A Formal Characterization of Semantic Pollution of Modal Proof Systems

by Robin Martinot

Paper currently under review

Abstract

This paper provides a formal characterization of the phenomenon of 'semantic pollution' relative to proof systems for modal logic and Kripke semantics. We propose that semantic pollution can be made precise by several properties of syntax occurring in proof systems. First, our base requirement for semantic pollution is given by the property of violating invariance results under Kripke model equivalences. On top of that, the distinction between local and global syntax, and between syntax that is dependent on and independent from the propositional valuation, induce four levels of semantic pollution: weak pollution, global pollution, local pollution and strong pollution. We analyze several main proof systems for (extensions of) modal logic in terms of these levels of pollution: the display calculus, the hybrid calculus and the labeled calculus. The results show that the display calculus is only weakly semantically polluted, while the hybrid calculus has formulas introducing several types of pollution. The only calculus that is strongly semantically polluted is the labeled calculus. These formal results are in line with general intuitions about semantic pollution. Additionally, our formal framework for semantic pollution is suitable for applications to other logics and semantics. Finally, we comment on the relation of our framework to the philosophical debate surrounding semantic pollution.

Keywords: semantic pollution, labeled calculi, proof systems for modal logic, philosophy of proof theory

1 Introduction

The phenomenon of semantic pollution has attracted growing attention in the search for effective proof systems for modal logic. At a high level, semantic pollution can be intuitively understood as the notion of importing a model-theoretic semantics into the proof-theoretic syntax. However, what this precisely entails remains largely ambiguous. In this work, we propose three formal measures of semantic pollution specific to modal logic and Kripke semantics, resulting in four distinct levels of pollution. We classify formulas from the display calculus, labeled calculus, and hybrid calculus according to these levels, and embed these findings within a broader philosophical discussion of semantic pollution. Our classification aligns with common intuitions about semantically polluted proof systems, providing a formal foundation for these intuitions. First, we review how semantic pollution has been discussed in the literature, followed by an outline of the structure of the paper.

1.1 Semantic pollution

Traditional natural deduction and sequent systems for modal logic have struggled to meet core requirements, such as analyticity and normalization of formal derivations (Negri, 2011). As a result, numerous generalizations and modifications of standard (sequent) calculi have been proposed, often by extending the proof-theoretic language. These include hypersequent calculi (introduced by Avron (1987) and Pottinger (1983)), tree-hypersequent or nested sequent calculi (see Brünnler (2010) and Poggiolesi (2009), predated by Bull (1992) and Kashima (1994)), display calculi (Belnap, 1982; Wansing, 1994) and labeled calculi (see e.g. (Negri, 2005), and dating back to Simpson (1994) and Kanger (1957)). The modal language itself has also been extended in hybrid logic (see e.g. (Braüner, 2010)), which solves the problem at the level of the object language. Labeled calculi, in particular, are considered to semantically pollute the modal language, by explicitly internalizing Kripke semantics into the proof system. The calculus introduces labels $x, y, z, \ldots$ and a 'forcing relation' ':' to accompany every modal formula occurring in the proof system, as well as relational atoms xRy as new primitive formulas. Consider for instance the rules for $\Box$ in the system **G3K** (Negri, 2005):

$$\frac{y : A, x : \Box A, xRy, \Gamma \Rightarrow \Delta}{x : \Box A, xRy, \Gamma \Rightarrow \Delta} L\Box \qquad \frac{xRy, \Gamma \Rightarrow \Delta, y : A}{\Gamma \Rightarrow \Delta, x : \Box A} R\Box$$

While some proof theorists claim labeled calculi provide desirable technical properties — such as "analyticity, applicability to proof search [and] the possibility to obtain direct completeness proofs" (Negri, 2011) — the literature also refers to them more apprehensively with philosophical concerns about semantic pollution. Consider for instance that "a philosophical objection to this kind of system is that it builds-in the (desired) semantics into the given syntax" (Braüner and Paiva, 2006); "[t]he use of a labeled calculus has been sometimes criticized, as mixing semantic elements into what should be a purely syntactic proof system" (Negri, 2011); "some proof-theorists are not satisfied with the idea of labels in proofs that would be seen as 'semantical pollution' because some ingredients of a labeled formalism resemble model-theoretic objects" (Marin, 2018); and "[s]ome have criticised this as a lack of syntactic purity, i.e. as the presence of "semantic pollution"; others defend it as allowing calculi for otherwise unmanageable logics" (Dyckhoff, 2016).

At this point, two main more elaborate analyses of semantic pollution can be found in the literature. The first is by Read (2015), who argues that the explicit encoding of Kripke semantics in labeled calculi is in fact a virtue, as opposed to tree-hypersequent or nested calculi, which attempt to 'obscure' the semantics in their structural syntax. A recent paper by De Martin Polo (2024) includes a more overarching insight into the types of labeling used by proof theorists, and outlines the current main philosophical attitudes towards using labels. Both authors argue that labeled calculi are suitable for *inferentialism*, the endeavour to specify the meanings of logical constants in terms of their rules of inference — this has been a main potential philosophical drawback of semantically polluted calculi. Thus, the philosophical debate so far seems tentatively ready to accept semantic pollution. However, in the interest of painting a complete picture, we believe that the philosophical views expressing caution with respect to semantic pollution at this point deserve more attention. Additionally, the notion of semantic pollution broadly as a relation between syntax and semantics, remains underspecified in studies so far.

Hence, we have sufficient reason to take a closer look at what semantic pollution could amount to, more formally. As far as we are aware, the only precisification of the notion of syntactic purity (as the counterpart to semantic pollution) has been put forward by Poggiolesi (2010), and appeals to the difference between *internal* calculi (where each element of the proof system has an interpretation as a formula of the logic) and *external* calculi (where some elements cannot be interpreted as a formula of the logic).[1] She proposes that a sequent calculus is syntactically pure if it does not "make use of explicit semantic elements", which are exactly elements that make the proof calculus external, i.e., that prevent translation of a sequent to "a formula equivalent to the sequent". Poggiolesi's account follows Avron (1996)'s requirements for 'good' proof systems. There, he suggests that "[a] sequent calculus should be independent of any particular semantic[s]. One should not be able to guess, just from the form of the structures which are used, the intended semantics of a given proof system". Poggiolesi calls this 'strong syntactic purity', since it is hardly satisfied by any proof system: already the classical propositional sequent calculus violates a reasonable interpretation of this type of syntactic purity. Her account then forms a compromised definition of 'weak syntactic purity'.

Taking Poggiolesi (2010)'s proposal as a starting point, we believe that the wide range of proof systems for modal logic could benefit from a more nuanced account of semantic pollution, and one that is motivated by a more semantic perspective. In attempting to provide one, we will define concrete measures (resulting in four levels) of semantic pollution of proof systems relative to the modal language, and provide a first comprehensive overview of the behaviour of several proof systems for (extensions of) modal logic under these measures. Since the debate on semantic pollution focuses almost entirely on logics with Kripke semantics, we restrict our analysis to this semantics and to proof systems for (extensions of) modal logic. Our framework, however, is broadly applicable to other logics and semantics, on which we briefly comment in the conclusion.

1.2 Outline of the paper

Before we dive into defining measures of semantic pollution, Section 2 will provide an overview of the formalities necessary for our characterization, paving the way to Section 3 and 4. A base requirement for semantic pollution will be defined in Section 3, which will also analyze several proof systems in terms of this requirement. Section 4 will then define four levels of semantic pollution and discuss its results, by building on the base requirement. Finally, Section 5 will provide an analysis of the philosophical debate surrounding semantic pollution.

Remark 1. There is a general distinction between 'bottom-up' and 'top-down' approaches to formalizing philosophical notions. Generally, bottom-up approaches focus on the use of this notion by experts in practice, and aim to provide a formalization that matches this practice closely. Top-down approaches, on the other hand, develop a framework based on theoretical (possibly idealized) principles intuitively underlying a notion. Practical examples may then instead be measured by the standard of this framework. Both approaches are valuable for different reasons. This paper starts 'bottom-up' with the intuitions on semantic pollution mentioned in Section 1.1; we aim to provide a formal framework capturing the idea that labeled calculi possess most semantic pollution, and that the usual (e.g. propositional) sequent calculus possesses none. The inspiration for and specification of our measures of semantic pollution in Section 3 and 4, however, also come with top-down influences on what we think makes up a 'semantic nature'.

[1]Although as mentioned by Lyon et al. (2023), "the proof-theoretic community lacks consensus on how ['internal' and 'external' calculus] should be precisely defined". See footnote 3 for an elaboration on the term 'formula interpretation'.

2 Semantic pollution as a formal property

This section provides the formal preliminaries to our characterization of semantic pollution. First, we define the various formal languages that we study in this paper. We then introduce notation for restricted languages and for results of semantic pollution and syntactic purity.

2.1 Proof-theoretic languages

Let $\mathsf{Prop} = \{p, q, r, ...\}$ be a set of propositional variables. From here on, the *object language* L will refer to the basic modal language, generated by the grammar:[2]

$$A ::= p \mid \top \mid \bot \mid A \wedge A \mid A \to A \mid \Box A \qquad \text{(where } p \in \mathsf{Prop}\text{)}$$

As its model theory, we take the usual Kripke semantics, concerning Kripke models $\mathsf{M} = (W, R, V)$ consisting of a set of states W, an accessibility relation $R \subseteq W \times W$ and a propositional valuation function $V : \mathsf{Prop} \to \mathcal{P}(W)$. L-formulas receive classical truth conditions relative to a pointed model, as specified e.g. in (Blackburn, De Rijke, and Venema, 2001).

A *proof-theoretic language* PL, generated by the grammar of a proof system, can then extend L by any new syntax.

Remark 2. It will not be possible for L-formulas to be semantically polluted: the same will hold for proof-theoretic formulas *translatable* to L.[3] Hence, if the translatability of a proof-theoretic formula depends on the background logic, then whether proof-theoretic formulas are semantically polluted does, too. This for instance concerns common extensions of the Gentzen calculus for modal logic, like the nested calculus or tree-hypersequent calculus, and the hypersequent calculus (and even the usual Gentzen arrow and comma). Relative to a background logic, these structural sequents commonly have intended translations in terms of, for instance, disjunction and the box operator. Our approach will render them automatically syntactically pure. However, such calculi have been speculated to possess (some version of) semantic pollution by Read (2015) and De Martin Polo (2024); hence, we will discuss them more in Section 5.

We now present the three proof-theoretic languages that this paper restricts itself to, as well as the truth conditions of the formulas they introduce in terms of Kripke semantics. In particular, the general notation PL for proof-theoretic language will have three instantiations:

- DL (the display language), based on the display calculus as in (Wansing, 1994).
- LL (the labeled language), based on the labeled calculus as in (Negri, 2005).
- HL (the hybrid language), based on propositional hybrid logic as in (Braüner, 2010), which is a logical extension of the basic modal language.

Both languages LL and HL introduce a set of variables, called 'labels' in labeled calculi and 'nominals' in hybrid logic. While labels are intuitively understood as naming states in a Kripke model, nominals were motivated by the want to formalize natural language sentences referring to specific

[2]For compactness, we leave out $\neg$, $\vee$, and $\Diamond$, but they may be assumed to exist explicitly as well.

[3]By 'translatable', we mean that there is some translation function t from proof-theoretic expressions A to expressions in the logic, such that t has suitable properties (like compositionality), and such that provability of A from B implies logical entailment of $t(A)$ from $t(B)$. We will use 'being translatable to the logic' synonymously with 'having a formula interpretation in the logic'.

time points or individuals. We will treat labels and nominals uniformly throughout the paper, and call them *name variables* for both languages, given by a set $\mathsf{Var} = \{a, b, c, ..., x, y, z, ...\}$.[4] Formulas including name variables will receive truth conditions in terms of Kripke models extended by an assignment function $\tau : \mathsf{Var} \to W$. Instead, the formulas of DL will receive truth conditions in terms of the usual (unextended) Kripke models.

M will always stand for a regular (unextended) Kripke model (W, R, V). A *model for a proof-theoretic language* PM will be a model that can provide truth conditions for the formulas in DL, LL or HL. This means that PM can either equal a regular Kripke model M (in case of DL), or an extended model (M, τ) (in case of LL and HL). Where it is relevant to know if we are talking about M or a pair (M, τ), we will always use the specific notation. Similarly, a *frame for a proof-theoretic language* PF will either stand for a regular Kripke frame $\mathsf{F} = (W, R)$ or for an extended frame (F, τ).

Now consider the proof-theoretic languages. DL extends L by forming the following grammar:

$$A ::= B \mid \mathbf{I} \mid (A \circ A) \mid *A \mid \bullet A \qquad \text{(where } B \in \mathsf{L})$$

That is, the display calculus adds new structural connectives $\mathbf{I}$, $\circ$, $*$, and $\bullet$ to the modal language. The resulting structures have two different translations into tense logic, depending on their position (antecedent or consequent) in a sequent. The operators $\mathbf{I}$, $\circ$ and $*$ have intended translations in terms of conjunction, disjunction, negation, truth and falsum (see (Wansing, 1994) for details). In consequent position, $\bullet A$ may be translated as $\Box A$. We will focus most, however, on the occurrences of $\bullet A$ in a proof system in *antecedent* position (and whenever we talk of the bullet operator from now on, we will assume this interpretation). In antecedent position, it has the intended translation and so the truth condition of a backwards diamond:

$$\mathsf{M}, w \vDash \bullet A \text{ iff } \exists v(Rvw \wedge \mathsf{M}, v \vDash A)$$

Second, the language LL uses the name variables in Var (as just defined above) to accompany object language formulas and to form new atomic formulas xRy. Thus, LL extends L by the following formulas:

$$A ::= B \mid x : B \mid xRy \qquad \text{(where } B \in \mathsf{L} \text{ and } x, y \in \mathsf{Var})$$

The truth conditions of the new proof-theoretic formulas are then defined as follows, for a Kripke model extended by assignment function τ.

$$\mathsf{M}, \tau, w \vDash x : A \text{ iff } \mathsf{M}, \tau(x) \vDash A$$
$$\mathsf{M}, \tau, w \vDash xRy \text{ iff } \tau(x)R\tau(y)$$

Finally, the language HL uses the name variables in Var to introduce various new operators. Some of them are closely related to the labeled calculus, and thus HL provides a good comparison for semantic pollution.[5] HL extends the modal language by the following grammar.

$$A ::= p \mid a \mid \top \mid \bot \mid A \wedge A \mid A \to A \mid \Box A \mid @_a A \mid \forall a A \mid \downarrow a A \qquad \text{(where } p \in \mathsf{Prop}, a \in \mathsf{Var})$$

[4] In the literature, it is common for nominals to use early-occurring alphabet letters $a, b, c, ...$, and for labels to use late-occurring alphabet letters $x, y, z, ...$ In accordance with this custom, our examples for HL will commonly use letters $a, b, c, ...$, and our examples for LL will commonly use letters $x, y, z, ...$

[5] Technically, as HL is primarily an object language, and secondarily a proof-theoretic language, results of semantic pollution will apply to this language in both roles. We thus take semantic pollution then generally to apply to any language *extending the basic modal language*, not just to proof-theoretic languages, although of course the use of languages in proof systems is what concerns the debate on semantic pollution most, and what we focus on here.

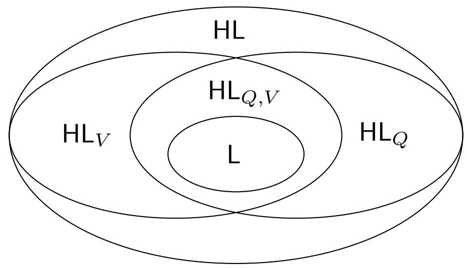

Figure 1: The restricted versions of the language HL.

The truth conditions of the hybrid formulas are as follows, again for a Kripke model extended by an assignment function τ. Let τ_a be the assignment function that agrees with τ on every nominal assignment, except possibly on a. Specifically, let $\tau_{[a \mapsto w]}$ be the assignment function that agrees with τ on every nominal assignment, except possibly on a, which is sent to world w.

$\mathsf{M}, \tau, w \models a$ iff $\tau(a) = w$

$\mathsf{M}, \tau, w \models @_a A$ iff $\mathsf{M}, \tau, \tau(a) \models A$

$\mathsf{M}, \tau, w \models \forall a A$ iff for any τ_a, $\mathsf{M}, \tau_a, w \models A$

$\mathsf{M}, \tau, w \models \downarrow a A$ iff $\mathsf{M}, \tau_{[a \mapsto w]}, w \models A$

2.2 Restricted versions of HL

There is now reason to consider several restricted versions of HL. This is because we are mainly interested in the effects of extending L by *individual* operators (that are, as far as possible, only applied to formulas in L), so that we can more easily compare their individual behaviour to that of L-formulas, and to that of other proof-theoretic operators. This is unproblematic in LL, as labels are already only applied to L-formulas (and relational atoms are nullary). In DL, although the display operators can stack, it is only $\bullet A$ that has no translation into L. We thus do not have to make any effort to consider $\bullet$ separately.

However, HL contains multiple hybrid operators $a, @_a, \forall a$ and $\downarrow a$, and the latter three are applied to the full language HL. In order to approach the effect of applying these operators just to L, we will try to limit combinations of hybrid operators only to the necessary ones (e.g., in order to have any effect at all, the operators $\forall a$ and $\downarrow a$ must at least combine with name variables a). Three restricted versions of HL (see Figure 1) will let us distinguish between different semantic pollution results, that will indicate from which language they are obtained. First, consider HL$_V$ (V for 'variable'), which reduces the set of name variables to a singleton.

$$A ::= p \mid a \mid \top \mid \bot \mid A \wedge A \mid A \to A \mid \Box A \mid @_a A \mid \forall a A \mid \downarrow a A \quad (\text{where } p \in \mathsf{Prop},\ a \in \mathsf{Var} \text{ for } |\mathsf{Var}| = 1)$$

Second, consider HL$_Q$, which excludes the operator $\forall a A$ from the language (Q for its 'quantifier'-like property).

$$A ::= p \mid a \mid \top \mid \bot \mid A \wedge A \mid A \to A \mid \Box A \mid @_a A \mid \downarrow a A \quad (\text{where } p \in \mathsf{Prop},\ a \in \mathsf{Var})$$

Finally, consider HL$_{Q,V}$, which combines the previous two changes.

$$A ::= p \mid a \mid \top \mid \bot \mid A \wedge A \mid A \to A \mid \Box A \mid @_a A \mid \downarrow a A \quad (\text{where } p \in \mathsf{Prop},\ a \in \mathsf{Var} \text{ for } |\mathsf{Var}| = 1)$$

2.3 Semantic pollution and syntactic purity of languages

We now need to specify a few more ingredients for setting up the property of semantic pollution. First, we will let any symbol C in a PL-grammar correspond to a *formula type*, which consists of all formulas whose outer logical symbol is C. Consider its precise definition.

Definition 1 (Context language). For operators or metavariables C from one of our proof-theoretic language PL, their possible *context languages*, denoted CL, are defined as follows.

- The context language for $C \in \mathsf{DL} \cup \mathsf{LL}$ is always L.
- The context language for $C \in \mathsf{HL}$ is HL, HL_V, HL_Q, $\mathsf{HL}_{Q,V}$ or L.

Definition 2 (Formula type). Consider an operator or metavariable C in one of our proof-theoretic languages PL. We define the *formula type* of C relative to a *context language* CL (possibly $\mathsf{PL} \neq \mathsf{CL}$), and indicate it by $C(\mathsf{CL})$. We distinguish three cases, classifying the formulas from DL, LL and HL.[6]

1. *C is a variable.*

$$p(\mathsf{CL}) = \mathsf{Prop}$$
$$a(\mathsf{CL}) = \mathsf{Var}$$

2. *C is a nullary operator.*

$$R(\mathsf{CL}) = \{xRy \mid x, y \in \mathsf{Var}\}$$
$$\alpha(\mathsf{CL}) = \{\alpha\} \text{ for } \alpha \in \{\mathbf{I}, \top, \bot\}$$

3. *C is an n-ary operator.*

$$C(\mathsf{CL}) = \{C(A_1, ..., A_n) \mid A_i \in \mathsf{CL} \text{ for } 1 \leqslant i \leqslant n\}$$

For example, $\bullet(p \vee \neg q) \in \bullet(\mathsf{L})$. Throughout the paper, we will propose a base requirement (denoted by BR) for semantic pollution, and four more elaborated levels of semantic pollution. Satisfaction of the base requirement and of a level of pollution will then be introduced to a proof-theoretic language, and in turn to a proof system, by a *formula type*. By focusing on formula types, the truth conditions of proof-theoretic formulas will be analyzed at the level of their main operator. This way, it is really the operators introduced by the proof system that get full responsibility for their semantic effects.

Remark 3. We will sometimes abuse terminology and say that a formula type or a proof-theoretic language itself is already semantically polluted — but it should be kept in mind that these statements in the end only serve as shorthand for saying that a *proof system* is semantically polluted (one introducing the particular proof-theoretic language and formula type).

Suppose we are given the definition for the base requirement (BR). Then its satisfaction is determined for a formula type $C(\mathsf{L})$ with respect to the *modal* context language L.

- The result po means $C(\mathsf{L})$ satisfies BR (po for (semantic) pollution).[7]
- The result pu means that $C(\mathsf{L})$ does not satisfy BR (pu for (syntactic) purity).

[6] For the first two cases, $C(\mathsf{PL}) = C(\mathsf{CL})$ for each $\mathsf{PL} \neq \mathsf{CL}$ (the context language does not matter).

[7] We will assume that $\bullet(\mathsf{L}) = \bullet(\mathsf{DL})$, as their results for the measures for semantic pollution will be the same.

For HL, $C(\mathsf{L})$ sometimes captures 'vacuous purity results'. Namely, its operators cannot always show their power when restricted to L. On the other hand, $C(\mathsf{HL})$ will sometimes fail to show the *individual* effect of C, as C can combine with any other operator of HL. Hence, we will introduce some weaker types of pollution that arise for context languages L' such that $\mathsf{L} \subset \mathsf{L}' \subseteq \mathsf{HL}$. po is then the strongest type of pollution for HL-operators. A result $\mathsf{po}_{\mathsf{L}'}$ will indicate that L' is the smallest language (among the ones we have selected) to satisfy BR. This means that the operator restricted to a smaller language will not satisfy BR, and will instead satisfy a purity result: several of such results will be proven in the Appendix. We use the following notation for weaker types of pollution.

- The result $\mathsf{po}_{Q,V}$ means that $C(\mathsf{HL}_{Q,V})$ satisfies BR.
- The result po_{V} means that $C(\mathsf{HL}_{V})$ satisfies BR.
- The result po_{Q} means that $C(\mathsf{HL}_{Q})$ satisfies BR.
- The result $\mathsf{po}_{\mathsf{HL}}$ means that $C(\mathsf{HL})$ satisfies BR.

It will follow from the definition of BR that po implies $\mathsf{po}_{Q,V}$. Furthermore, $\mathsf{po}_{Q,V}$ implies po_{V} and po_{Q}, and the latter two both imply $\mathsf{po}_{\mathsf{HL}}$, which is the weakest type of pollution.

2.4 Key aspects

Our method has a few characteristic properties, which we here emphasize for clarity. First, our approach is heavily *framework-dependent*. We see semantic pollution as dependent on an object language, a proof-theoretic language, and a particular semantics. We thus do not accommodate a notion of semantic pollution that is irrespective of this background context, and that only relies on proof-theoretic tools. We consider this to be natural: there is simply no absolute standard for syntax to be 'syntactic' (any expression can in principle included in a syntax) — hence, we need a baseline connection between an object language and a semantics in order to distinguish 'semantic' and 'syntactic' syntax in the proof-theoretic language.

Second, we define *'operator-level'* semantic pollution (by using formula types). One might instead focus on properties of individual concrete formulas (members of a formula type), or even of languages as an entirety. However, the interesting level of detail seems to concern operators that a proof-theoretic language introduces. This emphasizes that the object language itself is syntactically pure, and that the proof system introduces pollution by *adding* to this language.

Third, we see semantic pollution as something *static*: if a formula type satisfies it, then the proof system as a whole (and any formal proof using instances of the formula type) can be considered semantically polluted. This perspective ignores the behaviour of proof-theoretic syntax inside inference rules (i.e., the way formulas are *used* in proofs). For instance, if the use of a semantic formula in a proof is easily eliminable, this might take away from the level of semantic pollution. However, this encourages a minimal view of semantic pollution, as there often exist many translations between proof systems, where 'semantic' properties may be lost. We prefer to consider proof systems individually, and to evaluate their design. And perhaps especially when the uses of a 'semantic' formula type in a proof system are easily eliminable, the proof system should be considered semantically polluted: why should one introduce semantic notions into a language, if it is not even necessary?[8]

[8] See for more comments on this topic Section 5.2.

3 The base requirement: violating invariance results under model equivalences

This section will define the base requirement of semantically polluted formula types. By itself, this property provides enough information to conclude that a formula type is semantically polluted, and it specifies a level of pollution with respect to it. However, the next section will introduce two properties on top of the one defined here, that more clearly divide the different formula types among four levels of semantic pollution. The idea of the base requirement comes in when we consider the connection that the (classical) basic modal language L displays towards Kripke semantics as a syntactically pure baseline. That is, we can view the way that the basic modal language distinguishes Kripke models, by describing Kripke models with a certain level of detail, as a syntactically pure standard. The basic modal language partitions the space of pointed Kripke models based on what it can express about these models, by equating pairs (M, w) that it considers to be the same (i.e., that cannot be distinguished by formulas of L).

Formulas of a proof-theoretic language extending L may then be found to make *more* distinctions between pointed Kripke models than L, and create a more fine-grained partition. This tells us that the formula can express differences between two (modally equivalent) worlds in two models that no modal formula can. It can describe a Kripke model in a way that is unavailable to the modal language — and in this sense, the formula has a stronger connection to the semantics. We can check this by considering whether a proof-theoretic formula violates invariance results under Kripke model equivalences of L. As mentioned before, note that this is a semantic perspective and elaboration of the suggestion proposed in (Poggiolesi, 2010) that semantic pollution comes down to untranslatability to an object language. Namely, violating invariance results under model equivalences for L, is a way of establishing untranslatability to L. We will now formally define the property of violating invariance results under Kripke model equivalences.

3.1 Levels of satisfying the base requirement

There are three aspects affecting the level of satisfying the base requirement for formula types from HL, two for the formula types from LL and one for the formula type of DL.

Model equivalence (DL, LL, HL). This concerns baseline notions of equivalence for Kripke models. Generally, it is most difficult for a formula type to violate invariance results under model equivalences that reduce many models to each other, and this will indicate a higher level of pollution. We will use the symbol $\equiv$ as covering two notions of equivalence defined in (Blackburn, De Rijke, and Venema, 2001). That is, $\mathsf{M}, w \equiv \mathsf{M}', w'$ means either:

1. (M, w) is *isomorphic* to (M', w'), i.e. $\mathsf{M}, w \cong \mathsf{M}', w'$
2. (M, w) is *bisimilar* to (M', w'), i.e. $\mathsf{M}, w \leftrightarrows \mathsf{M}', w'$

Note that although we can make more distinctions by picking more notions of equivalence in between isomorphisms and bisimulations (such as generated submodels and disjoint unions), we believe too many equivalences will only obscure the results, and these two extremes already form a suitable representation of available equivalences and lead to variable results. Section 3.2 will make these notions precise.

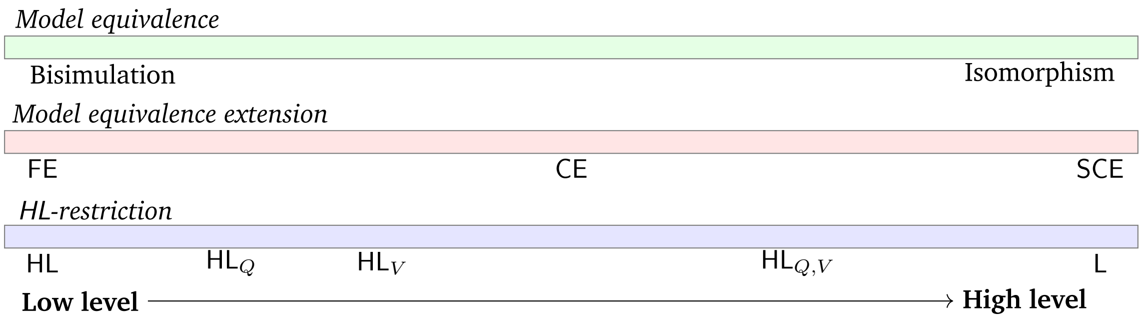

Figure 2: Three aspects affecting the level of satisfying the base requirement. 'Model equivalence' and 'model equivalence extension' apply to LL, DL and HL, while 'HL-restriction' applies only to HL.

Model equivalence extension (LL, HL). This concerns the extension of the two model equivalences to models with an assignment function. Three 'equivalence strengths' will be defined for extended models (M, τ), corresponding to the following notation:

1. Free extended (FE-)equivalence $\equiv_{\mathsf{FE}}$.
2. Constrained extended (CE-)equivalence $\equiv_{\mathsf{CE}}$.
3. Strongly constrained extended (SCE-)equivalence $\equiv_{\mathsf{SCE}}$.

It is more difficult for a formula type to violate invariance results under a stronger equivalence for extended models, which will indicate a higher level of pollution. We will use the symbol $\equiv_{\mathsf{P}}$ to cover the four equivalences defined above, combined with the three possible strengths just mentioned. Section 3.3 will make these notions precise.

Size of the context language (HL). It is most difficult for an HL-formula type to violate invariance results relative to a small context language (see Figure 1), and so this will indicate a higher level of pollution. Note that this does not apply to LL and DL.

Figure 2 visualizes the effect of these three properties on satisfying the base requirement. Now we will see the model equivalences and model equivalence extensions defined in detail.

3.2 Equivalences between regular Kripke models

Here we introduce isomorphisms and bisimulations for regular pointed Kripke models (M, w), after their definition in (Blackburn, De Rijke, and Venema, 2001). Consider isomorphisms, which reduce fewest pointed models to each other.

Definition 3 (Isomorphism). Two models $\mathsf{M} = (W, R, V), \mathsf{M}' = (W', R', V')$ are *isomorphic* if there is a function $f : M \to M'$ such that: $w \in V(p)$ iff $f(w) \in V'(p)$; wRv iff $f(w)R'f(v)$; and f is a bijection. Pairs $(w, f(w))$ indicate isomorphic worlds $(\mathsf{M}, w \cong \mathsf{M}', f(w))$.

If (M, w) and (M', w') are isomorphic, the frame structure and valuation of (M, w) and (M', w') are entirely alike, so that violating invariance results can only happen if a formula does something independent from the local frame and valuation structure that surrounds w and w'. For instance, unconstrained assignment functions τ can scan different parts of the same model, leading to different truth values of formulas including labels or nominals.

Although we do not include disjoint unions and generated submodels, note that potential reasons for violating invariance under one of these notions increase, as they reduce more models to each other. In disjoint unions, formulas that say something about the global model situation can additionally cause violation of invariance, since a disjoint union adds an entire model to an initial one, and a formula may detect this. In generated submodels, a formula may furthermore detect that local worlds preceding equivalent worlds disappear.

Both aspects are also captured by bisimulations. Additionally, a formula may violate invariance under bisimulations because it detects differences in the number of local successors of equivalent worlds, which can vary under bisimulation. Thus, bisimulations capture more types of semantic pollution than isomorphisms, which are more strict in what counts as pollution.

Definition 4 (Bisimulation). Let $\mathsf{M} = (W, R, V)$ and $\mathsf{M}' = (W', R', V')$. Then $Z \subseteq W \times W'$ is a bisimulation if: if wZw' then w and w' satisfy the same propositional letters; if wZw' and Rwv, then there exists a $v' \in M'$ such that vZv' and $Rw'v'$ *(forth)*; and if wZw' and $R'w'v'$, then there exists a $v \in M$ such that vZv' and Rwv *(back)*. Pairs (w, w') related by Z indicate bisimilar worlds $(\mathsf{M}, w \leftrightarrows \mathsf{M}', w')$.

3.3 Model equivalences for extended Kripke models

As mentioned, $\equiv_{\mathsf{P}}$ denotes one of three strengths with which to adapt isomorphisms and bisimulations to extended pointed models (M, τ, w). Consider the weakest extended equivalence.

Definition 5 (Free extended (FE-)equivalence). Two extended pointed models (M, τ, w) and (M', τ', w') are *free extended (FE-)equivalent* if their underlying pointed Kripke models are equivalent:

$$\mathsf{M}, \tau, w \equiv_{\mathsf{FE}} \mathsf{M}', \tau', w' \text{ iff } \mathsf{M}, w \equiv \mathsf{M}', w'$$

'Free' indicates that the extended equivalence poses no requirements on τ. This notion extends a regular equivalence simply by adding assignment functions on top of M and M'. It does not matter what these functions look like: any pair of functions τ, τ' added to M and M' will lead to the equivalence between (M, τ) and (M', τ'). This may make violations of invariance results by formulas including name variables rather unsurprising, since τ and τ' can map name variables to very different states in the equivalent models. Lack of surprise does not indicate lack of value, however: such results show exactly that τ scans a Kripke model in a way that is foreign to the modal language, a phenomenon that we aim for semantic pollution to capture.

We now define two ways of placing more restrictions on τ. A weak one connects the object language formulas satisfied at worlds $\tau(a)$ and $\tau'(a)$, and a strong one explicitly aligns equivalent worlds with the name variables they are assigned. They will indicate a higher level of semantic pollution, as a formula violating invariance results under these equivalences distinguishes even more Kripke models than formulas only violating invariance results under FE-equivalences.

Definition 6 (Constrained extended (CE-)equivalence). Two extended pointed models (M, τ, w) and (M', τ', w') are *constrained extended (CE-)equivalent* if their underlying pointed Kripke models are equivalent, and if object language formulas are invariant under name variable assignments.

$\mathsf{M}, \tau, w \equiv_{\mathsf{CE}} \mathsf{M}', \tau', w'$ if and only if:

1. $\mathsf{M}, w \equiv \mathsf{M}', w'$

2. For all name variables x and $A \in \mathsf{L}$: $\mathsf{M}, \tau(x) \models A$ iff $\mathsf{M}', \tau'(x) \models A$

'Constrained' indicates that the extended equivalence poses some requirements on τ. It still allows a name variable a to be assigned to non-equivalent worlds. The strongly constrained extended equivalence will not allow this; it is based on *hybrid bisimulations* (Blackburn, Benthem, and Wolter, 2006) (which is intended to provide an invariance result for the basic modal language extended with the satisfaction operator and nominals).

Definition 7 (Strongly constrained extended (SCE-)equivalence). Two extended pointed models (M, τ, w) and (M', τ', w') are *strongly constrained extended (SCE-)equivalent* if all states that are assigned a name variable are related by the equivalence, and if equivalent states are assigned the same name variables.

$\mathsf{M}, \tau, w \equiv_{\mathsf{SCE}} \mathsf{M}', \tau', w'$ if and only if:

1. $\mathsf{M}, w \equiv \mathsf{M}', w'$

2. For all name variables x, $\mathsf{M}, \tau(x) \equiv \mathsf{M}', \tau'(x)$

3. For all name variables x:

   (a) There is a unique v' such that $\mathsf{M}, \tau(x) \equiv \mathsf{M}', v'$, and

   (b) There is a unique v such that $\mathsf{M}, v \equiv \mathsf{M}', \tau'(x)$[9]

'Strongly constrained' indicates that the extended equivalence poses strong requirements on τ. Replacing $\equiv$ in these three definitions of FE-, CE- and SCE-equivalence by one of the equivalences from Section 3.2 leads to twelve notions of equivalence between extended models.

Note that strongly constrained extended equivalences require (by the second criterion) that the range of the functions τ, τ' is a subset of the states related by the equivalence. For isomorphisms, this is not a problem, because all states of the isomorphic models are by definition already related by the equivalence. For bisimulations, not all states in the model need to be bisimilar to some other state: this just means that for SCE-bisimulations, the range of the functions τ, τ' possibly needs to be restricted to a subset of all states (similar restrictions exist in case one wants to create SCE-disjoint unions or generated submodels).

3.4 Base requirement for semantic pollution

The work of the previous sections now comes together in the base criterion that any semantically polluted formula type will need to satisfy.

Definition 8 (Base requirement (BR)). Let C be an operator or metavariable of PL, and let CL be the *context language* of C. Let $\equiv_{\mathsf{P}}$ consist of an extended Kripke model equivalence with the following ingredients:

- A regular equivalence $E \in \{\cong, \leftrightarrows\}$ (as in Section 3.2)

- If PL equals LL or HL, a model equivalence extension $S \in \{\mathsf{FE}, \mathsf{CE}, \mathsf{SCE}\}$ (as in Section 3.3)

[9] Thus, the τ- and τ'-ranges of name variables are restricted to being isomorphic. Note that other equivalences can still also have an SCE-version, by the states in M and M′ that are not assigned any name variables.

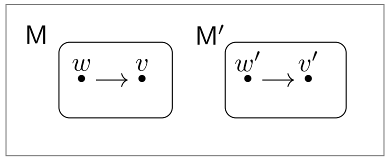

Figure 3: Example of the base requirement.

Then C satisfies $\mathsf{BR}_{\mathsf{CL},E,S}$, i.e., the *base requirement of semantic pollution relative to context language CL, equivalence E and strength S* if it violates invariance under $\equiv_{\mathsf{P}}$:

There are equivalent pointed models $\mathsf{PM}, w \equiv_{\mathsf{P}} \mathsf{PM}', w'$, and an $A \in C(\mathsf{CL})$ such that

1. $\mathsf{PM}, w \vDash A$
2. $\mathsf{PM}', w' \nvDash A$

This requirement gives us a level of pollution with respect to the aspects shown in Figure 2. We will discuss the results of BR with respect to the formula types of DL, LL and HL in the next section, using the variants for pollution (po) and purity (pu) as described in Section 2.3. In order to tear apart the formula types even more clearly, we will define more overarching levels of semantic pollution in Section 4. First, a small example of BR for an intuitive sense of it.

Example 3.1 (Example of the base requirement). Take the operator $@_a$ of HL, and suppose we want to show that it satisfies $\mathsf{BR}_{\mathsf{L},\cong,\mathsf{FE}}$. Figure 3 shows two models (M, τ) and (M', τ'), where we let $V(p) = w$, $V'(p) = w'$, $\tau(a) = w$, $\tau'(a) = v'$. Then (M, τ, w) and (M', τ', w') are FE-isomorphic, yet $\mathsf{M}, \tau, w \vDash @_a p$ and $\mathsf{M}', \tau', w' \nvDash @_a p$. Hence, $@_a(\mathsf{L})$ satisfies $\mathsf{BR}_{\mathsf{L},\cong,\mathsf{FE}}$.

Finally, we note that the base requirement is theoretically applicable to any other combination of object language and model-theoretic semantics. However, the details of making the requirement precise depend too much on the particular context for a useful generalization at this point.

3.5 Results

We treat the results per language DL, LL and HL. Most examples we provide are well-known and can already be found elsewhere, for instance in (Blackburn, De Rijke, and Venema, 2001), but we discuss them here for the first time within the context of semantic pollution.

Results of DL. From DL, we see that only $\bullet A$ as interpreted in the antecedent position of a sequent (see Section 2.1) satisfies the base requirement. By their translatability to L, $\mathbf{I}$, $A \circ A$ and $*A$ stay syntactically pure (see also Remark 2).

$\bullet A$ gives us simple results (see Table 1). Results for extended model equivalences do not apply to it, so we only consider the usual variants of Kripke model equivalences. There, we see that it only violates invariance under bisimulations, simply because it functions like a backwards diamond. This also indicates only a medium level of semantic pollution, as isomorphisms preserve its truth value, and they are harder to violate invariance under (note that a similar pattern holds if we included generated submodels and disjoint unions — $\bullet A$ is polluted with respect to the former, but pure with respect to the latter).

Equivalence	$\bullet A$
Isomorphism	pu
Bisimulation	po

Table 1: Base requirement results of the bullet operator.

Model equivalence extension	$x : A$	xRy
Free extended	po	po
Constrained extended	pu	po
Strongly constrained extended	pu	pu

Table 2: Base requirement results of LL-operators (results are the same for isomorphisms and bisimulations of the same model equivalence extension).

Results of LL**.** Consider now labeled formulas and relational atoms, that also give us rather clear-cut results (see Table 2). Interestingly, their semantic pollution is indifferent to the type of Kripke model equivalence chosen (bisimulation or isomorphism), for all model equivalence extensions. This indicates that τ rises above the differences in model properties that the modal language cannot see. Consider the following example for FE-isomorphisms (and so also bisimulations).

Example 3.2 ($x : p$, xRy)**.** Consider the models in Figure 3, and let $V(p) = \{w\}$, $V'(p) = \{w'\}$. Let $\tau(x) = w$, $\tau'(x) = v'$, $\tau(y) = v$ and $\tau'(y) = v'$. Then (M, τ, w) is FE-isomorphic (and observe, not CE- or SCE-isomorphic) to (M, τ', w'). However, $\mathsf{M}, \tau, w \vDash x : p$, $\mathsf{M}', \tau', w', \nvDash x : p$, and $\mathsf{M}, \tau, w \vDash xRy$, $\mathsf{M}', \tau', w', \nvDash xRy$.

Stronger model equivalence extensions reduce pollution, and also tear apart $x : A$ and xRy, as relational atoms are more semantically polluted than labeled formulas. Of course, this is by design of the model equivalence extensions: CE-equivalences purify $x : A$ by definition, while they do not yet tie down xRy (just edit the previous example by giving all worlds valuation $\{p\}$ — just because each world satisfies the same modal formulas, does not mean they have the same R-context).

Only SCE-equivalences tame xRy. To see this, suppose that $\mathsf{M}, \tau, w \vDash xRy$, and suppose there is a strong equivalence $\mathsf{M}, \tau, w \equiv \mathsf{M}', \tau', w'$. By assumption, $\tau(x)R\tau(y)$, and we know that the entire range of τ is included in the equivalence. Hence, there are worlds s' and t' in M' such that $\tau(x) \equiv s'$ and $\tau(y) \equiv t'$. We also know that equivalent worlds must satisfy the same name variables, so $s' = \tau'(x)$ and $t' = \tau'(y)$. Now to see that $s'Rt'$, suppose that our equivalence is a bisimulation (for isomorphisms, $s'Rt'$ is clear). By the forward condition, there exists $u' \in \mathsf{M}'$ such that $s'Ru'$ and $\tau(y) \equiv u'$. But as equivalent worlds must satisfy the same labels, $u' = \tau'(y)$, and so $u' = t'$, so that $s'Rt'$ (and $\mathsf{M}', \tau', w' \vDash xRy$).

Results of HL**.** The results of HL are a little more intricate. First consider the simple results of a as in Table 3, and note that it has the same results as xRy. The following example illustrates the cases of pollution.

Example 3.3 (a)**.** Take Figure 3 and let all worlds have valuation $\{p\}$. Let $\tau(a) = w$ and $\tau'(a) = v'$. Then (M, τ, w) and (M', τ', w') are FE- and CE-equivalent. Yet $\mathsf{M}, \tau, w \vDash a$ and $\mathsf{M}', \tau', w' \nvDash a$.

Clearly, strongly constrained extended equivalences will purify a by definition, because of the strict demand that equivalent worlds are assigned the same name variables.

Model equivalence extension	a
Free extended	po
Constrained extended	po
Strongly constrained extended	pu

Table 3: Base requirement results of a (results are the same for isomorphisms and bisimulations of the same model equivalence extension).

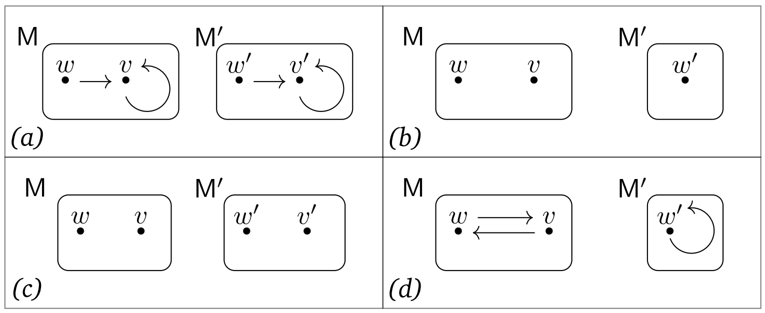

Figure 4: The extended models serving as proof for satisfying the base requirement. Assume that all worlds receive valuation $\{p\}$.

We then switch to the more graded results of the operators $@_a$, $\forall a$ and $\downarrow a$. Recall that we consider them with respect to several language fragments. Note from Table 5 that $@_a : A$ shows more pollution than $x : A$, as it retains pollution in a context language more powerful than L. The difference is seen in CE-equivalences, where $x : A$ is immediately turned pure, but $@_a A$ finds pollution even in the rather restricted language $\mathsf{HL}_{Q,V}$: consider the example below.

Example 3.4 ($@_a \Diamond a$)**.** Consider the pair of models in Figure 4(a). Let $\tau(a) = w$ and $\tau'(a) = v'$. Then (M, τ, v) is CE-isomorphic to (M', τ', v'). Yet $M, \tau, v \not\models @_a \Diamond a$, and $M', \tau', v' \models @_a \Diamond a$.

In a strongly constrained extended isomorphism (see Table 6), $@_a$ is finally turned pure. As Theorem 5 shows, in fact, all formulas in HL are invariant under SCE-isomorphisms.

But SCE-bisimulations still show a low level of pollution for $@_a A$ relative to HL_V, by interacting with $\forall a A$. Consider the example below.

Example 3.5 ($@_a(\forall a(a))$)**.** Consider the pair of models in Figure 4(b). Let $\tau(a) = w$ and $\tau'(a) = w'$. Then (M', τ', w') is SCE-bisimilar (M, τ, w). Yet $\mathsf{M}, \tau, w \not\models @_a \forall a(a)$, while $\mathsf{M}', \tau', w' \models @_a \forall a(a)$.

To see that for bisimulations, this is the highest level of pollution that $@_a A$ can have, Theorem 6 shows that all hybrid formulas in HL_Q are invariant under SCE-bisimulations. Thus, any smaller language than HL_V will purify $@_a A$.

Next, consider the operator $\forall a A$ separately. Just like in the example above, $\forall a$ is able to make restricted cardinality statements in a small language, and so is easily able to continuously satisfy the base requirement for non-isomorphism equivalences. An easy example works for bisimulations of all three strengths (in the language $\mathsf{HL}_{Q,V}$, accounting for the results in Table 4, 5, 6): simply take Example 3.5 for just the formula $\forall a(a)$. This is why the level of SCE-pollution of $@_a$ relative to HL_V has to be seen as rather low (it is mainly 'caused by' $\forall a A$). Note also that these results for

FE-**equivalence**	$@_a A$	$\forall a A$	$\downarrow a A$
Isomorphism	po	po(HL_Q)	po(HL_Q)
Bisimulation	po	po($\mathsf{HL}_{Q,V}$)	po($\mathsf{HL}_{Q,V}$)

Table 4: Base requirement results of HL-operators for free extended equivalences.

CE-**equivalence**	$@_a A$	$\forall a A$	$\downarrow a A$
Isomorphism	po($\mathsf{HL}_{Q,V}$)	po(HL_Q)	po(HL_Q)
Bisimulation	po($\mathsf{HL}_{Q,V}$)	po($\mathsf{HL}_{Q,V}$)	po($\mathsf{HL}_{Q,V}$)

Table 5: Base requirement results of HL-operators for constrained extended equivalences.

bisimulations indicate the highest level of pollution that $\forall a A$ can achieve, as it is invariant under L (the only smaller language than $\mathsf{HL}_{Q,V}$).

For isomorphisms, there is an example in the restricted language HL_Q that still works for $\forall a A$ in the FE and CE case (see Table 4, 5).

Example 3.6 ($\forall a(a \vee b)$)**.** Consider the pair of models in Figure 4(c). Let $\tau(a) = w$, $\tau(b) = w$, $\tau'(a) = w'$, $\tau'(b) = v'$. Then (M, τ, w) and (M', τ', w') are both FE- and CE-isomorphic. But $\mathsf{M}, \tau, w \models \forall a(a \vee b)$, while $\mathsf{M}', \tau', w' \not\models \forall a(a \vee b)$.

The fact that po(HL_Q) is the highest level of pollution of $\forall a A$ for FE-isomorphisms, is shown by Theorem 4 (showing that $\forall a A$ is invariant under FE-isomorphisms relative to HL_V). This result extends to CE-isomorphisms, which only pose more requirements on τ. However, as mentioned before, Theorem 5 shows that SCE-isomorphisms purify $\forall a A$ (see Table 6).

Finally, consider $\downarrow a A$. By resorting to reflexivity statements, it can satisfy the base requirement relative to FE- and CE-bisimulations and $\mathsf{HL}_{Q,V}$ (Table 4, 5), as in the following example.

Example 3.7 ($\downarrow a(\Diamond a)$)**.** Consider the pair of models in Figure 4(d). Let $\tau(a) = w$ and $\tau'(a) = w'$. Then (M, τ, w) and (M', τ', w') are FE- and CE-bisimilar. Yet $\mathsf{M}, \tau, w \not\models \downarrow a(\Diamond a)$, while $\mathsf{M}', \tau', w' \models \downarrow a(\Diamond a)$.

Just like for $\forall a A$, note that po($\mathsf{HL}_{Q,V}$) is the highest level of pollution for these categories, as $\downarrow a A$ is pure relative to L (the only smaller language than $\mathsf{HL}_{Q,V}$). For FE- and CE-isomorphisms, the following polluting example of $\downarrow a A$ is taken from HL_Q, a less restricted language, and so less polluting (Table 4, 5).

Example 3.8 ($\downarrow a(@_b a)$)**.** Consider the pair of models in Figure 4(c). Let $\tau(a) = w$, $\tau'(a) = w'$, $\tau(b) = w$, and $\tau'(b) = v'$. Then (M, τ, w) and (M', τ', w') are FE- and CE-isomorphic, yet $\mathsf{M}, \tau, w \models \downarrow a(@_b a)$, while $\mathsf{M}', \tau', w' \not\models \downarrow a(@_b a)$.

SCE-**equivalence**	$@_a A$	$\forall a A$	$\downarrow a A$
Isomorphism	pu	pu	pu
Bisimulation	po(HL_V)	po($\mathsf{HL}_{Q,V}$)	po(HL_V)

Table 6: Base requirement results of HL-operators for strongly constrained extended equivalences.

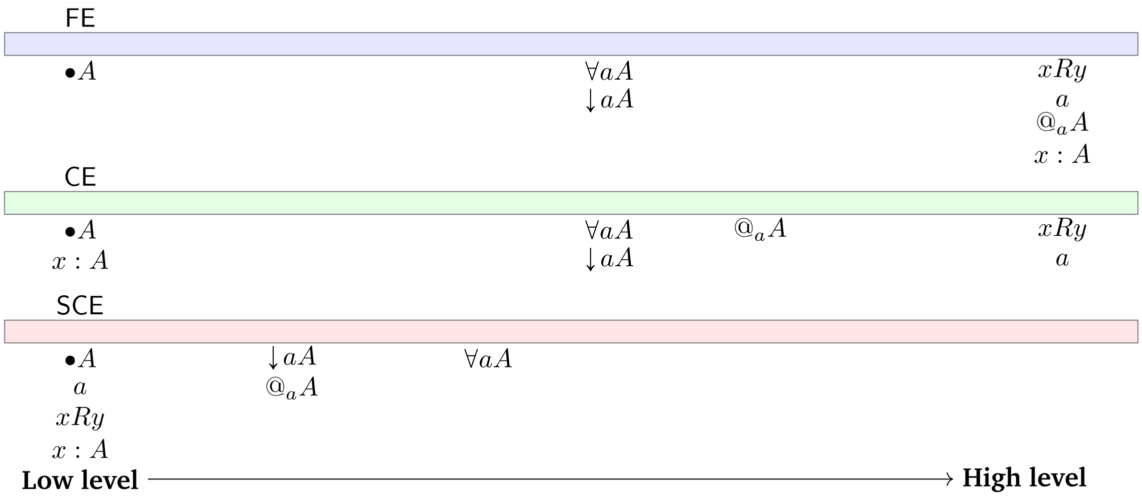

Figure 5: Summary of the results for (levels of) satisfaction of BR.[10]

To see that this $\mathsf{po}(\mathsf{HL}_Q)$ result is the highest level of pollution for FE-isomorphisms, see Theorem 3 (showing invariance of $\downarrow aA$ under FE-isomorphisms relative to HL_V). Again, this result extends to CE-isomorphisms. Finally, an example using $\forall aA$ comes back in, and gives a low pollution result for SCE-bisimulations relative to HL_V (Table 6).

Example 3.9 ($\downarrow a(\forall a(a))$)**.** Consider the pair of models in Figure 4(b). Let $\tau(a) = w$ and $\tau'(a) = w'$. Then (M, τ, w) and (M', τ', w') are SCE-bisimilar, yet $\mathsf{M}, \tau, w \not\models \downarrow a(\forall a(a))$, while $\mathsf{M}', \tau', w' \models \downarrow a(\forall a(a))$.

Theorem 6 then shows that (among others) $\downarrow aA$ is invariant under SCE-bisimulations relative to HL_Q, implying that $\mathsf{po}(\mathsf{HL}_V)$ is indeed the highest pollution level here for $\forall aA$. And once more, the purity result of $\forall aA$ for SCE-isomorphisms relative to HL is shown by Theorem 5.

3.6 Summing up

The base requirement for semantic pollution can be satisfied with various levels, depending on model equivalence, model equivalence extension (if applicable) and context language restriction (if applicable). These aspects already create a division in satisfaction of the base requirement for formula types in LL, DL and HL. The results are summarized in Figure 5. The lowest level of satisfaction of the base requirement is steadily provided by $\bullet A$. The highest level is displayed by xRy and a, which are only 'purified' for strongly constrained extended equivalences, showing the independence of τ from the basic modal language.

The level of semantic pollution of $@_a A$ is reduced significantly by strong equivalences, but never completely eliminated (unlike $x : A$), due to its interaction with other formulas in HL. The level of BR-satisfaction for $\downarrow aA$ and $\forall aA$ is relatively low, yet (especially for $\forall aA$) remains persistent

[10]Two formula types have a different level of pollution in the figure, if this difference exists with respect to isomorphisms, bisimulations, or both. Additionally, although $\bullet A$ does not need FE-, CE- and SCE-equivalences, its results can be seen as unchanged with respect to these notions.

throughout the different model equivalence extensions. Only for SCE-equivalences does $\downarrow aA$ get similar results to $@_a A$, and $\forall aA$ retains the highest level of pollution there. However, although low pollution levels remain with SCE-equivalences for $@_a A$, $\downarrow aA$ and $\forall aA$, note that these all rely on the workings of $\forall aA$. That is, they rely on the ability of $\forall$ to distinguish between a model cardinality of one and more than one. Clearly, we should emphasize that for $@_a A$ and $\downarrow aA$ this is a side effect of the context language HL_V, and this is a low level of BR-satisfaction. For $\forall aA$ itself, the property is more inherent and reflective of a proper remaining level of semantic pollution (relative to the minimal context language $\mathsf{HL}_{Q,V}$).

Furthermore, we see that the BR-results of $\downarrow aA$ for SCE-bisimulations hold with respect to a larger context language than those of $\forall aA$. This shows that sending a to the current world (as $\downarrow aA$ does) is something less invasive than sending it to *all* worlds (as $\forall aA$ does). Both operators change the assignment function, but $\downarrow aA$ changes it only by using local information, while $\forall aA$ changes it by using global information. Note that both formulas are still syntactically pure for isomorphisms. Thus, there must be some 'actual' difference between two equivalent models for $\forall aA$ and $\downarrow aA$ to pick up on. This is unlike more direct uses of the assignment function τ as in xRy, $x : A$ or a, which can more easily vary even for very similar models, but are also more easily be constrained again by stronger equivalences.

So far, we provided a general distinction between levels of BR-satisfaction. In the next section, we will highlight two aspects that should, according to us, count more strongly in determining the exact level of pollution. This provides a more intuitive division into various types of semantic pollution, and emphasizes the differences that the base requirement cannot capture.

4 Four levels of semantic pollution

The base requirement for semantic pollution is insightful, and distinguishes between some degrees of semantic pollution. However, we believe that the emphasis of two properties (underlying some of the results in the base requirement) should give rise to stronger forms of semantic pollution. They can be interpreted as being 'less modal' and 'more semantic' in two respects, and thus form more of an unnatural invasion into the modal language than lower forms of semantic pollution.

Globalness. The first is the property of globalness, or world invariance. It is well-known that "[m]odal satisfaction is intrinsically local: only the points accessible from the current state are relevant to truth or falsity" (Blackburn, De Rijke, and Venema, 2001). Another way to view locality is to recognize that the truth value of a modal formula can change in a model depending on the world of evaluation: if the latter changes, the context of accessible points may change as well. The property of globalness then becomes an 'unmodal' property, and has two corresponding conceptions: the satisfaction of a formula is global if points inaccessible from the current state are relevant to truth or falsity; or if its truth value is the same in a model for all worlds of evaluation. A well-known example of global formulas where these conceptions overlap is the *global modality* (Blackburn, De Rijke, and Venema, 2001; ten Cate, 2004):

Global diamond E: $\mathsf{M}, w \vDash \mathsf{E}B$ iff $\mathsf{M}, v \vDash B$ for *some* state v in M

Global box A: $\mathsf{M}, w \vDash \mathsf{A}B$ iff $\mathsf{M}, v \vDash B$ for *all* states v in M

The idea that inaccessible states affect the truth value of a formula can be made precise by the notion of violating invariance results under disjoint unions. Note that this is an instance of

BR, even though we only treated isomorphisms and bisimulations in Section 3 (namely, by taking as model equivalence not bisimulation or isomorphism, but disjoint union as in Definition 11.). Instead, world invariance as a measure of globalness is stricter than BR. We will show soon that BR is indeed implied by world invariance[11], and we take the latter as our conception of globalness. A border case that is considered local by this conception is $\forall a A$, that intuitively possesses some type of globalness. We will discuss this formula more in Section 4.2.

Besides straying from an intrinsically modal nature, global formulas may be seen as more semantic than local formulas. Local formulas, only taking into account accessible states, are simply blind to a certain part of a Kripke model. Global formulas, when their truth value depends on inaccessible states, can collect more information about the model — and by the property of world invariance, they lift information up to a state that the entire model finds itself in. That is, instead of forming truth relations with each world separately, these formulas provide a truth state for all worlds in a model at the same time.

Valuation independence. The second is the property of valuation independence. Generally, a modal formula "is valid on a frame when it is globally true, no matter what valuation is used. This concept allows modal languages to be viewed as languages for describing frames" (Blackburn, De Rijke, and Venema, 2001). This quote connects both globalness and valuation independence to semantic properties, but we here focus on valuation independence separately. The semantic nature of such formulas is then still best seen when first restricting to the modal language. Valuation independent modal formulas can be seen to describe the local frame structure of the world of evaluation w, as shown by the simple example $\Diamond\top$ ('w has a successor'). However, they are still instances of formula types that we do not regard as semantically polluted, because operators like $\Diamond$ are primarily intended to capture valuation dependent statements, and they clearly do not satisfy BR. On the other hand, if an operator introduced in the proof-theoretic language can *only* convey valuation independent information, then we consider it semantically polluted. In case such a formula is translatable to the modal language, its semantic nature is strengthened by the fact that it will be describing the frame. There can also be valuation independent formulas that are untranslatable to the modal language. They can still describe frame structure (such as a formula that is true when the world of evaluation has exactly two successors), but they can also concern other properties of worlds (concerning, for instance, an assignment function of an extended model). A more general view on the semantic nature of valuation independence then says that they have a stronger connection to the world of evaluation than the modal language.

We can thus consider formula types satisfying one of these properties as carrying a different type of semantic pollution as formula types just satisfying the base requirement. Furthermore, formula types satisfying both globalness and valuation independence can be thought of as possessing the strongest form of semantic pollution. The gradations of satisfying the base requirement can still show differences in semantic pollution within these four new categories.

4.1 Defining the levels

The properties of globalness and valuation independence are made precise by the following definition.

[11] Combined with valuation dependence or contingency.

Definition 9 (**Modal semantic properties**). Let C be an operator or metavariable of PL. The modal semantic properties are then defined with respect to PL.[12] The notion of globalness gives rise to two variants concerning the *relation of C to states*.

1. **Locality (LO).** There is an $A \in C(\mathsf{PL})$, a model PM and worlds w_1 and w_2 such that $\mathsf{PM}, w_1 \models A$ and $\mathsf{PM}, w_2 \not\models A$.

2. **Globalness (GL).** For each model PM and for all $A \in C(\mathsf{PL})$ it is the case that either:

   (a) $\mathsf{PM}, w \models A$ for all $w \in W$ (A is *globally true in PM*), or

   (b) $\mathsf{PM}, w \not\models A$ for all $w \in W$ (A is *globally false in PM*)

The second property, also with two variants, concerns the *relation of C to the valuation.*

1. **Valuation dependence (VD).** There is an $A \in C(\mathsf{PL})$, a pointed frame (PF, w), and two models $\mathsf{PM}, \mathsf{PM}'$ extending PF such that $\mathsf{PM}, w \models A$ and $\mathsf{PM}', w \not\models A$.

2. **Valuation independence (VI).** For each pointed frame (PF, w), and for all $A \in C(\mathsf{PL})$, it is the case that either:

   (a) $\mathsf{PM}, w \models A$ for all models PM over PF (we say A is *w-valid on PF*), or

   (b) $\mathsf{PM}, w \not\models A$ for all models PM over PF (we say A is a *w-contradiction on PF*)

As these properties will be imposed on top of the base requirement, they only function to further classify formula types that are already untranslatable to the modal language. Thus, now we can define the following four levels of semantic pollution (see Figure 6).

Definition 10 (**Levels of pollution**). Let C be an operator or metavariable of PL, and let BR be defined with respect to context language CL, model equivalence E and (if applicable) extended equivalence strength S. Then four levels of pollution are defined as follows.

1. C satisfies **weak semantic pollution** if it satisfies $\mathsf{BR}_{\mathsf{CL},E,S}$ + LO + VD.[13]

2. C satisfies **local semantic pollution** if it satisfies $\mathsf{BR}_{\mathsf{CL},E,S}$ + LO + VI.

3. C satisfies **global semantic pollution** if it satisfies $\mathsf{BR}_{\mathsf{CL},E,S}$ + GL + VD.

4. C satisfies **strong semantic pollution** if it satisfies $\mathsf{BR}_{\mathsf{CL},E,S}$ + GL + VI.

Blackburn, De Rijke, and Venema (2001) provide a simple proof that global diamond and global box are undefinable in the basic modal language. Similarly, we can provide an untranslatability result for the two properties of global semantic pollution (so that BR is actually a superfluous requirement here). Consider the definition of a disjoint union, followed by the theorem.

[12] This is a choice: they can also be defined with respect to L, for instance with our earlier motivation of capturing only the effects of C, and not interactions with other operators. However, for our operators, the results relative to the different available context languages remain the same. Additionally, as mentioned before, $\downarrow aA$ and $\forall aA$ simply act like A when $A \in \mathsf{L}$, so in order to see some more interesting examples highlighting the workings of $\forall$ and $\downarrow$, it is insightful to let $A \in \mathsf{HL}$, and so to let the context language generally be PL.

[13] Note that here, the only reason C is semantically polluted is that it satisfies some variant of BR.

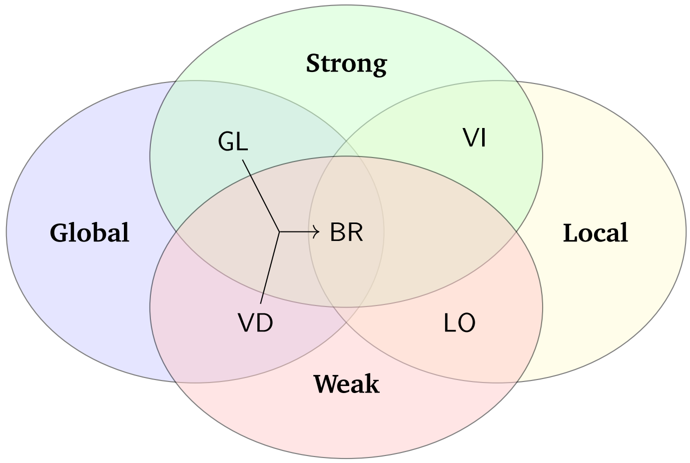

Figure 6: Levels of semantic pollution (the properties of GL together with VD imply the base requirement).

Definition 11 (**Disjoint union**). Two models are disjoint if their domains contain no common elements. For disjoint models $\mathsf{M}_i = (W_i, R_i, V_i)$ $(i \in I)$, their disjoint union is the structure $\biguplus_i \mathsf{M}_i = (W, R, V)$, where W is the union of the sets W_i, R is the union of the relations R_i, and for each proposition p, $V(p) = \bigcup_{i \in I} V_i(p)$. For each M_i, a world $w \in W_i$ is equivalent to its copy in $\biguplus_i M_i$.

Theorem 1. *If an operator or metavariable C of PL satisfies globalness and valuation dependence relative to CL, then it satisfies BR[CL, $\biguplus$, FE/CE].*[14]

Proof. Suppose that C satisfies globalness and valuation dependence relative to CL. Then there is an instance $A \in C(\mathsf{CL})$ and two models PM and PM′, such that A is globally true in PM and globally false in PM′. Now take an FE or CE- disjoint union PM $\biguplus$ PM′, which will be equivalent to both PM and PM′ (separately).[15] Since A is invariant under worlds, one of the following holds:

1. $\mathsf{PM} \biguplus \mathsf{PM}', w \vDash A$ for all $w \in W \cup W'$
2. $\mathsf{PM} \biguplus \mathsf{PM}', w \nvDash A$ for all $w \in W \cup W'$

Suppose wlog that the first case holds. Then there exists a world $w \in W'$ such that $\mathsf{PM}', w \nvDash A$, and such that $\mathsf{PM} \biguplus \mathsf{PM}', w \vDash A$. Thus, A violates invariance under disjoint unions. □

As this theorem gives the BR-result for disjoint unions, and disjoint unions are a specific type of bisimulation, the result is implied for bisimulations as well. However, it is not implied for isomorphisms. A counterexample is a global modality, for instance Ap or Ep, which are global and valuation dependent, yet invariant under isomorphisms.

Now we provide a few remarks concerned with the difference between formula types $C(\mathsf{PL})$ and

[14] Note that if we require that some $A \in C(\mathsf{PL})$ is contingent, then globalness by itself already implies BR. Without contingency, $\top$ satisfies globalness, and clearly does not satisfy BR (hence, we need valuation dependence).

[15] On the contrary, SCE-disjoint unions are only equivalent to one of PM or PM′, which prevents the proof from going through. As an example, consider the instance $@_a p$ of the global and valuation dependent $@_a(\mathsf{L})$. Under an SCE-disjoint union it will retain its truth or falsity, as a's range in the disjoint union must be the same as in the equivalent model.

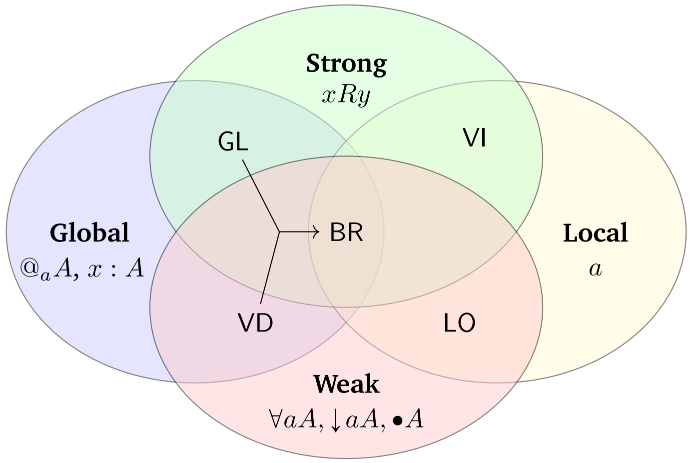

Figure 7: Levels of semantic pollution (the properties of GL together with VD imply the base requirement).

concrete instances $A \in C(\mathsf{PL})$ of a formula type. First, observe that global semantic pollution requires global matters to have full influence on the truth value of formulas for it to result in pollution. Suppose that a proof system introduces an atomic formula A translatable to the specific formula $\forall a(a \wedge p)$, or similarly to $xRy \wedge p$. The parts $\forall a(a)$ and xRy are global and violate invariance under (for instance) FE- and CE-disjoint unions (for these notions, just add disjoint unions to the regular model equivalences in Section 3). However, p introduces not only valuation dependence but also locality into A, so that both versions of A (as atomic formula types) are not globally semantically polluted. This means that covering up globalness with local 'camouflage' is considered to decrease semantic pollution, as this means that it is not the primary intent of a formula to convey just a global property. Analyzed at their main operator, our method still disects $\forall a(a \wedge p)$ and $xRy \wedge p$ into syntactically pure parts ($\wedge$ and p) and parts possessing semantic pollution ($\forall aA$, a, xRy).

As for the definition of local semantic pollution, note that it considers modal formulas that are local and valuation-independent (like $\Diamond\top$) to be syntactically pure. This is seemingly because of BR, which ensures that semantically polluted formulas are not translatable to the modal language. However, even without BR $\Diamond\top$ would not count as semantically polluted, as we measure pollution at the level of formula types. Clearly, $\Diamond(\mathsf{L})$ does not satisfy valuation independence. Still, for each modal instance of a local and valuation independent formula, a new primitive formula type A can be added to the modal language that has exactly the truth condition of this instance (such as 'w has a successor'). Requiring BR then means that formula types A that are translatable to concrete local, valuation independent modal formulas, are pure just like these modal formulas themselves — while without BR, such formula types A *are* semantically polluted (even though their translations are not). Hence, our definition of local pollution, including BR, says that a formula is only locally semantically polluted if you cannot in principle replace it in a formal proof by an expression of the object language.

The way that the modal language can describe a Kripke frame (by being valuation independent) is thus considered acceptable, and a syntactically pure baseline. This includes specific descriptions of R-depth, as the modal operators describe exactly one R-step. The 'height formulas' in (ten Cate and Koudijs, 2022) show that a formula $\Box^{n+1}\bot \wedge \Diamond^n\top$ says that w starts at least one R-path of

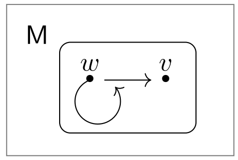

Figure 8: Example of locality.

length exactly n. The modal language cannot describe specific R-width, however, as a consequence of modalities describing 'at least one successor' or 'for all successors'. Specifying a precise number of successors thus provides more opportunities for introducing semantic pollution than specifying the length of an R-chain.

4.2 Results

Figure 7 shows the general division of the formula types in terms of the four levels of pollution. The results of locality and globalness are quite easily seen. Relative to their relevant proof-theoretic language, locality of $\bullet A$ should be obvious, as should globalness of xRy, $@_a A$ and $x : A$. Below are some examples of locality of the remaining HL-formulas.

Example 4.1 (Locality). Consider model M in Figure 8. Let $\tau(a) = w$. Then $\mathsf{M}, \tau, w \models \forall a(\Diamond a)$, as wherever a is sent to, w sees it. However, $\mathsf{M}, \tau, v \not\models \forall a(\Diamond a)$, as v does not see itself nor w. Additionally, $\mathsf{M}, \tau, w \models \downarrow a(\Diamond a)$, as w sees itself, while $\mathsf{M}, \tau, v \not\models \downarrow a(\Diamond a)$, as v does not see itself.[16] And clearly, $\mathsf{M}, \tau, w \models a$, while $\mathsf{M}, \tau, v \not\models a$.

The results of valuation independence are even quicker to see. It is easy to tell that only xRy and a have a truth value that does not vary under changing propositional valuations.

Finally, as promised, some reflection on the hybrid operator $\forall aA$. It does not satisfy our criterion of globalness, but intuitively, its truth value *is* affected by inaccessible worlds (A's truth is tested for a sent by τ to all worlds in turn). To give body to this intuition, note that $\forall aA$ would satisfy a criterion of globalness as follows, a variant of violating invariance under disjoint unions. Given a model PM, add a separate Kripke model (W, R, V) to it (essentially creating a disjoint union, but keeping τ of PM constant (in case of an extended model)). A formula type is then 'global' (i.e., able to be affected by inaccessible worlds) if its truth value can change under this model operation. Clearly, this is the case for $\forall a(\mathsf{HL}_{Q,V})$ (consider $\forall a(a)$), while it is not the case for $\downarrow a(\mathsf{HL}_Q)$ (as shown by an easy induction proof). This illustrates how $\downarrow aA$ is really just a specific subcase of $\forall aA$, sending a to the current world, instead of all worlds in the domain. Thus, $\downarrow aA$ in some sense possess more locality than $\forall aA$, so that $\forall aA$ may deserve a higher level of semantic pollution.

However, note that the latter idea for globalness would, as τ remains static, not consider any labeled formula to be global, which is undesirable. A variant that would attribute semantic pollution to labeled formulas is violating invariance results under CE-, FE- or SCE-disjoint unions — but clearly, this is just another variant of BR, which satisfaction we already require. On top of that, recall that BR itself already captures a difference in pollution level between $\forall aA$ and $\downarrow aA$ (see

[16] Locality of $\forall aA$ and $\downarrow aA$ is already clear from modal instances $\forall a(p)$, $\downarrow a(p)$, but hybrid local instances provide more insight.

Figure 5). Hence, after taking in the four levels of pollution defined here, more nuance within these levels can be found by looking at the levels of BR-satisfaction.

Finally, note the difference between $\forall aA$, $\downarrow aA$ on the one hand, and labels on the other: $\forall aA$ and $\downarrow aA$ are still *general* (it is unclear which world exactly a refers to), while name variables in $@_a A$, $x : A$ and xRy are always *specific* (they pinpoint particular points in a model). This corresponds intuitively to the idea that name variables in the latter operators are more semantically polluted than $\forall aA$ and $\downarrow aA$.

4.3 Four levels compared to the base requirement

To finish this section off, we highlight several observations that come out of the comparison of the four pollution levels to the results concerning BR of Section 3. First, *$\bullet A$ is a minimal example of semantic pollution.* It steadily has the lowest level of pollution, according to the four levels of pollution as well as BR. Within weak semantic pollution, it may be considered to possess a lower amount of pollution than $\forall aA$ and $\downarrow aA$. On our approach, it thus possesses (in our context) the least number of properties required to be semantically polluted. For BR, it only violates invariance under bisimulations (consider generated submodels), but it is neither global nor valuation-independent.

Second, *BR brings nuance within weak semantic pollution and global semantic pollution.* Within weak semantic pollution, BR shows that $\forall aA$ possesses the highest level of pollution in this category, leaving $\downarrow aA$ in between $\forall aA$ and $\bullet A$. Furthermore, while $@_a A$ and $x : A$ are put in the same category of global pollution, BR shows that $@_a A$ is able to express more semantic pollution because of its stronger context language.

Third, *the local-global distinction tears apart a and xRy.* While a and xRy seemed equal in their level of pollution with respect to BR, the four levels tear them apart. Their difference resides in that a retains locality, while xRy does not. In doing so, xRy becomes the maximal example of pollution among our collection of formulas.

We conclude that the four levels of semantic pollution create a helpful overview of the differences between our formula types, while BR remains a useful tool to gain more refined insights into semantic pollution results. The results show that display calculi are very weakly polluted, hybrid calculi possess an intermediate level of semantic pollution (with three different variants), and labeled calculi have the strongest level of pollution (with two different variants).

5 Philosophical views on semantic pollution

The previous sections gave us a definition of the intuitive phenomenon of semantic pollution of modal proof systems. As shown in Section 1.1, the discussion in the literature concerning philosophical suitability of semantically polluted proof systems is ongoing. Instead of aiming to solve the matter in this paper, we will propose that our characterization of semantic pollution is neutral with respect to the debate on suitability for inferentialism; that it emphasizes importance of the distinction between implicit and explicit proof systems; and that it is compatible with less often voiced reasons for desiring syntactic purity.

5.1 Suitability of proof systems for inferentialism

The literature discusses the suitability of proof systems with extended proof-theoretic syntax for inferentialism, i.e. the idea that the meaning of logical connectives is established by their inference rules, instead of model-theoretic semantics (see for instance Schroeder-Heister (2024)). Most notably Read (2015) and De Martin Polo (2024) have provided a philosophical defense of semantically polluted (in particular, labeled) calculi for inferentialism.

One view in this debate, as advocated by Read (2015) and De Martin Polo (2024), says that properties such as harmony (and separability, and others) are decisive in determining suitability for inferentialism, no matter the proof-theoretic language used. Then, labeled calculi are acceptable, as "[t]he labeled rules [...] for $\Box$ and $\Diamond$ are harmonious, that is, the introduction rules encapsulate the whole meaning of the modal operators" (De Martin Polo, 2024). Read (2015) emphasizes that "[t]he semantics lies in the shape of the rules", and so relational atoms and labels "need not be thought of as having any meaning themselves".

A strand of more philosophically-oriented proof theorists argue that proof systems should suitably correspond to our inferential practice (in order to be suitable for inferentialism). This idea also comes in independently for those who use a proof system to faithfully formalize informal reasoning. Its application to inferentialism for instance emerges as Steinberger (2011)'s Principle of Answerability: "[a]dherence to inferentialism importantly constrains one's choice of proof-theoretic frameworks and thus requires one to reject Carnap's amoralism about logic: the inferentialist must remain faithful to our ordinary inferential practice". Not only semantically polluted calculi are subject to this view, but also syntactically pure proof-theoretic languages: for instance multiple-conclusion calculi (see Steinberger (2011) and Restall (2005)), and hypersequent calculi (see e.g. Hjortland and Standefer (2018)). On this view, it is unclear where labeled calculi stand, although at first sight they seem a rather controversial idealization of modal reasoning in practice.[17] These arguments are burdened by the question of what an 'acceptable' idealization is: Dicher (2020) for instance claims that such a boundary is unhelpful, and any idealization should be acceptable.

Our characterization of semantic pollution in this paper is relatively neutral with respect to these aspects of inferentialism. Semantic pollution (in terms of satisfying the base requirement, valuation independence, or globalness) nor syntactic purity prevents or guarantees harmony. Concerning 'principles of answerability', our definition may at first sight be seen as a proposed 'border' for acceptable (semantically polluted) and unacceptable (syntactically pure) idealizations. Semantically polluted calculi analyze the use of a connective in a stronger language than the object language: such strong 'language contexts' are perhaps more prone to unacceptable idealization. However, we discourage such a strict view: strong languages are clearly not guaranteed to be separated from inferential practice (lots of natural logical languages differ in strength), while weaker languages can still be shaped artificially and lack correspondence to practice (just consider the debate on multiple-conclusion calculi (Steinberger, 2009)). In the end, we simply encourage inferentialists using syntax-rich proof systems to spell out the intended interpretation of the syntax in the context of inferential practice.

[17] In terms of the possible world interpretation, there are at least certainly counterexamples. For instance: "Why are horses necessarily mammals? Not because every horse is a mammal in every possible world. But because the property of being a horse bears a special relationship to the property of being a mammal." (Warmke, 2016) As for a temporal interpretation of labels, Arthur Prior (see (Blackburn, 2006)) for a long time considered the use of labels (in hybrid logic) to promote an unnatural 'reasoner-external' perspective, as opposed to our internal experience of time.

5.2 The relevance of distinguishing explicit and implicit proof systems

A further justification for the use of labeled rules in inferentialism draws on the distinction between explicit and implicit rules. In the literature, labeled systems are considered to incorporate the semantics 'explicitly'. Explicit semantic elements are made precise by Poggiolesi (2010) as the idea that sequents containing them are untranslatable to the modal language. Instead, systems such as nested sequents or tree-hypersequents (that do have an interpretation into the logic), then import semantic elements implicitly.[18] Authors often consider the implicit incorporation of semantic elements to be syntactically pure (Poggiolesi, 2010; Brünnler, 2010), while (explicit) labeled calculi are generally considered semantically polluted.

De Martin Polo (2024) and Read (2015) argue that the difference between explicitness and implicitness is less big than it seems. They argue that there is no actual semantic difference between them when considered closely — hence, labels should not be considered more semantically polluted than implicit calculi. The argument in (De Martin Polo, 2024) goes:

> "Read notes that in tree-hypersequent calculi, the semantic content is still explicitly present, but is indicated by the symbols "/" and ";" instead of R. Similarly to Boretti, he argues that the tree-hypersequent rules for necessity only encode the semantic structure of modal formulas in an opaque and disguised manner, thus simply obscuring the semantic apparatus that is more evident in the notation of labeled sequents [...] even though the apparatus of Kripke semantics is presented differently in tree-hypersequent systems than in labeled calculi, it is still (i) explicitly displayed (although obscured in an unconventional notation) and (ii) plays a fundamental role."

While we remain neutral on the difference between implicit and explicit proof systems regarding suitability for inferentialism, we here maintain that this difference *is* relevant regarding the notion of semantic pollution. As claimed above, it is true that the labeled calculus and the nested or tree-hypersequent calculus both arrange modal formulas into a graph structure: but this fact by itself is not enough to claim that their relation to the Kripke semantics is the same. In this paper, we suggest that semantic pollution arises from the *way* that the graph structure is described. If it is described by a language that can express more about Kripke models than the modal language can, then 'too much' detail about the semantics enters the language, and we talk of pollution, and of an 'explicit' calculus. If, like nested and tree-hypersequent calculi, the graph structure is described by formulas that have an interpretation in the logic (and so a truth condition like that of a logical formula), then this graph structure has no particular 'semantic' nature at all — in any case, it is no more semantic than the logic itself. This means that we do not think that "in tree-hypersequent calculi, the semantic content is still explicitly present". Rather, the semantics is described *through* the object language. Of course, it may still be the case that proof-theoretic syntax which has a logical counterpart, has a different informal meaning (e.g. in inferential practice) than its logical translation. However, our point here is that the relation to the model theory of the proof-theoretic syntax and its logical counterpart (by their truth conditions) is in fact the same.

Saying that "/" and ";" display Kripke semantics at the level of tree-hypersequents, is like saying that $\Diamond$ and $\Box$ display the semantics at the level of the basic modal language. Should we thus consider $\Diamond$ and $\Box$ as semantically polluted? It seems clear that this is not so. Simply as a consequence

[18] As mentioned, this is similar to the distinction between 'external' and 'internal' proof systems, see also (Lyon et al., 2023).

of its model-theoretic truth conditions, *any* syntactic element displays the semantics to some extent. The interesting question for semantic pollution is where the boundary lies: what syntactic elements do we consider pure (surely, the logical object language) and what syntactic elements do we consider impure (our proposal is found in the previous sections).

This relates to our conceptual understanding of the different ways in which proof calculi can describe Kripke frames (see also Poggiolesi and Restall (2012)): labeled systems can explicitly and globally describe a Kripke frame, while display calculi use a local perspective while allowing perspective switches along R-connections (i.e., switches between actual worlds). By incorporating the forward as well as backward perspective, display calculi can describe a Kripke frame better than the modal language. The tree-hypersequent or nested systems, on the other hand, through their logical interpretation, only possess the (syntactically purest) local 'forward' perspective: everything is encoded through uses of $\Box$.

Note also that this counters the 'proof-theoretic' idea that 'notational variants' of proof systems should have an equal level of semantic pollution.[19] There exist proof-theoretic translations between labeled sequents, nested sequents, (tree-)hypersequents and display sequents (see e.g. Ciabattoni et al. (2021) and Goré and Ramanayake (2014), and the hierarchy of translations defined in (Lyon et al., 2023)). Now consider for instance a labeled calculus that only allows labeled *tree* sequents (so that it is formally now 'just a notational variant' of a nested or tree-hypersequent calculus, in which all structures are already trees). Our proposal will still say that this labeled calculus is semantically polluted, whereas the nested and tree-hypersequent calculus is not: the labeled calculus still describes a tree using more expressive power than necessary. It uses relational atoms and labels that, by definition, can express more semantic information than the nested structure. And perhaps more importantly: if made true model-theoretically, the labeled tree sequent will actually indicate a tree-form in the Kripke model, by indicating the specific worlds and relations. A nested sequent or tree-hypersequent has a tree structure within the sequent, but its model-theoretic truth at a world (interpreted in terms of disjunction and box) does not enforce this world to be arranged in a tree inside the model. That is, even though the different proof systems may describe the same graph arrangement in a sequent, the labeled calculus does this in a *semantically polluted way*. This shows that proof translations do not always preserve philosophical values.

Hence, on our approach, the distinction between explicit and implicit notions matters when defining semantic pollution: it provides a natural formal definition supporting a distinction that philosophers and proof theorists already make intuitively — and we thus conclude that this distinction cannot be abolished on the grounds that their semantic content is the same.

5.3 Syntactic purity as an ideal of proof (systems)

Finally, we argue here that there exist clear cases where semantic pollution has philosophical harm, other than possible harm for inferentialism, and for closeness to inferential practice. First, motivations of aesthetics, or ideals of proof, are and have always been common in mathematical fields. Dawson (2006) presents an overview of reasons mathematicians have to reprove a theorem, among which "to employ reasoning that is simpler [...] than earlier proofs". Simplicity is a well-known ideal of proof, where we may distinguish between conceptual simplicity, and formal (computational) simplicity. In our analysis of proof systems, we can make a similar distinction. labeled

[19] See (French, 2019) for an analysis of when two logics can be said to be notational variants — note that it is unclear whether the intuitive use of notational variant here corresponds to this analysis, although this is not the place to further investigate this.

calculi may conceptually provide a simple way of analyzing inference rules (for those familiar with Kripke semantics). However, the ideal of simplicity may also manifest itself in the search for a proof system in as small a language as possible. I.e., as also noted by Lyon (2021, p. 112), we might be interested in finding out that there exists a satisfying proof system for modal logic that is restricted to the modal object language, without any 'brute force' or potential clutter from external syntax. No more than aesthetics (a desire for resource-minimality) is involved in this, and yet it is a common motivation in mathematics. From this point of view, semantic pollution is in fact undesirable. Incidentally, other reasons than aesthetics may still also apply: Lyon (2021) notes that a tree structure of sequents (that nested sequents or tree-hypersequents guarantee, but labeled sequents do not) can be necessary for certain proof-search algorithms. In this case, the excess of syntax in semantically polluted systems can even provide too much freedom for technical applications.

Another ideal of proof (systems) may be, given a model-theoretic semantics for it, that the proof system and the model theory are sufficiently conceptually separated. This is possibly simply what Avron (1996) meant with his claim that a proof calculus should be 'independent' from any particular semantics: the fact that we consider proof theory to be an activity that is somehow separate from model theory. Similarly, soundness and completeness proofs should then show us a 'valuable' insight: instead of giving us two notational variants of the same conceptual approach to a logic, they should connect a model-theoretic perspective to a (sufficiently different) proof-theoretic perspective. If not, the 'proof system' can be regarded merely as a systematization of semantic thought. Although not much more than aesthetics seems to validate these preferences, in a bottom-up approach to philosophy of proof theory, they should be taken seriously.

A different motivation for discarding semantic proof systems, in particular labeled systems, has to do with *impartiality* with respect to the background logic. Given a desire for a model-theoretic semantics, and the natural interpretation of labels into Kripke semantics, there is arguably a sense in which labeled calculi favor a classicist world-view. Classicists who want to reason with modalities may be happy to accept labels as concerning time or possibilities. Intuitionists who wish to assign some interpretation to labels and relational atoms, may struggle to find a satisfying one: labeled rules would ask them to quantify over worlds (or times), and explicitly refer to states other than the actual one. This simply may not be acceptable for them, even though modal reasoning should be a perfectly acceptable activity for intuitionists. An interesting open question relating to the latter two points (which we leave to future research) is how a proof system is *formally* independent or impartial from a particular semantics. That is, perhaps there is an interesting way to say that the relation of labeled calculi to Kripke semantics is more 'necessary' formally, than its relation to other types of semantics.

6 Conclusion

We have presented a characterization of semantic pollution of proof systems in terms of four levels of pollution. Our measures suggest that the nature of modal syntax lies in what it can express about Kripke models, its local view of a model, and direct interaction with basic propositions. Instead, the properties of higher expressivity than the modal language, a global view of a model, and valuation independence suggest semantic influence in proof-theoretic syntax.

Our results show that the display calculus is only weakly semantically polluted (by $\bullet A$). The hybrid calculus, on the other hand, introduces formula types that are weakly semantically polluted ($\forall aA, \downarrow aA$), but also ones that possess global ($@_a A$) and local (a) semantic pollution. Finally, the labeled calculus introduces a globally polluted formula type ($x : A$) and the only strongly semantically polluted formula type (xRy). In line with intuitions throughout the literature, then, the labeled calculus can be seen as possessing the highest level of semantic pollution (among the calculi that we studied). We concluded that the difference between explicit and implicit proof calculi is key in our characterization of semantic pollution, and that besides the virtues of polluted calculi, semantic pollution can just as well have technical and philosophical downsides.

We might also seek a more general analysis of semantic pollution as a distinction between syntactic proof systems and their semantics, with applications to all kinds of logics. Indeed, semantic pollution might occur in all logical areas where 'good' proof systems are hard to find. The measures of violating invariance results under model equivalences, and being valuation independent, seem relatively easily transferable to other logics. World invariance seems more tailored to a type of semantics with different 'points of evaluation', although extensions to intuitionistic logics seem natural. These quick considerations already give reason to think that certain proof systems for neighborhood semantics (similar to labeled systems) are semantically polluted (see (Dalmonte, Olivetti, and Negri, 2018), and similarly (Negri, 2016)), as well as a proof system for intuitionistic predicate logic (see (Baaz and Iemhoff, 2008)).

Besides the extension of semantic pollution to other logics, future research could analyze semantic pollution of modal proof systems for different types of proof systems than the ones we chose to study; or they could provide additional conceptions of semantic pollution that come with different measures. These directions could all provide us with a more fundamental understanding of the difference between syntax and semantics.

References

Avron, Arnon (1987). "A constructive analysis of RM". In: *The Journal of symbolic logic* 52.4, pp. 939–951.

— (1996). *The method of hypersequents in the proof theory of propositional non-classical logics*. na.

Baaz, Matthias and Rosalie Iemhoff (2008). "On Skolemization in constructive theories". In: *The Journal of Symbolic Logic* 73.3, pp. 969–998.

Belnap, Nuel D (1982). "Display logic". In: *Journal of philosophical logic*, pp. 375–417.

Blackburn, Patrick (2006). "Arthur Prior and hybrid logic". In: *Synthese* 150.3, pp. 329–372.

Blackburn, Patrick, Johan FAK van Benthem, and Frank Wolter (2006). *Handbook of modal logic*. Elsevier.

Blackburn, Patrick, Maarten De Rijke, and Yde Venema (2001). *Modal logic*. Vol. 53. Cambridge University Press.

Braüner, Torben (2010). *Hybrid logic and its proof-theory*. Vol. 37. Springer Science & Business Media.

Braüner, Torben and Valeria de Paiva (2006). "Intuitionistic hybrid logic". In: *Journal of Applied Logic* 4.3, pp. 231–255.

Brünnler, Kai (2010). "Nested sequents". In: *arXiv preprint arXiv:1004.1845*.

Bull, Robert A (1992). "Cut elimination for propositional dynamic logic without". In: *Mathematical Logic Quarterly* 38.1, pp. 85–100.

ten Cate, Balder (2004). *Model theory for extended modal languages*. University of Amsterdam.

ten Cate, Balder and Raoul Koudijs (2022). "Characterising Modal Formulas with Examples". In: *arXiv preprint arXiv:2206.06049*.

Ciabattoni, Agata et al. (2021). “Display to labeled proofs and back again for tense logics”. In: *ACM Transactions on Computational Logic (TOCL)* 22.3, pp. 1–31.

Dalmonte, Tiziano, Nicola Olivetti, and Sara Negri (2018). “Non-normal modal logics: Bi-neighbourhood semantics and its labelled calculi”. In: *Advances in Modal Logic 2018*.

Dawson, John W (2006). “Why do mathematicians re-prove theorems?” In: *Philosophia Mathematica* 14.3, pp. 269–286.

De Martin Polo, Fabio (2024). “Beyond Semantic Pollution: Towards a Practice-Based Philosophical Analysis of Labelled Calculi”. In: *Erkenntnis*, pp. 1–30.

Dicher, Bogdan (2020). “Hopeful monsters: a note on multiple conclusions”. In: *Erkenntnis* 85.1, pp. 77–98.

Dyckhoff, Roy (2016). “Intuitionistic decision procedures since Gentzen”. In: *Advances in proof theory*, pp. 245–267.

French, Rohan (2019). “Notational variance and its variants”. In: *Topoi* 38, pp. 321–331.

Goré, Rajeev and Revantha Ramanayake (2014). “Labelled tree sequents, tree hypersequents and nested (deep) sequents”. In: *Advances in modal logic*, pp. 279–299.

Hjortland, Ole and Shawn Standefer (2018). “Inferentialism, structure, and conservativeness”. In: *From Rules to Meanings*. Routledge, pp. 115–140.

Kanger, Stig (1957). “Provability in logic”. In.

Kashima, Ryo (1994). “Cut-free sequent calculi for some tense logics”. In: *Studia Logica*, pp. 119–135.

Lyon, Tim (2021). “Refining labelled systems for modal and constructive logics with applications”. In: *arXiv preprint arXiv:2107.14487*.

Lyon, Tim S et al. (2023). “Internal and external calculi: Ordering the jungle without being lost in translations”. In: *arXiv preprint arXiv:2312.03426*.

Marin, Sonia (2018). “Modal proof theory through a focused telescope”. PhD thesis. Université Paris Saclay.

Negri, Sara (2005). “Proof analysis in non-classical logics”. In: *Logic Colloquium*. Vol. 28, pp. 107–128.

— (2011). “Proof theory for modal logic”. In: *Philosophy Compass* 6.8, pp. 523–538.

— (2016). “Non-normal modal logics: a challenge to proof theory”. In: *The Logica Yearbook*, pp. 125–140.

Poggiolesi, Francesca (2009). “The method of tree-hypersequents for modal propositional logic”. In: *Towards mathematical philosophy*. Springer, pp. 31–51.

— (2010). *Gentzen calculi for modal propositional logic*. Vol. 32. Springer Science & Business Media.

Poggiolesi, Francesca and Greg Restall (2012). “Interpreting and applying proof theories for modal logic”. In: *New waves in philosophical logic*. Springer, pp. 39–62.

Pottinger, Garrel (1983). “Uniform, cut-free formulations of T, S4 and S5”. In: *Journal of Symbolic Logic* 48.3, p. 900.

Read, Stephen (2015). “Semantic pollution and syntactic purity”. In: *The Review of Symbolic Logic* 8.4, pp. 649–661.

Restall, Greg (2005). “Multiple conclusions”. In: *Logic, methodology and philosophy of science: Proceedings of the twelfth international congress*. Kings College Publications London, pp. 189–205.

Schroeder-Heister, Peter (2024). “Proof-Theoretic Semantics”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Summer 2024. Metaphysics Research Lab, Stanford University.

Simpson, Alex K (1994). “The proof theory and semantics of intuitionistic modal logic”. In.

Steinberger, Florian (2009). “Harmony and logical inferentialism”. PhD thesis. University of Cambridge.
— (2011). “Why conclusions should remain single”. In: *Journal of Philosophical Logic* 40.3, pp. 333–355.
Wansing, Heinrich (1994). “Sequent calculi for normal modal propositional logics”. In: *Journal of Logic and Computation* 4.2, pp. 125–142.
Warmke, Craig (2016). “Modal semantics without worlds”. In: *Philosophy Compass* 11.11, pp. 702–715.

7 Appendix

We here provide the various induction proofs referred to in the paper:

1. $\downarrow aA$ is invariant under *FE-isomorphisms* relative to HL_V (see Theorem 2 and 3).
2. $\forall aA$ is invariant under *FE-isomorphisms* relative to HL_V (see Theorem 2 and 4).
3. All formulas in HL are invariant under *SCE-isomorphisms* (see Theorem 5).
4. All formulas in HL_Q are invariant under *SCE-bisimulations* (see Theorem 6).

The first proof will set the stage for showing that $\downarrow aA$ is invariant under FE-isomorphisms relative to HL_V. Note that we cannot show that any hybrid formula A has this property, as the base case where A is a nominal a violates it. The reason we show the general Theorem 2 before Theorem 3 is that the case of $\Box A$ requires a to be mapped to an arbitrary y (instead of already sending a to the current world as $\downarrow aA$ does), in order to apply the induction hypothesis at R-reachable worlds. We use $\cong_{\mathsf{FE}}$ as a symbol for FE-isomorphism, and $\cong_{\mathsf{SCE}}$ for SCE-isomorphism.

Theorem 2. *Suppose $M, \tau, w \cong_{FE} M', \tau', w'$ and $M, \tau, y \cong_{FE} M', \tau', y'$. Then for all $A \in HL_V$: $M, \tau_{[a\mapsto y]}, w \vDash A$ iff $M', \tau'_{[a\mapsto y']}, w' \vDash A$.*

Proof. The proof proceeds by induction on A, with an induction hypothesis for B less complex than A. Assume $\mathsf{M}, \tau, w \cong_{\mathsf{FE}} \mathsf{M}', \tau', w'$, and $\mathsf{M}, \tau, y \cong_{\mathsf{FE}} \mathsf{M}', \tau', y'$. Consider the cases below (we omit the straightforward cases of conjunction and negation).

- **Base case.** Clearly, a proposition letter p has this type of invariance under FE-isomorphisms. For nominals, suppose $\mathsf{M}, \tau_{[a\mapsto y]}, w \vDash a$. Then $\tau_{[a\mapsto y]}(a) = w$ (and $y = w$). Now consider $\tau'_{[a\mapsto y']}$. As $\mathsf{M}, \tau, y \cong_{\mathsf{FE}} \mathsf{M}', \tau', y'$, and $y = w$, it holds that $w \cong_{\mathsf{FE}} y'$. As $\mathsf{M}, \tau, w \cong_{\mathsf{FE}} \mathsf{M}', \tau', w'$, it also holds that $w' = y'$. Hence, $\tau'_{[a\mapsto y']}(a) = w'$, and $\mathsf{M}', \tau'_{[a\mapsto y']}, w' \vDash a$.

- **Box.** Suppose $\mathsf{M}, \tau_{[a\mapsto y]}, w \vDash \Box A$. We need to show that for all v' such that $Rw'v'$, $\mathsf{M}', \tau'_{[a\mapsto y']}, v' \vDash A$. By the isomorphism, for every such v' there is an isomorphic world v such that wRv. By assumption, for these v, $\mathsf{M}, \tau_{[a\mapsto y]}, v \vDash A$. Then by the induction hypothesis, for each corresponding v' it holds that $\mathsf{M}', \tau'_{[a\mapsto y']}, v' \vDash A$. Hence, $\mathsf{M}', \tau'_{[a\mapsto y']}, w' \vDash \Box A$.

- **Satisfaction.** Suppose $\mathsf{M}, \tau_{[a\mapsto y]}, w \vDash @_a A$. Then $\mathsf{M}, \tau_{[a\mapsto y]}, \tau_{[a\mapsto y]}(a) \vDash A$. By assumption, $\tau_{[a\mapsto y]}(a)$ (that is, y) is isomorphic to y'. Then by the induction hypothesis, $\mathsf{M}', \tau'_{[a\mapsto y']}, \tau'_{[a\mapsto y']}(a) \vDash A$. That is, $\mathsf{M}', \tau'_{[a\mapsto y']}, w' \vDash @_a A$.

- **For all.** Suppose $\mathsf{M}, \tau_{[a\mapsto y]}, w \models \forall aA$. Then $\mathsf{M}, (\tau_{[a\mapsto y]})_a, w \models A$ for all $(\tau_{[a\mapsto y]})_a$. Note that this simply equals $\mathsf{M}, \tau_{[a\mapsto v]}, w \models A$ for all $v \in W$. By the isomorphism, for each $v \in W$ we have a world $v' \in W'$ such that $\mathsf{M}, \tau, v \cong_{\mathsf{FE}} \mathsf{M}', \tau', v'$. By the induction hypothesis, and by surjectivity of the isomorphism, we have $\mathsf{M}', \tau'_{[a\mapsto v']}, w' \models A$ for each $v' \in W'$. In other words, it holds that $\mathsf{M}', \tau'_a, w' \models A$ for all τ'_a. So, $\mathsf{M}', \tau', w' \models \forall aA$, and we can also rewrite this as $\mathsf{M}', \tau'_{[a\mapsto y']}, w' \models \forall aA$ (as the assignment of a does not matter).

- **Down-arrow.** Suppose $\mathsf{M}, \tau_{[a\mapsto y]}, w \models \downarrow aA$. Then $\mathsf{M}, (\tau_{[a\mapsto y]})_{[a\mapsto w]}, w \models A$. We can rewrite this as $\mathsf{M}, \tau_{[a\mapsto w]}, w \models A$. Now we apply the induction hypothesis to obtain $\mathsf{M}', \tau'_{[a\mapsto w']}, w' \models A$. Again, we can write this as $\mathsf{M}', (\tau'_{[a\mapsto y']})_{[a\mapsto w']}, w' \models A$, so that we obtain $\mathsf{M}', \tau'_{[a\mapsto y']}, w' \models \downarrow aA$.

□

Theorem 3. *$\downarrow aA$ is invariant under FE-isomorphisms relative to HL_V.*

Proof. Take Theorem 2 and consider the instance where $y = w$ and $y' = w'$. The theorem then says that for isomorphic states w and w', $\mathsf{M}, \tau, w \models \downarrow aA$ iff $\mathsf{M}', \tau', w' \models \downarrow aA$. □

Theorem 4. *$\forall aA$ is invariant under FE-isomorphisms relative to HL_V.*

Proof. We simply need to prove that for all models M, M', assignment functions τ, τ', worlds w, w' and variants τ_a and τ'_a, if $\mathsf{M}, \tau, w \cong_{\mathsf{FE}} \mathsf{M}', \tau', w'$ then

$$\mathsf{M}, \tau_a, w \models A \text{ for all } \tau_a \text{ iff } \mathsf{M}', \tau'_a, w' \models A \text{ for all } \tau'_a$$

So suppose $\mathsf{M}, \tau, w \cong_{\mathsf{FE}} \mathsf{M}', \tau', w'$, and that for some $A \in \mathsf{HL}_V$, $\mathsf{M}, \tau_a, w \models A$ for all τ_a. This equals $\mathsf{M}, \tau_{[a\mapsto v]}, w \models A$ for all $v \in W$. By the isomorphism, for each $v \in W$ there is a world $v' \in W'$ such that $\mathsf{M}, \tau, v \cong_{\mathsf{FE}} \mathsf{M}', \tau', v'$. By Theorem 2, and by surjectivity of the isomorphism, we obtain $\mathsf{M}', \tau'_{[a\mapsto v']}, w' \models A$ for each $v' \in W'$. Thus, $\mathsf{M}', \tau'_a, w' \models A$ for all τ'_a. □

Theorem 5. *All formulas $A \in \mathit{HL}$ are invariant under SCE-isomorphisms.*

Proof. The proof is by induction on A, with an induction hypothesis for formulas B less complex than A. We only treat the cases of nominals and of the hybrid operators (as the modal cases are straightforward). Suppose $\mathsf{M}, \tau, w \cong_{\mathsf{SCE}} \mathsf{M}', \tau', w'$.

- **Base case.** Suppose $\mathsf{M}, \tau, w \models a$, so that $\tau(a) = w$. By the requirements of SCE-isomorphisms, equivalent worlds satisfy the same nominals, and so $\tau'(a) = w'$, and $\mathsf{M}', \tau', w' \models a$.

- **Satisfaction.** Suppose $\mathsf{M}, \tau, w \models @_a A$. Then $\mathsf{M}, \tau, \tau(a) \models A$. By the induction hypothesis, and the fact that $\mathsf{M}, \tau, \tau(a) \cong_{\mathsf{SCE}} \mathsf{M}', \tau', \tau'(a)$, $\mathsf{M}', \tau', \tau'(a) \models A$. Hence, $\mathsf{M}', \tau', w' \models @_a A$.

- **For all.** Suppose $\mathsf{M}, \tau, w \models \forall aA$, so that $\mathsf{M}, \tau_{[a\mapsto v]}, w \models A$ for all $v \in W$. We have to show that $\mathsf{M}', \tau'_{[a\mapsto v']}, w \models A$ for all $v' \in W'$. By the isomorphism, for each $v' \in W'$ we have a world $v \in W$ such that $\mathsf{M}, \tau, v \cong_{\mathsf{SCE}} \mathsf{M}', \tau', v'$. Specifically, for each such pair of worlds, $\mathsf{M}, \tau_{[a\mapsto v]}, w \cong_{\mathsf{SCE}} \mathsf{M}', \tau'_{[a\mapsto v']}, w'$ holds: changing a's assignment does not break the SCE-isomorphism, as a is still sent to equivalent worlds. By the induction hypothesis, and by surjectivity of the isomorphism, we have $\mathsf{M}', \tau'_{[a\mapsto v']}, w' \models A$ for each $v' \in W'$. Hence, $\mathsf{M}', \tau', w' \models \forall aA$.

- **Down-arrow.** Suppose $\mathsf{M}, \tau, w \models \downarrow aA$, so $\mathsf{M}, \tau_{[a\mapsto w]}, w \models A$. Then, we also have a strong isomorphism $\mathsf{M}, \tau_{[a\mapsto w]}, w \cong_{\mathsf{SCE}} \mathsf{M}', \tau'_{[a\mapsto w']}, w'$, as a is still sent to equivalent worlds. By the induction hypothesis, $\mathsf{M}', \tau'_{[a\mapsto w']}, w' \models A$. Hence, $\mathsf{M}', \tau', w' \models \downarrow aA$.

□

Theorem 6. *All formulas $A \in HL_Q$ are invariant under SCE-bisimulations.*

Proof. This proof proceeds in the same way as that of Theorem 5. Note that the base case works even for bisimulations, as SCE-bisimulations still need equivalent worlds to satisfy the same nominals. We skip the $\forall aA$ case (which does not work anymore because of possible change in model cardinality over the equivalence), but the cases of nominals, $@_a A$ and $\downarrow aA$ go through as before. □