UvA-DARE (Digital Academic Repository)

Bisimulation for conditional modalities

Baltag, A.; Cinà, G.

DOI

[10.1007/s11225-017-9723-2](https://doi.org/10.1007/s11225-017-9723-2)

Publication date

2018

Document Version

Final published version

Published in

Studia Logica

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Baltag, A., & Cinà, G. (2018). Bisimulation for conditional modalities. *Studia Logica*, 106(1), 1–33. <https://doi.org/10.1007/s11225-017-9723-2>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

A. BALTAG
G. CINÀ

Bisimulation for Conditional Modalities

Abstract. We give a definition of bisimulation for conditional modalities interpreted on selection functions and prove the correspondence between bisimilarity and modal equivalence, generalizing the Hennessy–Milner Theorem to a wide class of conditional operators. We further investigate the operators and semantics to which these results apply. First, 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. Second, we treat relativized common knowledge, underlining how the same results still hold for a different modality in a different semantics. Third, we show the flexibility of the approach by generalizing it to multi-agent systems, encompassing the case of multi-agent plausibility models.

Keywords: Bisimulation, Conditional logic, Modal logic.

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 [1–3, 13, 14].

In this paper we attempt to understand what is *conditional* about conditional modalities, proposing a framework that covers all the aforementioned operators. The cornerstone of our approach is 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 [12, 22, 23, 25]; 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

Presented by **Heinrich Wansing**; *Received* April 18, 2016

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 establishing the correspondence between bisimilarity and modal equivalence, usually 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 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 way with respect to the features that the conditional modality can “detect”. This characteristic 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 **un-conditional modality uniformly across semantics**.

We use this list as a benchmark to assess the quality of a notion of bisimulation. In this paper we provide a notion of bisimulation for conditional modalities that complies with the list and prove the correspondence

¹See [10] for a standard reference in Modal Logic.

²For example, a non-structural notion of bisimulation for conditional belief on epistemic plausibility models was given in [13], but was regarded as problematic by Demey himself for the same reason.

between bisimilarity and modal equivalence for the semantics on selection functions.

In the rest of the paper we showcase the versatility of our framework along three directions of applications. First, we demonstrate that it applies to the same operator interpreted on different semantics (as for point 4 in our list), discussing how this approach provides a solid notion of bisimulation for conditional belief. Second, to display that our approach covers more than just conditional belief, we treat the case of another important operator, namely relativized common knowledge. Finally we explain how the central definition and results are amenable for a multi-agent generalization.

2. Bisimulation for Conditional Modalities

Consider the language $\mathcal{L}_{\rightsquigarrow}$ of conditional logic

$$\phi ::= p \mid \neg\phi \mid \psi \wedge \phi \mid \psi \rightsquigarrow \phi$$

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 ”. We also use the standard abbreviations for disjunction $\phi \vee \psi := \neg(\neg\phi \wedge \neg\psi)$, material implication $\phi \rightarrow \psi := \neg\phi \vee \psi$, equivalence $\phi \leftrightarrow \psi := (\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$, top (tautology) $\top := p \vee \neg p$ and bottom (contradiction) $\perp := p \wedge \neg p$ (where p is say the first atomic proposition according to some fixed enumeration of At).

As a semantics, we consider *selection functions* of type $W \times \wp(W) \rightarrow \wp(W)$, along the lines of [22]. Similar considerations can be cast in the more general framework proposed by Chellas in [12], but the generality of neighborhood selection functions is not really needed here, neither to prove our results nor to encompass the examples we mentioned; we thus limit ourselves to Lewis’ original proposal.

DEFINITION 1. A *conditional model* is a tuple $\mathcal{M} = \langle W, f, V \rangle$ with W a non-empty set of worlds, a function $f : W \times \wp(W) \rightarrow \wp(W)$ called *selection function* and $V : W \rightarrow \wp(At)$ 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 ‘relevant’ at w .³ For a given model $\mathcal{M}$, the semantics of the language is defined recursively via an interpretation function $\llbracket - \rrbracket_{\mathcal{M}} : \mathcal{L}_{\rightsquigarrow} \rightarrow \wp(W)$, where for the propositional part of the language the clauses are the usual ones and for conditionals we have the Stalnaker-Lewis semantics:

$$w \in \llbracket \psi \rightsquigarrow \phi \rrbracket_{\mathcal{M}} \quad \text{iff} \quad f(w, \llbracket \psi \rrbracket_{\mathcal{M}}) \subseteq \llbracket \phi \rrbracket_{\mathcal{M}}$$

This encodes the idea that “ ϕ is the case, conditional on ψ ” in a world w iff all the worlds in $\llbracket \psi \rrbracket_{\mathcal{M}}$ that are relevant at w according to f are worlds that satisfy ϕ . As customary, via the interpretation function $\llbracket - \rrbracket_{\mathcal{M}}$ we can define a satisfaction relation $\models \subseteq W \times \mathcal{L}_{\rightsquigarrow}$ putting $\mathcal{M}, w \models \psi$ iff $w \in \llbracket \psi \rrbracket_{\mathcal{M}}$; we will freely switch between the two notations.

To motivate our semantic clauses above, let us first recall that Gabbay [17] argues that our most general intuitions about non-monotonic derivations are captured by consequence relations $\vdash_{NM}$ satisfying the following three conditions, that he calls Reflexivity, Cut and Cautious Monotonicity:

- $\phi \vdash_{NM} \phi$
- $\phi \vdash_{NM} \psi$ and $(\phi \wedge \psi) \vdash_{NM} \theta$ entail $\phi \vdash_{NM} \theta$
- $\phi \vdash_{NM} \psi$ and $\phi \vdash_{NM} \theta$ entail $(\phi \wedge \psi) \vdash_{NM} \theta$

The Cut condition is obviously only a very special case of Gentzen’s Cut rule, and it is sometimes called Cautious Transitivity. We’ll adopt this last terminology, in order to avoid any confusions with the standard Cut rule. In terms of our conditional language, these requirements amount to claiming the validity of the following schemas:

- $\phi \rightsquigarrow \phi$ (Reflexivity)
- $((\phi \rightsquigarrow \psi) \wedge ((\phi \wedge \psi) \rightsquigarrow \theta)) \rightarrow (\phi \rightsquigarrow \theta)$ (Cautious Transitivity)
- $((\phi \rightsquigarrow \psi) \wedge (\phi \rightsquigarrow \theta)) \rightarrow ((\phi \wedge \psi) \rightsquigarrow \theta)$ (Cautious Monotonicity)

In terms of selection functions, the semantic clauses corresponding to these validities are:

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

³Where the vague term ‘relevant’ may assume different interpretations depending on the context: ‘similar’ in sphere models, ‘plausible’ in doxastic logic, ‘normal’ in default reasoning, and so on.

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

Indeed, it is easy to see that these clauses are exactly what is needed to validate the above three schemas. Moreover, they are more general than most other settings for conditional logic, conditional beliefs etc.⁴ Such clauses are in fact equivalent to our requirements on conditional models, which constitute a more compact presentation.

PROPOSITION 2. *Conditional models are exactly those satisfying Gabbay's requirements, when formulated in terms of selection functions.*

It is clear that Reflexivity is exactly our clause (1) in Definition 1; the following two lemmas show that, in the presence of Reflexivity, Cautious Transitivity and Cautious Monotonicity correspond to the two inclusions in our clause (2).

LEMMA 3. *Cautious Transitivity entails the left-to-right inclusion of condition (2) in Definition 1. In presence of Reflexivity, the latter condition entails Cautious Transitivity.*

PROOF. Suppose $X \subseteq Y$ and $f(w, Y) \subseteq X$. Substitute X' with $f(w, X)$ in the definition of Cautious Transitivity: 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 Cautious Transitivity 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: on 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 4. *Cautious Monotonicity entails the right-to-left inclusion of condition (2) in Definition 1. In presence of Reflexivity, the converse also holds.*

⁴One can show that Lewis' 'sphere models' are an example of conditional models. The later modification due to Grove [18], in order to model belief revision, is also a special case; interestingly, the appropriate selection function is suggested by Grove himself in [18] p. 159. As we will show, our clauses are weaker than the semantic requirements of conditional doxastic logic. A further example are the models for non-monotonic logics. Our conditions are more general than the models of, for example, the non-monotonic system P of Kraus, Lehmann and Magidor [21] or the conditional logic introduced by Halpern in [19].

PROOF. Suppose $X \subseteq Y$ and $f(w, Y) \subseteq X$. Replacing X' with $f(w, Y)$ in the definition of Cautious Monotonicity we obtain $f(w, Y) \subseteq X$ and $f(w, Y) \subseteq f(w, Y)$. The former is given by assumption and the latter is a tautology, so applying Cautious Monotonicity we can 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, obtaining $f(w, X \cap Y) \subseteq f(w, Y)$. ■

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 from the perspective of our conditional language. First we lay out some notation: given a relation $R \subseteq W \times W'$, $X \subseteq W$ and $X' \subseteq W'$ define

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

DEFINITION 5. (*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

- $V(w) = V(w')$,
- for all $X \subseteq W_1$ and $X' \subseteq W_2$ such that $Z[X] \subseteq X'$ and $Z^{-1}[X'] \subseteq X$ we have that for every $x \in f^1(w, X)$ there exists a $y \in f^2(w', X')$ (where f^2 is the selection function in $\mathcal{M}_2$) such that $(x, y) \in Z$, and vice versa.

The non-standard part of this definition, namely the quantification over subsets X and X' together with the additional requirement $Z[X] \subseteq X'$ and $Z^{-1}[X'] \subseteq X$, is meant to handle the precondition ψ in the conditional $\psi \rightsquigarrow \phi$. One would want the sets X and X' in the definition to be modally definable. However, to ensure that those sets are modally definable we would have to quantify over the formulas in the language and this would clash with the desideratum of having a structural bisimulation. Our solution is to replace “modally definable” with a structural condition that is close enough.⁵

⁵As a consequence of this quantification over subsets, the time needed to check for a bisimulation can be exponential on the size of the input models; this is however not surprising, since the bisimulation intends to capture an operator with preconditions.

The relation of (*conditional*) *bisimilarity* is defined as the existence of a conditional bisimulation: two states w and w' are bisimilar iff there exists a conditional bisimulation Z such that $(w, w') \in Z$. In other words, the relation of bisimilarity between models $\mathcal{M}_1$ and $\mathcal{M}_2$ is the union of all the bisimulation relations between these models. The next result implies that *bisimilarity is itself a bisimulation*, and hence it is the *largest bisimulation* between two given models.

PROPOSITION 6. *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 $X' \subseteq W_2$ are such that $\bigcup_{i \in I} Z_i[X] \subseteq X'$ and $(\bigcup_{i \in I} Z_i)^{-1}[X'] \subseteq X$.

To establish that $\bigcup_{i \in I} Z_i$ is a conditional bisimulation we need to show that for every $x \in f^1(w, X)$ there is $y \in f^2(w', X')$ 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 $x \in f^1(w, X)$ there is $y \in f^2(w', X')$ such that $(x, y) \in Z_i$. The latter fact entails $(x, y) \in \bigcup_{i \in I} Z_i$, we are done. The converse direction is proved symmetrically. ■

The last proposition secures only half of our second desideratum for a notion of bisimulation (see list in the “Introduction”). We postpone the matter of relational composition to the next subsection. The next thing to check is that our definition is suited to our conditional language: *bisimilar states satisfy the same conditional formulas*.

DEFINITION 7. ($\mathcal{L}_{\rightsquigarrow}$ -equivalence) We say that two worlds w, w' in conditional models $\mathcal{M}, \mathcal{M}'$ are $\mathcal{L}_{\rightsquigarrow}$ -equivalent iff they satisfy the same formulas in $\mathcal{L}_{\rightsquigarrow}$: i.e. for every $\phi \in \mathcal{L}_{\rightsquigarrow}$, $\mathcal{M}, w \models \phi$ holds iff $\mathcal{M}', w' \models \phi$ holds.

THEOREM 8. (Bisimilarity entails $\mathcal{L}_{\rightsquigarrow}$ -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 $\mathcal{L}_{\rightsquigarrow}$ -equivalent.*

PROOF. The proof is by induction on the structure of formulas; the case of $p, \neg, \wedge$ are treated as usual, we only show the case of the conditional modality. Suppose Z is a conditional bisimulation, $(w, w') \in Z$ and $\mathcal{M}_1, w \models \psi \rightsquigarrow$

ϕ . Note that by induction hypothesis on ψ we have that $\llbracket \psi \rrbracket_{\mathcal{M}_1}$ and $\llbracket \psi \rrbracket_{\mathcal{M}_2}$ satisfy the right requirements and therefore can act as X and X' in the preconditions of the bisimulation property. Because of $w \models \psi \rightsquigarrow \phi$ we have $f^1(w, \llbracket \psi \rrbracket_{\mathcal{M}_1}) \subseteq \llbracket \phi \rrbracket_{\mathcal{M}_1}$. Now consider $v' \in f^2(w', \llbracket \psi \rrbracket_{\mathcal{M}_2})$. By vice versa of the bisimulation property we know that there exists a $v \in f^1(w, \llbracket \psi \rrbracket_{\mathcal{M}_1})$ such that $(v, v') \in Z$. By assumption and induction hypothesis on ϕ we get $\mathcal{M}_2, v' \models \phi$. Since v' was generic we can conclude that $f^2(w', \llbracket \psi \rrbracket_{\mathcal{M}_2}) \subseteq \llbracket \phi \rrbracket_{\mathcal{M}_2}$, thus $\mathcal{M}_2, w' \models \psi \rightsquigarrow \phi$. For the converse use the other direction of the bisimulation property. ■

Our next theorem is the key result of this paper, providing a partial converse to the previous result. This is an analogue of the Hennessy–Milner–van Benthem theorem from modal logic, saying that on finite models bisimilarity completely captures $\mathcal{L}_{\rightsquigarrow}$ -equivalence.

THEOREM 9. ($\mathcal{L}_{\rightsquigarrow}$ -equivalence entails bisimilarity on finite models) *Given two finite conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$, if w and w' are $\mathcal{L}_{\rightsquigarrow}$ -equivalent then they are bisimilar.*

PROOF. We show that the relation Z of $\mathcal{L}_{\rightsquigarrow}$ -equivalence is a (conditional) bisimulation. First a preliminary observation. Suppose X and X' are 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 $\mathcal{L}_{\rightsquigarrow}$ -equivalent counterparts in X
- $W_1 \setminus (X \cup A)$

Notice how the conditions on X and X' ensure that the elements in A do not have any counterpart in $\mathcal{M}_2$: $a \in A$ cannot have a $\mathcal{L}_{\rightsquigarrow}$ -equivalent counterpart in X' , or otherwise a would be already in X ; on the other hand a cannot have an $\mathcal{L}_{\rightsquigarrow}$ -equivalent counterpart in $W_2 \setminus X'$ 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 $\mathcal{L}_{\rightsquigarrow}$ -equivalent to the elements outside X' , thus the elements in $X \cup A$ are also not $\mathcal{L}_{\rightsquigarrow}$ -equivalent to the elements outside X' . Since we are dealing with finite models we can enumerate the elements in $X \cup A$, call them $x_1, \dots, x_n$. Similarly, we can put the elements of $W_2 \setminus X'$ and $W_1 \setminus (X \cup A)$ all together in a finite list $y_1, \dots, y_m$. By our assumptions and definition of

the partition we know that every element in $X \cup A$ is not $\mathcal{L}_{\rightsquigarrow}$ -equivalent to any element in $W_2 \setminus X'$ or $W_1 \setminus (X \cup A)$. So 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 $x_i \models \psi_{ij}$ and $y_j \not\models \psi_{ij}$. 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 $(W_2 \setminus X') \cup (W_1 \setminus (X \cup A))$. Symmetrically, there must be a formula γ' that is true at $X' \cup A'$ and false at $W_1 \setminus X \cup (W_2 \setminus (X' \cup A'))$. Now consider the formula

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

Let us have a closer look at the extension $\llbracket \alpha \rrbracket_{\mathcal{M}_1}$ 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 $X \cup A$ and false in $W_1 \setminus (X \cup A)$. 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 interpretation of α in $\mathcal{M}_2$ is $X' \cup A'$.

Say now that $(w, w') \in Z$ and suppose Z does not satisfy the bisimulation property for sets X and X' : this means that there is an $x \in f^1(w, X)$ such that for all $y \in f^2(w', X')$ we have $(x, y) \notin Z$ (assumption $\star$).

We first prove that $f^1(w, \llbracket \alpha \rrbracket_{\mathcal{M}_1}) \subseteq X$. Let z be in $f^1(w, \llbracket \alpha \rrbracket_{\mathcal{M}_1})$. Since $f^1(w, \llbracket \alpha \rrbracket_{\mathcal{M}_1}) \subseteq \llbracket \alpha \rrbracket_{\mathcal{M}_1}$ by the first property of selection functions, we know that z must be either in X or in A . If $z \in f^1(w, \llbracket \alpha \rrbracket_{\mathcal{M}_1})$ is in A , since we know that elements in A are not $\mathcal{L}_{\rightsquigarrow}$ -equivalent to any element in W_2 , we can build a formula β that is false at z and true everywhere in W_2 , thus a fortiori in $f^2(w', \llbracket \alpha \rrbracket_{\mathcal{M}_2})$. This gives us the contradiction that we want: $w \models \neg(\alpha \rightsquigarrow \beta)$ and $w' \models \alpha \rightsquigarrow \beta$. We can thus conclude that $f^1(w, \llbracket \alpha \rrbracket_{\mathcal{M}_1}) \subseteq X$. This is enough to apply the second property of selection functions and conclude that $f^1(w, \llbracket \alpha \rrbracket_{\mathcal{M}_1}) = f^1(w, X)$.

This ensures that the element $x \in f(w, X)$ given by assumption $\star$ is indeed also in $f^1(w, \llbracket \alpha \rrbracket_{\mathcal{M}_1})$. If we now look at the set $f^2(w', \llbracket \alpha \rrbracket_{\mathcal{M}_2})$, repeating a reasoning similar to the one just outlined we can conclude that $f^2(w', \llbracket \alpha \rrbracket_{\mathcal{M}_2}) = f^2(w', X')$. By assumption we have that x is not $\mathcal{L}_{\rightsquigarrow}$ -equivalent to any $y \in f^2(w', X')$. We can thus build a formula β that is false at x and true everywhere in $f^2(w', \llbracket \alpha \rrbracket_{\mathcal{M}_2})$; this gives us the contradiction $w \models \neg(\alpha \rightsquigarrow \beta)$ and $w \models \alpha \rightsquigarrow \beta$. ■

2.1. Closure Under Composition

Closure under relational composition turns out to be more tricky: we need bisimulation to ‘transfer’ preconditions in a coherent manner. In this subsection we propose a sufficient condition to obtain closure under relational composition.

DEFINITION 10. Given a conditional model $\mathcal{M} = \langle W, f, V \rangle$, define $W^w = \bigcup_{Y \subseteq W} f(w, Y)$. The model $\mathcal{M}$ is *grounded* if, for any $X \subseteq W$, $X \cap W^w \neq \emptyset$ entails $f(w, X) \neq \emptyset$.

If $f(w, X)$ selects the worlds in X that are ‘relevant’ at w , the set W^w is the collection of all the relevant worlds for w , taking into account all possible preconditions.⁶ A conditional model is grounded when, given a precondition X that is consistent with the collection of all worlds relevant for w , the selection function returns a non-empty set of relevant worlds for w in X . The idea that conditioning with sets that are consistent with the current information should yield consistent results is widespread in Formal Epistemology, see for example Lewis in [22] and Board in [11]. The following equivalent definition of grounded models will be useful in later sections.

LEMMA 11. A model $\mathcal{M}$ is grounded iff, for any $x \in W$, $x \in W^w$ entails $f(w, \{x\}) \neq \emptyset$.

PROOF. The new condition is a special case of the main definition when instantiated to singletons, so one direction is given. For the right-to-left direction, suppose by contradiction that $X \cap W^w \neq \emptyset$ and $f(w, X) = \emptyset$. Let $x \in X \cap W^w$: we have $x \in W^w$ and thus $f(w, \{x\}) \neq \emptyset$. However, $\{x\} \subseteq X$ and $f(w, X) = \emptyset \subseteq \{x\}$ trigger the second condition on conditional models, which states that $f(w, X) = f(w, \{x\})$, contradiction. ■

DEFINITION 12. A conditional bisimulation $Z \subseteq W_1 \times W_2$ is *diffuse* if for every $x \in W_1$ there are $w \in W_1$ and $w' \in W_2$ such that $(w, w') \in Z$ and $x \in W_1^w$, and vice versa.

The idea of diffuse bisimulations is that every elements belongs to a set of relevant worlds which is connected to the other model.

DEFINITION 13. A relation $R \subseteq X \times Y$ is *two-ways surjective* 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 $(x, y) \in R$.

LEMMA 14. Any diffuse conditional bisimulation between grounded models is two-ways surjective.

⁶The notation is borrowed by Board, see [11, p. 56].

PROOF. Let $\mathcal{M}_1$ and $\mathcal{M}_2$ be such models and suppose $Z \subseteq W_1 \times W_2$ is a conditional bisimulation. Suppose moreover that Z is not two-ways surjective, 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: $Z[\{x\}] = \emptyset \subseteq \emptyset$ and $Z^{-1}[\emptyset] = \emptyset \subseteq \{x\}$.

Since the bisimulation is diffuse we know that there are $w \in W_1$ and $w' \in W_2$ such that $(w, w') \in Z$ and $x \in W_1^w$. From the latter fact we infer that $\{x\} \cap W_1^w \neq \emptyset$, thus by the fact that $\mathcal{M}_1$ is grounded we conclude that $f^1(w, \{x\}) \neq \emptyset$. Thanks to the assumption $f^1(w, \{x\}) \subseteq \{x\}$ on conditional models we can conclude $f^1(w, \{x\}) = \{x\}$. Since $(w, w') \in Z$, we must conclude that for every $z \in f^1(w, \{x\})$ there is a $y \in f^2(w', \emptyset)$ such that $(z, y) \in Z$. However, by the first condition on selection function we have $f^2(w', \emptyset) \subseteq \emptyset$, so there can be no counterpart for x , contradiction. The other direction is proved analogously. ■

PROPOSITION 15. *Restricted to any class of grounded models, the notion of diffuse 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 diffuse conditional bisimulations connecting them. To show that their relational composition $Z_1; Z_2$ is also a diffuse 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 two-ways surjective. 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 property of conditional bisimulation, suppose $(w, w'') \in Z_1; Z_2$. By definition it 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 $Z_1; Z_2[X] \subseteq X''$ and $(Z_1; Z_2)^{-1}[X''] \subseteq X$.

What we need to show is that for every $x \in f^1(w, X)$ there is a $z \in f^3(w'', X'')$ such that $(x, z) \in Z_1; Z_2$. Define

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

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 two-ways surjectivity 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 $Z_1; Z_2[X] \subseteq X''$ we have $z \in X''$. But then by the latter fact and $(x, z) \in Z_1; Z_2$, coupled with $(Z_1; Z_2)^{-1}[X''] \subseteq X$, 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 $(Z_1; Z_2)^{-1}[X''] \subseteq X$ 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 $x \in f^1(w, X)$ there is $y \in f^2(w', X')$ 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 $y \in f^2(w', X')$ there is $z \in f^3(w'', X'')$ such that $(y, z) \in Z_2$. Concatenating this with the previous result we get the desired conclusion: for every $x \in f^1(w, X)$ there is a $z \in f^3(w'', X'')$ such that $(x, z) \in Z_1; Z_2$. The converse is proved symmetrically.

It remains to show that $Z_1; Z_2$ is diffuse. Let $x \in W_1$, we need to find $w \in W_1$ and $w'' \in W_3$ such that $(w, w'') \in Z_1; Z_2$ and $x \in W_1^w$. Since Z_1 is diffuse, we know there are $w \in W_1$ and $w' \in W_2$ such that $(w, w') \in Z_1$ and $x \in W_1^w$. By Z_2 being two-ways surjective we know there is w'' such that $(w', w'') \in Z_2$, thus $(w, w'') \in Z_1; Z_2$. The converse is proved symmetrically. ■

PROPOSITION 16. *Restricted to grounded models and diffuse conditional bisimulations, the relation of bisimilarity 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 diffuse conditional bisimulation. The definition of diffuse conditional bisimulation is itself symmetric, hence the converse of a diffuse conditional bisimulation is always a diffuse conditional bisimulation; the symmetry for bisimilarity follows. As for transitivity, Proposition 15 ensures that if there are two diffuse conditional bisimulations Z_1 and Z_2 such that $(w, w') \in Z_1$ and $(w', w'') \in Z_2$ then there is a diffuse conditional bisimulation containing the pair (w, w'') , namely the relational composition $Z_1; Z_2$. ■

We will see that in the next two sections these restrictions vanish, because in those particular settings all models are grounded and all bisimulations are diffuse. In later sections we will encounter examples where the restriction does limit the scope of our results; we then characterize the grounded models and diffuse bisimulations in those particular contexts.

3. Plausibility Models

We now turn to applications, discussing our first example of conditional modality: *conditional belief* interpreted on plausibility models. Plausibility models are widely used in formal epistemology [4, 6], while their introduction can be traced back at least to [22]. They consist of a carrier set, to be understood as a collection of possible worlds, and a preorder for each world, intuitively representing how an agent ranks the possible scenarios in terms of plausibility, from the perspective of the current world.

DEFINITION 17. A *plausibility model* is a tuple $\mathcal{M} = \langle W, \{\leq_w\}_{w \in W}, V \rangle$ with W a non-empty set of worlds, a family of reflexive and transitive relations $\leq_w \subseteq W \times W$ and $V : W \rightarrow \wp(At)$ a valuation function.

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

$$Min_w(X) = \{v \in X \mid \neg \exists z \in X \text{ s.t. } z <_w v\}$$

We can think of $Min_w(X)$ as the set of most plausible worlds in X with respect to w .⁷ When we want to specify the ordering we write $Min_{\leq_w}(X)$.

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, usually written as $B^\psi\phi$. The standard belief operator can be defined via the conditional one as $B^\top\phi$. On plausibility models the semantic clause for belief and conditional belief are:

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

The notion of bisimulation for the standard belief operator on plausibility models, together with the corresponding Theorem, are both folklore.

DEFINITION 18. Given two plausibility models $\mathcal{M}_1$ and $\mathcal{M}_2$, a *plausibility B-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ s.t. if $(w, w') \in Z$ then

- $V(w) = V(w')$;
- for every $x \in Min_w W_1$ there is $y \in Min_{w'} W_2$ such that $(x, y) \in Z$, and vice versa.

⁷We sometimes omit the parenthesis in $Min(X)$ in what follows.

THEOREM 19. *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.*

The proof is a straightforward variation of the standard Hennessy–Milner argument.

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 20. A plausibility model $\mathcal{M}$ is *well-founded* if it contains no infinite descending chains.⁸

PROPOSITION 21. *Well-founded plausibility models are conditional models, where $f(w, X) = \text{Min}_w X$ for all w .*

PROOF. We need to check that the newly defined f fulfills the prerequisites of selection functions in Definition 1. The first condition on selection functions is fulfilled by the very definition of Min_w . For the second one, suppose $X \subseteq Y$, $\text{Min}_w Y \subseteq X$ and take $x' \in \text{Min}_w Y$. 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 $x' \in \text{Min}_w X$. For the other inclusion take $x' \in \text{Min}_w X$; we show x' is also minimal for Y . By contradiction, suppose there is $z \in Y \setminus X$ such that $z <_w x'$. Since we are in a well-founded model there must be a minimal element $z' \in \text{Min}_w Y$ such that $z' \leq_w z$; but by assumption $\text{Min}_w Y \subseteq X$, hence $z' \in X$ and $z' < x'$, contradicting the fact that x' is minimal in X . ■

Notice that, setting $f(w, X) = \text{Min}_w X$, the definition of the satisfaction relation for conditional belief becomes an instance of the satisfaction relation for conditional modalities given in Section 2. If we now replace the new f in Definition 5, we obtain a new notion of bisimulation for conditional belief on plausibility models.

⁸Equivalently, assuming the axiom of choice, if every non empty subsets has minimal elements.

DEFINITION 22. Given two plausibility models $\mathcal{M}_1$ and $\mathcal{M}_2$, a *plausibility CB-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ s.t. if $(w, w') \in Z$ then

- $V(w) = V(w')$,
- for all $X \subseteq W_1$ and $X' \subseteq W_2$ such that $Z[X] \subseteq X'$ and $Z^{-1}[X'] \subseteq X$ we have that for every $x \in \text{Min}_w X$ then there exists a $y \in \text{Min}_{w'} X'$ such that $(x, y) \in Z$, and vice versa.

Since finite plausibility models are well-founded, we can now transfer the results of Section 2 on the correspondence between bisimilarity and modal equivalence. Throughout this section and the following one we use ‘modal equivalence’ meaning with respect to the language of conditional belief.

THEOREM 23. *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 plausibility 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. First note that, with the current definition of f , the notation W^w trivializes: $W^w = \bigcup_{Y \subseteq W} f(w, Y) = \bigcup_{Y \subseteq W} \text{Min}_w(Y) = \bigcup_{\{x\} \subseteq W} \text{Min}_w(\{x\}) = W$. In other words, all the worlds in the model are relevant for every $w \in W$.

LEMMA 24. *Every well-founded plausibility model is a grounded conditional model and every plausibility CB-bisimulation is diffuse.*

PROOF. Given a well-founded plausibility model $\mathcal{M}$ and $X \subseteq W$, if $X \cap W^w \neq \emptyset$ then $X \cap W \neq \emptyset$ so actually $X \neq \emptyset$. So by well-foundedness $f(w, X) = \text{Min}X \neq \emptyset$. This shows that the model is grounded. For the second part of the claim, let $Z \subseteq W_1 \times W_2$ be a plausibility CB-bisimulation and $x \in W_1$. Since the bisimulation is non-empty, there are $(w, w') \in Z$ and furthermore $x \in W_1 = W_1^w$, hence Z is diffuse. ■

PROPOSITION 25. *On the class of well-founded plausibility models, 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. For the rest of this section we employ plausibility models

where $\leq_w$ is the same for all w (we thus remove the subscript). We begin with the operator of safe belief introduced in [5]:

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

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

PROPOSITION 26. *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 use rounded rectangles to set apart the worlds of the left-hand model), where we draw $a \leftarrow b$ to mean $a \leq b$. We omit reflexive arrows. We indicate within parenthesis the propositional atoms that are true at every world and with Z a CB-bisimulation between the two models:

To check that Z is a CB-bisimulation, notice that only three pairs of sets fulfill the right precondition: $(\{1\}, \{3, 4\})$, $(\{2\}, \{5\})$ and $(\{1, 2\}, \{3, 4, 5\})$. It is easy to see that the minimal elements of these pairs are connected by the bisimulation. 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, as the reader can check; 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, 5 and 4 such that $(1, 4) \in Z$ and $5 \leq 4$ but no world w such that $w \leq 1$ and $(w, 5) \in Z$. We now address the case of the strong belief operator, also introduced in [5].

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

PROPOSITION 27. *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:

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 28. 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 $v \in W$ such that $v \leq w$ there is a $v' \in W'$ such that $(v, v') \in Z$ and $v' \leq w'$
- for every $v' \in W'$ such that $v' \leq w'$ there is a $v \in W$ such that $(v, v') \in Z$ and $v \leq w$

PROPOSITION 29. *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 30. *On plausibility models, conditional belief is not definable in terms of the language containing the operators of safe belief and standard belief.*

PROOF. Suppose $B^{-p}q$ is definable by a formula α in the language of belief and safe belief. 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 29 they are modally equivalent in the language of belief and safe belief. Thus we can conclude $2 \models \alpha$ iff $4 \models \alpha$. But $2 \models B^{-p}q$ and $4 \not\models B^{-p}q$, contradiction. ■

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

4. Evidence Models

We now change the semantics of the belief operator to evidence models, showing how the passage to conditional belief in this different setting follows the same pattern as in plausibility models; this allows us to conclude that the generalization from un-conditional to conditional modality works uniformly across semantics (see item 4 in our checklist in the “Introduction”).

Evidence models, introduced in [8], 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.⁹

⁹Evidence models contain information that is lost in the transition to plausibility models; such information is captured by operators such as the evidence modality. See [6, 7] for a discussion on the relationship between the two kinds of models. The sphere systems

DEFINITION 31. An *evidence model* is a tuple $\mathcal{M} = \langle W, E, V \rangle$ with W a non-empty set of worlds, a function $E : W \rightarrow \wp(\wp(W))$ and $V : W \rightarrow \wp(At)$ a valuation function.

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 32. 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 $\{X_1, \dots, X_n\} \subseteq \mathcal{X}^X = \{Y \cap X \mid Y \in \mathcal{X}\}$ 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 belief and conditional belief on evidence models is:

- $\mathcal{M}, w \models B\phi$ iff for every w -scenario $\mathcal{X}$ we have $\mathcal{M}, v \models \phi$ for all $v \in \bigcap \mathcal{X}$
- $\mathcal{M}, w \models B^\psi \phi$ iff for every w - $\llbracket \psi \rrbracket$ -scenario $\mathcal{X}$ we have $\mathcal{M}, v \models \phi$ for all $v \in \bigcap \mathcal{X}^{\llbracket \psi \rrbracket}$

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

DEFINITION 33. Given two evidence models $\mathcal{M}_1$ and $\mathcal{M}_2$, an *evidence B-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ s.t. if $(w, w') \in Z$ then

- for all $p \in At$, $p \in V(w)$ iff $p \in V(w')$;
- 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 vice versa.

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

THEOREM 34. *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.*

Footnote 9 continued

of [18] 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 35. *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. Let $\mathcal{X}$ be a w - X -scenario. 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 $\mathcal{Y}$: $K_0, \dots, K_m$. 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 36. *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 $y \in Y \setminus X$. If no element $y \in \bigcap \mathcal{Y}^Y$ is in $Y \setminus X$ then $\bigcap \mathcal{X}^X = \bigcap \mathcal{Y}^Y$.*

PROOF. Say $y \in \bigcap \mathcal{Y}^Y$ and $y \notin Y \setminus X$. 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 $y \notin Y \setminus X$ for all $y \in \bigcap \mathcal{Y}^Y$, $\bigcap \mathcal{X}^X \supseteq \bigcap \mathcal{Y}^Y$. 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 $Y \setminus X$, 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 37. *Finite evidence models are conditional models, where*

$$f(w, X) = \bigcup \left\{ \bigcap \mathcal{X}^X \mid \mathcal{X} \text{ is a } w\text{-}X\text{-scenario} \right\}$$

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 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 35 we know there is a w - X -scenario $\mathcal{X}$ such that $\mathcal{X} \subseteq \mathcal{Y}$. By Lemma 36 either $x \in \bigcap \mathcal{X}^X$ or $x \in Y \setminus X$. 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 35 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 $Y \setminus X$ (that is, $\bigcap \mathcal{Y}^Y \subseteq X$), so by the second part of Lemma 36 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 $f(w, X) = \bigcup \{\bigcap \mathcal{X}^X \mid \text{for } \mathcal{X} \text{ } w\text{-}X\text{-scenario}\}$, 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 5, we obtain a new notion of bisimulation for conditional belief on evidence models.

DEFINITION 38. Given two evidence models $\mathcal{M}_1$ and $\mathcal{M}_2$, an *evidence CB-bisimulation* is a non-empty relation $Z \subseteq W_1 \times W_2$ s.t. if $(w, w') \in Z$ then

- $V(w) = V(w')$,
- for all $X \subseteq W_1$ and $X' \subseteq W_2$ such that $Z[X] \subseteq X'$ and $Z^{-1}[X'] \subseteq X$ 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 vice versa.

We can now specialize the results of Section 2: bisimilarity in the latter sense corresponds to modal equivalence on finite evidence models.

THEOREM 39. *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 evidence 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. Also in this context the definition of f renders the notation W^w trivial.

LEMMA 40. *For every evidence model $\mathcal{M}$ and $x \in W$, $f(w, \{x\}) = \{x\}$.*

PROOF. Given an evidence model $\mathcal{M}$, we check that $f(w, \{x\}) \neq \emptyset$ for all $w, x \in W$, for $f(w, X) = \bigcup \{\bigcap \mathcal{X}^X \mid \text{for } \mathcal{X} \text{ } w\text{-}X\text{-scenario}\}$. The claim then follows from the condition $f(w, \{x\}) \subseteq \{x\}$. 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. ■

We can thus derive that, for this particular f : $W^w = \bigcup_{Y \subseteq W} f(w, Y) = \bigcup_{\{x\} \subseteq W} f(w, \{x\}) = W$. In other words, all the worlds in the model are relevant for every $w \in W$, for every $w \in W$.

LEMMA 41. *Every evidence model is a grounded conditional model and every evidence CB-bisimulation is diffuse.*

PROOF. Thanks to the previous Lemma we can appeal to Lemma 11 and conclude that evidence models are grounded. For the second part of the claim, let $Z \subseteq W_1 \times W_2$ be a evidence CB-bisimulation and $x \in W_1$. Since the bisimulation is non-empty, there are $(w, w') \in Z$ and furthermore $x \in W_1 = W_1^w$, hence Z is diffuse. ■

PROPOSITION 42. *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.

PROPOSITION 43. *On evidence models, conditional belief is not definable in terms of the standard 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 an evidence 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 an evidence B-bisimulation. Given that α is a formula in the language of 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 $1\text{-}\llbracket p \rrbracket_{\mathcal{M}_1}$ -scenario $\mathcal{X} = \{\{2, 3\}, W_1\}$ and $2 \in \bigcap \mathcal{X}^{\llbracket p \rrbracket_{\mathcal{M}_1}}$ such that $2 \not\models q$. Hence we obtain a contradiction. ■

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 $1\text{-}\llbracket p \rrbracket_{\mathcal{M}_1}$ -scenario $\mathcal{X}$, and an element in $\bigcap \mathcal{X}^{\llbracket p \rrbracket_{\mathcal{M}_1}}$, 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 [8].

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 [8] 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 44. *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. ■

5. Relativized Common Knowledge

We now introduce a third example, the conditional modality known as relativized common knowledge, defined in [9]. Let $\mathcal{M} = \langle W, \{R_a\}_{a \in A}, V \rangle$ be a multi-agent Kripke model, where W is a non-empty set of worlds, relations $R_a \subseteq W \times W$ and $V : W \rightarrow \wp(At)$ a valuation function. Put $R := \bigcup_{a \in A} R_a$ and denote by R^+ its transitive closure. The operator of relativized common knowledge, denoted with $C(\phi, \psi)$, is meant to capture the intuition that every R -path which consists exclusively of ϕ -worlds ends in a world satisfying ψ . Formally:

$$\mathcal{M}, w \models C(\phi, \psi) \text{ iff } \mathcal{M}, v \models \psi \text{ for all } (w, v) \in (R \cap (W \times \llbracket \phi \rrbracket))^+$$

PROPOSITION 45. *Every Kripke model $\mathcal{M} = \langle W, \{R_a\}_{a \in A}, V \rangle$ can be converted into a conditional model, by taking $f(w, X) := \{v \mid (w, v) \in (R \cap (W \times X))^+\}$. Moreover, our semantics for conditionals for this f coincides with the above semantics for $C(\phi, \psi)$.*

PROOF. Again we check the prerequisites of selection functions in Definition 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 $X \subseteq Y$, $f(w, Y) = \{v \mid (w, v) \in (R \cap (W \times Y))^+\} \subseteq X$ 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: by reflexivity $w \in f(w, Y) \subseteq X$ so we also have a chain of length 0 contained in X , i.e. $w \in f(w, X)$. Suppose now $x \in f(w, X)$ 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 $f(w, X) \subseteq f(w, Y)$. ■

Replacing the new f in Definition 5, we obtain a new notion of bisimulation for conditional belief on plausibility models.

DEFINITION 46. 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

- $V(w) = V(w')$,
- for all $X \subseteq W_1$ and $X' \subseteq W_2$ such that $Z[X] \subseteq X'$ and $Z^{-1}[X'] \subseteq X$ we have that for every x such that $(w, x) \in (R_1 \cap (W_1 \times X))^+$ there exists a y such that $(w', y) \in (R_2 \cap (W_2 \times X'))^+$ such that $(x, y) \in Z$, and vice versa.

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

THEOREM 47. *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 they are RCK-bisimilar.*

The closure under unions also follows. As for composition, note that the notion of relevant worlds for w , indicated with W^w , starts to play a significant part, limiting the scope of our general results. Putting together the definition $W^w = \bigcup_{Y \subseteq W} f(w, Y)$ and $f(w, X) := \{v \mid (w, v) \in (R \cap (W \times X))^+\}$, W^w becomes the set of all the worlds reachable from w via an R -path (just substitute W for X in the definition of $f(w, X)$). Formally, $W^w = \{v \mid (w, v) \in R^+\}$. We can then characterize the grounded Kripke models.

PROPOSITION 48. *A Kripke model $\mathcal{M}$ is grounded iff, for every $w, x \in W$, if $(w, x) \in R^+$ then there is an agent a such that $(w, x) \in R_a$.*

PROOF. Let $\mathcal{M}$ be grounded. By Lemma 11 if $x \in W^w = \{v \mid (w, v) \in R^+\}$ then $f(w, \{x\}) = \{v \mid (w, v) \in (R \cap (W \times \{x\}))^+\} \neq \emptyset$. This entails that there is an edge $(w, x) \in R$, thus there must be an agent a such that $(w, x) \in R_a$.

For the other direction, let $X \subseteq W$ and $w \in W$ and suppose $X \cap W^w \neq \emptyset$. Then there is $x \in X$ such that $(w, x) \in R^+$. By our assumption on the model we know there is an agent a such that $(w, x) \in R_a$. This is enough to conclude $(w, x) \in R$ and thus $x \in f(w, X) = \{v \mid (w, v) \in (R \cap (W \times X))^+\}$, therefore $f(w, X) \neq \emptyset$. ■

In this context, a bisimulation $Z \subseteq W_1 \times W_2$ is diffuse if, for every $x \in W_1$, there are $w \in W_1$ and $w' \in W_2$ such that $(w, w') \in Z$ and x can be reached from w via an R -path (and vice versa).

PROPOSITION 49. *On grounded Kripke models, diffuse bisimulations are closed under relational composition.*

6. Generalization to Multi-agent Models

We have seen how our framework covers different conditional modalities, even when the same operator is interpreted on different semantics. Now we address the question: 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

$$\phi ::= p \mid \neg\phi \mid \psi \wedge \phi \mid \psi \rightsquigarrow_a \phi$$

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

DEFINITION 50. With the name *multi-agent conditional model* we indicate a tuple $\mathcal{M} = \langle W, A, \{f_a\}_{a \in A}, V \rangle$ with W a non-empty set of worlds, A a set of agents, $V : W \rightarrow \wp(At)$ a valuation function and for each agent $a \in A$ a selection function f_a satisfying the conditions listed in Definition 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. Instead of different agents, different labels could indicate different operators expressing distinct features of the models, depending on the interpretation. The semantics clause for the conditional modalities becomes:

$$\mathcal{M}, w \models \psi \rightsquigarrow_a \phi \quad \text{iff} \quad f_a(w, \llbracket \psi \rrbracket_{\mathcal{M}}) \subseteq \llbracket \phi \rrbracket_{\mathcal{M}}$$

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

DEFINITION 51. (*Multi-agent Conditional Bisimulation*) Given two multi-agent conditional models $\mathcal{M}_1$ and $\mathcal{M}_2$ based on 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

- $V(w) = V(w')$,
- for all $X \subseteq W_1$ and $X' \subseteq W_2$ such that $Z[X] \subseteq X'$ and $Z^{-1}[X'] \subseteq X$ we have that, for every $a \in A$, for every $x \in f_a^1(w, X)$ there exists a $y \in f_a^2(w', X')$ (where f^2 's are the selection functions in $\mathcal{M}_2$) such that $(x, y) \in Z$, and vice versa.

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

THEOREM 52. *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 multi-agent conditional 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 53. *Multi-agent conditional bisimulations are closed under arbitrary unions.*

The definitions of grounded models and diffuse bisimulation have to be generalized accordingly.

DEFINITION 54. Define $W_a^w = \bigcup_{Y \subseteq W} f_a(w, Y)$. A multi-agent conditional model is *grounded* if, for any $X \subseteq W$ and $a \in A$, $X \cap W_a^w \neq \emptyset$ entails $f_a(w, X) \neq \emptyset$.

DEFINITION 55. A multi-agent conditional bisimulation $Z \subseteq W_1 \times W_2$ is *diffuse* if for every $x \in W_1$ there are $a \in A$, $w \in W_1$ and $w' \in W_2$ such that $(w, w') \in Z$ and $x \in W_{1,a}^w$, and vice versa.

PROPOSITION 56. *Restricted to any class of multi-agent grounded models, the notion of multi-agent diffuse conditional bisimulation is closed under relational composition.*

6.1. Multi-agent Plausibility Models

We now turn to our fourth and last example, meant to display how the general definitions unfold in the multi-agent case. Our structure of choice is multi-agent plausibility models, a popular device used to model the knowledge and beliefs of different agents [4].

DEFINITION 57. A *multi-agent plausibility model* is a tuple $\mathcal{M} = \langle W, A, \{\leq_{a,w}, \sim_a\}_{a \in A, w \in W}, V \rangle$ with W a non-empty set of worlds, $\{\leq_{a,w}\}_{a \in A, w \in W}$ a family of reflexive and transitive relations $\leq_{a,w} \subseteq W \times W$ indexed by agents and worlds, $\{\sim_a\}_{a \in A}$ a family of “epistemic” equivalence relations $\sim_a \subseteq W \times W$ satisfying $\leq_a \subseteq \sim_a$, and $V : W \rightarrow \wp(At)$ a valuation function. A multi-agent plausibility model is *well-founded* if each relation $\leq_{a,w}$ is well-founded, i.e. $\forall X \subseteq W (X \neq \emptyset \Rightarrow \text{Min}_{\leq_{a,w}} X \neq \emptyset)$, where $\text{Min}_{\leq_{a,w}} X$ is the set of $\leq_{a,w}$ -minimal elements of X . For every $w \in W$ we write $[w]_{\sim_a} := \{v \in W \mid w \sim_a v\}$ for the associated equivalence class.

The semantics of the multi-agent belief and conditional belief operators on (well-founded) multi-agent plausibility models is given by:

- $\mathcal{M}, w \models B_a \phi$ iff for all $v \in \text{Min}_{\leq_{a,w}}([w]_{\sim_a})$ we have $\mathcal{M}, v \models \phi$
- $\mathcal{M}, w \models B_a^\psi \phi$ iff for all $v \in \text{Min}_{\leq_{a,w}}(\llbracket \psi \rrbracket_{\mathcal{M}} \cap [w]_{\sim_a})$ we have $\mathcal{M}, v \models \phi$

PROPOSITION 58. *Well-founded multi-agent plausibility models are multi-agent conditional models, where $f_a(w, X) = \text{Min}_{\leq_{a,w}}(X \cap [w]_{\sim_a})$.*

PROOF. We want to ascertain that the newly defined f_a fulfills the prerequisites of selection functions in Definition 1. The first condition is again given by definition. For the second one, suppose $X \subseteq Y$, $\text{Min}_{\leq_{a,w}}(Y \cap [w]_{\sim_a}) \subseteq X$ and consider a generic element x' in $\text{Min}_{\leq_{a,w}}(Y \cap [w]_{\sim_a})$. 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 $x' \in \text{Min}_{\leq_{a,w}}(X \cap [w]_{\sim_a})$. For the other inclusion take $x' \in \text{Min}_{\leq_{a,w}}(X \cap [w]_{\sim_a})$: 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 $z <_{a,w} x'$. Since $\leq_{a,w}$ is well-founded, there must be a minimal element $z' \in \text{Min}_{\leq_{a,w}}(Y \cap [w]_{\sim_a})$ such that $z' \leq_{a,w} z$; but by assumption $\text{Min}_{\leq_{a,w}}(Y \cap [w]_{\sim_a}) \subseteq X$, hence $z' \in X$. This gives us a $z' \in X \cap [w]_{\sim_a}$ such that $z' <_{a,w} x'$, contradicting the fact that x' is minimal in $X \cap [w]_{\sim_a}$. ■

Now that this step is secured, we can apply the general definition of multi-agent CB-bisimulation:

DEFINITION 59. 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

- $V(w) = V(w')$,
- for all $X \subseteq W_1$ and $X' \subseteq W_2$ such that $Z[X] \subseteq X'$ and $Z^{-1}[X'] \subseteq X$ we have, for all a , that for every $x \in \text{Min}_{\leq_{a,w}}(X \cap [w]_{\sim_a})$ then there exists a $y \in \text{Min}_{\leq'_{a,w'}}(X' \cap [w']_{\sim'_a})$ such that $(x, y) \in Z$, and vice versa (where $\leq'_{a,w'}$, $\sim'_a$ are the relations associated to a in $\mathcal{M}_2$).

Therefore all our results on the correspondence between bisimilarity and modal equivalence and closure under union do carry over to this setting. For closure under composition, notice that now $W_a^w = \bigcup_{Y \subseteq W} \text{Min}_{\leq_{a,w}}(Y \cap [w]_{\sim_a}) = [w]_{\sim_a}$ (just replace all the singletons for Y).

PROPOSITION 60. *Well-founded multi-agent plausibility models are grounded.*

PROOF. Let $X \cap W_a^w \neq \emptyset$. Then $X \cap [w]_{\sim_a} \neq \emptyset$. Since each relation $\leq_{a,w}$ is well founded, there will be minimal elements in $X \cap [w]_{\sim_a}$, thus $f_a(w, X) = \text{Min}_{\leq_{a,w}}(X \cap [w]_{\sim_a}) \neq \emptyset$. ■

In this setting a multi-agent plausibility CB-bisimulation $Z \subseteq W_1 \times W_2$ is diffuse if, for every $x \in W_1$, there are $w \in W_1$ and $w' \in W_2$ such that $(w, w') \in Z$ and x is in the information cell $[w]_{\sim_a}$ (and vice versa).

PROPOSITION 61. *On well-founded multi-agent plausibility models, diffuse multi-agent plausibility CB-bisimulation are closed under relational composition.*

7. Related Work

In light of the examples treated in the previous sections, one may wonder whether there are conditional modalities that do *not* fall under the scope of our framework. One example is 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 that goes back at least to [24].

$\mathcal{M}, w \models \phi \Rightarrow \psi$ iff for all v, v' we have that $\mathcal{M}, v \models \phi$ and $(w, v, v') \in R$ entail $\mathcal{M}, v' \models \psi$.

A possible selection function for this conditional modality could be

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

It is not hard to see, however, that such a selection function fails to satisfy the first requirement on conditional models: nothing in the definition of $f(w, X)$ ensures 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.

A different notion of bisimulation for conditional belief on multi-agent plausibility models was recently introduced in [3]. 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. But that analysis is confined to doxastic logic. Our approach has the following two distinctive features. First, the bisimulation for conditional belief stems from a general analysis of conditional modalities and it is not tailored to a specific application. This generality has the pleasant consequence that the key notions and proofs are 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 [3] depend crucially on the existence of the knowledge operator.¹⁰

A notion of bisimulation containing a quantification over subsets has been proposed originally in [20], adapted in [16] to epistemic lottery models and later again reshaped to work in the context of epistemic neighborhood models in [15]. 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.

Finally, we touch on the connection with the standard Hennessy–Milner result. Such result holds for an *un*-conditional modality, namely the box operator on Kripke models. For un-conditional modalities the proof of ‘modal equivalence entails bisimilarity’ simplifies considerably: it carries through

¹⁰Conversely, the undefinability result of Proposition 26 does not hold if we take knowledge into account, that is, we restrict the scope of belief to the current information cell.

with the usual technique just by assuming the finiteness of $f(w, X)$ for all w . When $f(w, X) = \{v | wRv\}$, where R is the relation of the Kripke model, we obtain a conditional model for the box operator; in this circumstance the finiteness of $f(w, X)$ for all w is precisely ‘finitely branching’.

8. Conclusion and Further Work

In this paper we proposed a general notion of bisimulation for conditional modalities interpreted on selection functions and proved general results including a Hennessy–Milner theorem. We applied this to a series of examples: plausibility models for conditional beliefs, evidence models, relativized common knowledge and multi-agent conditionals. We used these notions to obtain some new undefinability results.

The first open problem concerns the extension of our results on the closure under relational composition. Our results could be strengthened at the general level of conditional models or in the specific settings, where the selection functions may enjoy additional properties (e.g., the selection function for relativized common knowledge is fully monotonic).

The second open question concerns infinite models: does modal equivalence entail bisimilarity on some natural class of 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 obtain in the setting of conditional modalities. One natural example would be a version of the van Benthem characterization theorem.

Another group of questions arises from considering the new notion of bisimulation from a category-theoretic point of view. From this perspective bisimulations can be regarded as arrows in a suitable category of models. The closure under relational composition, together with the obvious fact that the identity relation is itself a bisimulation, ensures that we indeed obtain a category. This could enable a comparison between categories of models, for example between evidence and plausibility models, allowing for a systematic study of what has been called *tracking* [7], namely the matching of corresponding information dynamics in different classes of models.

Acknowledgements. We would like to thank Professor Johan van Benthem for his suggestions and encouragement. We also extend our gratitude to the

anonymous reviewers, their comments were extremely helpful in improving the paper.

Open Access. This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- [1] ANDERSEN, M. B., T. BOLANDER, and M. H. JENSEN, Don't plan for the unexpected: Planning based on plausibility models, *Logique et Analyse* 1:1, 2014.
- [2] ANDERSEN, M. B., T. BOLANDER, H. VAN DITMARSCH, and M. H. JENSEN, Bisimulation for single-agent plausibility models, in S. Cranefield, and A. Nayak (eds.), *AI 2013: Advances in Artificial Intelligence*, vol. 8272 of *Lecture Notes in Computer Science*, Springer International Publishing, New York, 2013, pp. 277–288.
- [3] ANDERSEN, M. B., T. BOLANDER, H. VAN DITMARSCH, and M. H. JENSEN, Bisimulation and expressivity for conditional belief, degrees of belief, and safe belief, *Synthese* 1–41, 2016.
- [4] BALTAG, A., and S. SMETS, Conditional doxastic models: A qualitative approach to dynamic belief revision, *Electronic Notes in Theoretical Computer Science* 165:5–21, 2006.
- [5] BALTAG, A., and S. SMETS, A qualitative theory of dynamic interactive belief revision, in M. Woolridge, G. Bonanno, and W. van der Hoek, (eds.), *Texts in Logic and Games*, vol. 3, Amsterdam University Press, Amsterdam, 2008, pp. 9–58.
- [6] VAN BENTHEM, J., Dynamic logic for belief revision, *Journal of Applied Non-Classical Logics* 17(2):129–155, 2007.
- [7] VAN BENTHEM, J., Tracking information, in K. Bimbo, (ed.), *J. Michael Dunn on Information Based Logics*, Springer, Dordrecht, 2016, pp. 363–389.
- [8] VAN BENTHEM, J., and E. PACUIT, Dynamic logics of evidence-based beliefs, *Studia Logica* 99(1–3):61–92, 2011.
- [9] VAN BENTHEM, J., J. VAN ELJCK, and B. KOOI, Logics of communication and change, *Information and Computation* 204(11):1620–1662, 2006.
- [10] BLACKBURN, P., M. DE RIJKE, and Y. VENEMA, *Modal Logic*, vol. 53, Cambridge University Press, Cambridge, 2002.
- [11] BOARD, O., Dynamic interactive epistemology, *Games and Economic Behavior* 49(1):49–80, 2004.
- [12] CHELLAS, B. F., Basic conditional logic, *Journal of Philosophical Logic* 4(2):133–153, 1975.
- [13] DEMEY, L., Some remarks on the model theory of epistemic plausibility models, *Journal of Applied Non-Classical Logics* 21(3–4):375–395, 2011.

- [14] DEMEY, L., 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*, 2012, pp. 73–78 .
- [15] VAN EIJCK, J., *Update, Probability, Knowledge and Belief*, 2015. Draft available from <http://homepages.cwi.nl/~jve/papers/15/pdfs/upkb.pdf>.
- [16] VAN EIJCK, J., and F. SCHWARZENTRUBER, Epistemic probability logic simplified, *Advances in Modal Logic* 10:158–177, 2014.
- [17] GABBAY, D. M., *Theoretical Foundations for Non-monotonic Reasoning in Expert Systems*, Springer, Berlin, 1985.
- [18] GROVE, A., Two modellings for theory change, *Journal of Philosophical Logic* 17(2):157–170, 1988.
- [19] HALPERN, J. Y., *Reasoning about uncertainty*, vol. 21, MIT Press Cambridge, 2003.
- [20] KOOI, B. P., *Knowledge, Chance, and Change*, PhD thesis, Groningen University, 2003.
- [21] KRAUS, S., D. LEHMANN, and M. MAGIDOR, Nonmonotonic reasoning, preferential models and cumulative logics, *Artificial Intelligence* 44(1):167–207, 1990.
- [22] LEWIS, D., *Counterfactuals*, Harvard University Press, Boston, 1973.
- [23] STALNAKER, R. C., A theory of conditionals, in N. Rescher, (ed.), *Studies in logical theory*, Blackwell, Oxford, 1968.
- [24] URQUHART, A., Semantics for relevant logics, *The Journal of Symbolic Logic* 37(01):159–169, 1972.
- [25] VELTMAN, F., *Logic for Conditionals*, PhD thesis, University of Amsterdam, 1985.

A. BALTAG, G. CINÀ
 ILLC
 University of Amsterdam
 Science Park 107
 Amsterdam 1098 XH
 The Netherlands
giovanni.cina88@gmail.com

A. BALTAG
thealexandrubaltag@gmail.com