



Logics of Synonymy

Levin Hornischer¹ 

Received: 30 October 2018 / Accepted: 14 November 2019 / Published online: 9 January 2020
© The Author(s) 2020

Abstract

We investigate synonymy in the strong sense of content identity (and not just meaning similarity). This notion is central in the philosophy of language and in applications of logic. We motivate, uniformly axiomatize, and characterize several “benchmark” notions of synonymy in the messy class of all possible notions of synonymy. This class is divided by two intuitive principles that are governed by a no-go result. We use the notion of a scenario to get a logic of synonymy (SF) which is the canonical representative of one division. In the other division, the so-called conceptivist logics, we find, e.g., the well-known system of analytic containment (AC). We axiomatize four logics of synonymy extending AC, relate them semantically and proof-theoretically to SF, and characterize them in terms of weak/strong subject matter preservation and weak/strong logical equivalence. This yields ways out of the no-go result and novel arguments—independent of a particular semantic framework—for each notion of synonymy discussed (using, e.g., Hurford disjunctions or homotopy theory). This points to pluralism about meaning and a certain non-compositionality of truth in logic programs and neural networks. And it unveils an impossibility for synonymy: if it is to preserve subject matter, then either conjunction and disjunction lose an essential property or a very weak absorption law is violated.

Keywords Synonymy · Content identity · Hyperintensionality · Truthmaker semantics · Conceptivism · Neural networks

1 Introduction

One of the most important problems in philosophy of language and related disciplines is to understand synonymy: when do two expressions mean the same thing? Similarly, in all instances of logical modeling it is crucial to become clear on which sentences

✉ Levin Hornischer
l.a.hornischer@uva.nl

¹ Institute for Logic, Language and Computation (ILLC), University of Amsterdam,
P.O. Box 94242, Amsterdam, 1090 GE, The Netherlands

should be considered to be equivalent: when a fallible cognitive agent is modeled, this is different from when, say, metaphysical grounding is modeled.

This problem is a persistently hard one since synonymy is such a multifaceted concept. In everyday language or in thesauri, we find many alleged synonyms that surely are synonymous in a wide range of contexts (contextual stability). However, on a less credulous and more critical stance, we usually can find contexts in which those synonyms differ in meaning (contextual flexibility). Indeed, for any two non-identical, logically atomic sentences it seems like we can almost always cook up a weird context in which they differ in meaning [30, 40].

There are more pairs of opposing features of synonymy other than “stability vs. flexibility”. On the one hand, we think that whether or not two sentences are synonymous is objectively and externally settled by the language and the world alone. On the other hand, there also is the intuitive idea (which is largely discredited in modern philosophy of language) that whether or not two expressions are synonymous also depends on subjective and internal attitudes of the speaker. Moreover, on the one hand, we sometimes think that usually synonymy respects logical equivalence (rendering it an at most intensional concept). On the other hand, there is reason to believe that classically equivalent sentences, like $p \vee \neg p$ and $q \vee \neg q$, are not synonymous since they are “about” different things (rendering synonymy a so-called hyperintensional concept). Furthermore, sometimes we think of two synonyms as being identical in meaning and sometimes as being only very similar in meaning. This list of opposing features could be extended much further.

In this paper, we’re concerned with synonymy in the strong sense of meaning (or content) *identity*—and not just meaning *similarity*.¹ That is, we’re interested in the notion of synonymy that we get when we adopt a critical stance and move to the more discerning side in the pairs of the opposing features. Many different semantic and proof-theoretic systems have been proposed to explicate such fine notions of synonymy—just to name a few: [1, 5, 10, 11, 13, 18, 20, 23, 35, 37, 41, 44, 46], and more will be mentioned below. However, these approaches differ tremendously and there is practically no consensus on which approach is correct. Because of this, we want to understand the intuitive notion (or family of notions) of strong synonymy *without* committing to a particular framework. To do so, we work with various logics (or axiomatizations) that attempt to capture the (or a) notion of synonymy. Thus, we’ll gain insights into synonymy directly and not into a particular framework representing synonymy.

The main contribution of this paper is to motivate, uniformly axiomatize, and characterize “benchmark” notions of synonymy in the messy class of possible notions of synonymy (as, e.g., provided by different frameworks). This helps to quickly identify not only the notion of synonymy of a given framework but also the precise synonymies that make it different from other frameworks. Moreover, this provides

¹But what do we mean by “meaning” or “content”? While we do have some intuitive grasp of these concepts, it is even more difficult to say precisely what kind of *entity* content is than saying what its *identity* is, that is, when two sentences should have the same content. This is why we focus on understanding the notion of synonymy rather than the ontology of content.

novel arguments and impossibility results for the various notions of synonymy that are *independent* of a particular conception of semantics.

Summary In Section 2, we show why finding the logical laws governing synonymy is problematic: Neither the famous possible worlds semantics nor any straightforward refinement of it can satisfy the fundamental principle of synonymy that being synonymous entails having the same subject matter. To understand this no-go result, we look, in Section 3, at a well-known logic satisfying that principle: the system of analytic containment (AC) with a sound and complete truthmaker semantics due to [23]. And in Section 4, we develop a formal notion of a scenario to see just how fine-grained can get the notion of synonymy in a straightforward refinement of possible world semantics. We axiomatize this notion as an extension SF of AC. In Section 5, we see that the two logics AC and SF are related by moving up one set-theoretic level: If we take sentences not to be true at a scenario but at sets of scenarios, we get a semantics that is extensionally equivalent to truthmaker semantics.

In Section 6, we investigate the lattice of conceptivist logics, that is, logics where synonymy (or equivalence) entails having the same atomic sentences—which hence can be regarded to satisfy the fundamental principle about synonymy. We show the main formal result of the paper: We identify various extensions of AC that correspond to the possible combinations of characterizing synonymy by weak/strong subject matter preservation and weak/strong logical equivalence. In Section 7, we thus can offer novel arguments for the various notions of synonymy (making use of, e.g., the Hurford constraint, a “truth plus subject matter” conception of meaning, or homotopy theory).

In Section 8, we discuss ways out of the no-go result: we identify the exact reason for the inconsistency and describe possible ways of adding more (intensional) structure to the notion of scenario that allows them to satisfy the fundamental principle about synonymy. Moreover, we state some consequences of the paradox: The inconsistency and the various arguments point to a pluralistic conception of meaning and to the non-compositionality of the notion of truth in logic programs or states of a neural network. In Section 9, we analyze this non-compositionality and generalize it to the following impossibility result for synonymy: if synonymy preserves subject matter, either some of our basic intuitions about conjunction and disjunction are violated or a very weak logical law—satisfied by most logics—is violated. We provide some linguistic and cognitive evidence against the law. This poses the problem of finding a logic violating this law and accounting for the evidence.

Most results are proven in an [Appendix](#) extending the method of proving completeness by normal forms. Thus, the characterization of the synonymies is proven uniformly and constructively. Moreover, it associates to each of these synonymies a characteristic notion of disjunctive normal form which can be seen as an invariant of the synonymy (across different, philosophically-laden theories and semantics for this synonymy).

Notation As common in the field, we only deal with sentences built in the usual way from a set of propositional letters p_0, p_1, \dots (called *atoms*) using the connectives \neg, \wedge, \vee . Variables for sentences are $\varphi, \psi, \chi, \dots$ and variables for atomic sentences

are p, q, r, \dots . A statement of the form $\varphi \equiv \psi$ is called *equivalential*. A logic of synonymy \mathcal{L} is a logic reasoning with equivalential statements: $\mathcal{L} \vDash \varphi \equiv \psi$ (respectively, $\mathcal{L} \vdash \varphi \equiv \psi$) means that under the semantics of \mathcal{L} , $\varphi \equiv \psi$ is valid (respectively, is derivable with the rules of \mathcal{L}).

To keep the framework as simple as possible, we don't add further operators like conditionals, modal operators, or even hyperintensional operators. We leave this to further research.

2 The Cause of the Problem: a No-go Result About Synonymy

In this section, we show why it is a real problem to find the appropriate axioms or laws governing synonymy in the strong sense of meaning identity: A fundamental principle about this notion is that being synonymous entails having the same subject matter. However, we argue that neither standard possible worlds semantics nor any straightforward refinement of it (with the same underlying idea) can satisfy this principle. To understand synonymy, it arguably is not enough to just provide a particular semantics with a notion of synonymy that satisfies the subject matter preservation principle. Rather, we need a *systematic* understanding of this impossibility result and how to avoid it—whence we need to analyze and compare different notions of synonymy.

As indicated in the introduction, we explore the possible logical laws of synonymy. That is, we want to know, for example, is it a general law of synonymy that p is synonymous to $p \vee (p \wedge q)$? Is $p \vee (p \wedge q)$ synonymous to $p \vee (p \wedge \neg q)$? What about $p \vee (p \wedge q)$ and $p \wedge (p \vee q)$?

Since we're interested in the general *logical* laws of synonymy, we're not so much interested in when logically atomic sentences are synonymous. Rather, we consider which logically complex sentences formed with these atomic sentences should be synonymous as a matter of a general law about synonymy.² So we assume we've fixed some (or, rather, any) theory about how to map logically atomic natural language sentences (that we're interested in) to propositional atoms such that atomic sentences with the same meaning are mapped to the same propositional atom and atomic sentences with different meaning are mapped to different atoms. Thus, distinct propositional atoms represent atomic sentences with distinct meaning.³ Whatever theory we've picked, our results can then be applied to see which complex sentences should be synonymous.

Let's illustrate this in three remarks. First, we could use some of the established theories to determine when two atomic sentences express the same proposition: possible world semantics, (exact) truthmaker semantics, two-dimensional semantics, structured propositions, impossible world semantics, etc. Then different propositional atoms represent different equivalence classes of atomic sentences (under the

²This is in line with the common conception that logic is topic-neutral and universal—and hence should be independent of particular interpretations of the atomic sentences.

³A more technical way to make this point is as follows: Map each atomic natural language sentence to a distinct propositional atom. Then write $p \approx q$ if p and q represent atomic sentences with the same meaning according to the theory we've fixed. Introduce a new set of propositional atoms p'_0, p'_1, \dots each representing a single \approx -equivalence class. Then work with p'_0, p'_1, \dots

equivalence relation of ‘expressing the same atomic proposition’ in the respective semantics). We’re then investigating what the appropriate notion of synonymy for complex sentences is—independent of what the respective semantics for atomic sentences says. Second, assume one is convinced by the view mentioned in the introduction that no two syntactically non-identical atomic sentences are completely identical in meaning (except, say, sentences with different allowed spellings or emphasis, etc.). Then every atomic sentence is assigned to a distinct propositional atom. Third, assume the meaning of atomic sentences is governed by a set of (defeasible) rules describing our semantic knowledge or contextual information (or both). This may include the rules ‘A sofa is a couch’ or—in a context far away from penguins, ostriches, and the like—‘A bird is an animal that can fly’. We then again assign every atomic sentence to a distinct propositional atom, but we add the synonymies ‘This is a couch \equiv This is a sofa’ or ‘This is a bird \equiv This is an animal that can fly’ as *non-logical* axiom to our logical axioms of synonymy. However, since we’re interested in the logical laws of synonymy, we won’t consider these additional non-logical axioms.

With these preliminaries at hand, we can get to arguably one of the most fundamental principles about synonymy: if two sentences are synonymous, they are about the same thing. In other words, if two sentences are about different things, there is a sense in which they are not synonymous—and hence shouldn’t be synonymous in the strong sense of meaning identity. As a slogan: synonymy entails subject matter identity.

To formulate this principle precisely, we need to specify what subject matter is. We could choose one of the reconstructions of this intuitive notion of subject matter [24, 31, 42, 58]. However, we can remain general and independent of particular reconstructions by working with the *syntactic reflection* of the subject matter of a sentence, that is, with the set of its atomic sentences. This is because we assumed that distinct propositional atoms represent distinct atomic propositions: so if two sentences are built from different atoms, one is about a proposition that the other isn’t about, whence they should have different subject matter. In particular, the most straightforward counterexample is excluded: we cannot have anymore two atomic sentences—like ‘The couch is black’ and ‘The sofa is black’—that express the same atomic proposition, and hence have the same subject matter, but still get assigned to different propositional atoms. Hence we have the following.

(S1) *Completeness of syntactic reflection.* If two sentences have the same subject matter, they have the same atomic sentences.

Alternatively, we can also think of (S1) as a logical—as opposed to a semantic—principle: if two sentences have different atoms, they have, as far as logic is concerned, different subject matter. (In Section 7, we also discuss the converse of (S1), that is, the soundness of syntactic reflection.) Thus, we can work with the following precise formulation of the fundamental principle that synonymy entails subject matter identity:

(P1) *Subject matter preserving.* If two sentences are synonymous, then the sets of their atomic sentences are identical.

When we look at the class of possible sets of logical laws governing synonymy below, this principle is conceptually useful in dividing the class into those sets of laws (or axioms) that satisfy the principle and those that don't. In particular, we'll consider what it takes to get into the class of subject matter preserving synonymies and what the benchmark representatives of this class are.

So what are notions of synonymy that satisfy the subject matter preservation principle (P1)? We immediately see that it cannot be provided by standard possible worlds semantics: For any distinct atoms p and q , the sentences $p \wedge \neg p$ and $q \wedge \neg q$ are equivalent in classical logic and hence not distinguishable by a possible world. So they are synonymous according to possible worlds semantics but they have different subject matter.

Thus, an advocate of possible worlds semantics may wonder whether it can be changed minimally and straightforwardly to obtain a notion of synonymy that satisfies principle (P1).

To answer this, let's recall the main idea of possible worlds semantics: Sentences are evaluated at possible worlds and the meaning of a sentence is its truth-value profile across all possible worlds. As far as semantics are concerned, possible worlds can be regarded as maximally consistent scenarios: every sentence is either true or false at a world (maximality) and no sentence is both true and false at a world (consistency).

Thus, arguably the most minimal and straightforward change to possible world semantics is to relax the minimality and/or consistency assumption. So possible worlds are relaxed to scenarios (or circumstances) where sentences can be not only (i) true, or (ii) false; but also (iii) not maximal, i.e., neither true or false, or (iv) not consistent, i.e., both true and false. Let's abbreviate these truth-values as follows: true t , false f , neither true nor false u (for *undecided*), and both true and false b (for *both*). Then there are three possibilities for the new sets of truth-values: $\{t, f, u\}$, $\{t, f, b\}$, and $\{t, f, u, b\}$.⁴

The meaning of a sentence still is its truth-value profile across all possible scenarios, though now using more than just the two truth-values true and false. For possible worlds, the truth-value of a complex sentence at a possible world is determined by the truth-value of its parts at that possible world according to the most straightforward logic for the set of truth-values $\{t, f\}$, i.e., according to classical logic. Thus, to change possible worlds semantics minimally, the truth-value of a complex sentence at a possible scenario is determined by the truth-value of its parts at that scenario according to the most straightforward logic for the chosen set of truth-values. Standardly, these are the following. For $\{t, f, u\}$ the truth-functions for the connectives \neg, \wedge, \vee are as in (strong) Kleene 3-valued logic (\mathcal{K}_3), for $\{t, f, b\}$ as in the Logic of

⁴Note that, by this line of thought, there is no reason to add more truth-values and/or interpret u and b differently than as 'neither true nor false' and 'both true and false', respectively. (Though, later we'll discuss such other options: for example, using five or more truth-values and interpreting u as meaninglessness [9] or off-topicness [3].) This is because we've added those truth-values (with their intended interpretation) that have to be added when relaxing one or both of the two conditions on truth-values in classical possible worlds, but there seems to be no straightforward motivation for further and/or differently interpreted truth-values.

Paradox (LP), and for $\{t, f, u, b\}$ as in First-Degree Entailment (FDE).⁵ See Section 4 and [49, sec. 7–8] for more details on these logics—though for now that is not needed.

To summarize, the most minimal and straightforward change of possible world semantics arguably is to relax possible worlds to scenarios where:

- (A1) Scenarios are structures where atomic sentences can be evaluated with truth-values in $\{t, f, u\}$, $\{t, f, b\}$, or $\{t, f, u, b\}$.
- (A2) The truth-value of a complex sentence at a scenario is determined by the truth-value of the sentence's parts at that scenario according to the truth-functions for the connectives as in K_3 , LP, or FDE, respectively.

In Section 4, we provide concrete examples of such scenarios and axiomatize the notion of synonymy that they provide. For now, let's consider whether this minimal change to possible world semantics can provide a notion of synonymy satisfying (P1).

The answer is no: Consider the two propositional sentences $\varphi := p$ and $\psi := p \vee (p \wedge q)$. Whatever set of truth-values one chooses ($\{t, f, u\}$, $\{t, f, b\}$, or $\{t, f, u, b\}$), the respective logic will evaluate φ and ψ to the same truth-value no matter how p and q were evaluated (this can easily be checked). Thus, there is no scenario in the adapted possible worlds semantics that can distinguish the two sentences, so they are synonymous according to this scenario semantics. However, the two sentences don't have the same atoms. Consequently, the adapted possible worlds semantics cannot provide a notion of synonymy satisfying (P1).

We can summarize the observations of this section as the following impossibility result. The natural principle about synonymy that is motivating any form of scenario- or circumstance-based semantics is the following.

- (P2) *Scenario respecting.* If there is no possible scenario or circumstance whatsoever in which two sentences differ in truth-value, then they are synonymous.

The above shows that there is no notion of scenario that satisfies all of (A1), (A2), (P1), and (P2).

This impossibility result shows that there is a real problem about finding out the logical laws governing synonymy: Whatever they are, we would like them to render the subject matter preservation principle (P1) true, but we also see that neither standard possible world semantics nor a conservatively refined scenario based semantics can yield such a notion of synonymy.

Thus, to get a subject matter preserving notion of synonymy we need to go a different way. There are many possibilities that have been suggested—they fall under the term *conceptivist logics*. However, they come in greatly varying intentions, semantic frameworks, and/or axiomatizations. This makes it hard to compare them and distill out the logical laws for synonymy. This is why we'll proceed in a different way: we look directly at possible sets of axioms about synonymy and compare these (e.g.,

⁵Note that the 'standardness' claim only concerns the truth-functions for negation, conjunction, and disjunction. It doesn't concern the truth-value of other connectives (like a conditional) or the correct choice of designated values to define validity, etc.

<p>(A1) $\varphi \equiv \neg\neg\varphi$</p> <p>(A2) $\varphi \equiv \varphi \wedge \varphi$</p> <p>(A3) $\varphi \wedge \psi \equiv \psi \wedge \varphi$</p> <p>(A4) $(\varphi \wedge \psi) \wedge \chi \equiv \varphi \wedge (\psi \wedge \chi)$</p> <p>(A5) $\varphi \equiv \varphi \vee \varphi$</p> <p>(A6) $\varphi \vee \psi \equiv \psi \vee \varphi$</p>	<p>(A7) $(\varphi \vee \psi) \vee \chi \equiv \varphi \vee (\psi \vee \chi)$</p> <p>(A8) $\neg(\varphi \wedge \psi) \equiv (\neg\varphi \vee \neg\psi)$</p> <p>(A9) $\neg(\varphi \vee \psi) \equiv (\neg\varphi \wedge \neg\psi)$</p> <p>(A10) $\varphi \wedge (\psi \vee \chi) \equiv (\varphi \wedge \psi) \vee (\varphi \wedge \chi)$</p> <p>(A11) $\varphi \vee (\psi \wedge \chi) \equiv (\varphi \vee \psi) \wedge (\varphi \vee \chi)$</p>
<p>(R1) $\frac{\varphi \equiv \psi}{\psi \equiv \varphi}$</p>	<p>(R3) $\frac{\varphi \equiv \psi}{\varphi \wedge \chi \equiv \psi \wedge \chi}$</p>
<p>(R2) $\frac{\varphi \equiv \psi \quad \psi \equiv \chi}{\varphi \equiv \chi}$</p>	<p>(R4) $\frac{\varphi \equiv \psi}{\varphi \vee \chi \equiv \psi \vee \chi}$</p>

Fig. 1 The system of analytic containment (AC) as presented in [23]

which is an extension of the other). This allows for direct comparison and avoids the commitments that come with opting for particular semantic systems.

We start in the next section by looking at a logic satisfying (P1) that recently gained much prominence. In Section 4, we develop and explore how close the scenario approach from this section comes to satisfying (P1). Against this background, we add and characterize further axiomatizations. Finally, in Section 9, we come full circle and both analyze and strengthen the impossibility result mentioned here.

3 Truthmaking Synonymy

We recall the logic of synonymy AC that satisfies (P1). Fine [23] provides a sound and complete truthmaker semantics for a system that aims to axiomatize the notion of analytic content containment: the content of a sentence φ is contained in the content of a sentence ψ and *vice versa*. The idea of the system originated in the work of Angell [1, 2].⁶ The system AC is given in Fig. 1.

By induction on derivations in AC, we immediately see that if $\text{AC} \vdash \varphi \equiv \psi$, then φ and ψ have the same atoms. Hence AC is a logic of synonymy that satisfies (P1).

As we will need it later on, we very briefly recap the truthmaker semantics of [23]. This semantics can be traced back to [27] and is defined as follows. A state model M is a triple $(S, \sqsubseteq, |\cdot|)$ where \sqsubseteq is a complete partial order⁷ on S and $|\cdot|$ maps atomic sentences p to pairs $(|p|^+, |p|^-)$ of non-empty subsets of S . (Elements of $|p|^+$ are called verifiers of p and elements of $|p|^-$ are called falsifiers of p .) We recursively define when a sentence φ is verified/falsified by a state $s \in S$ (in signs: $s \Vdash \varphi / s \dashv\vdash \varphi$).

⁶Also see [20, 23] for a good discussion of AC and similar systems and [28] for a sequent calculus for AC.

⁷So \sqsubseteq is reflexive, transitive, anti-symmetric, and every subset of S has a least upper bound.

- $s \Vdash p$:iff $s \in |p|^+$, and $s \dashv\vdash \varphi$:iff $s \in |p|^-$
- $s \Vdash \neg\varphi$:iff $s \dashv\vdash \varphi$, and $s \dashv\vdash \neg\varphi$:iff $s \Vdash \varphi$
- $s \Vdash \varphi \wedge \psi$:iff $\exists u, t \in S : u \Vdash \varphi \ \& \ t \Vdash \psi \ \& \ s = u \sqcup t$,⁸ and $s \dashv\vdash \varphi \wedge \psi$:iff $s \dashv\vdash \varphi$ or $s \dashv\vdash \psi$
- $s \Vdash \varphi \vee \psi$:iff $s \Vdash \varphi$ or $s \Vdash \psi$, and $s \dashv\vdash \varphi \vee \psi$:iff $\exists u, t \in S : u \dashv\vdash \varphi \ \& \ t \dashv\vdash \psi \ \& \ s = u \sqcup t$

The exact content of φ is $|\varphi| := \{s \in S : s \Vdash \varphi\}$, and the (replete) content of φ , denoted $[\varphi]$, is the convex closure of the complete closure of $|\varphi|$.⁹ We won't go into the philosophical difference between these two notions of content. For this see [23], and for the notion of synonymy induced by exact content see [25]. The soundness and completeness result of [23] states that $AC \vdash \varphi \equiv \psi$ if and only if $[\varphi] = [\psi]$ for all models.

4 Scenario Synonymy

To get a notion of synonymy satisfying (P2) and the two assumptions from Section 2, we formally describe scenarios and axiomatize “scenario synonymy”.

In Section 2, we straightforwardly relaxed possible worlds to scenarios in an attempt to obtain a semantics satisfying the subject matter preservation principle (P1). The idea behind such a scenario is that it is a (possibly inconsistent) representation of a part of the world or of a possible world. Before we present examples, let's first describe the two assumptions characterizing scenarios in more detail.

By (A1), every scenario s , whatever object it might be, determines a valuation $v_s : \{p_0, p_1, \dots\} \rightarrow T$, where T is one of the three possible sets of truth-values. Since we're interested in just how fine-grained this generalization of possible worlds semantics can become, we choose $T := \{t, fu, b\}$ (the other choices will be coarser or at best equally fine-grained). By (A2), this valuation extends to complex sentences according to FDE:

- $v_s \Vdash^+ p$:iff $v_s(p) \in \{t, b\}$, and $v_s \Vdash^- p$:iff $v_s(p) \in \{f, b\}$.
- $v_s \Vdash^+ \neg\varphi$:iff $v_s \Vdash^- \varphi$, and $v_s \Vdash^- \neg\varphi$:iff $v_s \Vdash^+ \varphi$.
- $v_s \Vdash^+ \varphi \wedge \psi$:iff $v_s \Vdash^+ \varphi$ and $v_s \Vdash^+ \psi$, and $v_s \Vdash^- \varphi \wedge \psi$:iff $v_s \Vdash^- \varphi$ or $v_s \Vdash^- \psi$.
- $v_s \Vdash^+ \varphi \vee \psi$:iff $v_s \Vdash^+ \varphi$ or $v_s \Vdash^+ \psi$, and $v_s \Vdash^- \varphi \vee \psi$:iff $v_s \Vdash^- \varphi$ and $v_s \Vdash^- \psi$.

Thus, we may define a model of *scenario semantics* as a pair (S, v) where S is a nonempty set whose elements are called scenarios and v is a function mapping scenarios $s \in S$ to valuations v_s . The positive content of a sentence φ is $[[\varphi]]^+ := \{s : v_s \Vdash^+ \varphi\}$ and the negative content is $[[\varphi]]^- := \{s : v_s \Vdash^- \varphi\}$. The content of φ is

⁸Where $u \sqcup t$ denotes the least upper bound of $\{u, t\}$ which exists by the completeness of the partial order.

⁹A set $T \subseteq S$ is convex if for all $s, t, u \in S$, if $s, t \in T$ and $s \sqsubseteq u \sqsubseteq t$, then $u \in T$. And a set $T \subseteq S$ is complete if for every nonempty $T_0 \subseteq T$ the least upper bound of T_0 is in T . The complete (resp. convex) closure of a set $T \subseteq S$ is the smallest set $T^* \subseteq S$ such that $T \subseteq T^*$ and T^* is complete (resp. convex).

$\llbracket \varphi \rrbracket := (\llbracket \varphi \rrbracket^+, \llbracket \varphi \rrbracket^-)$. Note that, while there is no restriction on (the set of) scenarios, their semantically relevant aspect is the valuation that they determine.

Let's consider two examples of such scenarios: First, if we want to give scenarios a metaphysical interpretation (also applicable in semantic paradoxes), we can interpret them as FDE-valuations or worlds [49, ch. 7–8].

Second, if we want to give scenarios a cognitive interpretation, we can take them to be sets of defeasible rules (i.e. rules that allow for exceptions) with which we conceptualize (a part of) the world. In other words, these rules are our knowledge base where we store both factual and semantic knowledge.¹⁰ Different knowledge bases correspond to different conceptualization (or representations). This explains what kind of object these scenarios are, but how do they determine a valuation? For this we use that such defeasible knowledge bases are paradigmatically modeled by programs in logic programming.¹¹ A logic program determines a three-valued model that is the canonical interpretation of the set of rules provided by the program [26]. Thus, a logic program s (modeling some knowledge base of a cognitive agent) determines a valuation (the intended model that the agent forms given her knowledge base).

This idea can be taken further. Logic programming has a neural interpretation: for every program there is a certain (artificial) neural network that computes the canonical interpretation of the program, that is, after starting in an initial state, the neural network will eventually reach a stable state that corresponds to the intended interpretation [55]. The idea is that, while the logic program is the symbolic representation of the knowledge base of the agent, the neural network is the (high-level) neural implementation of it. Each state of such a network can be described by a four-valued interpretation and could, roughly, be interpreted as how an agent currently cognizes the part of the world she is in. Thus, we can model knowledge bases as above as logic programs and take these as scenarios (determining the valuation given by their intended interpretation), but we can also model knowledge bases as neural networks and take their possible states as scenarios.¹² See [33] for elaborating this idea of grounding (cognitive) scenarios (or “worlds”) in states of neural networks (defining similarity between such scenarios, various modal operators, a counterfactual, and weaker versions of synonymy).

Let's now turn to the notion of synonymy provided by scenario semantics. Whatever instance we consider, scenarios satisfying the two assumptions *extensionally* act like FDE-models. (*Intensionally* they might be more complex than plain FDE-valuations; we discuss ways to intensionalize scenario semantics in Section 8.) We define the *canonical model of scenario semantics* as (S_c, v_c) where S_c is the set of

¹⁰To give some simple generic examples, factual knowledge might be the fact ‘water is H_2O ’ or the rule ‘if it rains, the street gets wet’. The rule is defeasible since there might be exceptions: e.g., a roof over the street. Semantic knowledge might be the rule ‘A sibling is a sister or a brother’.

¹¹Logic programming originated in the late 1960s and early 1970s [38]. It was developed precisely to deal with data- and knowledge bases in a defeasible way that is—unlike classical logic—more apt when dealing with real world data. Thus, logic programming also is very suitable to model cognitive and linguistic phenomena [7, 39, 55].

¹²Note that these cognitive examples of scenarios can but don't need to satisfy the second assumption on scenarios: they interpret atomic sentences but the logic for complex sentences can be chosen freely. In Section 8, we argue that, in fact, a more complicated logic than the straightforward one should be chosen.

valuations $v : \{p_0, p_1, \dots\} \rightarrow T$ and v_c is the identity function. Then the logic of synonymy obtained by scenario semantics can be characterized as follows (where $\varphi \Leftrightarrow_{\text{FDE}} \psi$ means that φ and ψ are equivalent in the logic FDE in the sense of having the same value under every valuation).

Theorem 1 *Let $\text{SF} := \text{AC} + \varphi \vee (\varphi \wedge \psi) \equiv \varphi$ (the name will be explained in Section 6). Then for all sentences φ and ψ , the following are equivalent*

- (i) $\llbracket \varphi \rrbracket = \llbracket \psi \rrbracket$ in any model of scenario semantics.
- (ii) $\llbracket \varphi \rrbracket = \llbracket \psi \rrbracket$ in the canonical model of scenario semantics.
- (iii) $\varphi \Leftrightarrow_{\text{FDE}} \psi$
- (iv) $\text{SF} \vdash \varphi \equiv \psi$.

Thus we can see that the two sentences φ and $\varphi \vee (\varphi \wedge \psi)$ that we used to show the inconsistency of the principles (P1), (P2), (A1), and (A2) really are the only source of the inconsistency. In Section 7, we present a possible argument in favor of scenario synonymy.

The first three items of the theorem are immediate, and the last one, which axiomatizes this logic of synonymy as an extension of AC, will be proven in the [Appendix](#).

5 Move-on Up: From Scenario to Truthmaking Synonymy

We show that truthmaking synonymy and scenario synonymy are related by moving up one set-theoretic level: moving from scenarios to sets of scenarios (as the entities at which sentences are evaluated) fine-grains scenario semantics to the level of truthmaking semantics.

By now, we know the proof-theoretic relationship between truthmaker synonymy (AC) and scenario synonymy ($\text{SF} = \text{AC} + \varphi \equiv \varphi \vee (\varphi \wedge \psi)$). We now want to see how the two are semantically related.

Analogous to the canonical model of [23], we define the *canonical scenario-based state model C* as follows.

Definition 1 In the canonical model of scenario semantics, the scenarios are the valuations which, in turn, correspond exactly to the subsets s of $L := \{p_0, \neg p_0, p_1, \neg p_1, \dots\}$.¹³ Thus, write $S := \mathcal{P}(L)$ (where \mathcal{P} is the powerset operator). Write $s_p := \{p\} \in S$, which is the scenario that makes p true and leaves anything else undetermined, and write $s_{\neg p} := \{\neg p\} \in S$. Define $C := (\mathcal{P}(S), \subseteq, |\cdot|_C)$, where $|p|_C := \{(V, F) \mid V = \{\{s_p\}\}, F = \{\{s_{\neg p}\}\}\}$. It's easily verified that C indeed is a state model.

It's straightforwardly checked that $[p \vee (p \wedge q)]_C = \{\{s_p\}, \{s_p, s_q\}\}$ and $[p]_C = \{\{s_p\}\}$. So the set of scenarios $\{s_p, s_q\} \in \mathcal{P}(S)$ verifies or makes true $p \vee (p \wedge q)$ but not p . Thus, if we move from scenarios making sentences true to sets of scenarios

¹³The correspondence between v and s is given by: $p \in s \ \& \ \neg p \notin s$ iff $v(p) = t$; $p \notin s \ \& \ \neg p \notin s$ iff $v(p) = u$; etc.

making sentences true, then we get a semantics that can distinguish between the scenario-synonymous sentences p and $p \vee (p \wedge q)$. In the [Appendix](#), we prove— analogously to [23]—the following result.

Corollary 1 *The following are equivalent:*

- (i) $\text{AC} \vdash \varphi \equiv \psi$.
- (ii) For every state model M , we have $[\varphi]_M = [\psi]_M$.
- (iii) $[\varphi]_C = [\psi]_C$ where C is the canonical scenario-based state model.

In other words, while scenario semantics can never get to the level of granularity achieved by truthmaking synonymy, any semantics that individuates content according to AC is equivalent—in terms of content individuation—to the sets-of-scenarios semantics.¹⁴ We'll next get to investigating the “space” between SF and AC.

6 Characterizing Benchmark Synonymies

So far, we've motivated and described two benchmarks, SF and AC, among the possible notions of synonymy. Now, we'll add more and thus get to the main formal result: We identify and uniformly axiomatize various logics of synonymy and characterize them by possible combinations of weak/strong subject matter preservation and weak/strong logical equivalence.

We first state the logics and the characterization, and then we'll put them into context. In the next section, we give arguments for the various benchmark notions of synonymy.

We first need some terminology. Recall that the set of atoms $At(\varphi)$ occurring in a sentence provides some information about the subject matter of the sentence. Now, we not only want to record which atoms occur in φ , but also whether they occur positively or negatively in φ . This is standardly done as follows.¹⁵

Definition 2 (Valence) The *valence* (positive or negative) of an atom p in a sentence φ is defined recursively by:

- (i) Only p occurs positively in p , and no atom occurs negatively in p .
- (ii) p occurs positively (negatively) in $\neg\varphi$ iff p occurs negatively (positively) in φ .
- (iii) p occurs positively (negatively) in $\varphi \circ \psi$ (for $\circ \in \{\wedge, \vee\}$) iff p occurs positively (negatively) in φ or ψ .

Note that p can occur in φ either not at all, or positively, or negatively, or both positively and negatively. We define

$$L(\varphi) := \{p : p \text{ occurs positively in } \varphi\} \cup \{\neg p : p \text{ occurs negatively in } \varphi\}.$$

¹⁴This yields an intuitive two-level picture of content where the dividing line lies somewhere between SF and AC. On the first granularity level, content can be modeled by first-order objects like scenarios or possible worlds. And on the second granularity level, content is modeled by second-order objects like sets of scenarios.

¹⁵See e.g. [23] or [21].

Thus, $L(\varphi)$ not only records which atoms occur in φ , but also whether they occur positively or negatively.

Next, we provide an equivalent perspective on valence that will be central to the proof of our characterization result. Again, we need some terminology. A *literal* is an atom or a negated atom. A sentence φ is in *disjunctive form* if it is a disjunction of conjunctions of literals. It is *standard* if the conjuncts and disjuncts are ordered according to a fixed order.¹⁶ Following [23], φ is *maximal* if whenever it contains a disjunct φ_0 and literal l (appearing as a conjunct of some disjunct), then it contains a disjunct φ_1 whose literals are exactly those of $\varphi_0 \wedge l$ (that is, φ_1 and $\varphi_0 \wedge l$ are identical modulo order and repeats). Fine [23] shows that every sentence φ is AC-provably equivalent to a unique standard maximal disjunctive form φ_{\max} .

Lemma 1 *For every sentence φ , $L(\varphi)$ is the set of literals that occur in φ_{\max} .*¹⁷

Thus, the definition of $L(\varphi)$ via valence or via the literals in φ_{\max} are equivalent. The former might be conceptually more useful (so we'll use it in the main text), while the later is more useful in proofs (so we'll use it in the [Appendix](#)).

Definition 3 The logics that we'll consider are axiomatized as follows.

$$\begin{aligned}
 \text{SFL} &:= \text{AC} \\
 \text{SF} &:= \text{AC} + \varphi \vee (\varphi \wedge \psi) \equiv \varphi \\
 \text{SFA} &:= \text{AC} + \varphi \vee (\varphi \wedge \psi) \equiv \varphi \vee (\varphi \wedge \neg\psi) \\
 \text{SCL} &:= \text{AC} + \varphi \equiv \varphi \vee (p \wedge \neg p) && \text{if } p, \neg p \in L(\varphi) \\
 \text{SCA} &:= \text{AC} + \varphi \equiv \varphi \vee (p \wedge \neg p) && \text{if } p \in \text{At}(\varphi)
 \end{aligned}$$

Their relationship is described in Fig. 2: all indicated containments are strict, and SF is incomparable to both SCA and SCL. We leave the search for the required counterexamples to the reader.

Their names come from the characterization theorem that we'll show next: each logic represents one combination of weak (Classical) or strong (FDE) logical equivalence and weak (Atomic) or strong (Literal) subject matter identity. So the first letter S stands for Synonymy, the second letter (F or C) abbreviates the kind of logical equivalence that the synonymy requires, and the third letter (L or A), if existing, abbreviates the kind of subject matter identity that the synonymy requires. This then explains the—perhaps surprising—axiomatization. (Recall, $\varphi \Leftrightarrow_{\text{FDE}} \psi$ denotes FDE-equivalence; similarly, $\varphi \Leftrightarrow_{\text{C}} \psi$ denotes classical equivalence.)

¹⁶More precisely, we fix a bijective enumeration e of literals, and we say that a conjunctive form $\varphi = l_1 \wedge \dots \wedge l_n$ is *standard* if its literals occur from left to right in e -increasing order without repeats, that is, $e(l_1) < \dots < e(l_n)$. We fix a bijective enumeration f of standard conjunctive forms, and we say that a disjunctive form $\varphi = \varphi_1 \vee \dots \vee \varphi_m$ is *standard* if its disjuncts are standard and f -increasing from left to right without repeats, that is, $f(\varphi_1) < \dots < f(\varphi_m)$.

¹⁷Proof: As [23] shows, AC-equivalence implies valence preservation, so, since $\text{AC} \vdash \varphi \equiv \varphi_{\max}$, we have for all p that p occurs positively (negatively) in φ iff p occurs positively (negatively) in φ_{\max} . The latter, by definition of valence and normal form, is equivalent to: p occurs in φ_{\max} ($\neg p$ occurs in φ_{\max}), and the claim follows.

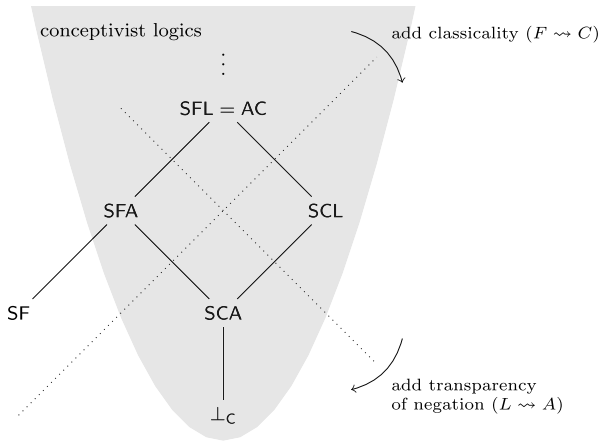


Fig. 2 The discussed benchmark logics of synonymy and the lattice of conceptivist logics

Theorem 2 (Characterization theorem) *We have for all sentences φ and ψ that*

- (i) $SF \vdash \varphi \equiv \psi$ iff $\varphi \Leftrightarrow_{FDE} \psi$
- (ii) $AC = SFL \vdash \varphi \equiv \psi$ iff $\varphi \Leftrightarrow_{FDE} \psi$ and $L(\varphi) = L(\psi)$
- (iii) $SFA \vdash \varphi \equiv \psi$ iff $\varphi \Leftrightarrow_{FDE} \psi$ and $At(\varphi) = At(\psi)$
- (iv) $SCL \vdash \varphi \equiv \psi$ iff $\varphi \Leftrightarrow_C \psi$ and $L(\varphi) = L(\psi)$
- (v) $SCA \vdash \varphi \equiv \psi$ iff $\varphi \Leftrightarrow_C \psi$ and $At(\varphi) = At(\psi)$.

We end this section with several remarks putting this result into context.

First, with scenario synonymy, we developed a logic of synonymy that exactly satisfies the principle (P2) under the two assumptions (A1) and (A2). To get a similar understanding for the class of logics satisfying (P1), we look at the newly added logics which all satisfy (P1). Such logics, where equivalence entails having the same atoms, have been called *conceptivist logics* by [53]. (For a short history and references see, e.g., [20].) The first such system has been suggested by [45, 47] to capture the idea that in many cases a valid inference shouldn't only preserve truth, but the conclusion should also be conceptually contained in the premises. The lattice of conceptivist equivalences is marked in gray in Fig. 2. The top element is the identity relation and the bottom element is the logic \perp_C where any two sentences with the same atoms are equivalent.¹⁸

Second, there are various related results. Part (i) is Theorem 1. Part (ii) is a characterization of AC proven by [21] and [23] formulated in terms of valence preservation. As we'll discuss below, we'll provide a novel proof. Concerning (iii), there are various logical systems whose notion of equivalence is characterized by SFA. One is

¹⁸Of course, this is not a serious notion of synonymy since, for example, the contradiction $p \wedge \neg p$ is synonymous to the tautology $p \vee \neg p$.

the first-degree fragment of the system of [16] with the “story semantics”. Another one is the formalization of Buddhist dialectics of [50] given by adding a fifth truth-value to the semantics of FDE that represents “emptiness” ([21] shows that equivalence in this system is FDE-equivalence plus having the same atoms). Concerning (iv), [52] shows that an equivalent characterization is that φ and ψ M_X -match in the category of classical proofs. We’ll come back to this below. Concerning (v), [53] shows that for all sentences φ and ψ that don’t contain the conditional \rightarrow , the sentence $\varphi \leftrightarrow \psi$ is provable in Parry’s logic of analytic containment iff φ and ψ are classically equivalent and have the same atoms.

Third, the proof of the theorem will take up most of the [Appendix](#). It defines, for each of the logics, a notion of a unique normal form and shows that two such forms are identical iff they have the properties of the characterization. This method is not only uniform across these logics, but it also is constructive since it provides an algorithm to decide equivalence. Moreover, it associates to each notion of synonymy a characteristic notion of normal form. In particular, the proof doesn’t require any (philosophically laden) ideas about semantics. (Because of these advantages we include our proof of (ii).)

Fourth, an intuitive conception of meaning takes the meaning of a sentence to consist of two components: a “truth component” that specifies the truth-conditions of the sentence, and an “aboutness component” that specifies what the sentence is about. Thus, synonymy according to this intuitive conception is given by “truth-based equivalence” plus “subject matter identity”. (Note how the two intuitive principles about synonymy mirror this distinction: (P2) is concerned with truth, while (P1) is concerned with subject matter.) Thus, the theorem provides the axiomatization for various choices of these two components.¹⁹

Fifth, this also provides an interpretation for the dotted lines in [Fig. 2](#). We get from logics above the dotted line labeled “add classicality” to the logics below by replacing the stricter truth-component of FDE-equivalence with the looser one of classical equivalence. Similarly, we get from logics above the “add transparency” line to the logics below by moving from having the same literals to the weaker requirement of having the same atoms. That is, we postulate that φ has the same subject matter as $\neg\varphi$. In other words, negation is subject matter transparent. This is a common principle if subject matter is understood as topic [31, 48].

Sixth, what’s beyond AC? There is, for example, isomorphism in the category of classical proofs which coincides with equivalence in multiplicative linear logic [17, 52], factual equivalence [15], equivalence in exact entailment [25], equivalence in some impossible world semantics [6], or syntactic identity. However, for reasons of space, we won’t further investigate those.

¹⁹In fact, not only our but also most other conceptivist logics eventually received an analysis in terms of a truth and a subject matter condition [20, 21]. This corroborates the two-component analysis of meaning: Any available notion of meaning identity that respects subject matter or topicality intuitions can be analyzed as having two components. (Of course this doesn’t mean that this is the only analysis as Priest’s formalization of Buddhist dialectics shows: it has a semantics with only truth-components but it also can be regarded under the two component analysis.)

7 Arguments for Various Notions of Synonymy

In this section, we present novel arguments for the various notions of synonymy discussed so far.

Argument for SF A reasonable notion of synonymy is to say that two sentences are synonymous if they, when uttered, always communicate the same thing. Denote this by \equiv_c . The argument claims that the intuitively correct notion of synonymy \equiv_c is, in fact, scenario synonymy.

Let's recap Hurford's constraint [34] which says that disjunctions where one disjunct entails the other are infelicitous—an often cited example is “John is American or Californian”.²⁰ Call such disjunctions *Hurford disjunctions*. An intuitive explanation for this constraint is the following. A Hurford disjunction $\varphi \vee \psi$ is equivalent to one of the disjuncts, say, φ . It's a pragmatic principle that you hence should utter the simpler φ and not the equivalent but more complicated $\varphi \vee \psi$. Hence $\varphi \vee \psi$ is infelicitous, while φ is not.

This has to be stated more carefully. The pragmatic principle can only be invoked when φ and $\varphi \vee \psi$ communicate the same thing: The principle works because you could have used either φ or $\varphi \vee \psi$ to communicate what you wanted to communicate, but you should choose φ over $\varphi \vee \psi$ on the grounds that φ is more concise. But if φ and $\varphi \vee \psi$ should communicate the same thing, the implication from ψ to φ needs to be “obvious”, because if the implication would require a very complicated classical logic computation, φ and $\varphi \vee \psi$ wouldn't communicate the same thing. Thus, the refined Hurford constraint with its explanation states the following.²¹ Assume $\varphi \vee \psi$ is a disjunction where ψ obviously implies φ . Then $\varphi \vee \psi$ is infelicitous because φ and $\varphi \vee \psi$ communicate the same thing, whence you should choose the simpler φ over the complicated $\varphi \vee \psi$.

Now for the argument. Assume the refined Hurford constraint with its explanation is correct. The sentence $\varphi \vee (\varphi \wedge \psi)$ is a disjunction where $(\varphi \wedge \psi)$ obviously implies φ (in the sense of the preceding paragraph) since it infers the conjunct from a conjunction. Hence, by assumption, φ and $\varphi \vee (\varphi \wedge \psi)$ need to communicate the same thing, whence $\varphi \equiv_c \varphi \vee (\varphi \wedge \psi)$. (So the Hurford constraint can be seen as a reason for coarse-graining synonymy.) Moreover, the axioms of AC represent necessary properties of \equiv_c , and there is no other reason to make more sentences \equiv_c -synonymous. Thus, \equiv_c is axiomatized by SF, and hence scenario synonymy is the intuitively correct notion of synonymy \equiv_c .

Argument for SFA Assume the two components conception of meaning, and assume the truth component is correctly spelled out as FDE-equivalence (motivated by scenarios). Moreover, assume that the aboutness component is given by some atom-based approach to subject matter which, as mentioned, is particularly well-suited if

²⁰For a discussion see e.g. [12, 14, 29, 36, 43, 54].

²¹ Another refinement (which is not relevant to the present argument) is that the weaker disjunct φ shouldn't get an exhaustive reading as $\varphi \wedge \neg\psi$ [12, 29]. Otherwise the disjunction is felicitous—an often cited example is “Mary read most or all of the books”.

subject matter is understood as topic [31]. That is, the subject matter of a complex sentence is given by merging the subject matter of the atoms of the sentence. For this notion of subject matter, the following is true.

(S2) *Soundness of syntactic reflection.* If two sentences have the same atomic sentences, they have the same subject matter.

Thus, by (S1), which we discussed in Section 2, two sentences have the same subject matter if and only if they have the same atoms. Hence, by the characterization theorem, synonymy in this intuitively correct conception of meaning is axiomatized by SFA.

Argument for AC The flip side of the argument just given is that subject matter identity has to be stronger than having the same atoms (which is in favor of AC): Assume one could argue that $\varphi \vee (\varphi \wedge \psi)$ is not synonymous to $\varphi \vee (\varphi \wedge \neg\psi)$ (as, for example, in truthmaker semantics). Then, by the characterization theorem, having the same atoms and FDE-equivalence are not enough for synonymy. Assume one sticks to the idea of two components semantics where truth-component identity shouldn't be spelled out in a way more fine-grained than FDE-equivalence. Then subject matter identity involves more than just having the same atoms (that is, the soundness of syntactic reflection (S2) fails). Since the two sentences differ only by a negation sign, this then seems to suggest that negation is not topic transparent, that is, that the subject matter of a statement and that of its negation may differ.²²

Argument for SCA We can take the same argument as for SFA but replace the choice for the truth-component: instead of FDE motivated by scenarios, take classical equivalence motivated as the conservative choice. Thus, SCA axiomatizes synonymy in the minimal modification of standard possible world semantics where synonymy entails subject matter identity.

Another way to bring this about is as follows. The semantics for weak Kleene logic uses three truth-values t, f, u , but u is interpreted as meaninglessness or off-topicness [3, 9]. Thus, when all atoms of a sentence φ have a classical truth-value (t or f), the whole sentence has a classical truth-value, but as soon as one atom of φ is u , the whole sentence is u . Hence two sentences φ and ψ are equivalent in weak Kleene logic iff they are classically equivalent and have the same atoms iff they are SCA-equivalent. So if we take weak Kleene to correctly describe reasoning preserving both truth and topic and adopt a two component view of meaning, then SCA is the correct logic of synonymy.

Taking up on footnote 21, a possible counterargument to SCA as the correct axiomatization of synonymy is as follows. The non-exhaustive reading of a Hurford disjunction is infelicitous, while the exhaustive reading is felicitous. So the two readings shouldn't be synonymous. In particular, $\varphi \vee (\varphi \wedge \psi)$ shouldn't be synonymous to $(\varphi \wedge \neg\psi) \vee (\varphi \wedge \psi)$. However, since these two sentences are classically equivalent and have the same atoms, they are equivalent according to SCA.

²²Unless conjunctions (or disjunctions) create opaque contexts which is not implausible as conjunction acts like a modal operator in truthmaker semantics (cf. the semantic clause for \wedge in Section 3).

Argument for AC and SCL An intuitive notion of synonymy is explanatory equivalence: any explanation for why φ is true can be “obviously” transformed into an explanation for why ψ is true. In other words, the two sentences don’t just agree on truth-conditions but also on how they are explained or proven. Inspired by the concept of a homotopy and the well-known idea that a logic yields a category where formulas are objects and proofs are morphisms, we may formalize this by two conditions: First, there is a proof from φ to ψ and a proof from ψ to φ (this allows transformation of an explanation for φ into one of ψ , and *vice versa*). Second, if we concatenate the proofs and move from φ to ψ and back to φ , we obtain a proof that is “essentially” like the identity proof that obtains φ from φ ; and similarly for ψ (this captures that the transformation is “obvious”).²³ Now, following [52], “essentially” can be spelled out in different ways: roughly, it could be exact identity or it could allow treating equally occurrences of the same atom (either all or only those of the same polarity). If we choose identity, the resulting notion of synonymy is that of isomorphism in the category of classical proofs as mentioned above [17, 52]. If we choose same-polarity occurrences, we get a notion of synonymy equivalent to AC [52]. If we choose all occurrences, we get a notion of synonymy with the same characterization as (and hence axiomatized by) SCL [52]. Thus, AC and SCL can be seen as two possible axiomatizations of explanatory synonymy.

8 Ways Out of the No-go Result and its Consequences

We mention ways out of the no-go result about synonymy from Section 2, and we state some of its consequences concerning a pluralistic conception of synonymy and the principle of compositionality in logic programs and neural networks.

Ways out We’ve started with two intuitive principles about synonymy in the strong sense of content identity: (P1) demanding synonymy to preserve having the same atoms, and (P2) demanding synonymy to respect scenarios. The two principles are inconsistent under the two assumptions (A1) and (A2) on scenarios (which a straightforward refinement of possible worlds semantics should satisfy). With the last sections, we can now understand this contradiction by identifying it and considering ways out. Indeed, we can precisely locate the inconsistency: Scenario synonymy exactly satisfies (P2), (A1), and (A2), and the instances of sentences violating (P1) can be traced back to exactly one axiom: $\varphi \equiv \varphi \vee (\psi \wedge \varphi)$.

So let’s wonder what modifications of (P1), (P2), (A1), or (A2) would render them consistent. We’ll consider three related options.

First, we weaken (P2): The conceptivist logics SCA, SCL, SFA, AC all satisfy (P1) and, as extensions of AC, they satisfy a weaker version of (P2): If there is no *set*

²³Figuratively speaking, the homotopy idea arises by thinking of proofs as paths through the space of formulas and demanding that the loop going from φ through ψ back to φ can be “continuously” deformed into the point φ (and similarly for ψ). Thus, various notions of continuity (i.e., of proof identity) correspond to various notions of “obviousness” and subject matter. We leave it to future research to explore these connections to homotopy theory.

of scenarios that verifies one sentence but not the other, then the sentences are synonymous (cf. Section 5).

Second, we keep (P1) and (P2) but change the notion of scenario whence modifying (A1) and/or (A2). This is done by most semantics for conceptivist logics. We consider three kinds of examples: (a) As mentioned, weak Kleene semantics uses three truth-values t, f, u —so validates (A1)—but u is interpreted as meaninglessness or off-topicalness. Hence conjunction and disjunction cannot be interpreted as lattice-like operations on the truth-values, whence (A2) is not satisfied.²⁴ (b) The truthmaker semantics for AC uses the four truth-values t, f, u, b —so, too, validates (A1)—but there can be states neither making φ nor ψ true while making $\varphi \wedge \psi$ true, whence (A2) is not satisfied. (c) The NC semantics for AC by [21] uses nine truth-values—so violates (A1)—but the truth-functions of the connectives on these truth-values is straightforward, whence it in spirit satisfies (A2).

Third, the general idea behind the preceding approach is to add more structure to the notion of a scenario and use this to make more fine-grained distinctions: in (a) the information what's on topic in a given scenario is added, in (b) mereological structure is added to the set of states, and in (c) more truth-values are added. Note an important difference: in (a) and (c) the scenarios themselves are 'locally' enriched by more structure while in (b) the set of scenarios is 'globally' enriched.

Let's consider more explicit ways of enriching scenarios and thus "intensionalize" scenario semantics. In Section 4, we described scenarios as representations of the world that extensionally act like FDE-models. So it could be that two scenarios determine the same valuation (extensionally identical) but they still differ in their representational or internal structure (intensionally different). Using this additional intensional structure, we could then weaken (P2) to demanding intensionally indistinguishable sentences to be synonymous. That is, if no scenario can find a difference between two sentences—neither based on its internal structure nor on the induced valuation—then they are synonymous. Here are three examples.

(a) As just indicated, a scenario might consist of a valuation plus a set of atomic sentences which are considered to be on topic in that scenario (or something similar to that effect). Then (intensional) scenario indistinguishability amounts to SFA.

(b) Two FDE-equivalent sentences φ and ψ might still be intensionally distinguished since the explanation for why φ is true according to the scenario or representation is different from why ψ is true.²⁵ Thus, two sentences are intensionally indistinguishable if they are FDE-equivalent and explanational synonymous as described in Section 6, whence this version of intensional scenario synonymy is axiomatized by AC or an even more fine-grained logic.

(c) Another, even more cognitive approach to weakening (P2) by intensionalizing scenario semantics is to read (P2) as "if no scenario can be imagined in which the two sentences differ, they are synonymous". Then scenario synonymy is moved from

²⁴Otherwise, since $t \vee u = u$ and $t \wedge u = u$, we'd have $u \leq t \leq u$, so $t = u$, contradiction.

²⁵Consider $\varphi = p$ and $\psi = p \vee (p \wedge q)$. Then the explanation for ψ might be: there is a fact p_0 that, according to the representation, implies $p \wedge q$, whence $p \vee (p \wedge q)$. And this is a different explanation from why φ is true: from the fact p_0 we get $p \wedge q$, whence p .

an extensional into an hyperintensional context, whence more is required for it to obtain. Using the logic of imagination of [4, 5], this can formally be spelled out as: for all sentences χ , $[\chi]\varphi$ is equivalent to $[\chi]\psi$. This roughly means that whenever we imagine a scenario constructed around making χ true, if φ will turn out to be true there, too, then also ψ will be true (and *vice versa*). What is the resulting notion of synonymy? If only possible worlds are used [5], it coincides with SCA (given same content implies having the same atoms). If also non-normal worlds are allowed [4], it depends on the assumptions about imagination but the resulting notion of synonymy will generally be very fine-grained. Thus, in line with the previous version, this version of intensionalized scenario synonymy, too, will be rather high up in the lattice of logics of synonymy.

Let's turn to the consequences of our results and the paradoxical nature of synonymy.

Pluralism According to pluralism about meaning, there are various, equally justified notions of meaning and thus logics of synonymy. Here are two ways how our results point in this direction.

First, assume that, despite the attempted reconciliation, it is not an option to give up on any of (P1), (P2), (A1), and (A2) in their original, non-weakened formulations. Say, because one is convinced that (P1) is a necessary truth and (P2) together with the two assumptions is key to how we evaluate complex sentences and use scenarios in our thinking (e.g. in counterfactual reasoning). Then, scenario synonymy and one of the conceptivist logics are equally justified notions of synonymy. Our intuitions in favor of the principles then come from distinct notions of meaning.

Second, and more importantly, we've seen several arguments motivating distinct notions of synonymy. That is, as soon as at least two of these arguments are accepted, pluralism holds. The characterization theorem associates these notions of synonymy to different choices of logic and subject matter. For example, the 'communicational synonymy' in the argument for SF is distinct from the 'explanational synonymy' in the argument for AC and SCL. That there is no universally correct choice can be made plausible by the fact that different domains of reasoning might demand different choices. For example, in some domains, maybe metaphysics or mathematics, it might be plausible to adopt a classical logic, while in other domains, maybe cognition or databases, it might be plausible to go four-valued. And in some domains, maybe topics of a discourse, it might be plausible that negation is transparent and hence same atoms should mean same subject matter, while in other domains, maybe epistemic notions of content, it might be that having the same literals is the right choice for tracking subject matter identity.

Non-compositionality Another consequence of our results is that synonymy cannot be spelled out via scenarios in what may be called a *straightforwardly compositional* way: Assume we take the common starting point and work with scenarios that assign (at most) four truth-values to atomic sentences—that is, we satisfy (A1). If we then build a semantics for complex sentences with these scenarios that satisfies the principles of synonymy (P1) and (P2), then (A2) fails: The truth-values of complex

sentences at scenarios are *not* determined by the truth-value of the sentences' parts at that scenario according to the straightforward truth-functions for the connectives.

Thus, there is a sense—namely that of violating (A2)—in which synonymy or content identity is not compositional. Let's now discuss the value of this insight.

On a positive note, this can be seen as an extension of a result from [40] showing that synonymy in the weaker sense of meaning similarity or resemblance is not compositional.

Moreover, we see that the above abstractly described failure of (A2) is found in many concrete examples: This applies, for instance, to the examples of scenarios mentioned in Section 4, that is, logic programs and their implementation in appropriate neural networks. But it also applies much more generally when we assign—e.g., in the quest for explainable artificial intelligence—to each activation state (or weight-setting) of a given neural network some human-interpretable atomic properties (e.g., “in this this activation state, the network recognizes a dog” or “in this weight-setting, the network accurately identifies stop signs”). In all these examples, we assign at most four truth-values to atoms, but to fully query logic programs or explain the network, we, at the very least, also need to understand when complex properties (formed from the atomic ones) are true according to a program or state. For this, it is natural to demand that programs and states can individuate complex properties built from different atomic properties: For example, when we consider the two complex properties ‘there is a dog’ and ‘there is a dog or there is both a dog and a cat’, we'd intuitively expect that in an activation state with the latter property the concept ‘cat’ is somehow “present”, while this is not necessary for states with the former property, whence the two properties intuitively should be distinguishable by some state. But then the notion of truth according to a logic program or (activation or weight) state of a neural network cannot be straightforwardly compositional: we cannot just use the straightforward truth-functions to determine the truth-value of a complex sentence. For example, to determine the truth-value of a conjunction we might have to take other states and programs into account—rendering ‘and’ a modal operator as in truth-maker semantics (also see the discussion at the end of Section 9). This is, to the best of our knowledge, a fresh perspective—that seems worth pursuing—on the correct logic and semantics for complex properties of logic programs and network states.

On a negative note, we may wonder: what good is the insight that there is no straightforwardly compositional semantics if we could still find a compositional semantics that is, if at all, only slightly less straightforward but satisfies (A1), (P1) and (P2). The following two examples exemplify two ways how this can be achieved. However, they show that the resulting compositionality is different from a straightforwardly compositional semantics in that it is not, what we may call, ‘purely extensional’.

First, assume scenarios use the three truth-values t , f , u and as truth-functions we don't use the straightforward ones of strong Kleene logic but rather the, if at all, slightly less straightforward ones of weak Kleene logic. As discussed above, this still satisfies (A1), (P1) and (P2), however, now the truth-values aren't ‘pure’ truth-values anymore but rather ‘composites’ of truth and topic (t is interpreted as true and on topic, f is interpreted as false and on topic, and u is interpreted as off topic). So, really, we added additional intensional structure—namely, topic—to the scenarios

and used this to satisfy the principles of synonymy. As described above, this is a promising way out of the inconsistency, but compositionality is restored by (covertly) intensional means. The semantics is not ‘purely extensionally’ compositional in the following sense: it uses extensional truth-functions but it interprets the truth-values as truth-topic composites which provide intensional structure.

Second, consider truthmaker semantics. Scenarios assign four truth-values to atomic sentences, but the truth-value of a complex sentence φ at a scenario s is not determined by a truth-function and the truth-values of φ ’s parts at s anymore (as it still was the case in the preceding example). Rather, the truth-value of complex sentences is determined modally, that is, by also taking into account scenarios other than s .²⁶ So, again as described above, this is a promising way out of the inconsistency, but compositionality is restored by the intensional means of using a modal semantics for conjunction and disjunction. Again the semantics is not ‘purely extensionally’ compositional: the truth-values are ‘pure’ but the connectives are modal and not extensional.

9 Impossibility Results for Subject Matter Preserving Synonymy

As just seen, a likely conclusion of the no-go result is that truth at scenarios is not straightforwardly compositional: the truth-value of a formula at a scenario cannot be determined from the truth-value of its atoms at the scenario alone in the straightforward way. We first analyze and then generalize this to an impossibility result for synonymy.

Analysis So what feature of a straightforwardly compositional semantics is the reason for the violation of the subject matter preservation principle? There are two evident potential reasons.

- (1) To determine the truth-value of a formula at a scenario one only needs to consider that scenario and no other scenarios. That is, the connectives have an extensional, non-modal semantics.²⁷
- (2) The truth-values can sensibly be ordered and conjunction and disjunction respect this order: the truth-value of a conjunction is the minimum of the truth-values of the conjuncts, and similarly the maximum for disjunctions.²⁸

We now formulate these ideas more formally, and then we argue that—maybe contrary to one’s first guess—(1) is not the real reason, but (2) is.

²⁶For example, the truth-value of $\varphi \wedge \psi$ at s is in $\{t, b\}$ if there are scenarios s' and s'' such that $s = s' \sqcup s''$ and the truth-value of φ at s' is in $\{t, b\}$ and the truth-value of ψ at s'' is in $\{t, b\}$.

²⁷This is what the logics used in (A2) satisfy (they have a straightforwardly compositional semantics). As already noted, truthmaker semantics doesn’t satisfy this (and it cannot be given a straightforwardly compositional semantics).

²⁸Again, this is what the logics used in (A2) satisfy (they have a straightforwardly compositional semantics). As already noted, weak Kleene logic doesn’t satisfy this (and it cannot be given a straightforwardly compositional semantics). Coming back to the weak Kleene logic example at the end of the previous section, one could speculate that an important feature of a ‘pure’ truth-value—as opposed to the ‘truth-topic composites’—is that it makes sense to order ‘pure’ truth-values.

Let’s define a simple but general framework to formulate the discussion. We work in the general semantic framework where formulas (here built from \neg, \wedge, \vee) get truth-values at states. So a *semantic model* is a triple $M = (S_M, T_M, V_M)$ where S_M is a non-empty set of states (with possible further structure) and T_M is a non-empty set of truth-values (with possible further structure), and for every formula φ and state $s \in S_M$, the semantic model determines the truth-value $V_M(\varphi, s) \in T_M$. A *semantics* \mathcal{S} is a non-empty class of semantic models. For example, truthmaker semantics is the class of semantic models (S_M, T_M, V_M) where S_M is a state space (in the sense of truthmaker semantics), T_M is the usual set of four truth-values, and V_M is recursively given by the clauses of truthmaker semantics.

We define a semantics \mathcal{S} to be *extensionally compositional* if, for all $M = (S_M, T_M, V_M) \in \mathcal{S}$ there is a function f such that for all formulas φ with atoms p_1, \dots, p_n and all states $s \in S_M$,

$$V_M(\varphi, s) = f(\varphi, V_M(p_1, s), \dots, V_M(p_n, s)).$$

This is the formalization of (1): The truth-value of φ at s is determined by the truth-value of φ ’s atoms at s , and no other state $s' \neq s$ is required for this.²⁹ (We might additionally demand that f can be chosen uniformly, but this weaker version will already be sufficient.)

We say that a semantics \mathcal{S} is *conjunction and disjunction conservative* if, for all $M = (S_M, T_M, V_M) \in \mathcal{S}$, the set T_M is a lattice (with operations \wedge and \vee) and for all states $s \in S_M$,

$$\begin{aligned} V_M(\varphi \wedge \psi, s) &= V_M(\varphi, s) \wedge V_M(\psi, s) \quad \text{and} \\ V_M(\varphi \vee \psi, s) &= V_M(\varphi, s) \vee V_M(\psi, s). \end{aligned}$$

This is the formalization of (2): Conjunction and disjunction of the language get interpreted by the corresponding standard functions on the truth-values.

Finally, we say that a semantics \mathcal{S} satisfies *subject matter preservation*, if for all formulas φ and ψ ,

$$\left(\forall M \in \mathcal{S} \forall s \in S_M : V_M(\varphi, s) = V_M(\psi, s) \right) \text{ implies } At(\varphi) = At(\psi).$$

We now show that (1) is not the reason for the inconsistency: We show that an extensionally compositional semantics doesn’t necessarily violate subject matter preservation (which it had to, if it were the reason). (This is unlike scenario semantics which is extensionally compositional and violates subject matter preservation.) Indeed, consider the trivial semantics consisting of only one semantic model $M_0 = (S_0, T_0, V_0)$ where S_0 consists of just one state s_0 , T_0 is simply the set of formulas, and $V_0(\varphi, s_0) = \varphi$. It is easily seen that this semantics is extensionally compositional and doesn’t violate (i.e., satisfies) the subject matter preservation. Of course, this semantics is no “serious” semantics, but it shows that being extensionally compositional doesn’t imply a violation of subject matter preservation.

Next we show that (2) is a reason for the inconsistency: We claim that synonymy in a conjunction and disjunction conservative semantics \mathcal{S} does not preserve subject

²⁹See [32] for a good discussion of compositionality at such a general level.

matter. Indeed, by the absorption laws that any lattice satisfies,³⁰ we have for any $M \in \mathcal{S}$ and $s \in S_M$,

$$V_M(p \vee (p \wedge q), s) = V_M(p, s) \vee (V_M(p, s) \wedge V_M(q, s)) = V_M(p, s),$$

but $At(p \vee (p \wedge q)) \neq At(p)$, so \mathcal{S} doesn't preserve subject matter.

Generalization In fact, a much weaker condition is enough to violate subject matter preservation. We say a semantics \mathcal{S} is *weakly absorptive* if for all formulas φ and ψ and all $M \in \mathcal{S}$ we have for all $s \in S_M$,

$$V_M(\varphi \wedge (\varphi \vee \psi), s) = V_M(\varphi \vee (\varphi \wedge \psi), s).$$

The name is to indicate that this is a weak version of the just mentioned absorption law of lattices. As lattice (-like) structures are so fundamental to logic, it's not surprising that most logics validate weak absorption. For example, even exact truthmaker semantics, which is much higher up than AC in the lattice of conceptivist synonymies, still is weakly absorptive.³¹ Hence all conceptivist logics containing AC—which includes all synonymies discussed so far—satisfy weak absorption.³²

We say a semantics \mathcal{S} is *order conservative*, if, for all $M \in \mathcal{S}$, the set of truth-values T_M has a partial order \leq and for all φ and ψ ,

$$V_M(\varphi \wedge \psi, s) \leq V_M(\varphi, s) \text{ and } V_M(\varphi, s) \leq V_M(\varphi \vee \psi, s).$$

Note that this is much weaker than being conjunction and disjunction conservative, since no lattice structure on the set of truth-values is assumed, let alone a homomorphism from formulas to truth-values. So this can be considered as a much weaker version of (2).

Still, these two weak properties yield the following simple but far-reaching impossibility result.

Theorem 3 (Impossibility result) *Synonymy in a weakly absorptive and order conservative semantics \mathcal{S} does not preserve subject matter.*

Proof For any $M \in \mathcal{S}$ and $s \in S_M$ we have by order conservativity:

$$V_M(p \wedge (p \vee q), s) \leq V_M(p, s) \leq V_M(p \vee (p \wedge q), s).$$

By weak absorption, $V_M(p \wedge (p \vee q), s) = V_M(p \vee (p \wedge q), s)$, whence $V_M(p \wedge (p \vee q), s) = V_M(p, s)$. But this violates subject matter preservation. \square

³⁰That is, $a \vee (a \wedge b) = a = a \wedge (a \vee b)$ for all elements a and b of the lattice.

³¹It is readily shown that for any state s in a state model M , we have $s \Vdash \varphi \wedge (\varphi \vee \psi)$ iff $s \Vdash \varphi \vee (\varphi \wedge \psi)$. (Contrast this with the fact that exact truthmaking doesn't satisfy the distributivity law [25].)

³² One logic that provides an exception is multiplicative linear logic: Conjunction (commonly written \otimes) can intuitively be interpreted, roughly, as there being enough resources to realize any of the disjuncts. Disjunction (commonly written \wp) can intuitively be interpreted, roughly, as there being enough resources to realize any of the disjuncts. So $\varphi \otimes (\varphi \wp \psi)$ requires at least two times φ , while $\varphi \wp (\varphi \otimes \psi)$ only requires at least one φ and one ψ (which is less if ψ requires less than φ). As mentioned, [17, 52] show that equivalence in multiplicative linear logic coincides with isomorphism in the category of classical proofs, and, in fact, [52] even explicitly mentions that $\varphi \wedge (\varphi \vee \psi)$ is not isomorphic to $\varphi \vee (\varphi \wedge \psi)$.

Thus, any subject matter preserving semantics either fails to be weakly absorptive or order conservative (or both). Let's consider both options.

Giving up order conservativeness: For example, truthmaker semantics (either exact or as for AC) is a subject matter preserving semantics and it indeed fails to be order preserving.³³ In contrast, as noted, any conceptivist logic containing AC (or exact equivalence) satisfies weak absorption. So any subject matter preserving semantics for these conceptivist logics is not order conservative, whence doesn't respect an essential property of conjunction and disjunction. This corollary can be summarized in the following slogan: Any subject matter preserving semantics that makes synonymous the few sentences dictated by exact equivalence, violates our intuitions about conjunction and disjunction.

Giving up weak absorption: If the impossibility result forces us to choose between giving up weak absorption or changing an essential property of conjunction and disjunction, it seems natural to, at the very least, consider giving up weak absorption (in contrast to what most logics do). Indeed, we'll present some linguistic and cognitive evidence against weak absorption.

Concerning the linguistic evidence, consider the following two sentences adapted from the famous Linda problem [57]:

- (a) Linda is a bank teller *or* (Linda is a bank teller *and* an activist).
- (b) Linda is a bank teller *and* (Linda is a bank teller *or* an activist).

Albeit a Hurford disjunction (cf. Section 4), sentence (a) seems to be a legitimate sentence in the context of the "Linda experiment" [57]: participants might ponder about it when they judge which of the two disjuncts is more likely. In contrast, sentence (b) is pragmatically very ill-behaved: it first makes a claim (that Linda is a bank teller) and then makes a second claim which is weaker than the first. Pragmatically, this doesn't make sense: a hearer will think that either the speaker shouldn't have made the first claim since it was false, or that the speaker shouldn't make the weaker claim to not be redundant. So it seems that while (a) is sometimes pragmatically legitimate, (b) never is, even though these two sentences are synonymous if weak absorption holds.³⁴

Here is some cognitive evidence against weak absorption. Very roughly, our cognition recognizes conjunctive features in a serial manner and disjunctive features in a parallel manner [56]. This suggests that $\varphi \wedge (\varphi \vee \psi)$ and $\varphi \vee (\varphi \wedge \psi)$ should play different cognitive roles: Assume we're presented with a few objects one of which has property $P(x)$. If we're asked whether there is an object x with the feature

³³A state can (exactly) make true a conjunction without (exactly) making true any of the conjuncts: consider three states $s \not\models p, q$ and $s' \models p$ and $s'' \models q$ such that $s = s' \sqcup s''$, whence $s \models p \wedge q$. So $\{1, b\} \ni V(p \wedge q, s) \not\subseteq V(p) \in \{0, n\}$.

³⁴Note that we only claim that the concrete sentences (a) and (b) provide counter-evidence to weak absorption. We don't claim that the following general principle is true: If φ and ψ are synonymous (in a given sense), then if φ makes sense in some context, also ψ makes sense in this context. If sentences (a) and (b) are inserted to this principle, it could of course be used (contrapositively) to argue against weak absorption. Though, it is doubtful whether this general principle is true. I'm grateful to an anonymous referee who pointed out that φ is synonymous to $\varphi \vee \varphi$ (even on very fine-grained notions of synonymy) but 'Which is true: φ or φ ?' makes sense while 'Which is true: φ ?' doesn't. We leave it as an open question to figure out which refined version of this principle can be used as an argument schema against acclaimed synonymies.

$P(x) \wedge (P(x) \vee Q(x))$, then the prediction is that we should—by default, i.e., without further reflecting on the question—serially go through the objects and for each check whether it has $P(x)$ and $P(x) \vee Q(x)$. If, on the other hand, we're asked whether there is an object x with the feature $P(x) \vee (P(x) \wedge Q(x))$, we would—by default—scan the scene in parallel and find the object with $P(x)$ from which we'd immediately conclude that $P(x) \vee (P(x) \wedge Q(x))$ holds since the first disjunct was confirmed. Thus, in this exceedingly idealized setting, the response times for these two sentences should be different, although they are equivalent according to weak absorption.³⁵

This poses some interesting, though highly speculative further questions. On the neural level, conjunctive features are realized by binding [19]: very roughly, a state where the neural network cognized that an object has two features P and Q somehow contains two parts—one signaling P and one signaling Q —that are bound together. So conjunction is much like the modal truthmaker conjunctions. However, since truthmaker semantics is weakly absorptive, this suggests that binding is more complicated than just 'merging' two states. What is this additional structure on the state space of the network? Or does disjunction also behave in a more complicated way—as, e.g., via a closure operator [51, ch. 12]?

We leave it to future research to answer these questions and to develop a logic that can account for the above linguistic and cognitive intuitions against weak absorption. (In light of footnote 32, multiplicative linear logic seems like a promising starting point.) If the outcome is order conservative and subject matter preserving, it would be a novel answer to the impossibility result.

Appendix : Proofs

In this appendix, we prove the theorems stated in the main text which essentially amounts to proving the characterization theorem (Theorem 2).

All our proofs are elementary. As mentioned, a methodological novelty is that we extend a technique used by [22, 23]: proving completeness results by developing an appropriate notion of a disjunctive normal form. That is, the idea of the proof is as follows: For each notion of synonymy, we find a corresponding notion of normal form that is (i) provably equivalent according to the synonymy and (ii) if two such forms satisfy the two characterizing properties, they are identical (modulo the order of literals and disjuncts). The theorems then follow: soundness is easy, and for completeness we move to the normal form of the two given sentences with the two properties; these normal forms hence have to be the identical, whence the original sentences are provably equivalent. The unifying, constructive, and theory-independent character of this proof has been discussed in Section 6.

³⁵ Although, given the special logical character of the sentences, it is not clear whether the "weirdness" of the sentence overrides the default response. Also note that, if at all, this can only work for sentences with small syntactic complexity: It couldn't be used to check whether or not two sentences involving, say, 100 disjunctions and/or conjunctions should be equivalent. For otherwise we wouldn't be able to immediately see whether we should go into 'conjunction/serial' mode or 'disjunction/parallel' mode when checking the sentence. This is analogous to why conjunctive/disjunctive feature detection also only works when presented with a small number of objects.

Recall Definition 3 collecting the logics that we'll be working with. If \mathcal{L} is one of these logic obtained from AC by adding an axiom $\varphi \equiv \psi$, then we refer to $\varphi \equiv \psi$ as the \mathcal{L} -axiom. Also note, by construction (see Definition 2), if φ is in disjunctive form, then $L(\varphi) = \{l : l \text{ literal in } \varphi\}$. So, by Lemma 1, for any φ , $L(\varphi) = L(\varphi_{\max})$. Thus, we can—as we find it more convenient in this appendix—work with the conception of $L(\varphi)$ as the set of literals of φ when φ is a disjunctive form, or as the set of literals of φ_{\max} when φ is arbitrary.

Step 1: Disjunctive Normal Forms

In this section, we do the first step: providing provably equivalent notions of normal form. In Section 6, we defined standard disjunctive normal forms. Now we define such a normal form for each logic. (Item (i) is due to [23].)

Definition 4 (Disjunctive forms) Let $\varphi = \varphi_1 \vee \dots \vee \varphi_n$ be in standard disjunctive normal form. We say

- (i) φ is *maximal* if for all $i \in \{1, \dots, n\}$, and for all literals l occurring in φ , there is a $j \in \{1, \dots, n\}$ such that $L(\varphi_i \wedge l) = L(\varphi_j)$. In other words, $\varphi_i \wedge l$ and φ_j have the same literals, that is, they are identical modulo order and repeats of the literals. (Roughly, this means that the set of disjuncts of φ is closed under adding literals of φ .) This will be the normal form of AC.
- (ii) φ is *minimal* if for all $i, j \in \{1, \dots, n\}$ with $i \neq j$,

$$L(\varphi_i) \not\subseteq L(\varphi_j) \text{ and } L(\varphi_j) \not\subseteq L(\varphi_i).$$

This will be the normal form of SF.

- (iii) φ is *maximal positive* if
 - (a) For every disjunct φ_i of φ , there is an $A \subseteq At(\varphi)$ and a minimal disjunct φ_0 of φ (i.e., there is no disjunct φ'_0 of φ such that $L(\varphi'_0) \subsetneq L(\varphi_0)$) such that $L(\varphi_i) = L(\varphi_0) \cup A$, and
 - (b) if φ_i is a disjunct of φ and $p \in At(\varphi)$, then $\varphi_i \wedge p$ is a disjunct of φ (modulo the order of the literals).

This will be the normal form of SFA.

- (iv) φ is *maximal literal-contradiction closed* if φ is maximal and if $p, \neg p \in L(\varphi)$, then $p \wedge \neg p$ is a disjunct of φ . This will be the normal form of SCL.
- (v) φ is *maximal atom-contradiction closed* if φ is maximal and if $p \in At(\varphi)$, then $p \wedge \neg p$ is a disjunct of φ . This will be the normal form of SCA.

As mentioned, [23] shows that every formula φ is AC-provably equivalent to a standard maximal disjunctive normal form φ_{\max} . We show the analogue for the new normal forms and extensions of AC. For this, we'll need the following replacement rule.

Lemma 2 (Replacement) For $\mathcal{C} \in \{\text{AC, SF, SFA}\}$, the following rule is \mathcal{C} -admissible, that is, if the premise is \mathcal{C} -derivable, then the conclusion is \mathcal{C} -derivable.

(When $\chi[\varphi]$ is a formula containing occurrences of φ , then $\chi[\psi]$ is the result of replacing all occurrences of φ by ψ .)

$$\frac{\varphi \equiv \psi}{\chi[\varphi] \equiv \chi[\psi]} (R)$$

Proof Most of the work has been done in [23]. It suffices to show that the following two rules are admissible.

$$\frac{\varphi \equiv \psi}{\chi[\varphi] \equiv \chi[\psi]} (PR) \qquad \frac{\varphi \equiv \psi}{\neg\varphi \equiv \neg\psi} (NR)$$

where in (PR) the occurrences of φ in $\chi[\varphi]$ are not in the scope of \neg . The admissibility of (PR) is shown for AC by [23] and the proof also works for SF and SFA. For (NR), the proof is by induction on the proof of $\varphi \equiv \psi$. All the cases corresponding to the axioms and rules of AC are dealt with in [23]. So we only need to consider the cases where $\varphi \equiv \psi$ is $\varphi_0 \vee (\varphi_0 \wedge \psi_0) \equiv \varphi_0$ or $\varphi_0 \vee (\varphi_0 \wedge \psi_0) \equiv \varphi_0 \vee (\varphi_0 \wedge \neg\psi_0)$. In these two cases we have to show that $\neg\varphi \equiv \neg\psi$ is derivable. Indeed, it is easy to check that we have

$$\begin{aligned} SF \vdash \neg(\varphi_0 \vee (\varphi_0 \wedge \psi_0)) &\equiv \neg\varphi_0 \\ SFA \vdash \neg(\varphi_0 \vee (\varphi_0 \wedge \psi_0)) &\equiv \neg(\varphi_0 \vee (\varphi_0 \wedge \neg\psi_0)) \end{aligned}$$

using the distributivity and de Morgan axioms. □

This also holds for SCL and SCA, but this will follow immediately from the characterization theorem for these logics and we won't need replacement for these logics in the proof. This is why we don't prove it directly here, although it's not too hard either.

Proposition 1 (Normal form for SF) *Every formula φ is SF-provably equivalent to a standard minimal disjunctive normal form φ_{\min} .*

Proof As mentioned, there is a formula φ' in standard disjunctive normal form that is AC-provably equivalent to φ and hence in particular SF-provably equivalent.

Next, we can delete—while preserving SF-provability—any disjunct φ_j occurring in φ' if there already is a disjunct φ_i in φ' with $L(\varphi_i) \subseteq L(\varphi_j)$. This is because if there are such φ_j and φ_i , then, without loss of generality, $\varphi_j = \varphi_i \wedge \chi$ and (using the underlining to increase readability)

$$\begin{aligned} \varphi' &= \varphi_1 \vee \dots \vee \underline{\varphi_i} \vee \dots \vee \underline{\varphi_j} \vee \dots \vee \varphi_n \\ &= \varphi_1 \vee \dots \vee \varphi_{i-1} \vee \underline{\varphi_i} \vee \varphi_{i+1} \vee \dots \vee \varphi_{j-1} \vee (\varphi_i \wedge \chi) \vee \varphi_{j+1} \vee \dots \vee \varphi_n \\ &\equiv \varphi_i \vee (\varphi_i \wedge \chi) \vee \varphi_1 \vee \dots \vee \varphi_{i-1} \vee \varphi_{i+1} \vee \dots \vee \varphi_{j-1} \vee \varphi_{j+1} \vee \dots \vee \varphi_n \\ &\equiv \underline{\varphi_i} \vee \varphi_1 \vee \dots \vee \varphi_{i-1} \vee \varphi_{i+1} \vee \dots \vee \varphi_{j-1} \vee \varphi_{j+1} \vee \dots \vee \varphi_n \\ &\equiv \varphi_1 \vee \dots \vee \underline{\varphi_i} \vee \dots \vee \varphi_{j-1} \vee \varphi_{j+1} \vee \dots \vee \varphi_n, \end{aligned}$$

where we essentially used commutativity and the axiom $\varphi \vee (\varphi \wedge \psi) \equiv \varphi$. Thus, we can reduce φ' to a provably equivalent formula φ^* in minimal disjunctive form.

Finally, by commutativity, associativity, and idempotence we can reorder φ^* to make it standard (without changing minimality). Thus, we get a formula φ_{\min} that is provably equivalent to φ and in standard minimal disjunctive form. \square

Proposition 2 (Normal form for SFA) *Every formula φ is SFA-provably equivalent to a standard maximal positive disjunctive normal form φ_{pos} .*

Proof As mentioned, φ is AC-provably (and hence SFA-provably) equivalent to a formula φ_{max} in maximal disjunctive normal form. Let $\varphi_1, \dots, \varphi_r$ be the minimal disjuncts of φ_{max} . Then every disjunct φ' of φ_{max} is of the form $\varphi' = \varphi_i \wedge L$ (modulo ordering) for an $i \leq r$ and a (possibly empty) set L of literals occurring in φ . By using replacement (Lemma 2), the SFA-axiom, and idempotence several times, we can SFA-provably replace each $\varphi_i \vee \varphi'$ by $\varphi_i \vee (\varphi_i \wedge \text{At}(L))$ and thus end up with a formula φ^* that still is SFA-provably equivalent to φ .³⁶ Clearly, φ^* satisfies (a), and it also satisfies (b): Let φ' be a disjunct of φ^* and $p \in \text{At}(\varphi^*)$. Then $\varphi' = \varphi_i \wedge \text{At}(L)$ for an $i \leq r$ and a set L of literals occurring in φ_{max} , and p occurs in a literal l_p of φ_{max} (since $\text{At}(\varphi^*) = \text{At}(\varphi_{\text{max}})$). Then, by the maximality of φ_{max} , $\varphi_i \wedge (L \cup \{l_p\})$ is (modulo order) a disjunct of φ_{max} . By our replacement process, $\varphi_i \wedge \text{At}(L \cup \{l_p\}) = \varphi_i \wedge (\text{At}(L) \cup \{p\}) = \varphi' \wedge p$ is a disjunct of φ^* . Consequently, φ^* is a maximal positive normal form of φ . \square

Proposition 3 (Normal form for SCL) *Every formula φ is SCL-provably equivalent to a maximal literal-contradiction closed disjunctive normal form φ_{lcl} .*

Proof Given φ , form φ_{max} (which can be done in AC which is contained in SCL). Then, for all $p, \neg p \in L(\varphi_{\text{max}})$, add the disjunct $p \wedge \neg p$ to φ_{max} . Call the result φ' which is still SCL-provably equivalent to φ by iterated application of the SCL-axiom (and applying the transitivity rule of AC). Then, again, form φ'_{max} which is the required φ_{lcl} . \square

Proposition 4 (Normal form for SCA) *Every formula φ is SCA-provably equivalent to a maximal atom-contradiction closed disjunctive normal form φ_{acl} .*

³⁶To be a bit more precise: Say $\varphi' = \varphi_i \wedge \bar{p} \wedge \neg q_1 \wedge \dots \wedge \neg q_m$. Then, by maximality, $\varphi_i \wedge \bar{p} \wedge \neg q_1 \wedge \dots \wedge \neg q_{m-1}$ is a disjunct of φ_{max} , too. By the SFA-axiom, SFA proves that

$$\begin{aligned} & (\varphi_i \wedge \bar{p} \wedge \neg q_1 \wedge \dots \wedge \neg q_{m-1}) \vee ((\varphi_i \wedge \bar{p} \wedge \neg q_1 \wedge \dots \wedge \neg q_{m-1}) \wedge \neg q_m) \\ \equiv & (\varphi_i \wedge \bar{p} \wedge \neg q_1 \wedge \dots \wedge \neg q_{m-1}) \vee ((\varphi_i \wedge \bar{p} \wedge \neg q_1 \wedge \dots \wedge \neg q_{m-1}) \wedge q_m), \end{aligned}$$

so we can replace the formula to the left of \equiv , which is modulo order a subformula of φ_{max} , by the formula to the right and obtain an SFA-equivalent formula φ_1 .

We continue this process with $\varphi_i \wedge (\bar{p} \wedge q_m) \wedge \neg q_1 \wedge \dots \wedge \neg q_{m-1}$ by using the disjunct $\varphi_i \wedge (\bar{p} \wedge q_m) \wedge \neg q_1 \wedge \dots \wedge \neg q_{m-2}$ that was in the original φ_{max} and still is in φ_1 . So we can SFA-provably replace $\neg q_{m-1}$ by q_{m-1} and obtain φ_2 . We continue until we replaced all the $\neg q_j$'s by q_j 's.

And if this replacement process applied to another $\varphi'' = \varphi_i \wedge \bar{p} \wedge \neg r$ also requires a disjunct $\varphi_i \wedge \bar{p} \wedge \neg q_1 \wedge \dots \wedge \neg q_k$, then we first add a copy of this disjunct to the current φ_j (which we SFA-provably can do by idempotence) and then use one of them to replace $\neg q_k$ with q_k .

To not be overly tedious, we omit a fully detailed proof of this fact.

Proof As in Proposition 3 except for adding the disjunct $p \wedge \neg p$ already if $p \in At(\varphi)$. □

In Corollary 2 below, we prove that all of these normal forms are unique. (In the case of AC this was, as mentioned, already proven by [23].)

Step 2: Characterizing Normal Forms

In this section, we do the second step of the proof: showing that if two normal forms of a synonymy satisfy the two characterizing properties of that synonymy, they are identical (modulo the order and repeats of literals and disjuncts).

We first prove a lemma that we'll need several times. Recall that $\varphi \Leftrightarrow_{FDE} \psi$ was defined as: for all four-valued valuations v , $v(\varphi) = v(\psi)$.

Lemma 3 *Let φ and ψ be two sentences in disjunctive normal form such that $\varphi \Leftrightarrow_{FDE} \psi$. Then:*

- (i) *For all disjuncts φ_0 of φ there is a disjunct ψ_* of ψ such that $L(\psi_*) \subseteq L(\varphi_0)$.*
- (ii) *For all disjuncts φ_i of φ that are not disjuncts of ψ , there are disjuncts φ_j and ψ_k such that $L(\varphi_j) = L(\psi_k) \subseteq L(\varphi_i)$ and φ_j is minimal in φ (i.e. no disjunct of φ is properly contained in φ_j).*

Proof *Ad (i).* Assume there is a disjunct φ_0 of φ such that for all disjuncts ψ_* of ψ we have $L(\psi_*) \not\subseteq L(\varphi_0)$. We show $\varphi \not\equiv_{FDE} \psi$.

Indeed, consider the four-valued valuation v defined by

$$v(p) := \begin{cases} 1, & \text{if } p \in L(\varphi_0) \text{ and } \neg p \notin L(\varphi_0) \\ 0, & \text{if } p \notin L(\varphi_0) \text{ and } \neg p \in L(\varphi_0) \\ b, & \text{if } p \in L(\varphi_0) \text{ and } \neg p \in L(\varphi_0) \\ u, & \text{if } p \notin L(\varphi_0) \text{ and } \neg p \notin L(\varphi_0). \end{cases}$$

(Cf. the correspondence between valuations and sets of literals introduced in Section 5.) Then $v(\varphi_0) \in \{1, b\}$ since each literal in φ_0 is either 1 or b under v . Hence $v(\varphi) \in \{1, b\}$ since φ is a disjunction with disjunct φ_0 .

Moreover, we claim that $v(\psi) \in \{0, u\}$. It suffices to show that for each disjunct ψ_* of ψ we have $v(\psi_*) \in \{0, u\}$. Indeed, pick such a ψ_* . By our assumption, there is a literal l in ψ_* that is not in φ_0 . If $l = q$ for an atomic sentence q , then $q \notin L(\varphi_0)$. So if $\neg q \in L(\varphi_0)$, then $v(q) = 0$, and if $\neg q \notin L(\varphi_0)$, then $v(q) = u$. Hence $v(q) = v(l) \in \{0, u\}$. Dually, if $l = \neg q$ for an atomic q , then $v(\neg q) = v(l) \in \{0, u\}$. Since $v(l) \in \{0, u\}$, also $v(\psi_*) \in \{0, u\}$ since ψ_* is a conjunction with conjunct l .

Ad (ii). Fix a φ_i that is not disjunct of ψ . Consider the disjuncts φ_* of φ such that $L(\varphi_*) \subseteq L(\varphi_i)$ and φ_* is not a disjunct of ψ . Since there are finitely many, we can pick a \subseteq -minimal one, say φ_i^* , that is, $L(\varphi_i^*) \subseteq L(\varphi_i)$ and φ_i^* is not a disjunct of ψ and

$$\forall j \leq n : L(\varphi_j) \subsetneq L(\varphi_i^*) \Rightarrow L(\varphi_j) \in C(\psi), \tag{1}$$

where $C(\psi)$ is the set of $L(\psi_0)$'s of disjuncts ψ_0 of ψ . Since $\varphi \Leftrightarrow_{\text{FDE}} \psi$, we have by (i) that there is a ψ_k ($k \leq m$) such that $L(\psi_k) \subseteq L(\varphi_i^*)$. Among the disjuncts of ψ with this property, we can choose a minimal one $L(\psi_k^*) \subseteq L(\psi_k)$, that is,

$$\forall j \leq m : L(\psi_j) \subseteq L(\psi_k^*) \text{ and } L(\psi_j) \subseteq L(\varphi_i^*) \Rightarrow L(\psi_j) = L(\psi_k^*). \tag{2}$$

Since $\psi \Leftrightarrow_{\text{FDE}} \varphi$, we have by (i) that there is a φ_r such that $L(\varphi_r) \subseteq L(\psi_k^*)$. Again, there is a minimal disjunct $L(\varphi_r^*) \subseteq L(\varphi_r)$ (so no disjunct of φ is properly contained in φ_r^*). Since $L(\varphi_r^*) \neq L(\psi_k)$ (otherwise φ_r^* would be in ψ), we have

$$L(\varphi_r^*) \subseteq L(\varphi_r) \subseteq L(\psi_k^*) \subseteq L(\psi_k) \subsetneq L(\varphi_i^*),$$

so by (1) we have $L(\varphi_r^*) \in C(\psi)$. So by (2), $L(\varphi_r^*) = L(\psi_k^*) \subseteq L(\varphi_i)$. □

Proposition 5 (Identity of SF-normal form) *Let φ and ψ be two sentences in standard minimal disjunctive normal form. Then*

$$\varphi \Leftrightarrow_{\text{FDE}} \psi \text{ iff } \varphi = \psi.$$

Proof For the non-trivial direction, write

$$\begin{aligned} \varphi &= \varphi_1 \vee \dots \vee \varphi_r \vee \varphi'_1 \vee \dots \vee \varphi'_s \\ \psi &= \psi_1 \vee \dots \vee \psi_u \vee \psi'_1 \vee \dots \vee \psi'_v, \end{aligned}$$

where $\varphi_1, \dots, \varphi_r$ are exactly those disjuncts of φ that are—modulo ordering—also disjuncts of ψ (so the remaining $\varphi'_1, \dots, \varphi'_s$ aren't disjuncts of ψ). Analogously, the unprimed disjuncts of ψ occur in φ , and the primed ones don't.

We claim that primed disjuncts are extensions of unprimed ones, that is, for all $j \leq s$, $\varphi'_j = \varphi_i \wedge L$ (modulo ordering) for an $i \leq r$ and a set of literals L (hence L is a subset of the literals occurring in φ). Analogously for ψ .

Indeed, fix a φ'_j . Since $\varphi \Leftrightarrow_{\text{FDE}} \psi$ and φ'_j is not in ψ , we have by Lemma 3(ii) that there are disjuncts φ_0 and ψ_* (primed or unprimed) such that $L(\varphi_0) = L(\psi_*) \subseteq L(\varphi'_j)$. Hence φ_0 is in ψ and $\varphi'_j = \varphi_0 \wedge L$ for $L := L(\varphi'_j) \setminus L(\varphi_0)$, which shows the claim.

Now, since φ is minimal, no disjunct can be the extension of another one, hence the set of primed disjuncts is empty. The same goes for ψ . Thus, φ and ψ really look like this:

$$\varphi = \varphi_1 \vee \dots \vee \varphi_r \quad \text{and} \quad \psi = \psi_1 \vee \dots \vee \psi_u,$$

and recall that the φ_i 's also occur as disjuncts in ψ and *vice versa*. Hence $\{L(\varphi_1), \dots, L(\varphi_r)\} = \{L(\psi_1), \dots, L(\psi_u)\}$. Since φ and ψ are standard, their order is fixed, so $\varphi = \psi$, as wanted. □

Proposition 6 (Identity of AC-normal form) *Let φ and ψ be two sentences in standard maximal disjunctive normal form. Then*

$$L(\varphi) = L(\psi) \text{ and } \varphi \Leftrightarrow_{\text{FDE}} \psi \text{ iff } \varphi = \psi.$$

Proof For the non-trivial direction, write $\varphi = \varphi_1 \vee \dots \vee \varphi_n$ and $\psi = \psi_1 \vee \dots \vee \psi_m$. By standardness, it suffices to show that

$$C(\varphi) := \{L(\varphi_1), \dots, L(\varphi_n)\} = \{L(\psi_1), \dots, L(\psi_m)\} =: C(\psi).$$

Assume for contradiction that there is a φ_i (for $i \leq n$) such that $L(\varphi_i) \notin C(\psi)$ (the other case is analogous). Since $\varphi \Leftrightarrow_{\text{FDE}} \psi$, we have by Lemma 3(ii) that there are $L(\varphi_j) = L(\psi_k) \subseteq L(\varphi_i)$. Write $L := L(\varphi_i) \setminus L(\varphi_j)$, so $L(\varphi_i) = L(\varphi_j) \cup L$. Since φ and ψ have the same literals, the literals in L also occur in ψ . So, since φ_j is a disjunct of ψ and ψ is in maximal normal form, $L(\varphi_j) \cup L$ is a disjunct of ψ , in contradiction to $L(\varphi_i) \notin C(\psi)$. \square

Proposition 7 (Identity of SFA-normal form) *Let φ and ψ be two sentences in standard maximal positive disjunctive normal form. Then*

$$At(\varphi) = At(\psi) \text{ and } \varphi \Leftrightarrow_{\text{FDE}} \psi \text{ iff } \varphi = \psi.$$

Proof For the non-trivial direction, write $\varphi = \varphi_1 \vee \dots \vee \varphi_n$ and $\psi = \psi_1 \vee \dots \vee \psi_m$. As in Proposition 6 before, it suffices to show that $C(\varphi) = C(\psi)$. Assume for contradiction that $L(\varphi_i)$ violates this claim. Since φ is maximally positive we have, by clause (iii)(a) of Definition 4, that $\varphi_i = \varphi_0 \wedge A$ for a minimal disjunct φ_0 of φ and an $A \subseteq At(\varphi)$.

We claim that $L(\varphi_0) \in C(\psi)$. If not, then by Lemma 3(ii) there are $L(\varphi_j) = L(\psi_k) \subseteq L(\varphi_0)$ with φ_j being minimal in φ . By minimality of φ_0 , $L(\varphi_0) = L(\varphi_j) = L(\psi_k) \in C(\psi)$.

Since $\varphi_i = \varphi_0 \wedge A$ with φ_0 being a disjunct of ψ and $A \subseteq At(\varphi) = At(\psi)$, we have, by clause (iii)(b) of Definition 4, that $\varphi_i \in C(\psi)$, contradiction. \square

Proposition 8 (Identity of SCL-normal form) *Let φ and ψ be two sentences in standard maximal literal-contradiction closed disjunctive normal form. Then*

$$L(\varphi) = L(\psi) \text{ and } \varphi \Leftrightarrow_{\text{C}} \psi \text{ iff } \varphi = \psi.$$

Proof For the non-trivial direction, assume $L(\varphi) = L(\psi)$ and $\varphi \Leftrightarrow_{\text{C}} \psi$. As before, given a disjunct φ_0 of $\varphi = \varphi_1 \vee \dots \vee \varphi_n$, we need to show that φ_0 also is a disjunct of $\psi = \psi_1 \vee \dots \vee \psi_m$. For contradiction, assume otherwise. Note that,

$$\text{for all } i \leq m, L(\psi_i) \not\subseteq L(\varphi_0), \tag{3}$$

because otherwise, by maximality and $L(\varphi) = L(\psi)$, φ_0 already is a disjunct of ψ . We consider two cases.

Case 1: φ_0 is not classically satisfiable. Then, since φ_0 is a conjunction of literals, there is $p \in At(\varphi_0)$ such that $p, \neg p \in \varphi_0$. Hence $p, \neg p \in L(\varphi) = L(\psi)$. Hence, since ψ is closed under literal-contradictions, $p \wedge \neg p$ is a disjunct of ψ that is contained in φ_0 , in contradiction to (3).

Case 2: φ_0 is classically satisfiable. We construct a classical valuation $v : \{p_0, p_1, \dots\} \rightarrow \{0, 1\}$ making φ_0 true (in signs, $v \models \varphi_0$) and every ψ_i false (in signs, $v \not\models \psi_i$), in contradiction to $\varphi \Leftrightarrow_{\text{C}} \psi$.

Since φ_0 is assumed to be satisfiable, fix a valuation v_0 making it true. We construct v inductively: We perform m steps (corresponding to ψ_1, \dots, ψ_m) such that at the end of step i we have a valuation v_i making φ_0 true and $\psi_1 \vee \dots \vee \psi_i$ false. We can then choose $v := v_m$.

Step 1. Find a literal l in ψ_1 which is not in φ_0 —which must exist since, by (3), $L(\psi_1) \not\subseteq L(\varphi_0)$. Define v_1 to be like v_0 except that it makes l false (then $v_1 \models \varphi_0$ and $v_1 \not\models \psi_1$, as required).

Step $i + 1$. Assume v_i has been constructed such that $v_i \models \varphi_0$ and $v_i \not\models \psi_1 \vee \dots \vee \psi_i$. We build v_{i+1} . We consider two cases.³⁷

Case (a). There is a literal l in $L(\psi_{i+1}) \setminus (L(\varphi_0) \cup L(\psi_1) \cup \dots \cup L(\psi_i))$. Then, as in step 1, define v_{i+1} to be like v_i except that it makes l false (then $v_{i+1} \models \varphi_0$ and $v_{i+1} \not\models \psi_1 \vee \dots \vee \psi_i$ and $v_{i+1} \not\models \psi_{i+1}$, as required).

Case (b). $L(\psi_{i+1}) \subseteq L(\varphi_0) \cup L(\psi_1) \cup \dots \cup L(\psi_i)$. For the remainder of the proof, we abuse notation and write the conjunctive normal form χ for the set $L(\chi)$. Thus, we have (where φ_0^c denotes all literals not in φ_0):

$$\psi_{i+1} = (\psi_{i+1} \cap \varphi_0) \cup (\psi_{i+1} \cap \psi_1 \cap \varphi_0^c) \cup \dots \cup (\psi_{i+1} \cap \psi_i \cap \varphi_0^c).$$

Since $\psi_{i+1} \not\subseteq \varphi_0$, at least one of the sets $\psi_{i+1} \cap \psi_j \cap \varphi_0^c$ (for $j \leq i$) is non-empty. Let $j_1, \dots, j_r \leq i$ be those j for which there is a literal $l_j \in \psi_{i+1} \cap \psi_j \cap \varphi_0^c \neq \emptyset$. Now, either there is a valuation w making all literals of the set $L := \{l_{j_1}, \dots, l_{j_r}\}$ false, or there isn't.

If there isn't, $l_{j_1} \vee \dots \vee l_{j_r}$ is a tautology, whence there is a p such that $p, \neg p \in L \subseteq \psi_{i+1}$. Hence, any valuation makes ψ_{i+1} false. Thus, we choose $v_{i+1} := v_i$ and we have $v_{i+1} \models \varphi_0$ and $v_{i+1} \not\models \psi_1 \vee \dots \vee \psi_i$ and $v_{i+1} \not\models \psi_{i+1}$, as required.

So assume there is a valuation $w \models \neg l_{j_1} \wedge \dots \wedge \neg l_{j_r}$. Define v_{i+1} to be like v_i except that on L it is like w . We claim v_{i+1} has the required properties. Indeed, $v_{i+1} \models \varphi_0$ because $v_i \models \varphi_0$ and v_{i+1} only differs from v_i on the set L which is disjoint from $At(\varphi_0)$ (since each $l \in L$ is in φ_0^c). Moreover, $v_{i+1} \not\models \psi_{i+1}$ since for $l_{j_1} \in L$ we have $l_{j_1} \in \psi_{i+1}$ (by construction of L) and $v_{i+1}(l_{j_1}) = w(l_{j_1}) = 0$. Finally, we fix an $j \leq i$ and show that $v_{i+1} \not\models \psi_j$. If $\psi_{i+1} \cap \psi_j \cap \varphi_0^c \neq \emptyset$. Then, by construction, $l_j \in L$, so $v_{i+1}(l_j) = w(l_j) = 0$, whence $v_{i+1} \not\models \psi_j$. So assume $\psi_{i+1} \cap \psi_j \cap \varphi_0^c = \emptyset$. Since $v_i \not\models \psi_j$, there is a literal $l \in \psi_j$ such that $v_i(l) = 0$. Moreover, $l \notin \varphi_0$ since $v_i \models \varphi_0$. Since $\psi_{i+1} \cap \psi_j \cap \varphi_0^c = \emptyset$ we hence have $l \notin \psi_{i+1} \supseteq L$. Hence $v_{i+1}(l) = v_i(l) = 0$, so $v_{i+1} \not\models \psi_j$. \square

Proposition 9 (Identity of SCA-normal form) *Let φ and ψ be two sentences in standard maximal atom-contradiction closed disjunctive normal form. Then*

$$At(\varphi) = At(\psi) \text{ and } \varphi \Leftrightarrow_C \psi \text{ iff } \varphi = \psi.$$

Proof For the non-trivial direction, assume $At(\varphi) = At(\psi)$ and $\varphi \Leftrightarrow_C \psi$.

We claim that $L(\varphi) = At(\varphi) \cup \neg At(\varphi)$ (where $\neg At(\varphi) := \{\neg p : p \in At(\varphi)\}$). Indeed, " \subseteq " is trivial, so assume $l \in At(\varphi) \cup \neg At(\varphi)$. Write p for the atom of the

³⁷We'll see below that it is not necessary, but we could consider separately the trivial case where already $v_i \not\models \psi_{i+1}$. There we set $v_{i+1} := v_i$. By assumption, we have $v_{i+1} \models \varphi_0$ and $v_{i+1} \not\models \psi_1 \vee \dots \vee \psi_i \vee \psi_{i+1}$, as required.

literal l . Then $p \in At(\varphi)$. Since φ is closed under atom-contradictions, $p \wedge \neg p$ is a disjunct of φ . Hence $p, \neg p \in L(\varphi)$, whence $l \in L(\varphi)$.

Hence $L(\varphi) = At(\varphi) \cup \neg At(\varphi) = At(\psi) \cup \neg At(\psi) = L(\psi)$. So by Proposition 8 (which applies since φ and ψ are atom- and, hence in particular, literal-contradiction closed), $\varphi = \psi$. □

Proofs of the Theorems

In this section, we put everything together and prove the theorems.

Theorem 2 (i). For all sentences φ and ψ , $SF \vdash \varphi \equiv \psi$ iff $\llbracket \varphi \rrbracket = \llbracket \psi \rrbracket$ (iff $\varphi \leftrightarrow_{FDE} \psi$).

Proof Soundness (left to right). It is readily checked that if $\varphi \equiv \psi$ is an axiom of SF, then $\varphi \leftrightarrow_{FDE} \psi$. Moreover, it is also readily checked that if $\varphi' \equiv \psi'$ is the result of applying one of the rules to $\varphi \equiv \psi$, and if $\varphi \leftrightarrow_{FDE} \psi$, that then also $\varphi' \leftrightarrow_{FDE} \psi'$.

Completeness (right to left). Assume $\varphi \leftrightarrow_{FDE} \psi$. By Proposition 1, there are normal forms φ_{min} and ψ_{min} of φ and ψ , respectively. By soundness,

$$\varphi_{min} \leftrightarrow_{FDE} \varphi \leftrightarrow_{FDE} \psi \leftrightarrow_{FDE} \psi_{min},$$

so, by Proposition 5, $\varphi_{min} = \psi_{min}$, so we can prove in SF that

$$\varphi \equiv \varphi_{min} \equiv \psi_{min} \equiv \psi,$$

so $\vdash_{SF} \varphi \equiv \psi$. □

Curiously, while *maximal* disjunctive forms were needed for the completeness of truthmaker semantics for AC, we need *minimal* disjunctive forms for the completeness of scenario semantics for SF.

Theorem 2 (ii). For all sentences φ and ψ , we have

$$AC \vdash \varphi \equiv \psi \text{ iff } \begin{cases} L(\varphi_{max}) = L(\psi_{max}) , \text{ and} \\ \varphi \leftrightarrow_{FDE} \psi \end{cases} .$$

Proof Let's start with the left-to-right direction. We show $AC \vdash \varphi \equiv \psi$ implies $\leftrightarrow_{FDE} \psi$ by contraposition (though it could also be shown directly by induction): If $\varphi \not\leftrightarrow_{FDE} \psi$, then, by Theorem 1, $SF \not\vdash \varphi \equiv \psi$, so in particular $AC \not\vdash \varphi \equiv \psi$ (since SF is an extension of AC). And $AC \vdash \varphi \equiv \psi$ implies $L(\varphi_{max}) = L(\psi_{max})$ because of the following. If $AC \vdash \varphi \equiv \psi$, then $AC \vdash \varphi_{max} \equiv \varphi \equiv \psi \equiv \psi_{max}$. So φ_{max} and ψ_{max} are two sentences in standard maximal normal form that are AC-equivalent to φ . We know that a sentence's standard maximal normal form is unique in AC. (A purely syntactic proof of this fact was given in [2], and [23] gave a semantic proof using his truthmaker semantics.) Hence $\varphi_{max} = \psi_{max}$, and, in particular, $L(\varphi_{max}) = L(\psi_{max})$. (Of course, this could also be shown directly by induction on AC-proofs.)

For the other direction, assume $L(\varphi_{max}) = L(\psi_{max})$ and $\varphi \leftrightarrow_{FDE} \psi$. We have $AC \vdash \varphi_{max} \equiv \varphi$ and $AC \vdash \psi_{max} \equiv \psi$. Moreover, we've seen in the left-to-right direction that AC-equivalence entails FDE-equivalence, so we have

$$\varphi_{max} \leftrightarrow_{FDE} \varphi \leftrightarrow_{FDE} \psi \leftrightarrow_{FDE} \psi_{max}.$$

Hence, applying Proposition 6 to φ_{\max} and ψ_{\max} we get that $\varphi_{\max} = \psi_{\max}$, whence AC indeed proves $\varphi \equiv \varphi_{\max} \equiv \psi_{\max} \equiv \psi$. □

Further equivalences are given by the following.

Corollary 1. The following are equivalent:

- (i) $AC \vdash \varphi \equiv \psi$.
- (ii) For every state model M , we have $[\varphi]_M = [\psi]_M$.
- (iii) $[\varphi]_C = [\psi]_C$ where C is the canonical scenario model based on sets of scenarios.

Proof (Sketch) (i) \Rightarrow (ii) is the soundness theorem of [23]. (ii) \Rightarrow (iii) is trivial. So it remains to show (iii) \Rightarrow (i).

Indeed, assume $[\varphi]_C = [\psi]_C$. Since φ is AC-provably equivalent to its standard maximal disjunctive form φ_{\max} , and ψ to ψ_{\max} , we have, by the just mentioned soundness, that $[\varphi_{\max}]_C = [\psi_{\max}]_C$.

Analogous to [23], it is straightforward to show the following. (a) For a conjunctive form $\varphi = l_1 \wedge \dots \wedge l_n$ we have, by induction on n , that $|\varphi|_C = \{L_\varphi\}$ where $L_\varphi = \{s_{l_1}, \dots, s_{l_n}\}$.³⁸ (b) For a disjunctive form $\varphi = \varphi_1 \vee \dots \vee \varphi_n$, $|\varphi|_C = \{L_{\varphi_1}, \dots, L_{\varphi_n}\}$. (c) If φ is a maximal disjunctive normal form, then $|\varphi|_C = [\varphi]_C$. (d) If φ and ψ are maximal standard disjunctive forms, then $[\varphi]_C = [\psi]_C$ iff $\varphi = \psi$.

Hence $\varphi_{\max} = \psi_{\max}$. Whence AC proves $\varphi \equiv \varphi_{\max} \equiv \psi_{\max} \equiv \psi$. □

Theorem 2 (iii). For all sentences φ and ψ , we have

$$SFA \vdash \varphi \equiv \psi \text{ iff } \begin{cases} At(\varphi) = At(\psi) , \text{ and} \\ \varphi \Leftrightarrow_{FDE} \psi \end{cases} .$$

Proof The left-to-right direction is immediate by induction on SFA-proofs: For SFA-axioms $\varphi \equiv \psi$ we have that $At(\varphi) = At(\psi)$ and $\varphi \Leftrightarrow_{FDE} \psi$, and these two properties are preserved by the SFA-rules.

For the other direction, assume $At(\varphi) = At(\psi)$ and $\varphi \Leftrightarrow_{FDE} \psi$. By Proposition 2, there are φ_{pos} and ψ_{pos} in standard maximal positive disjunctive form such that $SFA \vdash \varphi_{\text{pos}} \equiv \varphi$ and $SFA \vdash \psi_{\text{pos}} \equiv \psi$. Moreover, we've seen in the left-to-right direction that SFA-equivalence entails having the same atoms and FDE-equivalence, so we have

$$At(\varphi_{\text{pos}}) = At(\varphi) = At(\psi) = At(\psi_{\text{pos}}), \text{ and} \\ \varphi_{\text{pos}} \Leftrightarrow_{FDE} \varphi \Leftrightarrow_{FDE} \psi \Leftrightarrow_{FDE} \psi_{\text{pos}} .$$

Hence we have by Proposition 7 that $\varphi_{\text{pos}} = \psi_{\text{pos}}$. Hence SFA indeed proves $\varphi \equiv \varphi_{\text{pos}} \equiv \psi_{\text{pos}} \equiv \psi$. □

Theorem 2 (iv). For all sentences φ and ψ , we have

$$SCL \vdash \varphi \equiv \psi \text{ iff } \begin{cases} L(\varphi_{\max}) = L(\psi_{\max}) , \text{ and} \\ \varphi \Leftrightarrow_C \psi \end{cases} .$$

³⁸Recall from Section 5 that in the canonical scenario model we defined $s_l := \{l\}$ for literals l .

Proof The left-to-right direction is shown by induction on SCL-proofs: That $SCL \vdash \varphi \equiv \psi$ implies $\varphi \leftrightarrow_C \psi$ is immediate, so let's consider the subject matter condition. Concerning the axioms, we only need to check the new SCL-axiom. So given φ and $p, \neg p \in L(\varphi) = L(\varphi_{max})$ we need to show $L(\varphi) = L(\varphi \vee (p \wedge \neg p))$, that is, $L(\varphi_{max}) = L([\varphi \vee (p \wedge \neg p)]_{max})$. Indeed, we have $AC \vdash \varphi \equiv \varphi_{max}$, so by AC's rule for disjunction $AC \vdash \varphi \vee (p \wedge \neg p) \equiv \varphi_{max} \vee (p \wedge \neg p)$. Hence AC proves $[\varphi \vee (p \wedge \neg p)]_{max} \equiv \varphi \vee (p \wedge \neg p) \equiv \varphi_{max} \vee (p \wedge \neg p)$. Since AC-synonymy implies L -identity, $L([\varphi \vee (p \wedge \neg p)]_{max}) = L(\varphi_{max} \vee (p \wedge \neg p)) = L(\varphi_{max}) \cup \{p, \neg p\} = L(\varphi_{max})$. Concerning the rules, we only have the rules of AC of which we know that they preserve L -identity.

For the other direction, assume $L(\varphi_{max}) = L(\psi_{max})$ and $\varphi \leftrightarrow_C \psi$. By Proposition 3, there are φ_{lcl} and ψ_{lcl} in standard maximal literal-contradiction closed disjunctive form such that $SCL \vdash \varphi_{lcl} \equiv \varphi$ and $SCL \vdash \psi_{lcl} \equiv \psi$. Moreover, from the left-to-right direction we know

$$L(\varphi_{lcl}) = L(\varphi) = L(\psi) = L(\psi_{lcl}), \text{ and}$$

$$\varphi_{lcl} \leftrightarrow_C \varphi \leftrightarrow_C \psi \leftrightarrow_C \psi_{lcl}.$$

Hence we have by Proposition 8 that $\varphi_{lcl} = \psi_{lcl}$. So SCL indeed proves $\varphi \equiv \varphi_{lcl} \equiv \psi_{lcl} \equiv \psi$. □

Theorem 2 (v). For all sentences φ and ψ , we have

$$SCA \vdash \varphi \equiv \psi \text{ iff } \begin{cases} At(\varphi) = At(\psi) , \text{ and} \\ \varphi \leftrightarrow_C \psi \end{cases} .$$

Proof The left-to-right direction is immediate by induction on SCA-proofs.

For the other direction, assume $At(\varphi) = At(\psi)$ and $\varphi \leftrightarrow_C \psi$. By Proposition 4, there are φ_{acl} and ψ_{acl} in standard maximal atom-contradiction closed disjunctive form such that $SCA \vdash \varphi_{acl} \equiv \varphi$ and $SCA \vdash \psi_{acl} \equiv \psi$. Moreover, from the left-to-right direction we know

$$At(\varphi_{acl}) = At(\varphi) = At(\psi) = At(\psi_{acl}), \text{ and}$$

$$\varphi_{acl} \leftrightarrow_C \varphi \leftrightarrow_C \psi \leftrightarrow_C \psi_{acl}.$$

Hence we have by Proposition 9 that $\varphi_{acl} = \psi_{acl}$. So SCA indeed proves $\varphi \equiv \varphi_{acl} \equiv \psi_{acl} \equiv \psi$. □

Corollary 2 (Uniqueness of normal form) *Let \mathcal{C} be one of the systems AC, SFA, SCL, SCA, or SF. Then every sentence φ has a unique standard disjunctive normal form $\varphi_{\mathcal{C}}$ with the properties corresponding to the system \mathcal{C} (e.g. maximal, maximal positive, etc.).*

Proof Assume $\varphi_{\mathcal{C}}$ and $\varphi'_{\mathcal{C}}$ are standard normal forms of φ in the system \mathcal{C} . Then $\mathcal{C} \vdash \varphi_{\mathcal{C}} \equiv \varphi'_{\mathcal{C}}$. Apply the \mathcal{C} -soundness theorem and then the \mathcal{C} -characterization lemma to get $\varphi_{\mathcal{C}} = \varphi'_{\mathcal{C}}$. □

Acknowledgements For inspiring discussions and comments, I'm grateful to Franz Berto, Kit Fine, Peter Hawke, Johannes Korbmayer, Karolina Krzyzanowska, Michiel van Lambalgen, Hannes Leitgeb, Martin Lipman, Dean McHugh, Aybüke Özgün, Giorgio Sbardolini, Tom Schoonen, an anonymous referee of this journal, and the audiences of the workshop "Hyperintensional Logics and Truthmaker Semantics" at Ghent University and the Logic of Conceivability Seminar at the University of Amsterdam. This work is part of the research programme "Foundations of Analogical Thinking" with project number 322-20-017, which is financed by the Netherlands Organisation for Scientific Research (NWO).

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

References

1. Angell, R. (1977). Three systems of first degree entailment. *Journal of Symbolic Logic*, 47, 147.
2. Angell, R.B. (1989). Deducibility, entailment and analytic containment. In Norma, J., & Sylvan, R. (Eds.) *Directions in relevant logic*, Kluwer, Dordrecht, chap. 8 (pp. 119–144).
3. Beall, J.C. (2016). Off-topic: a new interpretation of weak-kleene logic. *The Australasian Journal of Logic* 13,(6).
4. Berto, F. (2017). Impossible worlds and the logic of imagination. *Erkenntnis*, 82(6), 1277–1297.
5. Berto, F. (2018). Aboutness in imagination. *Philosophical Studies*, 175(8), 1871–1886.
6. Berto, F., & Jago, M. (2018). Impossible worlds. In Zalta, E.N. (Ed.) *The stanford encyclopedia of philosophy*, *Metaphysics Research Lab, Stanford University*.
7. Besold, T.R., d'Avila Garcez, A., Stenning, K., van der Torre, L., van Lambalgen, M. (2017). Reasoning in non-probabilistic uncertainty: logic programming and neural-symbolic computing as examples. *Minds and Machines*, 27(1), 37–77.
8. Bochvar, D.A. (1938). On a three-valued logical calculus and its application to the analysis of the paradoxes of the classical extended functional calculus. *Matematicheskii Sbornik*, 4(46)(2), 287–308.
9. Bochvar, D.A. (1981). On a three-valued logical calculus and its application to the analysis of the paradoxes of the classical extended functional calculus. *History and Philosophy of Logic*, 2, 87–112. This is the English translation by M.Bergmann of the original [8].
10. Carnap, R. (1947). *Meaning and necessity*. Chicago: The University of Chicago Press.
11. Chalmers, D.J. (2006). The foundations of two-dimensional semantics. In Garcia-Carpintero, M., & Macia, J. (Eds.) *Two-dimensional semantics: foundations and applications* (pp. 55–140). New York: Oxford University Press.
12. Chierchia, G., Fox, D., Spector, B. (2009). Hurford's constraint and the theory of scalar implicatures: evidence for embedded implicatures. In Egré, P., & Magri, G. (Eds.) *Presuppositions and implicatures* (pp. 47–62). Cambridge: MIT Press.
13. Church, A. (1954). Intensional isomorphism and identity of belief. *Philosophical Studies*, 5, 65–73.
14. Ciardelli, I., & Roelofsen, F. (2017). Hurford's constraint, the semantics of disjunction, and the nature of alternatives. *Natural Language Semantics*, 25, 199–222.
15. Correia, F. (2016). On the logic of factual equivalence. *The Review of Symbolic Logic*, 9(1), 103–122.
16. Daniels, C. (1990). A note on negation. *Erkenntnis*, 32, 423–429.
17. Dovšen, K., & Petrić, Z. (2012). Isomorphic formulae in classical propositional logic. *Mathematical Logic Quarterly*, 58(1-2), 5–17.
18. Duží, M. (2017). If structured propositions are logical procedures then how are procedures individuated? *Synthese*, 1–35.
19. Feldman, J. (2013). The neural binding problem(s). *Cognitive Neurodynamics*, 7, 1–11.
20. Ferguson, T.M. (2014). A computational interpretation of conceptivism. *Journal of Applied Non-Classical Logics*, 24(4), 333–367.
21. Ferguson, T.M. (2016). Faulty belnap computers and subsystems of fde. *Journal of Logic and Computation*, 26(5), 1617–1636.
22. Fine, K. (1975). Normal forms in modal logic. *Notre Dame Journal of Formal Logic*, 16(2), 229–237.

23. Fine, K. (2016). Angellic content. *Journal of Philosophical Logic*, 45(2), 199–226.
24. Fine, K. (2016). Review of steve yablo's 'aboutness', draft retrieved from http://www.academia.edu/15555407/Review_of_Steve_Yablos_Aboutness (last checked 20 August 2019).
25. Fine, K., & Jago, M. (2017). Logic for exact entailment, draft retrieved from http://www.academia.edu/32032379/Logic_for_Exact_Entailment (last checked 20 August 2019).
26. Fitting, M. (1985). A kripke-kleene semantics for logic programs. *The Journal of Logic Programming*, 2, 295–312.
27. van Fraassen, B.C. (1969). Facts and tautological entailments. *The Journal of Philosophy*, 66(15), 477–487.
28. French, R. (2017). A simple sequent calculus for angell's logic of analytic containment. *Studia Logica*, 105(5), 971–994.
29. Gazdar, G. (1979). *Pragmatics: implicature, presupposition, and logical form*. New York: Academic Press.
30. Goodman, N. (1949). On likeness of meaning. *Analysis*, 10(1), 1–7.
31. Hawke, P. (2018). Theories of aboutness. *Australasian Journal of Philosophy*, 96(4), 697–723.
32. Hodges, W. (2001). Formal features of compositionality. *Journal of Logic, Language, and Information*, 10, 7–28.
33. Hornischer, L. (2017). Hyperintensionality and synonymy: a logical, philosophical, and cognitive investigation. Master's thesis, Institute for Logic, Language and Computation, Amsterdam, available at <https://www.illc.uva.nl/Research/Publications/Reports/MoL-2017-07.text.pdf>.
34. Hurford, J. (1974). Exclusive or inclusive disjunction. *Foundations of Language*, 11(3), 409–411.
35. Jago, M. (2014). *The impossible*. Oxford: Oxford University Press.
36. Katzir, R., Singh, R., et al. (2013). Hurford disjunctions: embedded exhaustification and structural economy. In Etxeberria, U. (Ed.) *Proceedings of Sinn und Bedeutung*, (Vol. 18 pp. 201–216).
37. King, J.C. (2016). Structured propositions. In Zalta, E.N. (Ed.) *The stanford encyclopedia of philosophy, Metaphysics Research Lab, Stanford University*, <https://plato.stanford.edu/archives/win2016/entries/propositions-structured/>.
38. Kowalski, R.A. (1988). The early years of logic programming. *CACM*, 31(1), 38–43.
39. van Lambalgen, M., & Hamm, F. (2005). *The proper treatment of events, explorations in semantics* Vol. 4. Oxford: Blackwell.
40. Leitgeb, H. (2008). An impossibility result on semantic resemblance. *Dialectica*, 62, 293–306.
41. Leitgeb, H. (2019). Hype: a system of hyperintensional logic (with an application to semantic paradoxes). *Journal of Philosophical Logic*, 48(2), 305–405.
42. Lewis, D.K. (1988). Statements partly about observation. In *Papers in philosophical logic* (pp. 125–155). Cambridge: Cambridge University Press.
43. Meyer, M.C. (2014). Deriving Hurford's constraint. In Snider, T., D'Antonio, S., Weigand, M. (Eds.), (Vol. 24 pp. 577–596). *Proceedings of semantics and linguistic theory: LSA and CLC Publications*.
44. Moschovakis, Y.N. (2006). A logical calculus of meaning and synonymy. *Linguistics and Philosophy*, 29(1), 27–89.
45. Parry, W.T. (1932). Implication. PhD thesis Harvard University, Cambridge, MA, unpublished.
46. Parry, W.T. (1933). Ein axiomsystem für eine neue art von implikation (analytische implikation). *Ergebnisse eines Mathematischen Colloquiums*, 4, 5–6.
47. Parry, W.T. (1968). The logic of c. i. lewis. In Schlipp, P.A. (Ed.) *The philosophy of C.I. Lewis*, Cambridge University Press, Cambridge (pp. 115–154).
48. Parry, W.T. (1989). Analytic implication: its history, justification and varieties. In Norman, J., & Sylvan, R. (Eds.) *Directions in relevant logic* (pp. 101–118). Dordrecht: Kluwer Academic.
49. Priest, G. (2008). *An introduction to non-classical logic*, 2nd edn. Cambridge: Cambridge University Press.
50. Priest, G. (2010). The logic of the catuskoti. *Comparative Philosophy*, 1(2), 24–54.
51. Restall, G. (2000). *An introduction to substructural logics*. London: Routledge.
52. Restall, G. (2018). Isomorphisms in a category of proofs, presented at CUNY Graduate Center in the Logic and Metaphysics Seminar on 9 April 2018, Handout available at: <http://consequently.org/handouts/isomorphisms-handout-cuny-2018.pdf>, (last checked 20 August 2019).
53. Routley, R. (1982). *Relevant logics and their rivals: the basic philosophical and semantical theory* Ridgeview, Atascadero, CA.
54. Simons, M. (2001). Disjunction and alternativeness. *Linguistics and Philosophy*, 24(5), 597–619.

55. Stenning, K., & van Lambalgen, M. (2008). Human reasoning and cognitive science. A Bradford book, MIT Press, Cambridge, Massachusetts.
56. Treisman, A.M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136.
57. Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological Review*, *90*(4), 293–315.
58. Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.