Journal of Philosophical Logic
https://doi.org/10.1007/s10992-024-09775-2

The Structure of Paradoxes in a Logic of Sentential Operators

Michał Walicki[1]

Received: 21 June 2022 / Accepted: 29 August 2024

Abstract

Any language $\mathcal{L}$ of classical logic, of first- or higher-order, is expanded with sentential quantifiers and operators. The resulting language $\mathcal{L}^+$, capable of self-reference without arithmetic or syntax encoding, can serve as its own metalanguage. The syntax of $\mathcal{L}^+$ is represented by directed graphs, and its semantics, which coincides with the classical one on $\mathcal{L}$, uses the graph-theoretic concepts of kernels and semikernels. Kernels provide an explosive semantics, while semikernels generalize this to situations where paradoxes do not lead to explosion, thus distinguishing them from contradictions. Paradoxes arise only at the metalevel due to specific interpretations of the operators, but they can be avoided: $\mathcal{L}^+$ can express paradoxes but remains free from them. For an expansion $\mathcal{L}^+$ of any FOL language $\mathcal{L}$, with the non-explosive semantics, a complete reasoning system is obtained by extending Gentzen's classical sequent calculus with two rules for the sentential quantifiers. Adding (cut) yields a complete system for the explosive semantics. The novel semantics and self-referential capabilities seem promising for a further extension of classical logic towards one capable also of consistently expressing its own syntax and truth theory.

Keywords Sentential operators · Semantic and intensional paradoxes · Classical logic · Explosive and non-explosive semantics · (semi)kernels of digraphs

1 Introduction

Arithmetic with the arithmetized syntax can serve as a base syntax theory over which truth theories can be formulated. By "AST" we denote first-order arithmetic with an effective coding of expressions of its language augmented by a truth predicate, $\mathcal{T}$, on natural numbers. For each natural number n, if n is the code of an expression ϕ, let $\ulcorner\phi\urcorner$ be the numeral of n in the language. Convention T – instances of the schema $\mathcal{T}(\ulcorner\phi\urcorner) \leftrightarrow \phi$, for all sentences ϕ – provides *prima facie* desirable truth-theoretic

✉ Michał Walicki
Michal.Walicki@uib.no

1 Department of Informatics, University of Bergen, Bergen, Norway

principles. However, as Tarski has shown in [44], it cannot hold unrestrictedly in a classical theory, on pain of inconsistency. Since it can not even be systematically approximated, as formal criteria provide no reliable guide by [25], the search for adequate restrictions on Convention T has attracted much attention, yielding diverse proposals and results.

This paper lets the syntax and truth theories loom in the background, at most hinting towards the possibility of approaching them differently. It presents only a new model of self-reference and paradox analysis, formulated without such theories by replacing predicates on coded syntax with sentential operators. This change can be initially motivated by the fact that such operators appear to be less paradox prone. Convention T reflects the intended application of predicate $\mathcal{T}$ as a meta-predicate on sentences, identifying the (codes of the) true ones. Unrestricted, it gives only one form of paradoxes arising with such meta-predicates modelled by predicates on the arithmetized syntax. When the latter express basic modal notions, elements of temporality, or just negation, the diagonalization lemma yields paradoxes also without Convention T, e.g., [12, 20, 26, 37]. The corresponding paradoxes do not arise in the operator setting. Certainly, operators handling enough syntax, especially substitution, bring paradoxes back, e.g., [17, 30], but they act then on open formulas, which must be left for future work. Here, they act only on the closed formulas (sentences). The operator approach has proven fruitful in modal logics, with its philosophical advantages over predicates reviewed in [19]. The paper gathers such indications into a logic of sentential operators, LSO, neither restricted to modalities nor relying on Kripke semantics. Unlike other philosophical discussions, e.g., [27, 42], it provides a fully formal account of operators. Major differences from AST are not limited to formalism; they also involve the perspective on the language.

The key distinction lies in the relationship between the object-language and the metalanguage. In AST, the coding of the syntax turns the latter into a subset of the former. In LSO, an object-language is any classical – propositional, first- or higher-order – language $\mathcal{L}$. It can be extended with sentential quantifiers and operators to a language $\mathcal{L}^+$. Occurrences of the operators mark then what we view as the metalanguage of $\mathcal{L}^+$ (not only of $\mathcal{L}$). For an operator K and sentence S, a sentential atom, *s-atom*, $K(S)$ is an atomic sentence of the metalanguage about S, $K(K(S))$ is such an atomic sentence about $K(S)$, etc. (The paper is thus written in the meta-metalanguage.)

Valuations of s-atoms determine relations on sentences (and possibly objects). Being such a relation is the only requirement on an operator, just as being a relation on objects is the only requirement on an interpretation of a predicate symbol in FOL. If $K(S)$ is true, we can read it as K being true about S, as S being said or known, and such choices may require additional axioms. The operators need not be truth-functional and can treat arguments purely syntactically, e.g., $K(S)$ may be true exactly when S is a universal sentence in prenex normal form. This opens a pathway towards a theory of syntax, as well as to modelling intensionality and modalities, but these topics are not explored here. LSO is a hyperintensional logic only to the extent that the operators can be opaque, failing to preserve logical equivalence of arguments, but it neither provides any intensional semantics nor considers the status of propositions or propositional attitudes. Propositions appear at most as mere shadows of sentences, while examples blur easily the borders between "says ϕ", "assumes ϕ", "claims ϕ", etc. The critical distinction is between the statements with and without the operators, between the metalevel and the object-level. Modalities, attitudes, or intensions can be handled by further axiomatizations of operators. Nonetheless, paradoxes of semantic and intensional character arise and are treated in LSO in the same way. The liar, saying only "Every sentence I am saying now is false", is significantly different neither from one not believing any of one's beliefs nor from a club whose members are all people not belonging to any club. Semantic and intensional paradoxes are caused by the same patterns, which will be noted in the text.

One of the consequences of the formal distinction between the object-language and the metalanguage is the difference between contradictions and paradoxes. Informally, we can differentiate sentences that are always false from those that initially seem plausible but lead to contradictions. In classical logic, both types are just contradictions, but LSO maintains this distinction. Unlike contradictions, paradoxes can have some valid interpretations. For example, Karen might claim that John always tells the truth, while John accuses Karen of always lying. The implied incompatibility does not preclude the event. In this sense, paradoxes can occur in the world, even if only at the metalinguistic level, which is still a part of the world. This distinguishes paradoxes from contradictions, which cannot occur empirically. One can assert both "*S*" and "not *S*", but simultaneously saying "*S*" and not saying "*S*" is impossible – not only according to Aristotle and common sense but even according to some dialetheists, who accept only certain contradictions, such as the liar both lying and not lying. This paradox, like others, occurs at the metalevel, attributing contradictory properties to the liar's statement. Adhering thus to classical logic and excluding contradictions in the world of objects, LSO admits them as consequences of paradoxes, which emerge only in the metalanguage.

In LSO, as for Tarski, *"[t]he appearance of an antinomy is [...] a symptom of disease"*, see [45]. It is not, however, any intrinsic disease of the language but one caused by a contagious valuation of s-atoms, an unfortunate interpretation of the operators. Let $K(S)$ stand for "Karen says *S*" and Kl for "Everything Karen says is false", i.e., $\forall\phi(K\phi \rightarrow \neg\phi)$. Karen saying Kl, $K(Kl)$, is intuitively not paradoxical, if she also says something true. However, Karen saying only Kl implies Kl and $\neg Kl$, witnessing to a paradox, the Karen-liar, that would not occur if she also said something true. In LSO, this paradox arises when K is true about Kl and false about all true sentences, but does not occur under other interpretations of K.

As the example suggests, besides certain interpretations of operators, paradox also requires sentential quantification. In [42], Stern observes that *"in the predicate setting the liar-like paradoxes depend crucially on the properties of the truth predicate or the modal predicates, whereas in the operator setting it is quantification simpliciter that leads to paradox"*. We prove that extending classical logic with either sentential quantifiers alone or operators alone does not lead to paradoxes. Only adding both does, and then it is not the quantification simpliciter that is to blame, but its combination with certain interpretations of the operators.

Upon hearing a paradox, one recognizes the implied contradiction but continues reasoning as usual, unaffected. This natural non-explosiveness is mirrored in LSO in a specific way.

The Karen-liar implies that Karen is lying and not lying, but neither snow ceases to be white nor anything about what John is saying follows. Paradoxes can be seen as such claims in the metalanguage causing the impossibility of consistently evaluating some sentences but not leading to any explosion. In LSO, the possibility of Karen saying only Kl is reflected by the models realizing such a situation, where only some contradictions follow, e.g. $Kl \wedge \neg Kl$, but not others. This element of relevance reflects the informal way of identifying a paradox. We do not explain the liar by deriving from it that snow is not white. We derive 'the relevant contradiction', that the liar is lying and not lying. In LSO, with the valuation of s-atoms where only $K(Kl)$ is true, basically only sentence Kl cannot be assigned any coherent value, while most others can. (Coherence will be defined as a local consistency, which may fail to extend to a consistent valuation of all sentences.) This indicates which s-atoms might need change to resolve the paradox, here, the atoms with the main operator K, affecting the value of Kl. Otherwise, the model allows John to say (or not) whatever he likes, and retains all usual interpretations of the object-language. Formally, a paradox in LSO consists of a set of sentences possessing a model, namely, a valuation of atoms satisfying these sentences, which cannot be extended to any consistent valuation of all sentences of the metalanguage due to the implied contradictions.

In this way, LSO models of theories involving paradoxes are partial. We say that a model of the Karen-liar, providing no value for Kl, does not *cover* Kl. (A model cannot be arbitrarily extended to cover the desired sentences. In this case, Karen's models have no extensions covering Kl.) The notion has no counterpart in classical logic but could be compared to an argument that is classically consistent provided that one ignores questions about Kl. Total models, covering the whole language, i.e., interpreting coherently all sentences, are special cases when the theory, with all its metalevel claims, is consistent. At this point, we can only signal that both are formalized using digraphs with sentences as vertices. So-called *semikernels*, originating from [28], represent partial coherent valuations and non-explosive semantics. Simplifying, we can think of a semikernel as a subset of sentences covering the sentences it contains and their negations, while satisfying all its members and (vacuously) all sentences it does not cover. A contradiction can be satisfied only vacuously by a model not covering it. This is a weak form of paraconsistency. For instance, sound reasoning derives both Kl and $\neg Kl$ from the Karen-liar without deriving any object-level contradictions or anything about John's statements. (The Karen-liar has namely partial models where John makes any statements whatsoever, while all models covering the sentences of the object-language are classical.) A paradox in LSO implies a contradiction but possesses models. These can satisfy a contradiction only vacuously, and hence cannot be extended to valuations of the whole language. Such total valuations are captured by *kernels*, that is, semikernels covering all sentences of the language. (The concept of a kernel originated in [47].) Valuating consistently all sentences, kernels are total models which yield the explosive semantics as a special case of the non-explosive one. Valid sentences must be satisfied also by these models, so the logic remains two-valued and retains all classical tautologies. Classical contradictions can be satisfied vacuously, but have no models covering them.

The close connections between the non-explosive (partial) and the explosive (total) semantics are reflected by the associated reasoning system. Surprisingly, reasoning about the partial models validates all classical rules of the sequent calculus. The reasoning system, denoted (with a sans-serif font) LSO, extends Gentzen's classical sequent system LK, e.g., [16, 46], with two rules for s-quantifiers. Sentential operators and s-quantification bring a flavour of higher-order, but neither the former nor the substitutional interpretation of the latter carries any set-theoretic connotations. Actually, LSO is sound and complete for the non-explosive semantics of the expansion $\mathcal{L}^+$ of any first-order language $\mathcal{L}$ (denoted by FOL$^+$). The fact that it merely extends the classical LK with two rules for the s-quantifiers may justify calling it "classical". However, it works for the non-explosive semantics, providing classical reasoning exploding from contradictions but not from paradoxes. The derivability of all sentences from a paradox, via the contradictions it implies, is prevented by the inadmissibility of (cut) $\frac{\Gamma\vdash\Delta,A \quad A,\Gamma\vdash\Delta}{\Gamma\vdash\Delta}$ at the metalevel. The explosive, "fully classical" semantics is captured by reasoning in LSO with unrestricted (cut). It allows to conjoin the derivation of some contradiction C from a paradox P with the derivation of any statement X from contradiction C, yielding a derivation of X from paradox P. In this way also (cut) reflects the difference between contradictions and paradoxes, or between the object-level, where it is admissible, and the metalevel, where it makes paradoxes explode. In this sense, LSO is non-transitive, like the systems in [10, 39, 40], but unlike there, (cut) does not trivialise the logic but merely transforms its non-explosive version into the explosive one.

Unlike the valuations of $\mathcal{L}$ atoms, *o-atoms*, which determine the values of all object-language sentences, some valuations of s-atoms yield paradoxes, which preclude a consistent valuation of some sentences. This is one of the features distinguishing the object-language from the metalanguage. Another is that some valuations of atoms may still leave some sentences of the metalanguage, like the truth-teller, without any definite value. Considering such phenomena as defects of the model may stem from conflating the metalanguage and the object-language, inherent in coding the former as a subset of the latter. In LSO, just as in natural language, these phenomena only display differences between the two language levels. Furthermore, they can be avoided by using the language (that is, valuating s-atoms) in a careful manner.

This poses a peculiar challenge. In AST, self-reference arises in specific theories (with a sufficiently strong arithmetic) via the diagonalization lemma. In LSO, self-reference is a built-in feature of the language. An s-quantifier in a sentence ranges over all sentences, including this very sentence. For example, sentence $\forall\phi(K\phi \to \neg\phi)$ includes the instance $K(\forall\phi(K\phi \to \neg\phi)) \to \neg\forall\phi(K\phi \to \neg\phi)$. Self-reference, allowing the expression of paradoxes, presents a threat to consistency in both cases. AST retains consistency by various restrictions on Convention T, while LSO by the adequate valuations of atoms. This leads to the question of whether the LSO language itself – not just some theory over it – is consistent; whether any valuation of atoms can be extended to a consistent valuation of all sentences of the language respecting its semantic restrictions. The existence of such valuations is ensured by the non-trivial Theorem 3.20.

The consistency of the language modelled in LSO reflects the informal intuition that paradoxes do not reside in the language as such but in the ways we use and discuss it. This model aligns with the diagnosis in [9], which states that paradoxes arise from accepting certain assumptions that, upon closer examination, reveal a contradiction.

LSO embodies this idea through paradoxes arising exactly due to the way operators are defined by valuations of s-atoms, that is, by their interpretations. These amount to specific axioms – assumptions or claims, with the difference between them being inconsequential in the present context.

This view opposes Tarski's diagnosis from [44], voiced also in [9] or more recently in [3, 5, 15, 29], according to which *"everyday language for which the normal laws of logic hold, must be inconsistent"*. This uncomfortable diagnosis arises from accepting predicates on arithmetized syntax as an adequate model of the (everyday) metalanguage, to which the inconsistency of unrestricted Convention T is then ascribed. In LSO, a paradox relies on some assumptions or claims, represented by the interpretation of the involved operators, and can be avoided by avoiding the unfortunate claims. If nobody claims to be lying, no liar paradox results, and the language remains consistent.

To this, however, one may want to object. Statements like L :"This sentence is false" cause trouble simply by existing; they don't have to be asserted to create a problem. Well, just by being there, the liar L makes a claim. Usually, it is represented by $L \leftrightarrow \neg\mathcal{T}(L)$ or simply $L \leftrightarrow \neg L$, expressing the pretense to truth of this unsatisfiable equivalence, the liar's semantic claim. (We might say "truth condition" if this expression did not carry too heavy connotations.) The liar L and its claim "L is true iff it is false" are distinct sentences, and the paradox amounts to the nonexistence of a boolean value for the former satisfying the latter. Colloquial interaction dismisses the liar so easily, in spite of its intriguing challenge, because regardless of the truth-value of the pronounced sentence, the claim hidden behind it is false. Its analysis, starting from the assumption that L must be true or false, arrives in each case at a contradiction, which seems then somehow inherent in the language. In a way, LSO turns such an analysis around, starting with the inadmissibility of the claim, which renders the truth-value of the liar insignificant. Formally, the liar in LSO is Karen who claims only to be always lying. This makes it clear that the source of paradox is the implausible metalevel claim – a valuation of s-atoms that cannot be extended consistently to the full language. Adjusting this valuation solves the problem. Although a paradox implies inconsistency, the mere expressibility of paradoxes does not; it merely allows an unconventional use of the language. Everyday language with the normal laws of logic may still be consistent.

In AST, consistency is ensured only by rejecting enough instances of Convention T. Such restrictions seem to have no counterpart in natural reasoning, where expressing a paradox does not lead to the explosive suspension of rational discourse. The non-explosiveness of LSO, retaining all other elements of classical reasoning, may be a more balanced response. Whether LSO can capture its own syntax and truth theory remains to be seen, and a comparison to such theories over AST must wait until this is clarified. The current differences, concerning self-reference and paradoxes, can be summarized as follows.

1. Self-reference in LSO is a feature of the language, while AST achieves it by specific theories.
2. The syntax coding in AST makes the metalanguage a subset of the object-language. In LSO, operators distinguish syntactically between the two without any coding. Semantically, they differ in that a valuation of atoms determines a valuation of

all sentences of the object-language but may yield paradoxes or fail to assign any truth-value to some sentences of the metalanguage.
3. Paradoxes in AST arise from extending Convention T to too many sentences. In LSO, they arise only in the metalanguage due to some interpretations of operators.
4. Paradoxes in AST are just special cases of contradictions, distinguished at most from outside of the system. In LSO, contradictions have no (covering) models, while paradoxes have some, though only partial ones that have no consistent extension to the full language.
5. This enables non-explosive LSO reasoning in the presence of paradoxes, incorporating the classical system LK. Limitations of (cut) reflect here the reluctance of informal reasoning to conclude anything from the liar except that it is lying and not lying. Unrestricted (cut) extends this to explosive reasoning, where the liar entails everything, as it does in AST.
6. The explosive semantics of LSO is a special case of the non-explosive one, and the two coincide for consistent theories. Of course, there is no such distinction in AST.
7. LSO appears to be capable of including modal operators, though this requires further work. AST can be combined with modal logic, but itself can only represent modalities by predicates.

Sections 2-6 contain the main presentation, while the appendices in Sections 7-8 include the proofs and needed technicalities. Section 2 introduces an expansion of any classical language with sentential quantifiers and operators, as well as the reasoning system LSO for such an expansion of any countable first-order language. Section 3 presents the semantics based on language graphs, their kernels, and semikernels, and briefly discusses their relation to the classical semantics. It contains the main results of the paper: the consistency of the full language $\mathcal{L}^+$ with sentential quantifiers and operators, as well as soundness and completeness of LSO for the non-explosive semantics of FOL$^+$. Section 4 relates the non-explosive semantics to the explosive one and shows that extending LSO with (cut) yields a sound and complete system LSO^c for the latter. Moreover, if a contradiction follows from a theory in LSO^c, some follows already in LSO: reasoning with (cut) does not introduce any new paradoxes but only makes paradoxes explode like usual contradictions. We comment on the apparent similarity and the actual differences between (the reasoning in) LSO and non-transitive systems, e.g., [10, 39, 40]. Section 5 analyzes a series of examples from the literature and outlines some extensions of LSO. Section 6 summarizes the encountered features of paradoxes.

2 Reasoning about Sentences

A classical (propositional, first- or higher-order) language $\mathcal{L}$ is first expanded to $\mathcal{L}^\Phi$ by adding sentential variables, *s-variables* Φ, which can be *s-quantified*. For example, $\forall\phi\forall x(A(x) \vee \phi)$ is a sentence, with s-quantifier $\forall\phi$ and o-quantifier $\forall x$. (The prefix "o-" restricts a syntactic category to the object-language $\mathcal{L}$, while "s-" to the sentences or the metalanguage.) To this, we add *operators* applicable to sentences, allowing the construction of sentences like $\forall x\forall\phi(A(x) \vee \phi \vee P(\phi))$, for an operator P. (In what follows, we will ignore the propositional case, which can be derived through obvious simplifications.) With X denoting a set of *o-variables* (used in $\mathcal{L}$) and $\mathbf{C}$ a set of s-constants, the formulas of $\mathcal{L}^+$ are defined inductively:

1. Every $\mathcal{L}$ formula is an $\mathcal{L}^+$ formula (the object-language is a subset of the meta-language).
2. All s-constants and s-variables, $\mathbf{C} \cup \Phi$, are $\mathcal{L}^+$ formulas.
3. If K is an n-ary operator and $\phi_1 ... \phi_n$ are $\mathcal{L}^+$ formulas, then so is $K(\phi_1 ... \phi_n)$.
4. If ϕ, ψ are $\mathcal{L}^+$ formulas, then so are $\neg\phi$ and $\phi \wedge \psi$.
5. If $x \in X$, $\phi \in \Phi$ and F is an $\mathcal{L}^+$ formula, then $\forall x F$ and $\forall\phi F$ are $\mathcal{L}^+$ formulas.

A *sentence* is a formula without any free (object or sentential) variables. S-constants $\mathbf{C}$ can serve as sentence names. In point 3, operators can also take terms from $\mathcal{L}$ as arguments, but we simplify the presentation skipping their straightforward treatment. Syntactically, open formulas can stand in the argument positions of the operators, but semantically, the operators are applied only to the appropriate closed instances. Atoms are divided into

(a) o-atoms $\mathbf{A}_X$, the atomic formulas of $\mathcal{L}$, and
(b) s-atoms, $\mathbf{A}^\circ_X = \mathbf{A}^+_X \backslash \mathbf{A}_X$, s-constants $\mathbf{C}$ or formulas of type 3 above (not s-variables Φ).

For a set M, by $\mathbf{T}_M$ we denote the free algebra of $\mathcal{L}$ terms over M, by $\mathbf{S}_M/\mathbf{S}^+_M$ all $\mathcal{L}/\mathcal{L}^+$ sentences over $\mathbf{T}_M$, and by $\mathbf{S}/\mathbf{S}^+$ all $\mathcal{L}/\mathcal{L}^+$ sentences. The superscript $\_^\circ$ marks the metalevel, added to the object-level $\mathcal{L}$ and yielding the resulting expansion $\_^+$, e.g., $\mathcal{L}^\circ = \mathcal{L}^+ \setminus \mathcal{L}$, $\mathbf{S}^\circ_M = \mathbf{S}^+_M \setminus \mathbf{S}_M$, etc.

By FOL$^+$ we denote any FOL language expanded as above. The reasoning system LSO for FOL$^+$, given below, extends LK with two rules for s-quantifiers. The basic syntax uses only $\{\wedge, \neg, \forall\}$, with other connectives and $\exists$, and rules for them, defined in the classical way. Sequents, written $\Gamma \Rightarrow \Delta$, are formed from countable sets $\Gamma \cup \Delta$ of $\mathcal{L}^+$ formulas. $\Gamma \vdash \Delta$ denotes provability of $\Gamma \Rightarrow \Delta$ in LSO.

$$(\mathrm{Ax}) \quad \Gamma \vdash \Delta \quad \text{for } \Gamma \cap \Delta \neq \emptyset$$

$$(\neg_L)\ \frac{\Gamma \vdash \Delta, A}{\neg A, \Gamma \vdash \Delta} \qquad (\neg_R)\ \frac{A, \Gamma \vdash \Delta}{\Gamma \vdash \Delta, \neg A}$$

$$(\wedge_L)\ \frac{A, B, \Gamma \vdash \Delta}{A \wedge B, \Gamma \vdash \Delta} \qquad (\wedge_R)\ \frac{\Gamma \vdash \Delta, A \qquad \Gamma \vdash \Delta, B}{\Gamma \vdash \Delta, A \wedge B}$$

$$(\forall_L)\ \frac{F(t), \Gamma, \forall x F(x) \vdash \Delta}{\Gamma, \forall x F(x) \vdash \Delta}\ \text{legal } t/x \text{ in } F \qquad (\forall_R)\ \frac{\Gamma \vdash \Delta, F(y)}{\Gamma \vdash \Delta, \forall x F(x)}\ \text{fresh } y$$

$$(\forall^+_L)\ \frac{F(S), \Gamma, \forall\phi F(\phi) \vdash \Delta}{\Gamma, \forall\phi F(\phi) \vdash \Delta}\ \text{any } S \in \mathbf{S}^+ \qquad (\forall^+_R)\ \frac{\Gamma \vdash \Delta, F(S) \quad \text{– for all } S \in \mathbf{S}^+}{\Gamma \vdash \Delta, \forall\phi F(\phi)}$$

The form of (Ax) makes weakening admissible, while using sets implicitly introduces rules for contraction and permutation, which can be added if sequences or multisets are used. The infinitary rule $(\forall^+_R)$ indicates a lack of compactness, but should not present any challenge for the intuition.[1] It reflects the unrestricted substitution class

[1] This infinitary rule could be replaced by a counterpart of $(\forall_R)$, introducing a fresh s-variable, provided that the semantics addresses not only the actual language, but also its expansions with s-constants.

for the interpretation of s-quantifiers, comprising all sentences of $\mathcal{L}^+$, unlike in other substitutional approaches to sentential quantification (e.g. [23], or more recent [2]), which restrict this class to avoid problematic self-reference.

Infinite sequents allow us to handle some cases of infinite axiomatizations. For instance, making only finitely many statements requires infinitely many premises to exclude making all others. This case has a finite representation using operator $\doteq$ for syntactic equality of sentences, *s-equality*. For instance, 'Karen saying only S' is expressible as $K(S) \wedge \forall\phi(K\phi \rightarrow \phi \doteq S)$, or $\forall\phi(K\phi \leftrightarrow \phi \doteq S)$, abbreviated by $K!\phi$. The following rules for s-equality suffice.

$$\text{(ref)}\ \frac{S \doteq S, \Gamma \vdash \Delta}{\Gamma \vdash \Delta} \qquad \text{(rep)}\ \frac{A(S), A(Q), S \doteq Q, \Gamma \vdash \Delta}{A(Q), S \doteq Q, \Gamma \vdash \Delta} \qquad \text{(neq)}\ \frac{\Gamma \vdash \Delta, Q \doteq S}{\Gamma \vdash \Delta}\ Q \not\equiv S.$$

We primarily consider sequents with no free s-variables, but they can be useful for handling s-equalities. Claims like 'property D holds for each sentence, possibly except $S_1, \ldots, S_n$' are now finitely expressible as $\forall\sigma(\sigma \doteq S_1 \vee \ldots \vee \sigma \doteq S_n \vee D(\sigma))$, allowing sometimes to establish $\forall\phi D(\phi)$ by a finite case analysis, instead of $(\forall^+_R)$:

$$(\forall^{\doteq}_R)\ \frac{\Gamma \vdash \Delta, D(S_1) \ldots \Gamma \vdash \Delta, D(S_n) \qquad \Gamma \vdash \Delta, D(\sigma), \sigma \doteq S_1, \ldots, \sigma \doteq S_n}{\Gamma \vdash \Delta, \forall\phi D(\phi)}\ \text{fresh } \sigma \in \Phi.$$

Properties of LSO, like soundness/completeness, the interaction of its object-level and metalevel, the role of (cut), will be discussed along with the corresponding semantic notions.

3 Semantics

We keep the presentation focused on FOL, but the semantic definitions and results of this section work with minimal adjustments for the propositional or higher-order classical logics. Roughly, any classical interpretation of $\mathcal{L}$ in a structure M is extended to an interpretation of $\mathcal{L}^+$ by providing a valuation of s-atoms and interpreting s-quantifiers substitutionally, e.g.:

$$M \models \forall\phi(P(\phi) \vee \phi) \Leftrightarrow \text{for all } S \in \mathbf{S}^+ : M \models P(S) \vee S. \tag{3.1}$$

The right-hand side has instances like $P(\forall\phi P(\phi) \vee \phi) \vee (\forall\phi(P(\phi) \vee \phi)$, involving the definiendum. Such circularities are handled formally using graphs and their (semi)kernels.

3.1 Kernels and Semikernels

By "graph" we mean a directed graph $G = (\mathbf{V}_G, \mathbf{E}_G)$, with vertex set $\mathbf{V}_G$ and $\mathbf{E}_G \subseteq \mathbf{V}_G \times \mathbf{V}_G$, dropping the subscript $\_G$ when an arbitrary or fixed graph is addressed. For a binary relation R, we let $R(x) = \{y \mid R(x, y)\}$ and extend this notation pointwise to sets, $R(X) = \bigcup_{x \in X} R(x)$. The converse $\mathbf{E}^-$ of the edge relation $\mathbf{E}$ is obtained by flipping the directions of all edges. A *neighbour* y of a vertex x is either its *out-neighbour*, with an edge from x, i.e., $y \in \mathbf{E}(x)$, or its *in-neighbour*, with an edge to x, $y \in \mathbf{E}^-(x)$. A *path* is a (typically finite) sequence of vertices $x_1x_2 \ldots x_n$, where each x_{i+1} is an out-neighbour of the preceding x_i. A path is *acyclic* if each vertex occurs at most once. A *kernel* (or *solution* in [47]) of G is a subset $K \subseteq \mathbf{V}$ such that $\mathbf{E}^-(K) = \mathbf{V} \setminus K$. This equality amounts to the requirement that a kernel K is

(a) independent, i.e., $\mathbf{E}^-(K) \subseteq \mathbf{V} \setminus K$ (no edges between vertices in K), and
(b) absorbing, i.e., $\mathbf{E}^-(K) \supseteq \mathbf{V} \setminus K$ (each vertex outside K has an edge to some vertex in K),

Equivalently, a kernel is an assignment $\kappa \in \mathbf{2}^{\mathbf{V}}$, for $\mathbf{2} = \{\mathbf{0}, \mathbf{1}\}$, such that

$$\forall x \in \mathbf{V} : \kappa(x) = \mathbf{1} \Leftrightarrow \forall y \in \mathbf{E}(x) : \kappa(y) = \mathbf{0}, \tag{3.2}$$

i.e., a vertex is $\mathbf{1}$ if all its out-neighbours are $\mathbf{0}$, and is $\mathbf{0}$ if at least one of its out-neighbours is $\mathbf{1}$. This allows to view each edge as the negation of its target, and branching as the conjunction of such negations. Given (3.2), the set $\{x \in \mathbf{V} \mid \kappa(x) = \mathbf{1}\}$ satisfies (a) and (b), while if K satisfies (a), (b) then $\kappa \in \mathbf{2}^{\mathbf{V}}$ given by $\kappa(x) = \mathbf{1} \Leftrightarrow x \in K$ satisfies (3.2). We therefore do not distinguish the two and by $sol(G)$ denote the set of kernels or such assignments. Graph G is *solvable* if $sol(G) \neq \emptyset$.

According to Richardson's theorem from [38], the absence of odd cycles ensures the solvability of graphs that have no infinite branchings or no infinite outgoing acyclic paths. In particular, a finite graph without odd cycles is solvable. Our graphs, introduced in the next subsection, are solvable, but model paradoxes by unsolvable subgraphs. They arise mainly due to the odd cycles, but also acyclic Yablo-like paradoxes are captured and discussed in Section 5.5.

A *sink* is a vertex with no outgoing edges, and sinks belong to every kernel by (b). One can think of them as the basic truths, independent from the rest of the graph.

A subset $L \subseteq \mathbf{V}$ *covers* vertices in L and those pointing to L, that is, $L \cup \mathbf{E}^-(L)$, denoted by $\mathbf{E}^-[L]$. The equation $\mathbf{E}^-(K) = \mathbf{V} \setminus K$ implies that a kernel K covers the entire graph. A valuation is *coherent* on vertices for which it satisfies (3.2), so a kernel represents a coherent valuation of all sentences. Our semantics also allows for more general situations that are only locally coherent and cannot be extended to the entire graph. In the absence of a kernel, a relevant part of the graph may still be coherently covered by a *semikernel* (see [28]), namely, a subset $L \subseteq \mathbf{V}$ such that

$$\mathbf{E}(L) \subseteq \mathbf{E}^-(L) \subseteq \mathbf{V} \setminus L. \tag{3.3}$$

The set of semikernels of G is denoted by $SK(G)$. In terms of assignments, a semikernel L determines a coherent valuation of vertices it covers, satisfying (3.2) for every $x \in \mathbf{E}^-[L]$: each statement denied by any true one (in L) is false (in $\mathbf{E}^-(L)$), while every false statement (in $\mathbf{E}^-(L)$) denies some true one (in L). Although locally consistent, a coherent situation represented by a semikernel can involve or even entail inconsistency.[2]

[2] The branch of argumentation theory arising from [13] shares only its origins in a similar reading of digraph (semi)kernels. Links to reference (or dependency) graphs, used for paradox analysis in [6, 7, 34], although closer, are not essential either. The relevance of kernels for paradox analysis, noticed in [11], was extended to semikernels and used in [14] for modelling non-explosiveness and then combined with propositional reasoning by resolution in [48, 50]. The present paper extends this to FOL and models of the whole language, not only of particular theories.

Every graph possesses a semikernel, since $\emptyset$ satisfies trivially (3.3). But semikernels of interest are nonempty, also in graphs not possessing any kernel.

Example 3.4 The propositions to the left can be represented by the graph to the right.

a_1 : This and the next sentence are false.

b_1 : The next sentence is false.

b_2 : The previous sentence is false.

a_2 : This and the previous sentence are false.

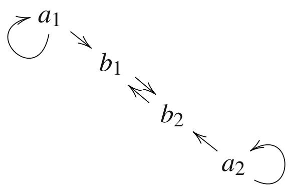

The graph has no kernel, witnessing to the involved paradox, but has two semikernels: $\{b_1\}$ and $\{b_2\}$ covering, respectively, $\{a_1, b_1, b_2\}$ and $\{b_1, b_2, a_2\}$, which can be seen as coherent subdiscourses, where truth-values can be assigned coherently, i.e., respecting (3.2).

3.2 Language Graphs

The graph in Example 3.4 represents a specific theory, while a language graph represents all sentences of $\mathcal{L}^+$, relatively to a given $\mathcal{L}$-*domain*, that is, a nonempty set M with a standard classical interpretation of $\mathcal{L}$-terms $\mathbf{T}_M$, but not of the predicate symbols. For each such a domain, there is one language graph $G_M(\mathcal{L}^+)$ representing the syntax of $\mathcal{L}^+$ over M, with sentences $\mathbf{S}_M^+$ as vertices. A (semi)kernel represents then a subset of sentences true in M under the valuation assigning $\mathbf{1}$ exactly to the *literals* (atoms or negated atoms) belonging to the subset. The structure of the graph, combined with the semikernel condition (3.3), ensures that the evaluation of the object-level sentences covered by L coincides with their classical semantics, generalizing this to self-referential sentences. We start by sketching the main ideas.

Each $S \in \mathbf{S}_M^+$ is the *source* (vertex reaching all others) of the subgraph $G_M(S)$ of $G_M(\mathcal{L}^+)$.

1. Dual literals form 2-cycles, e.g., $P(m) \leftrightarrows \neg P(m)$, for a predicate symbol P or, for an operator symbol K and sentences Q, R: $K(R, Q) \leftrightarrows \neg K(R, Q)$. Also s-constants $\mathbf{C}$ form such 2-cycles with their negations. Consequently, in each kernel exactly one of the literals is $\mathbf{1}$ and the other $\mathbf{0}$.
2. Out-branching represents conjunction (or universal quantification), and each edge negation of its target. For instance, the source of $G_M(\neg A)$ has a single outgoing edge $\neg A \rightarrow A\ldots$, while the source of $G_M(B \wedge C)$ has two: $\ldots B \leftarrow \neg B \leftarrow B \wedge C \rightarrow \neg C \rightarrow C\ldots$
   Subgraph $G_M(S)$ of each sentence S without s-quantifiers, in particular of each object-level sentence, can be seen as a *tree* (the source vertex S with no incoming edges and a unique path to every other vertex), except that instead of leaves (with no outgoing edges) there are atomic 2-cycles. It reminds of S's parse tree but, primarily, captures the semantics of the formula constructors ($\neg$, $\wedge$, $\forall$) in terms of kernels. Using (3.2), one can easily verify that the source $B \wedge C$ above is $\mathbf{1}$ exactly when $B = \mathbf{1} = C$. The graph for $A \to B$, obtained from $A \to B \Leftrightarrow \neg(A \wedge \neg B)$,

is $\begin{tikzcd}[column sep=small] \boxed{A \to B} \arrow[r] & \boxed{A \wedge \neg B} \arrow[r] \arrow[rr, bend right] & B\ldots & \boxed{\neg A} \arrow[r] & A\ldots \end{tikzcd}$ By (3.2) its source is $\mathbf{1}$ exactly when $A = \mathbf{0}$ or $B = \mathbf{1}$.

3. The source of the subgraph for a sentence with the universal o-quantifier has an edge to the negation of every instance, e.g., vertex $\forall x P(x)$ has an edge to $\neg P(m)$ for every $m \in M$.
   The collection $\{G_M(S) \mid S \in \mathbf{S}_M\}$ forms the subgraph $G_M(\mathcal{L})$ of $G_M(\mathcal{L}^+)$.
4. Similarly to the sentences starting with the universal o-quantifier, those starting with such an s-quantifier have an edge to the negation of every instance. This requires more care, and details are explained further down. The crucial point is that starting from, say, $S = \forall\phi(K(\phi) \wedge \phi)$, also $K(S) \wedge S$ is an instance of S. The tree is formed as for the sentences without s-quantifiers, but the instantiations are, so to speak, suspended until we reach the leaves of the tree. If sentence Q, instantiating an s-variable quantified in the original sentence, eventually becomes such a leaf, it is (identified with) the source of the subgraph $G_M(Q)$, as shown in (3.5) below for the edge from $\overline{S}$. ($\overline{X}$ stands for $\neg X$.) The literals have only 2-cycles to their duals, as shown for $K(S)$ and $K(Q)$.

$$\begin{tikzcd}
\forall\phi(K(\phi) \wedge \phi) \arrow[d] \arrow[rrd] \arrow[rrrrd, dotted] & & & & \\
\overline{K(S) \wedge S} \arrow[d] & & \overline{K(Q) \wedge Q} \arrow[d] & & {} \arrow[d, dotted] \\
K(S) \wedge S \arrow[d] \arrow[rd] & & K(Q) \wedge Q \arrow[d] \arrow[rd] & & \ldots \\
\overline{S} \arrow[uuu, bend left] & \overline{K(S)} \leftrightarrows K(S) & \overline{Q} & \overline{K(Q)} \leftrightarrows K(Q) & \ldots
\end{tikzcd} \tag{3.5}$$

By (3.2), the source $\forall\phi(K\phi \wedge \phi)$ is $\mathbf{1}$ iff, for each $X \in \mathbf{S}^+_M$, $\overline{K(X) \wedge X} = \mathbf{0}$, i.e., $K(X) \wedge X = \mathbf{1}$. This is, of course, falsified by any false sentence X, e.g., from the object-level $\mathcal{L}_M$. If Q in the drawing is an $\mathcal{L}_M$ sentence, vertex $\overline{Q}$ belongs to the subgraph $G_M(\mathcal{L})$ and has an edge to the source of the tree $G_M(Q)$ there. The whole graph $G_M(\mathcal{L}^+)$ has the three main subgraphs:

G1 is subgraph $G_M(\mathcal{L})$ for the object-language,
G2 extends G1 with 2-cycles for s-atoms $\mathbf{A}^\circ_M$, their propositional combinations, also with sentences from G1, and with the o-quantification of such combinations,
G3 contains $G_M(A)$ for each sentence A with s-quantifiers.

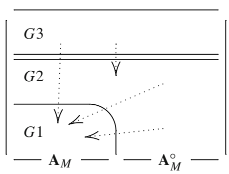

(3.6)

Subgraph G2, containing G1, is a collection of trees, with the atomic 2-cycles at all leaves. These 2-cycles provide only the possibility of different valuations of atoms, so subgraph G2 is essentially acyclic, following the inductive definition of the language. In subgraph G3, however, there is a path between each pair of source vertices, forming multiple cycles, as will be explained below. Dotted arrows indicate edges between these subgraphs, going only from G3 to G2/G1 and from G2 to G1, but never in the opposite directions. By our notational convention $G_M(\mathcal{L}^\circ) = G_M(\mathcal{L}^+) \setminus G1$.

Definition 3.7 The *language graph* $G_M(\mathcal{L}^+)$, for a language $\mathcal{L}^+$ and $\mathcal{L}$-domain M, is given by:

1. Vertices $\mathbf{V} = \mathbf{S}^+_M$.
2. Each atomic sentence $A \in \mathbf{A}^+_M$ has a 2-cycle to its negation: $A \leftrightarrows \overline{A}$.
3. Each nonatomic sentence $S \in \mathbf{S}^+_M$ is the source of the subgraph $G_M(S)$ with the outgoing edges:

$G_M(S)$: *source*		*with edges to the source of*:
$(a)\quad \neg F$	$\rightarrow$	$G_M(F)$,
$(b)\quad F_1 \wedge F_2$	$\rightarrow$	$G_M(\neg F_i),\ \ for\ i \in \{1, 2\}$,
$(c)\quad \forall x Fx$	$\rightarrow$	$G_M(\neg F(m)),\ \ for\ each\ m \in M$,
$(d)\quad \forall \phi F\phi$	$\rightarrow$	$G_M(\neg F(S)),\ \ for\ each\ S \in \mathbf{S}^+$.

When $\mathcal{L}$ is higher-order, the only difference is the domain M, containing required sets, with object quantifier(s) in point 3.(c) being those from $\mathcal{L}$. When it is inessential, we often drop M and write $G(\mathcal{L}^+)$ for $G_M(\mathcal{L}^+)$. By $\mathcal{LG}r(\mathcal{L}^+)$ we denote the class of all language graphs for language $\mathcal{L}^+$.

As can be expected (and as shown in Section 7.2), every $\mathcal{L}^+$ sentence has a prenex normal form, and its quantifier-free matrix a disjunctive normal form, yielding the normal form PDNF, used for graph construction. The universal and existential quantifiers give rise to the following branchings to the instantiations of the quantified o-variables by all elements $a, b, c\ldots$ of the domain, and of the s-variables by all $\mathbf{S}^+$.

$$
\begin{tikzcd}
 & \forall\phi D(\phi) \arrow[dl] \arrow[d] \arrow[dr] \arrow[drr, dotted] & & \\
\overline{D(S_1)} \arrow[d] & \overline{D(S_2)} \arrow[d] & \overline{D(\forall\phi D(\phi))} \arrow[d] & \ldots \arrow[d, dotted] \\
D(S_1) & D(S_2) & D(\forall\phi D(\phi)) & \ldots
\end{tikzcd}
\qquad
\begin{tikzcd}
 & \exists x D(x) \arrow[d] & & \\
 & \forall x \neg D(x) \arrow[dl] \arrow[d] \arrow[dr] \arrow[drr, dotted] & & \\
D(a) & D(b) & D(c) & \ldots
\end{tikzcd}
\tag{3.8}
$$

The right graph illustrates a simplification performed often without mention. Strictly speaking, $\forall x\neg D(x)$ has an edge to $\neg\neg D(a)$, for every a, forming a 3-path $\forall x\neg D(x) \rightarrow \overline{\overline{D(a)}} \rightarrow \overline{D(a)} \rightarrow D(a)$. Now, any *double edge* $x \rightarrow y \rightarrow z$, where x has no other out-neighbours and y no other neighbours, can be contracted by removing y and identifying $x = z$, virtually without changing (semi)kernels, Fact 7.2. Such 3-paths are usually contracted to single edges, here, identifying $\overline{\overline{D}}(a)$ and $D(a)$.

The quantifier prefix is converted to the graph by successively performing such instantiations and branchings until no quantified variables remain. At the end of each branch of instantiations of all variables quantified in the original sentence, there remains the subgraph for its instantiated DNF matrix, a *DNF-foot*. For example, DNF matrix $D(\phi, x) = (\neg\phi \wedge \neg Q(x)) \vee (\neg P(\phi) \wedge R)$, with a ground atom R, instantiated by $S \in \mathbf{S}^+$ and $a \in M$, yields DNF-foot $G_M(D(S,a))$:

$$\begin{tikzcd}
D(S,a) \arrow[rrr] & & & \circ_{\overline{D(S,a)}} \arrow[dll] \arrow[dr] & & \\
& C_1(S,a) \arrow[d] \arrow[dr] & & & C_2(S,a) \arrow[d] \arrow[dr] & \\
& S & Q(a) \arrow[d, shift left] & & P(S) \arrow[d, shift left] & \overline{R} \arrow[d, shift left] \\
& \neg S \arrow[u] & \overline{Q(a)} \arrow[u, shift left] & & \overline{P(S)} \arrow[u, shift left] & R \arrow[u, shift left]
\end{tikzcd} \tag{3.9}$$

The original closed atoms and the expressions resulting from the susbstitutions into the atoms of the original matrix are the *leaves* of the foot; here: S, $Q(a)$, $P(S)$, R. (Typically, they have outgoing edges and may be complex sentences.) Vertex $\circ_{\overline{D(S,a)}}$ is sentence $\neg C_1(S,a) \wedge \neg C_2(S,a) \in \mathbf{S}^+_M$, while $C_1(S,a)$ is sentence $\neg S \wedge \neg Q(a)$. For $L \in SK(G_M(D(S,a))) : D(S,a) \in L \Leftrightarrow \circ_{\overline{D(S,a)}} \notin L \Leftrightarrow C_1(S,a) \in L \vee C_2(S,a) \in L \Leftrightarrow \{\neg S, \neg Q(a)\} \subseteq L \vee \{\neg P(S), R\} \subseteq L$, reflecting the expected $D(S,a) = \mathbf{1} \Leftrightarrow (S = \mathbf{0} = Q(a)) \vee (P(S) = \mathbf{0} \wedge R = \mathbf{1})$.

The full graph $G_M(\mathcal{L}^+)$ has, besides the essentially acyclic G2 described above, also subgraph G3 containing subgraphs $G_M(S)$ for the sentences with s-quantifiers. These form the main source of complexity. An s-variable ϕ can occur in a sentence in a *nominal position*, that is, within the scope of some operator, or outside any such scope, in a *sentential position*, e.g., ϕ in $C_1 = \neg\phi \wedge \neg Q(x)$. For an s-quantified S, each $A \in \mathbf{S}^+$ instantiating, according to Definition 3.7.3.(d), an s-variable ϕ in a sentential position in S, becomes a leaf of $G_M(S)$, that is, of its DNF-foot (3.9). This leaf is the source of subgraph $G_M(A)$. In particular, sentence S also instantiates ϕ, and the resulting leaf is actually the source of this very $G_M(S)$, as in (3.5).[3] Every $A \in \mathbf{S}^+$, instantiating ϕ in a sentential position in S, either occurs as an internal vertex, i.e., on some path from the source S to some source of a DNF-foot, or not. In the former case, the leaf A is called an *internal leaf* of $G_M(S)$, possibly forming a cycle. In the latter case A occurs in $G_M(S)$ only as a leaf and is its *external leaf*, $extrG_M(S)$, as the source of its separate $G_M(A)$. If an external leaf is s-quantified, its subgraph instantiates its s-variables by all sentences, in particular by S, giving paths back to the source of $G_M(S)$. The subgraph $G_M(S)$ is given by

- the source S,
- all paths to all the leaves (at all its DNF-feet),
- 2-cycles at the atoms (occurring in S or instantiating the atomic subexpressions of S),
- cycles from the internal leaves of $G_M(S)$, and
- the external leaves, without their subgraphs.

Ignoring the cycles at the atoms and through the internal leaves, we speak occasionally about such subgraphs as if they were trees.

[3] To keep the tree analogy, especially the notion of leaves, we can think of each $G_M(S)$ as a tree, where leaves like A, S have double edges to the sources of $G_M(A)$, $G_M(S)$. The intermediate vertices duplicate then $\neg A$, $\neg S$, but this does not change (semi)kernels by Fact 7.3, allowing to identify vertices with identical out-neighbourhoods.

All sources of s-quantified sentences belong to one strongly connected component of $G_M(\mathcal{L}^+)$, which is the subgraph G3 from (3.6). Their leaves that are sentences $\mathbf{S}_M$ belong to the subgraph $G_M(\mathcal{L})$, but there are no edges returning from there to G3.

Cycles connecting subgraphs of distinct sentences arise only from sentences substituted for s-variables in sentential positions. Sentences substituted into nominal positions, i.e., into the scope of some operator, form only literals with 2-cycles to their duals, like $P(S) \leftrightarrows \overline{P(S)}$ in (3.9), resulting from substituting S into $P(\phi)$.[4] This difference is exemplified on drawing (3.5).

This completes our description of the language graphs representing the syntax of $\mathcal{L}^+$. The semantics, including the involved self-reference, is captured by their (semi)kernels.

3.3 Satisfaction Relation and the (semi)kernel Models

A (semi)kernel L of a language graph represents sentences that are satisfied under the valuation of atoms determined by the literals contained in L, as indicated by the evaluation under graph (3.9). Here is a more intricate example involving s-quantified sentences.

Example 3.10 Let $S_1, S_2, \ldots$ stand for all $\mathbf{S}^+$, except the iterated negations of the source sentences in each graph sketched below: $G(A_\forall)$, for $A_\forall = \forall\phi.\phi$, and $G(A_\exists)$, for $A_\exists = \exists\phi.\phi = \neg\forall\phi.\neg\phi$:

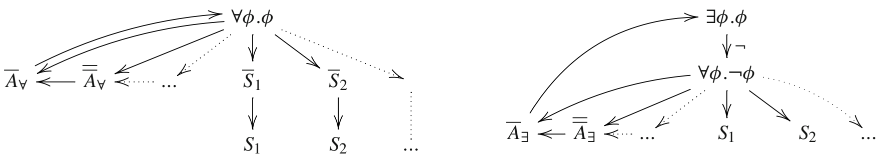

The drawings indicate only the essential aspects, ignoring other edges and cycles.

In the graph on the left, the source $\forall\phi.\phi$ is an internal leaf, while all S_i are external ones. Any $S_i \in \mathbf{S}^+$ valuated to $\mathbf{0}$ yields $\overline{S}_i = \mathbf{1}$ and $\forall\phi.\phi = \mathbf{0}$, but even if all $S_i = \mathbf{1}$, the mere cycles involving $\overline{A}_\forall$ and $\overline{\overline{A}}_\forall$ force $\forall\phi.\phi = \mathbf{0}$. To obtain a kernel, the odd cycle via $\overline{\overline{A}}_\forall$ must be broken, i.e., some of its vertices must have an out-neighbour $= \mathbf{1}$. If all $\overline{S}_i = \mathbf{0}$, this still happens when $\overline{A}_\forall = \mathbf{1}$, making $\overline{\overline{A}}_\forall = \mathbf{0} = \forall\phi.\phi$. We might say that the sentence "All sentences are true", $\forall\phi.\phi$, is a counterexample to its own truth.

In the dual situation of $G(A_\exists)$, a kernel requires breaking the odd cycle via $\forall\phi.\neg\phi$ and $\overline{A}_\exists$. This happens if any $S_i = \mathbf{1}$, making $\exists\phi.\phi = \mathbf{1}$, but even if all S_i were $\mathbf{0}$, the only way of breaking this odd cycle is with $\overline{\overline{A}}_\exists = \mathbf{1}$, hence $\exists\phi.\phi = \mathbf{1}$, which provides thus a witness to its own truth.

The truth-value of $\forall\phi.\phi$ depends on the value of this very sentence. Such circular dependencies can hardly be captured by any inductive semantic definitions, while

[4] These 2-cycles are formed only for atoms with the outermost operator. Substituting S into $P(\phi, Q(\phi))$ yields atom $P(S, Q(S))$ with 2-cycle to its dual $\overline{P(S, Q(S))}$. The inner $Q(S)$ does not obtain any edges to its dual $\overline{Q(S)}$ here, but only when atom $Q(S)$ occurs in a sentential position.

semikernels covering this sentence determine its value in line with informal understanding.

We define an $\mathcal{L}^+$-sequent $\Gamma \Rightarrow \Delta$ to be *valid*, $\Gamma \models \Delta$, iff in every language graph $G_M \in \mathcal{LG}r(\mathcal{L}^+)$ every relevant situation satisfies it. A *situation* is a semikernel L, it is *relevant* if it covers $\Gamma \cup \Delta$, i.e., $\Gamma \cup \Delta \subseteq \mathbf{E}^-[L] = \mathbf{E}^-(L) \cup L$, and it *satisfies* the sequent by making some $D \in \Delta$ true, i.e., $\Delta \cap L \neq \emptyset$, or some $G \in \Gamma$ false, i.e., $\Gamma \cap \mathbf{E}^-(L) \neq \emptyset$. (This is generalized to free object-variables $\mathcal{V}(\Gamma, \Delta)$, by considering all assignments $\alpha \in M^{\mathcal{V}(\Gamma,\Delta)}$, with $\alpha(S)$ denoting insertion of $\alpha(v)$ for each free object-variable v in formula S and $\alpha(\Gamma) = \{\alpha(S) \mid S \in \Gamma\}$. For free s-variables, all substitution instances with sentences from $\mathbf{S}^+$ must hold.)

$$\begin{aligned} \Gamma \models \Delta &\iff \forall G_M \in \mathcal{LG}r(\mathcal{L}^+)\ \forall L \in SK(G_M) : L \models \Gamma \Rightarrow \Delta, \text{ where} \\ L \models \Gamma \Rightarrow \Delta &\iff \forall \alpha \in M^{\mathcal{V}(\Gamma,\Delta)} : \\ &\qquad \alpha(\Gamma \cup \Delta) \subseteq \mathbf{E}^-[L] \to (\alpha(\Gamma) \cap \mathbf{E}^-(L) \neq \emptyset) \vee (\alpha(\Delta) \cap L \neq \emptyset). \end{aligned} \tag{3.11}$$

Semikernel models of $\Gamma \subseteq \mathbf{S}^+$ are pairs $(G, L) \in \mathcal{LG}r(\mathcal{L}^+) \times SK(G)$ satisfying $\Rightarrow F$, for all $F \in \Gamma$.

The covering condition $\alpha(\Gamma \cup \Delta) \subseteq \mathbf{E}^-[L]$ brings the important aspect of (ir)relevance. According to it, a semikernel satisfies every sequent/sentence that it does not cover. For instance, the empty semikernel satisfies every sequent and every contradiction $A \wedge \neg A$, but only vacuously, since it does not cover anything.[5] A nonempty semikernel also satisfies A and $\neg A$ for every sentence it does not cover. This makes LSO paraconsistent, but in a degenerate way, as contradictions are satisfied only vacuously, in situations ignoring them. Semikernels covering $A \wedge \neg A$ contain either A or $\neg A$, but never both, so a contradiction is unsatisfiable by any semikernel covering it.

In spite of the vacuous satisfiability of contradictions, even possible validity $\Gamma \models A \wedge \neg A$ for some Γ and A, contradictions entail everything and $\Gamma, A \wedge \neg A \models \Delta$ holds for every A, Γ, Δ. By the previous paragraph, every semikernel satisfying $A \wedge \neg A$ does so only vacuously, not covering $A \wedge \neg A$, so no Γ or Δ falsify any such validity.

Satisfiability of a contradiction is an oxymoron, so we should clarify this notion in the present context. An $S \subseteq \mathbf{S}^+$ is a *contradiction*, $S \in \mathbb{C}$, if it is not contained in any semikernel, i.e., $S \nsubseteq L$, for every language graph G and $L \in SK(G)$. Each classical contradiction is thus an LSO contradiction, and it can be satisfied only vacuously. As a typical example, $A \wedge \neg A$ has the graph $\boxed{A \wedge \neg A} \to \boxed{\neg A} \to A \ldots$ No semikernel L contains $A \wedge \neg A$. If it did then $\neg A \in \mathbf{E}(A \wedge \neg A) \subseteq \mathbf{E}^-(L)$ would force also $A \in L$, contradicting the independence of L.

Dually, $S \subseteq \mathbf{S}^+$ is a *tautology* if it is contained in every covering semikernel, $S \subseteq \mathbf{E}^-[L] \Rightarrow S \subseteq L$. Semikernels not covering S satisfy it vacuously, so a tautology is satisfied by every semikernel. Classical tautologies are LSO tautologies, so these definitions just generalize the classical notions (e.g., $\forall \phi(\neg(\phi \wedge \neg\phi))$, $\forall \phi, \psi((\neg\phi \wedge (\phi \vee \psi)) \to \psi)$ are tautologies). But as the satisfiable contradictions indicate, their status and role are a bit more complex than in classical logic.

[5] Its existence might be philosophically pleasing, e.g., as a representation of total ignorance, accepting both poles of each contradiction. Formally, it plays no role, as most theories also have nonempty models.

It might seem that a semikernel, not covering the whole graph, assigns a third value to all uncovered sentences, but this impression is more wrong than right. In spite of the partiality of the semikernel models and the involved paraconsistency, the semantics is two-valued, since by (3.11) each semikernel determines a boolean value of each sentence, perhaps vacuously by not covering it. If $\neg A$ is not covered by a semikernel L, then neither is A, and L satisfies both, so to say, independently of each other. Satisfaction of both A and $\neg A$ reflects the non-inductive, holistic character of the semantics. A semikernel satisfies simply all sentences it does not cover, in addition to the ones it contains. Although both uncovered A and $\neg A$ obtain thus value $\mathbf{1}$, the logic is not dialetheic, since neither A nor $\neg A$ obtain both values $\mathbf{1}$ and $\mathbf{0}$.

This way of satisfying contradictions underlies the use of semikernels to define models also for theories that, although locally coherent, perhaps even apparently plausible, lead to contradictions, as illustrated by the following example.

Example 3.12 The Karen-liar says only that everything said by Karen is false, $Kl = \forall\phi(K\phi \to \neg\phi)$. Semikernel $L = \{K(Kl)\} \cup \{\overline{K(S)} \mid S \not\equiv Kl\}$, capturing this situation, cannot be extended to any kernel because $K(Kl) = \mathbf{1}$ makes $\overline{K(Kl)} = \mathbf{0}$, while ${}_{X\wedge K(X)} = \mathbf{0}$, for $X \not\equiv Kl$, due to $\overline{K(X)} = \mathbf{1}$. The remaining unresolved 3-cycle $\overline{Kl} - Kl -_{Kl\wedge K(Kl)}$ prevents any extension of L to a kernel.

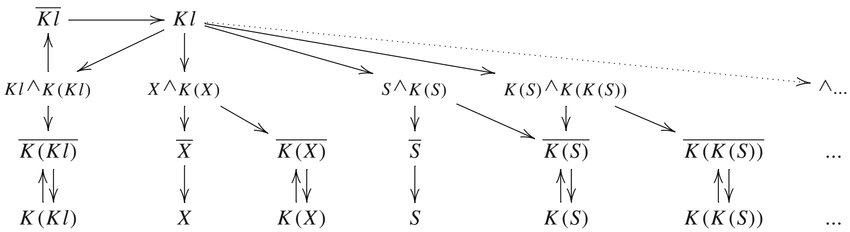

Semikernels provide thus models for some situations that are merely contradictions in classical logic, yielding non-explosive semantics. In the example above, semikernel L can be extended with many other sentences, for instance, with $J(A)$ or with $\neg J(A)$ but not with both. The paradox captured by L implies that Karen lies and does not lie, but, like in an informal analysis of the liar, virtually nothing beyond that. John may still say A and may not say A, but neither follows in LSO from Karen's statement. This suggests defining an *s-contradiction* (or a paradox) as a theory possessing semikernel models but no kernel models. Less formally, a paradox is an apparently meaningful set of statements (possessing a limited, semikernel model) that, at a closer analysis, displays a contradiction (hence possesses no kernel model). A 'closer analysis' means here not so much deriving all consequences (which would unveil the paradox if it were manageable), as rather expanding attention to the relevant statements not included in the original context but demonstrating the implied contradiction. E.g., trying to expand L with Kl or $\neg Kl$ fails, displaying only the contradiction $Kl \wedge \neg Kl$ which is, so to say, relevant for this paradox. A similar expansion with $J(A)$ would succeed, yielding a semikernel showing $K!(Kl) \not\models \neg J(A)$ and hence $K!(Kl) \not\models J(A) \wedge \neg J(A)$. This contradiction is not relevant for this paradox, while the relevant $Kl \wedge \neg Kl$ tells how to resolve it: valuate s-atoms so that Karen always lies or not, but not both.

Kernels are semikernels that cover the entire graph. Consequently, they cannot satisfy contradictions vacuously, yielding a semantics that explodes not only from contradictions but also from paradoxes. We denote it by $\models_c$, with the subscript suggesting its classical explosive nature, which will later be linked to (cut). An $\mathcal{L}^+$-sequent $\Gamma \Rightarrow \Delta$ is *c-valid*, $\Gamma \models_c \Delta$, iff in every language graph $G_M \in \mathcal{LG}r(\mathcal{L}^+)$ every kernel satisfies it. This merely specializes condition (3.11) by replacing semikernels with kernels:

$$\Gamma \models_c \Delta \;\Leftrightarrow\; \forall G_M \in \mathcal{LG}r(\mathcal{L}^+)\ \forall K \in sol(G_M)\ \forall \alpha \in M^{\mathcal{V}(\Gamma,\Delta)} : \\ \alpha(\Gamma) \cap \mathbf{E}^-(K) \neq \emptyset \vee \alpha(\Delta) \cap K \neq \emptyset. \tag{3.13}$$

A *kernel model* of $\Gamma \subseteq \mathbf{S}^+$ is its semikernel model (G, K) with $K \in sol(G)$.

3.4 The Graph versus Classical Semantics

By Definition 3.7, point 3, each complex sentence is a vertex with edges (or paths) to the sentence's components (or instances for quantified sentences), while the value of each vertex depends solely on the values of its out-neighbours by the semikernel condition (3.3), in terms of assignments (3.2). This is the compositional aspect of the semikernel semantics, determining the value of a complex sentence from the values of its out-neighbours, eventually, its components (or instances). This local compositionality interacts, however, with the circular dependencies and evades any usual inductive definition. The holistic element of condition (3.3) lies in the requirement it puts on *the whole* set of vertices (sentences), like the holistic element of consistency requiring simultaneous satisfaction of *all* axioms. Consequently, the non-vacuous satisfaction in (3.11) is given by the membership of the satisfied formulas in the satisfying (semi)kernel and not by an inductive definition (e.g., on the complexity of formulas), which seems rather unlikely to be possible for such circular phenomena.[6]

The restriction of kernel semantics (3.13) to the trees of subgraph G2 in (3.6), i.e., the sentences without s-quantifiers, coincides with the satisfaction defined inductively in the usual way. For each such a sentence A, its subgraph $G_M(A)$ is a tree, except that instead of leaves, there are literals with 2-cycles. Exactly one element from each such cycle is in each kernel, thus capturing the valuations of atoms. Every such tree has exactly one kernel for every valuation of its 2-cycles, obtained by inducing values from such a valuation upwards,[7] reflecting thus the inductive definition of satisfaction. Kernel K of subgraph G2 (in particular, of $G_M(\mathcal{L})$) represents exactly the formulas satisfied (in M) under valuation ρ of atoms given by intersecting each atomic 2-cycle with K, i.e., by setting $\rho(x) = \mathbf{1} \Leftrightarrow x \in K$, for $x \in \mathbf{A}^+_M$ ($x \in \mathbf{A}_M$ for $G_M(\mathcal{L})$).

Example 3.14 Inclusion of $\overline{P(S)}$ and R from (3.9) in a kernel K forces, by independence, $P(S)$ and $\overline{R}$ out of it. This, in turn, forces $C_2(S, a) \in K$ by absorption, so

[6] An inductive definition of a semantics, based on (another use of) graphs in [41], assigns to sentences with circular dependencies sets of equations that are formed, under the specified semantic restrictions, from such sets assigned to the sentence's components. This formally inductive definition merely relocates circularity and non-compositionality to solving sets of equations; hence, it does not seem to contradict the claimed unlikeliness.

[7] Incidentally, this is the first result about kernels from [47], where the concept was introduced.

$\circ_{\overline{D(S,a)}} \notin K$ and $D(S,a) \in K$. This implication from $\{\overline{P(S)}, R\} \subseteq K$ to $D(S,a) \in K$ reflects that from $\neg P(S) \wedge R$ to $D(S,a)$.

There is thus a bijection mapping a FOL structure (M, ρ), with an $\mathcal{L}$-domain M and $\rho \in \mathbf{2}^{\mathbf{A}_M}$, to the language graph with its kernel $(G_M(\mathcal{L}), K_\rho)$, where $A \in K_\rho \Leftrightarrow \rho(A) = \mathbf{1}$ for $\mathcal{L}_M$-atoms A. Then also, for all $S \in \mathbf{S}_M$ (with $\models$ denoting here the standard FOL satisfaction),

$$(M, \rho) \models S \Leftrightarrow S \in K_\rho. \tag{3.15}$$

Unlike kernels, containing *exactly* one element from every atomic 2-cycle, each semikernel contains *at most* one such element, representing thus a partial valuation of atoms. Since atomic 2-cycles have no outgoing edges, every such partial valuation forms a semikernel of the whole graph.

An analogous observation applies to subgraph $G(\mathcal{L})$. No edges go out of it to its complement $G(\mathcal{L}^\circ)$, so every kernel of $G(\mathcal{L})$ is a semikernel of the whole graph $G(\mathcal{L}^+)$. All trees of $G(\mathcal{L})$ have atomic 2-cycles at their leaves, so this subgraph has a unique solution for every choice from these 2-cycles. This reflects the elementary fact that every valuation of atoms determines, by inductive definition, the semantic values of all sentences in classical and most other logics.

The same holds for subgraph G2. It, too, is a collection of trees with the atomic 2-cycles; hence, it is uniquely solvable for every choice from these cycles (the valuation of atoms). It has no edges to its complement G3, hence each kernel of G2 is a semikernel of $G(\mathcal{L}^+)$. This shows the unsurprising fact that adding operators to a classical language does not create any paradoxes (allows to evaluate inductively all sentences under every valuation of all atoms, including the new s-atoms).

What remains is subgraph G3 which contains multiple cycles, many of which are odd. The reader may, even should, wonder if this subgraph hides perhaps some unavoidable paradoxes, making the whole language graph $G(\mathcal{L}^+)$ unsolvable. Its complexity makes its solvability far from obvious. The undefinability of truth underlies various forms of the claim, originating with Tarski's [44], that the natural language actually is inconsistent, e.g., [3, 5, 9, 15, 29]. The theorem is based on the availability of self-reference and the construction of the liar in a sufficiently strong theory. In LSO, self-reference is present in the very language and, as seen in Example 3.12, some valuations of s-atoms make the graph $G(\mathcal{L}^+)$ unsolvable. It might happen that the graph is unsolvable for every valuation of atoms, that is, the language itself is inconsistent.

The following Section 3.5 resolves this worry, showing that $G(\mathcal{L}^+)$ does have a kernel, that is, all sentences of $\mathcal{L}^+$ can be evaluated respecting the classical semantics. First, Theorem 3.16 states that, for a language $\mathcal{L}^\Phi$ with s-quantifiers but no operators, language graphs not only have kernels, but have unique one for every valuation of $\mathcal{L}$ atoms. (Thus, adding such quantifiers to any classical language does not create any paradoxes, just as adding only sentential operators does not. Paradoxes require both operators and s-quantifiers.) Solvability of graphs for the full language $\mathcal{L}^+$ follows, Theorem 3.20, but paradoxes become possible and if they occur then semikernel semantics enables more detailed analysis than kernel semantics.

Section 3.5 is rather technical, referring to even more technical proofs in the Appendix. One can safely skip it on a casual reading, going now to Section 4 with just the two mentioned theorems.

3.5 Solvability of $G(\mathcal{L}^{\Phi})$ and $G(\mathcal{L}^{+})$

In $\mathcal{L}^{\Phi}$, extending the object-language $\mathcal{L}$ with s-quantifiers but no operators, s-variables occur only in sentential positions. The only atoms are $\mathcal{L}$-atoms $\mathbf{A}$ (and possibly $\mathbf{C}$. When $\mathbf{A} = \emptyset$, the language $\emptyset^{\Phi}$ is that of quantified boolean sentences, QBS.) Given a domain M and $\rho \in \mathbf{2}^{\mathbf{A}_M}$, all $\mathcal{L}^{\Phi}$ sentences obtain unique values by a unique extension to a kernel $\hat{\rho}$ of the corresponding graph $G_M(\mathcal{L}^{\Phi})$. (Numbers in parentheses refer to the same statements with proofs in the Appendix.)

Theorem 3.16 (7.6) *For each $G_M(\mathcal{L}^{\Phi})$ and $\rho \in \mathbf{2}^{\mathbf{A}_M}$, there is a unique $\hat{\rho} \in sol(G_M(\mathcal{L}^{\Phi}))$ with $\hat{\rho}|_{\mathbf{A}_M} = \rho$.*

The proof relies on the next lemma, stating that a solution of $G^{-}_M(S)$ – denoting, for $S \in \mathbf{S}^{\Phi}_M \setminus \mathbf{S}_M$, vertices of $G_M(S)$ without those in its DNF-feet – depends on the valuation of $\mathbf{A}_M$, but not of external leaves $ext(G_M(S))$, p.14, as the second part of the lemma states. In a way, DNF-foot determines a boolean function, and the value of S depends on this function (and the valuation of $\mathbf{A}_M$) rather than on the values of external leaves, which span all possibilities. Valuation of $ext(G_M(S))$ affects of course the values in the DNF-feet, where they occur. For either sentence A from Example 3.10, the lemma implies that the value of the source vertex with its marked cycles, $G^{-}(A)$, is independent from the values of the external leaves $S_1, S_2, \ldots$

Lemma 3.17 (7.7) *For every $G_M(\mathcal{L}^{\Phi})$ and $A \in \mathbf{S}^{\Phi}_M$, each valuation ρ of atoms $\mathbf{A}_M$ and external leaves of $G_M(A)$, $\rho \in \mathbf{2}^{\mathbf{S}_M \cup ext(G_M(A))}$, has a unique extension to $\rho_A \in sol(G_M(A))$. Valuation of atoms, $\rho|_{\mathbf{A}_M}$, determines the restriction of ρ_A to $G^{-}_M(A)$, in particular, the value $\rho_A(A)$.*

Valuation of sentences $\mathbf{S}^{\Phi}_M \setminus \mathbf{S}_M$ does not have any standard definition, which is merely suggested by (3.1). By Theorem 3.16, such a valuation $\hat{\rho}$ is determined by $\rho \in \mathbf{2}^{\mathbf{A}_M}$, just as is the valuation of $\mathbf{S}_M$. Existence and uniqueness of $\hat{\rho}$ ensure well-definedness of (3.1), given by the following.

Definition 3.18 An $\mathcal{L}^{\Phi}_M$-sentence A is true in an $\mathcal{L}$ structure (M, ρ), i.e., an $\mathcal{L}$-domain M with $\rho \in \mathbf{2}^{\mathbf{A}_M}$, $(M, \rho) \models A$, iff $\hat{\rho}(A) = \mathbf{1}$ for the unique solution $\hat{\rho} \in sol(G_M(\mathcal{L}^{\Phi}))$ with $\hat{\rho}|_{\mathbf{S}_M} = \rho$.

An $\mathcal{L}$ structure (M, ρ) can be thus seen as an $\mathcal{L}^{\Phi}$ structure interpreting all $\mathcal{L}^{\Phi}$ sentences. Any $\Gamma \subseteq \mathbf{S}^{\Phi}$ determines a well-defined class $Mod(\Gamma) = \{(M, \rho) \mid \forall A \in \Gamma : (M, \rho) \models A\} = \bigcap_{A \in \Gamma} Mod(A)$ of $\mathcal{L}$ structures modelling Γ. The bijection (3.15) between FOL structures and graphs with kernels, mapping (M, ρ) to (G_M, K_{ρ}), extends to FOL^{Φ} by mapping (M, ρ) to $(G_M, \hat{\rho})$.

The hardly unexpected but significant Theorem 3.16 implies that a classical language $\mathcal{L}$ remains free from paradoxes, under every valuation of atoms, when extended with quantification over all sentences to $\mathcal{L}^{\Phi}$. In fact, by the following theorem, $\mathcal{L}^{\Phi}$ has the same expressive power as $\mathcal{L}$.

Theorem 3.19 **(7.11)** *For every $\Gamma \subseteq \mathcal{L}^{\Phi}$ there is a $\Gamma^{-} \subseteq \mathcal{L}$ with $Mod(\Gamma) = Mod(\Gamma^{-})$.*

In particular, quantification over all sentences in FOL$^{\Phi}$, extending FOL apparently as far as possibility of self-reference, reduces to propositional quantification.

Solvability of $G(\mathcal{L}^{+})$

Operators applied to sentences provide only fresh atoms, so one might think that everything works unchanged. It does, if only such operators are introduced without s-quantifiers. The language graph which is then, as for the object-language, a collection of trees with new s-atoms, is uniquely solvable for every valuation of atoms. However, the combination of operators with s-quantifiers changes things dramatically. For instance, blind ascriptions of truth or infinite conjunctions, namely claims of the form "All Ks are true", for an operator K, become expressible as $\forall\phi(K\phi \rightarrow \phi)$.[8] Technically, a more significant novelty is the dependence of the operators on their argument *sentences*, not only the boolean values of these sentences, and the possibility of violating semantic equivalence of arguments. Consequently, only even cycles may be broken without breaking the corresponding odd ones, leading to paradoxes. Unlike valuations of $\mathcal{L}$-atoms in a domain M, determining solutions of $G_M(\mathcal{L}^{\Phi})$, some valuations of s-atoms can make language graph $G_M(\mathcal{L}^{+})$ unsolvable, as illustrated by Example 3.12. $\mathcal{L}^{+}$ can thus express some paradoxes, but none are implied. They appear, as in the example, only due to unfortunate valuations of s-atoms. An expansion $\mathcal{L}^{+}$ of a classical language $\mathcal{L}$ remains consistent in the sense of the following theorem.

Theorem 3.20 *Every language graph $G_M(\mathcal{L}^{+})$ has a kernel.*

A proof can just repeat the proof of Theorem 3.16 treating all operators as, e.g., constantly true. Solvability of graphs $G(\mathcal{L}^{+})$ with a more flexible interpretation of operators can be obtained from Theorem 5.6. It states the preservation of solvability by *definitional extension*, that is, introduction of a fresh operator, say P, to a language $\mathcal{L}^{+}$, by a sentence of the form $\forall\phi(P(\phi) \leftrightarrow \forall\!\!\!\exists\psi F(\phi, \psi))$, with $\forall\!\!\!\exists \in \{\forall, \exists\}$ and an $\mathcal{L}^{+}$-formula (not containing P) $\forall\!\!\!\exists\psi F(\phi, \psi)$ with the free variables ϕ among those of $P(\phi)$. Starting now from, e.g., $\mathcal{L}^{\Phi}$, with no operators but with graphs solvable by Theorem 3.16, yields the solvability of graphs for its expansion with any, also infinite, number of operators.

4 Reasoning, (non-)explosiveness and (cut)

LSO provides a sound and complete reasoning for the semikernel semantics (3.11).

Theorem 4.1 **(8.1)**, **(8.3)** *$\Gamma \vdash \Delta \Leftrightarrow \Gamma \models \Delta$, for Γ, Δ over a countable $\mathcal{L}^{+} \in FOL^{+}$.*

[8] Their role for truth-theory has been discussed at least since Quine's [33]. When syntax is arithmetized, they become problematic due to the complications with controlling their interaction with the restrictions on Convention T, e.g., [18, 31]. A paradox in LSO, in turn, requires a sentence or s-variable to occur in both a sentential and a nominal position, as exemplified by such blind ascriptions.

Given the classical character of LSO and the non-explosiveness of the semikernel semantics, this can seem surprising, so we comment on it a bit closer.

The derivability $\Gamma \vdash A \wedge \neg A$ shows the impossibility of combining Γ with any coherent valuation of A. By soundness, $\Gamma \models A \wedge \neg A$, so every semikernel L satisfying Γ satisfies $A \wedge \neg A$. But L can do it only vacuously, so if Γ has any models, they are only semikernels not covering $A \wedge \neg A$.

Such Γs do not exist at the object-level, where entailing a contradiction amounts to being one. Paradoxes occur only at the metalevel, with the help of operators and s-quantifiers, and the contradictions they imply involve these operators as well.

For an illustration, we go back to Example 3.12 with Karen claiming to always lie. We can then prove that she does not, $K(Kl) \vdash \neg Kl$, but no paradox follows yet. If this is everything she says, then we can also prove that she is always lying, and $K!Kl \vdash Kl \wedge \neg Kl$ witnesses to a paradox, where $K!Kl$ abbreviates $\{K(Kl), \forall\phi(K\phi \to \phi \doteq Kl)\}$. (Section 5.1 contains more details.)

Provability $K!Kl \vdash Kl \wedge \neg Kl$ does not imply the nonexistence of a semikernel containing $K!Kl$, as $K!Kl \vdash \emptyset$ would do, but the nonexistence of such a semikernel covering also $Kl \wedge \neg Kl$. Most other contradictions and sentences are not derivable from $K!Kl$. As we saw under Example 3.12, $K!Kl \not\models J(A)$, which follows from the infinite branch of the attempted derivation ($X = J(A)$):

$$(\forall^+_L)\,[Kl/\phi]\ \dfrac{\dfrac{(\forall^+_L)\,[X/\phi]\ \dfrac{\dfrac{(\forall^+_L)\,[S_1/\phi]\ \dfrac{\dfrac{\begin{array}{c}\vdots\\ K!Kl \vdash X, K(X), K(S_1)\end{array} \qquad \dfrac{S_1 \doteq Kl, K!Kl \vdash X, K(X), S_1 \doteq Kl}{S_1 \doteq Kl, K!Kl \vdash X, K(X)}\ (\text{neq})\ S_1 \not\doteq Kl}{K(S_1) \to S_1 \doteq Kl, K!Kl \vdash X, K(X)}}{K!Kl \vdash X, K(X)} \qquad (\text{neq})\ X \not\doteq Kl\ \dfrac{X \doteq Kl, K!Kl \vdash X, X \doteq Kl}{X \doteq Kl, K!Kl \vdash X}}{K(X) \to X \doteq Kl, K!Kl \vdash X}}{Kl \doteq Kl, K!Kl \vdash X} \qquad K!Kl \vdash X, K(Kl)}{K(Kl) \to Kl \doteq Kl, K!Kl \vdash X}}{K!Kl \vdash X}$$

The branch keeps instantiating $\forall\phi(K\phi \to \phi \doteq Kl)$ by all sentences. In the limit, $K(S_i)$ for each $S_i \not\doteq Kl$ appears to the right of . Since no disjunction of X with $K(X)$ or other $K(S_i)$ is provable, the following countermodel results, reflecting $K!Kl \nvdash J(A)$ (extending L from Example 3.12):

$$Z = \{\overline{J(A)}, K(Kl), \forall\phi(K\phi \to \phi \doteq Kl)\} \cup \{\overline{K(S_i)} \mid S_i \not\doteq Kl\}.^{9} \tag{4.2}$$

In the same way, $K!Kl \nvdash \emptyset$. Its derivation would copy the attempted $K!Kl \vdash X$ above, removing all Xs and yielding the infinite branch with the countermodel $Z_1 = Z \setminus \{\overline{J(A)}\}$. The possible situation where Karen claims to be lying can be extended to one where Karen claims nothing else, Z_1, and then to Z where John is not saying A, so $K!Kl \nvdash J(A)$.

Now, a contradiction entails every sentence S, e.g., $Kl \wedge \neg Kl \vdash S$, reflecting the fact that it does not belong to any semikernel. Since a paradox entails some contradiction, like $K!Kl \vdash \neg Kl \wedge Kl$, using (cut) $\frac{\Gamma \vdash \Delta, A \quad A, \Gamma \vdash \Delta}{\Gamma \vdash \Delta}$ would yield $K!Kl \vdash S$. But semikernel (4.2) gives a countermodel $Z \not\models K!Kl \Rightarrow J(A)$, so (cut) is not sound.

[9] This shows only the crucial sentences in Z, which contains also others implied by these ones: $K(X) \to X \doteq Kl$ for every $X \in \mathbf{S}^+$, $\overline{Y \doteq Kl}$ for all $Y \not\doteq Kl$, and $Kl \doteq Kl$.

It is trivially admissible for the object-language, as long as only LK is used, but it changes the semantics for the whole LSO. The contradiction $Kl \wedge \neg Kl$, following from Karen's statement, is not 'discovered' in Z. A semikernel that is not a kernel represents a limited context which is coherent, that is, only locally consistent, without taking into account the whole language. Z provides a model for Karen saying only Kl, but inquiry into the truth of what she is saying, Kl or $\neg Kl$, expands this context to the point where the paradox – the impossibility of a valuation of Kl coherent with Z – is discovered.

To obtain a "fully classical", i.e., explosive logic, enabling derivation of everything from $K!Kl$ via the contradiction it entails, it suffices to add unrestricted (cut), yielding LSO^c, with $\vdash_c$ denoting provability. This brings us to the kernel semantics (3.13), for which LSO^c is sound and complete.

Theorem 4.3 **(8.4)** *For Γ, Δ over a countable $\mathcal{L}^+ \in FOL^+$: $\Gamma \models_c \Delta$ iff $\Gamma \vdash_c \Delta$.*

Making thus contradictions explode seems the only contribution of (cut) to LSO. Paradoxes discovered using (cut) can also be diagnosed without it, since if LSO^c derives some contradiction from a theory, so does LSO. By the following theorem, LSO derives then a contradiction of the specific form $\bot_Q = \bigvee_{S \in Q}(S \wedge \neg S)$, for a finite set of sentences Q, denoted by $Q \Subset \mathbf{S}^+$.

Theorem 4.4 **(7.20)** *For Γ, Δ over a countable $\mathcal{L}^+ \in FOL^+$,*
$(\forall Q \Subset \mathbf{S}^+ : \Gamma \nvdash \bot_Q) \Rightarrow (\forall \bot \in \mathbb{C} : \Gamma \nvdash_c \bot)$.

4.1 Non-transitivity

The unsoundness of (cut) for the semikernel semantics (3.11) arises from the subtle element of (ir)relevance: vacuous satisfaction of a formula by a semikernel not covering it. If $\{X\} \in \mathbb{C}$ then $\Gamma \vdash X$ either if $\Gamma \vdash \emptyset$, i.e., $\Gamma \in \mathbb{C}$, or if $\Gamma \notin \mathbb{C}$ but every semikernel containing Γ satisfies $\Gamma \Rightarrow X$ vacuously, by not covering X. In the latter case, semikernels containing Γ may cover other sentences, enabling $\Gamma \not\models \Delta$. The unsoundness of (cut) is limited to such cases, when countermodels to the conclusion satisfy premise(s) only vacuously.

Fact 4.5 *(a1) $\Gamma \models \Delta, A$ and (a2) $\Gamma, A \models \Delta$ and (c) $\Gamma \not\models \Delta$ iff there is a semikernel $L \not\models \Gamma \Rightarrow \Delta$ but none such can be extended to a semikernel $L' \supseteq L$ with $A \in \mathbf{E}^-[L']$.*

Proof $\Rightarrow$) (c) $\Gamma \not\models \Delta$ means that there is a semikernel $L \not\models \Gamma \Rightarrow \Delta$, i.e., $\Gamma \subseteq L$ and $\Delta \subseteq \mathbf{E}^-(L)$. Let L be any such. If L can be extended to L' covering A, then $\Gamma \subseteq L'$ and $\Delta \subseteq \mathbf{E}^-(L')$, while either (i1) $A \in \mathbf{E}^-(L')$, contradicting (a1) since $L' \not\models \Gamma \Rightarrow \Delta, A$, or (i2) $A \in L'$, contradicting (a2) as $L' \not\models \Gamma, A \Rightarrow \Delta$. Thus, L cannot be extended to cover A.

$\Leftarrow$) A semikernel $L \not\models \Gamma \Rightarrow \Delta$ establishes (c). Let L' be any semikernel covering Γ, Δ and A. By assumption $L' \models \Gamma \Rightarrow \Delta$ (otherwise it could not cover A), establishing both (a1) and (a2). □

The impossibility to extend a semikernel $L \not\models \Gamma \Rightarrow \Delta$ to a semikernel covering A means that L harbours a paradox: excluding both A and $\neg A$, it cannot be extended to a kernel, precluding a valuation of the entire language. In Fact 4.5, this applies to all semikernels falsifying the conclusion $\Gamma \Rightarrow \Delta$ of (cut), meaning that its negation is paradoxical but not contradictory. The conclusion itself need not be paradoxical, as semikernels satisfying it may admit extensions to kernels.

This inadmissibility of (cut) at the metalevel of LSO is reminiscent of the non-transitive consequence studied, e.g., in [10, 39, 40]. These non-transitive systems are fundamentally different from LSO, as they deal with a transparent truth predicate in the style of AST. However, the role of (cut), or rather of its absence, seems sufficiently similar in both to warrant a closer comment.

Relation $\models^{st}_{+}$ from [39], appearing most relevant for the comparison, concerns the so-called external or meta-inferential level: $A, \neg A \Rightarrow B$ is valid (vacuously), but (cut) in the form $\frac{\Gamma \vdash A\ ;\ \Gamma \vdash \neg A}{\Gamma \vdash B}$ is not admissible. More specifically, relation $\models^{st}_{+}$ holds for $\mathbf{1} \models^{st}_{+} \frac{\mathbf{1}}{\mathbf{2}}$ and $\frac{\mathbf{1}}{\mathbf{2}} \models^{st}_{+} \mathbf{0}$, with $\frac{\mathbf{1}}{\mathbf{2}}$ representing paradox. To prevent $\mathbf{1} \models^{st}_{+} \mathbf{0}$, transitivity is blocked precisely when the cut formula is a paradox (relatively to the context). Fact 3.17 in [39] excludes (cut) between $\Gamma \Rightarrow \Delta, A$ and $A, \Gamma \Rightarrow \Delta$ when A evaluates to $\frac{\mathbf{1}}{\mathbf{2}}$ in the models falsifying $\Gamma \Rightarrow \Delta$. This is reflected, if not exactly repeated, by our Fact 4.5, barring transitivity to the conclusions falsifiable only by the models that display their latent paradox under the extensions covering A. Despite conveying very similar informal messages, the two facts differ due to the different models of paradox.

This is a fundamental difference between the 3-valued semantics in [39, 40] and our 2-valued semantics. As noted earlier, even if a semikernel does not cover the whole graph, it still assigns a boolean value to all sentences. Semikernels do handle inconsistency by leaving it out, but in each relevant semikernel L, which covers the actual sequent, every sentence of the sequent is either true (in L) or false (in $\mathbf{E}^{-}(L)$). Kernels are just special cases that leave nothing out. Since they exist for language graphs, all sentences can obtain truth values. A paradox is not a third value but a failure to assign any of the two. Atomic claims are either true or false, while the unfortunate third value of paradox – or rather the unfortunate fact of not being amenable to evaluation – arises only from confused (compounds of) sentences. Semikernels allow Karen to say that she is always lying and even to say only that. That a paradox results is as unfortunate a consequence as is the contradiction that she is both telling the truth and not, which falls out of reasonable discourse, out of any model. It could be viewed as a third value, but it seems more accurate to see it as the impossibility of extending a bivalent valuation, given by a semikernel containing Karen's statements, to one determining also the truth of their consequences.

These differences in the semantics come clearly forth in the reasoning systems. The 3-sided sequent systems in [39], reflecting the 3-valued semantics, come in two variants, disjunctive and conjunctive, which can be related to the 2-sided ones in the expected ways but extend considerably their expressiveness. We do not dispute their merits but limit the comparison to the 2-sided system for $\models^{st}_{+}$. First, LSO restricted to the mere truth predicate is a trivial extension of LK admitting, besides unrestricted (cut), insertion of $\mathcal{T}$ around any sentence. ST reasoning from [40] almost coincides with the so-restricted LSO except that, using $\models^{st}_{+}$ with 3-valued models, ST does not admit (cut). Restrictions on (cut) are very similar in ST and LSO, guarding against applications over paradoxes. However, while ST needs such restrictions for reasoning with a truth predicate, responsible for paradoxes in AST, the restrictions on (cut) in LSO are used for reasoning with arbitrary operators. Limited to the mere object-language, LSO is just LK admitting (cut).

Perhaps the most significant difference emerges from admitting unrestricted (cut). Since we have both $\mathbf{1} \models^{st}_{+} \frac{\mathbf{1}}{\mathbf{2}}$ and $\frac{\mathbf{1}}{\mathbf{2}} \models^{st}_{+} \mathbf{0}$, (cut) trivializes $\models^{st}_{+}$ yielding $\mathbf{1} \models^{st}_{+} \mathbf{0}$. In LSO, (cut) does not trivialize the logic but only inconsistent theories, turning the non-explosive logic of semikernels into the explosive one of kernels. Paradoxical statements by Karen imply specific contradictions, but they do not affect John or the object-level. Unrestricted (cut), by making paradoxes entail everything, breaks this bond of relevance and reverts LSO to the classically explosive logic.

Given such diverging technical contexts, the restricted transitivity via paradoxes is indeed a striking similarity between the two approaches. Still, the differences in the scope of (cut)'s applicability and the dramatically different consequences of lifting its restrictions suggest that the significance of this similarity might be smaller than it initially appears.

5 Some Examples

This section gives some examples of paradox analysis in LSO. Even when approaches based on AST provide similar conclusions, the analyses differ. Cases with an intensional or modal character may be problematic in AST due to paradoxes arising easier with modalities-as-predicates than with modal operators. The finer distinctions of such an intensional kind in Section 5.2 are hardly expressible in AST. In general, the non-explosiveness of LSO leads to the derivability of only relevant consequences of paradoxes, which suggest specific changes for avoiding them.

5.1 Stating a Paradox is Possible, even if not Evaluating its Truth-Value

Karen saying that she always lies, $K(Kl)$ with $Kl = \forall\phi(K\phi \rightarrow \neg\phi)$, sometimes tells the truth:

$$\dfrac{\dfrac{\dfrac{\dfrac{\forall\phi(K\phi \rightarrow \neg\phi), K(Kl) \vdash Kl}{\neg Kl, \forall\phi(K\phi \rightarrow \neg\phi), K(Kl) \vdash} \qquad \forall\phi(K\phi \rightarrow \neg\phi), K(Kl) \vdash K(Kl)}{K(Kl) \rightarrow \neg Kl, \forall\phi(K\phi \rightarrow \neg\phi), K(Kl) \vdash}}{\forall\phi(K\phi \rightarrow \neg\phi), K(Kl) \vdash}\;(\forall^{+}_{L})\; Kl[Kl/\phi]}{K(Kl) \vdash \neg\forall\phi(K\phi \rightarrow \neg\phi)} \tag{5.1}$$

As noted by Prior in [32], Karen must then also sometimes lie:

$$\dfrac{\dfrac{\dfrac{Kl, K(Kl) \vdash K(Kl) \qquad \begin{matrix}(5.1)\\ \vdots \\ Kl, K(Kl) \vdash\end{matrix}}{\forall\phi(K\phi \rightarrow \phi), K(Kl) \rightarrow Kl, K(Kl) \vdash}}{\forall\phi(K\phi \rightarrow \phi), K(Kl) \vdash}\;(\forall^{+}_{L})\; Kl[Kl/\phi]}{K(Kl) \vdash \neg\forall\phi(K\phi \rightarrow \phi)}$$

The resulting Prior's theorem

$$K(\forall\phi(K\phi \rightarrow \neg\phi)) \rightarrow (\exists\phi(K\phi \wedge \phi) \wedge \exists\phi(K\phi \wedge \neg\phi)), \tag{5.2}$$

is so far no paradox, as Karen can also say other things. If she does not, what follows is not that she is saying two things, one true and one false, but a contradiction, signalling a paradox seen in Example 3.12. On the one hand, (5.1) gives $K(Kl) \vdash \neg Kl$. To obtain Kl requires capturing that she says nothing else, which amounts to the infinite set of negated atoms $\overline{L} = \{\overline{K(S)} \mid S \not\doteq Kl\}$.

$$
\frac{
 \dfrac{\begin{array}{c}(5.1)\\ \vdots \\ \overline{L}, K(Kl) \vdash \neg Kl\end{array}}{\overline{L}, K(Kl) \vdash K(Kl) \to \neg Kl}
 \qquad
 \dfrac{\dfrac{\overline{L}\setminus\{K(S)\}, K(Kl), K(S), S \vdash K(S)}{\overline{L}, K(Kl), K(S), S \vdash}\ \neg K(S)\in\overline{L}}{\overline{L}, K(Kl)\vdash K(S)\to\neg S}
 \qquad \text{– for each } S \neq Kl
}{\overline{L}, K(Kl)\vdash \forall\phi(K\phi\to\neg\phi)}\ (\forall^{+}_{R})
$$

S-equality allows a finite expression and proof of this fact, using $(\forall^{\doteq}_{R})$ instead of $(\forall^{+}_{R})$:

$$
\frac{
 \dfrac{\begin{array}{c}(5.1)\\ \vdots \\ K!Kl, K(Kl), Kl \vdash\end{array}}{K!Kl \vdash K(Kl) \to \neg Kl}
 \qquad
 \dfrac{\dfrac{\dfrac{\ldots, K\sigma \vdash \ldots, K\sigma \qquad \ldots, \sigma \doteq Kl, \ldots \vdash \sigma \doteq Kl}{\ldots, K\sigma \to \sigma \doteq Kl, K!Kl, K\sigma, \sigma \vdash \sigma \doteq Kl}}{K!Kl, K\sigma, \sigma \vdash \sigma \doteq Kl}\ (\forall^{\sigma}_{L})\ K!Kl[\sigma/\phi]}{K!Kl \vdash K\sigma \to \neg\sigma, \sigma \doteq Kl}
}{K!Kl \vdash \forall\phi(K\phi \to \neg\phi)}\ (\forall^{\doteq}_{R})
$$

The semikernel $L = \overline{L} \cup \{K(Kl)\}$ models Karen saying only that she is lying. The paradox amounts to the provability of contradiction, $K!Kl \vdash Kl \wedge \neg Kl$, and the impossibility of extending semikernel L to any covering Kl. Karen's paradoxical statement makes evaluation of its truth-value impossible. Semikernel L does not, however, validate any other statements. Most other facts – and contradictions – remain unprovable. Snow does not become non-white, while for any sentence S distinct from Kl, we have $K!Kl \vdash \neg K(S)$ but $K!Kl \nvdash K(S)$. Karen saying only Kl does not become Karen saying everything, as she does in AST (or LSO^{c}).

The Diagonalization Lemma

This beautiful and powerful result ensures that, for every formula with one free variable, $F(x)$, there is a sentence S such that $S \leftrightarrow F(\ulcorner S \urcorner)$ is provable (if the theory contains enough arithmetic). It gives then the equivalence $\neg\mathcal{T}(\ulcorner L \urcorner) \leftrightarrow L$, with Convention T yielding $L \leftrightarrow \mathcal{T}(\ulcorner L \urcorner)$. The situation in LSO is different, although variants of the lemma hold here, too. For instance:

Fact 5.3 *For every formula $F\phi$ with one free variable, s-variable ϕ, there is a sentence S, with a fresh operator K, such that $K_1 \models S \to F(S)$, for valuations K_1 making $K(S) = \mathbf{1}$.*

Taking $S = \forall\phi(K\phi \to F\phi)$ yields $K(S) \models S \to F(S)$, so that S says (or implies) "S is F". Just as the liar L is the paradigmatic paradoxical instance of the diagonalization lemma in AST, Karen is such an instance of this fact in LSO, with $F\phi = \neg\phi$ and Kl for S. Proof (5.1) gives then $K(S) \models F(S)$, hence $K(S) \models S \to F(S)$, and paradox results if K is true only about Kl. With the chosen S, no paradox arises unless ϕ has a sentential occurrence in $F\phi$, as S has only nominal ϕ under K, while a paradox requires both. Paradoxes arising potentially from self-reference in Fact 5.3 exemplify again dependence on the interpretation of operators.

5.2 Statements do not Affect Facts, but Tautologies can make a Difference

Just like the fact of making a statement should not be confused with the statement itself, thinking a thought should be distinguished from the thought's content, let alone its truth. When thoughts are confused with their truth-values, thinking may seem to force something outside one's thoughts. The following is quoted with inessential modifications after [1]:

"If K is thinking only 'Everything I am thinking now is false iff X is false', then X is true."

We take X as an arbitrary atom, while Asher, reading it as "Everything Tarski is thinking is false", notes: *"By reasoning that is valid in the simple theory of types, we conclude that Tarski was not able to think that snow is white, a bizarre and unwanted consequence of a logic for belief"*.

The relevant part of the graph is shown below. K is thinking $A = Kl \leftrightarrow \neg X$, where $Kl = \forall\phi(K\phi \to \neg\phi)$ is as in Example 3.12; each ${}_S\wedge_{K(S)}$ has edges to $\overline{K(S)}$ and $\overline{S}$, as shown for ${}_A\wedge_{K(A)}$.

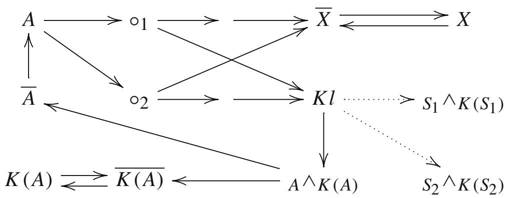

Each combination $\mathbf{2}^{\{K(A),X\}}$, of literals over $K(A)$ and X, forms a semikernel so $K(A) \not\models X$, since e.g., $\{K(A), \overline{X}\} \not\models X$. LSO proves something about the relations between the truth-values of A and of X, e.g., $A, K(A) \vdash X$, but this isn't as exciting as K's thought limiting Tarski's.

If A is the only thought of K, $K!A = \{K(A), \forall\psi(K\psi \to \psi \doteq A)\}$, then $K(S) = \mathbf{0} = {}_S\wedge_{K(S)}$ for each $S \not\doteq A$, leaving ${}_A\wedge_{K(A)}$ undetermined. We still have $K!A \nvdash X$, as $\{K(A), \overline{X}\}$ can be extended to a semikernel containing $K!A$. Something new follows now about the relation between truth of A and X. Of the two cycles via $A - Kl - {}_A\wedge_{K(A)} - \overline{A}$, the one via $\circ_2$ is even and the one via $\circ_1$, odd. Having ${}_S\wedge_{K(S)} = \mathbf{0}$, for all $S \not\doteq A$, the only way to break the latter is by $\overline{X} = \mathbf{0}$, i.e., $X = \mathbf{1}$. Now, this follows also assuming $\neg A$, i.e., $\neg A, K!(A) \vdash X$.

The two proofs give $A \vee \neg A, K!A \vdash X$. Having also $\vdash A \vee \neg A$, (cut) would yield $K!A \vdash X$. But semikernel $\{K!A, \overline{X}\}$ shows $K!A \not\models X$. To obtain X requires an additional assumption. The two proofs show that A's value does not matter. What matters is that it has one, $A \vee \neg A = \mathbf{1}$. This is guaranteed by $\vdash A \vee \neg A$ only in situations covering $A \vee \neg A$. The needed assumption is that in situations containing $K!A$ and covering X, also A can be evaluated, so that $A \vee \neg A = \mathbf{1}$. No semikernel extension of $\{K!A, \overline{X}\}$ allows this, while $K!A \not\models X$ makes this (cut) unsound.

K's thought A, "This thought is false iff X is", is true iff the equivalence it states is. X is not entailed by K thinking A but by this equivalence, $A \leftrightarrow (\neg A \leftrightarrow \neg X) \models X$, and vacuously so when $X = \mathbf{0}$. To this propositional essence, our semantics adds the possibility of non-satisfaction of X when K thinks only A, so that $K!A \not\models X$. This thought of K, being a contingent paradox, yields inconsistency when X is false, $K!A, \neg X \vdash A \wedge \neg A$. Hence, X must be true if K thinking only A is to form a non-paradoxical situation with A's truth or falsity, $A \vee \neg A,\ K!A \models X$, not merely because K is thinking A and nothing else.

Although tautological, $A \vee \neg A$ adds thus the nonempty assumption that A is not part of a paradox. (Most other tautologies are here irrelevant and adding them will not entail X.) In classical logic, this tacit assumption of non-paradoxicality coincides with consistency. Distinguishing the two and admitting the plausible event $\{K!A, \neg X\}$, LSO notices also that A becomes then paradoxical.

5.3 Metalanguage: Paradoxes and Indeterminate Statements

The non-explosiveness of paradoxes is one feature distinguishing the metalanguage from the object-language. Another is exemplified by sentences that remain indeterminate under some valuations of all atoms. John saying only that he always tells the truth, the framed $J(Jt)$ with $Jt = \forall\phi(J\phi \to \phi)$ on the drawing below, is the truth-teller. Each $_{\overline{X}}\wedge_{J(X)}$, for $X \not\equiv Jt$, is false due to John not saying X, while $\overline{J(Jt)} = \mathbf{0}$ leaves the 4-cycle $Jt -_{\overline{Jt}} \wedge_{J(Jt)} - \overline{\overline{Jt}}$ with one solution $Jt = \mathbf{1}$ and the other $Jt = \mathbf{0}$.

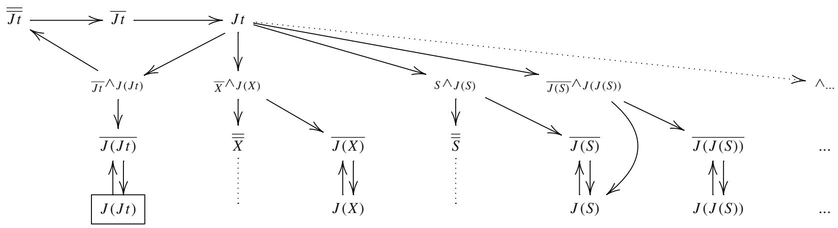

LSO admits semikernel models with sentences like Jt left without any value. This does not happen at the object-level $\mathcal{L}$, nor $\mathcal{L}^{\Phi}$, where each valuation of atoms determines the values of all sentences. Considering this metalevel phenomenon a flaw seems due to the identification of the metalanguage with (a subset of) the object-language by coding the former in the latter. In LSO such indeterminateness is simply another, besides paradoxes, feature distinguishing the two. Moreover, the difference between the innocent self-reference of the truth-teller and the vicious circularity of paradoxes, reflecting the indeterminacy of the former versus the impossibility of valuating the latter, is captured in language graphs by even versus (unresolvable) odd cycles.

The unproblematic status of the truth-teller amounts to the informal observation that it says nothing. Making no real claim, its truth or falsity makes no difference and can be chosen arbitrarily. A difference appears if it also says some false X, because then $_{\overline{X}}\wedge_{J(X)} = \mathbf{1}$ makes $Jt = \mathbf{0}$. But even if everything else John is saying is true, the 4-cycle $Jt -_{\overline{Jt}} \wedge_{J(Jt)} - \overline{\overline{Jt}}$ can still be solved with $Jt = \mathbf{0}$ and $_{\overline{Jt}}\wedge_{J(Jt)} = \mathbf{1}$. The mere claim of telling only the truth implies the consistency of this claim's falsehood, regardless of what other true or false things one may be saying.[10]

5.4 Modal Logics

The authors of [8] consider the following situation, with a peculiar pattern of reference.

(1) Ann believes that Bob assumes that (2).

(2) Ann believes that Bob's assumption is wrong.

The question whether (*) *Ann believes Bob's assumption to be wrong*, is answered by the following informal reasoning. (We insert (a), (b), and (c) to mark the assumptions used later.)

"If so, then in Ann's view, Bob's assumption, namely 'Ann believes that Bob's assumption is wrong', is right. But then Ann does not believe that Bob's assumption is wrong (a), which contradicts our starting supposition. This leaves the other possibility (b), that Ann does not believe that Bob's assumption is wrong. If this is so, then in Ann's view, Bob's assumption, namely 'Ann believes that Bob's assumption is wrong', is wrong (c). But then Ann does believe that Bob's assumption is wrong, so we again get a contradiction."

Whether statements form a paradox or not depends often on the representation, and the authors build an impressive machinery to ensure that these do. The following is only one possible way of capturing the situation in LSO. We do not worry about the distinction between 'believes' and 'assumes', central in [8], but denote Ann's thoughts by A, Bob's by B and let σ be what Ann believes to be Bob's assumption (2). This yields the following representation:

(1) $AB\sigma$ (2) σ, where $\sigma \leftrightarrow A(B\sigma \wedge \neg\sigma)$ and $A(B\sigma \wedge \neg\sigma) \overset{(d)}{\leftrightarrow} AB\sigma \wedge A\neg\sigma$.
As in normal modal logic, A distributes over conjunction and implication, (d). Equivalence (2), available to the agents, has any number of As (or Bs), in particular, $A(\sigma \leftrightarrow A(B\sigma \wedge \neg\sigma))$. Adding the assumptions from the informal argument yields, for all ϕ, ψ:

(a) $AA\phi \leftrightarrow A\phi$ (c) $\neg A\phi \rightarrow A\neg\phi$, for relevant ϕ
(b) $A\phi \vee \neg A\phi$ (d) $A(\phi \wedge \psi) \leftrightarrow (A\phi \wedge A\psi)$ and $A(\phi \rightarrow \psi) \rightarrow (A\phi \rightarrow A\psi)$

The question whether (*) *Ann believes Bob's assumption to be wrong*, asking apparently whether $A\sigma$ or $\neg A\sigma$ (or $A\neg\sigma$), since σ is what B assumes, asks equally whether $\neg\sigma$ or σ, as σ states exactly that Ann believes Bob's assumption to be wrong. Taking the former, LSO proves $AxA\sigma \wedge \neg A\sigma$ from $Ax = \{(1), (2), (a), (b), (c), (d)\}$, but we spare the reader the involved intricacies. In spite of the appearances, this paradox concerns only a single complex belief of Ann. No axioms about Bob are needed, and their irrelevance becomes apparent when we note that (1) simplifies (2) to a version of Knower's paradox, $\sigma \leftrightarrow A\neg\sigma$. Subtleties of the analyses in [8], intended

[10] Buridan's early proposal, that each statement claims its own truth in addition to whatever it may be saying, provides thus a 'solution' to the liar and similar paradoxes by making them false – but for the price of consistency of all statements being false. Earlier, Bradwardine maintained the falsity of paradoxes without this over-generalization, taking only some – self-negating – statements as also claiming their own truth, e.g., [35, 36]. These statements, only vaguely specified, can be now identified precisely by the presence of the (unresolvable) odd cycles.

for applications to multiagent games, result in a particular modal logic, while LSO is a general schema admitting various specializations and helping to unveil uniform patterns. Granting the ingenuity of the scenario, there seems to be nothing specifically intensional or modal about the paradox, at least, when represented as above. Reading $A\phi$ as 'A claims ϕ', the scenario becomes

(1) Ann claims that Bob claims that (2),
(2) Ann claims that Bob's claim is false,

yielding a paradox of semantic character, represented and analyzed in exactly the same way.

In general, modal paradoxes have a natural representation in LSO, with modalities axiomatized as appropriate operators. (The system LSO might need adjustments to handle such axiomatizations, especially, to cater for the necessitation rule.) Relating modal logic to LSO would take at least another paper, so we only comment on one more example, utilizing the difference between language graphs and Kripke frames, reflected here by that between sentences and propositions.

The formula $K\phi \wedge \forall\psi(K\psi \rightarrow \psi \doteq \phi)$, for Karen saying only ϕ, occurs as a subformula in the sentence $\forall\phi\Diamond\forall\psi(K\psi \leftrightarrow \psi \doteq \phi)$, considered by Kaplan in [21], if we write it equivalently as

(A) $\forall\phi\Diamond(K\phi \wedge \forall\psi(K\psi \rightarrow \psi \doteq \phi))$.

The important difference is that Kaplan uses = as the equality of propositions, instead of our syntactic equality $\doteq$. With s-quantifiers ranging over propositions, viewed as arbitrary subsets of possible worlds, a cardinality argument excludes an operator K satisfying (A), making it false rather than paradoxical. However, (A) says that for every ϕ it is possible for K to say ϕ and nothing else, which seems quite plausible in limited situations. According to [21], logic should not rule it out. Attempts to save (A) (for instance, by taking as propositions only some subsets, as in [24], or by restricting the principle of universal instantiation, e.g., [4]) leave the issue open.

The cardinality argument does not affect language graphs, where $\Diamond$ and K act on sentences and not any subsets. (A) is trivially satisfied by semikernel L containing $\forall\phi\Diamond(K!\phi)$ and atom $\Diamond(K!S_i)$, for each $S_i \in \mathbf{S}^+$. Ensuring modal content of $\Diamond$ and K, by closing L under appropriate modal axioms, does not change the situation and keeps (A) satisfiable, as the problem is due to the model of propositions and not modalities.

5.5 The Yablo Paradox

All our examples so far have involved circular paradoxes because they are most common and natural. The Yablo paradox, on the other hand, appears noncircular, unless one applies some esoteric notions of circularity. In LSO, it amounts simply to the graph cycles and Yablo graph $Y = (\omega, <)$ has none. The essential aspects of this paradox can be captured, e.g., by the following theory Y from [22]:

(a) a transitive binary relation R on a nonempty set of sentences,
(b) that has no endpoints, $\forall\alpha\exists\beta R(\alpha, \beta)$, and where
(B) operator P satisfies the formula: $\forall\alpha(P(\alpha) \leftrightarrow \forall\beta(R(\alpha, \beta) \rightarrow \neg P(\beta)))$.

A single sentence with a loop provides a model of R, and so does the ω-ordering, but no semikernel contains Y. The author observes that its "inconsistency [...] has nothing to do with truth, for it [...] arises irrespective of what P means: other than the Yablo scheme itself (B) and the auxiliary axioms (a), (b), no specific axioms for P are used in the deduction of the inconsistency.” Indeed, Y with variables ranging over objects rather than sentences is a contradiction and the inconsistency of Y has nothing to do with the truth predicate. By our definition of paradox, as the inclusion in a semikernel that cannot be extended to a kernel, Y is not a paradox but a contradiction.

A paradox can be obtained using a different formulation. For instance, let $P_1, P_2, \ldots$ be ω operators (persons), each P_i holding for (saying) exactly $\forall\phi(P_j\phi \to \neg\phi)$ for all $j > i$. Theory Γ, containing for each P_i s-atoms $\{P_i(\forall\phi(P_j\phi \to \neg\phi)) \mid j > i\}$ and negations of P_i applied to all other sentences, captures the Yablo-like situation, where everybody makes infinitely many claims, amounting to: everything everybody after me is saying is false.

This Γ, composed of s-literals, is a semikernel of any language graph. Schema (1) below, for subgraph $G(\forall\phi(P_1\phi \to \neg\phi))$, is equivalent to (2), with contracted double edges from $\circ_i$ to $\bullet_i$.

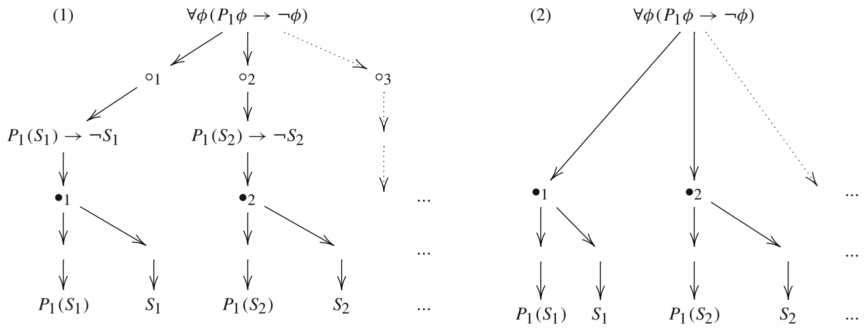

When $P_1(S_k) = \mathbf{0}$, the respective $\bullet_k = \mathbf{0}$. With the valuation of s-atoms given by Γ, this leaves only subgraph (2) with S_{n_i} identical to $\forall\phi(P_i\phi \to \neg\phi)$ for $i > 1$. The same happens for every P_i, leaving pattern (3) below, where dashed arrows mark the triple arrows which can be contracted to single ones, yielding Yablo graph $(\omega, <)$:

(3)

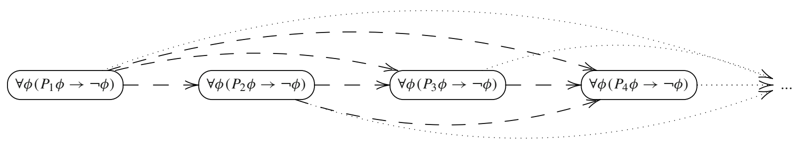

Theory Γ defines thus operators P_i forming a semikernel which cannot be extended to a kernel due to the unsolvability of subgraph (3).

The infinitary character of the argument is captured in LSO^c, formalizing the following proof:

(a) P_1 says something true if either what it says about P_2 is true, i.e., everything P_2 says is false, or else if P_2 says something true: $\Gamma \models (\exists\phi(P_2\phi \wedge \phi) \vee \forall\phi(P_2\phi \to \neg\phi)) \to \exists\phi(P_1\phi \wedge \phi)$

(b) The antecedent of (a) is a tautology, so P_1 says something true: $\Gamma \models \exists\phi(P_1\phi \wedge \phi)$.

(c) By the analogous argument, $\varGamma \models \exists\phi(P_i\phi \wedge \phi)$ for every $i \geq 1$, i.e., $\varGamma \models \neg\forall\phi(P_i\phi \to \neg\phi)$.
(d) Collecting all formulas from point (c) for $i \geq 2$, gives actually $\varGamma \models \forall\phi(P_1\phi \to \neg\phi)$.
(e) Contradiction of (b) and (d) shows paradoxicality of $\varGamma$.

We abbreviate $L_i = \forall\phi(P_i\phi \to \neg\phi)$, so that $\neg L_i \leftrightarrow \exists\phi(P_i\phi \wedge \phi)$, and let $L_i[\phi\backslash S]$ denote the instantiation of $\forall\phi$ in L_i by S. The following is a proof of (a):

$$
\dfrac{
\dfrac{
(\forall_R^+)\;\dfrac{
L_1[\phi\backslash S]\,(\forall_L^+)\;\dfrac{
\dfrac{
\dfrac{
\dfrac{\vdots\,(\mathrm{x})}{\varGamma, L_1, P_2S \vdash P_1S} \qquad \varGamma, L_1, P_2S, \neg S \vdash \neg S
}{\varGamma, L_1, P_2S, P_1S \to \neg S \vdash \neg S}\;(\to_L)
}{\varGamma, L_1, P_1S \to \neg S \vdash P_2S \to \neg S}
}{\varGamma, L_1 \vdash P_2S \to \neg S}
\qquad \ldots \vdash \ldots
}{\varGamma, L_1 \vdash L_2}
}{\varGamma, \neg L_2 \vdash \neg L_1}
\qquad
\dfrac{
\dfrac{
\dfrac{
\dfrac{\varGamma \ni P_1(L_2)}{\varGamma, L_2, L_1 \vdash P_1(L_2)} \qquad \dfrac{\varGamma, L_2, L_1 \vdash L_2}{\varGamma, L_2, L_1, \neg L_2 \vdash}
}{\varGamma, L_2, L_1, P_1(L_2) \to \neg L_2 \vdash}
}{\varGamma, L_2, L_1 \vdash}\;(\forall_L^+)\,L_1[\phi\backslash L_2]
}{\varGamma, L_2 \vdash \neg L_1}
}{\varGamma, \neg L_2 \vee L_2 \vdash \neg L_1}
$$

Branch (x), of the same kind for every S instantiating L_2 in the premises of $(\forall_R^+)$, has two cases. If $S = L_i$ for some i such that $P_2S \in \varGamma$, then also $P_1S \in \varGamma$, yielding an axiom. Otherwise, $\neg P_2S \in \varGamma$ yielding also an axiom.

(b) The extra assumption, $\neg L_2 \vee L_2$, of the resulting sequent is a tautology, hence $\varGamma \vdash \neg L_2 \vee L_2$. Applying (cut) to this and the result of (a) $\varGamma, \neg L_2 \vee L_2 \vdash \neg L_1$, yields $\varGamma \vdash_c \neg L_1$.

(c) In the same way, $\varGamma \vdash_c \neg L_i$ follows for each P_i, giving premises in the indicated branches of the following proof of point (d):

$$
(\forall_R^+)\;\dfrac{
(\to_R)\;\dfrac{\dfrac{\vdots}{\varGamma \vdash_c \neg L_2}}{\varGamma \vdash P_1(L_2) \to \neg L_2}
\qquad
\dfrac{\dfrac{\vdots}{\varGamma \vdash_c \neg L_3}}{\varGamma \vdash P_1(L_3) \to \neg L_3}
\qquad \cdots \quad \text{for each S} \in \mathbf{S}^+ \qquad
\dfrac{\dfrac{\varGamma \ni \neg P_1(S)}{\varGamma, P_1(S) \vdash \neg S}}{\varGamma \vdash P_1(S) \to \neg S}
}{\varGamma \vdash \forall\phi(P_1\phi \to \neg\phi)}
\tag{5.4}
$$

Dots $\cdots$ stand also for the branches instantiating $\forall\phi$ of the conclusion with sentences S other than $\forall\phi(P_i\phi \to \neg\phi)$, which terminate with axioms as shown to the right. Combined with (b), this yields (e) $\varGamma \vdash_c \forall\phi(P_1\phi \to \neg\phi) \wedge \neg\forall\phi(P_1\phi \to \neg\phi)$.

By Theorem 4.4, a proof of a contradiction from $\varGamma$ in LSO^c ensures also one without (cut). (Cut) was used only in the proof of (b)/(c), which is here proven without it:

$$
\dfrac{
\dfrac{
\dfrac{
\dfrac{\varGamma \ni P_1(L_2)}{\varGamma \vdash P_1(L_2)}
\qquad
\dfrac{
\dfrac{
\dfrac{
L_3[\phi\backslash L_1]\,(\forall_L^+)\;\dfrac{
\dfrac{
\dfrac{\varGamma \ni P_1(L_3)}{\varGamma, L_1, L_3 \vdash P_1(L_3)} \qquad \dfrac{\varGamma, L_1, L_3 \vdash L_3}{\varGamma, L_1, \neg L_3, L_3 \vdash}
}{\varGamma, L_1, P_1(L_3) \to \neg L_3, L_3 \vdash}
}{\varGamma, L_1, L_3 \vdash}
}{\varGamma, L_1 \vdash \neg L_3}
\qquad
\dfrac{\dfrac{\vdots}{\varGamma, L_1, L_i \vdash}}{\varGamma, L_1 \vdash \neg L_i}
\qquad \cdots \quad \forall i > 2
}{\varGamma, L_1 \vdash L_2}\;(\forall_R^+)
}{\varGamma, L_1, \neg L_2 \vdash}
}{\varGamma, L_1, P_1(L_2) \to \neg L_2 \vdash}
}{\varGamma, \forall\phi(P_1\phi \to \neg\phi) \vdash}\;(\forall_L^+)\,L_1[\phi\backslash L_2]
}{\varGamma \vdash \neg\forall\phi(P_1\phi \to \neg\phi)}
$$

In the same way $\neg L_i$ follows for every $i > 1$, and then L_1 by (5.4) with $\vdash_c$ replaced by $\vdash$.

5.6 Definitional Extensions and "Convention T"

As every kernel is a semikernel, the explosive kernel semantics is a special case of the non-explosive semantics of semikernels, and the two coincide on consistent theories. Theorem 4.1 can help checking whether we are in such a desirable situation, but some sufficient syntactic conditions would be desirable. A modest example can be a *definitional extension*. It extends a language $\mathcal{L}^+$ to $\mathcal{L}^P$ with a fresh n-ary operator symbol P, defined by a sentence of the form ($\exists\!\!\!\forall \in \{\forall, \exists\}$)

$$\forall\phi_1 ... \forall\phi_n (P(\phi_1, ..., \phi_n) \leftrightarrow \exists\!\!\!\forall\psi\, F(\phi_1, ..., \phi_n, \psi)), \tag{5.5}$$

where $\exists\!\!\!\forall\psi\, F(\phi_1, ..., \phi_n, \psi)$ is an $\mathcal{L}^+$-formula, with at most $\phi_1, ..., \phi_n$ free.

Definitional extension preserves solvability by the following theorem, according to which any solution of any graph $G_M(\mathcal{L}^+)$ can be extended to a solution of $G_M(\mathcal{L}^P)$.

Theorem 5.6 **(7.17)** *For every $\Gamma \subseteq \mathcal{L}^+$ and its definitional extension F, every kernel model of Γ can be extended to a kernel model of $\Gamma \cup F$.*

The proof amounts to the elimination of symbol(s) P, replacing each $P(S)$ by its definiens $\exists\!\!\!\forall\psi\, F(S, \psi)$. This operation, trivial in FOL, has to be performed recursively (e.g., $P(P(S))$ needs repeated replacements) on a cyclic graph and is given in Appendix.

As a special case, the counterpart of Convention T, $\forall\phi(\mathcal{T}\phi \leftrightarrow \phi)$, satisfies trivially (5.5).

Corollary 5.7 *Each kernel model of any $\Gamma \subseteq \mathcal{L}^+$ can be extended to a kernel model of $\Gamma \cup \{(\mathrm{T})\}$.*

Holding for FOL and higher-order classical logics, this does not contradict Tarski's undefinability theorem. On the one hand, the so-defined $\mathcal{T}$ is just the identity operator, not any truth predicate 'decoding' numbers (or names) as formulas. More significantly, the diagonalization lemma and Tarski's theorem rely on the availability of the substitution operation. Its definition has not been given in LSO and must wait until a future work addresses open formulas.

6 Summary

As in natural reasoning, a paradox in LSO is revealed by deriving from it a specific contradiction that is relevant to the statement of the paradox. From the liar, we deduce that it lies and does not lie, not that snow is white and not white. Most other contradictions do not follow from the liar in LSO with the non-explosive semantics. If explosiveness is desired also from paradoxes, it can be obtained by specializing semikernels to kernels and, in reasoning, by allowing unrestricted (cut) which enables chaining the derived contradiction with its arbitrary consequences.

Unlike in AST, expression of paradoxes in LSO does not require any coding of syntax, as self-reference is inherently available. Paradoxes arise from problematic interpretations of operators. From a purely syntactic perspective, they occur only at the metalevel, similarly to natural language. A paradox requires a sentence or sentences that involve both an operator and an s-quantifier, with both nominal and sentential occurrences of a variable within the same scope. Future results may refine these conditions and provide more specific syntactic guidelines for the paradox-free use of the language, but it is hard to expect any simple, sufficient, and necessary conditions that would ensure this.

A more specific and intuitive characterization of paradoxes is offered in terms of the language graphs. All finite paradoxes involve unresolvable odd cycles, representing negative self-reference of indirect liars. This is no longer merely an empirical generalization but an actual theorem, due to Richardson's [38]. According to it, every graph without infinite outgoing acyclic paths or infinite branchings – in particular, a finite graph – is solvable when it has no odd cycles. We can therefore conclude non-paradoxicality if such a graph remains after inducing values from a given valuation of atoms. Although such cases are rare, the crucial analysis can often be carried out with finitely representable schemas, as was done in our examples. More specific examples and patterns of paradoxes, in theories represented by digraphs closely related to the language graphs, can be found in [11, 14, 48, 50, 51].

All such examples contain either an odd cycle or a form of acyclic Yablo-like pattern. The latter was not explicitly defined here, but it was exemplified by the graph $(\omega, <)$, with each edge subdivided into three edges. These two patterns are the only ones that cause paradoxes in a large class of cases, according to a theorem from [49], where also the Yablo pattern is defined in terms of graphs. The plausible conjecture that they are necessary in all cases remains an open question. Proving this, or presenting a counterexample, would further confirm the usefulness of the approach based on digraph (semi)kernels for analyzing and classifying paradoxes.

The next step is to extend LSO to operators that act on open formulas. Operators that handle their arguments purely syntactically should suffice, but this may also require a form of quotation mechanism. A consistent theory of syntax, addressing in particular the substitution of formulas, can be expected to yield a consistent combination with LSO. Looking further ahead, a satisfactory truth theory for LSO within LSO should not be excluded, thus extending the presented view of the language, whose expressivity conflicts with neither its classical character nor its coherence.

7 Appendix: Language Graphs and (semi)kernels

7.1 Some Facts About (semi)kernels

The following equivalent semikernel condition is used in some proofs.

Fact 7.1 *For any* $L \subseteq \mathbf{V} : \mathbf{E}(L) \subseteq \mathbf{E}^-(L) \subseteq \mathbf{V} \setminus L \iff \mathbf{E}(L) \subseteq \mathbf{E}^-(L) \cap \mathbf{V} \setminus L$.

Proof If $\mathbf{E}(L) \subseteq \mathbf{E}^-(L) \subseteq \mathbf{V} \setminus L$ then $\mathbf{E}(L) \subseteq \mathbf{E}^-(L) = \mathbf{E}^-(L) \cap \mathbf{V} \setminus L$. If $\mathbf{E}(L) \subseteq \mathbf{E}^-(L) \cap \mathbf{V} \setminus L$ then $\mathbf{E}^-(L) \subseteq \mathbf{V} \setminus L$, for if some $x \in \mathbf{E}^-(L) \cap L$ then $\mathbf{E}(x) \cap L \neq \emptyset$, i.e., $\mathbf{E}(L) \not\subseteq \mathbf{V} \setminus L$. □

The two facts below imply equisolvability of graphs, showing that the two have *essentially the same solutions*: each solution of one can be extended to a solution of the other, and each solution of the other, restricted to the first, is its solution. These facts, applied implicitly on the drawings, justify the use of auxiliary duplicates of vertices $\mathbf{S}_M$, without affecting the solutions.

A path $a_0...a_k$ is *isolated* if $\mathbf{E}_G(a_i) = \{a_{i+1}\}$ for $0 \leq i < k$ and $\mathbf{E}^-_G(a_i) = \{a_{i-1}\}$ for $0 < i < k$. A *double edge*, introduced under (3.8), is an isolated path of length 2. *Contraction* of such an isolated path amounts to identifying the first and the last vertex, joining their neighbourhoods and removing the intermediate vertices, i.e., obtaining graph G' where $\mathbf{V}_{G'} = \mathbf{V}_G \setminus \{a_1...a_k\}$, $\mathbf{E}_{G'}(a_0) = \mathbf{E}_G(a_k)$ and $\mathbf{E}^-_{G'}(a_0) = \mathbf{E}^-_G(a_0) \cup \mathbf{E}^-_G(a_k) \setminus \{a_{k-1}\}$. The first fact is a trivial observation.

Fact 7.2 *If G' results from G by contracting an isolated path of even length, then* $\forall K' \in sol(G')\exists! K \in sol(G) : K' \subseteq K$, *and* $\forall K \in sol(G) : K \cap \mathbf{V}_{G'} \in sol(G')$.

The same holds if G' results from a transfinite number of such contractions, provided that no *ray*, i.e., an infinite outgoing path with no repeated vertex, is contracted to a finite path.

The second fact shows that identifying vertices with identical out-neighbourhoods preserves and reflects (semi)kernels. To define this operation, let $R_G \subseteq \mathbf{V}_G \times \mathbf{V}_G$ relate two vertices in G with identical out-neighbourhoods, i.e., $R_G(a, b) \Leftrightarrow \mathbf{E}_G(a) = \mathbf{E}_G(b)$. It is an equivalence, so let $G^{\downarrow}$ denote the quotient graph over equivalence classes, $[v] = \{u \in \mathbf{V}_G \mid R_G(v, u)\}$, with edges $\mathbf{E}_{G^{\downarrow}}([v], [u]) \Leftrightarrow \exists v \in [v], u \in [u] : \mathbf{E}_G(v, u)$. The operation can be iterated any number n of times, denoted by $G^{\downarrow n}$ and defined by: $G^{\downarrow 1} = G^{\downarrow}$ and $G^{\downarrow(n+1)} = (G^{\downarrow n})^{\downarrow}$. Vertices of $G^{\downarrow n}$ are taken as subsets of $\mathbf{V}_G$, $[u]^n = \{v \in \mathbf{V}_G \mid \exists i \leq n : R_{G^{\downarrow i}}([v]^i, [u]^i)\}$. For limit ordinals λ, $G^{\downarrow\lambda}$ is given by

$\mathbf{V}_{G^{\downarrow\lambda}} = \{[u]^\lambda \mid u \in \mathbf{V}_G\}$ where $[u]^\lambda = \bigcup_{i<\lambda}[u]^i = \{v \in \mathbf{V}_G \mid \exists i < \lambda : R_{G^{\downarrow i}}([u]^i, [v]^i)\}$ and

$\mathbf{E}_{G^{\downarrow\lambda}}([v]^\lambda, [u]^\lambda) \Leftrightarrow \exists n \in \lambda : \mathbf{E}_{G^{\downarrow n}}([v]^n, [u]^n)$.

Fact 7.3 *For every ordinal n, and SKr denoting either kernels or semikernels (sol or SK):*

(a) $K \in SKr(G) \Rightarrow \{[v]^n \mid v \in K\} \in SKr(G^{\downarrow n})$, *and*
(b) $K^{\downarrow n} \in SKr(G^{\downarrow n}) \Rightarrow \bigcup K^{\downarrow n} \in SKr(G)$.

Proof (1) The proof for $n = 1$ shows the claim also for every successor n.

(a) $K^{\downarrow} = \{[v] \mid v \in K\}$ is independent, for if $\mathbf{E}_{G^{\downarrow}}([v], [w])$ for some $[v], [w] \in K^{\downarrow}$, then $\mathbf{E}_G(v, w)$ for some $v \in [v], w \in [w]$. But then $v, w \in K$ contradicting independence of K – if $x \in K$ then $[x] \subseteq K$, since $\forall x, y \in [v] : \mathbf{E}_G(x) = \mathbf{E}_G(y)$, so $\mathbf{E}_G(x) \cap K = \emptyset \Leftrightarrow \mathbf{E}_G(y) \cap K = \emptyset$.

If $[v] \in \mathbf{V}_{G^{\downarrow}} \setminus K^{\downarrow}$, then $[v] \subseteq \mathbf{V}_G \setminus K \subseteq \mathbf{E}^-_G(K)$, so $\forall v \in [v]\, \exists w \in K : \mathbf{E}_G(v, w)$. Then $[w] \in K^{\downarrow}$ and $[v] \in \mathbf{E}^-_{G^{\downarrow}}([w]) \subseteq \mathbf{E}^-_{G^{\downarrow}}(K^{\downarrow})$. Thus $\mathbf{V}_{G^{\downarrow}} \setminus K^{\downarrow} \subseteq \mathbf{E}^-_{G^{\downarrow}}(K^{\downarrow})$, so $K^{\downarrow} \in sol(G^{\downarrow})$.

If $K \in SK(G)$ and $[v] \in \mathbf{E}^-_{G^{\downarrow}}(K^{\downarrow})$, i.e., for some $v \in [v], w \in K : v \in \mathbf{E}^-_G(w)$, then $[v] \subseteq \mathbf{E}^-_G(w)$ and $[w] \in K^{\downarrow}$, so $[v] \in \mathbf{E}^-_{G^{\downarrow}}(K^{\downarrow})$, i.e., $\mathbf{E}_{G^{\downarrow}}(K^{\downarrow}) \subseteq \mathbf{E}^-_{G^{\downarrow}}(K^{\downarrow})$, so $K^{\downarrow} \in SK(G^{\downarrow})$.

(b) $K = \bigcup K^{\downarrow} = \{v \in \mathbf{V}_G \mid [v] \in K^{\downarrow}\}$ is independent, for if $\mathbf{E}_G(v, u)$ for some $v, u \in K$, then also $\mathbf{E}_{G^{\downarrow}}([v], [u])$ contradicting independence of $K^{\downarrow}$. If $x \notin K$

then $[x] \notin K^{\downarrow}$, and since $\mathbf{E}_{G^{\downarrow}}([x],[v])$ for some $[v] \in K^{\downarrow}$, so for some $y \in [x]$ and $v \in [v] \subseteq K, \mathbf{E}_G(y, v)$. But since $\mathbf{E}_G(y) = \mathbf{E}_G(x)$, so also $\mathbf{E}_G(x, v)$, i.e., $x \in \mathbf{E}^-_G(K)$. Thus $\mathbf{V}_G \setminus K \subseteq \mathbf{E}^-_G(K)$, and $K \in sol(G)$.

If $K^{\downarrow} \in SK(G^{\downarrow})$ and $v \in \mathbf{E}_G(K)$, then $[v] \in \mathbf{E}_{G^{\downarrow}}(K^{\downarrow}) \subseteq \mathbf{E}^-_{G^{\downarrow}}(K^{\downarrow})$, i.e., $[v] \in \mathbf{E}^-_{G^{\downarrow}}([w])$ for some $[w] \in K^{\downarrow}$. Then $[w] \subseteq K$ and $[v] \subseteq \mathbf{E}^-_G([w])$, so that $\mathbf{E}_G(K) \subseteq \mathbf{E}^-_G(K)$.

(2) We show the claim for limit λ.

(a) If $K \in sol(G)$, let $K^{\downarrow\lambda} = \{[v]^{\lambda} \mid v \in K\}$. If $\mathbf{E}_{G^{\lambda}}([v]^{\lambda},[u]^{\lambda})$ for some $[v]^{\lambda},[u]^{\lambda} \in K^{\downarrow\lambda}$, i.e., $v, u \in K$, then for some $n \in \lambda : \mathbf{E}_{G^n}([v]^n,[u]^n)$, which means that $K^{\downarrow n} = \{[x]^n \mid x \in K\}$ is not a kernel of G^n, contrary to point (1). Hence $K^{\downarrow\lambda}$ is independent. If $[v]^{\lambda} \in \mathbf{V}_{G^{\lambda}} \setminus K^{\downarrow\lambda}$, then $[v]^{\lambda} \subseteq \mathbf{V}_G \setminus K$, so for any $v \in [v]$, there is a $u \in \mathbf{E}_G(v) \cap K$. Then also $[u]^{\lambda} \in \mathbf{E}_{G^{\lambda}}([v]^{\lambda}) \cap K^{\downarrow\lambda}$, hence $\mathbf{V}_{G^{\lambda}} \setminus K^{\downarrow\lambda} \subseteq \mathbf{E}^-_{G^{\lambda}}(K^{\downarrow\lambda})$, and $K^{\downarrow\lambda} \in sol(G^{\downarrow\lambda})$.

If $K \in SK(G)$, i.e., $\mathbf{E}_G(K) \subseteq \mathbf{E}^-_G(K)$ and $[v]^{\lambda} \in \mathbf{E}_{G^{\downarrow\lambda}}(K^{\downarrow\lambda})$, then for some $n \in \lambda : [v]^n \in \mathbf{E}_{G^{\downarrow}n}([w]^n)$ for some $[w]^n \in K^{\downarrow n}$. Then also $[w]^n \in \mathbf{E}^-_{G^{\downarrow n}}(K^{\downarrow n})$, as $K^{\downarrow n} \in SK(G^{\downarrow n})$ by IH, but then also $[w]^{\lambda} \in \mathbf{E}^-_{G^{\downarrow\lambda}}(K^{\downarrow\lambda})$. Thus $\mathbf{E}_{G^{\downarrow\lambda}}(K^{\downarrow\lambda}) \subseteq \mathbf{E}^-_{\downarrow\lambda}(K^{\downarrow\lambda})$.

(b) For a kernel $K^{\downarrow\lambda}$ of $G^{\downarrow\lambda}$, let $K = \bigcup K^{\downarrow\lambda} = \{v \in \mathbf{V}_G \mid [v]^{\lambda} \in K^{\downarrow\lambda}\}$. If $v \in \mathbf{E}^-_G(x)$ for some $x \in K$, then $v \notin K$ for if $v \in K$, i.e., $[v]^{\lambda} \subseteq K$, then $[v]^{\lambda} \in \mathbf{E}^-_{G^{\downarrow\lambda}}([x]^{\lambda}) \cap K^{\downarrow\lambda} \subseteq \mathbf{E}^-_{G^{\downarrow\lambda}}(K^{\downarrow\lambda}) \cap K^{\downarrow\lambda}$ contradicting independence of $K^{\downarrow\lambda}$. If $v \in \mathbf{V}_G \setminus K$, i.e., $[v]^{\lambda} \notin K^{\downarrow\lambda}$, then there is some $[u]^{\lambda} \in \mathbf{E}_{G^{\downarrow\lambda}}([v]^{\lambda}) \cap K^{\downarrow\lambda}$. Since $[u]^{\lambda} \in \mathbf{E}_{G^{\downarrow\lambda}}([v]^{\lambda})$, so for some $n < \lambda, [u]^n \in \mathbf{E}_{G^{\downarrow n}}([v]^n)$, that is, for some $u' \in [u]^n, u' \in \mathbf{E}_G(v)$. Since $[u]^{\lambda} \in K^{\downarrow\lambda}$, so $[u]^n \subseteq [u]^{\lambda} \subseteq K$, hence $v \in \mathbf{E}^-_G(K)$ and $K \in sol(G)$.

If $K^{\downarrow\lambda} \in SK(G^{\downarrow\lambda})$, independence of K follows as above. If $v \in \mathbf{E}_G(K)$, then $[v]^{\lambda} \in \mathbf{E}_{G^{\downarrow\lambda}}(K^{\downarrow\lambda}) \subseteq \mathbf{E}^-_{G^{\downarrow\lambda}}(K^{\downarrow\lambda})$, i.e., for some $n \in \lambda : [v]^n \in \mathbf{E}^-_{G^{\downarrow n}}(K^{\downarrow n})$. By IH, $[v]^n \subseteq \mathbf{E}^-_G(K^{\downarrow n}) \subseteq \mathbf{E}^-_G(K)$. Hence $\mathbf{E}_G(K) \subseteq \mathbf{E}^-_G(K)$. $\square$

7.2 Logical and Graph Equivalences

We formulate logical and some other notions of equivalence in terms of graphs. Two $\mathcal{L}^+_M$ sentences are equivalent, in $G_M(\mathcal{L}^+)$, if they belong to the same kernels, and $\mathcal{L}^+$ sentences are (logically) equivalent if they are so in every language graph. 0

$$\begin{array}{lll} \text{For a graph } G \text{ and } A, B \in \mathbf{V}_G : & A \overset{G}{\Longleftrightarrow} B \text{ iff } \forall K \in sol(G) : A \in K \Leftrightarrow B \in K, \\ \text{for } A, B \in \mathbf{S}^+_M : & A \overset{\mathcal{L}^+_M}{\Longleftrightarrow} B \text{ iff } A \overset{G_M(\mathcal{L}^+)}{\Longleftrightarrow} B, \\ \text{for } A, B \in \mathbf{S}^+ : & A \overset{\mathcal{L}^+}{\Longleftrightarrow} B \text{ iff } \forall M : A \overset{\mathcal{L}^+_M}{\Longleftrightarrow} B. \end{array} \tag{7.4}$$

A more specific equivalence will be used, corresponding to prenex operations. Each sentence can be written in PDNF, that is, prenex normal form with matrix in DNF. Two $\mathcal{L}^+_M$ sentences are *PDNF equivalent*, denoted by $A \overset{P}{\Leftrightarrow} B$, if they have (also) identical PDNFs. To show that PDNF equivalence implies $\mathcal{L}^+$ equivalence, we use a more structural notion of equivalence in a graph.

By $\mathbf{E}_G^*$ we denote the reflexive and transitive closure of $\mathbf{E}_G$ and by $\mathbf{E}_G^*(S)$, for $S \in \mathbf{V}_G$, the subgraph of G induced by all vertices reachable from S. A *common cut* of $A, B \in \mathbf{V}_G$ is a set of vertices $C \subseteq \mathbf{E}_G^*(A) \cap \mathbf{E}_G^*(B)$, such that every path leaving A and prolonged sufficiently far crosses C and so does every path leaving B. (C may intersect A and B and contain vertices on various cycles intersecting A and B.) We say that A and B are *cut equivalent*, $A \overset{\mathcal{C}}{\Leftrightarrow} B$, if there is a common cut C such that for every coherent (satisfying (3.2)) valuation of C, every coherent extension to $\{A, B\}$ forces identical value of A and B. Obviously, if $A \overset{\mathcal{C}}{\Leftrightarrow} B$ in a graph G, then also $A \overset{G}{\Leftrightarrow} B$, as each $K \in sol(G)$ determines a coherent valuation of every common cut of A and B. FOL tautologies/contradictions are in all/none kernels, implying graph equivalence. Contingent propositional/FOL equivalence $P(\overline{\phi}) \Leftrightarrow Q(\overline{\phi})$ implies cut equivalence, $P(\overline{\phi}) \overset{\mathcal{C}}{\Leftrightarrow} Q(\overline{\phi})$, with the shared atoms giving a possible common cut in the graphs (trees) of each instance. We show this for PDNF.

Fact 7.5 *For $A, B \in \mathbf{S}_M^+$ in $G_M(\mathcal{L}^+)$, if $A \overset{P}{\Leftrightarrow} B$ then $A \overset{\mathcal{C}}{\Leftrightarrow} B$, hence $A \overset{\mathcal{L}_M^+}{\Longleftrightarrow} B$.*

Proof Letting $G = G_M(\mathcal{L}^+)$ and assuming $sol(G) \neq \emptyset$, we verify standard prenex transformations, considering only s-quantifiers, as the object-quantifiers can be treated in the same way.

1. The claim holds trivially for B obtained by renaming bound s-variables (avoiding name clashes) in A, as the two have the same subgraph. This is also the case for the subgraphs of $A = \neg\forall\phi D(\phi)$ and $B = \exists\phi\neg D(\phi)$.
2. $A = (\forall\phi D(\phi)) \wedge C \overset{G}{\Leftrightarrow} \forall\phi(D(\phi) \wedge C) = B$, with no free occurrences of ϕ in C. On the schematic subgraph below, $X_i, X_j \ldots$ stand for all $\mathbf{S}_M^+$ and common cut is marked by the waved line.

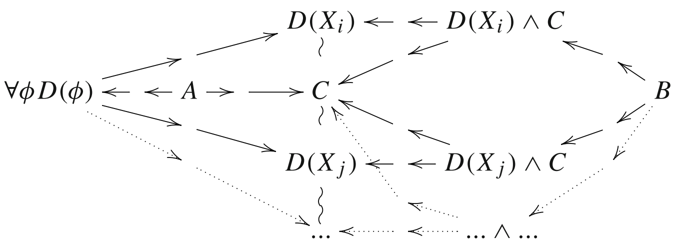

   Inspecting the graph, we see that, for any kernel K:
   $B \in K \Leftrightarrow \big((D(X_i) \wedge C) \in K \text{ for all } X_i\big) \Leftrightarrow \big(C \in K \wedge (D(X_i) \in K \text{ for all } X_i)\big) \Leftrightarrow A \in K$.
3. For $A = \neg\exists\phi D(\phi) \overset{G}{\Leftrightarrow} \forall\phi\neg D(\phi) = B$ the schematic subgraph is as follows:

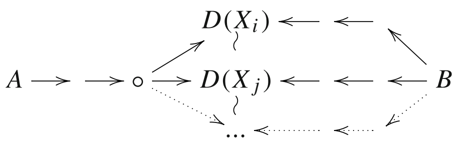

Obviously, for any kernel K : $A \in K \Leftrightarrow \circ \in K \Leftrightarrow (D(X_i) \notin K$ for all $X_i) \Leftrightarrow B \in K$. $\square$

Thus, every sentence in $\mathcal{L}^+$ has an $\overset{\mathcal{L}^+}{\Longleftrightarrow}$-equivalent PDNF sentence. A useful consequence is that, considering below solvability of $G_M(\mathcal{L}^\Phi)$ or $G_M(\mathcal{L}^+)$, we can limit attention to sentences in PDNF.

7.3 No Paradoxes in $\mathcal{L}^\Phi$ – solvability of $G(\mathcal{L}^\Phi)$

Expanding any classical language $\mathcal{L}$ with s-quantifiers to $\mathcal{L}^\Phi$ does not introduce any paradoxes. The following theorem shows a stronger claim, namely, for any domain M, each valution of $\mathcal{L}_M$ sentences (determined by a valuation of atoms $\mathbf{A}_M$) determines a valuation of all $\mathcal{L}^\Phi_M$ sentences.

Theorem 7.6 (3.16) *In any $G_M(\mathcal{L}^\Phi)$, each $\rho \in \mathbf{2}^{\mathbf{S}_M}$ has a unique extension to $\hat{\rho} \in sol(G_M(\mathcal{L}^\Phi))$.*

Proof Graph $G_M(\mathcal{L}^\Phi)$ consists of two subgraphs, the strong component with all s-quantified sentences, $G_M(\mathcal{L}^\Phi \backslash \mathcal{L}) = \bigcup_{A \in \mathbf{S}^\Phi_M \backslash \mathbf{S}_M} G_M(A)$, and the collection $G_M(\mathcal{L}) = \bigcup_{B \in \mathbf{S}_M} G_M(B)$ of trees for the object-language sentences (with 2-cycles at the atoms), with no edges from the latter to the former. Following footnote 3, we also view each $A \in \mathbf{S}^\Phi_M \backslash \mathbf{S}_M$ in $G_M(\mathcal{L}^\Phi \backslash \mathcal{L})$ as tree-like, with (single or double) edges from its external leaves $V \in ext(G_M(A))$, p.14, to the sources of $G_M(V)$, that are trees when $V \in \mathbf{S}_M$. By Lemma 3.17 below, valuation ρ of $\mathbf{S}_M = \mathbf{V}_{G_M(\mathcal{L})}$, determines a solution ρ^-_A of each $G^-_M(A)$ (the subgraph of $G_M(A)$ without its DNF-feet), compatible with every valuation of $ext(G_M(A))$. Hence, these can be combined into $\rho \cup \bigcup_{A \in \mathbf{S}^\Phi_M \backslash \mathbf{S}_M} \rho^-_A$ forcing value $\rho^-_V(V)$ at each $V \in ext(G_M(A))$, and thus determining solutions of all DNF-feet. Each $G_M(A)$ obtains thus a solution $\rho_A \supset \rho^-_A$, yielding a unique $\hat{\rho} = (\rho \cup \bigcup_{A \in \mathbf{S}^\Phi_M \backslash \mathbf{S}_M} \rho_A) \in sol(G_M(\mathcal{L}))$, extending ρ. $\square$

The missing lemma shows that for each sentence $A \in \mathbf{S}^\Phi_M \backslash \mathbf{S}_M$, solution of the subgraph of $G_M(A)$ without its DNF-feet, denoted by $G^-_M(A)$, depends on the valuation of $\mathbf{S}_M$, but not of its external leaves $ext(G_M(A))$, as the second part of the lemma states. Valuation of $ext(G_M(A))$ affects of course values in the DNF-feet in which they occur. In Lemma 3.17, a valuation of atoms $\mathbf{A}_M$ was assumed, but since it induces a valuation of all sentences $\mathbf{S}_M$, we now assume the latter.

Lemma 7.7 (3.17) *For each graph $G_M(\mathcal{L}^\Phi)$ and $A \in \mathbf{S}^\Phi_M$, each valuation ρ of $\mathbf{S}_M$ and of external leaves of $G_M(A)$, $\rho \in \mathbf{2}^{\mathbf{S}_M \cup ext(G_M(A))}$, has a unique extension to $\rho_A \in sol(G_M(A))$. The restriction $\rho|_{\mathbf{S}_M}$ determines the restriction of ρ_A to $G^-_M(A)$: if $\rho|_{\mathbf{S}_M} = \sigma|_{\mathbf{S}_M}$ then $\rho_A|_{G^-_M(A)} = \sigma_A|_{G^-_M(A)}$.*

Proof By Fact 7.5, we can limit attention to sentences in PDNF.

For $A \in \mathbf{S}^\Phi_M$, with the number $q(A) = n + 1$ of s-quantifiers and s-variables, and for n-sequence $\pi \in (\mathbf{S}^\Phi_M)^n$ of sentences instantiating the n s-quantified variables of A, the sources of all feet, $A(\pi S) = D(\pi S)$, $S \in \mathbf{S}^\Phi$, are grandchildren of vertex

$A(\pi) = \exists\!\!\!\forall\phi D(\pi\phi)$. (In the drawing, $\exists\!\!\!\forall = \exists$ and all feet have the common parent $\bullet$; when $\exists\!\!\!\forall = \forall$, their distinct parents are children of $A(\pi)$.) Each foot $A(\pi S)$ represents an application of the same boolean function $d^{\pi}(\phi) = D(\pi\phi)$, evaluating $D(\pi\phi)$ given valuation of its parameters π, ϕ and, possibly, some atoms $L_A \subset \mathbf{S}_M$ occurring in the original matrix $D(...)$. For any $\rho \in \mathbf{2}^{\mathbf{S}_M}$, L_A obtain fixed values so, considering d^{π}, we assume the effects of $\rho(L_A)$ taken into account.

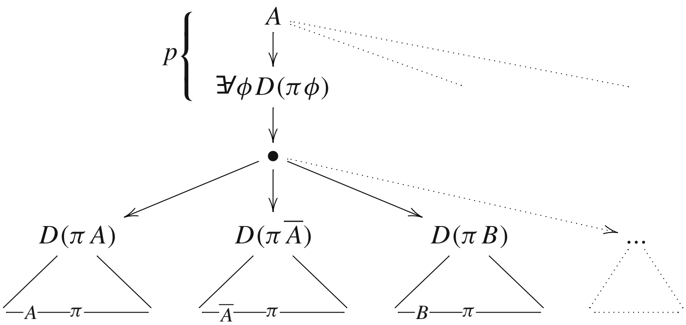

For a path π of sentences instantiating some initial quantifiers of A, let p be the resulting sequence of the successive instances ending with $A(\pi)$. Some 'sinks' of the feet are (or, following footnote 3, have double edges to) vertices from p. By the *internal* leaves of $A(\pi)$, $int(A(\pi))$, we mean those leaves of $G_M(A(\pi))$ which either occur on p or on some paths in this subgraph above its leaves. The *external* leaves are the remaining sentences, excluding atoms L_A, $ext(\pi) = \mathbf{S}_M^{\Phi} \setminus (int(\pi) \cup L_A)$. Branches from $\bullet$ instantiate ϕ with each $S \in \mathbf{S}^{\Phi}$, while ϕ occurs as an atom in DNF D, so each sentence $\mathbf{S}^{\Phi}$ from p is an internal leaf in some feet. If $A \in \mathbf{S}_M^{\Phi} \setminus \mathbf{S}^{\Phi}$, or once an o-quantifier of A is instantiated, sentences on p belong to $\mathbf{S}_M^{\Phi} \setminus \mathbf{S}^{\Phi}$.

1. Depending on whether $\exists\!\!\!\forall$ is $\forall$ or $\exists$, the value at vertex $\exists\!\!\!\forall\phi D(\pi\phi)$, as a function of values of its grandchildren, is either
(*) $\exists\phi D(\pi\phi) = \bigvee_{S\in\mathbf{S}^{\Phi}} d^{\pi}(S)$ or $\forall\phi D(\pi\phi) = \bigwedge_{S\in\mathbf{S}^{\Phi}} d^{\pi}(S)$.
We consider first the case when $|\pi| = q(A) - 1$, i.e., $A(\pi) = \exists\!\!\!\forall\phi D(\pi\phi)$ is the grandparent of the completely substituted (sources of) DNF-feet ($D(\pi A)$, $D(\pi B)$, etc., in the drawing).
Every valuation of sentences from π, abbreviated as $\alpha \in \mathbf{2}^{\pi}$, specializes function $d^{\pi}(\phi)$ to a unary boolean function $d^{\alpha(\pi)}(\phi) = D(\alpha(\pi)\phi)$, and (*) to
(**) $\exists\phi D(\alpha(\pi)\phi) = \bigvee_{S\in\mathbf{S}^{\Phi}} d^{\alpha(\pi)}(S)$ or $\forall\phi D(\alpha(\pi)\phi) = \bigwedge_{S\in\mathbf{S}^{\Phi}} d^{\alpha(\pi)}(S)$.

2. As a boolean function of one variable, $d^{\alpha(\pi)}(\phi)$ is either constant or not. If it is constant, i.e., $d^{\alpha(\pi)}(\phi) = d^{\alpha(\pi)}(\neg\phi)$, then $\exists\!\!\!\forall\phi D(\alpha(\pi)\phi)$ obtains this value in either case of (**). Otherwise, $d^{\alpha(\pi)}(\neg\phi) = \neg d^{\alpha(\pi)}(\phi)$ and, since for each $S \in \mathbf{S}^{\Phi}$ both $d^{\alpha(\pi)}(S)$ and $d^{\alpha(\pi)}(\neg S)$ enter the evaluation of (**), this yields constant $\mathbf{0}$ at their least common predecessor ($\bullet$ when $\exists\!\!\!\forall = \exists$ and $A(\pi)$ when $\exists\!\!\!\forall = \forall$). In this way, for every $\alpha \in \mathbf{2}^{\pi}$, $A(\pi)$ obtains a unique value $\alpha^{\uparrow}(A(\pi))$, induced from all $D(\alpha(\pi)S)$ by (**), but determined already by $d^{\alpha(\pi)}(\phi)$, independently from

(i) valuation $\alpha(A(\pi))$, i.e., if $\alpha_{\mathbf{0}}, \alpha_{\mathbf{1}} \in \mathbf{2}^{\pi, A(\pi)}$ differ only at $A(\pi)$, then $\alpha_{\mathbf{0}}^{\uparrow}(A(\pi)) = \alpha_{\mathbf{1}}^{\uparrow}(A(\pi))$, and

(ii) independently from the valuation of $ext(A(\pi))$, as each external S either enters both evaluation of $d^{\alpha(\pi)}(S)$ and of $d^{\alpha(\pi)}(\neg S)$, with jointly constant contribution to (**) as just explained, or else $S \in \pi$ has a value assigned by α.

By (i), given any $\rho \in \mathbf{2}^{\mathbf{S}_M \cup ext(A(\pi))}$, the cycles from the feet to p, including $A(\pi)$, admit a unique solution $\rho_{\pi,\alpha}$ to the subgraph $G_M(A(\pi))$ of $G_M(A)$, which does not depend on the initial value $\alpha(A(\pi))$. By (ii), also $\rho|_{ext(A(\pi))}$ is inessential, so if $\rho|_{\mathbf{S}_M} = \sigma|_{\mathbf{S}_M}$ then $\rho_{\pi,\alpha}(A(\pi)) = \sigma_{\pi,\alpha}(A(\pi))$.

3. This is the basis for the claim that for each A with $q(A) \geq 1$ and each path π from the source A with $|\pi| < q(A)$, each valuation of $\pi^- = \pi \setminus \{A(\pi)\}$ (treating π as a set) and of $\mathbf{S}_M$ determines a value of $A(\pi)$. We use its formulation above, namely: for each $\rho \in \mathbf{2}^{\mathbf{S}_M \cup ext(\pi)}$ and each $\alpha \in \mathbf{2}^\pi$, vertex $V = A(\pi)$ (above the sources of the feet) obtains a unique value $\alpha^\uparrow(V)$, which depends at most on valuation α (actually, only on its restriction $\alpha|_{\pi^-}$ to the internal leaves above V), but neither on the value (i) of $\alpha(V)$ nor (ii) of $\rho(X)$, for any $X \in ext(A(\pi))$.
The argument amounts to inducing the values bottom-up, starting with the feet's gradparents in point **2**. More formally, we show the claim by induction on $h - l$, where $h \geq 1$ is the distance of the source A from the sources of the feet and l the distance of V from the source A, $h > l \geq 0$. Point **2** gives the basis for $h - l = 1$.

4. The argument from point **2** works also in the induction step. For $0 \leq |\pi| = l < h - 1$, we have the following counterpart of the drawing from **2**, with $A(\pi) = \exists\!\!\!/\phi\exists\!\!\!/\psi D(\pi\phi\psi)$, where $\exists\!\!\!/\psi$ is the sequence of the remaining quantifiers, and ψ_1, ψ_2 at the bottom signal various substitutions for ψ.

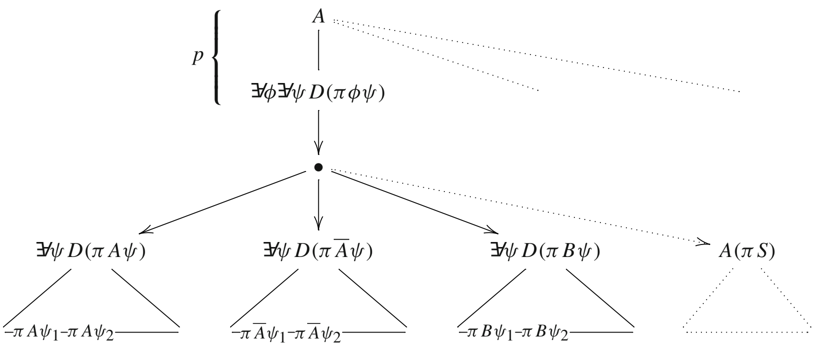

Given $\alpha \in \mathbf{2}^\pi$, IH applied to the lowest triangles in the drawing, i.e., subgraphs $G_M(A(\pi S))$ with sources $A(\pi S)$ for $S \in \mathbf{S}^\Phi$, gives to each $A(\pi S)$ a unique value, independent of the valuation of $ext(G_M(A(\pi S)))$. Consequently $A(\pi\phi)$, viewed as a boolean function, depends only on the values of π and ϕ. For each $\alpha \in \mathbf{2}^\pi$, it yields a unary function $d^{\alpha(\pi)}$ of the value of ϕ.
By the same argument and cases for $d^{\alpha(\pi)}$ as in point **2**, the value $\alpha^\uparrow(A(\pi))$, induced to the common grandparent of all $A(\pi S)$ under $\alpha \in \mathbf{2}^\pi$, is equal whether

$\alpha(A(\pi)) = \mathbf{1}$ or $\alpha(A(\pi)) = \mathbf{0}$, giving point (i) of induction. As for each $A(\pi S)$ its value under $\alpha^{\uparrow}$ is independent from the valuation of $ext(G_M(A(\pi S)))$ by IH, the induced value $\alpha^{\uparrow}(A(\pi))$ is independent from the valuation of $ext(A(\pi)) = \bigcap_{S \in \mathbf{S}_M} ext(A(\pi S))$, giving point (ii) of induction. Thus, $\alpha^{\uparrow}(A(\pi))$ is unique and independent from the valuation of $ext(A(\pi))$ and of $A(\pi)$, establishing the induction step.

5. Thus, the value of the source A is determined, for each $\rho \in \mathbf{2}^{\mathbf{S}_M \cup ext(G_M(A))}$, independently from $\rho|_{ext(G_M(A))}$. Starting now from A and using claim **3** downwards, the value of $A(S)$, for each $S \in \mathbf{S}^{\Phi}$, is determined by ρ and the value of A (independently from the valuation of $ext(A)$). Since A is determined by ρ, so is the value of $A(S)$. Proceeding inductively down the tree of $G_M(A)$, valuation ρ_A^- of $G_M(A)^-$ is seen determined by ρ, independently from the valuation of $ext(G_M(A))$. The latter, giving values to all external leaves, determines the values in all feet of $G_M(A)$, yielding a unique solution ρ_A of the whole graph $G_M(A)$, with $\rho_A^- \subset \rho_A$ and $\rho_A|_{\mathbf{S}_M \cup ext(G_M(A))} = \rho$. □

7.4 The Expressive Power of $\mathcal{L}^{\Phi}$ (and Quantified Boolean Formulas)

By Theorem 3.16/7.6 and Definition 3.18, each $\mathcal{L}$ structure (M, ρ), with $\rho \in \mathbf{2}^{\mathbf{A}_M}$, determines an $\mathcal{L}^{\Phi}$ structure $(G_M(\mathcal{L}^{\Phi}), \hat{\rho})$, with $\hat{\rho} \in sol(G_M(\mathcal{L}^{\Phi}))$. Hence, we do not distinguish between the two; "a structure M" means in this section a pair (M, ρ) or, equivalently, $(G_M(\mathcal{L}^{\Phi}), \hat{\rho})$.

By Theorem 3.19/7.11 below, expanding $\mathcal{L}$ to $\mathcal{L}^{\Phi}$ does not increase the expressive power, as the introduced s-quantification amounts to a complex form of quantification over boolean values. In models of $A = \forall\phi F(\phi)$, F is true for all sentences ϕ, including A itself. Guaranteeing a well-defined value for each sentence (in each structure), the theorem makes this "including itself" harmless, reducing $\forall\phi$ to the propositional quantifier. To verify A it suffices to verify $F(\phi)$ for $\phi = \mathbf{1}$ and $\phi = \mathbf{0}$. This follows provided that every sentential context $F(\phi)$ (having only ϕ free), is a congruence that preserves equivalence, i.e., such that for each pair of $\mathcal{L}^{\Phi}$ sentences A, B,

$$A \overset{\mathcal{L}^{\Phi}}{\Longleftrightarrow} B \text{ implies } F(A) \overset{\mathcal{L}^{\Phi}}{\Longleftrightarrow} F(B). \tag{7.8}$$

Given an internal equivalence $A \leftrightarrow B \overset{\mathcal{L}^{\Phi}}{\Longleftrightarrow} (A \wedge B) \vee (\neg A \wedge \neg B)$, it suffices that for every structure M, if $M \models A \leftrightarrow B$ then $M \models F(A) \leftrightarrow F(B)$, as is for instance the case for classical logic. Let $\top/\bot$ stand for an arbitrary tautology/contradiction in $\mathcal{L}$. 0

Fact 7.9 *For every $\mathcal{L}^{\Phi}$ formula $F(\phi)$ with only ϕ free and for every structure M:* $M \models \forall\phi F(\phi)$ *iff* $M \models F(\top) \wedge F(\bot)$*, and* $M \models \exists\phi F(\phi)$ *iff* $M \models F(\top) \vee F(\bot)$.

Proof If $M \models \forall\phi F(\phi)$ then, in particular, $M \models F(\top)$ and $M \models F(\bot)$, so $M \models F(\top) \wedge F(\bot)$. Conversely, assuming $M \models F(\top) \wedge F(\bot)$, let S be an arbitrary $\mathcal{L}^{\Phi}$-sentence. By Theorem 3.16/7.6, either $M \models S$ or $M \models \neg S$. If $M \models S$ then also

$M \models S \leftrightarrow \top$, hence $M \models F(S)$ by (7.8), since $M \models F(\top)$. If $M \not\models S$ then also $M \models S \leftrightarrow \bot$, hence $M \models F(S)$, since $M \models F(\bot)$. In either case $M \models F(S)$, and since S was arbitrary, $M \models \forall\phi F(\phi)$.

If $M \models \exists\phi F(\phi)$, let S be a sentence for which $M \models F(S)$. By Theorem 3.16/7.6, either $M \models S$ or $M \models \neg S$. In the first case $M \models S \leftrightarrow \top$ and in the latter $M \models S \leftrightarrow \bot$. Thus either $M \models F(\top)$ or $M \models F(\bot)$, hence $M \models F(\top) \vee F(\bot)$. Conversely, if $M \models F(\top) \vee F(\bot)$ then either $M \models F(\top)$ or $M \models F(\bot)$, so in each case $M \models \exists\phi F(\phi)$. $\square$

A special case is the language $\emptyset^{\Phi}$ of quantified boolean sentences, QBS. The unique solution of its also unique (due to the absence of any domain) language graph $G(\emptyset^{\Phi})$ contains exactly true QBS. The right-hand sides of the equivalences in Fact 7.9 reflect the standard semantics of QBS.

By Theorem 3.16, values of $\mathcal{L}$ sentences determine values of all $\mathcal{L}^{\Phi}$ sentences. Consequently, if structures M, N are elementarily equivalent in $\mathcal{L}$, $M \overset{\mathcal{L}}{\equiv} N$, they are so also in $\mathcal{L}^{\Phi}$, $M \overset{\mathcal{L}^{\Phi}}{\equiv} N$.

Fact 7.10 *For any $\mathcal{L}$ structures M and N, $M \overset{\mathcal{L}}{\equiv} N$ iff $M \overset{\mathcal{L}^{\Phi}}{\equiv} N$.*

Proof The non-obvious implication to the right follows by induction on the number of s-quantifiers. Let $M \overset{k}{\equiv} N$ denote that M and N model the same $\mathcal{L}^{\Phi}$ sentences with up to k s-quantifiers, so that $M \overset{\mathcal{L}}{\equiv} N$ corresponds to $M \overset{0}{\equiv} N$, giving the induction basis. For a PDNF sentence $A = \forall\phi\overline{\exists\forall\psi} D(\phi, \overline{\psi})$, where $|\overline{\psi}| = k \geq 0$, suppose that

(m) $M \models A$, i.e., for every $F \in \mathbf{S}^{\Phi} : M \models \overline{\exists\forall\psi} D(F, \overline{\psi})$, while

(n) $N \not\models A$, i.e., for some $F_0 \in \mathbf{S}^{\Phi} : N \not\models \overline{\exists\forall\psi} D(F_0, \overline{\psi})$.

F_0 has some s-quantifiers, as otherwise (m), (n) contradict IH, $M \overset{k}{\equiv} N$. Now $N \not\models \overline{\exists\forall\psi} D(F_0, \overline{\psi})$ implies $N \not\models \forall\phi\overline{\exists\forall\psi} D(\phi, \overline{\psi})$ yielding, by Fact 7.9, either $N \not\models \overline{\exists\forall\psi} D(\top, \overline{\psi})$ or $N \not\models \overline{\exists\forall\psi} D(\bot, \overline{\psi})$. For any $\mathcal{L}$ sentence $P_0 \Leftrightarrow \top$ in the former case, and $P_0 \Leftrightarrow \bot$ in the latter, $N \not\models \overline{\exists\forall\psi} D(P_0, \overline{\psi})$. This last sentence has k s-quantifiers so, by IH, $M \not\models \overline{\exists\forall\psi} D(P_0, \overline{\psi})$, which contradicts (m). An analogical argument shows the induction step for $A = \exists\phi\overline{\exists\forall\psi} D(\phi, \overline{\psi})$. $\square$

For any theory in $\mathcal{L}^{\Phi}$, Fact 7.9 makes it straightforward to construct a theory in $\mathcal{L}$ with the same model class. For any $\mathcal{L}^{\Phi}$ sentence A in PDNF, an $\mathcal{L}$ sentence A^{-}, with $Mod(A) = Mod(A^{-})$, is obtained replacing $\forall\phi F(\phi)$ by $F(\top) \wedge F(\bot)$ and $\exists\phi F(\phi)$ by $F(\top) \vee F(\bot)$. E.g., starting with $A = \forall\phi\exists\psi(C \wedge \phi) \vee (D \wedge \psi) \vee (\phi \wedge \psi)$, with $C, D \in \mathbf{S}$, one application of Fact 7.9 yields

$\exists\psi\big((C \wedge \top) \vee (D \wedge \psi) \vee (\top \wedge \psi)\big) \wedge \exists\psi\big((C \wedge \bot) \vee (D \wedge \psi) \vee (\bot \wedge \psi)\big)$,

which simplifies to: $\exists\psi(C \vee (D \wedge \psi) \vee \psi) \wedge \exists\psi(D \wedge \psi) \Longleftrightarrow \exists\psi(D \wedge \psi) \overset{7.9}{\Longleftrightarrow} (D \wedge \top) \vee (D \wedge \bot) \Longleftrightarrow D$,

so $Mod(A) = Mod(D)$. Proceeding thus by induction on the number of s-quantifiers (in PDNF of $\mathcal{L}^{\Phi}$ sentences), Fact 7.9 yields $\forall A \in \mathbf{S}^{\Phi} \exists A^{-} \in \mathbf{S} : Mod(A) = Mod(A^{-})$, establishing

Theorem 7.11 (3.19) *For every $\Gamma \subseteq \mathcal{L}^{\Phi}$ there is a $\Gamma^- \subseteq \mathcal{L}$ with $Mod(\Gamma) = Mod(\Gamma^-)$.*

7.5 Solvability of $G(\mathcal{L}^+)$ and of Definitional Extensions

The proof of Lemma 3.17/7.7 relies on each DNF-foot being a boolean function. The proof can be repeated, ensuring the absence of paradoxes in $\mathcal{L}^+$, for instance, if each operator is constant. Each language graph $G_M(\mathcal{L}^+)$ is thus trivially solvable, as stated by Theorem 3.20, although such interpretations of operators are hardly sufficient. We show a bit more general fact. A non-paradoxical language $\mathcal{L}^+$ remains free from paradoxes after a definitional extension, namely, an expansion with a fresh operator P defined by a sentence

$$\forall\phi(P(\phi) \leftrightarrow \exists\!\!\!\forall\psi F(\phi, \psi)), \tag{5.5}$$

where ϕ is a list of variables, one for each argument of P, and $\exists\!\!\!\forall\psi F(\phi, \psi)$ is an $\mathcal{L}^+$-formula (with no other free variables than ϕ). An expansion of $\mathcal{L}^+$ with P to $\mathcal{L}^P$ adds to each graph $G(\mathcal{L}^+)$ all complex sentences with P and atomic 2-cycles $P(S) \leftrightarrows \overline{P(S)}$, for every sentence S of $\mathcal{L}^P$. For the special extension (5.5), $G(\mathcal{L}^+)$ can be extended to $G(\mathcal{L}^P)$ by drawing a double edge from $P(S)$ to its defining sentence $\exists\!\!\!\forall\psi F(S, \psi)$, instead of an edge to $\overline{P(S)}$. Every kernel of such a graph determines a kernel of the graph with atomic 2-cycles where each $P(S)$ and $\exists\!\!\!\forall\psi F(S, \psi)$ obtain the same values, and vice versa. Lemma 7.16 below, giving immediately Theorem 5.6, shows that for any language graph G for $\mathcal{L}^+$, such an extension G^P, for P axiomatized by (5.5), preserves solutions of G. Its proof amounts to the elimination of symbol P, replacing each $P(S)$ by its definiens $\exists\!\!\!\forall\psi F(S, \psi)$. Such a replacement, trivial in FOL, must proceed recursively on a cyclic graph (e.g., $P(P(S))$ needs repeated replacements) and involves some technicalities. These end with the paragraph before Lemma 7.16.

The proof assumes a language graph G over some domain M, in which no two vertices have equal out-neighbourhoods. (If G contains such vertices, as language graphs typically do, their identification preserves essentially the solutions by Fact 7.3, and we apply the construction and fact below to the so quotiented G.) The graph $G^P = G_M(\mathcal{L}^P)$ contains G as an induced subgraph.

As the first step, we quotient atoms of G^P containing P. Let $\simeq$ be congruence on $\mathcal{L}^P_M$-sentences induced by the basic reflexive relation $P(S) \simeq_0 \exists\!\!\!\forall\psi F(S, \psi)$, for every $\mathcal{L}^P_M$-sentence S. For every operator Q distinct from P, we identify every two atoms $Q(A_1 ... A_n) \simeq Q(B_1 ... B_n)$ when $A_i \simeq B_i$ for $1 \leq i \leq n$. Each equivalence class contains an atom $Q(S_1 ... S_n)$ for some $S_i \in \mathbf{S}^+_M$, not containing any P, so in the following we can assume only such atoms present. It is a simple observation that quotient $q : G^P \to H$, where $\mathbf{E}_H(q(x)) = \{q(y) \mid y \in \mathbf{E}_{G^P}(x)\}$ in the resulting graph H, reflects kernels, so the preimage of any kernel of H is a kernel of G^P.

We now map $\gamma : H \to G$, performing a sequence of identifications $\gamma_i : H_{i-1} \to H_i$, for $0 < i \in \omega$ and $H_0 = H$. Each γ_i is identity on the subgraph G of H_i, identifying some vertices from $\mathbf{V}_i \setminus \mathbf{V}_G$ with some in $\mathbf{V}_G$. First, we identify $\gamma_1(P(S)) = \exists\!\!\!\forall\psi F(S, \psi)$, removing the double edge and the intermediate vertex $\bullet_{P(S)}$ between $P(S)$ and its definiens $\exists\!\!\!\forall\psi F(S, \psi)$, for $S \in \mathbf{S}^P_M$. Then $\gamma_{i+1}(v) = w$ when vertices

$v \in \mathbf{V}_i \setminus \mathbf{V}_G$ and $w \in \mathbf{V}_G$ have the same out-neighbourhood. More precisely, (7.12) below defines γ inductively, starting with $\mathbf{V}_0 = \mathbf{V}_H, \mathbf{E}_0 = \mathbf{E}_H$.

$$
\begin{array}{ll}
i=1, & \text{letting } Re_0 = \bigcup\{\{P(S), \bullet_{P(S)}\} \mid S \in \mathbf{S}^P_M, \{\bullet_{P(S)}\} = \mathbf{E}_0(P(S))\} \text{ define:} \\
& \gamma_1(v) = \begin{cases} \exists\!\!\!/\psi F(S,\psi), & \text{if } v = P(S) \text{ for any } S \in \mathbf{S}^P_M \\ v & \text{if } v \notin Re_0 \end{cases} \\
& \text{The resulting graph } H_1 \text{ is given by:} \\
& \mathbf{V}_1 = \mathbf{V}_0 \setminus Re_0, \text{ and } \mathbf{E}_1(v) = \{\gamma_1(w) \mid w \in \mathbf{E}_0(v)\} \setminus Re_0 \\
i+1, & \text{letting } Re_i = \{v \in \mathbf{V}_i \setminus \mathbf{V}_G \mid \exists w \in \mathbf{V}_G : \mathbf{E}_i(v) = \mathbf{E}_i(w)\} \text{ define:} \\
& \gamma_{i+1}(v) = \begin{cases} w \in \mathbf{V}_G \text{ such that } \mathbf{E}_i(v) = \mathbf{E}_i(w) & \text{if } v \in Re_i \\ v & \text{if } v \notin Re_i \end{cases} \\
& \text{The resulting graph } H_{i+1} \text{ is given by:} \\
& \mathbf{V}_{i+1} = \mathbf{V}_i \setminus Re_i \text{ and } \mathbf{E}_{i+1}(v) = \mathbf{E}_i(\gamma_{i+1}(v)) \setminus Re_i \\
\multicolumn{2}{l}{\text{Let } \gamma(v) = \gamma_n(v), \text{ for } v \in \mathbf{V}_H, \text{ where } n \in \omega \text{ is the least such that } \forall m > n : \gamma_m(v) = \gamma_n(v).}
\end{array}
\tag{7.12}
$$

The resulting function γ is well-defined by the assumption that G has no pair of vertices with identical out-neighbourhoods. For $A, B \in \mathbf{V}_H$ and $n \in \omega$, we denote by $A \sim_n B$ that $\gamma_n(A) = \gamma_n(B)$, and by $A \sim B$ that $\gamma(A) = \gamma(B)$, i.e., $\exists n \in \omega : A \sim_n B$. 0

Example 7.13 Let $P(\phi) \leftrightarrow \exists\psi(\phi \wedge \psi)$ and, for some $S \in \mathbf{S}_M$, consider vertex $P(P(S)) \in \mathbf{V}_H$. The relevant parts of graph H are sketched below with $A/_X$, $B/_X$, ... denoting the instances of the grandparent with A, B... substituted for the $\exists$-quantified ψ. The subscripts $_{L,R}$ mark these instantiations in the respective subgraphs, e.g., $A/_L = P(S) \wedge A$ and $A/_R = \exists\phi(S \wedge \phi) \wedge A$. Sentences A, B, ... (and $\overline{A}$, $\overline{B}$, ...) are duplicated in both subgraphs to increase readability, but they are actually the same vertices.

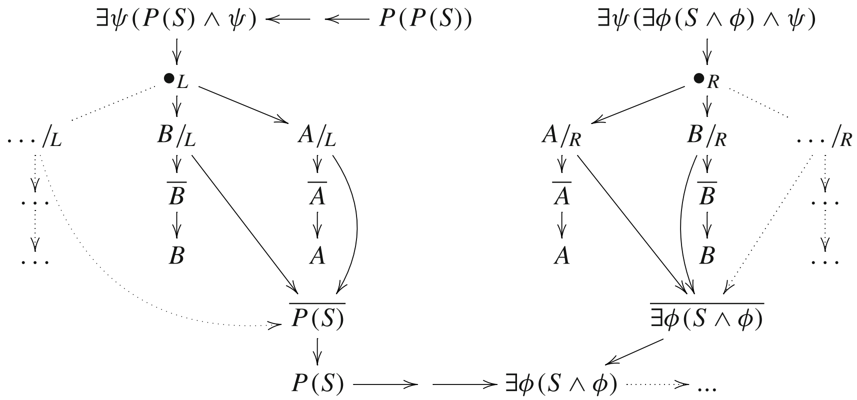

1. $P(P(S)) \sim_1 \exists\psi(P(S) \wedge \psi)$ and $P(S) \sim_1 \exists\psi(S \wedge \psi)$, hence $\mathbf{E}_1(\overline{P(S)}) = \{\gamma_1(P(S))\} = \{\exists\psi(S \wedge \psi)\} = \mathbf{E}_1(\overline{\exists\psi(S \wedge \psi)})$ and, consequently,
2. $\overline{P(S)} \sim_2 \overline{\exists\psi(S \wedge \psi)}$. Then, for each $A \in \mathbf{S}^P_M$, $\mathbf{E}_2(A/_L) = \{\overline{\exists\psi(S \wedge \psi)}, \overline{A}\} = \mathbf{E}_2(A/_R)$, so
3. $A/_L \sim_3 A/_R$, for every $A \in \mathbf{S}^P_M$.
4. Consequently, $\bullet_L \sim_4 \bullet_R$ and then

5. $\exists\psi(\exists\phi(S \wedge \phi) \wedge \psi)) \sim_5 \exists\psi(P(S) \wedge \psi) \sim_1 P(P(S))$, leaving only G's subgraph to the right. $\square$

The equivalence $\sim$ is a congruence on $\mathbf{V}_H$ in the sense that if all out-neighbours of A and B are $\sim$-equivalent then also $A \sim B$, i.e., for $\mathbf{E}_H(A) = \{A_i \mid i \in I\}$ and $\mathbf{E}_H(B) = \{B_i \mid i \in I\}$:

$$\text{if } (\forall i \in I : A_i \sim B_i) \text{ then } A \sim B. \tag{7.14}$$

This holds since each sentence subgraph $G_M(A)$ (tree $\mathrm{T}_M(A)$) has finite height $h(A)$, in particular distance from the source A to atoms $P(S)$ of $G_M(A)$ is at most $h(A)$. Hence, if $\forall i \in I : A_i \sim B_i$ then $\exists n \leq \max\{h(A), h(B)\} \forall i \in I : A_i \sim_n B_i$.[11] The equality $\gamma_n(A_i) = \gamma_n(B_i)$ implies, in turn, that $\mathbf{E}_n(A) = \{\gamma_n(A_i) \mid i \in I\} = \{\gamma_n(B_i) \mid i \in I\} = \mathbf{E}_n(B)$, which yields $A \sim_{n+1} B$. 0

Fact 7.15 *(a)* $\forall S \in \mathbf{S}_M^P \setminus \mathbf{S}_M \; \exists Q \in \mathbf{S}_M : Q \sim S$*, hence* $\Gamma(H) = G$*,*
(b) H *and* G *have essentially the same solutions,*
(c) Every solution of G *extends to a unique solution of* G^P*.*

Proof Point (a) is shown by induction on the number p of Ps in a sentence $S \in \mathbf{S}_M^P \setminus \mathbf{S}_M$.

1. If $p = 1$ and S is atomic, then $S = P(R)$ for some $R \in \mathbf{S}_M$, so $S \sim_1 \exists\!\!\!/\psi F(R, \psi) \in \mathbf{S}_M$.

2. If $p = 1$ and S is not atomic, we proceed by structural induction on S, with point **1** providing the basis and induction hypothesis IH_2:

i. $\bigwedge_{i \in I} S_i$, for finite I. By IH_2, for each S_i there is $Q_i \in \mathbf{S}_M$ with $S_i \sim Q_i$, so $\bigwedge_{i \in I} S_i \sim \bigwedge_{i \in I} Q_i$ by (7.14), and $\bigwedge_{i \in I} Q_i \in \mathbf{S}_M$.

ii. $\neg A$. By IH_2, $A \sim Q$ for some $Q \in \mathbf{S}_M$, so $\neg A \sim \neg Q$ by (7.14), while $\neg Q \in \mathbf{S}_M$.

iii. $S = \exists\!\!\!/\phi A(\phi)$, where ϕ does not occur under P, so that $S = \exists\!\!\!/\phi A(\phi, P(R))$, for some $R \in \mathbf{S}_M$ and context $A(\phi, \_)$ with no P. Since $P(R) \sim_1 \exists\!\!\!/\psi F(R) \in \mathbf{S}_M$, taking $Q = \exists\!\!\!/\phi A(\phi, \exists\!\!\!/\psi F(R)) \in \mathbf{S}_M$, we obtain $A(T, P(R)) \sim A(T, \exists\!\!\!/\psi F(R))$ for every $T \in \mathbf{S}_M^P$ by (7.14), i.e., for all grandchildren of S and Q. By (7.14), this yields $S \sim Q$.

iv. $S = \exists\!\!\!/\phi A(P(C(\phi)))$, i.e., S contains quantification into P, for some contexts $A(\_), C(\_)$ without any P, as $p = 1$. For grandchildren of S, namely, $A(P(C(T)))$ for all $T \in \mathbf{S}_M^P$, the equivalence $P(C(T)) \sim_1 \exists\!\!\!/\psi F(C(T), \psi)$ gives $A(P(C(T))) \sim A(\exists\!\!\!/\psi F(C(T), \psi))$ by (7.14). Sentences on the left, for all $T \in \mathbf{S}_M^P$, comprise all grandchildren of S, and those on the right all grandchildren of $Q = \exists\!\!\!/\phi A(\exists\!\!\!/\psi F(C(\phi), \psi) \in \mathbf{S}_M$, so $S \sim Q$ by (7.14).

3. For the induction step for $p > 1$, the two cases depend on whether P is nested or not.

i. If the number of Ps not nested under others is $n > 1$, consider all these highest Ps in $T_M(S)$, i.e., $S = C(P(A_1), ..., P(A_n))$, where $C(\_)$ contains no Ps. For $R = C(\exists\!\!\!/\psi F(A_1, \psi), ..., \exists\!\!\!/\psi F(A_n, \psi))$, $S \sim R$ by (7.14). R has $p - n < p$ Ps so, by IH, $R \sim Q$ for some $Q \in \mathbf{S}_M$. Hence $S \sim Q$.

[11] This implication fails in general graphs for $\sim$ defined in (7.12) from some basis $\sim_1$, when I is infinite and distance from $A_i, B_i, i \in I$, to relevant pairs $X \sim_1 Y$ is unbounded.

ii. If all Ps are nested under each other, then $S = C(P(A))$ for some context $C(\_)$ without any Ps, and with $p-1$ occurrences of P in A. $P(A) \sim_1 \exists\!\!\!\forall\psi F(A,\psi)$ and, by IH, $\exists\!\!\!\forall\psi F(A,\psi) \sim R$ for some $R \in \mathbf{S}_M$, so that also $P(A) \sim R$. Then $C(P(A)) \sim C(R)$, by (7.14) and $C(R) \in \mathbf{S}_M$, as required.
The equality $\Gamma(H) = G$ follows since each $S \in \mathbf{V}_H \setminus \mathbf{V}_G$ represents a sentence in $\mathbf{S}^P_M \setminus \mathbf{S}_M$.

(b) For $i \geq 0$, H_i in (7.12) is the quotient of H by $\sim_{1,\dots,}\sim_i$. By Fact 7.2, H_1 has essentially the same solutions as H. (No ray is contracted to a finite path, because the case $P(S) \sim_1 \exists\!\!\!\forall\psi F(S,\psi)$ is applied at most finitely many times along each path under each sentence Q, since Q contains at most finitely many nested Ps.) By Fact 7.3, the same holds for H_1 and every H_i, $i > 1$, including limits H_λ. Thus, H and $\Gamma(H) = G$ have essentially the same solutions.

(c) By the observation before this fact, quotient $G^P \to H$ reflects solutions, so that the preimage of every solution of H is a solution of G^P. Using (b), each solution of G extends to one for G^P. □

Let *definitional extension* refer to any well-ordered chain starting with any theory $\Gamma_0 \subseteq \mathcal{L}_0 \subseteq \mathcal{L}^+$ and adding, at step $i+1$, axiom (5.5) with a fresh operator $P \notin \mathcal{L}_i$ and $F(\phi,\psi) \in \mathcal{L}_i$, for language $\mathcal{L}_i$ of theory Γ_i obtained at step i. In the limits, the language and theory are unions of all steps. The following counterpart of model theoretic conservativeness of usual definitional extensions holds.

Lemma 7.16 *Each solution of a language graph $G_0 = G_M(\mathcal{L}_0)$ extends to a solution of the graph of its definitional extension.*

Proof Fact 7.15.(c) gives the claim for an extension with a single operator. By IH, definitional extension G_i of G_0 with $P_1, \dots, P_i$, preserves all solutions of G. Graph G_{i+1}, obtained now by adding P_{i+1}, whose definiens F_{i+1} can utilize P_j, $j \leq i$, preserves by Fact 7.15 solutions of G_i, and hence of G. This establishes successor step.

For any limit, the language $\mathcal{L}^\omega_M = \bigcup_{i\in\omega} \mathcal{L}^i_M$ extends the initial language $\mathcal{L}^0_M$ with ω operators $P_1, P_2, \dots$ introduced on the way. Its graph $G_\omega = \bigcup_{i\in\omega} G_i$, with unions of vertices and of edges, contains all double edges from the new operator's instances to their definienses. We repeat the proof with the unions of all equivalences used along the way. As the first step, let $\simeq^\omega$ be a congruence on $\mathcal{L}^\omega_M$-sentences induced from the relation $A \simeq^\omega B \Leftrightarrow \exists n \in \omega : A \simeq^n B$, where $\simeq^n$ is the congruence $\simeq$ on $\mathcal{L}^n_M$-sentences from step n. Identification of all atoms $Q(A_1 \dots A_k) \simeq^\omega Q(B_1 \dots B_k)$ when $A_i \simeq^\omega B_i$ for $1 \leq i \leq k$ gives a quotient H reflecting kernels as before. Each equivalence class contains an atom from $\mathcal{L}^0_M$. Let H be the resulting graph, and H_i its restriction to the subgraph induced by vertices of G_i (with the atoms identified as just described), so that $H = \bigcup_{i\in\omega} H_i$.

In the chain $G_0 = H_0 \subseteq H_1 \subseteq H_2 \subseteq \dots$, for each pair of subsequent $H_{i-1} \subseteq H_i$, the construction from (7.12) yields $\gamma^i : H_i \to H_{i-1}$ satisfying Fact 7.15. Composing $\gamma^1(\gamma^2(\dots(\gamma^{i-1}(\gamma^i(H_i)))\dots))$ gives surjective $\overline{\gamma}^i : H_i \to G_0$, where $\overline{\gamma}^j(H_i) = \overline{\gamma}^i(H_i)$ for any $j \geq i$. Hence, the union $\gamma^\omega = \bigcup_{i\in\omega} \gamma^i$ gives a surjective quotient $\gamma^\omega : H \to G_0$, reflecting solutions. □

A language with a solvable graph remains non-paradoxical under a definitional extension.

Theorem 7.17 (5.6) *For every $\Gamma \subseteq \mathcal{L}^+$ and its definitional extension F, every kernel model of Γ can be extended to a kernel model of $\Gamma \cup F$.*

One comment is appropriate here. The solvability of the language graphs, as given by the theorem above and summarized in Theorem 3.20, demonstrates the consistency of the language. This does not, of course, exempt one from proving it when a specific interpretation of operators is desired – its consistency must still be verified. This applies in particular to s-equality. While the solvability of language graphs with its intended interpretation can be expected, it has not been proven here.

7.6 (Cut) Preserves Consistency

An important theorem from [28] states that if every induced proper subgraph has a semikernel then the graph has a kernel. For language graphs, a more relevant question might be whether a kernel exists that contains a given theory, provided the existence of semikernels that extend the theory to some (finite) parts of the language. The following fact provides such a compactness-like statement for any language graph. (By $S \Subset X$, we denote that S is a finite subset of X.)

Fact 7.18 *For $\Gamma \subseteq \mathcal{L}^+$ and any language graph $G = G_M(\mathcal{L}^+)$, if for each $S \Subset \mathbf{S}^+_M$ there is a semikernel of G containing Γ and covering S, then G has a kernel containing Γ.*

Proof Let SK_Γ denote all semikernels of G containing Γ. For a finite set $X \Subset \mathbf{S}^+_M$, denote by

$SK_X = \{L \in SK_\Gamma \mid X \subseteq L\}$ – semikernels of G containing Γ and X, and
$SK^c_X = \{L \in SK_\Gamma \mid X \subseteq \mathbf{E}^-_G[L]\}$ – semikernels of G containing Γ and covering X.

The set $F = \{SK^c_X \mid X \Subset \mathbf{S}^+\}$ has the finite intersection property by the main assumption. Let $U \supseteq F$ be an ultrafilter on $\mathcal{P}(SK_\Gamma)$, existing by the ultrafilter lemma. Points **1-3** below show that $K = \{S \in \mathbf{S}^+_M \mid SK_S \in U\}$ is a kernel of G ($\__S$ abbreviating now $\__{\{S\}}$ for a single sentence S).

1. K covers G. If $S \in \mathbf{S}^+_M \setminus K$, then $SK_S \notin U$ so $\overline{SK_S} = \{L \in SK_\Gamma \mid S \notin L\} \in U$. If also $SK_{\neg S} \notin U$, then $\overline{SK_{\neg S}} = \{L \in SK_\Gamma \mid \neg S \notin L\} \in U$. Hence, if both $S \notin K$ and $\neg S \notin K$, then $nS = \overline{SK_S} \cap \overline{SK_{\neg S}} = \{L \in SK_\Gamma \mid S \notin L \wedge \neg S \notin L\} \in U$. As $SK_S \cup SK_{\neg S} = SK^c_{\neg S} \in F \subseteq U$, so $nS \cap SK^c_{\neg S} = \emptyset \in U$ contradicts U being an ultrafilter. Hence $S \in K$ or $\neg S \in K$ for each $S \in \mathbf{S}^+_M$. Since $SK_S \cap SK_{\neg S} = \emptyset$, also $S \notin K$ or $\neg S \notin K$, so $S \in K \Leftrightarrow \neg S \notin K$ for every sentence S.
2. Independence of K is shown for each kind of vertex $S \in K$ by showing $\mathbf{E}_G(S) \subseteq \mathbf{V}_G \setminus K$.

i. If $\neg S \in K$ then $\mathbf{E}_G(\neg S) = \{S\} \subset \mathbf{V}_G \setminus K$, by **1**, while for literals $S \in K \Leftrightarrow \mathbf{E}_G(S) = \{\overline{S}\} \subset \mathbf{V}_G \setminus K$ ($\overline{S}$ is dual literal to S). The two other main connectives in a nonatomic S give the cases:

ii. For a conjunction $\mathbf{E}_G(S_1 \wedge S_2) = \{\neg S_1, \neg S_2\}$, hence $SK_{S_1 \wedge S_2} \cap SK_{\neg S_i} = \emptyset$, for $i \in \{1, 2\}$, so either both $\neg S_i \notin K$ or else $S_1 \wedge S_2 \notin K$, i.e., $S_1 \wedge S_2 \in K \Rightarrow \mathbf{E}_G(S_1 \wedge S_2) \subset \mathbf{V}_G \setminus K$.

iii. For a $\forall$-quantified sentence $\mathbf{E}_G(\forall x F(x)) = \{\neg F(m) \mid m \in M\}$, so $SK_{\forall x F(x)} \cap SK_{\neg F(m)} = \emptyset$ for each $m \in M$, hence if $\forall x F(x) \in K$ then $\mathbf{E}_G(\forall x F(x)) \subset \mathbf{V}_G \setminus K$. (The same works for $\forall \phi F(\phi)$.)

3. We show $\mathbf{E}_G(K) \subseteq \mathbf{E}^-_G(K)$ by a similar case analysis.

i. For literals $S \in K \overset{\mathbf{1}}{\Leftrightarrow} \overline{S} \notin K$, so $S \in K \Leftrightarrow \mathbf{E}_G(S) = \{\overline{S}\} \subseteq \mathbf{E}^-_G(S) \subseteq \mathbf{E}^-_G(K)$.

ii. $\mathbf{E}_G(S_1 \wedge S_2) = \{\neg S_1, \neg S_2\}$ and $(SK_{S_1} \cap SK_{S_2}) \cap SK^c_{S_1 \wedge S_2} = SK_{S_1 \wedge S_2}$, so

(*) $(SK_{\neg S_1} \notin U \wedge SK_{\neg S_2} \notin U) \overset{\mathbf{1}}{\Leftrightarrow} (SK_{S_1} \in U \wedge SK_{S_2} \in U) \Leftrightarrow SK_{S_1 \wedge S_2} \in U \overset{\mathbf{1}}{\Leftrightarrow} SK_{\neg(S_1 \wedge S_2)} \notin U$.

Thus, if $(S_1 \wedge S_2) \in K$ then each $\neg S_i \notin K$ so $S_i \in K$ and $\mathbf{E}_G(S_1 \wedge S) \subseteq \mathbf{E}^-_G(K)$.

For negated conjunction: $\neg(S_1 \wedge S_2) \in K \overset{\mathbf{1}}{\Leftrightarrow} (S_1 \wedge S_2) \notin K \overset{*}{\Leftrightarrow} \neg S_1 \in K \vee \neg S_2 \in K \Rightarrow \mathbf{E}_G(\neg(S_1 \wedge S_2)) = \{S_1 \wedge S_2\} \subseteq \mathbf{E}^-_G(K)$.

iii. $\mathbf{E}_G(\forall x F(x)) = \{\neg F(m) \mid m \in M\} \subseteq \mathbf{V}_G \setminus K \Leftrightarrow \{F(m) \mid m \in M\} \subseteq K$. (The proof for s-quantified $\forall \phi F(\phi)$ is identical, so we write only $\forall x F(x)$.) For each $L \in SK^c_{\forall x F(x)}$:

$\forall x F(x) \in L \Leftrightarrow \{F(m) \mid m \in M\} \subseteq L \Leftrightarrow \{\neg F(m) \mid m \in M\} \subseteq \mathbf{V}_G \setminus L$.

Thus $SK_{\forall x F(x)} \cap SK_{\neg F(m)} = \emptyset$, for every $m \in M$, so

(*) $SK_{\forall x F(x)} \in U \Leftrightarrow \forall m \in M : SK_{\neg F(m)} \notin U \overset{\mathbf{1}}{\Leftrightarrow} \forall m \in M : SK_{F(m)} \in U$, which yields:

$\forall x F(x) \in K \Leftrightarrow SK_{\forall x F(x)} \in U \overset{*}{\Leftrightarrow} \forall m \in M : SK_{F(m)} \in U \Rightarrow \mathbf{E}_G(\forall x F(x)) \subseteq \mathbf{E}^-_G(K)$,

and for the negated quantifier:

$\neg(\forall x F(x) \in K) \overset{\mathbf{1}}{\Leftrightarrow} (\forall x F(x)) \notin K \Leftrightarrow SK_{\forall x F(x)} \notin U \overset{*}{\Leftrightarrow} \exists m \in M : SK_{\neg F(m)} \in U \Rightarrow \mathbf{E}_G(\neg \forall x F(x)) = \{\forall x F(x)\} \subseteq \mathbf{E}^-_G(\neg F(m)) \subseteq \mathbf{E}^-_G(K)$. □

From a syntactic perspective, a limitation of this fact is that it pertains to a single graph, $G_M(\mathcal{L}^+)$, rather than a theory, and it involves all $\mathbf{S}^+_M$ sentences, not just $\mathbf{S}^+$. The idea of a theory having a consistent extension to the entire language, given that it has such extensions to its finite parts, is better captured by the following lemma. It leads to the theorem that the LSO-unprovability of any (specific) contradiction implies the existence of a solvable graph, and thus, the unprovability of any contradiction even when using (cut). For a finite set $Q \Subset \mathbf{S}^+$, we let $\bot_Q = \bigvee_{S \in Q}(S \wedge \neg S)$.

Lemma 7.19 *For Γ, Δ over a countable $\mathcal{L}^+ \in FOL^+$,*
$(\forall Q \Subset \mathbf{S}^+ : \Gamma \nvdash \bot_Q) \Rightarrow \exists G, K \in sol(G) : \Gamma \subseteq K$.

Proof To use Fact 7.18, we construct a graph G over domain M with semikernels containing Γ and covering every finite subset of $\mathbf{S}^+_M$. Letting I index finite subsets of $\mathbf{S}^+$, the assumption gives a semikernel L_i of a graph G_i containing Γ and covering $\bot_i$, for every $i \in I$. Let G be the language graph over $M = \prod M_i$ with $f^M(\prod m_i) = \prod f^{M_i}(m_i)$ and constants $c^M = \prod c^{M_i}$ Define inductively the operation $\_^{\uparrow M}$, lifting terms $\mathbf{T}^+_{M_i} \to \mathcal{P}(\mathbf{T}^+_M)$ and sentences $\mathbf{S}^+_{M_i} \to \mathcal{P}(\mathbf{S}^+_M)$:

$m_i^{\uparrow M} = (m_i)^{\uparrow M} = \{n \in M \mid n_i = m_i\} = \{m_i\} \times \prod_{j \neq i} M_j$, for $m_i \in M_i$, and likewise

$(f(m_i))^{\uparrow M} = f^M(m_i^{\uparrow M}) = \{f^M(n) \in \mathbf{T}_M^+ \mid n \in m_i^{\uparrow M}\}$ and

$F(m_i)^{\uparrow M} = F(m_i^{\uparrow M}) = \{F(n) \in \mathbf{S}_M^+ \mid n \in m_i^{\uparrow M}\}$, for a formula $F(\phi_1 ... \phi_k)$ and $m_i \in M_i^k$.

Notation $F(m_i)$ implies that the only M_i elements are among m_i. In general, S_i denotes an $\mathcal{L}_{M_i}^+$ sentence with possibly some elements from M_i, that equals $S \in \mathbf{S}^+$ if no such elements occur. Then $S_i^{\uparrow M}$ is the set of sentences obtained by replacing each $m_i \in M_i$ by all $m_i^{\uparrow M}$. (S_i denotes also an $\mathcal{L}^+/\mathcal{L}_M^+$ sentence S with all terms interpreted in/projected onto M_i.) Some observations:

(a) $\bigcup\{m_i^{\uparrow M} \mid m_i \in M_i\} = M$ and $\bigcup\{S_i^{\uparrow M} \mid S_i \in \mathbf{S}_{M_i}^+\} = \mathbf{S}_M^+$, and $\bigcup\{F(S_i)^{\uparrow M} \mid S_i \in \mathbf{S}_{M_i}^+\} = \{F(S) \mid S \in \mathbf{S}_M^+\}$.

(b) $\forall S_i, R_i \in \mathbf{S}_{M_i}^+ : S_i \not\doteq R_i \Rightarrow S_i^{\uparrow M} \cap R_i^{\uparrow M} = \emptyset$ ($\doteq$ modulo renaming of bound variables)

(c) For $S \in \mathbf{S}^+$, $S^{\uparrow M} = \{S\}$, e.g., $(\forall x P(x))^{\uparrow M} = \{\forall x P(x)\}$, so for $S \in \mathbf{S}^+$ we identify $S = S^{\uparrow M}$.

Setting (*) $L_i^{\uparrow M} = \bigcup\{S_i^{\uparrow M} \subseteq \mathbf{S}_M^+ \mid S_i \in L_i\}$, for each $i \in I$, we also have:

(d) $S_i \notin L_i \Leftrightarrow \forall X \in L_i : S_i \not\doteq X \overset{\textbf{(b)}}{\Rightarrow} \forall X \in L_i : S_i^{\uparrow M} \cap X^{\uparrow M} = \emptyset \Rightarrow S_i^{\uparrow M} \cap L_i^{\uparrow M} = \emptyset$. This implies: $S(m) \in L_i^{\uparrow M} \Rightarrow S(m_i) \in L_i$ and, as a special case

(e) $S \notin L_i \Rightarrow S^{\uparrow M} = S \notin L_i^{\uparrow M}$, for $S \in \mathbf{S}^+$.

By the main assumption $\Gamma \subseteq L_i$ for each i, hence $\Gamma \subset L_i^{\uparrow M}$. We show that $L_i^{\uparrow M} \in SK(G)$.

1. That $L_i^{\uparrow M}$ is absorbing, $\mathbf{E}_G(L_i^{\uparrow M}) \subseteq \mathbf{E}_G^-(L_i^{\uparrow M})$, follows by considering cases of its vertices:

1.i. If $P(m) \in L_i^{\uparrow M}$, then $\mathbf{E}_G(P(m)) = \{\overline{P(m)}\} \subseteq \mathbf{E}_G^-(P(m)) \subseteq \mathbf{E}_G^-(L_i^{\uparrow M})$.

1.ii. For a negated $\neg S \in L_i^{\uparrow M}$, we have some $\neg S_i \in L_i$ by (*), hence $S_i \in \mathbf{E}_{G_i}^-(L_i)$ since L_i is a semikernel, and show $S \in \mathbf{E}_G^-(L_i^{\uparrow M})$ for $\mathbf{E}_G(\neg S) = \{S\}$ by cases of S. Where relevant, we mark possible $n \in M$ occurring in the considered sentences as extra parameters.

1. Negated atom $\neg P(n) \in L_i^{\uparrow M} \Rightarrow \mathbf{E}_G(\neg P(n)) = \{P(n)\} \subset \mathbf{E}_G^-(\neg P(n)) \subset \mathbf{E}_G^-(L_i^{\uparrow M})$.
2. $\neg\neg S \in L_i^{\uparrow M} : \neg S_i \in \mathbf{E}_{G_i}^-(L_i) \Rightarrow S_i \in L_i \overset{(*)}{\Rightarrow} S \in L_i^{\uparrow M} \Rightarrow \neg S \in \mathbf{E}_G^-(L_i^{\uparrow M})$.
3. $\neg(S \wedge R) \in L_i^{\uparrow M} : S_i \wedge R_i \in \mathbf{E}_{G_i}^-(L_i) \Rightarrow \neg S_i \in L_i \vee \neg R_i \in L_i \overset{(*)}{\Rightarrow} \neg S \in L_i^{\uparrow M} \vee \neg R \in L_i^{\uparrow M} \Rightarrow S \wedge R \in \mathbf{E}_G^-(L_i^{\uparrow M})$.
4. $\neg\forall x F(x, n) \in L_i^{\uparrow M} : \forall x F(x, n_i) \in \mathbf{E}_{G_i}^-(L_i) \Rightarrow \exists m_i \in M_i : \neg F(m_i, n_i) \in L_i \overset{(*)}{\Rightarrow} \neg F(m_i^{\uparrow M}, n_i^{\uparrow M}) \subseteq L_i^{\uparrow M} \Rightarrow \neg F(m, n) \in L_i^{\uparrow M} \Rightarrow \forall x F(x, n) \in \mathbf{E}_G^-(L_i^{\uparrow M})$.

5. $\neg\forall\phi F(\phi,n) \in L_i^{\uparrow M} : \forall\phi F(\phi,n_i) \in \mathbf{E}^-_{G_i}(L_i) \Rightarrow \exists S_i \in \mathbf{S}^+_{M_i} : \neg F(S_i,n_i) \in L_i \stackrel{(*)}{\Rightarrow} \neg F(S_i^{\uparrow M}, n_i^{\uparrow M}) \subseteq L_i^{\uparrow M} \Rightarrow \neg F(S,n) \in L_i^{\uparrow M} \Rightarrow \forall\phi F(x,n) \in \mathbf{E}^-_G(L_i^{\uparrow M})$.

1.iii. $S_1 \wedge S_2 \in L_i^{\uparrow M} \stackrel{(*)}{\Rightarrow} (S_1 \wedge S_2)_i \in L_i \Rightarrow \{(\neg S_1)_i, (\neg S_2)_i\} \subset \mathbf{E}^-_{G_i}(L_i) \Rightarrow \{(S_1)_i, (S_2)_i\} \subseteq L_i \Rightarrow$

$\stackrel{(*)}{\Rightarrow} (S_1)_i^{\uparrow M} \cup (S_2)_i^{\uparrow M} \subset L_i^{\uparrow M} \Rightarrow \{S_1, S_2\} \subset L_i^{\uparrow M} \Rightarrow \{\neg S_1, \neg S_2\} \subset \mathbf{E}^-_G(L_i^{\uparrow M})$.

1.iv. $\forall x F(x,n) \in L_i^{\uparrow M} \stackrel{(*)}{\Rightarrow} \forall x F(x,n_i) \in L_i \Leftrightarrow \{\neg F(m_i,n_i) \mid m_i \in M_i\} \subseteq \mathbf{E}^-_{G_i}(L_i) \Leftrightarrow \{F(m_i,n_i) \mid m_i \in M_i\} \subseteq L_i$

$\stackrel{(*)}{\Rightarrow} \bigcup\{(F(m_i^{\uparrow M}, n_i^{\uparrow M})) \mid m_i \in M_i\} \stackrel{\mathbf{(a)}}{=} \{F(m, n_i^{\uparrow M}) \mid m \in M\} \subseteq L_i^{\uparrow M} \Rightarrow$

$\Rightarrow \{F(m,n) \mid m \in M\} \subseteq L_i^{\uparrow M} \Rightarrow \mathbf{E}_G(\forall x F(x,n)) = \{\neg F(m,n) \mid m \in M\} \subseteq \mathbf{E}^-_G(L_i^{\uparrow M})$

1.v. $\forall\phi F(\phi,n) \in L_i^{\uparrow M} \stackrel{(*)}{\Rightarrow} \forall\phi F(\phi,n_i) \in L_i$

$\Leftrightarrow \{\neg F(S_i,n_i) \mid S_i \in \mathbf{S}^+_{M_i}\} \subseteq \mathbf{E}^-_{G_i}(L_i) \Leftrightarrow \{F(S_i,n_i) \mid S_i \in \mathbf{S}^+_{M_i}\} \subseteq L_i$

$\stackrel{(*)}{\Rightarrow} \bigcup\{F(S_i^{\uparrow M}, n_i^{\uparrow M}) \mid S_i \in \mathbf{S}^+_{M_i}\} = \{F(S, n_i^{\uparrow M}) \mid S \in \mathbf{S}^+_M\} \subseteq L_i^{\uparrow M}$

$\Rightarrow \{F(S,n) \mid S \in \mathbf{S}^+_M\} \subseteq L_i^{\uparrow M} \Rightarrow \mathbf{E}_G(\forall\phi F(\phi,n)) = \{\neg F(S,n) \mid S \in \mathbf{S}^+_M\} \subseteq \mathbf{E}^-_G(L_i^{\uparrow M})$

2. Also independence of $L_i^{\uparrow M}$, $\mathbf{E}_G(L_i^{\uparrow M}) \subseteq \mathbf{V}_G \setminus L_i^{\uparrow M}$, follows by considering its possible vertices:

2.i. $P(m) \in (L_i)^{\uparrow M} \stackrel{(d)}{\Rightarrow} P(m_i) \in L_i \Rightarrow \overline{P(m_i)} \notin L_i \stackrel{(d)}{\Rightarrow} \overline{P(m)} \notin L_i^{\uparrow M}$

2.ii. $\neg S(m) \in L_i^{\uparrow M} \stackrel{(d)}{\Rightarrow} \neg S_i(m_i) \in L_i \Rightarrow S_i(m_i) \in \mathbf{E}^-_{G_i}(L_i) \Rightarrow S_i(m_i) \notin L_i$

$\stackrel{(d)}{\Rightarrow} S_i(m_i)^{\uparrow M} \cap L_i^{\uparrow M} = \emptyset \stackrel{S(m) \in S(m_i)^{\uparrow M}}{\Longrightarrow} S(m) \in \mathbf{V}_G \setminus L_i^{\uparrow M}$.

2.iii. $(S^1(m) \wedge S^2(n)) \in L_i^{\uparrow M} \stackrel{(d)}{\Rightarrow} (S^1(m_i) \wedge S^2(n_i)) \in L_i \Rightarrow$

$\Rightarrow \{\neg S^1(m_i), \neg S^2(n_i)\} \subset \mathbf{E}^-_{G_i}(L_i)$

$\Rightarrow \{\neg S^1(m_i), \neg S^2(n_i)\} \cap L_i = \emptyset$

$\Rightarrow (\neg S^1(m_i^{\uparrow M}) \cup \neg S^2(n_i^{\uparrow M})) \cap L_i^{\uparrow M} = \emptyset \Rightarrow \mathbf{E}_G(S^1(m) \wedge S^2(n)) \cap L_i^{\uparrow M} = \emptyset$

The last implication holds by $(\neg S^1(m_i^{\uparrow M}) \cup \neg S^2(n_i^{\uparrow M})) \supset \{\neg S^1(m), \neg S^2(n)\} = \mathbf{E}_G(S^1(m) \wedge S^2(n))$.

2.iv. $\forall x F(x,n) \in L_i^{\uparrow M} \stackrel{\mathbf{(d)}}{\Rightarrow} \forall x F(x,n_i) \in L_i \Rightarrow$

$\Rightarrow \{\neg F(m_j,n_i) \mid m_j \in M_i\} \cap L_i = \emptyset$

$\stackrel{\mathbf{(d)}}{\Rightarrow} \bigcup\{\neg F(m_j^{\uparrow M}, n_i^{\uparrow M}) \mid m_j \in M_i\} \cap L_i^{\uparrow M} = \emptyset$

$\stackrel{\mathbf{(a)}}{\Rightarrow} \bigcup\{(\neg F(m, n_i^{\uparrow M})) \mid m \in M\} \cap L_i^{\uparrow M} = \emptyset$

$\stackrel{n \in n_i^{\uparrow M}}{\Longrightarrow} \{\neg F(m,n) \mid m \in M\} \cap L_i^{\uparrow M} = \emptyset \Rightarrow \mathbf{E}_G(\forall x F(x,n)) \cap L_i^{\uparrow M} = \emptyset$

2.v. $\forall\phi F(\phi,n) \in L_i^{\uparrow M} \stackrel{\mathbf{(d)}}{\Rightarrow} \forall\phi F(\phi,n_i) \in L_i \Rightarrow$

$\Rightarrow \{\neg F(S_j, n_i) \mid S_j \in \mathbf{S}^+_{M_i}\} \subseteq \mathbf{E}^-_{G_i}(L_i)$
$\Rightarrow \{\neg F(S_j, n_i) \mid S_j \in \mathbf{S}^+_{M_i}\} \cap L_i = \emptyset$
$\overset{\textbf{(d)}}{\Rightarrow} \bigcup\{(\neg F(S_j, n_i))^{\uparrow M} \subset \mathbf{S}^+_M \mid S_j \in \mathbf{S}^+_{M_i}\} \cap L_i^{\uparrow M} = \emptyset$
$\overset{\textbf{(a)}}{\Rightarrow} \bigcup\{\neg F(S, n) \subset \mathbf{S}^+_M \mid S \in \mathbf{S}^+_M\} \cap L_i^{\uparrow M} = \emptyset \Rightarrow \mathbf{E}_G(\forall\phi F(\phi, n)) \cap L_i^{\uparrow M} = \emptyset.$

Thus, for every L_i, its lifting $L_i^{\uparrow M}$ is a semikernel of G (by Fact 7.1), containing Γ and covering $\bot_i$. To use Fact 7.18, we only have to show that every finite $Q \Subset \mathbf{S}^+_M$ is covered by some $L_i^{\uparrow M}$.

3. For an arbitrary $S \in \mathbf{S}^+$, $\bot_{\{S\}} = S \wedge \neg S$ is covered by $L_{\{S\}}$ which does not contain it. Hence $(S \wedge \neg S) \in \mathbf{E}^-_{G_{\{S\}}}(L_{\{S\}})$, so either $\neg S \in L_{\{S\}}$ or $\neg\neg S \in L_{\{S\}}$. In the former case $S \in \mathbf{E}_{G_{\{S\}}}(L_{\{S\}}) \subseteq \mathbf{E}^-_{G_{\{S\}}}(L_{\{S\}})$ while in the latter $S \in L_{\{S\}}$. In each case, $L_{\{S\}}^{\uparrow M}$ covers $S = S^{\uparrow M}$.

For an arbitrary $S \in \mathbf{S}^+_M \setminus \mathbf{S}^+$, i.e., $S = Fm$, where $m \in M$ are all M elements occurring in S, contradiction $N(F) = \forall x \neg(Fx \wedge \neg Fx) \wedge \neg\forall x \neg(Fx \wedge \neg Fx)$ has, in some graph G_i, a countermodel L_i covering it by $N(F) \in \mathbf{E}^-_{G_i}(L_i)$. Thus either $\neg\forall x \neg(Fx \wedge \neg Fx) \in L_i$ or $\neg\neg\forall x \neg(Fx \wedge \neg Fx) \in L_i$, but since the former sentence is a contradiction, the latter is the case. Then $\forall x \neg(Fx \wedge \neg Fx) \in L_i$, hence also $\neg(Fm_i \wedge \neg Fm_i) \in L_i$ for every $m_i \in M_i$. Since then $(Fm_i \wedge \neg Fm_i) \in \mathbf{E}^-_{G_i}(L_i)$, for every $m_i \in M_i$, either $\neg\neg Fm_i \in L_i$ or $\neg Fm_i \in L_i$. In the former case $Fm_i \in L_i$, while in the latter, $Fm_i \in \mathbf{E}^-_{G_i}(L_i)$. Thus, L_i covers Fm_i, containing either Fm_i or $\neg Fm_i$, for every $m_i \in M_i$. By **(c)**, $Fm \in L_i^{\uparrow M}$ or $\neg Fm \in L_i^{\uparrow M}$, which means that $L_i^{\uparrow M}$ covers Fm, for every $m \in M$.

These arguments for single sentences are extended to an arbitrary $Q = \{F^1m^1, ..., F^km^k\} \Subset \mathbf{S}^+_M$, by considering semikernel L_i which covers, without containing, $\bot_Q = N(F^1) \vee ... \vee N(F^k)$, and existing by the assumption $\Gamma \nvdash \bot_Q$. The subgraph of $\bot_Q$ has the form

$$
\begin{array}{ccc}
 & \bot_Q & \\
 & \downarrow & \\
 \swarrow & \bullet & \searrow \\
 & \vdots & \\
N(F^1) & \cdots & N(F^k)
\end{array}
$$

hence $\bot_Q \in \mathbf{E}^-_{G_Q}(L_Q)$ implies $\bullet \in L_Q$, so $N(F^j) \in \mathbf{E}^-_{G_Q}(L_Q)$, for $1 \leq j \leq k$. By the argument for a single $N(F)$, this implies that L_Q covers every F^jm^j. □

If LSO does not derive from Γ contradiction $\bot_Q = \bigvee_{S \in Q}(S \wedge \neg S)$, for any $Q \Subset \mathbf{S}^+$, there is thus a graph with a kernel containing Γ. Soundness, Theorem 4.3, implies then $\Gamma \nvdash_c \bot$ for each $\bot \in \mathbb{C}$. Thus, if LSO^c proves a contradiction from a theory Γ, then so does LSO.

Theorem 7.20 (4.4) *For Γ, Δ over a countable $\mathcal{L}^+ \in FOL^+$,*
$(\forall Q \Subset \mathbf{S}^+ : \Gamma \nvdash \bot_Q) \Rightarrow (\forall \bot \in \mathbb{C} : \Gamma \nvdash_c \bot).$

8 Appendix: Soundness and Completeness

Facts 8.1 and 8.3 below show soundness and completeness of LSO for the semikernel semantics from (3.11), establishing Theorem 4.1. Theorem 4.3 shows these properties for LSO with (cut) for the kernel semantics (3.13).

Fact 8.1 (4.1) *The rules of* LSO *are sound and invertible for (3.11).*

Proof Given an arbitrary language graph G^+ (over an arbitrary domain M), soundness for each rule follows by showing that any semikernel L covering the conclusion satisfies it, assuming validity of the premise(s), while invertibility by showing that any semikernel L covering (each) premise satisfies it, assuming validity of the rule's conclusion.

1. $(\wedge_R)$. For soundness, assume $\Gamma \models \Delta, A_1$ and $\Gamma \models \Delta, A_2$, and let semikernel L cover the rule's conclusion, under a given $\alpha \in M^{\mathcal{V}(\Gamma,\Delta)}$. Assume that $\alpha(\Gamma) \subseteq L$, $\alpha(\Delta) \subseteq \mathbf{E}^-(L)$ and $\alpha(A_1 \wedge A_2) \in \mathbf{E}^-(L)$ – if not, then $L \models_\alpha \Gamma \Rightarrow \Delta, A_1 \wedge A_2$, as desired. Since $\alpha(A_1 \wedge A_2) \in \mathbf{E}^-(L)$ and $\mathbf{E}(\alpha(A_1 \wedge A_2)) = \{\neg\alpha(A_1), \neg\alpha(A_2)\}$ so, for some $i \in \{1, 2\}$, $\neg\alpha(A_i) \in L$, and then $\alpha(A_i) \in \mathbf{E}^-(L)$, contradicting the assumption $\Gamma \models \Delta, A_i$.
   For invertibility, let $\Gamma \models \Delta, A_1 \wedge A_2$ and L cover A_1 (or A_2) under α. If (*) $\alpha(\Gamma) \subseteq L$ and $\alpha(\Delta \cup \{A_1\}) \subseteq \mathbf{E}^-(L)$, then $L' = L \cup \{\neg\alpha(A_1)\}$ is a semikernel, since $\mathbf{E}(\neg\alpha(A_1)) = \{\alpha(A_1)\} \subseteq \mathbf{E}^-(L) \subseteq \mathbf{E}^-(L')$. L' covers also $\alpha(A_1 \wedge A_2) \in \mathbf{E}^-(\neg\alpha(A_1))$. Thus L' covers the conclusion, while $\alpha(\Gamma) \cap \mathbf{E}^-(L') = \emptyset$ and $\alpha(\Delta \cup \{A_1 \wedge A_2\}) \cap L' = \emptyset$, so $L' \not\models \Gamma \Rightarrow \Delta, A_1 \wedge A_2$, contrary to $\Gamma \models \Delta, A_1 \wedge A_2$. Hence (*) fails, so $\alpha(\Gamma) \cap \mathbf{E}^-(L) \neq \emptyset$ or $\alpha(\Delta \cup \{A_1\}) \cap L \neq \emptyset$, yielding the claim.
   Assignments to free FOL-variables do not affect the argument, so covering by L below is to be taken relatively to a given α, which we do not mention, except for $(\forall_R)$.
2. $(\wedge_L)$. For soundness, assume $\Gamma, A_1, A_2 \models \Delta$, let semikernel L cover the rule's conclusion, $\Gamma \subseteq L$ and $\Delta \subseteq \mathbf{E}^-(L)$. If $A_1 \wedge A_2 \in L$, then $\mathbf{E}(A_1 \wedge A_2) = \{\neg A_1, \neg A_2\} \subseteq \mathbf{E}^-(L)$, so $\mathbf{E}(\{\neg A_1, \neg A_2\}) = \{A_1, A_2\} \subseteq L$, contradicting $\Gamma, A_1, A_2 \models \Delta$. Thus $A_1 \wedge A_2 \in \mathbf{E}^-(L)$ and $L \models \Gamma, A_1 \wedge A_2 \Rightarrow \Delta$.
   For invertibility, assume $\Gamma, A_1 \wedge A_2 \models \Delta$, let semikernel L cover the rule's premise, and assume $\Gamma \subseteq L$ and $\Delta \subseteq \mathbf{E}^-(L)$. If $A_1, A_2 \in L$, which is the only way L can contradict $\Gamma, A_1, A_2 \models \Delta$, then $\{\neg A_1, \neg A_2\} \subseteq \mathbf{E}^-(L)$, and $L' = L \cup \{A_1 \wedge A_2\}$ is also a semikernel:
   $$\mathbf{E}(L') = \mathbf{E}(L \cup \{A_1 \wedge A_2\}) = \mathbf{E}(L) \cup \mathbf{E}(\{A_1 \wedge A_2\}) \subseteq \mathbf{E}^-(L) \cup \{\neg A_1, \neg A_2\} \subseteq$$
   $$\subseteq \mathbf{E}^-(L) \subseteq \mathbf{V} \setminus (L \cup \{A_1 \wedge A_2\}).$$
   The last inclusion follows because $\mathbf{E}^-(L) \subseteq \mathbf{V} \setminus L$ and $A_1 \wedge A_2 \notin \mathbf{E}^-[L]$, since $A_1 \wedge A_2 \in L$ contradicts $\Gamma, A_1 \wedge A_2 \models \Delta$ (as $\Gamma \subseteq L$ and $\Delta \subseteq \mathbf{E}^-(L)$), while $A_1 \wedge A_2 \in \mathbf{E}^-(L)$ contradicts independence of L, implying $\neg A_i \in L$ (for $i = 1$ or $i = 2$), while $\neg A_i \in \mathbf{E}^-(L)$ since $A_i \in L$.
   Since $L' \not\models \Gamma, A_1 \wedge A_2 \Rightarrow \Delta$ contradicts the assumption, either $A_1 \notin L$ or $A_2 \notin L$, and $L \models \Gamma, A_1 \wedge A_2 \Rightarrow \Delta$ as desired.
3. $(\neg_R)$. For soundness, assume $\Gamma, A \models \Delta$, let semikernel L cover the rule's conclusion, and assume $\Gamma \subseteq L$ and $\Delta \subseteq \mathbf{E}^-(L)$. If $\neg A \in L$, we are done, while

if $\neg A \in \mathbf{E}^-(L)$ then $A \in L$, which contradicts the assumption, since now $\Gamma \cup \{A\} \subseteq L$ and $\Delta \subseteq \mathbf{E}^-(L)$.

For invertibility, assuming $\Gamma \models \Delta, \neg A$, let L cover the rule's premise, $\Gamma \subseteq L$ and $\Delta \subseteq \mathbf{E}^-(L)$. If $A \in L$ then $\neg A \in \mathbf{E}^-(L)$ and $L \not\models \Gamma \Rightarrow \Delta, \neg A$, contradicting the assumption. Hence $A \in \mathbf{E}^-(L)$, as required for $L \models \Gamma, A \Rightarrow \Delta$.

4. $(\neg_L)$. For soundness, assume $\Gamma \models \Delta, A$, let L cover the rule's conclusion, $\Gamma \subseteq L$ and $\Delta \subseteq \mathbf{E}^-(L)$. If $\neg A \in \mathbf{E}^-(L)$, we are done, while if $\neg A \in L$ then $A \in \mathbf{E}(\neg A) \subseteq \mathbf{E}^-(L)$, contradicting the assumption, since now $\Gamma \cup \{A\} \subseteq L$ and $(\Delta \cup \{A\}) \subseteq \mathbf{E}^-(L)$.

   For invertibility, assume $\Gamma, \neg A \models \Delta$, let L cover the rule's premise, $\Gamma \subseteq L$ and $\Delta \subseteq \mathbf{E}^-(L)$. If $A \in \mathbf{E}^-(L)$ then $L' = L \cup \{\neg A\}$ is a semikernel, because L is and $\mathbf{E}(\neg A) = \{A\} \subseteq \mathbf{E}^-(L)$. But L' contradicts the assumption, so $A \in L$, as required for $L \models \Gamma \Rightarrow \Delta, A$.

5. $(\forall_L)$. For soundness, assume $F(t), \Gamma, \forall x F(x) \models \Delta$ and let L cover the rule's conclusion. If $\forall x F(x) \notin L$, i.e., $\forall x F(x) \in \mathbf{E}^-(L)$, then $(\Gamma \cup \{\forall x F(x)\}) \cap \mathbf{E}^-(L) \neq \emptyset$, so $L \models \Gamma, \forall x F(x) \Rightarrow \Delta$. If $\forall x F(x) \in L$ then also $F(t) \in L$, since $\neg F(t) \in \mathbf{E}(\forall x F(x)) \subseteq \mathbf{E}^-(L)$ and $\mathbf{E}(\neg F(t)) = \{F(t)\}$. As L covers the premise, either $\Gamma \cap \mathbf{E}^-(L) \neq \emptyset$, since $F(t) \notin \mathbf{E}^-(L)$, or $\Delta \cap L \neq \emptyset$. Either case yields the claim for L, so $\Gamma, \forall x F(x) \models \Delta$.

   For invertibility, assuming $\Gamma, \forall x F(x) \models \Delta$, any L covering the premise of the rule covers also its conclusion, yielding the claim.

6. $(\forall_R)$. For soundness, let (*) $\Gamma \models \Delta, F(y)$, with eigenvariable y, and L cover the rule's conclusion, under a given assignment α to $\mathcal{V}(\Gamma, \Delta, \forall x F(x)) \not\ni y$. Assume also $\alpha(\Gamma) \subseteq L$ and $\alpha(\Delta) \subseteq \mathbf{E}^-(L)$. If $\alpha(\forall x F(x)) \notin L$ then $\alpha(\forall x F(x)) \in \mathbf{E}^-(L)$ and some $\alpha(\neg F(m)) \in L$, since $\mathbf{E}(\alpha(\forall x F(x))) = \{\alpha(\neg F(m)) \mid m \in M\}$. Extending α with $\alpha(y) = m$, we obtain $L \not\models_\alpha \Gamma \Rightarrow \Delta, F(y)$, contrary to (*). Thus, $\alpha(\forall x F(x)) \in L$ and $L \models_\alpha \Gamma \Rightarrow \Delta, \forall x F(x)$.

   For invertibility, if $L \not\models_\alpha \Gamma \Rightarrow \Delta, F(y)$, for $\alpha(y) = m$, i.e., $\alpha(\Gamma) \subseteq L$, $\alpha(\Delta) \subseteq \mathbf{E}^-(L)$ and $\alpha(F(m)) \in \mathbf{E}^-(L)$, then $L' = L \cup \{\alpha(\neg F(m))\}$ is a semikernel, because L is and $\mathbf{E}(\alpha(\neg F(m))) = \{\alpha(F(m))\} \subseteq \mathbf{E}^-(L) \subseteq \mathbf{E}^-(L')$. L' covers the conclusion since $\alpha(\forall x F(x)) \in \mathbf{E}^-(\alpha(\neg F(m)))$, but $L' \not\models_\alpha \Gamma \Rightarrow \Delta, \forall x F(x)$.

7. $(\forall_L^+)$. The argument repeats that for $(\forall_L)$. For soundness, assume $\Gamma, F(S), \forall \phi F(\phi) \models \Delta$ and let L cover the rule's conclusion. If $\forall \phi F(\phi) \notin L$ then $\forall \phi F(\phi) \in \mathbf{E}^-(L)$, yielding $L \models \Gamma, \forall \phi F(\phi) \Rightarrow \Delta$. If $\forall \phi F(\phi) \in L$ then also $F(S) \in L$, since $\neg F(S) \in \mathbf{E}(\forall \phi F(\phi)) \subseteq \mathbf{E}^-(L)$ and $\mathbf{E}(\neg F(S)) = \{F(S)\}$. Thus L covers also the premise, hence, either $\Gamma \cap \mathbf{E}^-(L) \neq \emptyset$, since $F(S) \notin \mathbf{E}^-(L)$, or $\Delta \cap L \neq \emptyset$. Either case yields the claim for L, so $\Gamma, \forall \phi F(\phi) \models \Delta$.

   For invertibility, assuming $\Gamma, \forall \phi F(\phi) \models \Delta$, any L covering the premise of the rule covers also its conclusion, yielding the claim.

8. $(\forall_R^+)$. For soundness, let $\Gamma \models \Delta, F(S)$ for every S, and L cover the rule's conclusion. If $\forall \phi F(\phi) \in L$ then L satisfies the rule's conclusion. If $\forall \phi F(\phi) \notin L$ then $\forall \phi F(\phi) \in \mathbf{E}^-(L)$ and some $\neg F(S) \in L$, since $\mathbf{E}(\forall \phi F(\phi)) = \{\neg F(S) \mid S \in \mathbf{S}^+\}$. Now L covers also the premise $\Gamma \Rightarrow \Delta, F(S)$ and $F(S) \notin L$, hence either $\Gamma \cap \mathbf{E}^-(L) \neq \emptyset$ or $\Delta \cap L \neq \emptyset$. Each case yields the claim that L satisfies the conclusion.

For invertibility, assume $\Gamma \models \Delta, \forall\phi F(\phi)$, and let L cover a premise $\Gamma \Rightarrow \Delta, F(S)$. If (*) $\Gamma \subseteq L$ and $\Delta \cup \{F(S)\} \subseteq \mathbf{E}^-(L)$ then $L' = L \cup \{\neg F(S)\}$ is a semikernel, since $\mathbf{E}(\neg F(S)) = \{F(S)\} \subseteq \mathbf{E}^-(L) \subseteq \mathbf{E}^-(L')$. L' covers also $\forall\phi F(\phi)$, since $\forall\phi F(\phi) \in \mathbf{E}^-(\neg F(Q))$, for every $Q \in \mathbf{S}^+$, in particular, $\forall\phi F(\phi) \in \mathbf{E}^-(\neg F(S))$. Thus $\Gamma \cup \Delta \cup \{F(S), \forall\phi F(\phi)\} \subseteq \mathbf{E}^-[L']$, while $\Gamma \cap \mathbf{E}^-(L') = \emptyset$, $\Delta \cap L' = \emptyset$ and $\forall\phi F(\phi) \notin L'$, contrary to $\Gamma \models \Delta, \forall\phi F(\phi)$. Hence (*) fails, so either $\Gamma \cap \mathbf{E}^-(L) \neq \emptyset$ or $(\Delta \cup \{F(S)\}) \cap L \neq \emptyset$, yielding the claim.

9. S-equality rules are sound and invertible, because atoms occurring in the premises but not in the conclusions are redundant – under the intended intepretation of $\doteq$. E.g., $S \doteq S$ in the premise of (ref) is contained in every covering semikernel (for the intended intepretation of $\doteq$), hence satisfaction of the premise implies satisfaction of the conclusion. Conversely, satisfaction of the conclusion by any semikernel allows to extend the conclusion with a vertex, like $S \doteq S$, that either already is in the semikernel (reflecting the intended interpretation of $\doteq$), or is an atom uncovered by it. An analogous argument works for (rep) and (neq). □

The following simple consequence of Definition 3.7 is used in the completeness proof below.

Fact 8.2 *In any graph G^+_M, the forms of a nonatomic sentence $X \in \mathbf{S}^+_M$ and of its in- and out-neighbours are related as follows:*

1. $\mathbf{E}^-(X) = \{\neg X\}$ *– when X does not start with $\neg$,*
2. $\mathbf{E}^-(\neg X) = \{\neg\neg X\} \cup \{X \wedge S \mid S \in \mathbf{S}^+_M\} \cup \{\forall(\phi.D(\phi) \mid \exists S \in \mathbf{S}^+_M : D(S) = X\} \cup \ldots$
$\ldots \cup \{\forall x.D(x) \mid \exists t \in \mathbf{T}_M : D(t) = X\}$
3. *when X does not start with $\neg$, then each out-neighbour of X does,*
4. $\mathbf{E}(\neg X) = \{X\}$.

For an atomic X: $\mathbf{E}^-(X) = \{\neg X\} = \mathbf{E}(X)$ and $\mathbf{E}^-(\neg X) = \{X\} = \mathbf{E}(\neg X)$.

The proof of completeness can apply the standard techniques because proofs in LSO, even if infinite, are well-founded trees with axioms as leaves. A few adjustments are required to handle deviations from LK. One is the infinitary rule ($\forall^+_R$), which is needed because substitution of fresh eigenvariables for s-variables, although sound, does not necessarily lead to a countermodel in an unsuccessful derivation, since s-variables are not sentences. (Replacing ($\forall^+_R$) by a usual $\forall_R$-rule using eigen-variables would yield a complete system for a modified notion of $\models$, admitting expansions of the language with new s-constants.) In the proof, we ensure that not only all formulas are processed and all terms are substituted by ($\forall_L$), but also all sentences are substituted by ($\forall^+_L$). The absence of the subformula property, due to the substitution of all sentences for s-variables, is addressed by retaining the principal formula from the conclusion in all its premises in a bottom-up construction of the derivation tree. A specific case of violating this property is a cyclic branch, with the same sequent appearing repeatedly. Any nonaxiomatic (e.g., cyclic) branch provides a countermodel.

Fact 8.3 (4.1) *For Γ, Δ over a countable $\mathcal{L}^+ \in FOL^+$, $\Gamma \nvdash \Delta \Rightarrow \exists G^+ \exists L \in SK(G^+) : L \not\models \Gamma \Rightarrow \Delta$.*

Proof We fix an enumeration E^+ of all sentences $\mathbf{S}^+$ so that each occurs infinitely often, an enumeration $E_T = t_1, t_2...$ of terms $\mathbf{T}_X$ so that each occurs infinitely often, and an enumeration $E_S = S_1, S_2...$ of FOL$^+$ formulas without free s-variables and with operators applied only to sentences $\mathbf{S}^+$, so that each occurs infinitely often. (FOL variables, requiring special care, are treated in the standard way and ignored below; we keep also an enumeration of eigenvariables). We enumerate all triples $\langle S_i, t_j, S_k\rangle \in E_S \times E_T \times E^+$, with each $\langle S_i, t_j, \_\rangle$ and $\langle S_i, \_, S_k\rangle$ occurring infinitely often. This is interleaved with an enumeration of all pairs $E_S \times E_S$, with each pair occurring infinitely often.

1. We construct a derivation tree, starting with the root $\Gamma \vdash \Delta$, which is to be proven. An *active* sequent – initially, only the root – is a nonaxiomatic leaf of the tree constructed bottom-up so far. We proceed along the enumeration of the triples and pairs considering, for each $\langle S_i, t_j, S_k\rangle$, the cases of active occurrences (in the active sequents) of S_i. Pairs $\langle S_i, S_j\rangle$ are considered in point **v**.

i. If $S_i \in \mathbf{A}^+$, or S_i has no active occurrences, proceed to the next triple.

ii. Otherwise, proceed retaining S_i from the active sequent, which instantiates the conclusion of the relevant rule, in the new leaves obtained from the rule's premises. For instance, if $S_i = A \wedge B$ then every active sequent of the form $\Gamma', A \wedge B, \Gamma'' \vdash \Delta$ is replaced by

$$\frac{A, B, \Gamma', A \wedge B, \Gamma'' \vdash \Delta}{\Gamma', A \wedge B, \Gamma'' \vdash \Delta}$$

while every active sequent of the form $\Gamma \vdash \Delta', A \wedge B, \Delta''$ by

$$\frac{\Gamma \vdash A, \Delta', A \wedge B, \Delta'' \qquad \Gamma \vdash B, \Delta', A \wedge B, \Delta''}{\Gamma \vdash \Delta', A \wedge B, \Delta''}$$

In the same way, for other connectives.

iii. If $S_i = \forall x D(x)$, every active sequent of the form $\Gamma', \forall x D(x), \Gamma'' \vdash \Delta$, is replaced by the derivation with a new leaf adding $D(t_j)$ to its antecedent

$$\frac{D(t_j), \Gamma', \forall x D(x), \Gamma'' \vdash \Delta}{\Gamma', \forall x D(x), \Gamma'' \vdash \Delta}.$$

Every active sequent of the form $\Gamma \vdash \Delta', \forall x D(x), \Delta''$ is replaced by

$$\frac{\Gamma \vdash D(c), \Delta', \forall x D(x), \Delta''}{\Gamma \vdash \Delta', \forall x D(x), \Delta''}$$

where c is a fresh eigenvariable c.

iv. If $S_i = \forall\phi D(\phi)$ then replace every active sequent of the form $\Gamma', \forall\phi D(\phi), \Gamma'' \vdash \Delta$ by

$$\frac{D(S_k), \Gamma', \forall\phi D(\phi), \Gamma'' \vdash \Delta}{\Gamma', \forall\phi D(\phi), \Gamma'' \vdash \Delta}$$

while every active sequent of the form $\Gamma \vdash \Delta', \forall\phi D(\phi), \Delta''$ by the infinitely branching derivation with a new leaf sequent for each $S_n \in \mathbf{S}^+$:

$$\frac{\Gamma \vdash D(S_1), \Delta', \forall\phi D(\phi), \Delta'' \qquad \Gamma \vdash D(S_2), \Delta', \forall\phi D(\phi), \Delta'' \qquad \text{... – for all } S_n \in \mathbf{S}^+}{\Gamma \vdash \Delta', \forall\phi D(\phi), \Delta''}$$

v. For a pair $\langle S_i, S_j\rangle$, we apply rules for $\doteq$. If $S_i \not\doteq S_j$, we add atom $S_i \doteq S_j$ to the consequent of every active sequent. Otherwise, we add it to the antecedent. Finally, for each active sequent containing $S_i \doteq Q$ in its antecedent, along with any formula $A(S_i)$, we add to it $A(Q)$.

2. A branch gets closed when its leaf is an axiom, and the tree is obtained as the ω-limit of this process. If all branches are closed (finite), the derivation yields a proof of the root.
Otherwise, an infinite branch allows us to construct a countermodel of all sequents on this branch, including the root sequent. (Such an infinite branch can represent a finite process of derivation terminating with a nonaxiomatic sequent, which remains unchanged in an infinite tail of the branch. It can also be cyclic. These special cases are treated uniformly with an infinite branch without any repeated sequents.)

3. The claim is that if β is an infinite branch, with β'_L/β'_R all formulas occurring in β on the left/right of $\vdash$, then there is a language graph G with a semikernel L' such that $\beta'_L \subseteq L'$ and $\beta'_R \subseteq \mathbf{E}^-(L')$. The rest of the proof establishes this claim.
Absence of any axiom in β implies that $\beta'_L \cap \beta'_R = \emptyset$, which is often applied implicitly. $\beta'_L = \beta_L \cup Eq_L$, where Eq_L are $\doteq$-atoms $S \doteq S$ occurring on the left. $\beta'_R = \beta_R \cup Eq_R$, where Eq_R are $\doteq$-atoms occurring on the right, with $\overline{Eq}_R$ denoting the set of their negations.
If β contains any FOL-atoms, construct first a FOL structure M, giving a countermodel to $(\beta_L \cap \mathbf{S}_M) \Rightarrow (\beta_R \cap \mathbf{S}_M)$, in the standard way. Otherwise, set $M = \emptyset$. Let $G = G_M$ (when $M = \emptyset$, this is the graph for QBS). We show that (def) $L = \beta_L \cup (\mathbf{E}(\beta_R) \cap \mathbf{E}^-(\beta_R))$ is a semikernel of G, with $\beta_R \subseteq \mathbf{E}^-(L)$ and $\beta_L \subseteq L$. Then $L' = Eq_L \cup \overline{Eq}_R \cup L$ is a required semikernel of G.

4. First, $\doteq$-atoms can be treated separately. Since $\beta'_L \cap \beta'_R = \emptyset$, each $\doteq$-atom $A \in Eq_R$ has the form $S \doteq T$ for syntactically distinct sentences, while each such atom in Eq_L has the form $S \doteq S$. Any semikernel of G, in particular L, can be extended to semikernel $L' = L \cup Eq_L \cup \overline{Eq}_R$, as the added atoms have only 2-cycles to their dual literals in G, by Definition 3.7. Thus $\mathbf{E}(Eq_L \cup \overline{Eq}_R) = \emptyset$, while $Eq_R \subseteq \mathbf{E}^-(L') \cap (\mathbf{V} \setminus L')$.

5. To show $L \in SK(G)$, we show first $\beta_R \subseteq \mathbf{E}^-(L)$, which follows from definitions of L and G by considering the cases for $A \in \beta_R$. Use of Fact 8.2/Definition 3.7 is marked by superscript $\_^{8.2}$.

i. If $A \in \mathbf{A}^+$ then $\mathbf{E}(A) \overset{8.2}{=} \{\neg A\} \overset{8.2}{=} \mathbf{E}^-(A)$, so $\neg A \in L$ by (def) and $A \overset{8.2}{\in} \mathbf{E}^-(L)$.

ii. If $A = \neg C$ then $C \in \beta_L \subseteq L$, so $A \overset{8.2}{\in} \mathbf{E}^-(L)$.

iii. If $A = C \wedge D$ then $C \in \beta_R$ (or $D \in \beta_R$), so $\neg C \overset{8.2}{\in} \mathbf{E}^-(C) \cap \mathbf{E}(C \wedge D) \subseteq \mathbf{E}^-(\beta_R) \cap \mathbf{E}(\beta_R) \subseteq L$, and thus $A = C \wedge D \overset{8.2}{\in} \mathbf{E}^-(\neg C) \subseteq \mathbf{E}^-(L)$. (The case of $D \in \beta_R$ is analogous.)

iv. If $A = \forall x.D(x)$ then $D(c) \in \beta_R$, for some $c \in M$, so $\neg D(c) \overset{8.2}{\in} \mathbf{E}^-(D(c)) \cap \mathbf{E}(\forall x.D(x)) \subseteq L$, and $A \overset{8.2}{\in} \mathbf{E}^-(\neg D(c)) \subseteq \mathbf{E}^-(L)$.

v. If $A = \forall \phi.D(\phi)$ then $D(S) \in \beta_R$ for some $S \in \mathbf{S}^+$, so $\neg D(S) \overset{8.2}{\in} \mathbf{E}^-(D(S)) \cap \mathbf{E}(\forall \phi.D(\phi)) \subseteq L$, and $A \in \mathbf{E}^-(\neg D(S)) \subseteq \mathbf{E}^-(L)$.

6. We show $\mathbf{E}(L) \subseteq \mathbf{E}^-(L) \cap (\mathbf{V} \setminus L)$, partitioning $L = \beta_L \cup Z$, where $Z = (\mathbf{E}(\beta_R) \cap \mathbf{E}^-(\beta_R)) \setminus \beta_L$, and establish first $\mathbf{E}(\beta_L) \subseteq \mathbf{E}^-(L) \cap (\mathbf{V} \setminus L)$, considering cases of $A \in \beta_L$.

i. For atoms $A \in \mathbf{A}^+$, $A \in \beta_L \subseteq L$ and $A \notin \beta_R$ imply $\neg A \notin \beta_L$ and, since $\mathbf{E}(\neg A) \overset{8.2}{=} \{A\}$, $\neg A \notin \mathbf{E}^-(\beta_R)$. Thus $\mathbf{E}(A) \overset{8.2}{=} \{\neg A\} \subseteq \mathbf{E}^-(A) \cap \mathbf{V} \setminus L \subseteq \mathbf{E}^-(L) \cap \mathbf{V} \setminus L$.

ii. $A = \neg C \in \beta_L$ implies $C \in \beta_R$, so $\mathbf{E}(A) \overset{8.2}{=} \{C\} \subseteq \beta_R \subseteq \mathbf{E}^-(L)$ by **5**.
We show $\mathbf{E}(A) \subset \mathbf{V} \setminus L$. $C \notin \beta_L$ since $\beta_L \cap \beta_R = \emptyset$. Suppose $C \in \mathbf{E}(\beta_R) \cap \mathbf{E}^-(\beta_R)$. If $C = \neg D$ then $\neg D \in \mathbf{E}^-(\beta_R)$, i.e., $\mathbf{E}(\neg D) \overset{8.2}{=} \{D\} \subset \beta_R$, while $A = \neg C = \neg\neg D \in \beta_L$ implies also $\neg D \in \beta_R$ and $D \in \beta_L$, contradicting $\beta_L \cap \beta_R = \emptyset$.
Otherwise, i.e., if C does not start with $\neg$, then for any $F \in \beta_R$ for which $C \in \mathbf{E}(F)$, Fact 8.2.(3-4) forces $F = \neg C = A$, contradicting $\beta_R \cap \beta_L = \emptyset$.

iii. $A = B \wedge C \in \beta_L$ implies $\{B, C\} \subset \beta_L$ and $\{\neg B, \neg C\} \cap \beta_L = \emptyset$, so $\mathbf{E}(B \wedge C) \overset{8.2}{=} \{\neg B, \neg C\} \subseteq \mathbf{V} \setminus \beta_L$ and $\mathbf{E}(B \wedge C) = \{\neg B, \neg C\} \overset{8.2}{\subseteq} \mathbf{E}^-(\{B, C\}) \subseteq \mathbf{E}^-(\beta_L)$. If, say, $\neg B \in \mathbf{E}^-(\beta_R)$, then $B \in \beta_R$ would contradict $\beta_L \cap \beta_R = \emptyset$. The same if $\neg C \in \mathbf{E}^-(\beta_R)$. Thus, $\mathbf{E}(B \wedge C) \subseteq \mathbf{E}^-(L) \cap \mathbf{V} \setminus L$.

iv. $A = \forall\phi D(\phi) \in \beta_L \Rightarrow \{D(S) \mid S \in \mathbf{S}^+\} \subseteq \beta_L$, so $\mathbf{E}(\forall\phi D(\phi)) \overset{8.2}{=} \{\neg D(S) \mid S \in \mathbf{S}^+\} \overset{8.2}{\subseteq} \mathbf{E}^-(\{D(S) \mid S \in \mathbf{S}^+\}) \subseteq \mathbf{E}^-(\beta_L) \subseteq \mathbf{E}^-(L)$.
If any $\neg D(S) \in L$ then either $\neg D(S) \in \beta_L$, so $D(S) \in \beta_R$, or $\neg D(S) \in \mathbf{E}(\beta_R) \cap \mathbf{E}^-(\beta_R)$, which implies $D(S) \in \beta_R$, since $\mathbf{E}(\neg D(S)) \overset{8.2}{=} \{D(S)\}$. In either case, $D(S) \in \beta_R$ contradicts $\beta_L \cap \beta_R = \emptyset$. Thus $\mathbf{E}(\forall\phi.D(\phi)) \subseteq \mathbf{V} \setminus L$.

v. For $A = \forall x.D(x)$, the argument is as in **iv**. $\forall x.D(x) \in \beta_L$ implies $\{D(t) \mid t \in \mathbf{T}_M\} \subseteq \beta_L$, so $\mathbf{E}(\forall x D(x)) \overset{8.2}{=} \{\neg D(t) \mid t \in \mathbf{T}_M\} \overset{8.2}{\subseteq} \mathbf{E}^-(\{D(t) \mid t \in \mathbf{T}_M\}) \subseteq \mathbf{E}^-(\beta_L) \subseteq \mathbf{E}^-(L)$.
If any $\neg D(t) \in L$, then either $\neg D(t) \in \beta_L$, so $D(t) \in \beta_R$, or $\neg D(t) \in \mathbf{E}(\beta_R) \cap \mathbf{E}^-(\beta_R)$, which implies $D(t) \in \beta_R$, since $\mathbf{E}(\neg D(t)) \overset{8.2}{=} \{D(t)\}$. In either case, $D(t) \in \beta_R$ contradicts $\beta_L \cap \beta_R = \emptyset$. Thus $\mathbf{E}(\forall x.D(x)) \subseteq \mathbf{V} \setminus L$.

7. Also each sentence $S \in Z = (\mathbf{E}(\beta_R) \cap \mathbf{E}^-(\beta_R)) \setminus \beta_L$ satisfies $\mathbf{E}(S) \subseteq \mathbf{E}^-(L) \cap (\mathbf{V} \setminus L)$:

i. If $S \in Z$ does not start with $\neg$, then $\mathbf{E}^-(S) \overset{8.2}{=} \{\neg S\}$, so $\neg S \in \beta_R$, implying $S \in \beta_L$, so $S \notin Z$.

ii. If $S = \neg A \in Z \subseteq \mathbf{E}^-(\beta_R)$ then $\mathbf{E}(\neg A) \overset{8.2}{=} \{A\} \subset \beta_R \overset{\mathbf{5}}{\subseteq} \mathbf{E}^-(L)$. If $A \in Z$, then it starts with $\neg$ by 7.i, i.e., $A = \neg B$ and $\mathbf{E}(\neg B) \overset{8.2}{=} \{B\} \subset \beta_R$. Since also $A \in \beta_R$ so $B \in \beta_L$, contradicting $\beta_L \cap \beta_R = \emptyset$. Hence $A \notin Z$ and $A \notin \beta_L$ (since $A \in \beta_R$), i.e., $A \notin L = Z \cup \beta_L$, so that $\mathbf{E}(\neg A) = \{A\} \subseteq \mathbf{V} \setminus L$.

By **6** and **7**, $\mathbf{E}(L) = \mathbf{E}(\beta_L) \cup \mathbf{E}(Z) \subseteq \mathbf{E}^-(L) \cap (\mathbf{V} \setminus L)$, so $L \in SK(G)$ by Fact 7.1. $\square$

Unlike in variants of circular proof theory, an infinite branch gives always a rise to a countermodel. A paradigmatic example of a cyclic proof, with the same sequent reappearing infinitely often in a branch, can be the attempted derivation of $\forall\phi.\phi$:

$$(\forall^+_L)\ \dfrac{\dfrac{\vdots}{\vdash \forall\phi.\phi} \qquad \dfrac{A \vdash \qquad \vdash A}{\vdash \neg A \wedge A} \qquad \dfrac{\vdots}{\vdash S}\ \text{... for all } S \in \mathbf{S}^+}{\vdash \forall\phi.\phi}$$

Any sentence $\neg A \wedge A$ gives a counterexample when any of its branches does not terminate with an axiom, providing a countermodel. Such is, in particular, the leftmost branch where $\forall\phi.\phi$ instantiates ϕ in the root sentence and is expanded further, giving a copy of the whole tree and, eventually, a special branch β with $\beta_L = \emptyset$ and $\beta_R = \{\forall\phi.\phi\}$. This infinite branch provides also a countermodel, with $\forall\phi.\phi = \mathbf{0}$. This looks strange, but is verified by inspecting graph $G_M(A_\forall)$ in Example 3.10, according to which $\forall\phi.\phi$ does act as a witness to its own falsity.

A branch with a repeated sequent can be terminated, with the conclusion of unprovability, if one can verify that subsequent substitutions, higher up in the derivation, will also yield the same sequent. A single repetition is not enough, as it may be due to a specific substitution, while subsequent ones might yield new sequents.

The remaining fact is soundness and completeness with (cut) for the kernel semantics (3.13).

Theorem 8.4 **(4.3)** *For Γ, Δ over a countable $\mathcal{L}^+ \in FOL^+$, $\Gamma \models_c \Delta$ iff $\Gamma \vdash_c \Delta$.*

Proof Soundness and invertibility follow by essentially the same argument as in Fact 8.1, with some simplifications due to each kernel $K \in sol(G)$ covering the entire graph, $\mathbf{E}^-_G(K) = \mathbf{V}_G \setminus K$. We fix an arbitrary graph G and show each case for an arbitrary fixed $K \in sol(G)$.

1. $(\wedge_R)$. For soundness, let $\Gamma \models_c \Delta, A_1$, $\Gamma \models_c \Delta, A_2$, and $\alpha \in M^{\mathcal{V}(\Gamma,\Delta)}$. Then $\alpha(\Gamma) \cap \mathbf{V} \setminus K \neq \emptyset$ or $\alpha(\Delta) \cap K \neq \emptyset$, in which case also conclusion is satisfied under α, or else $\{\alpha(A_1), \alpha(A_2)\} \subseteq K$. Then $\{\neg\alpha(A_1), \neg\alpha(A_2)\} \subseteq \mathbf{V} \setminus K$, and hence $\alpha(A_1 \wedge A_2) \in K$ since $\mathbf{E}(\alpha(A_1 \wedge A_2)) = \{\neg\alpha(A_1), \neg\alpha(A_2)\}$.
   For invertibility, let $\Gamma \models_c \Delta, A_1 \wedge A_2$. If $\alpha(\Gamma) \cap \mathbf{V} \setminus K \neq \emptyset$ or $\alpha(\Delta) \cap K \neq \emptyset$, then K satisfies also both premises under α. If neither is the case, then $\alpha(A_1 \wedge A_2) \in K$, hence $\mathbf{E}(\alpha(A_1 \wedge A_2)) = \{\neg\alpha(A_1), \neg\alpha(A_2)\} \subseteq \mathbf{V} \setminus K$ and, $\mathbf{E}(\neg\alpha(A_i)) = \alpha(A_i) \in K$, for $i \in \{1, 2\}$, hence $K \models_\alpha \Gamma \Rightarrow \Delta, A_i$.
   Assignments to free FOL-variables do not affect the arguments below. They are relative to a given α, which we do not mention, except for $(\forall_R)$. In each case, we assume that $\Gamma \subseteq K$ and $\Delta \subseteq \mathbf{V} \setminus K$ focusing on the active/principal formulas.
2. $(\wedge_L)$. For soundness, assuming $\Gamma, A_1, A_2 \models_c \Delta$ (and $\Gamma \subseteq K$ and $\Delta \subseteq \mathbf{V} \setminus K$), $A_i \in \mathbf{V} \setminus K$, for $i = 1$ or $i = 2$. Then $\neg A_i \in K$, since $\mathbf{E}(\neg A_i) = A_i$, and $A_1 \wedge A_2 \in \mathbf{E}^-(\neg A_i) \subseteq \mathbf{E}^-(K) \subseteq \mathbf{V} \setminus K$. Thus $K \models \Gamma, A_1 \wedge A_2 \Rightarrow \Delta$.
   For invertibility, assume $\Gamma, A_1 \wedge A_2 \models_c \Delta$. If $A_1, A_2 \in K$, which is the only way K can contradict $\Gamma, A_1, A_2 \models_c \Delta$, then $\mathbf{E}(A_1 \wedge A_2) = \{\neg A_1, \neg A_2\} \subseteq \mathbf{E}^-(\{A_1, A_2\}) \subseteq \mathbf{E}^-(K) = \mathbf{V} \setminus K$, and hence $A_1 \wedge A_2 \in K$, contradicting $K \models \Gamma, A_1 \wedge A_2 \Rightarrow \Delta$.
3. $(\neg_R)$. For soundness, $\Gamma, A \models_c \Delta$ implies $A \in \mathbf{V} \setminus K$, so $\neg A \in K$, since $\mathbf{E}(\neg A) = A$.
   For invertibility, $\Gamma \models_c \Delta, \neg A$ implies $\neg A \in K$, so $\mathbf{E}(\neg A) = \{A\} \subseteq \mathbf{V} \setminus K$ and $K \models \Gamma, A \Rightarrow \Delta$.
4. $(\neg_L)$. For soundness, $\Gamma \models_c \Delta, A$, implies $A \in K$ hence $\neg A \in \mathbf{V} \setminus K$ and $K \models \Gamma, \neg A \Rightarrow \Delta$.
   For invertibility, $\Gamma, \neg A \models_c \Delta$ implies $\neg A \in \mathbf{V} \setminus K$, hence $\mathbf{E}(\neg A) = A \in K$ and $K \models \Gamma \Rightarrow \Delta, A$.

5. $(\forall_L)$. For soundness, assume $F(t), \Gamma, \forall x F(x) \models_c \Delta$. If $\forall x F(x) \notin K$, i.e., $\forall x F(x) \in \mathbf{V} \setminus K$, then $(\Gamma \cup \{\forall x F(x)\}) \cap \mathbf{E}^-(K) \neq \emptyset$, so $K \models \Gamma, \forall x F(x) \Rightarrow \Delta$. If $\forall x F(x) \in K$ then also $F(t) \in K$, since $\neg F(t) \in \mathbf{E}(\forall x F(x)) \subseteq \mathbf{V} \setminus K$, so $\mathbf{E}(\neg F(t)) \cap K \neq \emptyset$ while $\mathbf{E}(\neg F(t)) = \{F(t)\}$. Thus either $\Gamma \cap (\mathbf{V} \setminus K) \neq \emptyset$ or $\Delta \cap K \neq \emptyset$, yielding $K \models \Gamma, \forall x F(x) \Rightarrow \Delta$.
Invertibility follows by weakening since $\Gamma, \forall x F(x) \models_c \Delta$ implies $F(t), \Gamma, \forall x F(x) \models_c \Delta$.

6. $(\forall_R)$. For soundness, let (*) $\Gamma \models_c \Delta, F(y)$, with eigenvariable $y \notin \mathcal{V}(\Gamma, \Delta)$, and $\alpha(\Gamma) \subseteq K$ and $\alpha(\Delta) \subseteq \mathbf{V} \setminus K$. If $\alpha(\forall x F(x)) \notin K$ then $\alpha(\forall x F(x)) \in \mathbf{E}^-(K)$ and some $\alpha(\neg F(m)) \in K$, since $\mathbf{E}(\alpha(\forall x F(x))) = \{\alpha(\neg F(m)) \mid m \in M\}$. Extending α with $\alpha(y) = m$ yields $\alpha(\Gamma) \subseteq K$ and $\alpha(\Delta, F(y)) \subseteq \mathbf{V} \setminus K$, contrary to (*). Hence, $\alpha(\forall x F(x)) \in K$.
For invertibility, if $\alpha(\Gamma) \subseteq K$ and $\alpha(\Delta, F(y)) \subseteq \mathbf{V} \setminus K$, for $\alpha(y) = m$, then $\neg\alpha(F(m)) \in K$, since $\mathbf{E}(\neg\alpha(F(m))) = \alpha(F(m)) \in \mathbf{V} \setminus K$. Then $\alpha(\forall x F(x)) \in \mathbf{E}^-(\neg\alpha(F(m))) \subseteq \mathbf{E}^-(K) = \mathbf{V} \setminus K$, giving $\alpha(\Gamma) \subseteq K$ and $\alpha(\Delta, \forall x Fx) \subseteq \mathbf{V} \setminus K$, which contradicts $K \models \Gamma \Rightarrow \Delta, \forall x Fx$.

7. $(\forall_L^+)$. The argument repeats that for $(\forall_L)$. For soundness, assume $\Gamma, F(S), \forall\phi F(\phi) \models_c \Delta$. If $\forall\phi F(\phi) \in K$ then also $F(S) \in K$, since $\neg F(S) \in \mathbf{E}(\forall\phi F(\phi)) \subseteq \mathbf{E}^-(K)$ and $\mathbf{E}(\neg F(S)) = \{F(S)\}$. Hence, as $\Gamma \subseteq K$ and $\Delta \subseteq \mathbf{V} \setminus K$, it holds $\forall\phi F(\phi) \in \mathbf{V} \setminus K$ and $K \models \Gamma, \forall\phi F(\phi) \Rightarrow \Delta$.
For invertibility, assuming $\Gamma, \forall\phi F(\phi) \models_c \Delta$, weakening yields $F(S), \Gamma, \forall\phi F(\phi) \models_c \Delta$.

8. $(\forall_R^+)$. For soundness, let $\Gamma \models_c \Delta, F(S)$ for every $S \in \mathbf{S}^+$. If $\forall\phi F(\phi) \notin K$ then $\forall\phi F(\phi) \in \mathbf{E}^-(K)$ and some $\neg F(S) \in K$, since $\mathbf{E}(\forall\phi F(\phi)) = \{\neg F(S) \mid S \in \mathbf{S}^+\}$. Since $F(S) \notin K$, either $\Gamma \cap \mathbf{E}^-(K) \neq \emptyset$ or $\Delta \cap K \neq \emptyset$, contradicting the assumption $\Gamma \subseteq K$ and $\Delta \subseteq \mathbf{V} \setminus K$. Hence $\forall\phi F(\phi) \in K$ and K satisfies the rule's conclusion.
For invertibility, assume $\Gamma \models_c \Delta, \forall\phi F(\phi)$, and $\forall\phi F(\phi) \in K$. If for some $S \in \mathbf{S}^+$, $F(S) \in \mathbf{V} \setminus K$, then $\neg F(S) \in K$ since $\mathbf{E}(\neg F(S)) = \{F(S)\} \subseteq \mathbf{V} \setminus K$. But $\neg F(S) \in \mathbf{E}(\forall\phi F(\phi))$, contradicting independence of K. Hence, $F(S) \in K$ for all $S \in \mathbf{S}^+$.

9. The rules for $\doteq$ are sound and invertible – for the intended interpretation of s-equality – by the same argument as in the case of semikernels, point 9 in the proof of Fact 8.1.
This concludes the proof of soundness. For completeness, we modify the construction from the proof of Fact 8.3, by interleaving the enumeration of all triples $E_S \times E_T \times E^+$ and pairs $E_S \times E_S$ with enumeration E'_S of all FOL$^+$ formulas without free s-variables, where each such formula occurs only once. Following this interleaved enumeration yields now a new case, **1.vi**, of an $A \in E'_S$, in which we expand each active sequent $\Gamma \Rightarrow \Delta$ with the premises of (cut) over A, i.e., with $\Gamma \Rightarrow \Delta, A$ and $\Gamma, A \Rightarrow \Delta$. A semikernel falsifying any one of them, falsifies the conclusion. Given an infinite nonaxiomatic branch β, a language graph G_M is obtained as in the proof of Fact 8.3, over domain M consisting of free variables and ground terms used in the standard construction of a FOL countermodel for $\beta \cap \mathbf{S}_M$. Point **3** of the proof of Fact 8.3 shows β to determine a semikernel K of G_M, falsifying each sequent on β. Now, β contains one of the premises of an

application of (cut) for each $A \in E'_S = \mathbf{S}^+_M$. As every sentence $\mathbf{S}^+_M$ occurs thus in β_L or β_R, semikernel K covers all $\mathbf{S}^+_M$, so it is a kernel of G_M. $\square$

Acknowledgements The author gratefully acknowledges partial financial support from L. Meltzers Høyskolefond. Lavinia Picollo deserves thanks for many valuable comments and detailed corrections.

Funding Open access funding provided by University of Bergen (incl Haukeland University Hospital)

References

1. Asher, N. (1990). Intentional paradoxes and an inductive theory of propositional quantification. In R. Parikh (Eds.), *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*, New York: Morgan Kaufmann
2. Azzouni, J. (2001). Truth via anaphorically unrestricted quantifiers. *Journal of Philosophical Logic, 30*, 329–354.
3. Azzouni, J. (2003). The strengthened liar, the expressive strength of natural languages, and regimentation. *The Philosophical Forum, 34*(3–4), 329–350.
4. Bacon, A., Hawthorne, J., & Uzquiano, G. (2016). Higher-order free logic and the Prior-Kaplan paradox. *Canadian Journal of Philosophy, 46*(4–5), 493–541.
5. Båve, A. (2012). On using inconsistent expressions. *Erkenntnis, 77*, 133–148.
6. Beringer, T., & Schindler, T. (2017). A graph-theoretic analysis of the semantic paradoxes. *The Bulletin of Symbolic Logic, 23*(4), 442–492.
7. Bolander, T. (2003). *Logical theories for agent introspection*. PhD thesis, IMM, Technical University of Denmark
8. Brandenburger, A., & Keisler, H. J. (2006). On possible worlds and related notions. In V. F. Hendricks, & S. A. Pedersen (Eds.), *Studia Logica* (pp. 211–240). Special Issue: Ways of Worlds II.
9. Chihara, C. (1974). The semantic paradoxes: a diagnostic investigation. *The Philosophical Review, 88*(4), 590–618.
10. Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2013). Reaching transparent truth. *Mind, 122*(488), 841–866.
11. Cook, R. (2004). Patterns of paradox. *The Journal of Symbolic Logic, 69*(3), 767–774.
12. Deutsch, H. (2010). Diagonalization and truth functional operators. *Analysis, 70*(2), 215–217.
13. Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence, 77*, 321–357.
14. Dyrkolbotn, S., & Walicki, M. (2014). Propositional discourse logic. *Synthese, 191*(5), 863–899.
15. Eklund, M. (2002). Inconsistent languages. *Philosophy and Phenomenological Research, 64*(2), 251–275.
16. Gentzen, G. (1935). Untersuchungen über das logische Schließen. *Mathematische Zeitschrift, 39*, 176–210. [translation in M. E. Szabo (Eds.) (1969). *The Collected Papers of Gerhard Gentzen*. North-Holland Publ. Co.].
17. Grim, P. (1993). Operators in the paradox of the knower. *Synthese, 44*, 409–428.
18. Halbach, V. (1999). Disquotationalism and infinite conjunctions. *Mind, 108*, 1–22.
19. Halbach, V. (2009). On the benfits of a reduction of modal predicates to modal operators. In A. Hieke, & H. Leitgeb (Eds.), *Reduction, Abstraction, Analysis*, pages 323–333. Ontos Verlag.
20. Horsten, L., & Leigeb, H. (2001). No future. *Journal of Philosophical Logic, 30*(3), 259–265.

21. Kaplan, D. (1995). A problem in possible world semantics. In Sinott-Armstrong, W., Raffman, D., & Asher, N. (Eds.), *Modality, Morality and Belief: Essays in Honor of Ruth Barcan Marcus*, pp. 41–52. Cambridge University Press.
22. Ketland, J. (2005). Yablo's paradox and ω-inconsistency. *Synthese, 149*, 295–302.
23. Kripke, S. (1976). Is there a problem about substitutional quantification? In G. Evans, & J. McDowell (Eds.), *Truth and Meaning* (pp. 325–419). Oxford: Clarendon Press.
24. Kripke, S. (2011). A puzzle about time and thought. In *Philosophical Troubles. Collected Papers*, vol. I. Oxford University Press.
25. McGee, V. (1992). Maximal consistent sets of instances of Tarski's schema (T). *Journal of Philosophical Logic, 21*, 235–241.
26. Montague, R. (1963). Syntactical treatment of modality, with corollaries on reflection principles and finite axiomatizability. *Acta philosophica Fennica, 16*, 153–167.
27. Mulligan, K. (2010). The truth predicate vs. the truth connective. On taking connectives seriously. *Dialectica, 64*, 565–584.
28. Neumann-Lara, V. (1971). Seminúcleos de una digráfica. Technical report, Anales del Instituto de Matemáticas II, Universidad Nacional Autónoma México.
29. Patterson, D. (2009). Inconsistency theory of semantic paradox. *Philosophy and Phenomenological Research, 79*(2), 387–422.
30. Perlis, D. (1988). Languages with self-reference II: knowledge, belief, and modality. *Arificial Intelligence, 34*, 179–212.
31. Picollo, L., & Schindler, T. (2018). Disquotation and infinite conjunctions. *Erkenntnis, 83*, 899–928.
32. Prior, A. N. (1961). On a family of paradoxes. *Notre Dame Journal of Formal Logic, 2*(1), 16–32.
33. Quine, W. V. O. (1970). *Philosophy of logic*. Cambridge: Harvard University Press.
34. Rabern, L., Rabern, B., & Macauley, M. (2013). Dangerous reference graphs and semantic paradoxes. *Journal of Philosophical Logic, 42*, 727–765.
35. S. Rahman, T. Tulenheimo, & E. Genot (Eds.) (2008). *Unity, Truth and the Liar*. Springer.
36. Read, S. (2002). The liar paradox from John Buridan back to Thomas Bradwardine. *Vivarium, 40*(2), 189–218.
37. Reinhardt, W. N. (1980). Necessity predicates and operators. *Journal of Philosophical Logic, 9*(4), 437–450.
38. Richardson, M. (1953). Solutions of irreflexive relations. *The Annals of Mathematics, Second Series, 58*(3), 573–590.
39. Ripley, D. (2012). Conservatively extending classical logic with transparent truth. *The Review of Symbolic Logic, 5*(2), 354–378.
40. Ripley, D. (2013). Paradoxes and failures of cut. *Australasian Journal of Philosophy, 91*(1), 139–164.
41. Rossi, L. (2019). A unified theory of truth and paradox. *The Review of Symbolic Logic, 12*(2), 209–254.
42. Stern J. (2011). The truth about predicates and connectives. In A. Reboul (Eds), *Philosophical papers dedicated to Kevin Mulligan*, pp. 1–14. Genéve.
43. M. E. Szabo (Eds.) (1969). *The Collected Papers of Gerhard Gentzen*. North-Holland Publ. Co.
44. Tarski, A. (1956). The concept of truth in formalised languages. In *Logic, Semantics, Metamathematics, papers from 1923 to 1938*. Oxford at the Clarendon Press. [translation of the Polish original from 1933 by J. H. Woodger].
45. Tarski, A. (1969). Truth and proof. *Scientific American*, pp. 63–77.
46. Troelstra, A. S., & Schwichteberg, H. (1996). *Basic Proof Theory*. Cambridge University Press.
47. von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
48. Walicki, M. (2017). Resolving infinitary paradoxes. *Journal of Symbolic Logic, 82*(2), 709–723.
49. Walicki, M. (2019). Kernels of digraphs with finitely many ends. *Discrete Mathematics, 342*, 473–486.
50. Walicki, M., & Dyrkolbotn, S. (2022). Paraconsistent resolution. *Australasian Journal of Logic, 19*(3), 96–123.
51. Wolfgram, J., & Cook, R. (2022). Strong homomorphisms, category theory and semantic paradox. *The Review of Symbolic Logic, 15*(4), 1–24.