ψ" is the case in a world w iff all the ψ worlds that are relevant to w according to f are worlds that satisfy ϕ.

We now turn to the definition of bisimulation for conditional modalities, that is, the notion that is supposed to capture when two models are indistinguishable in the eyes of the language we just introduced. First we lay out some notation: given a relation $R \subseteq W \times W'$, $X \subseteq W$ and $X' \subseteq W'$ define

- $R[X] = \{y \in W' | \exists x \in X, (x, y) \in R\}$
- $R^{-1}[X] = \{x \in W | \exists y \in X', (x, y) \in R\}$

Definition 2.2 (Conditional Bisimulation)**.** Given two conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$, a *conditional bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ such that if $(w, w') \in Z$ then

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$,
- for all $X \subseteq W$ and $X' \subseteq W'$ such that $R[X] \subseteq X'$ and $R^{-1}[X'] \subseteq X$ we have that for every $x \in f(w, X)$ there exists a $y \in f'(w', X')$ (where f' is the selection function in $\mathcal{M}_2$) such that $(x, y) \in Z$, and viceversa.

Theorem 2.3 (Bisimilarity entails modal equivalence)**.** *Given two conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$, if $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a conditional bisimulation, then w and w' are modally equivalent with respect to the logic of conditionals.*

Proof. Suppose that Z is a conditional bisimulation, $(w, w') \in Z$ and $\mathcal{M}_1, w \vDash \psi \rightsquigarrow \phi$. Note that by induction hypothesis on ψ we have that $[\![\psi]\!]_{\mathcal{M}_1}$ and $[\![\psi]\!]_{\mathcal{M}_2}$ satisfy the right requirements and therefore can act as X and X' in the preconditions of the bisimulation property. Because of $w \vDash \psi \rightsquigarrow \phi$ we have $f(w, [\![\psi]\!]_{\mathcal{M}_1}) \subseteq [\![\phi]\!]_{\mathcal{M}_1}$. Now consider $v' \in f'(w', [\![\psi]\!]_{\mathcal{M}_2})$. By viceversa of the bisimulation property we know that there exists a $v \in f(w, [\![\psi]\!]_{\mathcal{M}_1})$ such that $(v, v') \in Z$. By assumption and induction hypothesis on ϕ we get $\mathcal{M}_2, v' \vDash \phi$. Since v' was generic we can conclude that $f'(w', [\![\psi]\!]_{\mathcal{M}_2}) \subseteq [\![\phi]\!]_{\mathcal{M}_2}$, thus $\mathcal{M}_2, w' \vDash \psi \rightsquigarrow \phi$. For the converse use the other direction of the bisimulation property. □

Note that this proof does not make use of the properties of selection functions. However, they will turn out to be crucial in the following Theorem.

Theorem 2.4 (Modal equivalence entails bisimilarity on finite models)**.** *Given two finite conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$, if w and w' are modally equivalent then $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a conditional bisimulation.*

Proof. We show that the relation of modal equivalence is a conditional bisimulation, we will call it Z for short.

First a preliminary observation. Call X and X' the two sets satisfying $Z[X] \subseteq X'$ and $Z^{-1}[X'] \subseteq X$. We show how to build a formula α that plays the role of X and X' as precondition. Notice that we can divide the domain of $\mathcal{M}_1$ into three disjoint parts

- X
- A, the set of elements having some modally equivalent counterparts in X, but having no modally equivalent counterpart in $\mathcal{M}_2$
- $W_1 \backslash X \cup A$

Notice how the conditions on X and X' enforce the property of A: an element in A cannot have a counterpart in X', or otherwise would be already in X; an element of A cannot have a modally equivalent counterpart in $W_2 \backslash X'$ or X itself would violate the first precondition. A symmetric partition can be defined on the model $\mathcal{M}_2$, switching the roles of X and X'; we will indicate with A' the corresponding region in $\mathcal{M}_2$.

Since the image of X under Z lies within X', we know that the elements in X are not modally equivalent to the elements outside X', thus the elements in $X \cup A$ are also not modally equivalent to the elements outside X'. Since we are dealing with finite models we can enumerate the elements in $X \cup A$, call them $x_1, \ldots, x_n$, and the elements in $W_2 \backslash X' \cup W_1 \backslash X \cup A$, call them $y_1, \ldots, y_m$. By our assumptions and definition of the partition we know that for each i and j, with $1 \leq i \leq n$ and $1 \leq j \leq m$, there is a formula ψ_{ij} such that $x_i \vDash \psi_{ij}$ and $y_j \nvDash \psi_{ij}$. We can thus construct a formula

$$\gamma := \bigvee_{1 \leq i \leq n} \bigwedge_{1 \leq j \leq m} \psi_{ij}$$

that is true at each x_i in $X \cup A$ and false at each y_j in $(W_2 \backslash X') \cup (W_1 \backslash X \cup A)$. Symmetrically, there must be a formula γ' that is true at $X' \cup A'$ and false at

$W_1 \backslash X \cup W_2 \backslash X' \cup A'$. Now consider the formula

$$\alpha := \gamma \vee \gamma'$$

Let us have a closer look at the extension of α in $\mathcal{M}_1$. We have that γ' is false outside X, hence its extension lies within X. As for γ, we know it is true at $X \cup D$ and false in $W_1 \backslash X \cup A$. Thus the extension of $\gamma \vee \gamma'$, and therefore of the formula α itself, is $X \cup A$. We can make an analogous argument to show that the extension of α in $\mathcal{M}_2$ is $X' \cup A'$.

Say now that $(w, w') \in Z$ and reasoning by contradiction suppose Z does not satisfy the bisimulation property for sets X and X'. This means that there is an $x \in f(w, X)$ such that for all $y \in f'(w', X')$ we have $(x, y) \notin Z$.

Consider a generic element x' in $f(w, [\![\alpha]\!]_{\mathcal{M}_1})$. Since $f(w, [\![\alpha]\!]_{\mathcal{M}_1}) \subseteq [\![\alpha]\!]_{\mathcal{M}_1}$ by the first property of selection functions, we know that x' must be either in X or in A. If there is an element $x' \in f(w, [\![\alpha]\!]_{\mathcal{M}_1})$ in A, since we know that elements in A are not modally equivalent to any element in W_2, we can build a formula β that is false at x' and true everywhere in W_2, thus a fortiori in $f(w', [\![\alpha]\!]_{\mathcal{M}_2})$. This gives us the contradiction that we want: $w \vDash \neg(\alpha \rightsquigarrow \beta)$ and $w' \vDash \alpha \rightsquigarrow \beta$. We can thus assume that $f(w, [\![\alpha]\!]_{\mathcal{M}_1}) \subseteq X$. This is enough to apply the second property of selection functions and conclude that $f(w, [\![\alpha]\!]_{\mathcal{M}_1}) = f(w, X)$.

This ensures that the element $x \in f(w, X)$ for which the condition fails is indeed also in $f(w, [\![\alpha]\!]_{\mathcal{M}_1})$. If we now look at the set $f'(w', [\![\alpha]\!]_{\mathcal{M}_2})$, repeating a reasoning similar to the one just outlined we can conclude that the elements in $f'(w', [\![\alpha]\!]_{\mathcal{M}_2})$ are either in $f'(w', X')$ or in A'. Note that in this step we use the other inclusion of the second property. By assumption we have that x is not modally equivalent to any $y \in f'(w', X')$ and by definition x is not m.e. to any element in A' (because elements in the latter set have no m.e. counterpart in the first model). We can thus build a formula β that is false at x and true everywhere in $f'(w', [\![\alpha]\!]_{\mathcal{M}_2})$; this gives us the contradiction $w \vDash \neg(\alpha \rightsquigarrow \beta)$ and $w \vDash \alpha \rightsquigarrow \beta$. □

2.1 Back to unary modalities

The reader may now wonder if and how these results hold for the *un*-conditional version of such modalities. Each selection function $f : W \times \wp(W) \to \wp(W)$ specifies a unary modality when we fix the second input to W. Indicating the unary modality with $\top \rightsquigarrow \phi$, we get the semantic clause

$$\mathcal{M}, w \vDash \top \rightsquigarrow \phi \quad \text{iff} \quad f(w, W) \subseteq [\![\phi]\!]_{\mathcal{M}}$$

Accordingly, the precondition on subsets trivializes and the correct notion of bisimulation becomes

Definition 2.5 (Unary bisimulation)**.** Given two conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$, a *unary bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ such that if $(w, w') \in Z$ then

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$,

- for every $x \in f(w, W)$ there exists a $y \in f'(w', W')$ (where f' is the selection function in $\mathcal{M}_2$) such that $(x, y) \in Z$, and viceversa.

We then obtain analogous version of the previous results.

Theorem 2.6 (Bisimilarity entails modal equivalence)**.** *Given two conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$, if $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a unary bisimulation, then w and w' are modally equivalent with respect to the logic containing only the unary modality.*

Proof. Proved via a straightforward induction on the structure of the formula. □

Now that we are free from preconditions, however, we can generalize the other direction.

Theorem 2.7 (Modal equivalence entails bisimilarity)**.** *Given two conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$ such that*

- *$f(w, W)$ is finite for all $w \in W$*
- *$f'(w', W')$ is finite for all $w' \in W'$*

if w and w' are modally equivalent in the language of the unary modality then $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a unary bisimulation.

Proof. This proof essentially follows the original idea of the Hennessy-Milner result. Suppose modal equivalence is not a unary bisimulation, then there is, say, $x \in f(w, W)$ such that for all $y \in f'(w', W')$ we have $(x, y) \notin Z$. Since $f'(w', W') = \{y_1, \dots, y_n\}$ we have finitely many formulas ψ_i, with $1 \leq i \leq n$, such that $x \not\vDash \psi_i$ and $y_i \vDash \psi_i$. Then $w \not\vDash \vee_{1 \leq i \leq n} \psi_i$ and $w' \vDash \vee_{1 \leq i \leq n} \psi_i$, contradiction. □

This last result clarifies why the finiteness conditions that we have in the conditional case can be relaxed in the un-conditional case, allowing for wider classes of models. We will see in section 5 how this result specializes to the original Hennessy-Milner theorem.

2.2 Closure under union and relational composition

In this section we display some properties of the newly defined notion of bisimulation: closure under arbitrary unions and, restricted to grounded models, closure under relational composition.

Proposition 2.8. *Conditional bisimulations are closed under unions.*

Proof. Given a family of conditional bisimulations $\{Z_i \subseteq W_1 \times W_2\}_{i \in I}$, consider their union $\bigcup_{i \in I} Z_i$. Suppose $(w, w') \in \bigcup_{i \in I} Z_i$ and two sets $X \subseteq W_1$ and $X' \subseteq W_3$ are such that

1. $\bigcup_{i \in I} Z_i[X] \subseteq X'$,

2. $\bigcup_{i\in I} Z_i^{-1}[X'] \subseteq X$,

To establish that $\bigcup_{i\in I} Z_i$ is a conditional bisimulation we need to show that for every $x \in f(w,X)$ there is $y \in f'(w',X')$ such that $(x,y) \in \bigcup_{i\in I} Z_i$. Notice that from $(w,w') \in \bigcup_{i\in I} Z_i$ we can deduce that there is an index i for which $(w,w') \in Z_i$. We also know that

1. $\{y|\exists x \in X(x,y) \in Z_i\} = Z_i[X] \subseteq \bigcup_{i\in I} Z_i[X] \subseteq X'$,
2. $\{x|\exists y \in X'(x,y) \in Z_i\} = Z_i^{-1}[X'] \subseteq \bigcup_{i\in I} Z_i^{-1}[X'] \subseteq X$,

Therefore X and X' also satisfy the preconditions for the relation Z_i: applying the property of conditional bisimulation we obtain that for every $x \in f(w,X)$ there is $y \in f'(w',X')$ such that $(x,y) \in Z_i$. But the latter fact entails $(x,y) \in \bigcup_{i\in I} Z_i$, we are done. The converse direction is proved symmetrically. □

Definition 2.9. A relation $R \subseteq X \times Y$ is *total* if for every $x \in X$ there is a $y \in Y$ such that $(x,y) \in R$ and for every $y \in Y$ there is an $x \in X$ such that $(x,y) \in R$.

Definition 2.10. A conditional model is *grounded* if, for all w and $\{x\}$ in the domain of the model, $f(w,\{x\}) \neq \emptyset$.

Lemma 2.11. *Every conditional bisimulation between two grounded models is total.*

Proof. For $\mathcal{M}_1$ and $\mathcal{M}_2$ grounded conditional models, suppose $Z \subseteq W_1 \times W_2$ is a conditional bisimulation: then there must be a pair $(w,w') \in Z$. Suppose moreover that Z is not total, say because there is an $x \in W_1$ with no counterpart in W_2. Take $\{x\}$ and $\emptyset$ and notice that they fulfill the preconditions of the property of conditional bisimulation:

- $Z[\{x\}] = \emptyset \subseteq \emptyset$
- $Z^{-1}[\emptyset] = \emptyset \subseteq \{x\}$

Then we must conclude that for every $x \in f(w,\{x\})$ there is a $y \in f'(w',\emptyset)$ such that $(x,y) \in Z$. Since by assumption $f(w,\{x\}) \neq \emptyset$, there must be some $z \in f'(w',\emptyset)$. However, by the first condition on selection function we have $f'(w',\emptyset) \subseteq \emptyset$, so there cannot be a $k \in f'(w',\emptyset)$ such that $(z,k) \in Z$, contradiction. The other direction is proved analogously. □

Proposition 2.12. *Restricted to any class of grounded models, the notion of conditional bisimulation is closed under relational composition.*

Proof. Suppose $\mathcal{M}_1$, $\mathcal{M}_2$ and $\mathcal{M}_3$ are three grounded models and $Z_1 \subseteq W_1 \times W_2$ and $Z_2 \subseteq W_2 \times W_3$ are two conditional bisimulations connecting them. To show that their relational composition $Z_1;Z_2$ is also a conditional bisimulation we first need to show that it is not empty. By Z_1 being not empty we know that there is $(w,w') \in Z_1$. By the previous Lemma we know that Z_1 and Z_2 are total. The latter fact ensures that there is some w'' such that $(w', w'') \in Z_2$, thus $(w, w'') \in Z_1; Z_2$.

For the last property, suppose $(w, w'') \in Z_1; Z_2$. By definition this means that there is a w' such that $(w, w') \in Z_1$ and $(w', w'') \in Z_2$. Now consider two sets $X \subseteq W_1$ and $X'' \subseteq W_3$ such that

1. $Z_1; Z_2[X] \subseteq X''$,
2. $Z_1; Z_2^{-1}[X''] \subseteq X$,

What we need to show is that for every $x \in f(w, X)$ there is a $z \in f''(w'', X'')$ such that $(x, z) \in Z_1; Z_2$. Define

$$X' := \{y \in W_2 | \exists x \in X, (x, y) \in Z_1 \quad \text{or} \quad \exists z \in X'', (y, z) \in Z_2\}$$

We check that

- $Z_1[X] \subseteq X'$,
- $Z_1^{-1}[X'] \subseteq X$,

The first item holds by definition of X'. For the second one suppose $(x, y) \in Z_1$ and $y \in X'$. By totality of Z_2 we know that there is a z such that $(y, z) \in Z_2$, hence $(x, z) \in Z_1; Z_2$. By definition of X' we can now make a case distinction. In the first case there is an element $x' \in X$ such that $(x', y) \in Z_1$. We can then conclude that $(x', z) \in Z_1; Z_2$ and thus by assumption 1 we have $z \in X''$. But then by the latter fact and $(x, z) \in Z_1; Z_2$, coupled with assumption 2, we can infer that $x \in X$. In the second case we have that there is a $z' \in X''$ such that $(y, z') \in Z_2$. This gives us immediately that $(x, z') \in Z_1; Z_2$ and thus by assumption 2 we can again conclude $x \in X$.

Since X and X' fulfill the preconditions of the property of conditional bisimulation for Z_1, we can deduce that for every $x \in f(w, X)$ there is $y \in f'(w', X')$ such that $(x, y) \in Z_1$. We can now repeat the same proof strategy for X' and X'' and apply the property of Z_2 to obtain that for every $y \in f'(w', X')$ there is $z \in f''(w'', X'')$ such that $(y, z) \in Z_2$. Concatenating this with the previous result we get the desired conclusion: for every $x \in f(w, X)$ there is a $z \in f''(w'', X'')$ such that $(x, z) \in Z_1; Z_2$. The converse is proved symmetrically. □

Proposition 2.13. *Restricted to any class of grounded models, the relation of bisimilarity defined as*

two states w and w' are bisimilar iff there exists a conditional bisimulation Z such that $(w, w') \in Z$

is an equivalence relation.

Proof. We need to show that the relation of bisimilarity is reflexive, symmetric and transitive. For reflexivity, it is immediate to see that the identity relation is a conditional bisimulation. The definition of conditional bisimulation is itself symmetric, hence the converse of a conditional bisimulation is always a conditional bisimulation; the symmetry for bisimilarity follows. As for transitivity, Proposition 3.9 ensures that if there are two conditional bisimulations Z_1 and Z_2 such that $(w, w') \in Z_1$ and $(w', w'') \in Z_2$ then there is a conditional bisimulation containing the pair (w, w''), namely the relational composition $Z_1; Z_2$. □

2.3 Models for non-monotonic logics

In this section we compare our conditions on selection functions with well-known semantic requirements for the models of non-monotonic logics. In [12] Gabbay argues that our intuitions about non-monotonic derivations are captured by consequence relations $\vdash_{NM}$ which satisfy three properties known as Reflexivity, Cut and Cautious Monotonicity:

- $\phi \vdash_{NM} \phi$ (Reflexivity)
- $\phi \vdash_{NM} \psi$ and $(\phi \wedge \psi) \vdash_{NM} \theta$ entail $\phi \vdash_{NM} \theta$ (Cut)
- $\phi \vdash_{NM} \psi$ and $\phi \vdash_{NM} \theta$ entail $(\phi \wedge \psi) \vdash_{NM} \theta$ (Cautious Monotonicity)

The semantic translation of such properties on selection functions is the following.

- $f(w, X) \subseteq X$ (Reflexivity)
- $(f(w, Y) \subseteq X$ and $f(w, X \cap Y) \subseteq X'$ entail $(f(w, Y) \subseteq X'$ (Cut)
- $(f(w, Y) \subseteq X$ and $f(w, Y) \subseteq X'$ entail $f(w, X \cap Y) \subseteq X'$ (Cautious Monotonicity)

It is clear that Reflexivity is exactly our first requirement on conditional models. The other two properties correspond to the two inclusion of our second condition on selection functions.

Lemma 2.14. *Cut entails the left-to-right inclusion in the second condition on selection functions. In presence of Reflexivity, the converse also holds.*

Proof. Suppose $X \subseteq Y$ and $f(w, Y) \subseteq X$. Substitute X' with $f(w, X)$ in the definition of Cut: the premises are now $f(w, Y) \subseteq X$, which we have by assumption, and $f(w, X \cap Y) = f(w, X) \subseteq f(w, X)$, which is trivially the case. By Cut we can then conclude $f(w, Y) \subseteq f(w, X)$, as desired.

For the other direction, assume $(f(w, Y) \subseteq X$ and $f(w, X \cap Y) \subseteq X'$. To conclude $(f(w, Y) \subseteq X'$ it is enough to derive $(f(w, Y) \subseteq f(w, X \cap Y)$. Notice now that Y and $X \cap Y$ satisfy the antecedent of the second condition: one one hand $X \cap Y \subseteq Y$ by definition, on the other hand $(f(w, Y) \subseteq X \cap Y$ follows from our assumption $(f(w, Y) \subseteq X$ and Reflexivity $(f(w, Y) \subseteq Y$. Thus applying the second condition we obtain $(f(w, Y) \subseteq f(w, X \cap Y)$ and we are done. □

Lemma 2.15. *Cautious Monotonicity entails the right-to-left inclusion in the second condition on selection functions. In presence of Reflexivity, the converse also holds.*

Proof. Suppose $X \subseteq Y$ and $f(w,Y) \subseteq X$. Replacing X' with $f(w,Y)$ in the definition of Cautious Monotonicity we can check that both premises are given: $f(w,Y) \subseteq X$, by assumption, and $f(w,Y) \subseteq f(w,Y)$. We can thus conclude $f(w,X) = f(w,X \cap Y) \subseteq X' = f(w,Y)$.

For the converse, assume $(f(w,Y) \subseteq X$ and $f(w,Y) \subseteq X'$. To obtain $f(w,X \cap Y) \subseteq X'$ it is enough to show $f(w,X \cap Y) \subseteq f(w,Y)$. Notice that we have $(f(w,Y) \subseteq X \cap Y$, by assumption $(f(w,Y) \subseteq X$ and Reflexivity $(f(w,Y) \subseteq Y$. Coupled with $X \cap Y \subseteq Y$, we are in position to use the right-to-left inclusion in the second condition, thus obtaining $f(w,X \cap Y) \subseteq f(w,Y)$. $\square$

Thus our conditional models are exactly those satisfying Gabbay's requirements, when those are formulated in terms of selection functions.

2.4 Sphere models for belief revision

Proposed first by Lewis [16] and later modified by Grove [14] the so called 'sphere models' are one of the most important models used in the analysis of conditionals, belief revision and theory change. In this subsection we show that such models are an example of our conditional models.

Definition 2.16 (adapted from [14])**.** A *sphere model* is a tuple $\mathcal{M} = \langle W, E, V\rangle$ with W a non-empty set of worlds, $V : W \to \wp(At)$ a valuation function and $S : \wp(W) \to \wp(\wp(W))$ a function assigning to each set of worlds a so-called *system of spheres* satisfying the following requirements for every $X \subseteq W$

1. $S(X)$ is totally ordered by inclusion: if $U, V \in S(X)$ then either $U \subseteq V$ or $V \subseteq U$;
2. X is the minimum of $S(X)$;
3. $W \in S(X)$;
4. for any subset $Y \subseteq W$, if there is $U \in S(X)$ such that $U \cap Y \neq \emptyset$ then there is a smallest set $V \in S(X)$ with such a property, call it $c_X(Y)$.

Proposition 2.17. *Sphere models are conditional models, where* $f(w,X) = c_{\{w\}}(X) \cap X$ *for all* w.[3]

Proof. We need to check that the newly defined f fulfills the prerequisites of selection functions in Definition 2.1. The first condition is obviously given by $f(w,X) = c_{\{w\}}(X) \cap X \subseteq X$. For the second one, suppose $X \subseteq Y$, $f(w,Y) = (c_{\{w\}}(Y) \cap Y) \subseteq X$ and take $x' \in f(w,Y)$. By assumption $f(w,Y) = (c_{\{w\}}(Y) \cap Y) \subseteq X$ we can readily conclude that $x' \in X$. On the other hand, $X \subseteq Y$ entails that $c_{\{w\}}(Y) \subseteq c_{\{w\}}(X)$: the smallest set having non-empty intersection with X also has non empty intersection with Y, and $c_{\{w\}}(Y)$ is minimal with such property. Thus $x' \in c_{\{w\}}(X)$ and we can conclude $x' \in c_{\{w\}}(X) \cap X = f(w,X)$. For the other inclusion consider $x' \in c_{\{w\}}(X) \cap X = f(w,X)$. Since $x' \in X$ we

[3] Grove himself suggests such a function in [14] p. 159.

get $x' \in Y$. By $f(w,Y) = (c_{\{w\}}(Y) \cap Y) \subseteq X$ and the fact that $c_{\{w\}}(Y) \cap Y$ is non-empty we can infer that X has non-empty intersection with $c_{\{w\}}(Y)$, thus $c_{\{w\}}(X) \subseteq c_{\{w\}}(Y)$, and therefore $x' \in c_{\{w\}}(Y)$. This suffices to conclude $x' \in c_{\{w\}}(Y) \cap Y = f(w,Y)$. □

We have seen that sphere models and selection function models for non-monotonic reasoning do fall under the scope of our analysis. The language we were considering was a language with an additional binary operator $\psi \rightsquigarrow \phi$, meant to capture conditionals or defeasible entailment. Now we proceed to show that our notion of conditional modality is general enough to encompass well-known modal operators with independent origins and motivations, namely conditional belief and relativized common knowledge.

3 Conditional belief on plausibility models

Plausibility models are widely used in formal epistemology [5, 20], while their introduction can be traced back at least to [16]. In their simplest form they come in the shape of a set equipped with a preorder, intuitively representing the possible scenarios and how an agent ranks them in terms of plausibility, whence the name.

Definition 3.1. A *plausibility model* is a tuple $\mathcal{M} = \langle W, \leq, V \rangle$ with W a non-empty set of worlds, a reflexive and transitive relation $\leq \subseteq W \times W$ and $V : W \to \wp(At)$ a valuation function.

The strict relation $<$ is defined as usual from $\leq$. Given a set $X \subseteq W$, let

$$Min(X) = \{v \in X | \neg \exists w \in X \quad \text{s.t.} \quad w < v\}$$

We can think of $Min(X)$ as the set of most plausible worlds in X.[4] When we want to specify the ordering we write $Min_{\leq}(X)$.

Among the variety of operators that are studied in the setting of plausibility models, a prominent part is played by the operator of *conditional belief*. This operator is an example of a conditional modality and it is usually written as $B^{\psi}\phi$. The standard belief operator can be defined as $B^{\top}\phi$.

The semantics for the propositional part of the language is the usual one, while for belief and conditional belief we have the clauses:

- $\mathcal{M}, w \vDash B\phi$ iff for all $v \in Min(W)$ we have $\mathcal{M}, v \vDash \phi$
- $\mathcal{M}, w \vDash B^{\psi}\phi$ iff for all $v \in Min(\llbracket\psi\rrbracket)$ we have $\mathcal{M}, v \vDash \phi$

The notion of bisimulation for the standard belief operator on plausibility models is the following.

[4] We mostly omit the parenthesis in $Min(X)$ in what follows.

Definition 3.2. Given two plausibility models $\mathcal{M}_1$ and $\mathcal{M}_2$, a *plausibility B-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ such that if $(w, w') \in Z$ then

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$;
- for every $x \in MinW_1$ there is $y \in MinW_2$ such that $(x, y) \in Z$, and viceversa.

The following result is folklore.

Theorem 3.3. *Bisimilarity with respect to plausibility B-bisimulation entails modal equivalence with respect to the language with only the belief operator. On models having finitely many minimal elements, modal equivalence with respect to the latter language entails bisimilarity for plausibility B-bisimulation.*

3.1 Plausibility CB-bisimulation

To obtain a bisimulation for conditional belief on plausibility models we show how the latter are an instance of conditional models; this move will indicate a systematic way to specialize the results of Section 2 to this particular context.

Definition 3.4. A plausibility model $\mathcal{M}$ is *well-founded* if it contains no infinite descending chains.[5]

Proposition 3.5. *Well-founded plausibility models are conditional models, where $f(w, X) = MinX$ for all w.*

Proof. We need to check that the newly defined f fulfills the prerequisites of selection functions in Definition 2.1. The first condition on selection functions is fulfilled by the very definition of Min. For the second one, suppose $X \subseteq Y$, $MinY \subseteq X$ and take $x' \in MinY$. Since $X \subseteq Y$, if there is no element below x' in Y then a fortiori there is no element below it in the subset X, thus in this circumstance $x' \in MinX$. For the other inclusion take $x' \in MinX$; we show x' is also minimal for Y. By contradiction, suppose there is $z \in Y \setminus X$ such that $z < x'$. Since we are in a well-founded model there must be a minimal element $z' \in MinY$ such that $z' \leq z$; but by assumption $MinY \subseteq X$, hence $z' \in X$ and $z' < x'$, contradicting the fact that x' is minimal in X. □

Notice that, setting $f(w, X) = MinX$, the definition of the satisfaction relation for conditional belief becomes an instance of the satisfaction relation for conditional modalities given in Section 2. With the new f we can see in retrospect that Theorem 3.3 follows from our framework from the considerations of Subsection 2.1. We also obtain a new notion of bisimulation for conditional belief on plausibility models, following the footprint of Definition 2.5.

Definition 3.6. Given two plausibility models $\mathcal{M}_1$ and $\mathcal{M}_2$, a *plausibility CB-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ such that if $(w, w') \in Z$ then

[5] Equivalently, assuming choice, if every non empty subsets has minimal elements.

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$,
- for all $X \subseteq W$ and $X' \subseteq W'$ such that $R[X] \subseteq X'$ and $R^{-1}[X'] \subseteq X$ we have that for every $x \in MinX$ then there exists a $y \in MinX'$ such that $(x, y) \in Z$, and viceversa.

We can now conclude that this notion of bisimilarity corresponds to modal equivalence on finite plausibility models. Throughout this section and the following one we use 'modal equivalence' meaning with respect to the language of conditional belief.

Theorem 3.7. *Given two plausibility models $\mathcal{M}_1$ and $\mathcal{M}_2$, if $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a plausibility CB-bisimulation, then w and w' are modally equivalent. On finite models, if w and w' are modally equivalent then $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a plausibility CB-bisimulation.*

We can also import the results concerning the closure under union and relational composition.

Lemma 3.8. *Every plausibility model is a grounded model.*

Proof. Given a plausibility model $\mathcal{M}$, it is enough to check that, having defined $f(w, X) = MinX$, we have $f(w, \{x\}) \neq \emptyset$ for all $w, x \in W$. But clearly $Min\{x\} = \{x\}$, hence the condition is fulfilled. □

Proposition 3.9. *The notion of plausibility CB-bisimulation is closed under arbitrary unions and relational composition.*

3.2 Undefinability

In this subsection we put the new notion of bisimulation to use, addressing the problem of inter-definability between conditional belief and other widely-used operators. We begin with the operator of safe belief as introduced in [6]:

Safe belief: $\mathcal{M}, w \vDash [\leq]\phi$ iff for all $v \leq w$ we have $\mathcal{M}, v \vDash \phi$.

The dual operator is customarily defined as $\langle\leq\rangle\phi := \neg[\leq]\neg\phi$.

Proposition 3.10. *On plausibility models, safe belief is not definable in terms of the conditional belief operator.*

Proof. Suppose $\langle\leq\rangle p$ is definable by a formula α in the language of conditional belief. Consider the two models depicted on the left and right side of the following picture. We indicate within parenthesis the propositional atoms that are true at every world and with Z a CB-bisimulation between the two models (we omit reflexive arrows):

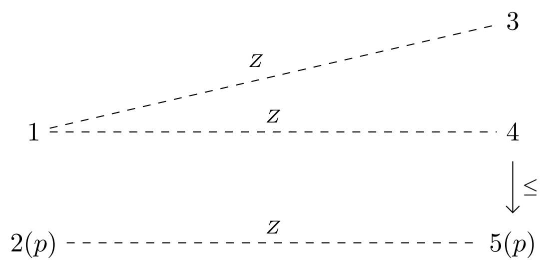

Given that α is a formula in the language of conditional belief, it will be invariant between states that are bisimilar according to a CB-bisimulation. However, $\langle\leq\rangle p$ is true in the second model at 4 but false in the first model at 1; contradiction. □

Notice that the CB-bisimulation Z of this counterexample is not a bisimulation for safe belief, since it fails to satisfy the zig-zag condition: there are worlds 1, 3 and 4 such that $(1,3) \in Z$ and $4 \leq 3$ but no world w such that $w \leq 1$ and $(w,4) \in Z$.

We now address the case of the strong belief operator, also introduced in [6].

Strong belief: $\mathcal{M}, w \vDash Sb\phi$ iff there is $k \in W$ such that $\mathcal{M}, k \vDash \phi$ and for all v, v' if $\mathcal{M}, v \vDash \phi$ and $\mathcal{M}, v' \vDash \neg\phi$ then $v \leq v'$.

Proposition 3.11. *On plausibility models, strong belief is not definable in terms of the conditional belief operator.*

Proof. Again, suppose Sbp is definable by a formula α in the language of conditional belief. Consider the two models displayed below, where Z a CB-bisimulation and the propositional variables are attached to worlds as before (reflexive and transitive arrows omitted):

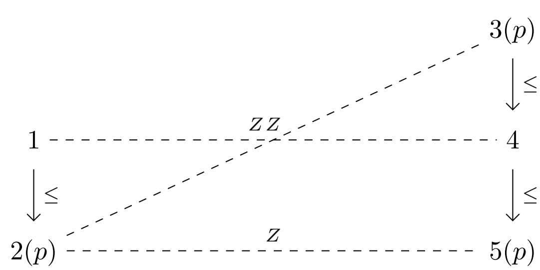

The formula α in the language of conditional belief will be invariant between states that are bisimilar according to a CB-bisimulation; nevertheless, Sbp is true in the first model at 1 but false in the second model at 4, thus α will be true in one world and not in the other: contradiction. □

We now turn our attention to the definability of the conditional belief operator itself. We first warm up with a definition and two auxiliary observations.

Definition 3.12. A BSB-bisimulation, a bisimulation for standard belief and safe belief, is a B-bisimulation satisfying an additional condition, namely the usual zig-zag condition for the $\leq$ relation: given two plausibility models $\mathcal{M}$ and $\mathcal{M}'$ and two worlds w and w' in the respective models, if $(w, w') \in Z$ then

- for every v in $\mathcal{M}$ such that $v \leq w$ there is a v' in $\mathcal{M}'$ such that $(v, v') \in Z$ and $v' \leq w'$
- for every v' in $\mathcal{M}'$ such that $v' \leq w'$ there is a v in $\mathcal{M}$ such that $(v, v') \in Z$ and $v \leq w$

Proposition 3.13. *On plausibility models, if two states w and w' are in a BSB-bisimulation then they are modally equivalent with respect to the language containing the belief and safe belief operators.*

Proof. Straightforward induction on the complexity of the formula. □

Proposition 3.14. *On plausibility models, conditional belief is not definable in terms of the language containing the operators of safe belief and standard belief.*

Proof. Suppose $B^{\neg p}q$ is definable by a formula α in the language of belief, safe belief and public announcement. Consider the two models displayed below, where Z a BSB-bisimulation and the propositional variables are attached to worlds as before:

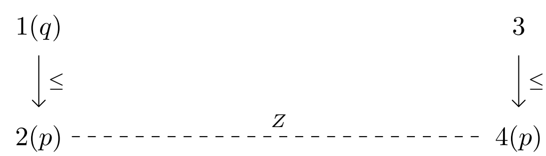

Since 2 and 4 are in a BSB-bisimulation, by Proposition 3.13 they are modally equivalent in the language of belief and safe belief. Thus we can conclude $2 \vDash \alpha$ iff $4 \vDash \alpha$. But $2 \vDash B^{\neg p}q$ and $4 \vDash \neg B^{\neg p}q$, contradiction. □

Notice that the bisimulation used in this counterexample is not a plausibility CB-bisimulation.

4 Conditional belief on evidence models

Evidence models, introduced in [23], are structures capturing the evidence available to an agent in different possible worlds. The evidence available at a world w is represented via a family of sets of possible worlds: intuitively each set in the family constitutes a piece of evidence that the agent can use to draw conclusions at w. They constitute a generalization over plausibility models, but can be collapsed to plausibility models by considering the specialization preorder induced by the sets of evidence, however not without loss of information.[6]

Definition 4.1. An *evidence model* is a tuple $\mathcal{M} = \langle W, E, V \rangle$ with W a non-empty set of worlds, a function $E : W \to \wp(\wp(W))$ and $V : W \to \wp(At)$ a valuation function. A *uniform evidence model* is an evidence model where E is a constant function; in this case we denote it with $\mathcal{E}$.

We indicate with $E(w)$ the set of subsets image of w. We furthermore assume $W \in E(w)$ and $\emptyset \notin E(w)$ for all $w \in W$.

The last requirement ensures that at every possible world the agents has trivial evidence, namely the whole set W, and does not have inconsistent evidence, i.e. the empty set.

Definition 4.2. A *w-scenario* is a maximal family $\mathcal{X} \subseteq E(w)$ having the finite intersection property (abbreviated in 'f.i.p.'), that is, for each finite subfamily $\{X_1, \dots, X_n\} \subseteq \mathcal{X}$ we have $\bigcap_{1 \leq i \leq n} X_i \neq \emptyset$. Given a set $X \subseteq W$ and a collection $\mathcal{X} \subseteq E(w)$, the latter has the f.i.p. relative to X if for each finite subfamily $\{X_1, \dots, X_n\} \subseteq \mathcal{X}^X = \{Y \cap X | Y \in \mathcal{X}\}$ we have $\bigcap_{1 \leq i \leq n} X_i \neq \emptyset$. We say that $\mathcal{X}$ is an *w-X-scenario* if it is a maximal family with the f.i.p. relative to X.

The semantics for the propositional part of the language is the usual one, while for belief and conditional belief we have the clauses:

- $\mathcal{M}, w \vDash B\phi$ iff for every w-scenario $\mathcal{X}$ we have $\mathcal{M}, v \vDash \phi$ for all $v \in \bigcap \mathcal{X}$
- $\mathcal{M}, w \vDash B^{\psi}\phi$ iff every w-$[\![\psi]\!]$-scenario $\mathcal{X}$ we have $\mathcal{M}, v \vDash \phi$ for all $v \in \bigcap \mathcal{X}^{[\![\psi]\!]}$

The notion of bisimulation for the standard belief operator on evidence models establishes a connection between the scenarios of two models:

Definition 4.3. Given two evidence models $\mathcal{M}_1$ and $\mathcal{M}_2$, an *evidence B-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ such that if $(w, w') \in Z$ then

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$;
- then for every w-scenario $\mathcal{X}$ and $x \in \bigcap \mathcal{X}$ there is a w'-scenario $\mathcal{Y}$ and $y \in \mathcal{Y}$ such that $(x, y) \in Z$, and viceversa.

The following result can be proven via the standard line of reasoning.

Theorem 4.4. *Bisimilarity with respect to evidence B-bisimulation entails modal equivalence with respect to the language with only the belief operator. On finite models, modal equivalence with respect to the latter language entails bisimilarity for evidence B-bisimulation.*

[6]See [20, 21] for a discussion on the relationship between the two kinds of models. The sphere systems of [14] also constitute an example of neighborhood models with a close tie to relational structures.

4.1 Evidence CB-bisimulation

We first show that finite evidence models are an example of conditional models by means of two auxiliary lemmas.

Lemma 4.5. *On finite models, suppose $Y \supseteq X$. Then for every w-X-scenario $\mathcal{X}$ there is a w-Y-scenario $\mathcal{Y}$ such that $\mathcal{X} \subseteq \mathcal{Y}$. Conversely, for every w-Y-scenario $\mathcal{Y}$ there is a w-X-scenario $\mathcal{X}$ such that $\mathcal{X} \subseteq \mathcal{Y}$.*

Proof. Clearly $\mathcal{X}$ already has the f.i.p. relative to Y. Enumerate the sets K in $E(w)$ (there are finitely many), then proceed following the enumeration: if $K \in \mathcal{X}$ or $\mathcal{X} \cup \{K\}$ has the f.i.p. relative to Y then put K in $\mathcal{Y}$, otherwise not. Because of the first condition we get $\mathcal{X} \subseteq \mathcal{Y}$, while from the second one we obtain that $\mathcal{Y}$ is a w-Y-scenario.

For the second claim, enumerate the sets in $\mathcal{Y}$: $K_0, \ldots, K_m$. Construct $\mathcal{X}$ in stages beginning from $\mathcal{X}_0 = \emptyset$ and putting $\mathcal{X}_{n+1} = \mathcal{X}_n \cup \{K_n\}$ if $\bigcap \mathcal{X}_n^X \cap K_n \neq \emptyset$. Clearly $\mathcal{X} \subseteq \mathcal{Y}$. To see that $\mathcal{X}$ is maximal with the f.i.p. relative to X suppose that there is $K \notin \mathcal{X}$ such that $\bigcap \mathcal{X}^X \cap K \neq \emptyset$. By construction, if $\bigcap \mathcal{X}^X \cap K \neq \emptyset$ and $K \notin \mathcal{X}$ then $K \notin \mathcal{Y}$, hence by the maximality of $\mathcal{Y}$ it must be that $\bigcap \mathcal{Y}^Y \cap K = \emptyset$. Since $\bigcap \mathcal{X}^X \subseteq \bigcap \mathcal{Y}^Y$ by construction we get a contradiction. Therefore $\mathcal{X}$ is maximal with the f.i.p. relative to X. □

Lemma 4.6. *On finite models, if $Y \supseteq X$, then for every w-X-scenario $\mathcal{X}$ and w-Y-scenario $\mathcal{Y}$ such that $\mathcal{X} \subseteq \mathcal{Y}$ if $y \in \bigcap \mathcal{Y}^Y$ then either $y \in \bigcap \mathcal{X}^X$ or $y \in Y \backslash X$. If no element $y \in \bigcap \mathcal{Y}^Y$ is in $Y \backslash X$ then $\bigcap \mathcal{X}^X = \bigcap \mathcal{Y}^Y$.*

Proof. Say $y \notin Y \backslash X$. Then, since $y \in Y$, it must be that $y \in X$. Since $y \in \bigcap \mathcal{Y}^Y$ we have that $y \in K$ for all $K \in \mathcal{Y}$, and hence $y \in K$ for all $K \in \mathcal{X}$, so $y \in \bigcap \mathcal{X}^X$. We can thus conclude that if $y \notin Y \backslash X$ for all $y \in \bigcap \mathcal{Y}^Y$ then $\bigcap \mathcal{X}^X \supseteq \bigcap \mathcal{Y}^Y$. For the other inclusion suppose $z \in \bigcap \mathcal{X}^X$ but not in $\bigcap \mathcal{Y}^Y$. Then there must be $K \in \mathcal{Y}$ such that $K \notin \mathcal{X}$ and $z \notin K$. By maximality of $\mathcal{X}$ it must be that K has empty intersection with $\bigcap \mathcal{X}^X$. Under the assumption that no element $y \in \bigcap \mathcal{Y}^Y$ is in $Y \backslash X$, the latter fact entails that $\bigcap \mathcal{Y}^Y$ must be empty, contradiction. Hence there can be no element z that is in $\bigcap \mathcal{X}^X$ but not in $\bigcap \mathcal{Y}^Y$, thus $\bigcap \mathcal{X}^X = \bigcap \mathcal{Y}^Y$. □

Proposition 4.7. *Finite evidence models are conditional models, where*

$$f(w, X) = \bigcup \{ \bigcap \mathcal{X}^X | \textit{for} \quad \mathcal{X} \quad w\text{-}X\textit{-scenario} \}$$

Proof. The satisfaction of the first property is ensured by the definition of $\mathcal{X}^X$: since each $\bigcap \mathcal{X}^X$ lies within X, the big union will also be contained in X.

For the second property suppose $Y \supseteq X$ and $f(w, Y) \subseteq X$. If $x \in f(w, Y)$ then there is a w-Y-scenario $\mathcal{Y}$ such that $x \in \bigcap \mathcal{Y}^Y$. By Lemma 4.5 we know there is a w-X-scenario $\mathcal{X}$ such that $\mathcal{X} \subseteq \mathcal{Y}$. By Lemma 4.6 either $x \in \bigcap \mathcal{X}^X$ or $x \in Y \backslash X$. But the latter cannot be because $x \in X$ by assumption, so $x \in \bigcap \mathcal{X}^X$. Then we can conclude that $x \in f(w, X)$.

Now for the other direction. If $x \in f(w, X)$ then there is a w-X-scenario $\mathcal{X}$ such that $x \in \bigcap \mathcal{X}^X$. By Lemma 4.5 there is a w-Y-scenario $\mathcal{Y}$ such that $\mathcal{X} \subseteq \mathcal{Y}$. Because $f(w, Y) \subseteq X$ we can infer that there is no element $y \in \bigcap \mathcal{Y}^Y$ that is in $Y \backslash X$ (that is, $\bigcap \mathcal{Y}^Y \subseteq X$), so by the second part of Lemma 4.6 we can conclude that $\bigcap \mathcal{X}^X = \bigcap \mathcal{Y}^Y$. This gives us $x \in \bigcap \mathcal{Y}^Y$ and thus $x \in f(w, Y)$. □

Notice that, setting $f(w, X) = \bigcup \{\bigcap \mathcal{X}^X | \text{for} \quad \mathcal{X} \quad w\text{-}W\text{-scenario}\}$, the definition of the satisfaction relation for conditional belief on evidence models becomes an instance of the satisfaction relation for conditional modalities given in Section 2. Replacing the new f in Definition 2.5, we obtain a new notion of bisimulation for conditional belief on evidence models.

Definition 4.8. Given two evidence models $\mathcal{M}_1$ and $\mathcal{M}_2$, an *evidence CB-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ such that if $(w, w') \in Z$ then

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$,
- for all $X \subseteq W$ and $X' \subseteq W'$ such that $R[X] \subseteq X'$ and $R^{-1}[X'] \subseteq X$ we have that for every w-X-scenario $\mathcal{X}$ and $x \in \bigcap \mathcal{X}^X$ there is a w'-X'-scenario $\mathcal{Y}$ and $y \in \bigcap \mathcal{Y}^{X'}$ such that $(x, y) \in Z$, and viceversa.

We can now import the previous results on conditional models. First of all, bisimilarity in the latter sense corresponds to modal equivalence on finite evidence models.

Theorem 4.9. *Given two evidence models $\mathcal{M}_1$ and $\mathcal{M}_2$ if $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is an evidence CB-bisimulation, then w and w' are modally equivalent. On finite models, if w and w' are modally equivalent then $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is an evidence CB-bisimulation.*

As for plausibility models, we can infer the results concerning the closure under union and relational composition.

Lemma 4.10. *Every evidence model is a grounded conditional model.*

Proof. Given an evidence model $\mathcal{M}$, it is enough to check that, having defined $f(w, X) = \bigcup \{\bigcap \mathcal{X}^X | \text{for} \quad \mathcal{X} \quad w\text{-}X\text{-scenario}\}$, we have $f(w, \{x\}) \neq \emptyset$ for all $w, x \in W$. It is enough to show that there exist a w-$\{x\}$-scenario $\mathcal{X}$, then by the f.i.p. relative to $\{x\}$ we know that every element of $\mathcal{X}$ must contain x, thus $x \in \bigcap \mathcal{X}^{\{x\}}$ and $f(w, X)$ is not empty. To find the desired w-$\{x\}$-scenario $\mathcal{X}$, take the family of all the sets in $E(w)$ containing x. This family is non-empty, since $W \in E(w)$ for every w in the domain of the model. Clearly this family is maximal with the f.i.p. relative to $\{x\}$ (not only, it is the only one), so we are done. □

Proposition 4.11. *The notion of evidence CB-bisimulation is closed under arbitrary unions and relational composition.*

4.2 Undefinability

Thanks to the now clearly defined bisimulation for conditional belief, we can give a precise argument for the undefinability of conditional belief in terms of standard belief over evidence models.

Proposition 4.12. *On evidence models, conditional belief is not definable in terms of the normal belief operator.*

Proof. Suppose $B^p q$ is definable by a formula α in the language of standard belief. Consider the two models depicted on the left and right side of the following picture, where we indicate within parenthesis the propositional atoms that are true at every world and with Z a B-bisimulation between the two models:

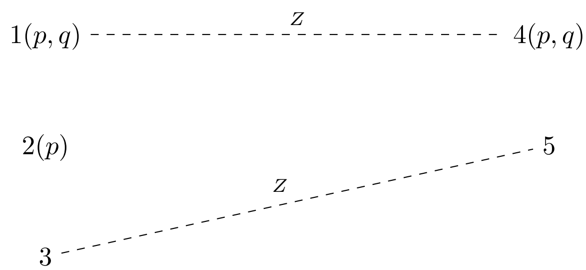

The evidence available at each world is: $E(1) = \{\{1\}, \{3\}, \{2, 3\}, W_1\}$, $E(4) = \{\{4\}, \{5\}, W_2\}$, $E(2) = \{W_1\}$, $E(3) = \{\{3\}, W_1\}$, $E(5) = \{\{5\}, W_2\}$. The reader can check that the relation Z is a B-bisimulation. Given that α is a formula in the language of normal belief, it will be invariant between states that are bisimilar according to a B-bisimulation. However, $B^p q$ is true in the second model at 4 but false in the first model at 1: there is a 1-$[\![p]\!]_{\mathcal{M}_1}$-scenario $\mathcal{X} = \{\{2, 3\}, W_1\}$ and $2 \in \bigcap \mathcal{X}^{[\![p]\!]_{\mathcal{M}_1}}$ such that $2 \not\models q$. Hence we obtain a contradiction. $\square$

Note that the relation Z is not an evidence CB-bisimulation: the sets of worlds satisfying p in the two models satisfy the prerequisites, they are sent into each other by Z, but fail with respect to the main property, since there is a 1-$[\![p]\!]_{\mathcal{M}_1}$-scenario $\mathcal{X}$, and an element in $\bigcap \mathcal{X}^{[\![p]\!]}$, namely 2, that has no bisimilar counterpart in the second model. Another important operator to describe the features of evidence models is the so called evidence modality [23].

Evidence modality: $\mathcal{M}, w \models \Box\phi$ iff there is $K \in E(w)$ such that, for all $v \in K$, $\mathcal{M}, v \models \phi$.

It was shown in [23] that, on evidence models, standard belief cannot be defined in terms of the evidence modality. Since standard belief is definable in terms of conditional belief, we can conclude that also conditional belief is not definable via the evidence modality. Here we show that also the converse is the case.

Proposition 4.13. *On evidence models, the evidence modality is not definable in terms of the conditional belief operator.*

Proof. Suppose $\Box p$ is definable by a formula α in the language of conditional belief. Consider the two models depicted on the left and right side of the following picture, where we indicate within parenthesis the propositional atoms that are true at every world and with Z a CB-bisimulation between the two models:

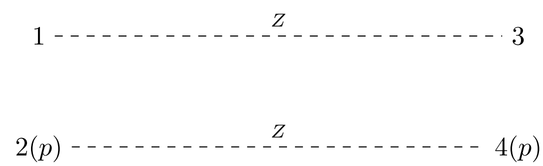

We take both models to be uniform, where $E_1 = \{\{1\}, \{2\}, W_1\}$ and $E_2 = \{W_2\}$. The reader can check that with this evidence the relation Z is a CB-bisimulation. Given that α is a formula in the language of normal belief, it will be invariant between states that are bisimilar according to a CB-bisimulation. Nevertheless, $\Box p$ is true in the first model at 1 but false in the second model at 3: in the first model there is an evidence set contained in the extension of p, namely $\{2\}$, while there is no such set in the second model; contradiction. $\Box$

5 Relativized common knowledge

In this section we depart from the conditional belief operator, to study how this general framework is also applicable to another conditional modality, the so-called relativized common knowledge operator, defined in [22] and studied in [24].

Definition 5.1. A *Kripke model* is a tuple $\mathcal{M} = \langle W, R, V\rangle$ with W a non-empty set of worlds, a relation $R \subseteq W \times W$ and $V : W \to \wp(At)$ a valuation function.

Definition 5.2. Given a relation $R \subseteq W \times W$, denote with R^* its reflexive and transitive closure.

The operator of relativized common knowledge, denoted with $C(\phi, \psi)$, is meant to capture the intuition that every path which consists exclusively of ϕ-worlds ends in a world satisfying ψ. The informal notion of path is captured in the following definition via the reflexive transitive closure.

$$\mathcal{M}, w \vDash C(\phi, \psi) \text{ iff } \mathcal{M}, v \vDash \phi \text{ for all } (w, v) \in (R \cap (W \times [\![\phi]\!]))^*$$

where $\mathcal{M}$ is a Kripke model. The most direct way to encompass this operator is to design a selection function for it.

Proposition 5.3. *Kripke models are conditional models, where* $f(w, X) = \{v | (w, v) \in (R \cap (W \times X))^*\}$ *for all* w.

Proof. Again we check the prerequisites of selection functions in Definition 2.1. Clearly all the worlds reachable with a path in X will also lie in X, hence the first condition on selection functions is given. For the second one, suppose $X \subseteq Y$, $f(w, Y) = \{v | (w, v) \in (R \cap (W \times Y))^*\} \subseteq X$ and take $x' \in f(w, Y)$. Hence there is a chain of Y-worlds leading to x'. We show $x' \in f(w, X)$ by induction on the length of the chain. The base case: since $(w, w) \in f(w, Y) \subseteq X$ so we also have a chain of length 0 contained in X, i.e. $(w, w) \in f(w, X)$. For the inductive step suppose $x \in f(w, X)$ for all $x \in f(w, Y)$ reachable with a chain of Y-worlds of length $\leq n$. Now say $x' \in f(w, Y)$ is reachable with a chain of Y-worlds of length $n + 1$. By $x' \in f(w, Y) \subseteq X$ we know that also $x' \in X$, thus the whole chain is in X and $x' \in f(w, X)$. For the other inclusion, it is straightforward to see that $X \subseteq Y$ immediately entails $f(w, X) \subseteq f(w, Y)$. □

Setting $f(w, X) = \{v | (w, v) \in (R \cap (W \times X))^*\}$, we can repeat the chain of reasoning from Section 2.

Definition 5.4. Given two Kripke models $\mathcal{M}_1$ and $\mathcal{M}_2$, a *bisimulation for relativized common knowledge* or *RCK-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ such that if $(w, w') \in Z$ then

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$,
- for all $X \subseteq W$ and $X' \subseteq W'$ such that $R[X] \subseteq X'$ and $R^{-1}[X'] \subseteq X$ we have that for every $x \in \{v | (w, v) \in (R \cap (W \times X))^*\}$ then there exists a $y \in \{v | (w, v) \in (R \cap (W \times X'))^*\}$ such that $(x, y) \in Z$, and viceversa.

where the last condition essentially states that for every state reachable via a X-path in the first model there must be a corresponding state in the second model reachable via a X'-path.

We can now derive our previous results for this specific setting. In this section and the following one we use 'modal equivalence' meaning with respect to the language containing only the usual propositional connectives and the relativized common knowledge operator.

Theorem 5.5. *Given two Kripke models $\mathcal{M}_1$ and $\mathcal{M}_2$, if $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a RCK-bisimulation, then w and w' are modally equivalent. On finite models, if w and w' are modally equivalent then $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a RCK-bisimulation.*

Proposition 5.6. *RCK-bisimulations are closed under arbitrary unions.*

Lemma 5.7. *Every Kripke model is a grounded conditional model for $f(w, X) = \{v | (w, v) \in (R \cap (W \times X))^*\}$.*

Proof. By reflexive closure we always have $(w, w) \in f(w, X)$ for all X, hence also $(w, w) \in f(w, \{x\})$. □

Proposition 5.8. *RCK-bisimulations are closed under relational composition.*

5.1 Recovering the original Hennessy-Milner Theorem

Now that we are dealing with Kripke models we can finally show how to obtain the original Hennessy-Milner Theorem in our framework. Given a Kripke model $\mathcal{M}$, define the selection function for the *conditional box* $\Box^{\psi}\phi$ as:

$$f(w, X) = \{v | v \in X, (w, v) \in R\}$$

whose satisfaction relation is clearly

$$\begin{aligned} \mathcal{M}, w \models \Box^{\psi}\phi \quad & \text{iff} \quad \mathcal{M}, v \models \psi \quad \text{and} \quad (w, v) \in R \quad \text{entail} \quad \mathcal{M}, v \models \phi \\ & \text{iff} \quad f(w, [\![\psi]\!]) \subseteq [\![\phi]\!] \end{aligned}$$

Proposition 5.9. *Kripke models are conditional models, where* $f(w, X) = \{v | v \in X, (w, v) \in R\}$ *for all* w.

Proof. A straightforward check of the conditions of Definition 2.1. □

We can thus directly infer all the results from Section 2, as the reader can expect. It is easy to see that in this case the unary modality is just the usual $\Box$ operator of basic modal logic and the unary bisimulation is the standard notion of bisimulation for that language. When we instantiate Theorem 2.7 we obtain

Theorem 5.10 (Hennessy-Milner Theorem)**.** *Given two Kripke models* $\mathcal{M}_1$ *and* $\mathcal{M}_2$ *such that*

- $f(w, W) = \{v | (w, v) \in R\}$ *is finite for all* $w \in W$ *(R is image-finite)*
- $f'(w', W') = \{v' | (w', v') \in R'\}$ *is finite for all* $w' \in W'$ *(R' is image-finite)*

if w *and* w' *are modally equivalent in the language of basic modal logic then* $(w, w') \in Z \subseteq W_1 \times W_2$*, where* Z *is a unary bisimulation.*

Note that, defining $f(w, X) = \{v | v \in X, (w, v) \in R\}$, we can extract the precondition X and show that the conditional language is reducible to the unconditional one:

$$\begin{aligned} \mathcal{M}, w \models \Box^{\psi}\phi \quad & \text{iff} \quad \{v | v \in [\![\psi]\!], (w, v) \in R\} \subseteq [\![\phi]\!] \\ & \text{iff} \quad \{v | (w, v) \in R\} \cap [\![\psi]\!] \subseteq [\![\phi]\!] \\ & \text{iff} \quad \{v | (w, v) \in R\} \subseteq [\![\phi]\!] \cup \overline{[\![\psi]\!]} \\ & \text{iff} \mathcal{M}, w \models \Box(\psi \to \phi) \end{aligned}$$

6 Generalization to multi-agent models

We now address the matter: can we extend the analysis of Section 3 to cover the multi-agent case? Given a set of agents A, the language we are interested in will look like

$$\phi ::= \quad p \,|\, \neg\phi \,|\, \psi \wedge \phi \,|\, \psi \rightsquigarrow_a \phi$$

where $\rightsquigarrow_a$ will denote the modality for agent a. This leads to an easy generalization of conditional models.

Definition 6.1. A *multi-agent conditional model* $\mathcal{M}$ is a tuple $\langle W, A, \{f_a\}_{a\in A}, V\rangle$ with W a non-empty set of worlds, A a set of agents, $V : W \to \wp(At)$ a valuation function and for each agent a selection function f_a satisfying the conditions listed in Definition 2.1.

The set of agents is nothing more than a set of labels for different selection functions, co-existing in the same models but essentially independent from each other. Depending on the interpretation, such labels may not refer to different agents but, for example, to different doxastic conditional modalities for a single agent. The semantics clause for the conditional modalities will be:

$$\mathcal{M}, w \vDash \psi \rightsquigarrow_a \phi \quad \text{iff} \quad f_a(w, [\![\psi]\!]_{\mathcal{M}}) \subseteq [\![\phi]\!]_{\mathcal{M}}$$

for every $a \in A$. Likewise, the bisimulation can also be relativized in the same fashion.

Definition 6.2 (Multi-agent Conditional Bisimulation)**.** Given two multi-agent conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$ based n the same set of agents, a *multi-agent conditional bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ such that if $(w, w') \in Z$ then

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$,
- for all $X \subseteq W$ and $X' \subseteq W'$ such that $R[X] \subseteq X'$ and $R^{-1}[X'] \subseteq X$ we have that, for every $a \in A$, for every $x \in f_a(w, X)$ there exists a $y \in f'_a(w', X')$ (where f' is the selection function in $\mathcal{M}_2$) such that $(x, y) \in Z$, and viceversa.

The proofs of the following results are a straightforward generalization of the proofs of the analogous single-agent statements.

Theorem 6.3. *Given two multi-agent conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$, if $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a multi-agent conditional bisimulation, then w and w' are modally equivalent with respect to the logic of conditionals.*

On finite models, if w and w' are modally equivalent then $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a multi-agent conditional bisimulation.

Proposition 6.4. *Multi-agent conditional bisimulations are closed under arbitrary unions.*

Definition 6.5. A multi-agent conditional model is *grounded* if, for all $w \in W$, $a \in A$ and $\{x\}$ in the domain of the model, $f_a(w, \{x\}) \neq \emptyset$.

Proposition 6.6. *Restricted to any class of multi-agent grounded models, the notion of multi-agent conditional bisimulation is closed under relational composition.*

In the following subsection we exemplify multi-agent conditional models by means of a well-known structure called *multi-agent plausibility model* or *epistemic plausibility model*. The reader will see that an analogous generalization can be carried over for evidence models and Kripke models, by introducing evidence functions labelled by agents in the former and by employing relations labelled by agents in the latter.

We devote another subsection to a further possible generalization, namely conditional modalities for *groups* of agents. The prime example here is group conditional belief.

6.1 Conditional belief on multi-agent plausibility models

The multi-agent generalization of plausibility models, epistemic plausibility model are widely used in Formal Epistemology to capture knowledge and beliefs in a multi-agent system.

Definition 6.7. A relation $\leq \subseteq W \times W$ is a *well-preorder* if it is reflexive, transitive and if there exist minimal element for every non-empty subset of W. Given a reflexive and transitive relation $\leq \subseteq W \times W$, we stipulate $\sim$ to be its reflexive closure and, for a given $w \in W$, define its associated *comparability class* to be $[w]_\sim = \{v | v \leq w \quad \text{or} \quad w \leq v\}$. Call $\leq$ *locally well-preordered* if its restriction to each comparability class is well-preordered.

Definition 6.8. A *multi-agent plausibility model* or *epistemic plausibility model* is a tuple $\mathcal{M} = \langle W, A, \{\leq_a\}_{a \in A}, V \rangle$ with W a non-empty set of worlds, $\{\leq_a\}_{a \in A}$ a family of reflexive, transitive and locally well-preordered relations $\leq_a \subseteq W \times W$ indexed by agents and $V : W \to \wp(At)$ a valuation function.

This structure is given these two names because in fact the epistemic indistinguishability relation $\sim_a$ can be defined as the reflexive closure of $\leq_a$. The semantics of the belief and conditional belief operators are then relativized to a specific agent and comparability class:

- $\mathcal{M}, w \vDash B_a \phi$ iff for all $v \in Min([w]_{\sim_a})$ we have $\mathcal{M}, v \vDash \phi$
- $\mathcal{M}, w \vDash B_a^\psi \phi$ iff for all $v \in Min([\![\psi]\!] \cap [w]_{\sim_a})$ we have $\mathcal{M}, v \vDash \phi$

Proposition 6.9. *Multi-agent plausibility models are multi-agent conditional models, where* $f_a(w, X) = Min_{\leq_a}(X \cap [w]_{\sim_a})$.

Proof. We want to ascertain that the newly defined f_a fulfills the prerequisites of selection functions in Definition 2.1. The first condition is again given by the very definition of Min. For the second one, suppose $X \subseteq Y$, $Min_{\leq_a}(Y \cap [w]_{\sim_a}) \subseteq X$ and consider a generic element x' in $Min_{\leq_a}(Y \cap [w]_{\sim_a})$. Clearly from $X \subseteq Y$ we have $X \cap [w]_{\sim_a} \subseteq Y \cap [w]_{\sim_a}$. Thus since $x' \in X$ and then there is no element below x' in $Y \cap [w]_{\sim_a}$ then a fortiori there is no element below it in the subset $X \cap [w]_{\sim_a}$, hence $x' \in Min_{\leq_a}(X \cap [w]_{\sim_a})$. For the other inclusion consider $x' \in Min_{\leq_a}(X \cap [w]_{\sim_a})$; we show x' is also minimal within $Y \cap [w]_{\sim_a}$. By contradiction suppose this is not the case: then there is $z \in Y \cap [w]_{\sim_a}$ such that

$z <_a x'$. Since $\leq_a$ is locally well-preordered, there must be a minimal element $z' \in Min_{\leq_a}(Y \cap [w]_{\sim_a})$ such that $z' \leq_a z$; but by assumption $Min_{\leq_a}(Y \cap [w]_{\sim_a}) \subseteq X$, hence $z' \in X$. This gives us a $z \in X \cap [w]_{\sim_a}$ such that $z' <_a x'$, contradicting the fact that x' is minimal in $X \cap [w]_{\sim_a}$. □

Now that this step is secured, we can reproduce the reasoning of Section 3 to obtain a definition of multi-agent CB-bisimulation (clearly this latter definition will have to impose the same requirements on all the f_a's).

Definition 6.10. Given two multi-agent plausibility models $\mathcal{M}_1$ and $\mathcal{M}_2$, a *multi-agent plausibility CB-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ such that if $(w, w') \in Z$ then

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$,
- for all $X \subseteq W$ and $X' \subseteq W'$ such that $R[X] \subseteq X'$ and $R^{-1}[X'] \subseteq X$ we have that for every a and $x \in f_a(w, X) = Min_{\leq_a}(X \cap [w]_{\sim_a})$ then there exists a $y \in f'_a(w, X) = Min_{\leq'_a}(X' \cap [w]_{\sim_a})$ such that $(x, y) \in Z$, and viceversa (where $\leq'_a$ is the relation associated to a in $\mathcal{M}_2$).

The results on the correspondence between bisimilarity and modal equivalence, closure under union and closure under relational compositions do carry over to this setting. One direction holds by a straightforward generalization of the single agent case.

Theorem 6.11 (Bisimilarity entails modal equivalence)**.** *Given two multi-agent plausibility models $\mathcal{M}_1$ and $\mathcal{M}_2$, if $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a multi-agent plausibility CB-bisimulation, then w and w' are modally equivalent.*

The converse can even be generalized to (pre)image finite models.

Definition 6.12. A multi-agent plausibility model $\mathcal{M}$ is *image finite* if every world has finitely many successors. It is *(pre)image finite* if every world has finitely many successors and finitely many predecessors.

An immediate consequence is that in (pre)image-finite multi-agent plausibility models each information cell $[w]_{\sim_a}$ is finite.

Theorem 6.13 (Modal equivalence entails bisimilarity on finite models)**.** *For any two (pre)image-finite multi-agent plausibility models $\mathcal{M}_1$ and $\mathcal{M}_2$, if w and w' are modally equivalent then $(w, w') \in Z \subseteq W_1 \times W_2$, where Z is a multi-agent plausibility CB-bisimulation.*

Proof. Supposing w and w' are modally equivalent, we want to show that modal equivalence (we call it Z for short) is a multi-agent plausibility CB-bisimulation. Suppose it is not: then there are $X \subseteq W$ and $X' \subseteq W'$ satisfying the requirements, an agent a and $x \in Min_{\leq_a}(X \cap [w]_{\sim_a})$ such that for all $y \in Min_{\leq'_a}(X' \cap [w']_{\sim_a})$ we have $(x, y) \notin Z$. Since the information cells are finite, both $X \cap [w]_{\sim_a}$ and $X' \cap [w']_{\sim_a}$ will be finite. We can then repeat the reasoning of Theorem 2.4, considering partitions on $[w]_{\sim_a}$ and $[w']_{\sim_a}$ and so on. In fact, the latter sets play the part of finite models in this proof. □

Proposition 6.14. *Multi-agent plausibility CB-bisimulations are closed under arbitrary unions and under relational compositions.*

Proof. Follows from the fact that multi-agent plausibility models are grounded, as $f_a(w, \{x\})$ is equal to $\{x\}$ for all a. $\square$

6.2 Group beliefs

As a further question pertaining to the multi-agent realm, we may ask whether this approach suggests how to handle group beliefs, that is, beliefs of a given group of agents. Contrary to knowledge, beliefs of different agents may be mutually inconsistent and thus the problem of aggregating different agents' beliefs into a consistent set is in general non trivial. We display one solution and sketch another, arguing that they are all amenable to the analysis that we showcased in the previous sections.

In the setting of multi-agent plausibility models, one may want to regard the procedure of belief aggregation as a function F that receives the plausibility orderings $\{\leq_b\}_{b\in G}$ of the agents as inputs and returns as output a single plausibility ordering $F(G) = \leq_G$ for the group.[7]

Definition 6.15. Given a set W, call Σ the set of binary relations on W that are reflexive and transitive.

Definition 6.16. Given a set of agents $G \subseteq A$, define $\sim_G = \bigcap_{b\in G} \sim_b$, the intersection of the equivalence relations of the agents in the group.

Definition 6.17. A *group plausibility model* is a tuple $\mathcal{M} = \langle W, A, \{\leq_a\}_{a\in A}, V, F\rangle$ such that $\mathcal{M} = \langle W, A, \{\leq_a\}_{a\in A}, V\rangle$ is a multi-agent plausibility model and $F : \wp(A) \to \Sigma$ is a function called *doxastic aggregation function* such that:

- for all $a \in A$, $F(\{a\}) = \leq_a$
- for all $G \subseteq A$, $F(G)$ is locally well-preordered with respect to $\sim_G$.

Some concrete proposals for protocols tailored to belief merge were put forward in [7]. The choice of the aggregator F is natural in some special settings, for example when we assume that the plausibility orderings of the agents all originate from a "common prior" ordering $\leq$, that is, each agent's plausibility is just the restriction of $\leq$ to her information cells. This kind of assumption is common in situations where we assume a fixed and objective notion of rationality that dictates the 'objective' plausibility ordering $\leq$ of a set of possible worlds.[8] This assumption allows us to intersect the information cells of the agents in the group B and calculate the resulting group plausibility $\leq_G$ with no obstacle of

[7]This procedure bears some similarities with the aggregation of preferences studied by Social Choice Theory, although typically in that context preferences are assumed to be linear orders. See [17] for a classic primer in Social Choice Theory.

[8]See [4] for a notable example in Game Theory and [13] for a recent application of this idea to plausibility orderings.

inconsistent beliefs. Thus in this particular case $F(G)$ is just given by intersecting the plausibility orderings of the agents in the group.

Once F is fixed, the beliefs of the group are then computed from the plausibility ordering $F(G) = \leq_G$ with the usual semantics in terms of minimal elements:

- $\mathcal{M}, w \vDash B_G\phi$ iff for all $v \in Min_{\leq_G}(W)$ we have $\mathcal{M}, v \vDash \phi$
- $\mathcal{M}, w \vDash B_G^{\psi}\phi$ iff for all $v \in Min_{\leq_G}([\![\psi]\!])$ we have $\mathcal{M}, v \vDash \phi$

We do not dwell on the properties that such an aggregator F should have, but point out that, as long as the semantics of group beliefs is given in the customary way from the new relation $\leq_G$, our analysis of the notion of bisimulation can be extended to cover this case.

Observation 6.18. *Group plausibility models are multi-agent plausibility models.*

Proof. Given a group plausibility model relative to a set of agents A, define a multi-agent plausibility model based on the set of agents $\wp(A)$. This is possible due to the protocol F, which provides a legitimate plausibility ordering for each $G \subseteq A$. □

Combining this observation with Proposition 6.9 we immediately obtain for group plausibility models all the results mentioned in the previous subsection.

Another possibility is to use a multi-agent version of evidence models, where each agents has a different set of evidences available at each world. To obtain the group beliefs for the group $G \subseteq A$ we first compute the evidence available to G at w, simply taking the union of all the evidences possessed by the agents in G at w: $E_G(w) = \bigcup_{a \in G} E_a(w)$. The beliefs of the group G can then be computed from E_G using scenarios as in Section 4. Since the information to define E_G is already contained in a multi-agent evidence model and the semantics of the belief operator is of the same kind, this extension still falls within the scope of our approach.

7 Related work

We begin by mentioning a conditional modality that does not fall under the scope of our framework, relevant implication. The proponents of this connective intend to overcome the counterintuitive properties of material implication by a notion of entailment that consider the relevance of the antecedent with respect to the consequent. This particular kind of entailment, that we will denote with $\phi \Rightarrow \psi$ is interpreted on ternary relations with a semantics going back at least to [19].

$$\mathcal{M}, w \vDash \phi \Rightarrow \psi \text{ iff } \forall v, v' \ \mathcal{M}, v \vDash \phi \text{ and } (w, v, v') \in R \text{ entail } \mathcal{M}, v' \vDash \psi.$$

A possible selection function for this conditional modality could be

$$f(w, X) = \{v' | \exists v(w, v, v') \in R, v \in X\}$$

It is not hard to see, however, that such a selection function fails to satisfy the first requirement on conditional models: there is nothing in the definition of $f(w, X)$ to ensure that $v' \in X$. This observation aligns with the intentions of the advocates of relevant implications, who contemplate the possibility of $p \Rightarrow p$ failing at some worlds.

Another notion of bisimulation for conditional belief on multi-agent plausibility models is given in [3]. In said paper the authors prove the correspondence between bisimilarity and modal equivalence, respectively for the languages containing conditional belief and knowledge, safe belief and knowledge, degrees of belief and knowledge. The approach put forward in this paper has the following two distinctive features. First, the bisimulation for conditional belief stems from a general analysis of conditional modalities and it is not tied to a specific model or operator. As we have showed, despite the level of generality the proofs are still relatively simple and transparent. Second, the notion of bisimulation for conditional belief offered here is modular, in the sense that it can be merged with other conditions when we consider languages with additional operators. In contrast, some results in the aforementioned paper depend crucially on the existence of the knowledge operator.[9]

A notion of bisimulation containing a quantification over subsets has been proposed originally in [15], adapted in [26] to epistemic lottery models and later again reshaped to work in the context of epistemic neighborhood models in [25]. Such bisimulations were introduced to deal with probabilities and weights, not conditional modalities. The main difference with the present approach lies in the structure of the quantification. In our case the zig and zag conditions both share the same preconditions, a universal quantification over pairs of subsets satisfying certain prerequisites. In the aforementioned papers each direction has a $\forall\exists$ quantification, stating that for each subset in the first model (usually within the current information cell) there exists a subset in the second model fulfilling certain properties. The connection between these different notions will be subject of further work.

8 Conclusion and further work

We have given a general definition of bisimulation for conditional modalities interpreted on selection functions and proved a HM result, closure under unions and closure under relational composition. To highlight the generality of this framework we have considered a series of examples. First, we described how to specialize the HM result to the case of the unconditional modalities. Second, we argued that our models encompass both Grove's sphere models and Gabbay's

[9] Conversely, the undefinability result of Proposition 3.10 does not hold if we take knowledge into account, that is, we restrict the scope of belief to the current information cell.

models for non-monotonic reasoning. Third, we have shown how this approach provides a solid definition of bisimulation for conditional belief, remarking how such notion abstracts away from standard belief in a coherent way across semantics. Fourth, we treated another conditional operator, relativized common knowledge, underlining how the same results still hold for a different modality with a different semantics. Finally, we showed the flexibility of the approach by generalizing it to multi-agent systems, on one hand encompassing the case of multi-agent plausibility models and on the other hand suggesting how to cover group conditional belief.

The first open question along this line of enquiry concerns infinite models: does modal equivalence entail bisimilarity on some infinite conditional models? We have seen an example of how, in the case of multi-agent plausibility models, the particular structure of the model can determine this answer, but we do not have an answer in the general case yet. We may furthermore ask how many 'classical' results of the model theory for basic modal logic we can we obtain in the setting of conditional modalities. One natural example would be a version of the Van Benthem characterization theorem.

Another group of question arises from considering the new notion of bisimulation from a category-theoretic point of view. From this perspective bisimulations (or simulations, if we drop the viceversa condition) can be regarded as the arrows for a suitable category of models, e.g. the category of plausibility models and plausibility CB-bisimulations. The closure under relational composition, together with the obvious fact that the identity relation is itself such a bisimulation, ensures that the conditions of a category are met. This enables a comparison between categories of models, for example between evidence and plausibility models, allowing for a systematic study of what has been called *tracking* [21], namely the matching of corresponding information dynamics in different classes of models.

References

[1] Mikkel Birkegaard Andersen, Thomas Bolander, and Martin Holm Jensen. Don't plan for the unexpected: Planning based on plausibility models. *Logique et Analyse*, 1(1), 2014.

[2] Mikkel Birkegaard Andersen, Thomas Bolander, Hans van Ditmarsch, and Martin Holm Jensen. Bisimulation for single-agent plausibility models. In Stephen Cranefield and Abhaya Nayak, editors, *AI 2013: Advances in Artificial Intelligence*, volume 8272 of *Lecture Notes in Computer Science*, pages 277–288. Springer International Publishing, 2013.

[3] Mikkel Birkegaard Andersen, Thomas Bolander, Hans van Ditmarsch, and Martin Holm Jensen. Bisimulation and expressivity for conditional belief, degrees of belief, and safe belief. *Synthese*, Forthcoming.

[4] Robert J Aumann. Agreeing to disagree. *The annals of statistics*, pages 1236–1239, 1976.

[5] Alexandru Baltag and Sonja Smets. Conditional doxastic models: A qualitative approach to dynamic belief revision. *Electronic Notes in Theoretical Computer Science*, 165:5–21, 2006.

[6] Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In M. Woolridge G. Bonanno, W. van der Hoek, editor, *Texts in Logic and Games*, volume 3, pages 9–58. Amsterdam University Press, 2008.

[7] Alexandru Baltag and Sonja Smets. Protocols for belief merge: Reaching agreement via communication. *Logic Journal of IGPL*, page jzs049, 2013.

[8] Patrick Blackburn, Maarten De Rijke, and Yde Venema. *Modal Logic*, volume 53. Cambridge University Press, 2002.

[9] Brian F Chellas. Basic conditional logic. *Journal of philosophical logic*, 4(2):133–153, 1975.

[10] Lorenz Demey. Some remarks on the model theory of epistemic plausibility models. *Journal of Applied Non-Classical Logics*, 21(3-4):375–395, 2011.

[11] Lorenz Demey. Looking for the right notion of epistemic plausibility model. *Logic, Philosophy and History of Science in Belgium II. Proceedings of the Young Researchers Days 2010*, pages 73–78, 2012.

[12] Dov M Gabbay. *Theoretical foundations for non-monotonic reasoning in expert systems*. Springer, 1985.

[13] Roosmarijn Goldbach. *Modelling Democratic Deliberation*. Master thesis, ILLC, 2015.

[14] Adam Grove. Two modellings for theory change. *Journal of philosophical logic*, 17(2):157–170, 1988.

[15] Barteld Pieter Kooi. *Knowledge, chance, and change*. PhD thesis, Groningen University, 2003.

[16] David Lewis. Counterfactuals. *Harvard University Press*, 1973.

[17] Amartya K. Sen. Social choice theory. In K. J. Arrow and M. D. Intriligator, editors, *Handbook of Mathematical Economics*, volume 3. North-Holland, 1986.

[18] Robert C Stalnaker. A theory of conditionals. In N. Rescher, editor, *Studies in logical theory*. Blackwell, Oxford, 1968.

[19] Alasdair Urquhart. Semantics for relevant logics. *The Journal of Symbolic Logic*, 37(01):159–169, 1972.

[20] Johan Van Benthem. Dynamic logic for belief revision. *Journal of applied non-classical logics*, 17(2):129–155, 2007.

[21] Johan van Benthem. Tracking information. In K. Bimbo, editor, *Mike Dunn Volume, Outstanding Contributions to Logic*. Springer Science, Dordrecht, to appear.

[22] Johan van Benthem and Barteld Kooi. Reduction axioms for epistemic actions. *AiML-2004: Advances in Modal Logic, number UMCS-04-9-1 in Technical Report Series*, pages 197–211, 2004.

[23] Johan van Benthem and Eric Pacuit. Dynamic logics of evidence-based beliefs. *Studia Logica*, 99(1-3):61–92, 2011.

[24] Johan Van Benthem, Jan Van Eijck, and Barteld Kooi. Logics of communication and change. *Information and computation*, 204(11):1620–1662, 2006.

[25] Jan van Eijck. Update, probability, knowledge and belief. Draft available from homepages.cwi.nl/ jve/papers/15/pdfs/upkb.pdf, December 2015.

[26] Jan van Eijck and François Schwarzentruber. Epistemic probability logic simplified. *Advances in Modal Logic*, 10:158–177, 2014.

[27] Frank Veltman. Logic for conditionals. *PhD thesis, University of Amsterdam*, 1985.