Modal Interpolation via Nested Sequents

Melvin Fitting[a,1], Roman Kuznets[b,c,2,3,*]

[a]*Department of Computer Science (Emeritus), The Graduate School and University Center (CUNY)*
[b]*Institut für Informatik und angewandte Mathematik, Universität Bern*
[c]*Institut für Computersprachen, Technische Universität Wien*

Abstract

The main method of proving the Craig Interpolation Property (CIP) constructively uses cut-free sequent proof systems. Until now, however, no such method has been known for proving the CIP using more general sequent-like proof formalisms, such as hypersequents, nested sequents, and labelled sequents. In this paper, we start closing this gap by presenting an algorithm for proving the CIP for modal logics by induction on a nested-sequent derivation. This algorithm is applied to all the logics of the so-called modal cube.

Keywords: Craig interpolation, nested sequent, structural proof theory, modal logic
2010 MSC: 03B45, 03C40, 03F07

1. Introduction

Suppose, for the moment, that we identify a logic with a set of formulas L, ignoring semantics, proof systems, and so on. Then L has the *Craig interpolation property* (CIP) if, whenever $A \supset B \in \mathsf{L}$, there exists a formula C such that $A \supset C, C \supset B \in \mathsf{L}$, where C is in the "common language" of A and B. What "common language" means is situation-dependent: it can mean having shared propositional variables, or individual variables, or modal operators, or nominals, etc. Whether a logic has the CIP is an important characteristic of the logic (for an overview of the problems and complexity of interpolation, as well as a history of the subject, see [6]).

In addition to knowing whether a logic has the CIP, it is useful to be able to prove the property *constructively*.[4] For instance, in [12], the existence of fixed points in the logic of provability GL was established constructively using the CIP of GL, and in [3], that method was made the basis of a program to compute fixed points for GL. Historically, constructive proofs of the CIP for L make use of a *cut-free* proof system for L, typically a sequent calculus or tableau system. Unfortunately, sequent calculi for modal logics whose semantics involves symmetry are hard to come by, though sometimes *ad hoc* methods have been devised. Actually, cut-freeness is not the central issue. What is important is that we have a proof procedure that has the subformula property and in which polarity of subformulas is preserved during the course of a proof. (Both are violated by the cut rule.) Over the years generalizations of sequent and tableau systems have been developed, capable of handling a greater number of logics in a uniform way: among them, *nested sequents* ([1])

---

and *prefixed tableaus* ([2, 11]). Nested sequents are in the tradition allowing so-called *deep reasoning* in the course of a proof. Prefixed tableaus allow a kind of representation of possible worlds directly in the proof, with accessibility captured using simple syntactic machinery. It turns out that nested sequents and prefixed tableaus are related much as ordinary sequents and tableaus are—loosely, one is the other upside-down ([4]).

In this paper, we show that nested sequents can be used in a natural way to prove the CIP constructively for all the logics in the *modal cube*. There is no appeal to *ad hoc* methods; the work is essentially uniform across the whole family. While our primary interest is in interpolants between two formulas or, more generally, two sets of formulas, we develop something broader: we introduce the notion of interpolant between sets of formulas within a nested structure. Interpolants created at this level cannot be just formulas. Instead, our interpolants are Boolean combinations of formulas within the same nested structure. Thus, we actually work with a more general notion of interpolant. Interpolation in the usual sense is a special case.

We assume we have a countable set Prop of propositional variables, fixed throughout this paper. Our modal language $\mathcal{L}$ is built up from these propositional variables, together with $\top$ and $\bot$, in the usual way, using propositional connectives $\neg$, $\wedge$, $\vee$, and $\supset$ and modal operators $\square$ and $\Diamond$. We omit parentheses when it will not lead to confusion. For a modal formula A, by $\mathrm{Prop}(A)$ we mean the set of propositional variables that occur in A. We note that all logics considered here are monomodal, though proof systems of the sort we use do exist for multimodal logics as well.

Definition 1.1 (CIP for modal logics). A modal logic L has the *CIP* if for any formulas A and B such that $A \supset B \in \mathsf{L}$, there exists an *interpolant* C such that

$$A \supset C \in \mathsf{L}, \qquad C \supset B \in \mathsf{L}, \qquad \text{and} \qquad \mathrm{Prop}(C) \subseteq \mathrm{Prop}(A) \cap \mathrm{Prop}(B) \ .$$

Thus, *common language* for modal logics simply means having common propositional variables.

Unlike with ordinary sequent calculi, there are cut-free nested sequent systems for each normal modal logic formed from any combination of axioms d, t, b, 4, and 5. The difficulty in extracting interpolants from nested sequent proofs lies in the presence in nested sequents of an additional structural connective that corresponds to $\square$ the same way that comma in ordinary sequents corresponds to $\wedge$ in the antecedent or to $\vee$ in the consequent of a two-sided sequent. This additional nested structure has to be reflected in the interpolation process, and this is the source of most of the technical complexity to be found here.

2. Nested Sequent Calculus for the Basic Normal Monomodal Logic K

Definition 2.1 (Logic K). The *minimal normal monomodal logic* K is the logic of all Kripke frames. It is axiomatized by

- an arbitrary complete set of axioms of classical propositional logic (in the monomodal language $\mathcal{L}$),

- the rule modus ponens,

- the normality axiom k: $\square(A \supset B) \supset (\square A \supset \square B)$,

- the necessitation rule: $\dfrac{A}{\square A}$.

We identify the logic K with the set of its theorems and write $\mathsf{K} \vdash A$ instead of $A \in \mathsf{K}$.

Before presenting a nested sequent calculus for K, we need to define *nested sequents*, the objects of the derivation in such a calculus, and *contexts*, the tools necessary to describe rules of such a calculus. We define a grammar for nested objects that covers both sequents and contexts and explain how to distinguish between them.

Definition 2.2 (Nested objects). We define *nested objects* Φ according to the following grammar:

$$\Phi ::= \varepsilon \mid \Phi, A \mid \Phi, \{\,\} \mid \Phi, [\Phi] \qquad ,$$

where ε is the empty sequence, $A \in \mathcal{L}$ is a formula, $\{\,\}$ is the *hole* symbol, brackets in $[\Phi]$ are called a *structural box*, and comma is the operation of appending an element to the end of a sequence. Thus, a nested object is a sequence of formulas, of occurrences of the hole symbol, and of nested objects within structural boxes.

It is trivial to define the number of occurrences of the hole symbol in a given nested object, which is why we omit the formal definition.

Definition 2.3 (Nested sequents, contexts, and multicontexts). A nested object without holes is called a *nested sequent*, or, simply, a *sequent*. A nested object with exactly one hole is called a *context*. A nested object with more than one hole is called a *multicontext*. In this paper, we do not use multicontexts, although they do play a significant role elsewhere. Thus, from now on, all nested objects are assumed to be either sequents or contexts. Since there is exactly one hole in a context, we call it *the* hole. We use Γ, Δ, $\ldots$, possibly with sub- and/or superscripts, to denote sequents, $\Gamma\{\,\}$, $\Delta\{\,\}$, $\ldots$, possibly with sub- and/or superscripts, to denote contexts, and Φ, Ψ, $\ldots$, possibly with sub- and/or superscripts, to denote nested objects that can be either sequents or contexts.

One way of looking at a nested sequent is to consider a tree of ordinary one-sided sequents, i.e., of sequences of formulas, which we call here *shallow sequents* to avoid ambiguity. Each structural box in the nested sequent corresponds to a child in the tree. Nested sequent calculi are designed to use the mechanism of deep inference, where rules are applied at an arbitrary node of this tree, i.e., arbitrarily deep in the nested structure of the sequent. The hole in a context provides a reference to the place, or to the node in the tree, at which the rule should be applied.

Definition 2.4 (Inserting a sequent into a nested object). The insertion of a sequent Δ into a nested object Φ is obtained by performing the following action on Φ: if Φ contains the hole, replace it with Δ; otherwise, do not do anything. The result of such an insertion is denoted $\Phi\{\Delta\}$.

When we use the notation $\Phi\{\Delta\{\Gamma\}\}$, it should be read as follows: first, the sequent Γ is inserted into the context $\Delta\{\,\}$; second, the resulting sequent $\Delta\{\Gamma\}$ is inserted into the nested object Φ.

Figure 1 recalls the *uniform notation*, which is typical of tableau calculi and which we use in the paper.

α	α_1	α_2
$A \wedge B$	A	B
$\neg(A \vee B)$	$\neg A$	$\neg B$
$\neg(A \supset B)$	A	$\neg B$

β	β_1	β_2
$A \vee B$	A	B
$\neg(A \wedge B)$	$\neg A$	$\neg B$
$A \supset B$	$\neg A$	B

ν	ν_0
$\Box A$	A
$\neg \Diamond A$	$\neg A$

π	π_0
$\Diamond A$	A
$\neg \Box A$	$\neg A$

Figure 1: Uniform notation.

Definition 2.5 (Nested sequent calculus NK). The nested sequent calculus NK for the modal logic K can be found in Figure 2. This calculus is a hybrid of the multiset-based version from [1], where formulas are restricted to the negation normal form, of the sequence-based version from [8], where formulas are also restricted to the negation normal form, and of the set-based version from [4], where formulas are unrestricted. None of the three versions uses Boolean constants $\bot$ and $\top$, necessitating an addition of the rules id$_\top$ and id$_{\neg\bot}$ for handling these. According to the uniform notation, the α and β rules in Figure 2 encode three rules each and the ν and π rules encode two rules each. In [1, 8], the ν and π rules are called the $\Box$ and k rules respectively.

$$\text{id}_\mathsf{P}\ \overline{\Gamma\{P, \neg P\}} \qquad \text{id}_\top\ \overline{\Gamma\{\top\}} \qquad \text{id}_{\neg\perp}\ \overline{\Gamma\{\neg\perp\}}$$

$$\neg\neg\ \frac{\Gamma\{A\}}{\Gamma\{\neg\,\neg\,A\}} \qquad \beta\ \frac{\Gamma\{\beta_1, \beta_2\}}{\Gamma\{\beta\}} \qquad \alpha\ \frac{\Gamma\{\alpha_1\}\quad \Gamma\{\alpha_2\}}{\Gamma\{\alpha\}}$$

$$\text{ctr}\ \frac{\Gamma\{A, A\}}{\Gamma\{A\}} \qquad \text{exch}\ \frac{\Gamma\{\Delta, \Sigma\}}{\Gamma\{\Sigma, \Delta\}} \qquad \nu\ \frac{\Gamma\{[\nu_0]\}}{\Gamma\{\nu\}} \qquad \pi\ \frac{\Gamma\{[\Delta, \pi_0]\}}{\Gamma\{[\Delta], \pi\}}$$

Figure 2: Rules of the nested sequent calculus NK.

Example 2.6. A nested sequent derivation of $\Box(P \supset \Box Q) \supset \big(\Box(\Box Q \supset R) \supset \Box(P \supset R)\big)$ can be found in Figure 3.

Figure 3: An NK derivation of $\Box(P \supset \Box Q) \supset \big(\Box(\Box Q \supset R) \supset \Box(P \supset R)\big)$.

To understand in which sense this nested calculus represents the modal logic, we need a translation from sequents to formulas of this modal logic. Intuitively, our sequents are one-sided (i.e., comprising only the consequent part of a two-sided sequent) with two structural connectives: structural disjunction (comma ,) and structural necessitation modality (bracket $[\cdot]$). This intuition is formalized using the notion of a *corresponding formula*. Since this notion does not make sense for contexts, we use the following subgrammar that defines sequents:

$$\Gamma ::= \varepsilon \mid \Gamma, A \mid \Gamma, [\Gamma]\ .$$

This subgrammar will also be used in proofs by induction on the structure of a given sequent.

Definition 2.7 (Corresponding formula of a nested sequent). For any sequent Γ, its *corresponding*

4

formula $\underline{\Gamma}$ is defined by induction on the construction of the sequent:

$$\underline{\varepsilon} := \bot; \qquad \underline{\Gamma, A} := \begin{cases} \underline{\Gamma} \vee A & \text{if } \Gamma \neq \varepsilon, \\ A & \text{otherwise}; \end{cases} \qquad \underline{\Gamma, [\Delta]} := \begin{cases} \underline{\Gamma} \vee \Box \underline{\Delta} & \text{if } \Gamma \neq \varepsilon, \\ \Box \underline{\Delta} & \text{otherwise}. \end{cases}$$

The following theorem follows from the results in [1, 4, 8].

Theorem 2.8 (Completeness of the calculus NK). *For any sequent Γ, we have $\mathsf{NK} \vdash \Gamma$ iff $\mathsf{K} \vdash \underline{\Gamma}$ iff $\Vdash \underline{\Gamma}$, where $\Vdash$ expresses validity in the class of all Kripke models.*

3. Formulating Interpolation for Nested Sequents

Our goal is to prove the CIP by induction on the nested sequent derivation. To achieve this, it is necessary to decide what being an interpolant of a nested sequent means and what kind of object an interpolant of a nested sequent is. Since the answers to these questions have turned out to be quite non-trivial, we first explain how we arrived at these answers by using an analogy with the formalism closest to nested sequents: namely, shallow one-sided sequents, i.e., multisets of formulas. We show why for nested sequents it is not possible to directly use the definitions of interpolation and interpolant known from shallow sequents and explain in which sense our definitions conform to the idea of interpolation and how they are formally related to the CIP, whose formulation is independent of the proof system used.

Definition 3.1 (Interpolant of a shallow one-sided sequent). Let a multiset of formulas Δ be partitioned into two multisets $L(\Delta)$ and $R(\Delta)$. Whenever safe, we omit the parentheses. We call $L\Delta \mid R\Delta$ a *split* of Δ. An interpolant of the split $L\Delta \mid R\Delta$ of the sequent Δ is a formula C such that

$$\neg C \Vdash L\Delta, \qquad C \Vdash R\Delta \ ,$$

and C only uses propositional variables common to $L\Delta$ and $R\Delta$.

For a given shallow one-sided sequent calculus, this statement is usually formulated syntactically in terms of derivability in the calculus; however, as we explain later, interpolants of nested sequents cannot be restricted to the formula format. Since an introduction of yet another calculus just for the purpose of formulating the interpolation statement syntactically does not seem reasonable, our interpolation statements for nested sequents are formulated semantically, and the statement above is the closest analog.

Example 3.2. Before giving formal definitions, let us consider a simple example to motivate our choices. A shallow sequent $Q, R, P, \neg P$ is derivable by id_P. One of the possible splits is $Q, R, \neg P \mid P$, and the interpolant prescribed to such a split by most sequent-based algorithms is P.

Let us now consider the sequent $\Delta = [Q, R], [P, \neg P]$, derivable by the same rule id_P. It should be possible to split this sequent in the same way: $[Q, R], [\neg P] \mid [P]$. However, the information that P and $\neg P$ are in the same structural box, which makes the sequent derivable, is lost in such a split. Thus, instead of creating the split by physically moving formulas in the sequent, we indicate to which side of the split each formula should belong by assigning it a *bias* in the form of a superscript. The above-mentioned split is then represented by the *biased sequent* $\widetilde{\Delta} = [Q^\ell, R^\ell], [P^\Re, \neg P^\ell]$. The left (right) side of such a biased sequent is obtained by dropping all right-biased (left-biased) formulas and erasing the biases of the remaining formulas:

$$L\widetilde{\Delta} = [Q, R], [\neg P] \ , \quad R\widetilde{\Delta} = [\varepsilon], [P] \ .$$

Note that the structural boxes are not removed even if they become empty.

Let us simplify the example further. Suppose we want to find an interpolant of $[P, \neg P]$ biased in the following way: $[P^\Re, \neg P^\ell]$. The first intuition is that $[P]$ should be an interpolant, by analogy with the interpolant P of the corresponding split $\neg P \mid P$ in the shallow case. Unfortunately, such an interpolant would not satisfy the left part of the interpolation definition: $\neg\left([P]\right) \nVdash [\neg P]$, i.e., $\neg \Box P \nVdash \Box \neg P$. In fact, no modal formula C exists such that $\neg C \Vdash [\neg P]$ and $C \Vdash [P]$. Indeed, the existence of such C would imply that $\mathsf{K} \vdash \Box \neg P \vee \Box P$, which is not the case.

5

This observation may seem to be the end of the quest for interpolation via nested sequents. But what it means in reality is that either negation of a nested interpolant should not be equated to negating its corresponding formula or that the interpolation statement should not be about corresponding formulas, at least not directly. We have implemented our method of proving the CIP in both of the approaches. We, however, only present the latter in the paper because the rules for interpolant construction are more intuitive if the interpolation statement is separated from corresponding formulas. By contrast, the negation operation for interpolants that is necessary for our method is not natural for nested sequents. This operation and our whole interpolation method are inspired by Fitting's prefixed sequent calculus from [4]. We now give the necessary formal definitions.

Definition 3.3 (Biased nested object). We define *biased nested objects* $\widetilde{\Phi}$ as nested objects within which each formula is assigned exactly one of the two *biases*: the superscripts $\cdot^{\ell}$ or $\cdot^{\Re}$. *Biased nested sequents*, *biased contexts*, *biased multicontexts*, and inserting a biased sequent into a biased nested object are defined the same way as in Definitions 2.3 and 2.4. When it is not important whether a nested object is biased, we use the same notation for biased ones. To emphasize that the nested object is biased, we put a tilde over it. Moreover, if we use the same letter with and without a tilde, e.g., both $\widetilde{\Phi}$ and Φ, it means that Φ can be obtained from $\widetilde{\Phi}$ by erasing all biases. In this case, Φ is the *unbiased version* of $\widetilde{\Phi}$ and $\widetilde{\Phi}$ is a *biased version* of Φ.

For a biased nested object $\widetilde{\Phi}$, its *left side* $L\widetilde{\Phi}$ (*right side* $R\widetilde{\Phi}$) is obtained by dropping all right-biased (left-biased) formulas and erasing the biases of the remaining formulas:

Example 3.4.

$$L\left(E^{\Re}, A^{\ell}, \left[B^{\ell}, \left[C^{\ell}\right], \left[F^{\Re}\right], \left[G^{\Re}, D^{\ell}, H^{\Re}\right]\right]\right) = A, [B, [C], [\varepsilon], [D]] \ ,$$
$$R\left(E^{\Re}, A^{\ell}, \left[B^{\ell}, \left[C^{\ell}\right], \left[F^{\Re}\right], \left[G^{\Re}, D^{\ell}, H^{\Re}\right]\right]\right) = E, [[\varepsilon], [F], [G, H]] \ .$$

Splitting an ordinary sequent plays the role of disjunction: the sequent is equivalent to the disjunction of the left and right sides of the split. As noted above, this property is not preserved if we interpret nested sequents via their corresponding formulas. We now define an alternative semantics for nested sequents that supports this disjunctive view of splits achieved by biasing formulas. We call this semantics *decorative* because an interpretation of a nested object viewed as a tree with sequences of formulas in its nodes is achieved by decorating each node with a world from a given Kripke model. Before giving a formal definition, which is quite technical, let us illustrate what we mean with an example:

Example 3.5. Let us call a node of the sequent tree of a nested object its *sequent node*. Here is a decoration:

$$u, E^{\Re}, A^{\ell}, \left[v, B^{\ell}, \left[w_1, C^{\ell}\right], \left[w_2, F^{\Re}\right], \{\ \}_v, \left[w_3, G^{\Re}, D^{\ell}, H^{\Re}\right]\right] \ . \tag{1}$$

It decorates the biased context

$$E^{\Re}, A^{\ell}, \left[B^{\ell}, \left[C^{\ell}\right], \left[F^{\Re}\right], \{\ \}, \left[G^{\Re}, D^{\ell}, H^{\Re}\right]\right] \ , \tag{2}$$

which can be obtained by simply erasing all decorating worlds from (1). The world u decorates the root of the sequent tree and is called the *root of the decoration*. The only child of the root sequent node is decorated by v; this is the node with the hole $\{\ \}$ in it. The three children of the sequent node with the hole are decorated by w_1, w_2, and w_3 in this order. In the linear notation of nested sequents, the decorating worlds are added at the beginning of each sequent node, and the hole is (redundantly) indexed with v, which decorates the sequent node with the hole. The final condition is that the tree of decorations should be embeddable into the Kripke model (W, R, V), which means uRv, vRw_1, vRw_2, and vRw_3 in case of (1).

Definition 3.6 (Decorated nested object). Let $\mathcal{M} = (W, R, V)$ be a Kripke model with the set of worlds $W \neq \varnothing$, the accessibility relation $R \subseteq W \times W$, and the propositional valuation $V \colon \text{Prop} \to 2^W$. We define $\mathcal{M}$-*decorated nested objects* Φ^* as (possibly biased) nested objects with each sequent node decorated by a world from W in such a way that the world decorating a child of a given sequent node is R-accessible from the world decorating the sequent node itself. We often omit the mention of the model.

The decoration of each sequent node delimited by brackets is added right after the opening bracket. The decoration of the root of the sequent tree is placed at the beginning of the whole nested object and is called the *root of the decoration*. For a decorated context, we write the world decorating the sequent node with the hole as a subscript of the hole. While not necessary, this is notationally convenient.

We use superscripts $^*, ^\star$, etc. to denote decorated objects. If we use the same letter with and without such a superscript, e.g., both Φ^* and Φ, it means that Φ can be obtained from Φ^* by removing all decorations. In this case, Φ^* is a *decoration* of Φ.

The root of a decoration Φ^* is denoted $r(\Phi^*)$. We denote Φ^* with the root removed $t(\Phi^*)$ and call it the *tail of* Φ^*. Thus, $\Phi^* = r(\Phi^*), t(\Phi^*)$.

Remark 3.7. Note that (1) is also a biased version of the decorated context

$$u, E, A, [v, B, [w_1, C], [w_2, F], \{\ \}_v, [w_3, G, D, H]]\ . \tag{3}$$

(Strictly speaking, we have not defined biased versions of decorated nested objects, but the definition is literally the same as for biased versions of nested objects.) Note further that (3) is a decoration and (2) is a biased version of the same context

$$E, A, [B, [C], [F], \{\ \}, [G, D, H]]\ .$$

This commutative diagram should come as no surprise. Biasing and decorating a nested object are independent of each other: the former affects only formulas; the latter affects only sequent nodes. Thus, given a decorated biased nested object, it does not matter whether we first erase biases and then remove decorations or first remove decorations and then erase biases. The result is going to be the same. In our notational scheme, $\widetilde{\Phi}^*$ is a decoration of $\widetilde{\Phi}$, which is a biased version of Φ. $\widetilde{\Phi}^*$ is also a biased version of Φ^*, which is a decoration of Φ. We do not have to deal with alternative biases of the same nested object, so the use of the tilde presents no problems. As we just discussed, the use of $\widetilde{\Phi}^*$ and Φ^* presupposes that the former is a biased version of the latter, i.e., that the decoration is the same in both cases. If we need to describe different decorations, we write $\widetilde{\Phi}^*$ and $\Phi^\star$.

The decoration of a nested sequent plays the role of a valuation in propositional classical logic or of an interpretation and valuation in first-order logic. Now we have to define what it means for a sequent to be true under a given decoration. While we could define decorated sequents via a formal grammar and give the definition of truth by induction on that grammar, we find the following definition more direct and intuitive without any loss of rigor.

Definition 3.8 (True sequent decoration). Let $\mathcal{M} = (W, R, V)$ be a Kripke model. An $\mathcal{M}$-decorated sequent Γ^* (possibly biased) is *true* if at least one formula from Γ^* holds at the world of $\mathcal{M}$ that decorates the sequent node of this formula. If the sequent is biased, formula biases are ignored. We write $\Vdash \Gamma^*$ to denote that Γ^* is true.

The crucial idea behind this definition is that decorations interpret nested sequents as disjunctions of member formulas, similar to the standard interpretation of one-sided shallow sequents. However, in the shallow case, all formulas are evaluated at the same world of the Kripke model, making it possible to write the intended disjunction as one formula to be evaluated at this world and, consequently, to express the interpolation statement on the formula level. In the nested case, by contrast, formulas from different sequent nodes are evaluated at different worlds of the Kripke model, as specified by a decoration. Since such a disjunction cannot be expressed within the object language, the interpolation statement for nested sequents cannot be stated on the formula level.

Example 3.9. Both the $\mathcal{M}$-decorated sequent $\Gamma^* = u, E, A, [v, [w_1, C], [w_2], B, [w_3, G, D]]$ and its biased version $\widetilde{\Gamma}^* = u, E^\Re, A^\ell, [v, [w_1, C^\ell], [w_2], B^\ell, [w_3, G^\Re, D^\ell]]$ are true iff

$$\mathcal{M}, u \Vdash E \quad \text{or} \quad \mathcal{M}, u \Vdash A \quad \text{or} \quad \mathcal{M}, w_1 \Vdash C \quad \text{or} \quad \mathcal{M}, v \Vdash B \quad \text{or} \quad \mathcal{M}, w_3 \Vdash G \quad \text{or} \quad \mathcal{M}, w_3 \Vdash D\ .$$

In addition, $L\widetilde{\Gamma}^* = u, A, [v, [w_1, C], [w_2], B, [w_3, D]]$ is true iff

$$\mathcal{M}, u \Vdash A \quad or \quad \mathcal{M}, w_1 \Vdash C \quad or \quad \mathcal{M}, v \Vdash B \quad or \quad \mathcal{M}, w_3 \Vdash D$$

and $R\widetilde{\Gamma}^* = u, E, [v, [w_1], [w_2], [w_3, G]]$ is true iff $\mathcal{M}, u \Vdash E \quad or \quad \mathcal{M}, w_3 \Vdash G$. Thus, $\Vdash \widetilde{\Gamma}^*$ iff $\Vdash L\widetilde{\Gamma}^*$ or $\Vdash R\widetilde{\Gamma}^*$.

Remark 3.10. Strictly speaking, $L\widetilde{\Gamma}^*$ and $R\widetilde{\Gamma}^*$ are ambiguous notations. For instance, $L\widetilde{\Gamma}^*$ can be read as $\left(L\widetilde{\Gamma}\right)^*$ instead of $L\left(\widetilde{\Gamma}^*\right)$ as intended. However, these two readings always produce the same result.

This example shows a general property that substantiates the earlier claim that the decorative semantics is disjunctive with respect to the left–right biasing of a sequent.[5]

Fact 3.11. *For a decorated biased sequent $\widetilde{\Gamma}^*$,*

$$\Vdash \widetilde{\Gamma}^* \qquad \Longleftrightarrow \qquad \Vdash L\widetilde{\Gamma}^* \quad or \quad \Vdash R\widetilde{\Gamma}^* \ .$$

We have defined how to split a nested sequent by biasing formulas and how to evaluate the whole sequent and its sides by decorating sequent nodes. It remains to define interpolants and extend the decorative semantics to them. After that, we will be able to formulate the interpolation theorem.

Let us start by returning to Example 3.2 and show that the predicted interpolant is indeed an interpolant with respect to decorative semantics. Until a formal definition of interpolant is given in Definition 3.23, the discussion must remain on an informal level. The goal of this example, therefore, is to explain intuitively how our notion of interpolant resolves the problems outlined in Example 3.2.

Example 3.12. To show that $[P]$ is an interpolant of $[P^{\Re}, \neg P^{\ell}]$, we need to show that $\neg [P] \Vdash [\neg P]$ and $[P] \Vdash [P]$. We interpret these statements as follows: given any model $\mathcal{M} = (W, R, V)$ and any $\mathcal{M}$-decoration of $[P^{\Re}, \neg P^{\ell}]$,

- if the interpolant is true *with respect to this decoration*, the right side of the decorated sequent is true;

- if the interpolant is false *with respect to this decoration*, the left side of the decorated sequent is true.

Given the decoration of a biased sequent, we know how to decorate the sides of the biased sequent. What we have not defined yet is how to transfer the decoration of the biased sequent to its interpolant. Any $\mathcal{M}$-decoration of $[P^{\Re}, \neg P^{\ell}]$ has the form $w, [u, P^{\Re}, \neg P^{\ell}]$, where wRu. The corresponding decorations of the left and right sides are $w, [u, \neg P]$ and $w, [u, P]$ respectively. Since the suggested interpolant $[P]$ coincides with the right side, its *corresponding* decoration should be the same as that of the right side, i.e., $w, [u, P]$. Now our interpolation statement takes the form

$$\nVdash w, [u, P] \quad \Rightarrow \quad \Vdash w, [u, \neg P] \qquad and \qquad \Vdash w, [u, P] \quad \Rightarrow \quad \Vdash w, [u, P] \ .$$

The second implication is trivial. Let us verify the first. For a decorated sequent to be false, all formulas in it must be false at the appropriate worlds:

$$\nVdash w, [u, P] \quad \Rightarrow \quad \mathcal{M}, u \nVdash P \quad \Rightarrow \quad \mathcal{M}, u \Vdash \neg P \quad \Rightarrow \quad \Vdash w, [u, \neg P] \ .$$

Now we need to answer the question, what the interpolant of $[Q^{\ell}, R^{\Re}], [P^{\Re}, \neg P^{\ell}]$ should be. Any $\mathcal{M}$-decoration of this sequent has the form $w, [v, Q^{\ell}, R^{\Re}], [u, P^{\Re}, \neg P^{\ell}]$ with wRv and wRu. If we try to use the same interpolant $[P]$, it is not immediately clear what its corresponding decoration should be. Or, given the discussion above, it is not clear how to justify decorating the child node of such an interpolant with u rather than v. Our solution is very simple. To avoid the ambiguity, we match the sequent-tree structure of the interpolant to the sequent-tree structure of the given biased sequent. For the interpolant

---

[5]We formulate this property as a fact rather than a lemma because its proof is sufficiently transparent.

$[\varepsilon], [P]$, common sense suggests using v, the decoration of the first structural box of the sequent, for the first structural box $[\varepsilon]$ of the interpolant and using u, the decoration of the second structural box of the sequent, for the second structural box $[P]$ of the interpolant: $w, [v], [u, P]$. The sequence format of nested sequents has been chosen specifically to make it possible to refer to structural boxes as the first, the second, the second within the first, etc. We leave it to the reader to verify that

$$\nVdash w, [v], [u, P] \quad \Rightarrow \quad \Vdash w, [v, Q], [u, \neg P] \qquad \text{and} \qquad \Vdash w, [v], [u, P] \quad \Rightarrow \quad \Vdash w, [v, R], [u, P] \ .$$

This example explains our intuition that the interpolant should share the nested structure of the sequent it interpolates. However, we know from propositional interpolation that two-premise sequent rules require taking conjunctions and disjunctions of the interpolants of the premises. When interpolants are formulas, such operations present no problems. Since our interpolants, due to their nested structure, must be nested sequents, we are forced to allow interpolants to be conjunctions and disjunctions of nested sequents. We still require all members of these conjunctions and disjunctions to have the same nested structure as the sequent being interpolated. We realize this restriction via the notion of skeletons.

Definition 3.13 (Skeleton). The *skeleton* Φ° of a (possibly biased and/or decorated) nested object Φ is obtained by removing all (biased) formulas.

Example 3.14. We have $\left(A, \left[B, [C, E], [D] \right] \right)^\circ = \left[[], [] \right]$ and $\left(w, A^\ell, \left[u, B^\Re, [v, C^\Re, \{ \}_v], [v', D^\ell] \right] \right)^\circ = w, \left[u, [v, \{ \}_v], [v'] \right]$, where $[]$ stands for $[\varepsilon]$ (the two notations will be used interchangeably).

Definition 3.15 (Generalized nested sequent). We define *structure-preserving Boolean combinations of* (*decorated*) *nested sequents*, or, simply, (*decorated*) *generalized sequents* $\mho$, and their skeletons $\mho^\circ$ as follows:

- for any (decorated) sequent Γ, we say that $\mho = \Gamma$ is a (decorated) generalized sequent with $\mho^\circ := \Gamma^\circ$;

- if $\mho_1$ and $\mho_2$ are (decorated) generalized sequents with $\mho_1^\circ = \mho_2^\circ$, then $(\mho_1 \mathbin{\varoslash} \mho_2)$ and $(\mho_1 \mathbin{\varobslash} \mho_2)$ are also (decorated) generalized sequents and $(\mho_1 \mathbin{\varoslash} \mho_2)^\circ = (\mho_1 \mathbin{\varobslash} \mho_2)^\circ := \mho_1^\circ = \mho_2^\circ$.

The operations $\varoslash$ and $\varobslash$ are purely syntactic and are called *external disjunction and conjunction* respectively.

As before, we denote a decorated generalized sequent by $\mho^*$ if $\mho$ is the generalized sequent obtained by removing all decorations from $\mho^*$. In this case, $\mho^*$ is called a *decoration* of $\mho$.

The above definition of a decoration requires a trivial correctness check to show that external disjunction and conjunction can be applied to two generalized sequents whenever they can be applied to their decorations. Part of this correctness would be useful as a stand-alone fact.

Fact 3.16. *If a decorated generalized sequent $\mho^*$ is a decoration of a generalized sequent $\mho$, then the decorated generalized sequent $(\mho^*)^\circ$ is a decoration of the generalized sequent $\mho^\circ$.*

If decorated generalized sequents $\mho_1^*$ and $\mho_2^*$ are decorations of generalized sequents $\mho_1$ and $\mho_2$ respectively and $(\mho_1^*)^\circ = (\mho_2^*)^\circ$, then $\mho_1^\circ = \mho_2^\circ$. Moreover, $\mho_1^* \mathbin{\varoslash} \mho_2^*$ and $\mho_1^* \mathbin{\varobslash} \mho_2^*$ are decorations of $\mho_1 \mathbin{\varoslash} \mho_2$ and $\mho_1 \mathbin{\varobslash} \mho_2$ respectively.

The external disjunction $\varoslash$ and conjunction $\varobslash$ on generalized sequents play the role of the disjunction and conjunction respectively of formula interpolants for ordinary sequents:

Definition 3.17 (True generalized-sequent decoration). Just like in the case of decorated sequents, we write $\Vdash \mho^*$ to denote that the decorated generalized sequent $\mho^*$ is *true*.

- If $\mho^* = \Gamma^*$ for some decorated sequent Γ^*, then $\Vdash \mho^*$ means that $\Vdash \Gamma^*$.

- If $\mho^* = \mho_1^* \mathbin{\varoslash} \mho_2^*$, then $\Vdash \mho^*$ means that $\Vdash \mho_1^*$ or $\Vdash \mho_2^*$.

- If $\mho^* = \mho_1^* \otimes \mho_2^*$, then $\Vdash \mho^*$ means that $\Vdash \mho_1^*$ and $\Vdash \mho_2^*$.

By now, we have defined skeletons of multiple types of objects. Skeletons of sequents, biased sequents, and generalized sequents are sequents, while those of contexts and biased contexts are contexts. Likewise, skeletons of decorated sequents, decorated biased sequents, and decorated generalized sequents are decorated sequents, while those of decorated contexts and decorated biased contexts are decorated contexts.

Definition 3.18 (Structural equivalence). Two objects, each belonging to one of the categories mentioned in the paragraph above (not necessarily to the same category), are called *structurally equivalent* if their skeletons coincide. Structural equivalence is denoted by $\sim$. For brevity's sake, we also call structurally equivalent decorations *matching*.

Definition 3.19 (Shallowness). A sequent, a biased sequent, or a generalized sequent is called *shallow* if its skeleton is ε.
A context or a biased context is called *shallow* if its skeleton is $\{\ \}$.
A decorated sequent, a decorated biased sequent, or a decorated generalized sequent is called *shallow* if its skeleton is w for some world w.
A decorated context or a decorated biased context is called *shallow* if its skeleton is $w, \{\ \}_w$.

In particular, we require that an interpolant of a biased sequent $\widetilde{\Gamma}$ be a generalized sequent $\mho$ structurally equivalent to it: $\mho \sim \widetilde{\Gamma}$.
We now define formally the logical consequence to be used in the interpolation statement.

Definition 3.20 (Decorative consequence). Let $\mho_1$ and $\mho_2$ be structurally equivalent generalized sequents. We say that $\mho_1$ *decoratively implies* $\mho_2$, written $\mho_1 \Vdash \mho_2$, if $\Vdash \mho_1^*$ impies $\Vdash \mho_2^*$ for arbitrary matching decorations $\mho_1^*$ and $\mho_2^*$. We say that the negation of $\mho_1$ *decoratively implies* $\mho_2$, written $\neg \mho_1 \Vdash \mho_2$, if $\nVdash \mho_1^*$ implies $\Vdash \mho_2^*$ for arbitrary matching decorations $\mho_1^*$ and $\mho_2^*$. Note that each nested sequent can be viewed as a generalized sequent. Thus, this definition is applicable to nested sequents too.

It should be pointed out that $\neg$ in the definition above is not an operation on generalized sequents: we do not define a generalized sequent $\neg \mho_1$. Rather, it is a notation for assuming that decorations of $\mho_1$ are false instead of true.[6]
The following fact shows that our definition of logical consequence is not degenerate: namely, for any pair of structurally equivalent generalized sequents, there exist matching decorations. In other words, $\mho_1$ can never vacuously decoratively imply $\mho_2$.

Fact 3.21. *Let $\mho_1$ and $\mho_2$ be structurally equivalent generalized sequents. There exist decorations of $\mho_1$, and for any decoration $\mho_1^*$, there exists a unique matching decoration $\mho_2^*$.*

Definition 3.22 (Propositional variables of nested objects and generalized sequents). Let Φ be a possibly decorated and/or biased nested object or a possibly decorated generalized sequent. The *set of propositions of Φ* is denoted $\mathrm{Prop}(\Phi)$ and defined as follows: a propositional variable $P \in \mathrm{Prop}(\Phi)$ iff P occurs in some (biased) formula from Φ.

With all this machinery in place, we are finally ready to formulate the interpolation statement:

Definition 3.23 (Interpolant of a biased nested sequent). An *interpolant* of a biased sequent $\widetilde{\Delta}$ is a generalized sequent $\mho$ such that

(A) $\mho \sim \widetilde{\Delta}$,   (B) $\neg \mho \Vdash L\widetilde{\Delta}$,   (C) $\mho \Vdash R\widetilde{\Delta}$,   and   (D) $\mathrm{Prop}(\mho) \subseteq \mathrm{Prop}(L\widetilde{\Delta}) \cap \mathrm{Prop}(R\widetilde{\Delta})$.

We write $\widetilde{\Delta} \longleftarrow \mho$ to denote the fact that $\mho$ is an interpolant of $\widetilde{\Delta}$.

---

[6]It is possible to define a transformation of $\mho_1$ that would yield a generalized sequent whose decoration is true whenever the matching decoration of $\mho_1$ is false. However, the only motivation for introducing such an operation would be to mimic this semantics of negation on decorations, which can be done directly.

For parts (B) and (C) of the above definition not to be vacuous, the following trivial correctness check is needed: for any $\mho$ that satisfies (A), we have that $\mho \sim L\widetilde{\Delta} \sim R\widetilde{\Delta}$.

Example 3.24. Let us now return to Example 3.2, which presents problems for the use of corresponding formulas, and show that the intuitively fitting interpolant from Example 3.12 is indeed an interpolant according to our decorative semantics: namely, $\left[P^{\Re}, \neg P^{\ell}\right] \longleftarrow [P]$. The conditions (A) and (D) are fulfilled because $[P] \sim \left[P^{\Re}, \neg P^{\ell}\right]$ and $\mathrm{Prop}\left([P]\right) \subseteq \mathrm{Prop}(\underbrace{[\neg P]}_{L[P^{\Re},\neg P^{\ell}]}) \cap \mathrm{Prop}(\underbrace{[P]}_{R[P^{\Re},\neg P^{\ell}]})$. It remains to show that, for any Kripke model $\mathcal{M} = (W, R, V)$ and any $\mathcal{M}$-decoration of $\left[P^{\Re}, \neg P^{\ell}\right]$, i.e., for any $w, \left[v, P^{\Re}, \neg P^{\ell}\right]$ with $w, v \in W$ such that wRv,

$$\nVdash w, [v, P] \implies \Vdash w, [v, \neg P] \qquad \text{and} \qquad \Vdash w, [v, P] \implies \Vdash w, [v, P] \ .$$

The second implication is trivial. To show the first implication, it is sufficient to note that

$$\nVdash w, [v, P] \implies \mathcal{M}, v \nVdash P \implies \mathcal{M}, v \Vdash \neg P \implies \Vdash w, [v, \neg P] \ .$$

While our semantics for interpolants fits well with our intuition on what the interpolant should be, one may ask why we call it an interpolant if its semantics is so different from the standard one. We argue that the differences are not that significant. While our semantics is not equivalent to the standard one, it *is* semi-equivalent: refutability of a sequent coincides with respect to the two semantics (see Theorem 3.26). In other words, while the two semantics may differ with respect to satisfiability, they are equivalent as far as validity is concerned. Given that interpolation statements are exclusively about validity, it plays no role which of the two semantics should be used in the context of the CIP.

The following auxiliary fact is used to demonstrate the connection:

Fact 3.25. *Let $\mathcal{M} = (W, R, V)$ be a Kripke model.*

1. *Any $\mathcal{M}$-decoration of the empty sequent ε has the form w for some $w \in W$ and is always false: $\nVdash w$.*
2. *For arbitrary $\mathcal{M}$-decorations Δ^* and $\Pi^\star$ with $r(\Delta^*)Rr(\Pi^\star)$ and any formula A, we have that Δ^*, A is an $\mathcal{M}$-decoration of Δ, A and that $\Delta^*, [\Pi^\star]$ is an $\mathcal{M}$-decoration of $\Delta, [\Pi]$. Both Δ^*, A and $\Delta^*, [\Pi^\star]$ have the same root as Δ^*.*
3. *Any $\mathcal{M}$-decoration of Δ, A has the form Δ^*, A for some $\mathcal{M}$-decoration Δ^*, and*

$$\Vdash \Delta^*, A \qquad \Longleftrightarrow \qquad \Vdash \Delta^* \quad or \quad \mathcal{M}, r(\Delta^*) \Vdash A \ .$$

4. *Any $\mathcal{M}$-decoration of $\Delta, [\Pi]$ has the form $\Delta^*, [\Pi^\star]$ for some $\mathcal{M}$-decorations Δ^* and $\Pi^\star$ such that $r(\Delta^*)Rr(\Pi^\star)$, and*

$$\Vdash \Delta^*, [\Pi^\star] \qquad \Longleftrightarrow \qquad \Vdash \Delta^* \quad or \quad \Vdash \Pi^\star \ .$$

Theorem 3.26 (Relationship between decorations and corresponding formulas). *Let Γ be a sequent and $\mathcal{M} = (W, R, V)$ be a Kripke model. Then for any world $w \in W$,*

$$\mathcal{M}, w \nVdash \underline{\Gamma} \qquad \Longleftrightarrow \qquad \nVdash \Gamma^* \text{ for some } \mathcal{M}\text{-decoration } \Gamma^* \text{ with } r(\Gamma^*) = w \ .$$

PROOF. The statement is proved for an arbitrary $w \in W$ by induction on the construction of Γ. We only show the case for $\Gamma = \Delta, [\Pi]$. Then $\mathsf{K} \vdash \underline{\Gamma} \leftrightarrow \underline{\Delta} \vee \Box \underline{\Pi}$ (note that $\underline{\Gamma} \neq \underline{\Delta} \vee \Box \underline{\Pi}$ for $\Delta = \varepsilon$). For any $w \in W$, by soundness of Kripke semantics, Fact 3.25, and induction hypothesis,

$$\mathcal{M}, w \nVdash \underline{\Gamma} \iff \mathcal{M}, w \nVdash \underline{\Delta, [\Pi]} \iff \mathcal{M}, w \nVdash \underline{\Delta} \text{ and } \mathcal{M}, w \nVdash \Box \underline{\Pi} \iff$$

$$\mathcal{M}, w \nVdash \underline{\Delta} \text{ and } \mathcal{M}, v \nVdash \underline{\Pi} \text{ for some } v \in W \text{ with } wRv \iff$$

$$\nVdash \Delta^* \text{ for some } \Delta^* \text{ with root } w \text{ and } \nVdash \Pi^\star \text{ for some } \Pi^\star \text{ with root } v \text{ such that } wRv \iff$$

$$\nVdash \Delta^*, [\Pi^\star] \text{ for some decoration } \Delta^*, [\Pi^\star] \text{ of } \Delta, [\Pi] \text{ with root } w. \quad \square$$

Corollary 3.27 (Completeness with respect to decorative semantics)**.** *A nested sequent is derivable in* NK *iff all its decorations are true.*

Definition 3.28 (Corresponding formula of a generalized sequent).

- If a generalized sequent $\mho = \Gamma$ for some sequent Γ, then $\underline{\mho} := \underline{\Gamma}$.

- If $\mho_1 \sim \mho_2$, then $\underline{\mho_1 \varowedge \mho_2} := \underline{\mho_1} \vee \underline{\mho_2}$ and $\underline{\mho_1 \varolessthan \mho_2} := \underline{\mho_1} \wedge \underline{\mho_2}$.

Corollary 3.29. *Let $\mho$ be a generalized sequent and $\mathcal{M} = (W, R, V)$ be a Kripke model. For any $w \in W$,*

$$\mathcal{M}, w \nVdash \underline{\mho} \qquad \Longleftrightarrow \qquad \nVdash \mho^* \text{ for some } \mathcal{M}\text{-decoration } \mho^* \text{ with } r(\mho^*) = w \ .$$

In particular, the decorative semantics fully coincides with the semantics of corresponding formulas on shallow sequents:

Corollary 3.30. *Let $A_1, \ldots, A_n$ for some $n \geq 0$ be a shallow sequent and $\mathcal{M} = (W, R, V)$ be a Kripke model. Then for any world $w \in W$, we have $\mathcal{M}, w \Vdash \underline{A_1, \ldots, A_n}$ iff $\Vdash w, A_1, \ldots, A_n$.*

PROOF. By Theorem 3.26, the corresponding formula of $A_1, \ldots, A_n$ is false at w iff some decoration of the sequent with root w is false. But there is only one such decoration: $w, A_1, \ldots, A_n$. Hence, the corresponding formula is false at w iff this decoration is false. Equivalently, the corresponding formula is true at w iff this decoration is true. □

Remark 3.31. In the general case, the corresponding formula is true at w iff *all* decorations with root w are true. Thus, having one true decoration is not sufficient to make the corresponding formula true.

Corollary 3.32. *Let $\mho$ be a generalized sequent with $\mho^{\circ} = \varepsilon$ and $\mathcal{M} = (W, R, V)$ be a Kripke model. Then for any world $w \in W$, there exists a unique $\mathcal{M}$-decoration $\mho^*$ with $r(\mho^*) = w$ and for this decoration*

$$\mathcal{M}, w \Vdash \underline{\mho} \qquad \Longleftrightarrow \qquad \Vdash \mho^* \ .$$

Corollary 3.33. *Let a generalized sequent $\mho$ be an interpolant of a shallow biased sequent $\widetilde{\Delta}$. Then*

$$(\text{B}') \ \mathsf{K} \vdash \neg \underline{L\widetilde{\Delta}} \supset \underline{\mho}, \qquad (\text{C}') \ \mathsf{K} \vdash \underline{\mho} \supset \underline{R\widetilde{\Delta}}, \qquad \text{and} \qquad (\text{D}') \ \mathrm{Prop}(\mho) \subseteq \mathrm{Prop}(L\widetilde{\Delta}) \cap \mathrm{Prop}(R\widetilde{\Delta}) \ .$$

Thus, for the corresponding split $L\widetilde{\Delta} \mid R\widetilde{\Delta}$ of the shallow sequent Δ, a formula interpolant of the split can be obtained by taking the corresponding formula of the generalized-sequent interpolant $\mho$ of $\widetilde{\Delta}$.

PROOF. Being structurally equivalent to $\widetilde{\Delta}$, our interpolant $\mho$ must be shallow. Clearly, $\mathrm{Prop}(\mho) = \mathrm{Prop}(\underline{\mho})$. Thus, the statement (D') follows from (D) in Definition 3.23.

To prove (B'), suppose $\mathcal{M}, w \Vdash \neg \underline{\mho}$ for some Kripke model $\mathcal{M} = (W, R, V)$ and some $w \in W$. Then $\mathcal{M}, w \nVdash \underline{\mho}$. By Corollary 3.32, there exists a unique $\mathcal{M}$-decoration $\mho^*$ of the shallow $\mho$ with root w and, for this decoration, $\nVdash \mho^*$. Since $\neg \mho \Vdash L\widetilde{\Delta}$ by statement (B) of Definition 3.23, we conclude $\Vdash w, L\widetilde{\Delta}$ for the matching $\mathcal{M}$-decoration $w, L\widetilde{\Delta}$ of $L\widetilde{\Delta}$. By Corollary 3.30 applied to the shallow sequent $L\widetilde{\Delta}$, we have $\mathcal{M}, w \Vdash \underline{L\widetilde{\Delta}}$. We have demonstrated that the formula $\underline{L\widetilde{\Delta}}$ holds whenever the formula $\underline{\mho}$ does not. By completeness, $\mathsf{K} \vdash \neg \underline{L\widetilde{\Delta}} \supset \underline{\mho}$. The proof of (C') is analogous. □

4. Biasing a Nested Derivation

Our goal is to interpolate biased sequents by induction on the nested sequent derivation, but the calculus NK presented in Definition 2.5 is for sequents, not for biased sequents. We now repair this mismatch by presenting a corresponding proof system BNK for biased sequents.

Definition 4.1 (Biased nested sequent calculus BNK). The rules of the biased nested sequent calculus BNK can be found in Figure 4.

$$\mathsf{id}_\mathsf{P}^{lr}\ \overline{\widetilde{\Gamma}\{P^\ell, \neg P^\Re\}} \qquad \mathsf{id}_\mathsf{P}^{rl}\ \overline{\widetilde{\Gamma}\{P^\Re, \neg P^\ell\}} \qquad \mathsf{id}_\mathsf{P}^{ll}\ \overline{\widetilde{\Gamma}\{P^\ell, \neg P^\ell\}} \qquad \mathsf{id}_\mathsf{P}^{rr}\ \overline{\widetilde{\Gamma}\{P^\Re, \neg P^\Re\}}$$

$$\mathsf{id}_\top^{l}\ \overline{\widetilde{\Gamma}\{\top^\ell\}} \qquad \mathsf{id}_\top^{r}\ \overline{\widetilde{\Gamma}\{\top^\Re\}} \qquad \mathsf{id}_{\neg\bot}^{l}\ \overline{\widetilde{\Gamma}\{\neg\bot^\ell\}} \qquad \mathsf{id}_{\neg\bot}^{r}\ \overline{\widetilde{\Gamma}\{\neg\bot^\Re\}}$$

$$\neg\neg^{l}\ \frac{\widetilde{\Gamma}\{A^\ell\}}{\widetilde{\Gamma}\{\neg\neg A^\ell\}} \qquad \neg\neg^{r}\ \frac{\widetilde{\Gamma}\{A^\Re\}}{\widetilde{\Gamma}\{\neg\neg A^\Re\}} \qquad \mathsf{ctr}^{l}\ \frac{\widetilde{\Gamma}\{A^\ell, A^\ell\}}{\widetilde{\Gamma}\{A^\ell\}} \qquad \mathsf{ctr}^{r}\ \frac{\widetilde{\Gamma}\{A^\Re, A^\Re\}}{\widetilde{\Gamma}\{A^\Re\}}$$

$$\alpha^{l}\ \frac{\widetilde{\Gamma}\{\alpha_1^\ell\} \quad \widetilde{\Gamma}\{\alpha_2^\ell\}}{\widetilde{\Gamma}\{\alpha^\ell\}} \qquad \alpha^{r}\ \frac{\widetilde{\Gamma}\{\alpha_1^\Re\} \quad \widetilde{\Gamma}\{\alpha_2^\Re\}}{\widetilde{\Gamma}\{\alpha^\Re\}} \qquad \beta^{l}\ \frac{\widetilde{\Gamma}\{\beta_1^\ell, \beta_2^\ell\}}{\widetilde{\Gamma}\{\beta^\ell\}} \qquad \beta^{r}\ \frac{\widetilde{\Gamma}\{\beta_1^\Re, \beta_2^\Re\}}{\widetilde{\Gamma}\{\beta^\Re\}}$$

$$\mathsf{adtr}_\mathsf{ff}^{lr}\ \frac{\widetilde{\Gamma}\{A^\Re, B^\ell\}}{\widetilde{\Gamma}\{B^\ell, A^\Re\}} \qquad \mathsf{adtr}_\mathsf{ff}^{rl}\ \frac{\widetilde{\Gamma}\{A^\ell, B^\Re\}}{\widetilde{\Gamma}\{B^\Re, A^\ell\}} \qquad \mathsf{adtr}_\mathsf{ff}^{ll}\ \frac{\widetilde{\Gamma}\{A^\ell, B^\ell\}}{\widetilde{\Gamma}\{B^\ell, A^\ell\}} \qquad \mathsf{adtr}_\mathsf{ff}^{rr}\ \frac{\widetilde{\Gamma}\{A^\Re, B^\Re\}}{\widetilde{\Gamma}\{B^\Re, A^\Re\}}$$

$$\mathsf{adtr}_\mathsf{f[]}^{l}\ \frac{\widetilde{\Gamma}\{A^\ell, [\widetilde{\Sigma}]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma}], A^\ell\}} \qquad \mathsf{adtr}_\mathsf{f[]}^{r}\ \frac{\widetilde{\Gamma}\{A^\Re, [\widetilde{\Sigma}]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma}], A^\Re\}} \qquad \mathsf{adtr}_\mathsf{[]f}^{l}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}], A^\ell\}}{\widetilde{\Gamma}\{A^\ell, [\widetilde{\Sigma}]\}} \qquad \mathsf{adtr}_\mathsf{[]f}^{r}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}], A^\Re\}}{\widetilde{\Gamma}\{A^\Re, [\widetilde{\Sigma}]\}}$$

$$\mathsf{adtr}_\mathsf{[][]}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Delta}], [\widetilde{\Sigma}]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma}], [\widetilde{\Delta}]\}}$$

$$\nu^{l}\ \frac{\widetilde{\Gamma}\{[\nu_0^\ell]\}}{\widetilde{\Gamma}\{\nu^\ell\}} \qquad \nu^{r}\ \frac{\widetilde{\Gamma}\{[\nu_0^\Re]\}}{\widetilde{\Gamma}\{\nu^\Re\}} \qquad \pi^{l}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}, \pi_0^\ell]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma}], \pi^\ell\}} \qquad \pi^{r}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}, \pi_0^\Re]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma}], \pi^\Re\}}$$

Figure 4: Rules of the biased nested sequent calculus BNK.

Before stating and proving the formal correspondence between NK and BNK, we describe how the latter was obtained from the former. Since each formula in a biased sequent is biased, all the rules of NK with exactly one principal formula must be duplicated into two versions differing in its bias. In either case, each active formula of the rule is biased the same way as the principal one, while all the biases of the side formulas remain unchanged. The name of such a BNK-rule is formed by adding a superscript l or r to the name of the corresponding NK-rule to encode the bias of the principal formula. It remains to describe what happens with the rules id_P and exch. The zero-premise NK-rule id_P has two principal formulas, P and $\neg P$, which can be biased in four different ways, yielding four BNK-rules: $\mathsf{id}_\mathsf{P}^{lr}$, $\mathsf{id}_\mathsf{P}^{rl}$, $\mathsf{id}_\mathsf{P}^{ll}$, and $\mathsf{id}_\mathsf{P}^{rr}$. The superscript of these rules encodes first the bias of P and then the bias of $\neg P$. For instance, $\mathsf{id}_\mathsf{P}^{rl}$ is the version with $P^\Re$ and $\neg P^\ell$. The situation with the NK-rule exch is slightly more complicated. As discussed in Example 3.12, in order to be able to match the decoration of the interpolant and the decoration of a biased sequent it interpolates, the correspondence between the structural boxes of the interpolant and of the biased sequent is maintained. This correspondence is read from the order of structural boxes, meaning that an interpolant must be changed whenever structural boxes are rearranged within the biased sequent. Since exch rearranges structural boxes in an unpredictable way, in BNK we use simpler *adjacent transposition* adtr rules, which are special cases of exch. They are sufficient because exch permutes elements in one of the sequent nodes, and it is well known that any permutation can be represented as a composition of adjacent transpositions. Since (biased) sequents consist of formulas and structural boxes, there are four types of adjacent transpositions to consider: a formula with a formula, a formula with a structural box, a structural box with a formula, and a structural box with a structural box. In addition, we have to consider all possible biases of all principal formula(s).

Altogether this yields nine adtr rules. In the rule's name, the subscript and the superscript state the type of the conclusion and the biases of the principal formula(s) respectively.

We now show that the biased sequent calculus BNK is equivalent to the nested sequent calculus NK.

Theorem 4.2 (Equivalence between BNK and NK). *For any biased version $\widetilde{\Gamma}$ of a nested sequent Γ,* BNK $\vdash \widetilde{\Gamma}$ *iff* NK $\vdash \Gamma$.

PROOF. The direction from left to right is trivial since the unbiased version of each BNK-rule is the corresponding NK-rule (unbiased versions of adtr rules correspond to exch). Thus, erasing all biases in a BNK-derivation of $\widetilde{\Gamma}$ yields an NK-derivation of Γ.

For the direction from right to left, given an NK-derivation $\mathcal{D}$ of Γ, consider an arbitrary biased version $\widetilde{\Gamma}$ of Γ. First replace each application of exch in $\mathcal{D}$ by an equivalent sequence of applications of adjacent transpositions, yielding an NK-derivation $\mathcal{D}'$ of Γ where all instances of exch have one of the following forms:

$$\frac{\Delta\{A, B\}}{\Delta\{B, A\}} \quad , \quad \frac{\Delta\{A, [\Pi]\}}{\Delta\{[\Pi], A\}} \quad , \quad \frac{\Delta\{[\Pi], A\}}{\Delta\{A, [\Pi]\}} \quad , \quad \text{and} \quad \frac{\Delta\{[\Sigma], [\Pi]\}}{\Delta\{[\Pi], [\Sigma]\}} \quad . \tag{4}$$

It suffices to bias the formulas at the conclusion Γ of $\mathcal{D}'$ according to $\widetilde{\Gamma}$ and then bias all the remaining formulas in $\mathcal{D}'$ in such a way that it becomes a BNK-derivation. We leave the details to the reader. $\qquad\square$

Example 4.3. We illustrate this biasing process by applying Theorem 4.2 to the derivation in Figure 3. We remove the last rule from the derivation to have two formulas in the conclusion, bias the formulas, and propagate the bias bottom up through the derivation. The resulting biased derivation can be found in Figure 5. The labels $(\star)$ and $(\star\star)$ on the applications of ν in Figure 3 are preserved: both become applications of ν^r. The labels $(\dagger)$, $(\dagger\dagger)$, $(\dagger\dagger\dagger)$, and $(\dagger\dagger\dagger\dagger)$ on the applications of the rule exch in Figure 3 are transferred to the corresponding adtr rules in Figure 5. $(\dagger\dagger)$ and $(\dagger\dagger\dagger)$ were already adjacent transpositions and, thus, only require one adtr rule each. Each of $(\dagger)$ and $(\dagger\dagger\dagger\dagger)$ can be emulated by two adjacent transpositions.

$$
\begin{array}{c}
\text{id}_\mathsf{P}^\mathsf{rl} \; \dfrac{}{[\neg P^\mathfrak{R}, R^\mathfrak{R}, [Q^\mathfrak{R}, \neg Q^\ell]]} \\[4pt]
\pi^\mathsf{l} \; \dfrac{}{[\neg P^\mathfrak{R}, R^\mathfrak{R}, [Q^\mathfrak{R}], \neg\square Q^\ell]} \\[4pt]
\text{adtr}_{\mathsf{f}[\,]}^\mathsf{l} \; \dfrac{}{[\neg P^\mathfrak{R}, R^\mathfrak{R}, \neg\square Q^\ell, [Q^\mathfrak{R}]]} \;(\dagger\dagger) \\[4pt]
\nu^\mathsf{r} \; \dfrac{}{[\neg P^\mathfrak{R}, R^\mathfrak{R}, \neg\square Q^\ell, \square Q^\mathfrak{R}]} \;(\star)
\end{array}
\qquad
\begin{array}{c}
\text{id}_\mathsf{P}^\mathsf{rr} \; \dfrac{}{[\neg P^\mathfrak{R}, R^\mathfrak{R}, \neg R^\mathfrak{R}, \neg\square Q^\ell]} \\[4pt]
\text{adtr}_{\mathsf{ff}}^\mathsf{lr} \; \dfrac{}{[\neg P^\mathfrak{R}, R^\mathfrak{R}, \neg\square Q^\ell, \neg R^\mathfrak{R}]} \;(\dagger\dagger\dagger)
\end{array}
$$

Left branch:
$$
\text{id}_\mathsf{P}^\mathsf{lr} \; \dfrac{}{[P^\ell, \neg P^\mathfrak{R}, R^\mathfrak{R}, \neg(\square Q \supset R)^\mathfrak{R}]}
\qquad
\text{adtr}_{\mathsf{ff}}^\mathsf{rl} \; \dfrac{}{[\neg P^\mathfrak{R}, P^\ell, R^\mathfrak{R}, \neg(\square Q \supset R)^\mathfrak{R}]} \;(\dagger)
\qquad
\text{adtr}_{\mathsf{ff}}^\mathsf{rl} \; \dfrac{}{[\neg P^\mathfrak{R}, R^\mathfrak{R}, P^\ell, \neg(\square Q \supset R)^\mathfrak{R}]} \;(\dagger)
$$

$$
\alpha^\mathsf{r} \; \dfrac{}{[\neg P^\mathfrak{R}, R^\mathfrak{R}, \neg\square Q^\ell, \neg(\square Q \supset R)^\mathfrak{R}]}
$$

$$
\begin{array}{c}
\alpha^\mathsf{l} \; \dfrac{}{[\neg P^\mathfrak{R}, R^\mathfrak{R}, \neg(P \supset \square Q)^\ell, \neg(\square Q \supset R)^\mathfrak{R}]} \\[4pt]
\beta^\mathsf{r} \; \dfrac{}{[P \supset R^\mathfrak{R}, \neg(P \supset \square Q)^\ell, \neg(\square Q \supset R)^\mathfrak{R}]} \\[4pt]
\pi^\mathsf{r} \; \dfrac{}{[P \supset R^\mathfrak{R}, \neg(P \supset \square Q)^\ell], \neg\square(\square Q \supset R)^\mathfrak{R}} \\[4pt]
\text{adtr}_{\mathsf{f}[\,]}^\mathsf{r} \; \dfrac{}{\neg\square(\square Q \supset R)^\mathfrak{R}, [P \supset R^\mathfrak{R}, \neg(P \supset \square Q)^\ell]} \\[4pt]
\pi^\mathsf{l} \; \dfrac{}{\neg\square(\square Q \supset R)^\mathfrak{R}, [P \supset R^\mathfrak{R}], \neg\square(P \supset \square Q)^\ell} \\[4pt]
\text{adtr}_{\mathsf{f}[\,]}^\mathsf{l} \; \dfrac{}{\neg\square(\square Q \supset R)^\mathfrak{R}, \neg\square(P \supset \square Q)^\ell, [P \supset R^\mathfrak{R}]} \;(\dagger\dagger\dagger\dagger) \\[4pt]
\text{adtr}_{\mathsf{ff}}^\mathsf{lr} \; \dfrac{}{\neg\square(P \supset \square Q)^\ell, \neg\square(\square Q \supset R)^\mathfrak{R}, [P \supset R^\mathfrak{R}]} \;(\dagger\dagger\dagger\dagger) \\[4pt]
\nu^\mathsf{r} \; \dfrac{}{\neg\square(P \supset \square Q)^\ell, \neg\square(\square Q \supset R)^\mathfrak{R}, \square(P \supset R)^\mathfrak{R}} \;(\star\star) \\[4pt]
\beta^\mathsf{r} \; \dfrac{}{\neg\square(P \supset \square Q)^\ell, \square(\square Q \supset R) \supset \square(P \supset R)^\mathfrak{R}}
\end{array}
$$

Figure 5: BNK derivation of $\neg\square(P \supset \square Q)^\ell, \square(\square Q \supset R) \supset \square(P \supset R)^\mathfrak{R}$ that results from biasing the derivation from Figure 3.

5. More on Decorative Semantics

This section mostly contains auxiliary technical lemmas. Since it is a new semantics, we need to define what it means to be semantically equivalent with respect to this semantics. Semantical equivalence is needed to replace an interpolant with an equivalent one in a requisite form.

Definition 5.1 (Decorative equivalence). Two generalized sequents are called (*decoratively*) *equivalent*, written $\mho \dashv\vdash \mho'$, if $\mho \Vdash \mho'$ and $\mho' \Vdash \mho$ and if, in addition, $\text{Prop}(\mho) = \text{Prop}(\mho')$.

The usual properties of (internal) conjunction and disjunction with respect to logical consequence on formulas, including commutativity, associativity, and distributivity, transfer to external conjunction and disjunction with respect to decorated logical consequence on generalized sequents. As with other statements classified as facts in this paper, the proofs are omitted because they are simple and/or standard.

Fact 5.2 (Properties of $\oslash$ and $\otimes$). *Let $\mho$, $\mho_1$, $\mho_2$, $\mho_1'$, $\mho_2'$, and $\mho_3$ be generalized sequents.*

1. $\dashv\vdash$ *is an equivalence relation on generalized sequents.*
2. *If $\mho_1 \otimes \mho_2$ is defined, then $\mho_1 \otimes \mho_2 \Vdash \mho_i$ for each $i = 1, 2$.*
3. *If $\mho_1 \oslash \mho_2$ is defined, then $\mho_i \Vdash \mho_1 \oslash \mho_2$ for each $i = 1, 2$.*
4. *If $\mho_i \Vdash \mho$ for each $i = 1, 2$, then $\mho_1 \oslash \mho_2$ is defined and $\mho_1 \oslash \mho_2 \Vdash \mho$.*
5. *If $\mho \Vdash \mho_i$ for each $i = 1, 2$, then $\mho_1 \otimes \mho_2$ is defined and $\mho \Vdash \mho_1 \otimes \mho_2$.*
6. *$\mho_1 \otimes \mho_2$ is defined iff $\mho_2 \otimes \mho_1$ is defined, and $\mho_1 \otimes \mho_2 \dashv\vdash \mho_2 \otimes \mho_1$ when both are defined.*
7. *$(\mho_1 \otimes \mho_2) \otimes \mho_3$ is defined iff $\mho_1 \otimes (\mho_2 \otimes \mho_3)$ is, and $(\mho_1 \otimes \mho_2) \otimes \mho_3 \dashv\vdash \mho_1 \otimes (\mho_2 \otimes \mho_3)$ when both are.*
8. *$\mho \otimes (\mho_1 \oslash \mho_2)$ is defined iff $(\mho \otimes \mho_1) \oslash (\mho \otimes \mho_2)$ is, and $\mho \otimes (\mho_1 \oslash \mho_2) \dashv\vdash (\mho \otimes \mho_1) \oslash (\mho \otimes \mho_2)$ when both are.*
9. *If $\mho_1 \dashv\vdash \mho_1'$ and $\mho_2 \dashv\vdash \mho_2'$, then $\mho_1 \oslash \mho_2$ is defined iff $\mho_1' \oslash \mho_2'$ is and $\mho_1 \oslash \mho_2 \dashv\vdash \mho_1' \oslash \mho_2'$ when both are.*

Properties of $\oslash$ analogous to Properties 6–9 of $\otimes$ also hold.

This fact shows that parentheses in external conjunctions and disjunctions of generalized sequents can be omitted using the standard propositional conventions. In trivial cases, we omit proofs that objects exist.

Definition 5.3 (Void nested objects and singleton sequents). We call a nested object *void* if its construction does not involve the clause Φ, C. A sequent is called *singleton* if its construction involves exactly one use of the Φ, C clause.

Definition 5.4 (SDNF and SCNF). Let Λ_{ij} be pairwise structurally equivalent singleton sequents for $1 \le i \le n$ and $1 \le j \le m_i$. Generalized sequents $\bigoslash_{i=1}^{n} \bigotimes_{j=1}^{m_i} \Lambda_{ij}$ and $\bigotimes_{i=1}^{n} \bigoslash_{j=1}^{m_i} \Lambda_{ij}$ are said to be in a *singleton disjunctive normal form* (SDNF) and *singleton conjunctive normal form* (SCNF) respectively.

In order to show that any generalized sequent has decoratively equivalent SDNF and SCNF representations, we need several auxiliary statements.

Fact 5.5. $\nVdash \Delta^*$ *for any decoration of a void sequent Δ.*

Lemma 5.6 (Sequent as an external disjunction of singleton sequents). *For any sequent Δ, there exists $n \ge 1$ and singleton sequents $\Lambda_1, \ldots, \Lambda_n$ structurally equivalent to Δ such that $\Delta \dashv\vdash \bigoslash_{i=1}^{n} \Lambda_i$.*

PROOF. The idea of the proof is that, because of Fact 5.5, arbitrary void nested structure can be added to a sequent without changing the truth or falsity of any of its decorations (we cannot say that the result of such additions is decoratively equivalent because it is not structurally equivalent). The proof is constructive

and proceeds by induction on the construction of Δ. We state the steps for all the cases and give a detailed proof for the case of $\Delta = \Delta_1, [\Delta_2]$.

Case $\Delta = \varepsilon$. Use $n := 1$ and $\Lambda_1 := \bot$.

Case $\Delta = \Delta_1, C$. By induction hypothesis, $\Delta_1 \dashv\vdash \widetilde{\bigotimes_{i=1}^{m}} \Lambda_i'$ for some singleton sequents Λ_i' structurally equivalent to Δ_1. Then $\Delta \dashv\vdash (\Delta^\circ, C) \otimes \widetilde{\bigotimes_{i=1}^{m}} \Lambda_i'$.

Case $\Delta = \Delta_1, [\Delta_2]$. By induction hypothesis, $\Delta_1 \dashv\vdash \widetilde{\bigotimes_{i=1}^{m}} \Lambda_i'$ and $\Delta_2 \dashv\vdash \widetilde{\bigotimes_{j=1}^{k}} \Lambda_j''$ for some singleton sequents Λ_i' structurally equivalent to Δ_1 and some singleton sequents Λ_j'' structurally equivalent to Δ_2. We add void sequents to blow up the structures of each Λ_i' and Λ_j'' to make them structurally equivalent to Δ without affecting the truth of the decorations of these singleton sequents: each Λ_i' is replaced with $\Lambda_i', [\Delta_2^\circ]$, and each Λ_j'' is replaced with $\Delta_1^\circ, [\Lambda_j'']$. We now show that $\Delta \dashv\vdash \widetilde{\bigotimes_{i=1}^{m}} (\Lambda_i', [\Delta_2^\circ]) \otimes \widetilde{\bigotimes_{j=1}^{k}} (\Delta_1^\circ, [\Lambda_j''])$. By induction hypothesis, a propositional variable occurs in Δ_1 iff it occurs in at least one of Λ_i' and occurs in Δ_2 iff it occurs in at least one of Λ_j''. Thus, a propositional variable occurs in Δ iff it occurs in $\widetilde{\bigotimes_{i=1}^{m}} (\Lambda_i', [\Delta_2^\circ]) \otimes \widetilde{\bigotimes_{j=1}^{k}} (\Delta_1^\circ, [\Lambda_j''])$.

Since $\Lambda_i', [\Delta_2^\circ] \sim \Delta_1^\circ, [\Lambda_j''] \sim \Delta_1^\circ, [\Delta_2^\circ] \sim \Delta_1, [\Delta_2] = \Delta$ for each $i = 1, \ldots, m$ and each $j = 1, \ldots, k$, it remains to show the decorative equivalence. Consider a Kripke model $\mathcal{M} = (W, R, V)$, a world $w \in W$, and arbitrary matching decorations of Δ, of $\Lambda_i', [\Delta_2^\circ]$, and of $\Delta_1^\circ, [\Lambda_j'']$. By Fact 3.25, they must have the form $\Delta_1^*, [\Delta_2^\star]$, the form $(\Lambda_i')^*, [(\Delta_2^\circ)^\star]$, and the form $(\Delta_1^\circ)^*, [(\Lambda_j'')^\star]$ respectively, where $\Delta_1^* \sim (\Lambda_i')^* \sim (\Delta_1^\circ)^*$ and $\Delta_2^\star \sim (\Delta_2^\circ)^\star \sim (\Lambda_j'')^\star$. By Facts 3.25 and 5.5, we have that $\Vdash \Delta_1^*, [\Delta_2^\star]$ iff $\Vdash \Delta_1^*$ or $\Vdash \Delta_2^\star$, that $\Vdash (\Lambda_i')^*, [(\Delta_2^\circ)^\star]$ iff $\Vdash (\Lambda_i')^*$, and that $\Vdash (\Delta_1^\circ)^*, [(\Lambda_j'')^\star]$ iff $\Vdash (\Lambda_j'')^\star$. By induction hypothesis, $\Vdash \Delta_1^*, [\Delta_2^\star]$ iff $\Vdash \Delta_1^*$ or $\Vdash \Delta_2^\star$ iff $\Vdash (\Lambda_i')^*$ for some $1 \le i \le m$ or $\Vdash (\Lambda_j'')^\star$ for some $1 \le j \le k$ iff $\Vdash (\Lambda_i')^*, [(\Delta_2^\circ)^\star]$ for some $1 \le i \le m$ or $\Vdash (\Delta_1^\circ)^*, [(\Lambda_j'')^\star]$ for some $1 \le j \le k$ iff $\Vdash \widetilde{\bigotimes_{i=1}^{m}}((\Lambda_i')^*, [(\Delta_2^\circ)^\star]) \otimes \widetilde{\bigotimes_{j=1}^{k}}((\Delta_1^\circ)^*, [(\Lambda_j'')^\star])$. □

Fact 5.7 (Conversion to SDNF and to SCNF). *Any generalized sequent $\mho$ can be converted to a SDNF and a SCNF, i.e, there exist $\mho_1$ and $\mho_2$ in a SDNF and in a SCNF respectively such that $\mho \dashv\vdash \mho_1$ and $\mho \dashv\vdash \mho_2$.*

The following lemma implies that interpolants can always be converted to the required form.

Lemma 5.8 (Interpolant transformation). *If $\mho_1$ interpolates $\widetilde{\Gamma}$, so does any generalized sequent $\mho_2$ decoratively equivalent to $\mho_1$.*

PROOF. It follows from the definition of equivalence that $\mathrm{Prop}(\mho_2) = \mathrm{Prop}(\mho_1) \subseteq \mathrm{Prop}(L\widetilde{\Gamma}) \cap \mathrm{Prop}(R\widetilde{\Gamma})$ and also that $\mho_2 \sim \mho_1 \sim \widetilde{\Gamma}$. For arbitrary matching decorations $\mho_2^*$, $L\widetilde{\Gamma}^*$, and $R\widetilde{\Gamma}^*$, there exists a unique matching decoration $\mho_1^*$. If $\Vdash \mho_2^*$, then $\Vdash \mho_1^*$ by equivalence, and, consequently, $\Vdash R\widetilde{\Gamma}^*$ because $\mho_1 \Vdash R\widetilde{\Gamma}$. If $\nVdash \mho_2^*$, then $\nVdash \mho_1^*$ by equivalence, and, consequently, $\Vdash L\widetilde{\Gamma}^*$ because $\neg\mho_1 \Vdash L\widetilde{\Gamma}$. □

Lemma 5.9.   1. *If $\Omega^\circ = \Gamma\{[\Delta]\}$, then $\Omega = \Lambda\{[\Pi]\}$ where $\Lambda^\circ\{\} = \Gamma\{\}$ and $\Pi^\circ = \Delta$.*
   2. *If $\Omega^\circ = \Gamma\{[\Delta], [\Theta]\}$, then $\Omega = \Lambda\{[\Pi], A_1, \ldots, A_n, [\Sigma]\}$ where $\Lambda^\circ\{\} = \Gamma\{\}$, $\Pi^\circ = \Delta$, $\Sigma^\circ = \Theta$, and $n \ge 0$.*

PROOF. Both statements can be proved by induction on the structure of Ω. We only show the case of $\Omega = \Omega_1, [\Omega_2]$ for the second statement. Let $\Omega = \Omega_1, [\Omega_2]$ and $\Omega^\circ = \Gamma\{[\Delta], [\Theta]\}$. There are three subcases depending on where $[\Delta], [\Theta]$ occurs in $\Omega^\circ = \Omega_1^\circ, [\Omega_2^\circ]$: it can be within Ω_1°, be within Ω_2°, or be the last two elements in the sequence Ω°. In the first subcase, $\Gamma\{\} = \Gamma_1\{\}, [\Omega_2^\circ]$ and $\Omega_1^\circ = \Gamma_1\{[\Delta], [\Theta]\}$. In the second subcase, $\Gamma\{\} = \Omega_1^\circ, [\Gamma_2\{\}]$ and $\Omega_2^\circ = \Gamma_2\{[\Delta], [\Theta]\}$. Both situations are handled by straightforward uses of the induction hypothesis for Ω_1 and Ω_2 respectively. The only subcase of interest is the third one,

where $\Gamma\{\,\} = \Xi_1, \{\,\}$, $\Omega_1^\circ = \Xi_1, [\Delta]$, and $\Omega_2^\circ = \Theta$. It follows that $\Omega_1 = \Upsilon_1, [\Pi], A_1, \dots, A_n$ where $\Upsilon_1^\circ = \Xi_1$, $\Pi^\circ = \Delta$, and $n \geq 0$. For $\Lambda\{\,\} := \Upsilon_1, \{\,\}$ and $\Sigma := \Omega_2$,

$$
\begin{aligned}
\Lambda^\circ\{\,\} \quad &= \quad \Upsilon_1^\circ, \{\,\} \quad = \quad \Xi_1, \{\,\} \quad = \quad \Gamma\{\,\} \ , \\
\Sigma^\circ \quad &= \quad \Omega_2^\circ \quad = \quad \Theta \ , \\
\Lambda\{[\Pi], A_1, \dots, A_n, [\Sigma]\} \quad &= \quad \Upsilon_1, [\Pi], A_1, \dots, A_n, [\Sigma] \quad = \quad \Omega_1, [\Omega_2] \quad = \quad \Omega \ . \ \square
\end{aligned}
$$

Definition 5.10 (Inserting a sequent decoration into a context decoration). Let $\Gamma^*\{\,\}_w$ and Δ^* be $\mathcal{M}$-decorations of a (biased) context $\Gamma\{\,\}$ and a (biased) sequent Δ respectively, such that $r(\Delta^*) = w$. Then the insertion of Δ^* into $\Gamma^*\{\,\}_w$ is defined by $\Gamma^*\{\Delta^*\}_w := \Gamma^*\{t(\Delta^*)\}$ where the insertion in the right side is done as if $\Gamma^*\{\,\}$ were a non-decorated context.

Fact 5.11 (Properties of decoration insertion). *If $\Gamma^*\{\Delta^*\}_w$ is defined for $\mathcal{M}$-decorations $\Gamma^*\{\,\}_w$ and Δ^* of a (biased) context $\Gamma\{\,\}$ and a (biased) sequent Δ respectively, it is an $\mathcal{M}$-decoration of $\Gamma\{\Delta\}$ with the same root as $\Gamma^*\{\,\}_w$. In addition, in this case,*

$$
\Vdash \Gamma^*\{\Delta^*\}_w \qquad \Longleftrightarrow \qquad \Vdash \Gamma^*\{w\}_w \quad or \quad \Vdash \Delta^* \ .
$$

Fact 5.12 (Decoration of an insertion as an insertion of decorations). *Let Θ^* be a decoration of a (biased) sequent $\Gamma\{\Delta\}$. Then there exist decorations $\Gamma^*\{\,\}_w$ of $\Gamma\{\,\}$ and Δ^* of Δ with $r(\Delta^*) = w$ such that $\Gamma^*\{\Delta^*\}_w$ is defined and is equal to Θ^*.*

6. Proof of the Interpolation Theorem

Algorithm 6.1 (Interpolation algorithm). Following the style of sequent-based proofs of interpolation, we present the interpolation algorithm as a biased sequent calculus supplied with interpolant management machinery. For the rules $\mathsf{adtr}_{[][]}$, ν^{l}, and ν^{r}, the interpolant must be in either a SDNF or a SCNF. In addition, the rule ν^{l} (ν^{r}) requires a particular order of conjuncts (disjuncts) within each disjunct (conjunct) of the SDNF (SCNF). The algorithm can be found in Figure 6. For every interpolant-handling rule in Figure 6, dropping the interpolants in the conclusion and the premise(s) always yields the corresponding BNK-rule. Whenever the interpolant is represented as a conjunction of disjunctions of sequents or a disjunction of conjunctions of sequents, it is assumed to be in a SCNF or a SDNF respectively; in particular, the sequents are assumed to be singleton.

To apply the interpolation algorithm from Figure 6 to a given BNK-derivation $\mathcal{D}$, assign interpolants to all the leaves of $\mathcal{D}$, i.e., to the zero-premise rules id, and propagate the interpolant assignment downwards to the root of $\mathcal{D}$. To show that the generalized sequent thus assigned to the conclusion of $\mathcal{D}$ is its interpolant, we need to show, for each interpolant-handling rule from Figure 6, that, given an interpolant for each premise of the rule, we **can** apply the algorithm and the result **will be** an interpolant for the conclusion of the rule.

To demonstrate the former, we need to show that any given interpolant of the premise of the rules $\mathsf{adtr}_{[][]}$, ν^{l}, and ν^{r} can be transformed into the required form. To demonstrate the latter, we need to show, assuming the given interpolant(s) for the premise(s) to be in the prescribed form, that the constructed interpolant is a well-defined generalized sequent structurally equivalent to the biased sequent at the conclusion of the rule (part (A) of Definition 3.23), that the propositional variables of the constructed interpolant are common to the left and right sides of the biased sequent (part (D) of Definition 3.23), and that the interpolant decoratively implies the right side of the biased sequent and the negation of the interpolant decoratively implies the left side of it (parts (B) and (C) of Definition 3.23).

We break the proof into three lemmas: the first about the structure of the interpolant and the applicability of the rules, the second about propositional variables, and the third about decorative consequences. But before proving these lemmas and stating the interpolation theorem as their corollary, let us demonstrate the algorithm using an example. The reader can independently verify that every line of this example is a valid interpolation statement.

$$\mathsf{id}_\mathsf{P}^{\mathsf{lr}}\ \overline{\widetilde{\Gamma}\{P^\ell, \neg P^\Re\} \longleftarrow \widetilde{\Gamma}^\circ\{\neg P\}} \qquad\qquad \mathsf{id}_\mathsf{P}^{\mathsf{rl}}\ \overline{\widetilde{\Gamma}\{P^\Re, \neg P^\ell\} \longleftarrow \widetilde{\Gamma}^\circ\{P\}}$$

$$\mathsf{id}_\mathsf{P}^{\mathsf{ll}}\ \overline{\widetilde{\Gamma}\{P^\ell, \neg P^\ell\} \longleftarrow \widetilde{\Gamma}^\circ\{\bot\}} \qquad \mathsf{id}_\top^{\mathsf{l}}\ \overline{\widetilde{\Gamma}\{\top^\ell\} \longleftarrow \widetilde{\Gamma}^\circ\{\bot\}} \qquad \mathsf{id}_{\neg\bot}^{\mathsf{l}}\ \overline{\widetilde{\Gamma}\{\neg\bot^\ell\} \longleftarrow \widetilde{\Gamma}^\circ\{\bot\}}$$

$$\mathsf{id}_\mathsf{P}^{\mathsf{rr}}\ \overline{\widetilde{\Gamma}\{P^\Re, \neg P^\Re\} \longleftarrow \widetilde{\Gamma}^\circ\{\top\}} \qquad \mathsf{id}_\top^{\mathsf{r}}\ \overline{\widetilde{\Gamma}\{\top^\Re\} \longleftarrow \widetilde{\Gamma}^\circ\{\top\}} \qquad \mathsf{id}_{\neg\bot}^{\mathsf{r}}\ \overline{\widetilde{\Gamma}\{\neg\bot^\Re\} \longleftarrow \widetilde{\Gamma}^\circ\{\top\}}$$

$$\alpha^{\mathsf{l}}\ \frac{\widetilde{\Gamma}\{\alpha_1^\ell\} \longleftarrow \mho_1 \quad \widetilde{\Gamma}\{\alpha_2^\ell\} \longleftarrow \mho_2}{\widetilde{\Gamma}\{\alpha^\ell\} \longleftarrow \mho_1 \ovee \mho_2} \qquad\qquad \alpha^{\mathsf{r}}\ \frac{\widetilde{\Gamma}\{\alpha_1^\Re\} \longleftarrow \mho_1 \quad \widetilde{\Gamma}\{\alpha_2^\Re\} \longleftarrow \mho_2}{\widetilde{\Gamma}\{\alpha^\Re\} \longleftarrow \mho_1 \owedge \mho_2}$$

$$\mathsf{adtr}_{[\,][\,]}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Delta}], [\widetilde{\Sigma}]\} \longleftarrow \bigowedge_{i=1}^{n} \bigovee_{j=1}^{m_i} \Lambda_{ij}\{[\Delta_{ij}], \mathbf{A}_{ij}, [\Sigma_{ij}]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma}], [\widetilde{\Delta}]\} \longleftarrow \bigowedge_{i=1}^{n} \bigovee_{j=1}^{m_i} \Lambda_{ij}\{[\Sigma_{ij}], \mathbf{A}_{ij}, [\Delta_{ij}]\}}$$

$$\nu^{\mathsf{l}}\ \frac{\widetilde{\Gamma}\{[\nu_0^\ell]\} \longleftarrow \bigovee_{i=1}^{n} \left(\bigowedge_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \owedge \bigovee_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right)}{\widetilde{\Gamma}\{\nu^\ell\} \longleftarrow \bigovee_{i=1}^{n} \left(\widetilde{\Gamma}^\circ\{\Diamond\left(\bigwedge_{k=1}^{l_i} A_{ik}\right)\} \owedge \bigovee_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\} \right)}$$

$$\nu^{\mathsf{r}}\ \frac{\widetilde{\Gamma}\{[\nu_0^\Re]\} \longleftarrow \bigowedge_{i=1}^{n} \left(\bigovee_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \ovee \bigowedge_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right)}{\widetilde{\Gamma}\{\nu^\Re\} \longleftarrow \bigowedge_{i=1}^{n} \left(\widetilde{\Gamma}^\circ\{\Box\left(\bigvee_{k=1}^{l_i} A_{ik}\right)\} \ovee \bigowedge_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\} \right)}$$

For all the remaining rules, i.e., for the rules $\neg\neg^{\mathsf{l}}$, $\neg\neg^{\mathsf{r}}$, $\mathsf{ctr}^{\mathsf{l}}$, $\mathsf{ctr}^{\mathsf{r}}$, β^{l}, β^{r}, $\mathsf{adtr}_{\mathsf{ff}}^{\mathsf{lr}}$, $\mathsf{adtr}_{\mathsf{ff}}^{\mathsf{rl}}$, $\mathsf{adtr}_{\mathsf{ff}}^{\mathsf{ll}}$, $\mathsf{adtr}_{\mathsf{ff}}^{\mathsf{rr}}$, $\mathsf{adtr}_{\mathsf{f}[\,]}^{\mathsf{l}}$, $\mathsf{adtr}_{\mathsf{f}[\,]}^{\mathsf{r}}$, $\mathsf{adtr}_{[\,]\mathsf{f}}^{\mathsf{l}}$, $\mathsf{adtr}_{[\,]\mathsf{f}}^{\mathsf{r}}$, π^{l}, and π^{r}, the given interpolant for the premise is used as an interpolant for the conclusion.

Figure 6: Interpolation algorithm for the calculus BNK. Interpolants for the premises of $\mathsf{adtr}_{[\,][\,]}$ and ν^{r} must be in a SCNF; an interpolant for the premise of ν^{l} must be in a SDNF. For these three rules, we require that $\widetilde{\Gamma}\{\ \} \sim \Lambda_{ij}\{\ \} \sim \Pi_{ik}\{\ \}$ for all suitable i, j, and k. Finally, $\mathbf{A}_{ij}$ in $\mathsf{adtr}_{[\,][\,]}$ must be shallow sequents for all suitable i and j.

Example 6.2. Let us apply the algorithm to the derivation of $\neg\Box(P \supset \Box Q)^\ell, \Box(\Box Q \supset R) \supset \Box(P \supset R)^{\Re}$ from Figure 5. The result can be found in Figure 7. Since most of the rules do not require the interpolant to be changed, we only explain those steps of the algorithm that do, starting from the three leaves of the derivation tree. The leftmost leaf, an application of the rule $\mathsf{id}_\mathsf{P}^{\mathsf{lr}}$, has the conclusion $\widetilde{\Gamma}_1\{P^\ell, \neg P^{\Re}\}$ with $\widetilde{\Gamma}_1 = [\{\ \}, R^{\Re}, \neg(\Box Q \supset R)^{\Re}]$. Given that $\widetilde{\Gamma}_1^\circ = [\{\ \}]$, we assign to this conclusion an interpolant $\widetilde{\Gamma}_1^\circ\{\neg P\} = [\neg P]$. Similarly, the middle leaf, an application of the rule $\mathsf{id}_\mathsf{P}^{\mathsf{rl}}$ with the conclusion $\widetilde{\Gamma}_2\{Q^{\Re}, \neg Q^\ell\}$ where $\widetilde{\Gamma}_2 = [\neg P^{\Re}, R^{\Re}, [\{\ \}]]$, is assigned an interpolant $\widetilde{\Gamma}_2^\circ\{Q\} = [[Q]]$, where $\widetilde{\Gamma}_2^\circ = [[\{\ \}]]$. Finally, the rightmost leaf, an application of the rule $\mathsf{id}_\mathsf{P}^{\mathsf{rr}}$ with the conclusion $\widetilde{\Gamma}_3\{R^{\Re}, \neg R^{\Re}\}$ where $\widetilde{\Gamma}_3 = [\neg P^{\Re}, \{\ \}, \neg\Box Q^\ell]$, is assigned an interpolant $\widetilde{\Gamma}_3^\circ\{\top\} = [\top]$, where $\widetilde{\Gamma}_3^\circ = [\{\ \}]$.

There are two α rules in the derivation. The first, α^{r}, requires taking the external conjunction of $[\Box Q]$ and $[\top]$. The second rule, α^{l}, produces the external disjunction of $[\neg P]$ with the external conjunction $[\Box Q] \oslash [\top]$. In both cases, the result is clearly well formed. In Lemma 6.4(ii), we show that external conjunctions and disjunctions are always applicable to interpolants in the premises of α rules.

It remains to explain how the ν rules are applied (both are ν^{r} rules). First, in $(\star)$, from the interpolant $[[Q]]$ of $\widetilde{\Gamma}_4\{Q^{\Re}\} = [\neg P^{\Re}, R^{\Re}, \neg\Box Q^\ell, [Q^{\Re}]]$ with $\widetilde{\Gamma}_4\{\} = [\neg P^{\Re}, R^{\Re}, \neg\Box Q^\ell, \{\}]$, we construct an interpolant of $\widetilde{\Gamma}_4\{\Box Q^{\Re}\} = [\neg P^{\Re}, R^{\Re}, \neg\Box Q^\ell, \Box Q^{\Re}]$. The interpolant of the premise $[[Q]]$ is in the prescribed SCNF with its only disjunct of its only conjunct being $\Pi\{Q\}$ for $\Pi\{\} = [\{\ \}] \sim [\neg P^{\Re}, R^{\Re}, \neg\Box Q^\ell, \{\}] = \widetilde{\Gamma}_4\{\}$. To compute the interpolant $[\Box Q]$ of the conclusion, we take $\bigvee\{Q\} = Q$, prefix it with a $\Box$, and insert $\Box Q$ into $\widetilde{\Gamma}_4^\circ\{\} = [\{\ \}]$.

Finally, we discuss $(\star\star)$. The preceding rule $\mathsf{adtr}_{\mathsf{ff}}^{\mathsf{lr}}$ yields the interpolant $[\neg P] \oslash ([\Box Q] \oslash [\top])$ of the premise $\widetilde{\Gamma}_5\{[P \supset R^{\Re}]\} = \neg\Box(P \supset \Box Q)^\ell, \neg\Box(\Box Q \supset R)^{\Re}, [P \supset R^{\Re}]$ with $\widetilde{\Gamma}_5\{\} = \neg\Box(P \supset \Box Q)^\ell, \neg\Box(\Box Q \supset R)^{\Re}, \{\ \}$. To construct an interpolant of the conclusion $\widetilde{\Gamma}_5\{\Box(P \supset R)^{\Re}\} = \neg\Box(P \supset \Box Q)^\ell, \neg\Box(\Box Q \supset R)^{\Re}, \Box(P \supset R)^{\Re}$, we first need to convert $[\neg P] \oslash ([\Box Q] \oslash [\top])$ to a SCNF using the constructive method from Fact 5.7. The result, $([\neg P] \oslash [\Box Q]) \oslash ([\neg P] \oslash [\top]) = (\Pi_{1,1}\{[\neg P]\} \oslash \Pi_{1,2}\{[\Box Q]\}) \oslash (\Pi_{2,1}\{[\neg P]\} \oslash \Pi_{2,2}\{[\top]\})$ with $\Pi_{1,1}\{\} = \Pi_{1,2}\{\} = \Pi_{2,1}\{\} = \Pi_{2,2}\{\} = \{\} \sim \widetilde{\Gamma}_5\{\}$, is also an interpolant of the premise by Lemma 5.8. Since $\widetilde{\Gamma}_5^\circ = \{\ \}$, the interpolant of the conclusion is $\widetilde{\Gamma}_5^\circ\{\Box(\neg P \vee \Box Q)\} \oslash \widetilde{\Gamma}_5^\circ\{\Box(\neg P \vee \top)\} = \Box(\neg P \vee \Box Q) \oslash \Box(\neg P \vee \top)$.

Figure 7: Application of Algorithm 6.1 to the BNK-derivation of $\neg\Box(P \supset \Box Q)^\ell, \Box(\Box Q \supset R) \supset \Box(P \supset R)^{\Re}$ from Figure 5.

Let us now prove that the interpolation algorithm works for all biased derivations. We start with the simplest of the three lemmas that states that the propositional variables of the interpolant produced by the

algorithm are always common to the left and right sides of the biased sequent being interpolated.

Lemma 6.3. *For every rule in Figure 6, if the propositional variables of the given interpolant in each premise are common to the left and right sides of the biased sequent in this premise, then the propositional variables of the newly constructed interpolant for the conclusion are common to the left and right sides of the biased sequent in the conclusion.*

PROOF. We leave this as an exercise to be checked by the reader with the help of the fact that

$$\mathrm{Prop}\left(L(\widetilde{\Gamma}\{\widetilde{\Pi}\})\right) = \mathrm{Prop}(L\widetilde{\Gamma}\{\ \}) \cup \mathrm{Prop}(L\widetilde{\Pi}) \qquad \text{and} \qquad \mathrm{Prop}\left(R(\widetilde{\Gamma}\{\widetilde{\Pi}\})\right) = \mathrm{Prop}(R\widetilde{\Gamma}\{\ \}) \cup \mathrm{Prop}(R\widetilde{\Pi}) \ . \ \square$$

In the second lemma, we demonstrate all the necessary properties of the structure of interpolants:

Lemma 6.4.   (i) *To apply the algorithm to the rule* $\mathsf{adtr}_{[][]}$, *rule* ν^l, *or rule* ν^r, *a given interpolant for the premise of the rule must be in the required form (see Figure 6). If the given interpolant is not in this form, it can be efficiently converted to a decoratively equivalent generalized sequent that is.*

   (ii) *The object suggested by the algorithm as an interpolant for the conclusion of all zero-premise rules, as well as for the rule* α^l, *rule* α^r, *rule* $\mathsf{adtr}_{[][]}$, *rule* ν^l, *and rule* ν^r, *is always a well-formed generalized sequent (provided, for the rule* $\mathsf{adtr}_{[][]}$, *rule* ν^l, *and rule* ν^r, *that a given interpolant for the premise of the rule is in the required form).*[7]

   (iii) *Each generalized sequent suggested by the algorithm for the conclusion of any rule is structurally equivalent to the biased sequent from this conclusion.*

PROOF. For each rule, we need to prove all statements applicable to it. For brevity, we say "interpolant" instead of "suggested interpolant."

(ii) **and** (iii) **for** $\mathsf{id}_\mathsf{P}^\mathsf{lr}$, $\mathsf{id}_\mathsf{P}^\mathsf{rl}$, $\mathsf{id}_\mathsf{P}^\mathsf{ll}$, $\mathsf{id}_\mathsf{P}^\mathsf{rr}$, $\mathsf{id}_\top^\mathsf{l}$, $\mathsf{id}_\top^\mathsf{r}$, $\mathsf{id}_{\neg\perp}^\mathsf{l}$, **and** $\mathsf{id}_{\neg\perp}^\mathsf{r}$. The details are left to the reader.[8]

(iii) **for** $\neg\neg^\mathsf{l}$, $\neg\neg^\mathsf{r}$, ctr^l, ctr^r, β^l, β^r, $\mathsf{adtr}_\mathsf{ff}^\mathsf{rl}$, $\mathsf{adtr}_\mathsf{ff}^\mathsf{rl}$, $\mathsf{adtr}_\mathsf{ff}^\mathsf{ll}$, $\mathsf{adtr}_\mathsf{ff}^\mathsf{rr}$, $\mathsf{adtr}_\mathsf{f[]}^\mathsf{l}$, $\mathsf{adtr}_\mathsf{f[]}^\mathsf{r}$, $\mathsf{adtr}_\mathsf{[]f}^\mathsf{l}$, $\mathsf{adtr}_\mathsf{[]f}^\mathsf{r}$, π^l, **and** π^r. For each rule, given that the interpolant remains unchanged, to show the structural equivalence of a premise $\widetilde{\Gamma}\{\widetilde{\Pi}\}$ and a conclusion $\widetilde{\Gamma}\{\widetilde{\Lambda}\}$ of each rule, it is sufficient to show that $\widetilde{\Pi} \sim \widetilde{\Lambda}$. The details are left to the reader.

(ii) **and** (iii) **for** α^l **and** α^r. For α^l, since $\widetilde{\Gamma}\{\alpha_1^\ell\} \sim \widetilde{\Gamma}\{\alpha_2^\ell\} \sim \widetilde{\Gamma}\{\alpha^\ell\}$, it follows that $\mho_1 \sim \mho_2$. Thus, $\mho_1 \otimes \mho_2$ is defined. Further, $\mho_1 \otimes \mho_2 \sim \mho_1 \sim \widetilde{\Gamma}\{\alpha_1^\ell\} \sim \widetilde{\Gamma}\{\alpha^\ell\}$. The case of α^r is analogous.

(i), (ii), **and** (iii) **for** $\mathsf{adtr}_{[][]}$. Let us start with transforming a given interpolant $\mho$ for a premise $\widetilde{\Gamma}\{[\widetilde{\Delta}], [\widetilde{\Sigma}]\}$ into the required form. By Fact 5.7 and Lemma 5.8, another interpolant $\mho'$ of $\widetilde{\Gamma}\{[\widetilde{\Delta}], [\widetilde{\Sigma}]\}$ can be constructed that is in a SCNF. Further, every member sequent Ω of the interpolant $\mho'$ is structurally equivalent to $\widetilde{\Gamma}\{[\widetilde{\Delta}], [\widetilde{\Sigma}]\}$, meaning that $\Omega^\circ = \widetilde{\Gamma}^\circ\{[\widetilde{\Delta}^\circ], [\widetilde{\Sigma}^\circ]\}$. Thus, by Lemma 5.9.2, $\Omega = \Lambda\{[\Pi], \mathbf{A}, [\Theta]\}$ where $\Lambda^\circ\{\ \} = \widetilde{\Gamma}^\circ\{\ \}$, $\Pi^\circ = \widetilde{\Delta}^\circ$, $\Theta^\circ = \widetilde{\Sigma}^\circ$, and $\mathbf{A}$ is a shallow sequent[9]. It follows that $\Lambda\{\ \} \sim \widetilde{\Gamma}\{\ \}$, which completes the proof of (i). Each member sequent $\Lambda_{ij}\{[\Sigma_{ij}], \mathbf{A}_{ij}, [\Delta_{ij}]\}$ of the interpolant for the conclusion corresponds to the member sequent $\Lambda_{ij}\{[\Delta_{ij}], \mathbf{A}_{ij}, [\Sigma_{ij}]\}$ of the given interpolant of $\widetilde{\Gamma}\{[\widetilde{\Delta}], [\widetilde{\Sigma}]\}$, where $\Lambda_{ij}\{\ \} \sim \widetilde{\Gamma}\{\ \}$. Since $\Lambda_{ij}\{[\Delta_{ij}], \mathbf{A}_{ij}, [\Sigma_{ij}]\} \sim \widetilde{\Gamma}\{[\widetilde{\Delta}], [\widetilde{\Sigma}]\}$ by assumption, $[\Delta_{ij}], \mathbf{A}_{ij}, [\Sigma_{ij}]$ must be structurally equivalent to $[\widetilde{\Delta}], [\widetilde{\Sigma}]$, meaning that $\Delta_{ij} \sim \widetilde{\Delta}$ and $\Sigma_{ij} \sim \widetilde{\Sigma}$. It immediately follows that $\Lambda_{ij}\{[\Sigma_{ij}], \mathbf{A}_{ij}, [\Delta_{ij}]\} \sim \widetilde{\Gamma}\{[\widetilde{\Sigma}], [\widetilde{\Delta}]\}$. This completes the proof of (ii) and (iii).[10]

(i), (ii), **and** (iii) **for** ν^l **and** ν^r. We consider ν^l in detail, leaving ν^r to the reader. We start with transforming a given interpolant $\mho$ for a premise $\widetilde{\Gamma}\{[\nu_0^\ell]\}$ into the required form. By Fact 5.7 and Lemma 5.8, another

---

[7]For the remaining rules, the suggested interpolant is well-formed because it is the same as for the premise.

[8]It is worth noting that all such steps produce singleton sequents, which are both in a SDNF and in a SCNF. Strictly speaking, this observation is not needed to prove the lemma but is useful for implementation.

[9]$\mathbf{A}$ either is the empty sequent or consists of exactly one formula because $\mho'$ is in a singleton CNF.

[10]A note for the implementation: this interpolant for the conclusion is in a SCNF whenever the given interpolant for the premise is, because $\Lambda_{ij}\{[\Sigma_{ij}], \mathbf{A}_{ij}, [\Delta_{ij}]\}$ is a singleton sequent whenever $\Lambda_{ij}\{[\Delta_{ij}], \mathbf{A}_{ij}, [\Sigma_{ij}]\}$ is.

interpolant $\mho'$ of $\widetilde{\Gamma}\{[\nu_0^\ell]\}$ can be constructed that is in a SDNF. Every member sequent Ω of the interpolant $\mho'$ is structurally equivalent to $\widetilde{\Gamma}\{[\nu_0^\ell]\}$, meaning that $\Omega^\circ = \widetilde{\Gamma}^\circ\{[\varepsilon]\}$. Thus, by Lemma 5.9.1, $\Omega = \Sigma\{[\Theta]\}$ where $\Sigma^\circ\{\,\} = \widetilde{\Gamma}^\circ\{\,\}$ and $\Theta^\circ = \varepsilon$, making Θ a shallow sequent. It follows that $\Sigma\{\,\} \sim \widetilde{\Gamma}\{\,\}$. Given that Ω is a singleton sequent, Θ either is ε or consists of a single formula A. Using Fact 5.2 and Lemma 5.8, we construct an interpolant $\mho''$ of $\widetilde{\Gamma}\{[\nu_0^\ell]\}$ by moving all member sequents $\Omega = \Sigma\{[\varepsilon]\}$ to the end of each disjunct of $\mho'$ in a SDNF. This completes the proof of item (i) for ν^l. Since each member sequent of the interpolant for the conclusion is either $\Lambda_{ij}\{\varepsilon\}$ with $\Lambda_{ij}\{\,\} \sim \widetilde{\Gamma}\{\,\}$ or $\widetilde{\Gamma}^\circ\{B\}$, it is always structurally equivalent to $\widetilde{\Gamma}\{\nu^\ell\}$. This completes the proof of (ii) and (iii) for ν^l.[11] $\hfill\square$

It only remains to show that the well-formed generalized sequents proposed by the algorithm satisfy the appropriate decorative consequences. The algorithm often prescribes that interpolants should be kept unchanged. In all such cases, the proof that the interpolant continues to satisfy the decorative consequences is based on the same idea, which we formulate as a separate lemma:

Lemma 6.5. $L\widetilde{\Sigma} \Vdash L\widetilde{\Pi} \implies L(\widetilde{\Gamma}\{\widetilde{\Sigma}\}) \Vdash L(\widetilde{\Gamma}\{\widetilde{\Pi}\}).$ $R\widetilde{\Sigma} \Vdash R\widetilde{\Pi} \implies R(\widetilde{\Gamma}\{\widetilde{\Sigma}\}) \Vdash R(\widetilde{\Gamma}\{\widetilde{\Pi}\}).$

PROOF. We show the statement for the left biases, leaving the other case to the reader. First , $\widetilde{\Sigma} \sim L\widetilde{\Sigma}$ and $\widetilde{\Pi} \sim L\widetilde{\Pi}$. Further, $L\widetilde{\Sigma} \Vdash L\widetilde{\Pi}$ implies that $L\widetilde{\Sigma} \sim L\widetilde{\Pi}$. By transitivity, we get $\widetilde{\Pi} \sim \widetilde{\Sigma}$ and, consequently, $\widetilde{\Gamma}\{\widetilde{\Pi}\} \sim \widetilde{\Gamma}\{\widetilde{\Sigma}\}$. Finally, we conclude that $L(\widetilde{\Gamma}\{\widetilde{\Pi}\}) \sim L(\widetilde{\Gamma}\{\widetilde{\Sigma}\})$.

Consider arbitrary matching $\mathcal{M}$-decorations Θ^* of $L(\widetilde{\Gamma}\{\widetilde{\Sigma}\})$ and Ξ^* of $L(\widetilde{\Gamma}\{\widetilde{\Pi}\})$. By Fact 5.12, we have $\Theta^* = L\widetilde{\Gamma}^*\{L\widetilde{\Sigma}^*\}_w$ and $\Xi^* = L\widetilde{\Gamma}^*\{L\widetilde{\Pi}^*\}_w$ for some $\mathcal{M}$-decorations $L\widetilde{\Gamma}^*\{\,\}_w$, $L\widetilde{\Sigma}^*$, and $L\widetilde{\Pi}^*$ with $r(L\widetilde{\Sigma}^*) = r(L\widetilde{\Pi}^*) = w$. It is clear that such a decomposition of the matching decorations Θ^* and Ξ^* produces the same decoration of $L\widetilde{\Gamma}\{\,\}$ and matching decorations $L\widetilde{\Sigma}^*$ and $L\widetilde{\Pi}^*$.

Assume $\Vdash \Theta^*$. By Fact 5.11, either $\Vdash L\widetilde{\Gamma}^*\{w\}_w$ or $\Vdash L\widetilde{\Sigma}^*$. Since by the assumption $L\widetilde{\Sigma} \Vdash L\widetilde{\Pi}$ and since $L\widetilde{\Pi}^*$ matches $L\widetilde{\Sigma}^*$, either $\Vdash L\widetilde{\Gamma}^*\{w\}_w$ or $\Vdash L\widetilde{\Pi}^*$. Thus, $\Vdash \Xi^*$ by Fact 5.11. $\hfill\square$

Corollary 6.6. *Let* $L\widetilde{\Sigma} \Vdash L\widetilde{\Pi}$ *and* $R\widetilde{\Sigma} \Vdash R\widetilde{\Pi}$. *If* $\widetilde{\Gamma}\{\widetilde{\Sigma}\} \longleftarrow \mho$, *then* $\neg\mho \Vdash L(\widetilde{\Gamma}\{\widetilde{\Pi}\})$ *and* $\mho \Vdash R(\widetilde{\Gamma}\{\widetilde{\Pi}\})$.

The following two lemmas are to be used for the rules α^l and α^r and for the rule $\mathsf{adtr}_{[][]}$ respectively:

Lemma 6.7. $\Vdash \Delta^*\{w, \alpha_1\}_w$ *and* $\Vdash \Delta^*\{w, \alpha_2\}_w$ $\iff$ $\Vdash \Delta^*\{w, \alpha\}_w$.

Lemma 6.8. $\Vdash \Lambda^*\{w, [\Delta^*], \mathbf{A}, [\Sigma^*]\}_w$ $\iff$ $\Vdash \Lambda^*\{w, [\Sigma^*], \mathbf{A}, [\Delta^*]\}_w$.[12]

PROOF. Since the world assigned to each bracket is moved along with the bracket, each formula is evaluated at the same world as before and the truth of the decoration is not affected. The rest is left to the reader. $\hfill\square$

Lemma 6.9. *Given arbitrary interpolant(s) of the premise(s) of any rule from Figure 6 in the form prescribed for it/them for this rule, for the generalized sequent $\mho$ suggested by the algorithm as an interpolant for the conclusion $\widetilde{\Delta}$ of the rule, $\neg\mho \Vdash L\widetilde{\Delta}$ and $\mho \Vdash R\widetilde{\Delta}$.*

PROOF. Note that by Lemma 6.4, $\mho$ is a well-defined generalized sequent structurally equivalent to $\widetilde{\Delta}$ and, hence, to both $L\widetilde{\Delta}$ and $R\widetilde{\Delta}$.

Cases $\mathsf{id}_\mathsf{P}^\mathsf{lr}$ and $\mathsf{id}_\mathsf{P}^\mathsf{rl}$ are similar. For $\mathsf{id}_\mathsf{P}^\mathsf{lr}$, we need to show that $\neg\widetilde{\Gamma}^\circ\{\neg P\} \Vdash L\widetilde{\Gamma}\{P\}$ and $\widetilde{\Gamma}^\circ\{\neg P\} \Vdash R\widetilde{\Gamma}\{\neg P\}$. We only show the first of these decorative consequences, leaving the second and the case of $\mathsf{id}_\mathsf{P}^\mathsf{rl}$ to the reader. Consider arbitrary matching $\mathcal{M}$-decorations $(\widetilde{\Gamma}^\circ)^*\{w, \neg P\}_w$ of $\widetilde{\Gamma}^\circ\{\neg P\}$ and $L\widetilde{\Gamma}^\star\{w, P\}_w$ of $L\widetilde{\Gamma}\{P\}$. Assume now that $\nVdash (\widetilde{\Gamma}^\circ)^*\{w, \neg P\}_w$. By Fact 5.11, we can see that $\nVdash w, \neg P$: i.e, $\mathcal{M}, w \nVdash \neg P$. Thus, $\mathcal{M}, w \Vdash P$: i.e., $\Vdash w, P$. Now $\Vdash L\widetilde{\Gamma}^\star\{w, P\}_w$ follows from Fact 5.11.

---

[11] A note for the implementation: this generalized sequent is in the same form (SDNF or SCNF) as the given interpolant because $\Lambda_{ij}\{\varepsilon\}$ is a singleton sequent whenever $\Lambda_{ij}\{[\varepsilon]\}$ is and because inserting the formula B into the void context $\widetilde{\Gamma}^\circ\{\,\}$ always produces a singleton sequent.

[12] These decorated sequents are not, in general, decoratively equivalent because they need not be structurally equivalent.

Cases $\mathsf{id}_\mathsf{P}^{\mathsf{ll}}$, $\mathsf{id}_\top^{\mathsf{l}}$, $\mathsf{id}_{\neg\bot}^{\mathsf{l}}$, $\mathsf{id}_\mathsf{P}^{\mathsf{rr}}$, $\mathsf{id}_\top^{\mathsf{r}}$, **and** $\mathsf{id}_{\neg\bot}^{\mathsf{r}}$ are even simpler and are left to the reader.

Cases $\neg\neg^{\mathsf{l}}$ **and** $\neg\neg^{\mathsf{r}}$. By Corollary 6.6, it is sufficient to note that $A \Vdash \neg\neg A$ (and $\varepsilon \Vdash \varepsilon$).

Cases β^{l} **and** β^{r}. By Corollary 6.6, it is sufficient to note that $\beta_1, \beta_2 \Vdash \beta$.

Cases $\mathsf{ctr}^{\mathsf{l}}$ **and** $\mathsf{ctr}^{\mathsf{r}}$. By Corollary 6.6, it is sufficient to note that $A, A \Vdash A$.

Cases $\mathsf{adtr}_{\mathsf{ff}}^{\mathsf{ll}}$ **and** $\mathsf{adtr}_{\mathsf{ff}}^{\mathsf{rr}}$. By Corollary 6.6, it is sufficient to note that $A, B \Vdash B, A$.

Cases $\mathsf{adtr}_{\mathsf{ff}}^{\mathsf{lr}}$ **and** $\mathsf{adtr}_{\mathsf{ff}}^{\mathsf{rl}}$. By Corollary 6.6, it is sufficient to note that $A \Vdash A$ and $B \Vdash B$.

Cases $\mathsf{adtr}_{\mathsf{f}[]}^{\mathsf{l}}$, $\mathsf{adtr}_{\mathsf{f}[]}^{\mathsf{r}}$, $\mathsf{adtr}_{[]\mathsf{f}}^{\mathsf{l}}$, **and** $\mathsf{adtr}_{[]\mathsf{f}}^{\mathsf{r}}$. By Corollary 6.6, it is sufficient to note that $A, [\Delta] \dashv\vdash [\Delta], A$ and $\Lambda \Vdash \Lambda$, where $\{\Delta, \Lambda\} = \{L\widetilde{\Sigma}, R\widetilde{\Sigma}\}$, with the exact choice depending on the rule.

Cases π^{l} **and** π^{r}. By Corollary 6.6, it is sufficient to show that $[\Delta, \pi_0] \Vdash [\Delta], \pi$ and $\Lambda \Vdash \Lambda$, where $\{\Delta, \Lambda\} = \{L\widetilde{\Sigma}, R\widetilde{\Sigma}\}$, with the exact choice depending on the rule. The latter consequence is trivial. To show the former, let $\mathcal{M} = (W, R, V)$. Consider arbitrary matching $\mathcal{M}$-decorations $w, [\Delta^*, \pi_0]$ of $[\Delta, \pi_0]$ and $w, [\Delta^*], \pi$ of $[\Delta], \pi$. For $v = r(\Delta^*)$ we have wRv. Thus,

$$\Vdash w, [\Delta^*, \pi_0] \quad\Longrightarrow\quad \Vdash \Delta^* \text{ or } \mathcal{M}, v \Vdash \pi_0 \quad\Longrightarrow\quad \Vdash \Delta^* \text{ or } \mathcal{M}, w \Vdash \pi \quad\Longrightarrow\quad \Vdash w, [\Delta^*], \pi \ .$$

Cases α^{l} **and** α^{r}. We show the former, leaving the latter to the reader. Let $\widetilde{\Gamma}\{\alpha_i^\ell\} \longleftarrow \mho_i$ for each $i = 1, 2$. We need to show that

$$\neg(\mho_1 \oslash \mho_2) \Vdash L\widetilde{\Gamma}\{\alpha\} \qquad \text{and} \qquad \mho_1 \oslash \mho_2 \Vdash R\widetilde{\Gamma}\{\varepsilon\} \ .$$

We start with the first decorative consequence. Consider arbitrary matching $\mathcal{M}$-decorations $\mho_1^* \oslash \mho_2^*$ of $\mho_1 \oslash \mho_2$ and $L\widetilde{\Gamma}^*\{w, \alpha\}_w$ of $L\widetilde{\Gamma}\{\alpha\}$ (see Fact 5.12), where $\mho_1^*$ and $\mho_2^*$ match $L\widetilde{\Gamma}^*\{w, \alpha\}_w$. Assume that $\nVdash \mho_1^* \oslash \mho_2^*$. Then $\nVdash \mho_i^*$ for each $i = 1, 2$. Now $\Vdash L\widetilde{\Gamma}^*\{w, \alpha_i\}_w$ follows from $\neg\mho_i \Vdash L\widetilde{\Gamma}\{\alpha_i\}$ for each $i = 1, 2$. Finally, $\Vdash L\widetilde{\Gamma}^*\{w, \alpha\}_w$ follows by Lemma 6.7.

For the second consequence, consider arbitrary matching $\mathcal{M}$-decorations $\mho_1^* \oslash \mho_2^*$ of $\mho_1 \oslash \mho_2$ and $R\widetilde{\Gamma}^*\{w\}_w$ of $R\widetilde{\Gamma}\{\varepsilon\}$, where $\mho_1^*$ and $\mho_2^*$ match $R\widetilde{\Gamma}^*\{w\}_w$. Assume that $\Vdash \mho_1^* \oslash \mho_2^*$. Then $\Vdash \mho_i^*$ for some $i = 1, 2$. Now $\Vdash R\widetilde{\Gamma}^*\{w\}_w$ follows from $\mho_i \Vdash R\widetilde{\Gamma}\{\varepsilon\}$.

Case $\mathsf{adtr}_{[][]}$. Let $\widetilde{\Gamma}\{[\widetilde{\Delta}], [\widetilde{\Sigma}]\} \longleftarrow \bigotimes_{i=1}^{n} \bigoslash_{j=1}^{m_i} \Lambda_{ij}\{[\Delta_{ij}], \mathbf{A}_{ij}, [\Sigma_{ij}]\}$, where $\Lambda_{ij}\{\ \} \sim \widetilde{\Gamma}\{\ \}$ and $\mathbf{A}_{ij}$ is a shallow sequent for each $1 \leq i \leq n$ and each $1 \leq j \leq m_i$. The following two statements need to be demonstrated:

$$\neg\bigotimes_{i=1}^{n} \bigoslash_{j=1}^{m_i} \Lambda_{ij}\{[\Sigma_{ij}], \mathbf{A}_{ij}, [\Delta_{ij}]\} \Vdash L\widetilde{\Gamma}\{[L\widetilde{\Sigma}], [L\widetilde{\Delta}]\} \text{ and } \bigotimes_{i=1}^{n} \bigoslash_{j=1}^{m_i} \Lambda_{ij}\{[\Sigma_{ij}], \mathbf{A}_{ij}, [\Delta_{ij}]\} \Vdash R\widetilde{\Gamma}\{[R\widetilde{\Sigma}], [R\widetilde{\Delta}]\} \ .$$

The proofs use Lemma 6.8. The details are left to the reader.

Cases ν^{l} **and** ν^{r}. These are the most crucial cases because they require removing a structural box from the interpolant's structure, a non-trivial modification. We only show the former case, leaving the latter to the reader. Let $\widetilde{\Gamma}\{[\nu_0^\ell]\} \longleftarrow \bigoslash_{i=1}^{n}\left(\bigotimes_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \oslash \bigoslash_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\}\right)$, where the generalized sequent is in a SDNF and $\widetilde{\Gamma}\{\ \} \sim \Lambda_{ij}\{\ \} \sim \Pi_{ik}\{\ \}$ for each $1 \leq i \leq n$, each $1 \leq j \leq m_i$, and each $1 \leq k \leq l_i$. Each context $\Pi_{ik}\{\ \}$ is void because $\Pi_{ik}\{[A_{ik}]\}$ is a singleton sequent. Thus, $\Pi_{ik}\{\ \} = \widetilde{\Gamma}^\circ\{\ \}$. We need to show that

$$\neg\bigoslash_{i=1}^{n}\left(\widetilde{\Gamma}^\circ\{\Diamond \bigwedge_{k=1}^{l_i} A_{ik}\} \oslash \bigoslash_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\}\right) \Vdash L\widetilde{\Gamma}\{\nu\} \quad\text{and}\quad \bigoslash_{i=1}^{n}\left(\widetilde{\Gamma}^\circ\{\Diamond \bigwedge_{k=1}^{l_i} A_{ik}\} \oslash \bigoslash_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\}\right) \Vdash R\widetilde{\Gamma}\{\varepsilon\} \ .$$

We start with showing the contraposition of the first consequence. For $\mathcal{M} = (W, R, V)$, consider arbitrary matching $\mathcal{M}$-decorations of

$$\bigoslash_{i=1}^{n}\left(\widetilde{\Gamma}^\circ\{\Diamond \bigwedge_{k=1}^{l_i} A_{ik}\} \oslash \bigoslash_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\}\right) \qquad \text{and} \qquad L\widetilde{\Gamma}\{\nu\} \ .$$

By Fact 5.12 and considering that $\Lambda_{ij}\{\ \} \sim \widetilde{\Gamma}\{\ \} \sim L\widetilde{\Gamma}\{\ \} \sim \widetilde{\Gamma}^\circ\{\ \}$, they must have the form

$$\bigotimes_{i=1}^{n} \left((\widetilde{\Gamma}^\circ)^*\{w, \Diamond \bigwedge_{k=1}^{l_i} A_{ik}\}_w \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}^*\{w\}_w \right) \qquad \text{and} \qquad L\widetilde{\Gamma}^*\{w, \nu\}_w$$

respectively. Assume that $\nVdash L\widetilde{\Gamma}^*\{w, \nu\}_w$. By Fact 5.11, $\nVdash L\widetilde{\Gamma}^*\{w\}_w$ and $\nVdash w, \nu$. The latter means that $\mathcal{M}, w \nVdash \nu$. Then there exists $v \in W$ such that wRv and $\mathcal{M}, v \nVdash \nu_0$. Thus, $\nVdash w, [v, \nu_0]$, and it follows by Fact 5.11 that $\nVdash L\widetilde{\Gamma}^*\{w, [v, \nu_0]\}_w$ for a decoration of $L\widetilde{\Gamma}\{[\nu_0]\}$. It is easy to see that for $\Pi_{ik}^*\{\ \}_w := (\widetilde{\Gamma}^\circ)^*\{\ \}_w$ for each $1 \le i \le n$ and each $1 \le k \le l_i$, this last decoration matches the decoration $\bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}^*\{w, [v, A_{ik}]\}_w \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}^*\{w, [v]\}_w \right)$ of $\bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right)$. We get

$$\Vdash \bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}^*\{w, [v, A_{ik}]\}_w \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}^*\{w, [v]\}_w \right) \text{ from } \neg \bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right) \Vdash L\widetilde{\Gamma}\{[\nu_0]\}.$$

Thus, there must exist $1 \le L \le n$ such that $\Vdash \bigotimes_{k=1}^{l_L} \Pi_{Lk}^*\{w, [v, A_{Lk}]\}_w \oslash \bigotimes_{j=1}^{m_L} \Lambda_{Lj}^*\{w, [v]\}_w$. Given that $\Pi_{Lk}^*\{w\}_w$ and $w, [v]$ are void, using Fact 5.11, we can equivalently say that $\mathcal{M}, v \Vdash A_{Lk}$ for each $1 \le k \le l_L$ and $\Vdash \Lambda_{Lj}^*\{w\}_w$ for each $1 \le j \le m_L$. It follows that $\mathcal{M}, v \Vdash \bigwedge_{k=1}^{l_L} A_{Lk}$. Since wRv, it follows that $\mathcal{M}, w \Vdash \Diamond \bigwedge_{k=1}^{l_L} A_{Lk}$, which, by Fact 5.11, is sufficient to conclude that $\Vdash (\widetilde{\Gamma}^\circ)^*\{w, \Diamond \bigwedge_{k=1}^{l_L} A_{Lk}\}_w$. It follows that $\Vdash (\widetilde{\Gamma}^\circ)^*\{w, \Diamond \bigwedge_{k=1}^{l_L} A_{Lk}\}_w \oslash \bigotimes_{j=1}^{m_L} \Lambda_{Lj}^*\{w\}_w$, and, finally, that $\Vdash \bigotimes_{i=1}^{n} \left((\widetilde{\Gamma}^\circ)^*\{w, \Diamond \bigwedge_{k=1}^{l_i} A_{ik}\}_w \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}^*\{w\}_w \right)$.

For the second consequence, we omit some details. For $\mathcal{M} = (W, R, V)$, consider arbitrary matching $\mathcal{M}$-decorations $\bigotimes_{i=1}^{n} \left((\widetilde{\Gamma}^\circ)^*\{w, \Diamond \bigwedge_{k=1}^{l_i} A_{ik}\}_w \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}^*\{w\}_w \right)$ of $\bigotimes_{i=1}^{n} \left((\widetilde{\Gamma}^\circ)^*\{\Diamond \bigwedge_{k=1}^{l_i} A_{ik}\} \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}^*\{\varepsilon\} \right)$ and $R\widetilde{\Gamma}^*\{w\}_w$ of $R\widetilde{\Gamma}\{\varepsilon\}$. Assume that $\Vdash \bigotimes_{i=1}^{n} \left((\widetilde{\Gamma}^\circ)^*\{w, \Diamond \bigwedge_{k=1}^{l_i} A_{ik}\}_w \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}^*\{w\}_w \right)$. Then there must exist $1 \le L \le n$ such that $\Vdash \Lambda_{Lj}^*\{w\}_w$ for each $1 \le j \le m_L$ and $\Vdash (\widetilde{\Gamma}^\circ)^*\{w, \Diamond \bigwedge_{k=1}^{l_L} A_{Lk}\}_w$. Given that $(\widetilde{\Gamma}^\circ)^*\{\ \}_w$ is void, it follows that $\mathcal{M}, w \Vdash \Diamond \bigwedge_{k=1}^{l_L} A_{Lk}$. Then there exists $v \in W$ such that wRv and $\mathcal{M}, v \Vdash \bigwedge_{k=1}^{l_L} A_{Lk}$. In particular, $\mathcal{M}, v \Vdash A_{Lk}$ for each $1 \le k \le l_L$. It follows that $\Vdash \Pi_{Lk}^*\{w, [v, A_{Lk}]\}_w$ for each $1 \le k \le l_L$ where $\Pi_{ik}^*\{\ \}_w := (\widetilde{\Gamma}^\circ)^*\{\ \}_w$ for each $1 \le i \le n$ and each $1 \le k \le l_i$. Further, $\Vdash \Lambda_{Lj}^*\{w\}_w$ clearly implies $\Vdash \Lambda_{Lj}^*\{w, [v]\}_w$ for each $1 \le j \le m_L$. Overall, we conclude that $\Vdash \bigotimes_{k=1}^{l_L} \Pi_{Lk}^*\{w, [v, A_{Lk}]\}_w \oslash \bigotimes_{j=1}^{m_L} \Lambda_{Lj}^*\{w, [v]\}_w$, and, hence, $\Vdash \bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}^*\{w, [v, A_{ik}]\}_w \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}^*\{w, [v]\}_w \right)$. This last decoration of $\bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right)$ matches the decoration $R\widetilde{\Gamma}^*\{w, [v]\}_w$ of $R\widetilde{\Gamma}\{[\varepsilon]\}$. Now $\Vdash R\widetilde{\Gamma}^*\{w, [v]\}_w$ follows from $\bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \oslash \bigotimes_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right) \Vdash R\widetilde{\Gamma}\{[\varepsilon]\}$. Since $w, [v]$ is void, we conclude that $\Vdash R\widetilde{\Gamma}^*\{w\}_w$. $\qquad\square$

Theorem 6.10 (Interpolation theorem for K). *For any biased sequent $\widetilde{\Gamma}$, derivable in* BNK*, Algorithm 6.1 finds an interpolant $\Im$ of $\widetilde{\Gamma}$.*

PROOF. Follows from Lemmas 6.3, 6.4, and 6.9. □

Corollary 6.11 (Interpolation Theorem). *The modal logic* K *has the CIP.*

PROOF. Let K ⊢ $A \supset B$. Then by completeness of NK, clearly NK ⊢ $\neg A \vee B$, and also NK ⊢ $\neg A, B$. Thus, BNK ⊢ $\neg A^\ell, B^\Re$ by Theorem 4.2. By Theorem 6.10, $\neg A^\ell, B^\Re \longleftarrow \underline{\mho}$ for some interpolant $\underline{\mho}$. By Corollary 3.33, the formula $\underline{\mho}$ contains only propositional variables common to $\neg A$ and B, equivalently common to A and B and, in addition, K ⊢ $\neg\neg A \supset \underline{\mho}$ and K ⊢ $\underline{\mho} \supset B$. Thus, K ⊢ $A \supset \underline{\mho}$ and $\underline{\mho}$ is an interpolant of A and B.

7. Dealing with Other Modal Logics from the "Modal Cube"

In this section, we extend our methods of proving the CIP to all the logics from the so-called modal cube. Given the detailed presentation of the method for the logic K, we only outline the necessary changes while omitting most of the details and proofs.

Definition 7.1 (Modal cube). Identifying each logic with its set of theorems, we define the modal logics of the *modal cube* to be extensions of K with any combination of the following axioms:

$$\mathsf{d}:\ \ \Box\bot \supset \bot\ , \qquad \mathsf{t}:\ \ \Box A \supset A\ , \qquad \mathsf{b}:\ \ \Diamond\Box A \supset A\ ,$$
$$\mathsf{4}:\ \ \Box A \supset \Box\Box A\ , \qquad \mathsf{5}:\ \ \neg\Box A \supset \Box\neg\Box A\ .$$

The modal cube consists of 15 logics depicted in Figure 8. The names of the logics are traditional (according to one of the multiple existing traditions). We do not explain the naming scheme here in detail, referring the reader instead to the article "Modal Logic" in Stanford Encyclopedia of Philosophy [7, Sect. 8]. The general idea of (most of) the names is that D in the name of the logic means that d is an axiom of the logic, etc. An edge joining two logics in Figure 8 means that the logic to the right or above (or both) extends the logic to the left or below (or both). Given that there are 32 ways to extend K with a subset of the 5 axioms stated in Definition 7.1 but that there are only 15 logics in Figure 8, it follows that some logics in the modal cube have alternative axiomatizations. Not all such axiomatizations have straightforward translations into nested sequent systems that we are going to describe next. However, we are primarily interested in whether a given logic has the CIP rather than in the fine details of which axiomatization of the logic is better suited for proving it has. Thus, we simply work with maximal axiomatizations of each logic.
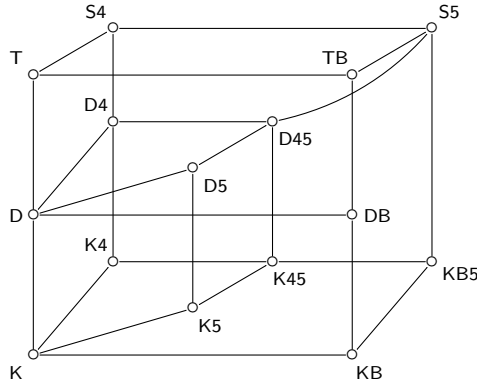

Figure 8: The modal cube.

Definition 7.2 (Maximal axiomatization). The *maximal axiomatization* of a logic from the modal cube consists of all the axioms and inference rules of K (see Definition 2.1) and all the extending axioms from Definition 7.1 that are derivable in the logic, with the following exception: the axiom d is not part of the axiomatization whenever t, of which d is an instance, is derivable.

24

Definition 7.3 (Kripke models for the modal cube). Each axiom from Definition 7.1 corresponds to a restriction on the accessibility relation. For d, accessibility must be *serial*: i.e., for each world w, there exists a world v such that wRv. For t, accessibility must be *reflexive*. For b, accessibility must be *symmetric*. For 4, accessibility must be *transitive*. Finally, for 5, accessibility must be *Euclidean*: i.e., vRu whenever, for some world w, wRv and wRu. A Kripke model $\mathcal{M} = (W, R, V)$ is called *serial* (*reflexive, symmetric, transitive*, or *Euclidean*) if its accessibility relation R is. Let L be a logic from the modal cube. A Kripke model $\mathcal{M} = (W, R, V)$ is called an L-*model* if R satisfies all the requirements that correspond to the additional axioms in the maximal axiomatization of L.

Definition 7.4 (Nested calculi for the modal-cube logics). For each of the modal-cube logics, we define a nested sequent calculus as the extension of the calculus NK with those nested rules from Figure 9 that correspond to the axioms from the maximal axiomatization of the logic. For instance, the nested rule b is added to the nested calculus whenever the Hilbert axiom b is part of the maximal axiomatization of the logic. Note that the presence of the axiom 5 in the maximal axiomatization necessitates the addition of all three rules 5a, 5b, and 5c to the nested calculus. We denote the nested calculus for a logic L by prepending its name with N. For instance, the nested calculus for the logic D45 is called ND45.

$$
\text{d}\ \frac{\Gamma\{[\pi_0]\}}{\Gamma\{\pi\}}
\qquad
\text{t}\ \frac{\Gamma\{\pi_0\}}{\Gamma\{\pi\}}
\qquad
\text{b}\ \frac{\Gamma\{[\Sigma], \pi_0\}}{\Gamma\{[\Sigma], \pi]\}}
\qquad
\text{4}\ \frac{\Gamma\{[\Sigma], \pi\}}{\Gamma\{[\Sigma], \pi\}}
$$

$$
\text{5a}\ \frac{\Gamma\{[\Sigma], \pi\}}{\Gamma\{[\Sigma], \pi]\}}
\qquad
\text{5b}\ \frac{\Gamma\{[\Sigma], [\Pi, \pi]\}}{\Gamma\{[\Sigma], \pi], [\Pi]\}}
\qquad
\text{5c}\ \frac{\Gamma\{[\Sigma, [\Pi, \pi]]\}}{\Gamma\{[\Sigma, [\Pi], \pi]\}}
$$

Figure 9: Nested rules for logics built from axioms d, t, b, 4, and 5.

Theorem 7.5 (Completeness of the nested calculi for the modal-cube logics). *For any logic L from the modal cube, for any sequent Γ, we have $\text{NL} \vdash \Gamma$ iff $\text{L} \vdash \underline{\Gamma}$ iff $\Vdash_{\text{L}} \underline{\Gamma}$, where $\Vdash_{\text{L}}$ denotes validity for L-models.*

PROOF. It follows from the results in [1, 4, 8]. $\square$

It immediately follows from this completeness theorem and Theorem 3.26 that

Corollary 7.6 (Completeness with respect to decorations for the modal-cube logics). *Let L be a logic from the modal cube. A nested sequent is derivable in NL iff all its $\mathcal{M}$-decorations are true for all L-models $\mathcal{M}$.*

The decorative consequence is a logical consequence, i.e., is based on the underlying semantics. To define the decorative consequence and interpolants for a logic L from the modal cube, we restrict the class of Kripke models used in Definitions 3.20 and 3.23 to L-models, and use $\Vdash_{\text{L}}$ instead of $\Vdash$. We also write $\widetilde{\Delta} \xleftarrow{\text{L}} \mho$ to denote the fact that $\mho$ is an L-interpolant of $\widetilde{\Delta}$ rather than a K-interpolant we have been discussing so far.

Corollary 7.7. *For any logic L from the modal cube, let a generalized sequent $\mho$ be an L-interpolant of a shallow biased sequent $\widetilde{\Delta}$. Then*

$$(\text{B}')\ \text{L} \vdash \neg \underline{L\widetilde{\Delta}} \supset \underline{\mho}, \qquad (\text{C}')\ \text{L} \vdash \underline{\mho} \supset \underline{R\widetilde{\Delta}}, \qquad and \qquad (\text{D}')\ \text{Prop}(\underline{\mho}) \subseteq \text{Prop}\left(L\widetilde{\Delta}\right) \cap \text{Prop}\left(R\widetilde{\Delta}\right)\ .$$

Thus, for the split $L\widetilde{\Delta} \mid R\widetilde{\Delta}$ of the shallow sequent Δ, which corresponds to the biasing in $\widetilde{\Delta}$, a formula L-interpolant of the split can be obtained by taking the corresponding formula of the generalized-sequent L-interpolant $\mho$ of $\widetilde{\Delta}$.

PROOF. The proof is obtained by restricting the proof of Corollary 3.33 to L-models. $\square$

25

$$
\mathsf{d}^{\mathsf{l}}\ \frac{\widetilde{\Gamma}\{[\pi_0^\ell]\}}{\widetilde{\Gamma}\{\pi^\ell\}}
\qquad
\mathsf{d}^{\mathsf{r}}\ \frac{\widetilde{\Gamma}\{[\pi_0^\Re]\}}{\widetilde{\Gamma}\{\pi^\Re\}}
\qquad
\mathsf{t}^{\mathsf{l}}\ \frac{\widetilde{\Gamma}\{\pi_0^\ell\}}{\widetilde{\Gamma}\{\pi^\ell\}}
\qquad
\mathsf{t}^{\mathsf{r}}\ \frac{\widetilde{\Gamma}\{\pi_0^\Re\}}{\widetilde{\Gamma}\{\pi^\Re\}}
$$

$$
\mathsf{b}^{\mathsf{l}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}],\pi_0^\ell\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma},\pi^\ell]\}}
\qquad
\mathsf{b}^{\mathsf{r}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}],\pi_0^\Re\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma},\pi^\Re]\}}
\qquad
4^{\mathsf{l}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma},\pi^\ell]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma}],\pi^\ell\}}
\qquad
4^{\mathsf{r}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma},\pi^\Re]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma}],\pi^\Re\}}
$$

$$
5\mathsf{a}^{\mathsf{l}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}],\pi^\ell\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma},\pi^\ell]\}}
\qquad
5\mathsf{a}^{\mathsf{r}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}],\pi^\Re\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma},\pi^\Re]\}}
\qquad
5\mathsf{b}^{\mathsf{l}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}],[\widetilde{\Pi},\pi^\ell]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma},\pi^\ell],[\widetilde{\Pi}]\}}
\qquad
5\mathsf{b}^{\mathsf{r}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma}],[\widetilde{\Pi},\pi^\Re]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma},\pi^\Re],[\widetilde{\Pi}]\}}
$$

$$
5\mathsf{c}^{\mathsf{l}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma},[\widetilde{\Pi},\pi^\ell]]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma},[\widetilde{\Pi}],\pi^\ell]\}}
\qquad
5\mathsf{c}^{\mathsf{r}}\ \frac{\widetilde{\Gamma}\{[\widetilde{\Sigma},[\widetilde{\Pi},\pi^\Re]]\}}{\widetilde{\Gamma}\{[\widetilde{\Sigma},[\widetilde{\Pi}],\pi^\Re]\}}
$$

Figure 10: Biased rules for logics built from the axioms d, t, b, 4, and 5.

Definition 7.8 (Biased nested sequent calculi for the modal-cube logics). Let L be a logic from the modal cube. Its biased nested sequent calculus BNL is obtained by extending BNK with the biased versions x^{l} and x^{r} from Figure 10 of each nested sequent rule x added to NK in NL.

Theorem 7.9 (Equivalence between BNL and NL)**.** *Let L be a logic from the modal cube. For any biased version $\widetilde{\Gamma}$ of a nested sequent Γ, we have $\mathsf{BNL} \vdash \widetilde{\Gamma}$ iff $\mathsf{NL} \vdash \Gamma$.*

PROOF. The proof is analogous to that of Theorem 4.2. $\qquad\square$

Lemma 7.10 (Interpolant transformation for the modal-cube logics)**.** *For any logic L from the modal cube, if $\mho_1$ is an L-interpolant of $\widetilde{\Gamma}$, so is any generalized sequent $\mho_2$ decoratively equivalent to $\mho_1$.*

PROOF. The proof is obtained by restricting the proof of Lemma 5.8 to $\mathcal{M}$-decorations for L-models $\mathcal{M}$. $\quad\square$

Algorithm 7.11 (Interpolation algorithm for the modal-cube logics). We present the algorithm as a biased sequent calculus supplied with interpolant-handling machinery. It is required for the rules d^{l} and d^{r} that the interpolant be in a SCNF or a SDNF respectively and that disjuncts (conjuncts) within each conjunct (disjunct) of the SDNF (SCNF) be in a particular order. For a modal-cube logic L from the modal cube, the algorithm consists of all the interpolant-handling rules from Figure 6 as well as all the interpolant-handling rules from Figure 11 that correspond to the rules for L from Figure 10. Whenever the interpolant is represented as a conjunction of disjunctions of sequents or a disjunction of conjunctions of sequents, it is assumed to be in a SCNF or a SDNF respectively: i.e., the sequents are assumed to be singleton.

Lemma 6.3 also holds for all the interpolant-handling rules from Figure 11. We provide more details regarding the expansion of Lemma 6.4 to these new rules:

Lemma 7.12. (i) *If a given interpolant for the premise of the rule d^{l} or rule d^{r} from Figure 11 is not in the required form, it can be efficiently converted to a decoratively equivalent generalized sequent that is.*

(ii) *The object suggested by the algorithm as an interpolant for the conclusion of the rule d^{l} and rule d^{r} is always a well-formed generalized sequent, provided that a given interpolant for the premise of the rule is in the required form.*

(iii) *Each generalized sequent suggested by the algorithm for the conclusion of any rule from Figure 11 is structurally equivalent to the biased sequent from this conclusion.*

$$
\mathsf{d}^{\mathsf{l}} \ \frac{\widetilde{\Gamma}\{[\pi_0^\ell]\} \ \longleftarrow \ \bigotimes_{i=1}^{n} \left(\bigveeobig_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \ \varovee \ \bigveeobig_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right)}{\widetilde{\Gamma}\{\pi^\ell\} \ \longleftarrow \ \bigotimes_{i=1}^{n} \left(\widetilde{\Gamma}^\circ\{\Box \left(\bigvee_{k=1}^{l_i} A_{ik} \right)\} \ \varovee \ \bigveeobig_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\} \right)}
$$

$$
\mathsf{d}^{\mathsf{r}} \ \frac{\widetilde{\Gamma}\{[\pi_0^{\Re}]\} \ \longleftarrow \ \bigveeobig_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \ \varoovee \ \bigotimes_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right)}{\widetilde{\Gamma}\{\pi^{\Re}\} \ \longleftarrow \ \bigveeobig_{i=1}^{n} \left(\widetilde{\Gamma}^\circ\{\Diamond \left(\bigwedge_{k=1}^{l_i} A_{ik} \right)\} \ \varoovee \ \bigotimes_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\} \right)}
$$

For all the remaining rules, i.e., for the rules t^{l}, t^{r}, b^{l}, b^{r}, $\mathsf{4}^{\mathsf{l}}$, $\mathsf{4}^{\mathsf{r}}$, $\mathsf{5a}^{\mathsf{l}}$, $\mathsf{5a}^{\mathsf{r}}$, $\mathsf{5b}^{\mathsf{l}}$, $\mathsf{5b}^{\mathsf{r}}$, $\mathsf{5c}^{\mathsf{l}}$, and $\mathsf{5c}^{\mathsf{r}}$, the given interpolant for the premise is used as an interpolant for the conclusion.

Figure 11: Interpolant-handling rules for logics built from the axioms d, t, b, $\mathsf{4}$, and $\mathsf{5}$. An interpolant in the premise of the rule d^{l} (d^{r}) must be in a SCNF (SDNF). For both rules, we require that $\widetilde{\Gamma}\{\,\} \sim \Lambda_{ij}\{\,\} \sim \Pi_{ik}\{\,\}$ for all suitable i, j, and k.

PROOF. For the rules d^{l} and d^{r}, all the three statements follow from those for ν^{r} and ν^{l} respectively in Lemma 6.4, where the same interpolant transformation is used. Indeed, the premises of all the four rules are structurally equivalent and the conclusions of these rules are structurally equivalent. For all the other rules, only (iii) is applicable and its proof is trivial since the interpolant remains unchanged and the biased sequent in the conclusion remains structurally equivalent to the one in the premise.

Lemma 7.13. *For a logic* L *from the modal cube,*
$$L\widetilde{\Sigma} \Vdash_{\mathsf{L}} L\widetilde{\Pi} \quad \Longrightarrow \quad L\left(\widetilde{\Gamma}\{\widetilde{\Sigma}\}\right) \Vdash_{\mathsf{L}} L\left(\widetilde{\Gamma}\{\widetilde{\Pi}\}\right) \qquad and \qquad R\widetilde{\Sigma} \Vdash_{\mathsf{L}} R\widetilde{\Pi} \quad \Longrightarrow \quad R\left(\widetilde{\Gamma}\{\widetilde{\Sigma}\}\right) \Vdash_{\mathsf{L}} R\left(\widetilde{\Gamma}\{\widetilde{\Pi}\}\right).$$

PROOF. The proof is obtained by restricting the proof of Lemma 6.5 to $\mathcal{M}$-decorations for L-models $\mathcal{M}$. □

Corollary 7.14. *For a logic* L *from the modal cube, let* $L\widetilde{\Sigma} \Vdash_{\mathsf{L}} L\widetilde{\Pi}$ *and* $R\widetilde{\Sigma} \Vdash_{\mathsf{L}} R\widetilde{\Pi}$. *If* $\widetilde{\Gamma}\{\widetilde{\Sigma}\} \xleftarrow{\mathsf{L}} \mho$, *then* $\neg\mho \Vdash_{\mathsf{L}} L(\widetilde{\Gamma}\{\widetilde{\Pi}\})$ *and* $\mho \Vdash_{\mathsf{L}} R(\widetilde{\Gamma}\{\widetilde{\Pi}\})$.

Lemma 7.15. *For any rule from Figure 6 or Figure 11 for a logic* L *from the modal cube, if, for each premise of the rule, an* L-*interpolant is given in the required form, the generalized sequent* $\mho$ *suggested by the algorithm for the conclusion* $\widetilde{\Delta}$ *of the rule satisfies* $\neg\mho \Vdash_{\mathsf{L}} L\widetilde{\Delta}$ *and* $\mho \Vdash_{\mathsf{L}} R\widetilde{\Delta}$.

PROOF. Note that by Lemmas 6.4 and 7.12, $\mho$ is a well-defined generalized sequent structurally equivalent to $\widetilde{\Delta}$ and, hence, to both $L\widetilde{\Delta}$ and $R\widetilde{\Delta}$. For all the interpolant-handling rules from Figure 6, the argument is the same as in the proof of Lemma 6.9 except that $\Vdash$ is replaced by $\Vdash_{\mathsf{L}}$. For the remaining rules, it is important to remember that $\mathcal{M}, v \Vdash \pi_0$ and wRv imply $\mathcal{M}, w \Vdash \pi$ for any $\mathcal{M} = (W, R, V)$.

Cases t^{l} and t^{r}. By Corollary 7.14, it is sufficient to note that $\pi_0 \Vdash_{\mathsf{L}} \pi$ for any logic L validating t.

Cases b^{l} and b^{r}. By Corollary 7.14, it is sufficient to note that $[\Lambda], \pi_0 \Vdash_{\mathsf{L}} [\Lambda, \pi]$ for any L validating b.

Cases $\mathsf{4}^{\mathsf{l}}$ and $\mathsf{4}^{\mathsf{r}}$. By Corollary 7.14, it is sufficient to note that $[\Lambda, \pi] \Vdash_{\mathsf{L}} [\Lambda], \pi$ for any L validating $\mathsf{4}$.

Cases $\mathsf{5a}^{\mathsf{l}}$, $\mathsf{5a}^{\mathsf{r}}$, $\mathsf{5b}^{\mathsf{l}}$, $\mathsf{5b}^{\mathsf{r}}$, $\mathsf{5c}^{\mathsf{l}}$, and $\mathsf{5c}^{\mathsf{r}}$. By Corollary 7.14, it is sufficient to show that

$$[\Lambda], \pi \Vdash_{\mathsf{L}} [\Lambda, \pi] \ , \qquad [\Lambda], [\Theta, \pi] \Vdash_{\mathsf{L}} [\Lambda, \pi], [\Theta] \ , \quad \text{and} \qquad \left[\Lambda, [\Theta, \pi]\right] \Vdash_{\mathsf{L}} \left[\Lambda, [\Theta], \pi\right]$$

for any L validating $\mathsf{5}$. We only show the last statement, leaving the other two to the reader. Consider any matching $\mathcal{M}$-decorations $w, \left[\Lambda^*, [\Theta^\star, \pi]\right]$ and $w, \left[\Lambda^*, [\Theta^\star], \pi\right]$ of $\left[\Lambda, [\Theta, \pi]\right]$ and $\left[\Lambda, [\Theta], \pi\right]$ respectively

for some L-model $\mathcal{M} = (W, R, V)$, which must be Euclidean. Let $v = r(\Lambda^*)$ and $u = r(\Theta^\star)$. Assume that $\Vdash w, \left[\Lambda^*, [\Theta^\star, \pi]\right]$, i.e., that $\Vdash \Lambda^*$, or $\Vdash \Theta^\star$, or $\mathcal{M}, u \Vdash \pi$. We have wRv and vRu by the definition of decorations. Since vRv and uRv by the Euclideanity of R, we have that uRz implies vRz. Consequently,

$$\Vdash \Lambda^* \text{ or } \Vdash \Theta^\star \text{ or } \mathcal{M}, u \Vdash \pi \quad \Longrightarrow \quad \Vdash \Lambda^* \text{ or } \Vdash \Theta^\star \text{ or } (\exists z)(\mathcal{M}, z \Vdash \pi_0 \text{ and } uRz) \quad \Longrightarrow$$
$$\Vdash \Lambda^* \text{ or } \Vdash \Theta^\star \text{ or } (\exists z)(\mathcal{M}, z \Vdash \pi_0 \text{ and } vRz) \quad \Longrightarrow \quad \Vdash \Lambda^* \text{ or } \Vdash \Theta^\star \text{ or } \mathcal{M}, v \Vdash \pi \quad \Longrightarrow$$
$$\Vdash w, \left[\Lambda^*, [\Theta^\star], \pi\right] \ .$$

Cases $\mathsf{d^l}$ and $\mathsf{d^r}$. Just like the cases of ν^l and ν^r, these are the crucial cases because they require significant modifications to the structure of the interpolant. We show the case of $\mathsf{d^l}$, leaving $\mathsf{d^r}$ to the reader. For any modal-cube logic L validating d, the class of L-models consists exclusively of serial models. Assume that $\widetilde{\Gamma}\{[\pi_0^\ell]\} \xleftarrow{\ \mathsf{L}\ } \bigotimes_{i=1}^{n} \left(\bigovee_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right)$, where the generalized sequent is in a SCNF and $\widetilde{\Gamma}\{\ \} \sim \Lambda_{ij}\{\ \} \sim \Pi_{ik}\{\ \}$ for each $1 \leq i \leq n$, each $1 \leq j \leq m_i$, and each $1 \leq k \leq l_i$. Just like in the case of ν^r, each $\Pi_{ik}\{\ \} = \widetilde{\Gamma}^\circ\{\ \}$. We need to show that

$$\neg \bigotimes_{i=1}^{n} \left(\widetilde{\Gamma}^\circ\{\Box \bigvee_{k=1}^{l_i} A_{ik}\} \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\} \right) \Vdash_\mathsf{L} L\widetilde{\Gamma}\{\pi\} \quad \text{and} \quad \bigotimes_{i=1}^{n} \left(\widetilde{\Gamma}^\circ\{\Box \bigvee_{k=1}^{l_i} A_{ik}\} \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\} \right) \Vdash_\mathsf{L} R\widetilde{\Gamma}\{\varepsilon\} \ .$$

We start with showing the first consequence. Let $\mathcal{M} = (W, R, V)$ be an L-model. Consider arbitrary matching $\mathcal{M}$-decorations $\bigotimes_{i=1}^{n} \left((\widetilde{\Gamma}^\circ)^*\{w, \Box \bigvee_{k=1}^{l_i} A_{ik}\}_w \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}^*\{w\}_w \right)$ of $\bigotimes_{i=1}^{n} \left(\widetilde{\Gamma}^\circ\{\Box \bigvee_{k=1}^{l_i} A_{ik}\} \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}\{\varepsilon\} \right)$ and $L\widetilde{\Gamma}^*\{w, \pi\}_w$ of $L\widetilde{\Gamma}\{\pi\}$. Assume that $\nVdash \bigotimes_{i=1}^{n} \left((\widetilde{\Gamma}^\circ)^*\{w, \Box \bigvee_{k=1}^{l_i} A_{ik}\}_w \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}^*\{w\}_w \right)$. Then there exists $1 \leq L \leq n$ such that $\nVdash \Lambda_{Lj}^*\{w\}_w$ for each $1 \leq j \leq m_L$ and $\nVdash (\widetilde{\Gamma}^\circ)^*\{w, \Box \bigvee_{k=1}^{l_L} A_{Lk}\}_w$. It follows that $\mathcal{M}, w \nVdash \Box \bigvee_{k=1}^{l_L} A_{Lk}$. Then there exists $v \in W$ such that wRv and $\mathcal{M}, v \nVdash \bigvee_{k=1}^{l_L} A_{Lk}$. In particular, $\mathcal{M}, v \nVdash A_{Lk}$ for each $1 \leq k \leq l_L$. Since each context $\Pi_{Lk}\{\ \} = \widetilde{\Gamma}^\circ\{\ \}$ is void, it follows that $\nVdash \Pi_{Lk}^*\{w, [v, A_{Lk}]\}_w$, where $\Pi_{ik}^*\{\ \}_w := (\Gamma^\circ)^*\{\ \}_w$ for each $1 \leq i \leq n$ and each $1 \leq k \leq l_i$. Further, since $w, [v]$ is void, it follows that $\nVdash \Lambda_{Lj}^*\{w, [v]\}_w$ for each $1 \leq j \leq m_L$. We conclude that $\nVdash \bigovee_{k=1}^{l_L} \Pi_{Lk}^*\{w, [v, A_{Lk}]\}_w \oslash \bigovee_{j=1}^{m_L} \Lambda_{Lj}^*\{w, [v]\}_w$; hence, $\nVdash \bigotimes_{i=1}^{n} \left(\bigovee_{k=1}^{l_i} \Pi_{ik}^*\{w, [v, A_{ik}]\}_w \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}^*\{w, [v]\}_w \right)$. This last decoration of $\bigotimes_{i=1}^{n} \left(\bigovee_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right)$ matches the decoration $L\widetilde{\Gamma}^*\{w, [v, \pi_0]\}_w$ of $L\widetilde{\Gamma}\{[\pi_0]\}$. Thus, $\Vdash L\widetilde{\Gamma}^*\{w, [v, \pi_0]\}_w$ follows from $\neg \bigotimes_{i=1}^{n} \left(\bigovee_{k=1}^{l_i} \Pi_{ik}\{[A_{ik}]\} \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}\{[\varepsilon]\} \right) \Vdash_\mathsf{L} L\widetilde{\Gamma}\{[\pi_0]\}$.

$$\Vdash L\widetilde{\Gamma}^*\{w, [v, \pi_0]\}_w \quad \Longrightarrow \quad \Vdash L\widetilde{\Gamma}^*\{w\}_w \text{ or } \mathcal{M}, v \Vdash \pi_0 \quad \Longrightarrow$$
$$\Vdash L\widetilde{\Gamma}^*\{w\}_w \text{ or } \mathcal{M}, w \Vdash \pi \quad \Longrightarrow \quad \Vdash L\widetilde{\Gamma}^*\{w, \pi\}_w$$

because wRv. Note that this consequence does not require the explicit use of seriality.

Now we show the contraposition of the second consequence. For any L-model $\mathcal{M} = (W, R, V)$, we know that R is serial. Consider arbitrary matching $\mathcal{M}$-decorations $\bigotimes_{i=1}^{n} \left((\widetilde{\Gamma}^\circ)^*\{w, \Box \bigvee_{k=1}^{l_i} A_{ik}\}_w \oslash \bigovee_{j=1}^{m_i} \Lambda_{ij}^*\{w\}_w \right)$

of $\bigotimes_{i=1}^{n} \left(\widetilde{\Gamma}^{\circ} \{ \square \bigvee_{k=1}^{l_i} A_{ik} \} \otimes \bigotimes_{j=1}^{m_i} \Lambda_{ij} \{ \varepsilon \} \right)$ and $R\widetilde{\Gamma}^* \{w\}_w$ of $R\widetilde{\Gamma} \{ \varepsilon \}$. Assume that $\nVdash R\widetilde{\Gamma}^* \{w\}_w$. By seriality of R, there exists $v \in W$ such that wRv. Then, $\nVdash R\widetilde{\Gamma}^* \{w, [v]\}_w$, which is a decoration of $R\widetilde{\Gamma} \{ [\varepsilon] \}$. For $\Pi_{ik}^* \{ \ \}_w := (\widetilde{\Gamma}^{\circ})^* \{ \ \}_w$ for each $1 \leq i \leq n$ and each $1 \leq k \leq l_i$ this last decoration matches the decoration $\bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}^* \{w, [v, A_{ik}]\}_w \otimes \bigotimes_{j=1}^{m_i} \Lambda_{ij}^* \{w, [v]\}_w \right)$ of $\bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik} \{ [A_{ik}] \} \otimes \bigotimes_{j=1}^{m_i} \Lambda_{ij} \{ [\varepsilon] \} \right)$. Therefore,

$\nVdash \bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik}^* \{w, [v, A_{ik}]\}_w \otimes \bigotimes_{j=1}^{m_i} \Lambda_{ij}^* \{w, [v]\}_w \right)$ because $\bigotimes_{i=1}^{n} \left(\bigotimes_{k=1}^{l_i} \Pi_{ik} \{ [A_{ik}] \} \otimes \bigotimes_{j=1}^{m_i} \Lambda_{ij} \{ [\varepsilon] \} \right) \Vdash_{\mathsf{L}} R\widetilde{\Gamma} \{ [\varepsilon] \}$.

Thus, there exists $1 \leq L \leq n$ such that $\nVdash \bigotimes_{k=1}^{l_L} \Pi_{Lk}^* \{w, [v, A_{Lk}]\}_w \otimes \bigotimes_{j=1}^{m_L} \Lambda_{Lj}^* \{w, [v]\}_w$. We have $\mathcal{M}, v \nVdash A_{Lk}$ for each $1 \leq k \leq l_L$ and $\nVdash \Lambda_{Lj}^* \{w\}_w$ for each $1 \leq j \leq m_L$. It follows that $\mathcal{M}, v \nVdash \bigvee_{k=1}^{l_L} A_{Lk}$. Since wRv, we have $\mathcal{M}, w \nVdash \square \bigvee_{k=1}^{l_L} A_{Lk}$. Given that $(\widetilde{\Gamma}^{\circ})^* \{w\}_w$ is void, we conclude $\nVdash (\widetilde{\Gamma}^{\circ})^* \{w, \square \bigvee_{k=1}^{l_L} A_{Lk} \}_w$. Overall, $\nVdash (\widetilde{\Gamma}^{\circ})^* \{w, \square \bigvee_{k=1}^{l_L} A_{Lk}\}_w \otimes \bigotimes_{j=1}^{m_L} \Lambda_{Lj}^* \{w\}_w$, and, finally, $\nVdash \bigotimes_{i=1}^{n} \left((\widetilde{\Gamma}^{\circ})^* \{w, \square \bigvee_{k=1}^{l_i} A_{ik} \}_w \otimes \bigotimes_{j=1}^{m_i} \Lambda_{ij}^* \{w\}_w \right)$. $\square$

Theorem 7.16 (Interpolation theorem for the modal-cube logics). *Let L be a logic from the modal cube. For any biased sequent $\widetilde{\Gamma}$, derivable in BNL, Algorithm 7.11 finds an interpolant $\mho$ of $\widetilde{\Gamma}$.*

PROOF. Follows from Lemmas 6.3 (extended to the new steps from Figure 11), 6.4, 7.12, and 7.15. $\square$

Corollary 7.17 (Interpolation Theorem). *All logics from the modal cube have the CIP.*

PROOF. Let $\mathsf{L} \vdash A \supset B$ for some logic L from the modal cube. By completeness of NL, clearly $\mathsf{NL} \vdash \neg A \vee B$ and $\mathsf{NL} \vdash \neg A, B$. Thus, $\mathsf{BNL} \vdash \neg A^\ell, B^\Re$ by Theorem 7.9. By Theorem 7.16, $\neg A^\ell, B^\Re \xleftarrow{\mathsf{L}} \mho$ for some interpolant $\mho$. By Corollary 7.7, the formula $\underline{\mho}$ contains only common propositional variables of $\neg A$ and B, i.e., of A and B, and, in addition, $\mathsf{L} \vdash \neg\neg A \supset \underline{\mho}$ and $\mathsf{L} \vdash \underline{\mho} \supset B$. Thus, $\mathsf{L} \vdash A \supset \underline{\mho}$ and $\underline{\mho}$ is an interpolant of A and B for the logic L. $\square$

8. Future Work

It would be interesting to extend our method to first-order-based logics and see where exactly the method breaks for those logics that are known not to have the CIP. Another natural development is to adapt our method to labelled sequents: labelled sequents are known to be more general than nested sequents. Further, our method heavily relies on the classical nature of the underlying logics because Brünnler's nested sequent calculus we use is not suitable, for example, for intuitionistic-based logics. There are several recently developed versions of nested sequents adapted for intuitionistic logic of various flavors, notably by Fitting [5], by Goré et al. [9], and by Straßburger [13]. Thus, it is natural to see whether our method can be extended to such intuitionistic nested sequents.

References

[1] Brünnler, K., Jul. 2009. Deep sequent systems for modal logic. Archive for Mathematical Logic 48 (6), 551–577.

[2] Fitting, M., Apr. 1972. Tableau methods of proof for modal logics. Notre Dame Journal of Formal Logic 13 (2), 237–247.

[3] Fitting, M., Feb. 1996. A program to compute Gödel-Löb fixpoints. Bulletin of the EATCS 58, 118–130.

[4] Fitting, M., Mar. 2012. Prefixed tableaus and nested sequents. Annals of Pure and Applied Logic 163 (3), 291–313.

[5] Fitting, M., 2014. Nested sequents for intuitionistic logics. Notre Dame Journal of Formal Logic 55 (1), 41–61.

[6] Gabbay, D. M., Maksimova, L., 2005. Interpolation and Definability: Modal and Intuitionistic Logic. Oxford Science Publications.

[7] Garson, J., 2013. Modal logic. In: Zalta, E. N. (Ed.), The Stanford Encyclopedia of Philosophy, spring 2013 Edition. URL http://plato.stanford.edu/archives/spr2013/entries/logic-modal/

[8] Goetschi, R., Kuznets, R., Sep. 2012. Realization for justification logics via nested sequents: Modularity through embedding. Annals of Pure and Applied Logic 163 (9), 1271–1298.

[9] Goré, R., Postniece, L., Tiu, A., 2010. Cut-elimination and proof search for bi-intuitionistic tense logic. In: Beklemishev, L., Goranko, V., Shehtman, V. (Eds.), Advances in Modal Logic, Volume 8. College Publications, pp. 156–177.

[10] Marin, S., Straßburger, L., 2014. Label-free modular systems for classical and intuitionistic modal logics. In: Goré, R., Kooi, B., Kurucz, A. (Eds.), Advances in Modal Logic, Volume 10. College Publications, pp. 387–406.

[11] Massacci, F., Apr. 2000. Single step tableaux for modal logics: Computational properties, complexity and methodology. Journal of Automated Reasoning 24 (3), 319–364.

[12] Sambin, G., Valentini, S., Aug. 1982. The modal logic of provability, the sequential approach. Journal of Philosophical Logic 11 (3), 311–342.

[13] Straßburger, L., 2013. Cut elimination in nested sequents for intuitionistic modal logics. In: Pfenning, F. (Ed.), Foundations of Software Science and Computation Structures, 16th International Conference, FOSSACS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16–24, 2013, Proceedings. Vol. 7794 of Lecture Notes in Computer Science. Springer, pp. 209–224.