



UvA-DARE (Digital Academic Repository)

Compositionality in Context

Baltag, A.; van Benthem, J.; Westerståhl, D.

DOI

[10.1007/978-3-031-24117-8_21](https://doi.org/10.1007/978-3-031-24117-8_21)

Publication date

2023

Document Version

Submitted manuscript

Published in

Samson Abramsky on Logic and Structure in Computer Science and Beyond

[Link to publication](#)

Citation for published version (APA):

Baltag, A., van Benthem, J., & Westerståhl, D. (2023). Compositionality in Context. In A. Palmigiano, & M. Zadrzadeh (Eds.), *Samson Abramsky on Logic and Structure in Computer Science and Beyond* (pp. 773-812). (Outstanding Contributions to Logic; Vol. 25). Springer. https://doi.org/10.1007/978-3-031-24117-8_21

General rights

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

Disclaimer/Complaints regulations

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

Compositionality in Context

Alexandru Baltag¹, Johan van Benthem², and Dag Westerståhl³

¹ILLC, University of Amsterdam

²ILLC, University of Amsterdam, and Dept. of Philosophy,
Tsinghua University

³Dept. of Philosophy, Stockholm University and Dept. of
Philosophy, Tsinghua University

March 27, 2021

Abstract

Compositionality is a principle used in logic, philosophy, mathematics, linguistics, and computer science for assigning meanings to language expressions in a systematic manner following syntactic construction, thereby allowing for a perspicuous algebraic view of the syntax-semantics interface. Yet the status of the principle remains under debate, with positions ranging from compositionality always being achievable to its having genuine empirical content. This paper attempts to sort out some major issues in all this from a logical perspective. First, we stress the fundamental harmony between Compositionality and its apparent antipode of Contextuality that locates meaning in interaction with other linguistic expressions and in other settings than the actual one. Next, we discuss basic further desiderata in designing and adjudicating a compositional semantics for a given language in harmony with relevant contextual syntactic and semantic cues. In particular, in a series of concrete examples in the realm of logic, we point out the dangers of over-interpreting compositional solutions, the ubiquitous entanglement of assigning meanings and the key task of explaining given target inferences, and the dynamics of new language design, illustrating how even established compositional semantics can be rethought in a fruitful manner. Finally, we discuss some fresh perspectives from the realm of game semantics for natural and formal languages, the general setting for Samson Abramsky's influential work on programming languages and process logics. We highlight outside-in coalgebraic perspectives on meanings as finite or infinitely unfolding behavior that might challenge and enrich current discussions of compositionality.

1 The makings of a semantics: compositionality, contextuality, and criteria for judging design

The three intertwined topics to be addressed in this paper are introduced in the following characteristically clear passage from Samson Abramsky’s paper [3]:

“It is, or should be, an aphorism of semantics that *the key to compositionality is parameterization*. Choosing the parameters aright allows the meaning of expressions to be made sensitive to their contexts, and hence defined compositionally. While this principle could, in theory, be carried to the point of trivialization, in practice the identification of the right form of parameterization does usually represent some genuine insight into the structure at hand, (o.c., p. 19).”

Samson Abramsky is a champion of the methodology of compositionality in the semantics of programming languages [7, 2], game semantics of logical systems [6, 3, 8], and more generally the category-theoretic foundations of logic and computation. In this survey and discussion paper, we look at compositionality in its broadest sense, from philosophy and linguistics to logic and computer science. We discuss what the principle says, first proceeding abstractly in broad mathematical terms, starting from the account in Hodges [44, 45] and adding some general observations of our own. Here and throughout this article, we will highlight the strategy of achieving compositionality by adding new parameters of evaluation, reflecting the spirit of the above quote and linking up with Frege’s celebrated principle of Contextuality. Following the abstract analysis, we consider how compositionality fares in a number of case studies, and what further basic features are entangled with compositionality and contextuality when giving a logical semantics for a language. Finally, we will discuss what we can learn from them about further characteristics of well-behaved semantics. But before we do all this, an introduction is in order to the main themes as we see them.

Compositionality. The principle of Compositionality says that the meaning of a linguistic expression is a function of the meanings of its parts, plus the syntactic mode of composition of these parts. The most striking feature of this formulation is its extreme generality: it says nothing about what meanings are or what syntax (the notion of ‘parts’) should look like. It should be seen primarily, we suggest, as a way of organizing one’s ideas about the syntax-semantics interface of a language in a perspicuous algebraic format. The principle occurs widely in logic, linguistics, philosophy, computer science, and cognitive science. But its status remains under debate. Is compositionality a refutable empirical claim, a methodological recommendation, or a bit of both? Can it always be satisfied in setting up a semantics, or is it a design constraint with real empirical bite? Camps keep forming around these issues, books keep getting written. One purpose of this paper is to add some concrete considerations to these debates stemming from analyzing a number of case studies in logical semantics.

A good thing? Here is a question of motivation that should come first. Why is Compositionality so important? A number of different virtues have been touted

in the literature: it facilitates learnability, it underlies our ability to produce infinitely many meaningful or even true new sentences, it explains successful linguistic communication, it is a sufficient condition for computationally efficient interpretability of natural and formal languages; see [61] for a survey and evaluation of these claims. Indeed, the inductive, or rather recursive character, of the usual compositional semantics, makes a language computably learnable. And from another computational standpoint, compositionality facilitates a divide and conquer strategy for the analysis and synthesis of programs, [47], a virtue also emphasized in Abramsky’s work. Finally, compositionality is often seen as guaranteeing the existence of an algebraic semantics validating perspicuous algebraic laws of reasoning with the key notions of a language, [10]. Our purpose in this paper is not to adjudicate these claimed virtues, or even to survey all results and approaches. We refer to [47] in the “Handbook of Logic and Language” for a survey of compositionality in logic, linguistics, and computer science which is up to date until 2000. For the period after that, the reader may consult [60, 61], and the “Oxford Handbook of Compositionality”, [42].

In this article, we concentrate on two methodological virtues. Compositionality is a method for laying out meanings in a perspicuous manner, or if you prefer a mentalistic picture: for organizing our thoughts perspicuously, as far as expressed in language. Moreover, following this method facilitates the design of logical proof systems that govern reasoning with these meaningful notions.

Simple and complex. It is also good to realize what Compositionality does not say. It is often thought that it is all about constructing *complex meanings* out of *simple ones*. This view is reinforced by the syntax of logical languages, where one tends to view atomic formulas as expressing simple facts about concrete objects, while further logical operators add complexity. However, we see all of this as wrong, or at least confused. Compositional design makes no assumptions about the complexity of basic meanings, and it can equally well lead from complex to complex meanings, perhaps even from complex to simple ones with expressions like “ignore” or “delete”. The neutrality of Compositionality with respect to complexity of meanings will show in particular with the substitutability versions to be discussed in Section 2. Finally, on the issue of simplicity, there may also be a confusion at work. In a relative sense, composition of meanings may indeed start from simpler atomic parts, where the simplicity just means that these parts are not analyzed any further. This is the simplicity of abstraction. But in concrete instances, atomic sentences can have very complex meanings, dwarfing the marginal complexity added by the compositional construction.¹

Perhaps the simplest way of illustrating all these issues is with the concrete case of Boolean algebra, whose language has an algebraic semantics that is a

¹For instance, an atomic predicate like ‘friendly person’ may have a very complex meaning such as ‘likely to be more pleasant and helpful than the average person in the reference group’, and then the additional complexity of understanding the meaning of, say, ‘two friendly persons’ seems marginal. A similar point is made by Dummett in [26] about Russell’s ‘logical atomism’ versus Frege’s view where basic expressions could have very complex meanings. In the background, Dummett also distinguishes two complementary directions of ‘recognition’ and ‘analysis’ of meanings that are relevant to our topic of compositionality and contextuality.

showcase of compositionality. In a Boolean set algebra, all algebraic terms denote sets: there is absolutely no way of saying that a complex term denotes a more complex set than a Boolean variable. Moreover, while some Boolean algebras indeed consist of simple objects, such as the two numbers 0, 1, another perfectly legitimate semantic structure for this language is the Lindenbaum algebra of propositional logic, whose objects are defined as equivalence classes of the relation of provable equality from the principles of Boolean Algebra. Clearly, the latter objects are defined by reference to the whole language, and much more complex than the truth table operations involved in computing their compositions. But Compositionality does not favor either model as an interpretation for the Boolean language.² This observation brings us to our next theme.

Contextuality. The principle of compositionality is usually ascribed to Frege, who introduced recursive syntax and matching meanings in modern logic and studies of linguistic meaning generally, [26, 60], though this attribution has also been questioned, [46]. However this may be, Frege also stated a different influential insight in [30, x], namely, his well-known *Context Principle*:

“never . . . ask for the meaning of a word in isolation, but only in the context of a sentence.”

This view of meaning as interaction with the rest of the language seems very close to certain modern views. One is reminded of Zellig Harris’s dictum

“To know the meaning of a word, look at the company it keeps”,

but also—with a wider notion of context to which we will return—of Abramsky’s insistence on finding the right contextual parameters. This external view of meaning may seem at odds with the intuitively more internal perspective of compositionality. But this is a misconception. The preceding discussion of simplicity versus complexity of meanings pointed at a peaceful co-existence: basic unanalyzed parts could actually have meanings dependent on their role in the context of the whole language. And as we explain in the next section, Hodges has provided an elegant more precise mathematical way of resolving the tension between compositionality and fregean contextuality.

The natural, and perhaps even inevitable, interplay of compositionality and contextuality is taken for granted in this paper. Many of our concrete examples are about the search for the right contextual parameters beyond the immediate syntax of the sentence and the actual occasion of its use, that make compositionality possible. Still, an important clarification needs to be made here, since the term “context” tends to be overused, and carries at least two different senses. In the *semantic sense of contextuality*, expressions may get their meanings in rich models with ‘indices of evaluation (points, worlds, situations, etc.)’ that pack all the necessary information for a compositional modus operandi. A simple example is modal logic, where interpreting expressions at one point or world may require information about the truth values of their parts at all other points.

²Similar points apply to model-theoretic semantics, the vehicle for our later examples. This is related to algebraic semantics via a representation theory that is beyond our scope here.

Abramsky’s quote also holds for the whole history of logical semantics: indices of evaluation develop as required by the needs of compositionality.

However, there is also a stronger sense of what may be called *syntactic contextuality*, where the meaning of an expression at some point may depend on the meanings of arbitrary expressions of the language at this, or other points. The latter version seems more in line with Frege’s Context Principle. This stronger version, too, occurs naturally in logical semantics: especially when we consider meaning assignments that have to respect the intuitive validity of given inference patterns: conclusions from sentences do not need to be syntactic parts. Also, in our later discussion of game semantics, meanings in language games may be strategies that record what an agent would say under various circumstances, which can be syntactically quite different alternatives to what is actually said. This paper keeps an open eye to both senses of contextuality.

The content of this paper. We will start by presenting a general mathematical perspective on compositionality due to Hodges (Section 2). Next, we present some classical compositional semantics in Section 3, as material for a supplementary analysis in terms of what we call “currying” (Section 4). After that, we turn to concrete case studies of semantics for logical systems in order to enrich the picture of what compositional semantics involves. Considering instances from modal and first-order logic, we draw some general heuristic lines in Section 5, including the dangers of over-interpreting received ‘solutions’ and the role of impossibility results. In Section 6, we discuss a particularly pervasive feature of compositional semantics: its entanglement with other desiderata, in particular, prior intuitions of valid inference, and in Section 7, the design of new languages based on compositional analysis of some initial given language. In both cases, we shake received wisdom a bit, using case studies of generalized assignment semantics and a logic of dependence to reopen discussion on semantic choices that may have seemed settled historically once and for all. This concludes what might be called the classical part of our presentation.

Coda: game perspectives. While most themes in this article fit well with classical philosophical and linguistic discussions of compositionality, we add one more twist. In Section 8, we discuss a number of themes emanating from the long-standing use of *games* in logic, linguistics and computer science. In particular, in Abramsky’s semantics of programming languages and concurrent processes, computation has come to mean producing, not just transitions from input to outputs, but complex ongoing *behavior*, that can be infinite just as well as finite. We present some basic ideas from game models for computing, including more radical co-algebraic versions where behavior is not built up from inside in terms of basic building blocks, but is only observable from the outside. This might well affect received understandings of semantics more broadly, also in philosophy and linguistics, as it suggests that meanings are associated with long-term use, a very radical form of ‘dynamic semantics’. In particular, sentence structure and its compositional meaning might just be the ‘tip of an iceberg’, representing a certain top level of formulating a slice of potentially much more complex behavior, or in a more mentalist picture: of thoughts. While our

aim here is not to endorse these radical consequences, we do feel they are an exciting complement to the more traditional literature on compositionality.

Finally, since even the radicalism of game semantics is well within the mind set of logicians and philosophers, in Section 9, we briefly remind the reader of the larger world around us, including perspectives on compositionality from such paradigms as distributional semantics and empirical cognitive science.

2 The minimal machinery of compositionality

Compositionality has invited general formulations that admit of mathematical definitions and results. Approaches include universal algebra (cf. [20] on algebraizability of logics), abstract recursion theory (compositionality as computability, [59]) and in particular, category theory, with Samson Abramsky’s long-standing work as a prominent instance. Bits and pieces of relevant formal theory can also be found elsewhere, for instance in the extensive body of work on translations between logical systems, [24], [79], especially, when viewing giving a semantics as a form of translation between object and meta-language.

Probably the most careful available general abstract analysis of compositionality is the one spelled out by Wilfrid Hodges, [44, 45]. To fix ideas, and for later reference and discussion, we give a quick overview in this section.

2.1 Compositionality defined

A set E of syntactically well-formed expressions is given, construed as *terms* in a partial term algebra $\mathbf{S} = (E, A, \alpha)_{\alpha \in \Sigma}$, generated from the atoms (lexical items) in a subset A of E by the partial operations in Σ .³ Adding a set of variables we get the usual notion of a *polynomial* $\pi[x_1, \dots, x_n]$ in the term algebra.⁴

A *semantics* is simply a map μ from E to some non-empty set M of values (‘ μ ’ for ‘meaning’).⁵ No constraints are placed on M ; it could contain truth values, possible worlds, individual objects or assignments, sets of these, elements in some given algebra, sets of formulas, sets of proofs, sets of strategies, etc. Importantly, μ can be partial: if its domain X consists of expressions whose meaning is not in doubt, the issue may be how to extend μ from X to all of E .⁶

³The partial term algebra approach is taken from [44]; the analysis presented in [45] uses a more abstract notion of a *constituent structure*.

⁴In Hodges’ original set-up there is also a homomorphism from E to surface strings, needed when the syntax allows (structural or lexical) ambiguities. Since our examples will mainly come from logic, we ignore the distinction between terms and strings here.

⁵The ‘switcher semantics’ of Kathrin Glüer and Peter Pagin generalizes the Hodges set-up to a *set* of semantic functions, with switching between them governed by linguistic context. [62] is an application to the semantics of quotation; a full presentation is [32]. We will not use this generalization here.

⁶Hodges’ set-up generalizes classical algebraic accounts of syntax and semantics, such as [55], in the following ways: (1) the syntax algebra is partial rather than many-sorted; (2) the meaning function μ can be partial as well; (3) there is no given semantic algebra (although if μ is compositional, such an algebra is induced on M); (4) meaning is not compositional by definition but a property that a semantics can have or not have.

Next, given the function μ , expressions/terms can be classified in terms of meaningfulness and of sameness of meaning: for $e, f \in E$,

- (1) $e \sim_\mu f$ iff for every polynomial $\pi[x]$ we have $\pi[e] \in X \Leftrightarrow \pi[f] \in X$
- (2) $e \equiv_\mu f$ iff $e, f \in X$ and $\mu(e) = \mu(f)$

Think of the equivalence classes of \sim_μ as semantic categories: sameness of category as preservation of μ -meaningfulness under substitutions. When $X = E$ they are also syntactic categories: reflecting the analysis of grammaticality in categorial grammars in terms of preservation under substitutions.

Here \equiv_μ is a partial equivalence relation (a *synonymy*) on E with domain X . This is all we need to define the substitution version of compositionality:

Definition 1 (substitutional compositionality)

A partial equivalence relation \equiv on E with domain X (so $X = \{e : e \equiv e\}$) is *compositional* if whenever $e_i \equiv f_i$ for $i = 1, \dots, n$, and $\pi[e_1, \dots, e_n], \pi[f_1, \dots, f_n]$ are both in X , we have $\pi[e_1, \dots, e_n] \equiv \pi[f_1, \dots, f_n]$. If \equiv is \equiv_μ for some semantics μ with domain X , we say that μ is *s-compositional*.⁷

The more familiar functional version of compositionality is as follows.

Definition 2 (functional compositionality)

A function μ is *f-compositional* if for every operation $\alpha \in \Sigma$ there is a corresponding operation r_α on M such that if $\alpha(e_1, \dots, e_k) \in X$, then $\mu(\alpha(e_1, \dots, e_k)) = r_\alpha(\mu(e_1), \dots, \mu(e_k))$.

F-compositionality of μ entails s-compositionality but presupposes, in contrast with the latter, that $X = \text{dom}(\mu)$ is closed under subterms. But when that holds, the two are equivalent. In this case—for instance, when μ is total—we can drop the s- and f- prefixes.

These classical formulations of compositionality assign meanings directly to linguistic expressions. No input from the linguistic or extra-linguistic context is made explicit. But taking contextual factors into account is in no way contrary to compositionality. The next subsection reviews how Hodges combines Frege’s Context Principle with compositionality. In Section 4 we discuss compositionality with explicit contextual parameters.

Digression (dependence). Definitions 1 and 2 can be seen as expressing notions of *dependence*. In particular, s-compositionality resembles a semantic intuition of dependence, as fixing values for an expression by values for its components. And f-compositionality resembles a widespread alternative intuition of dependence as definability by some explicit function. Notions of dependence

⁷This is of course similar to \equiv being a *congruence relation*; indeed, when X is closed under subterms the two notions coincide (see [82]). But in important applications, the domain of μ is not closed under subterms, and s-compositionality as defined here is then the most general notion of compositionality.

in logic will be discussed in Sections 6, 7 below. On the other hand, compositionality also involves *independence*, since meanings of expressions do not need meanings of syntactic material not occurring in the expression. But our later discussion of Contextuality in this section will show how the two extremes meet.

2.2 Contextuality and the Lifting Lemma

With the preceding in place, Hodges defines a third equivalence relation among terms, which is total even when \equiv_μ is partial. (We use ‘ F ’ for ‘Frege’.)

Definition 3 (fregean values)

(3) $e \equiv_\mu^F f$ iff $e \sim_\mu f$ and for all $\pi[x]$, if $\pi[e] \in X$ then $\mu(\pi[e]) = \mu(\pi[f])$

The equivalence class $|e|_\mu^F$ of e under \equiv_μ^F is called the *fregean value* of e .

Essentially, the fregean value of an expression e outside X is given by its contribution to the meanings (given by μ) of complex expressions in X containing e . This can be seen as a precise implementation of fregean Contextuality.

Say that μ is *cofinal* if every term in E is a subterm of some term in X .

Lemma 4 (Hodges’ Lifting Lemma)

Suppose that μ is cofinal, $\pi[e_1, \dots, e_n] \in E$, and also $e_i \equiv_\mu^F f_i$ for $i = 1, \dots, n$. Then $\pi[f_1, \dots, f_n] \in E$ and $\pi[e_1, \dots, e_n] \equiv_\mu^F \pi[f_1, \dots, f_n]$.

In particular, \equiv_μ^F is always compositional when μ is cofinal, so the (total) semantics $|\cdot|_\mu^F$ is compositional. And if two expressions of the same category (related by \sim_μ) have *different* fregean values, this is witnessed by a corresponding difference of μ -values of complex expressions containing them. Hodges calls the combination of these two properties *full abstraction*. Moreover, the fregean values are *unique* in the sense that, if another total semantics ν has the corresponding properties with respect to μ , then $\equiv_\nu = \equiv_\mu^F$.⁸

The relation \equiv_μ^F refines \equiv_μ in the sense that if $e, f \in X$, then $e \equiv_\mu^F f$ implies $e \equiv_\mu f$ (use the unit polynomial). If μ was s-compositional to begin with, and is also *husserlian* in the sense that $e \equiv_\mu f$ implies $e \sim_\mu f$,⁹ then the converse holds as well. In that case one can take (by choosing suitable representatives) the total fregean semantics to *extend* the partial μ to all terms.

Example 5 (First-order logic)

Let E be the set of first-order formulas (in some given signature) and $X = \text{dom}(\mu)$ the subset of sentences, for which μ provides the usual values: with a

⁸More precisely, consider these properties of a total semantics ν for E : (i) If $e \equiv_\nu f$ and $\pi[e] \in X$, then $\pi[f] \in X$; (ii) If $e \equiv_\nu f$ and $\pi[e], \pi[f] \in X$, then $\pi[e] \equiv_\mu \pi[f]$; (iii) If $e \not\equiv_\nu f$, then there is $\pi[x]$ s.t. either exactly one of $\pi[e], \pi[f]$ is in X , or both are and $\pi[e] \not\equiv_\mu \pi[f]$. It is clear that, if ν and ν' have properties (i)–(iii), then $\equiv_\nu = \equiv_{\nu'}$, and also that the fregean semantics has these properties.

⁹Hodges introduces this term in view of what Husserl writes about ‘Bedeutungskategorien’ in *Logische Untersuchungen*.

fixed model $\mathcal{M} = (M, I)$, $\mu(\varphi) = 1$ iff $\mathcal{M} \models \varphi$.¹⁰ μ is cofinal, s-compositional (replacing a true (false) *subsentence* of a sentence with another true (false) sentence doesn't change the truth value in \mathcal{M}), and, for $\varphi, \psi \in E$, $\varphi \sim_\mu \psi$ iff $FV(\varphi) = FV(\psi)$, so μ is trivially husserlian. By the above, the fregean semantics can be taken to extend μ to all formulas.

But the fregean values of formulas with free variables, i.e. their equivalence classes under \equiv_μ^F , aren't the usual tarskian values. Finding interesting semantic values is “where semantic theory takes off” [45, p. 270]. In the case of FOL, we already know what these values should look like: the function ν defined by

$$\nu(\varphi) = \langle FV(\varphi), [\varphi]_{\mathcal{M}} \rangle$$

where $[\varphi]_{\mathcal{M}} = \{s \in M^{FV(\varphi)} : \mathcal{M}, s \models \varphi\}$, has the desired properties (listed in footnote 8), and is essentially the usual compositional tarskian semantics.¹¹

Example 6 (Independence-friendly logic)

The language of Independence-Friendly Logic (IF) (in one of its versions, see [41] or [54]) is like that of FOL but now with additional quantifiers

$$\exists x / \{y_1, \dots, y_n\}$$

read intuitively as ‘there is an x which is independent of y_1, \dots, y_n such that ...’. The initial semantics μ for *sentences* was given in terms of games with imperfect information, but there is no obvious extension to formulas with free variables (since these games do not have an obvious notion of subgame). Nevertheless, μ is cofinal, husserlian, and can be shown to be compositional for sentences. So again the total fregean extension *exists*. But this time there is no well-known way to find ‘natural’ semantic values.

What Hodges does in [43] is recursively define a new satisfaction relation $\mathcal{M}, S \models \varphi$, where S is a *set* of assignments (now often called a *team*), such that

$$\nu(\varphi) = \langle FV(\varphi), \{S \subseteq M^{FV(\varphi)} : \mathcal{M}, S \models \varphi\} \rangle$$

has the desired properties, so its associated synonymy is the fregean synonymy.

Sets of assignments can capture the (in)dependencies between variables characteristic of IF logic. The Dependence Logic of [73] uses essentially the same satisfaction relation as Hodges but for a language with *dependence atoms*, which

¹⁰Everything we say below generalizes to the case when \mathcal{M} is instead a parameter of the relevant functions; see also Section 4.2.

¹¹We tend to believe since Tarski's 1933 truth definition that this is the ‘right’ semantics. But at that time, introducing *assignments* was a non-trivial achievement, not something one could simply read off an intuitive notion of truth for sentences. Instead of considering assignments in $M^{FV(\varphi)}$ we could use (as Tarski did) assignments in M^{Var} to *all* variables. The reason for incorporating $FV(\varphi)$ in the meaning of φ is to ensure the Husserl property: formulas with distinct free variables should differ in meaning (cf. φ and $\varphi \wedge x = x$).

Hodges takes ν to be the ‘usual’ semantics in the 2011 paper, but his 2001 paper uses (with \mathcal{M} as a parameter): $\nu'(\varphi) = \langle FV(\varphi), [\varphi]_{\mathcal{M}} \rangle$, where $[\varphi]_{\mathcal{M}} = \{\psi : \varphi \sim_\mu \psi \ \& \ \mathcal{M} \models \forall \bar{x}(\varphi \leftrightarrow \psi)\}$ (\bar{x} lists the variables in $FV(\varphi)$). Of course $\equiv_\nu = \equiv_{\nu'}$, and ν' is definable from \equiv_ν (since $\nu'(\varphi) = \langle FV(\varphi), \{\psi : \psi \equiv_\nu \varphi\} \rangle$); so maybe ν is more deserving of the label ‘usual’.

obviates the need for slashed quantifiers. A different approach to the use of dependence atoms will be discussed in Section 7.2 below.

2.3 Discussion

Hodges’ analysis is attractive for its sparseness, and seems to lay bare the essence of compositionality. Even so, questions remain. The initial μ -values represent a free initial choice, crucially affecting the resulting semantics, so the analysis does not produce the compositional fregean values out of nothing. And perhaps most importantly, the fregean values are defined holistically using the behavior of the whole language. But this is also a major virtue of the above style of analysis. Compositionality and Contextuality become two sides of the same coin, what we called earlier the ‘internal’ and ‘external’ perspectives on meaning are in harmony. One is reminded of the harmony in logicism between defining natural numbers as equivalence classes of equally large sets under bijections and as structured objects generated by arithmetical operations.¹²

Hodges’ analysis of compositionality remains at the abstract level of the fregean values. As we have seen, however, the tarskian semantics for FOL and the new team semantics for IF logic cannot be read off these general considerations, which only tell us under which circumstances fully abstract semantics *exist*. The main issue then becomes: what are interesting or useful denotations, and what are criteria for judging that? To get more of a grip on these issues, in what follows, we will discuss compositionality as a design principle in actual cases, identifying further strategies in designing semantic denotations, and eventually finding additional constraints on what counts as a good semantics.

3 Some classical challenges

Much of the history of logical semantics can be told as a series of responses to challenges to compositionality. In one version of this script, compositionality has to defend itself against various attacks from the (presumably, evil) realm of contextuality. As will be clear, this is not at all our view, since we stressed how compositionality and contextuality can co-exist without friction. However, there are still non-trivial challenges in how to *manage* this co-existence, and how much contextuality needs to be taken on board. In this section, we briefly survey a few earlier examples, while adding some new ones.

Intensions. Perhaps the first example is Frege’s introduction of the distinction between *Bedeutung* and *Sinn* in response to the Morning Star – Evening Star Paradox. The Morning Star is the Evening Star, but the discovery by the Babylonians of this identity was not the discovery of the trivial identity that The Morning Star is the Morning Star. Thus, substitution of identicals fails in

¹²More generally, Contextuality sounds more like a category-theoretic slogan: ‘know an object by its interactions with similar objects’, and as is well-known, such external characterizations often match more set-theoretic descriptions in terms of internal structure.

intensional contexts created by verbs like “discover”. In modern terms, Frege’s solution introduced *intensions* of expressions as denotations, in addition to the usual extensions, whenever compositional interpretation of a complex modal expression requires this. The same strategy can be seen at work in Montague’s famous analysis of natural language, where predicate meanings are intensionalized throughout to deal compositionally with constructions like “John seeks a unicorn”, or “The temperature is ninety but it is rising”. In contemporary modal logic, intensions are often identified with functions from possible worlds (indices of evaluation interpreted in whatever way is germane to the purpose at hand) to truth values, objects, or other entities in the relevant models.

Higher types. A classical case of ‘compositionalization’ is Montague’s use of generalized quantifiers to interpret noun phrases (a style of analysis from Categorical Grammar). This elegantly replaced earlier analyses like Russell’s in [68], which (in effect) blamed their non-compositionality on the ‘misleading’ subject-predicate form of natural languages. With complex denotations in a type hierarchy for linguistic expressions, subject-predicate form could be maintained and compositionality restored.

Assignments in logic and language. Yet another well-known set of compositionality issues occurs with the phenomena of anaphora and binding in natural and formal languages. Tarski’s semantics introduced the new parameter of *variable assignments*, in addition to models and truth, to account for the semantics of formulas with free variables. This move to an enriched setting returned in later work on natural language, but now in new ways.

DRT and dynamic semantics. In particular, the Discourse Representation Theory of [49] analyzes anaphoric relationships in a more linguistically sensitive manner than Montague as driving a process of building successive discourse representation structures for linguistic expressions and texts. Typically, discourse representation structures involve markers for objects representing the anaphoric structure of the text. And when it turned out that the process of building such representations was not compositional, compositional solutions were found after all. For instance, in the dynamic semantics of [36] expressions no longer denote sets of assignments as with FOL, but sets of pairs of assignments, viewed as state transitions in an evaluation process.¹³

To be sure, there is also an independent non-methodological motivation for adopting a dynamic semantics, namely, providing an account of meaning that records more features of actual language use, as in Austin’s dictum:

“Words are what words do”, [13].

Current dynamic semantics also take on board information update by speech acts, or information exchange between participants in communication, [58].

It should be noted that, as we shall see again and again in this paper, compositionality by itself does not force a choice here. A dynamic semantics need not replace existing semantics involving more static denotations. It might also

¹³For the full story of compositionalizing DRT, see [47].

be seen as modelling other aspects of language use such as the fine-structure of evaluation, or in other versions of dynamic semantics, the process of communication. Moreover, mathematically, both static and dynamic structure may be of interest. For instance, while the standard semantics for FOL involves sets (i.e., unary properties) of assignments, its dynamic semantics involves binary relations between assignments. Thus, we now have two sorts of algebras: cylindric algebras of standard denotations and relational algebras for dynamic denotations. The two realms have interesting connections, and developing them in tandem throws more light upon language than choosing just one.

At this point, we reach a bridge to another field. The dynamic semantics solution was inspired by the Hoare-style semantics of imperative programming languages, where programs denote the pairs of input and output states of their successful terminating executions. These and other parallels with computer science, known since the 1970s, bring us to a general issue that we think should inform modern discussions of compositionality more generally: semantic ideas about programming languages and computation can enrich more traditional discussions of compositionality in philosophy and linguistics.

Semantics of programs. The issues and techniques used in analyzing compositionality are not confined to logic and linguistics. Striking parallels were observed in [48] with the semantics of programming languages, where compositional design, in various guises, had been a desideratum in the structured programming methodology of Dijkstra, Hoare, and Scott-Strachey; see [38]. For instance, Montague-style intensionalization was shown to be a good strategy for dealing with the semantics of arrays and other data structures. [47] discusses program semantics in terms of mathematical results about initial algebras and polynomials in universal algebra. The earlier-noted coexistence of static and dynamic perspectives returns in computer science as the contrast between *denotational* and *operational* semantics for a programming language, where one often wants both styles plus an analysis connecting the two, [6]. For Abramsky’s powerful category-theoretic approach to compositional analysis and design of programming languages, see Section 8. In this richer setting, computing generalizes to interactive *game play*, which brings us to our last example.

IF logic revisited. Games are a persistent theme in logic and natural language. A famous compositionality dispute concerns Jaakko Hintikka’s game-theoretic perspective on natural languages, first put forward in [40] and developed further in [41]. Hintikka argued that his IF logic is necessarily non-compositional, a claim that was refuted by Hodges’ compositional team semantics for IF logic in the preceding section. Even so, there was also a second claim that game-theoretic semantics fits natural language better than tarskian semantics, since it proceeds ‘outside-in’, something that also fits with Abramsky’s view of computation as unfolding interactive agency. Again this leads us to criteria for a ‘good’ compositional semantics. Is team semantics also a refutation of Hintikka’s second claim? We shall return to these issues in the sections to come.

4 Currying

The compositional semantics mentioned so far, as well as many other cases, all work on a schema of (a) *introducing new parameters of evaluation*, and (b) *giving semantic truth conditions* for the given language in terms of these. This is obvious in the, by now, standard semantics of first-order logic or modal logic. But (a) and (b) keep arising. One more recent example is ‘public announcement logic’ for information update, [66], [15], where the language contains a new kind of dynamic modalities that refer to what is true in models different from the initial one. A typical PAL truth condition now reads as follows:

$$(4) \quad \mathcal{M}, s \models [!\varphi]\psi \text{ iff } \mathcal{M}, s \models \varphi \text{ implies } \mathcal{M}|_{\varphi}, s \models \psi$$

where $\mathcal{M}|_{\varphi}$ is the model \mathcal{M} *restricted* to the states in which φ holds. In contrast with other familiar logical systems, here the current model is a parameter that cannot be held constant in the truth definition (just as the assignment parameter cannot be held constant in the FOL clauses for quantified formulas).

When more ‘active’ parameters are added, is the resulting parametric semantics compositional? To even ask this question, we a notion of compositionality that takes the presence of parameters into account.

4.1 Contextual compositionality

The following generalization of compositionality has been around in the literature for a while (see e.g. [60]). Let μ be a function taking expressions in E (terms in our syntactic algebra \mathbf{S}) as arguments and in addition a sequence $\bar{p} = p_1, \dots, p_m$ of parameters, say, $p_i \in P_i$ for $i = 1, \dots, m$. Again, μ can be partial in the first argument, so it is a function from $X \times \bar{P}$ to M , where $X \subseteq E$ and $\bar{P} = P_1 \times \dots \times P_m$.

Definition 7 (contextual compositionality)

- (a) μ is (contextually) *f-compositional* if, for every operation $\alpha \in \Sigma$ there is a corresponding operation r_α on M such that, whenever $\alpha(e_1, \dots, e_k) \in X$ and $\bar{p} \in \bar{P}$, then $\mu(\alpha(e_1, \dots, e_k), \bar{p}) = r_\alpha(\mu(e_1, \bar{p}), \dots, \mu(e_k, \bar{p}), \bar{p})$.
- (b) μ is (contextually) *s-compositional* if, whenever $\mu(e_i, \bar{p}) = \mu(f_i, \bar{p})$ for $i = 1, \dots, n$, and $\pi[e_1, \dots, e_n]$ and $\pi[f_1, \dots, f_n]$ are both in X , we have $\mu(\pi[e_1, \dots, e_n], \bar{p}) = \mu(\pi[f_1, \dots, f_n], \bar{p})$.¹⁴

Again, (a) and (b) are equivalent when X is closed under subterms.

Ordinary compositionality is a special case of contextual compositionality. *Given* some parameters ranging over given contextual features, the semantic value function μ can be compositional, or not. In fact, very often it is *not* compositional in the above sense, even though μ is recursively defined. But

¹⁴The literature also considers a stronger version, without \bar{p} as an argument of r_α . However (despite some claims to the contrary), this version has no corresponding natural substitutional version. In particular, it is *not* equivalent to the statement that whenever $\mu(e_i, \bar{p}) = \mu(f_i, \bar{q})$ for $i = 1, \dots, n$, and $\pi[e_1, \dots, e_n], \pi[f_1, \dots, f_n] \in X$, $\mu(\pi[e_1, \dots, e_n], \bar{p}) = \mu(\pi[f_1, \dots, f_n], \bar{q})$.

in this case, strict (ordinary or contextual) compositionality can be enforced by lambda abstraction over a suitable set of parameters, under very general circumstances. We now proceed to make this observation precise.¹⁵

4.2 A currying recipe

Currying arguments of an n -place function F where $n > 1$ is essentially just abstracting over them. For example, if $F: P_1 \times \dots \times P_n \rightarrow M$ and the second and fourth arguments are curried, the result is the 2-place function $F': P_1 \times P_3 \rightarrow M^{P_2 \times P_4}$ defined by

$$F'(p_1, p_3)(p_2, p_4) = F(p_1, \dots, p_n)$$

So if we let $\mu(\varphi, \mathcal{M}, s) = 1 \Leftrightarrow \mathcal{M}, s \models \varphi$ in the PAL example above, then μ is not (contextually) compositional, but the function obtained by currying both \mathcal{M} and s is compositional; this will follow from Fact 9 below.¹⁶

Similarly, as we have seen, currying the assignment argument, but not the model argument, in the standard Tarskian definition of satisfaction in first-order logic yields a compositional semantic function whose values, relative to a model, are (characteristic functions of) sets of assignments. Let us look at one more example to see the general pattern.

Example 8 (Propositional Dynamic Logic)

In the two-component language of PDL, [39], μ must interpret both formulas and programs, the latter as (characteristic functions of) binary relations between states. That is, $\mu(\varphi, \mathcal{M}, w_1, w_2) = 1 \Leftrightarrow w_1 = w_2 \ \& \ \mathcal{M}, w_1 \models \varphi$, and $\mu(\rho, \mathcal{M}, w_1, w_2) = 1 \Leftrightarrow w_1 R_\rho w_2$. Some relevant defining clauses:

- $\mu([\rho]\varphi, \mathcal{M}, w_1, w_2) = 1$ iff
 $w_1 = w_2 \ \& \ \forall w' (\mu(\rho, \mathcal{M}, w_1, w') = 1 \rightarrow \mu(\varphi, \mathcal{M}, w', w') = 1)$
- $\mu(\delta_1; \delta_2, \mathcal{M}, w_1, w_2) = 1$ iff
 $\exists w (\mu(\delta_1, \mathcal{M}, w_1, w) = 1 \ \& \ \mu(\delta_2, \mathcal{M}, w, w_2) = 1)$
- $\mu(\rho^*, \mathcal{M}, w_1, w_2) = 1$ iff
 $\exists n \exists z_0 \dots \exists z_n (z_0 = w_1 \ \& \ z_n = w_2 \ \& \ \forall i (i < n \rightarrow \mu(\rho, \mathcal{M}, z_i, z_{i+1}) = 1))$

Again, μ is not contextually compositional, but currying w_1 and w_2 restores compositionality. The reason is that w_1 and w_2 are ‘shifted’ in the inductive definition of μ , but \mathcal{M} is not.

¹⁵The role of currying for compositionality was first discussed in [52], generalized in [63] and [83], and is elaborated still further here.

¹⁶There is a slight technical complication. Let Mod be the class of PAL-models of the form $\mathcal{M} = (W, \{R_a\}_{a \in A}, V)$, and let $|\mathcal{M}| = W$. μ as described is really a (total) function from $\{(\varphi, \mathcal{M}, s) : \varphi \in E \ \& \ \mathcal{M} \in Mod \ \& \ s \in |\mathcal{M}|\}$, which is not strictly speaking a cartesian product. It is trivial but slightly cumbersome to redefine the domain so that it becomes such a product. In the following we ignore this subtlety.

The general pattern behind these examples is this. μ is defined by an inductive parametric ‘truth definition’ with (apart from clauses for atoms) one clause for each syntactic rule $\alpha \in \Sigma$, which we write schematically as follows:

$$(5) \quad \mu(\alpha(e_1, \dots, e_k), \mathcal{M}, \bar{p}) = u \text{ iff } \Phi_\alpha[\dots, \mu(e_i, \mathcal{M}, \bar{t}_i[\bar{p}']), \dots, \mathcal{M}, \bar{p}, u]$$

Here $\bar{p} = p_1, \dots, p_m$ are parameters and \mathcal{M} is the relevant model parameter; it really should be one of the p_j but is made explicit here to make the format more recognizable. Φ_α is written in some suitable set-theoretic metalanguage.¹⁷ The complex terms of the form $\mu(e_i, \mathcal{M}, \bar{t}_i[\bar{p}'])$ that may occur in Φ_α are to be understood as follows. $\bar{t}_i[\bar{p}']$ is a sequence of m (possibly) complex terms $t_{i1}[\bar{p}'], \dots, t_{im}[\bar{p}']$, in which parameters among $\bar{p}' = p'_1, \dots, p'_k$ (for some k') may occur. Here each p'_l is either one of p_1, \dots, p_m or a new parameter, which is bound inside Φ_α (by a quantifier or other variable-binding operator).

To help in formulating a general observation, let us say that a parameter p_j is *fixed* in Φ_α , if whenever p_j occurs in a term $\mu(\alpha_i, \mathcal{M}, \bar{t}_i[\bar{p}'])$, we have $t_{ij}[\bar{p}'] = p_j$. If that is not the case, we say that p_j is *shifted*.

It should be fairly clear that this format covers most common truth definitions, in particular those exemplified above with the logical systems FOL, PAL, and PDL.¹⁸ Now we have the following fact.

Fact 9

Let the semantics μ be defined with a recursive clause of the form (5) for each syntactic operator.

- (a) *If some parameter in some clause is shifted, then (under very general circumstances) μ is not contextually compositional.*
- (b) *If μ' is obtained by currying all shifted parameters, and possibly other parameters as well, then μ' is contextually compositional. In particular, if all parameters are curried, or if no parameter is shifted, μ' is compositional.*

Summing up, in most recursive truth definitions familiar from logic and formal semantics, the parametric meaning function is not strictly speaking compositional, but currying shifted parameters ensures compositionality, although at the cost of using higher-order semantic values.

4.3 Connections: currying and fregean semantics

The compositional tarskian semantics for FOL is obtained by currying the assignment parameter. But it also resulted—up to synonymy—with the Lifting

¹⁷It contains names for the various terms on the right-hand side of (5), but for ease of reading we haven’t distinguished objects from their names here.

¹⁸For example, consider the second PDL-clause above. E can be taken to be the union of the disjoint sets of formulas and programs, inductively defined from atomic proposition letters and program symbols, so μ is total. There is a syntactic operation $\alpha \in \Sigma$ such that $\alpha(\delta_1, \delta_2) = \delta_1; \delta_2$. The semantic clause then has the form $\mu(\alpha(\delta_1, \delta_2), \mathcal{M}, w_1, w_2) = 1$ iff $\Phi_\alpha[\mu(\delta_1, \mathcal{M}, \bar{t}_1[\bar{w}']), \mu(\delta_2, \mathcal{M}, \bar{t}_2[\bar{w}'])]$ where $\bar{w}' = w_1, w_2, w$, w is quantified, and $\bar{t}_1[\bar{w}'] = t_{11}[\bar{w}'], t_{12}[\bar{w}']$, where $t_{11}[\bar{w}'] = w_1$ and $t_{12}[\bar{w}'] = w$; similarly for $\bar{t}_2[\bar{w}']$. Since $t_{12}[\bar{w}']$ is not w_2 , and $t_{21}[\bar{w}']$ is not w_1 , both w_1 and w_2 are shifted in the clause.

Lemma in Section 2.2. There is in fact a natural connection between the two methods that may be worth spelling out.

In this subsection, let μ be a total reursively defined parametric semantics, and let μ' result from currying all the parameters: $\mu'(e)(\bar{p}) = \mu(e, \bar{p})$. As we know, μ' is compositional. For simplicity, we also assume that μ is ‘husserlian’, in the sense that if $\mu(e, \bar{p}) = \mu(f, \bar{p})$, then $e \sim f$.¹⁹

Next, for each tuple \bar{p} , define the unary function $\mu_{\bar{p}}$ by setting

$$\mu_{\bar{p}}(e) = \mu(e, \bar{p})$$

The function $\mu_{\bar{p}}$ is in general *not* compositional. In fact, we have:

Fact 10

μ is (contextually) compositional iff each $\mu_{\bar{p}}$ is compositional.

Hodges’ applications of the Lifting Lemma concern extending a partial compositional semantics to a total one. But the Lemma holds also when the given semantics is total, except in this case it has effect only when that semantics is *not* compositional, as with the $\mu_{\bar{p}}$. Recall the fregean synonymies $\equiv_{\mu_{\bar{p}}}^F$ from Definition 3. The following holds in general.

Fact 11

$$\equiv_{\mu'} = \bigcap_{\bar{p} \in \bar{P}} \equiv_{\mu_{\bar{p}}}^F$$

This connects the curried compositional semantics with the fregean one presented in Section 2, be it in a rather weak way. Here is a concrete example showing how, with some extra assumptions, the connection becomes much stronger.

Say that our language *contains propositional logic* if it has a category of formulas (using φ, ψ, \dots for these) whose values under μ , relative to a model \mathcal{M} and other parameters, are 0 or 1, and μ respects tautological consequence.²⁰ Also, a *universal operator* U is a unary formula operator such that $\mu(U\varphi, \mathcal{M}, \bar{p}) = 1$ iff for all \bar{q} in \mathcal{M} , $\mu(\varphi, \mathcal{M}, \bar{q}) = 1$. Now fix \mathcal{M} and parameters \bar{p}_0 , and let $\mu_0 = \mu_{\mathcal{M}, \bar{p}_0}$. Again, μ_0 is usually not compositional.

Fact 12

If μ contains propositional logic and has a universal operator, then

$$\equiv_{\mu'} = \equiv_{\mu_0}^F$$

So in this case the curried semantics *is* the fregean semantics (up to synonymy). FOL is an example satisfying the assumptions; another is intensional

¹⁹We can write $e \sim f$ rather than $e \sim_{\mu} f$, since μ is total, so $e \sim f$ just means that e and f have the same category in the syntactic algebra \mathbf{S} .

²⁰That is, writing $\llbracket \varphi \rrbracket_{\mathcal{M}}^{\mu}$ for $\{\bar{p} : \mu(\varphi, \mathcal{M}, \bar{p}) = 1\}$, we have:

$$\text{If } \Gamma \vdash_{\text{PL}} \varphi, \text{ then for all } \mathcal{M}, \bigcap_{\psi \in \Gamma} \llbracket \psi \rrbracket_{\mathcal{M}}^{\mu} \subseteq \llbracket \varphi \rrbracket_{\mathcal{M}}^{\mu}.$$

logic with a universal modality. Relative to a model \mathcal{M} , the compositional semantics whose values are sets of assignments is—up to synonymy—what you get by applying Hodges’ construction to the non-compositional semantics μ_0 , where s_0 is a fixed assignment and $\mu_0(\varphi) = 1$ iff $\mathcal{M}, s_0 \models \varphi$.

5 General issues in compositional semantics

We now move to state some rules of thumb that seem to apply when searching for a good compositional semantics. Our discussion results in a check-list of points to keep in mind that occur in many concrete cases.

5.1 The danger of over-interpretation: Do Not Invert

A striking phenomenon found with many proposed compositional semantics is a tendency to invert the moral of the solution. For instance, in propositional modal logic, propositions denote intensions, i.e., functions from indices to truth values. It then seems reasonable to identify propositions with sets of worlds. And once this is done, it is tempting to identify the set of propositions with the complete set of all sets of worlds. But this is unwarranted. The compositional solution merely says that propositions correspond to some sets of worlds: letting in all of them is a major step to a second-order perspective on the semantics.

Here, and in many similar well-known semantics, a good recommendation is the following counsel of caution:

- (I) *Do Not Invert.*

Indeed, for modal logic, there is a philosophical rationale to merely correlating the intuitive notion of a proposition with a set of worlds without making a total converse identification that would produce huge amounts of undefinable and unusable propositions.²¹ And there is also a clear mathematical rationale for the same reticence. The representation theory of modal algebras works with ‘general frames’ that have merely a designated set of sets of worlds as algebraic values, not necessarily the full power set.²² Similar points can be made about assignments in first-order semantics and functions from variables to objects, where there is no need to assume that the full function space of all maps from variables to objects is available as indices of evaluation. The latter case study will be discussed at length in Section 6.3 below, after first pointing to some further reasons for heeding the Do Not Invert injunction.

²¹Likewise, in current ‘set liftings’ of possible worlds to ‘state-based’ hyper-intensional semantics, there is no need to assume that all sets are available as states, or structured situations.

²²Stated in other terms, when quantifying over propositions in modal logic, Henkin models may be preferable to full second-order models, cf. [12].

5.2 Avoiding triviality

Every proposal for compositionality, even when of the non-inverted kind, needs to face a threat of triviality. A compositional semantics for any given language can almost always be found with enough industry in manipulating semantics or syntax. Thus, a very general precaution is this:

(II) *Don't make too many, or too few, meaning distinctions.*

There might be pre-theoretic intuitions about *how many* meanings should be available, or put in terms of expressive power, (at least) how many distinctions the given language should be able to make. At the extreme end of making too few distinctions: if every expression means the same, compositionality holds trivially. And at the other end: if all expressions mean different things, compositionality is again trivial. Semantics of the latter kind have in fact been proposed, for example, by making the expression itself a part of its meaning. This holds for one suggestion in [84]; see [81] for further discussion.

We note in passing that the last observation also cautions us that compositionality in itself is no guarantee for learnability, communicability, or any other of the good properties linguistic meanings are supposed to have. A language with no synonymy relation except identity may well be completely unlearnable. Compositionality is at best a *necessary* condition.

5.3 The role of impossibility results

Avoiding triviality is good, but how? The following positive constraints are sometimes at work in the search for a new semantics to replace an old one, e.g. because it was partial, or not compositional.

(III) *For some subset X of expressions, the proposed semantics should agree with the benchmark given by some prior semantics for X .*

(IV) *The values of the proposed semantics should be of a specified kind.*

(III) applies to Hodges' set-up in Section 2. It can be taken in a strong and a weak sense: (III.i) the new semantic *values* of expressions in the set X should be the same as the old ones, or (III.ii) the new semantics should preserve given *synonymies*, or non-synonymies, in X . Together with (IV), mathematical impossibility results can sometimes guide the search for a good semantics. We look at three familiar examples.

Example 13 (IF logic in a bit more detail)

Let E be the set of IF-formulas, X the subset of IF-sentences, and $\mu_{\mathcal{M}}(\varphi)$ the truth value of a sentence φ in a model \mathcal{M} according to Hintikka's game-theoretic semantics. As we saw in Example 6, the existence of a compositional (fregean) semantics for all of E which agrees with $\mu_{\mathcal{M}}$ on sentences (in the sense of (III.i)) is then guaranteed. [23] showed that here constraint (IV) makes a difference: there is no compositional semantics for E which agrees with $\mu_{\mathcal{M}}$ on sentences

and whose values on formulas are sets of assignments in \mathcal{M} . In particular, it showed that if ν is total, compositional, and agrees with μ on sentences, then for any $n \geq 2$ there is a model \mathcal{M} with n elements such that the number of distinct values $\nu_{\mathcal{M}}(\varphi(x))$ of IF-formulas with one free variable is (much) greater than 2^n . Thus, these values cannot all be subsets of M .²³

This impossibility result in a sense vindicates Hintikka’s claim that IF logic has no compositional semantics. But note how strong constraint (III) is here. In effect, it builds the expressive power on sentences of IF logic (which is that of Existential Second-Order logic), and hence the IF consequence relation, into any candidate semantics. If the main goal is to extend FOL in order to deal explicitly with (in)dependence between variables, that may already seem too much. We sketch an alternative, more sensitive approach in Section 7.2.

Finally, observe that if we change the force of (III) to just require agreement with μ on the set Y of FOL-formulas—which holds for the team semantics of IF logic—there is no similar impossibility result. In fact, various uninteresting total compositional extensions of μ , satisfying this version of (III) and also (IV), then exist. Note that Y is not cofinal, rather, E is an *end-extension* of Y . As [44] observes, it follows, since μ is also husserlian, that the *one-point extension* μ^1 of μ , which gives formulas in Y the old values, but all formulas in $E - Y$ the same distinct value $*$, is in fact the fregean extension (up to synonymy).²⁴ Of course, no one would consider μ^1 a useful semantics. Moral: by themselves, conditions (III) and (IV) do not guarantee non-triviality.

Example 14 (Modal logic)

Next, let E be the set of basic modal formulas, generated by

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi$$

Here we might take (III) to require that \neg and \wedge preserve their standard meaning. A problem with this, which did not arise in the IF example, is that we haven’t made clear what a *model* is supposed to be. Until we do, neither is the ‘standard’ meaning of \neg and \wedge clear. For example, if we generalize the two-valued PL semantics to n -valued semantics, which is a way to implement (IV), there are many candidates for a standard meaning of these connectives. But this is actually not a problem, since there are well-known impossibility results: most modal logics are not determined by any finite-valued truth tables, [14]. Historically, a many-valued approach to modal logic was indeed attempted (by Łukasiewicz), interpreting $\Diamond\varphi$ using a separate truth value for ‘possible’, but the impossibility results put an end to that endeavor.

Note, however, that we have now moved from semantic constraints to logical ones. If a logic is seen as a set of (valid) formulas closed under certain inference rules, or more generally a consequence relation, an impossibility result of the kind mentioned says that no compositional semantic function of a particular

²³[31] generalizes this result to infinite models; this needs an extra requirement on ν , to do with how $\nu_{\mathcal{M}}$ relates to $\nu_{\mathcal{M}'}$ when \mathcal{M} and \mathcal{M}' are isomorphic.

²⁴Without the Husserl property, a total compositional extension can also be shown to exist, but this is less trivial; see [82].

kind is *sound and complete* for that logic. This is a strong interpretation of the following kind of constraint that we formulate more generally and loosely:

(V) *A new compositional semantics must respect some given (non-)inferences.*

We will return to this constraint in much greater detail in the next section.

Example 15 (Intuitionistic logic)

The case of intuitionistic propositional logic IPL is similar to the preceding example, except there is no addition to the syntax, we still have: $\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid (\varphi \rightarrow \varphi)$. Classical models (bivalent truth value assignments) are unacceptable because they yield validities deemed false. And already [33] proved that no finite-valued truth tables determine IPL.²⁵ Again, this is a strong use of (V). A weaker use might be to require e.g. that $\varphi \wedge \neg\varphi$ is not valid, or that IPL validity agrees with PL validity for formulas without \vee and \rightarrow . All of these constraints concern inference. Semantic requirements like the above (III) seem harder to make sense of in this case, absent a common notion of model.

The negative results mentioned in the last two examples ruled out a particular kind of compositional semantics, showing that finite ‘matrices’ do not determine these logics. On the positive side, study of the ‘matrix method’ using infinite matrices led to (compositional) *algebraic* semantics; see [34], [14].

Coda: other perspectives.²⁶ Our discussion highlighted inference, but other technical perspectives exist. [75] discussed a compositional extension problem similar to the one in Section 2.2: when can a new syntactic operation be added to a homomorphically interpreted syntactic algebra satisfying certain equational constraints, in a way that does not ‘disturb’ the given interpretation? The setting of the results there was classical model theory; it would be interesting to rethink them in the more abstract Hodges framework.

5.4 Discussion

In order to evaluate a particular compositional solution, more specific versions of (III) – (V) must come into play. If we view Compositionality as an empirical claim, one consideration is evident: does the proposed semantics mirror how the given empirical phenomenon works? An instance of this is Hintikka’s earlier-mentioned criticism of proposed compositional semantics for his IF logic. On his view, it is an evident empirical feature of natural language interpretation

²⁵Interestingly, both IPL and basic modal logic still stay close to finite truth tables in that they have the *finite model property*: each satisfiable formula is true at some point in some finite model. Locally, a set of finite truth tables still works.

²⁶Other constraints than (III) – (V) have been considered in the linguistic literature, from strict requirements on syntax—e.g. that the only operation is concatenation, possibly also ‘wrapping’, or that grammar rules should be context-free—to specific ideas about the corresponding semantics—e.g. that concatenation should correspond to function application. ([84] achieves this, at the cost of using non-wellfounded sets, but also of making the meaning function 1-1, thereby violating (II).) These constraints are closely tied to specific theories of the syntax-semantics interface in natural language, and will not be discussed here.

that it proceeds outside–in, so any compositional semantics suggesting that the process works inside–out is off the mark. We return to the interesting contrast between outside–in and inside–out accounts in game semantics in Section 8, but in principle we find this sort of disagreement entirely legitimate.

So, did Hodges’ team semantics refute Hintikka’s claim that IF logic is non-compositional? As we have seen, this is no simple yes/no question, and pondering answers highlights issues about what a good semantics is meant to do. A strict application of (III) and (IV) rules out team semantics, but we have noted that one may keep (IV) (semantic values as sets of assignments) while relenting a bit on (III), and still get a compositional semantics. This however changes both syntax and logic, not much, but enough to make the logic decidable; Section 7.2 below has details. And the outside–in vs. inside–out issue is yet a further aspect. So the jury seems to be out on the merits of Hintikka vs. Hodges. The best answer for now is: It depends.

In general, the plausibility of type (IV) constraints must come from independent prior intuitions about what meanings are supposed to be (or do). For Tarski’s assignments, a credible case can be made that these correspond to our intuitive picture of the contexts in which linguistic utterances takes place; cf. Discourse Representation Theory, or the situation-semantic view of utterance situations as containing realistically interpreted ‘anchors’ [17]. But, say, whether modal intensions match Frege’s intuitive pre-theoretical notion of Sinn is already a much more debatable issue.

Finally, more general methodological considerations about utility, fruitfulness, or explanatory power are clearly relevant too when evaluating a proposed compositional semantics. Making these precise raises difficult perennial questions in the general philosophy of science, which we shall not go into here. Instead, in the next section we continue with major logical dimensions in judging proposed compositional semantics, in terms of language design, inference, and general views of what a semantics should deliver.

6 Compositional semantics and inference

6.1 Empirical perspectives

Requirements (V) and the closely related (III) above have a solid tradition. Around 1980, Barbara Partee emphasized that the purpose of a semantics of natural language is not just to explain how expressions acquire their meanings, but also to account for our prior intuitions about valid or invalid inferences. Language users have intuitions about inference (one reason why we can teach logic courses without mass revolts), though they will usually be partial, rather than total. Any proposed semantics should then at least agree with these intuitions, while providing a plausible extrapolation to inferences not covered by the original intuitions. In Partee’s terms:

“Inferences are part of the data”.²⁷

We have already discussed different ways of taking this. One is that particular concrete sentences are to imply each other, say, “John walks slowly” should imply “John walks”. Similarly for preservation of (non-)synonymies. Of course, which sentences are actually synonymous is always up for debate. Some scholars tend to the view that, with the right sense of “synonymous”, there are almost no non-trivial synonymies in natural languages. As we have seen, this trivializes the compositionality problem. We need notions of synonymy that, ideally, support the good work compositionality is supposed to do. But there are still many choices. One might be to identify it more or less with logical equivalence, but more fine-grained (‘hyperintensional’) variants also make sense. For example, one might well consider the formulas p and $p \wedge p$ to be synonymous, but deny that $\neg(p \wedge q)$ is synonymous with $\neg p \vee \neg q$.²⁸

However, preserving particular (non-)synonymies may seem too modest, and it also does not suggest general design features for a semantics. Intuitions in the semantic literature are often stated in more general schematic terms such as the following principle of upward Monotonicity for “all”:

From any sentence of the form “All A are B ”, one may infer that
“All A are C ”, if C is a ‘weaker’ predicate implied by B .

This principle is intelligible to language users, it covers infinitely many cases, and it puts clear constraints on the semantic interpretation of the quantifier expression “all”, [65]. Likewise, intuitively, some schema may be judged invalid, in which case the proposed semantics should produce at least one counter-example (and probably even more: a general explanation of the non-validity).

[75] outlines a formal approach to the preservation of (non)inference in terms of the existence, or lack thereof, of compositional translations of the source language, equipped with a given partial inference relation, and a partial non-inference relation, into a target formal language with standard notions of consequence and non-consequence. The author shows that the risk of triviality lurks here as well: without significant constraints, such translations may always exist.

It would be interesting to analyze the entanglement of compositionality and intuitions about inference in natural language at the abstraction level of Hodges’ minimal requirements for compositionality presented in Section 2 above. However, we leave this matter for further investigation.

Contextuality. As a final comment, the present topic of preserving intuitions about valid and invalid inference patterns in compositional semantics again demonstrates our theme of the entanglement of compositionality with *contextuality*. Inference patterns have infinitely many instances going far beyond single expressions of the language. And more importantly, inferences from a single given sentence will usually go beyond subformulas, or even more general syntactic ‘parts’ of that formula. In other words, looking for a compositional se-

²⁷We leave aside the trickier philosophical issue of what it means for a formal semantics to explain the prior inferential intuitions.

²⁸[65, Chapter 11.3] has a detailed discussion of various notions of synonymy.

mantics under inferential constraints on the valid reasoning to come out of the semantics takes on board contextuality, in both senses discussed in Section 1.

6.2 A logical systems perspective

Any semantics supports a notion of universal validity and valid consequence. Often, it supports even more than one natural candidate, as different notions of consequence can be defined over the same class of models with respect to the same language, [19]. While in cases such as the semantics of classical first-order logic, people often just accept the received ‘standard’ logical system, another viewpoint is possible. One might think that the valid consequences to come out of the semantics are themselves a criterion for judging whether the proposed compositional solution achieves what we want it to do. This is of course just the intent of Point (V) in the general discussion of Section 5, which highlighted the important, though usually only partial, nature of prior inferential constraints. We now take this theme a bit further in terms of logical systems.

A compositional semantics must always be analyzed critically for the validities it produces: as long as there is some latitude in setting up the semantics, these are not ‘forced’ uniquely. Especially with the earlier-mentioned schematic view on inferential constraints, Point (V) leads us to thinking in terms of full-fledged *logical systems*. And in such a perspective, further system considerations may come in beyond accounting for individual valid or non-valid schemata, especially when prior intuitions are incomplete, and admit of precisification and extrapolation. Such considerations may be general system properties of intelligibility blocking opaque or ad-hoc ways of coding up given inferential intuitions. But also further criteria such as desirable theoretical meta-properties of the system, or more practical computational complexity of the proposed proof system can play a legitimate role. We now proceed to a case study for this system-oriented style of thinking, which also brings further issues of its own.

6.3 Case study: Tarski semantics revisited

The Do Not Invert maxim (I) also applies to the semantics of first-order logic. While predicate-logical formulas need variable assignments as indices of evaluation, the requirement of compositionality per se does not force us to assume that the full space of all functions from variables to objects is available as assignments. Letting go of this inversion leads to the broader class of ‘generalized assignment models’

$$\mathcal{M} = (D, I, A)$$

consisting of a first-order $\mathcal{M} = (D, I)$ plus a set A of ‘available’ assignments. Motivations for this move include algebraic simplicity [57], but also the natural phenomenon of dependence, which is beyond the range of standard first-order semantics: with ‘gaps’ in the function space, changes in the value for one variable may be correlated with changes in value for another variable. Standard

Tarski models are not lost, being the special case where $A = D^{VAR}$. In these generalized models, we set

$$\mathcal{M}, s \models \exists x\varphi \text{ iff there exists } t \in A \text{ with } s =^x t \text{ and } \mathcal{M}, t \models \varphi$$

where $s =^x t$ says that $s(y) = t(y)$ for all variables y distinct from x .

The result is a perfectly compositional semantics for the first-order language. Even so, it has some unusual features. For instance, given the preceding stipulation, the truth value of a formula φ may depend on values that the current assignment gives to variables that do not occur in φ ; see below.²⁹

First-order logic deconstructed. The motivation for the preceding semantics had to do with the theme of inference. It might be expected that the basic logic of quantification is simple, like that of the Boolean operations, and the undecidability of first-order logic then comes as a surprise. So, what are its causes? An analysis requires decoupling the goal of a compositional semantics from the high cost in complexity of the notion of validity. In the above analysis, the latter complexity arises from the negotiable mathematical ‘wrappings’ of Tarski’s definition, namely, the assumption of a regular second-order mathematical object, viz. the presence of the full function space D^{VAR} . After all, we know from Gödel’s theorems that the first-order theory of regular mathematical structures may be complex. Here is how this deconstruction works.

Fact 16

The logic CRS consisting of the first-order validities on generalized assignment models has a simple complete axiomatization and is decidable.

On this basis, the semantic content of additional first-order validities can be determined. A non-valid principle in generalized assignment models such as $\exists x\exists y\varphi \rightarrow \exists y\exists x\varphi$ expresses *independence* of the variables x, y . More precisely, it imposes the following condition on the set A of admissible assignments: whenever $s =^x t =^y u$, there exists an assignment v with $s =^y v =^x u$. This imposes an existential confluence property on A that facilitates embedding of undecidable tiling problems, [18]. A body of theory exists for CRS, [11], including connections with the Guarded Fragment of first-order logic, arguably making the above compositional semantics ‘useful’ and ‘interesting’—if only, as a way of throwing new light on what makes standard Tarski semantics tick.

Discussion. The preceding is not a defense of CRS, which for us just serves as a case study of ‘loosening up’ the uniqueness of a compositional semantics and the tight fit with one logic to be validated by it. There are many more examples of this phenomenon. But there is still more to this particular case study.

New aspects of compositionality? As a final point of interest, some features of CRS suggest further aspects to compositionality not brought to light so far.

²⁹More abstract modal models generalize these first-order models still further to ‘states’ and appropriate relations between them for making compositional interpretation work.

Consider the following objection to CRS semantics. *Locality* fails, in that truth values of formulas φ may depend on values of variables that do not occur in φ . Locality is a requirement close to compositionality, though one not implied by our analysis so far: the semantics of CRS is perfectly algebraic. However, Locality is reinstated in the system LFD to be discussed in Section 7.2: there, values of the current assignment for the variables in a formula determine its truth value. Failures of Locality occur in other logics, too, such as IF logic or extensions of PAL such as ‘arbitrary announcement logic’. Repairs are sometimes made, but not always, since interesting new phenomena may come to light. For instance, the failure of Locality in IF logic models the natural phenomenon of *signaling* in games. This was initially treated as a problem, but is now seen as a desirable characteristic feature of the IF language, see [54].

7 The other side of the coin: language design

A compositional semantics is usually chosen so as to match a given language. But this language need not be sacrosanct: something may give on both sides.

7.1 Patterns and issues in language change

Changing the original language. Compositional solutions have also changed initial ideas for the syntax of natural language, [46], and the same is true for programming languages. However, this phenomenon seems largely a matter of avoiding getting trapped in ill-formulated problems in the first place.

Improved fit to the chosen models. A more interesting role for language redesign arises even when a compositional semantics is successful by the criteria discussed so far. Once an appealing semantics is there, with well-motivated structures and truth conditions, its very success generates a question. Is the original language the best medium for describing these structures, bringing out their crucial structure? Or can it be extended with notions that are interpreted in the already established compositional style?

CRS is a case in point. With restricted sets of assignments A , there is now room for significant *polyadic quantifiers* $\exists \mathbf{x}\varphi$, which say that there is an available assignment different from the current one at most in its values for the \mathbf{x} such that φ holds. Given the failure of $\exists x\exists y\varphi \rightarrow \exists y\exists x\varphi$, this no longer reduces to iterated single quantifiers as was the case on standard models.

In practice, this now means that CRS has a larger repertoire for formalizing concrete arguments couched in natural or scientific language—and it would be of interest to see how it fares on traditional formalization projects carried out at a time when FOL was still the only game in town.

There are many further examples of the continuing interplay of language design and semantic analysis. For instance, in the semantics of temporal expressions in natural language, flows of time were introduced which then themselves suggested new temporal and modal operators, whether or not realized already

in natural language, [35]. Similar language redesign occurs in temporal logics for computational processes whose expressive power may go beyond the original needs of some specific programming practice, [72]. For another example, logics interpreted over information states tend to ‘split’ classical vocabulary, witness the Boolean and compositional conjunctions of relevant or linear logic, which reflect the new options available in the semantics. A final example is epistemic logic with its new notions of common and distributed knowledge that go beyond the original epistemic language originating in philosophy, [67].

Language redesign is an ongoing process bringing to light legitimate options. In particular, when a compositional semantics has been found introducing a family of new parameters, one may or may not make these parameters explicit in the syntax of the logic. And thus, for instance, in the 1970s, a long debate raged as to whether all first-order properties of temporal flows should be made explicit in tense logics for describing discourse in natural language, [74].

Implicit versus explicit. This language redesign process is a force for new system building, but it may take some time to kick in before the conservativeness inherent in the program ‘giving a semantics’ for a given initial language is overcome. For instance, when one only thinks of finding a ‘topological semantics’ for the basic language of intuitionistic or modal logic, it may take a long time before one sees that formalizing elementary reasoning in mathematical topology quickly needs notions beyond those simple languages, and the same is even more true when topology gets mixed with geometry. And indeed, it took a very long time before such extensions started being studied: cf. [9].

Language design interacts with valid inference. Likewise, one may view the contrast between intuitionistic logic and epistemic logic as being one of how much structure of information models is made explicit in the language. This again illustrates how issues of language design are entangled with the validities one endorses. Intuitionistic logic is weaker than classical logic, but the very absence of classical laws like Excluded Middle encodes significant features of an information- or knowledge-based view of truth: silence is informative. By contrast, epistemic logic has explicit operators for knowledge of agents, but with these in place, the base logic can remain classical, based on the traditional notion of truth. A general study of the ‘implicit’ versus the ‘explicit’ methodology, and the many new questions to which this gives rise, is made in [79].

Natural versus formal language. Of course, it is easier to redesign formal languages than natural languages that have evolved through history in a community of users.³⁰ Even so, one might speculate to which extent natural language is malleable, too, perhaps even by theoretical semantic considerations.

³⁰Programming languages are an intermediate case: they have been designed at some point in history, and can be redesigned in principle, but once they have an established community of users, redesign becomes harder, and ad-hoc patches may be the norm.

7.2 Case study: from CRS to modal dependence logic

Several of the above themes concerning language redesign are clarified by the following case study, which also shows how ideas in existing established semantics can be reassembled in surprising ways, and offers yet another illustration of the entanglement of compositional interpretation with considerations of inference.

Explicit dependence atoms. Language redesign can have many motivations. A concrete case continues with the earlier topic of generalized first-order semantics in Section 6.3. As an analysis of dependence, the logic CRS is ‘implicit’, [79]. It merely reinterprets the existing language of first-order logic, and locates information about (in-)dependence of variables in failures of classical laws. An alternative would be to introduce explicit *vocabulary for dependence* on the same models. For this purpose, one can use atoms D_{xy} for global dependence, interpreted in generalized assignment models \mathcal{M} as saying that, if two assignments $s, t \in A$ agree on the value of x , they also agree on that of y . To find a compositional semantics for this language, [73] adopts Hodges’ ‘team semantics’ for Hintikka’s IF logic (cf. Section 2.2). The result is a second-order dependence logic quantifying over sets of assignments and lifting propositional connectives to this setting, which is undecidable, indeed non-axiomatizable, though some fragments are better-behaved. It is often suggested in the recent literature that this second-order system is the favored outcome of compositional analysis.

A modal approach. However, our earlier point returns. Compositionality can often be achieved in various ways, and the complexity of the resulting logic may guide our choice. The following logic of functional dependence LFD takes *local*, not global, dependence as its basic notion, where $D_x^s y$ holds if any assignment t agreeing with s on the value of x also agrees on that of y . LFD uses the following syntax: $\varphi ::= Px_1 \dots x_k \mid D_X y \mid \neg\varphi \mid \varphi \wedge \psi \mid D_X \varphi$, where X is a finite set of variables. Models are still generalized assignment models $\mathcal{M} = (D, I, A)$, where I maps predicate letters P into predicates of the same arity over D . The two highlights of the truth definition are as follows:

$$\begin{aligned} \mathcal{M}, s &\models D_X y \text{ iff } D_X^s y \\ \mathcal{M}, s &\models D_X \varphi \text{ iff for all } t \in A \text{ with } s =_X t, \mathcal{M}, t \models \varphi \end{aligned}$$

Here $s =_X t$ means that s and t agree on the variables in X . So, evaluation takes place, modal-style, at single assignment inside sets of assignments, where the basic quantifier or modality $D_X \varphi$ says that the current values of the variables in the set X fix the truth of the formula φ . Global dependence and global modalities are definable from local ones using the fact that the special case D_\emptyset expresses the universal modality quantifying over all assignments in A .

As mentioned, *locality* holds in LFD, in the following sense: Define the *free variables* of a formula by the following recursion: (a) $free(Px_1 \dots x_k) = \{x_1, \dots, x_k\}$, (b) $free(\neg\varphi) = free(\varphi)$, (c) $free(\varphi \wedge \psi) = free(\varphi) \cup free(\psi)$, (d) $free(D_X \varphi) = free(D_X \varphi) = X$.³¹ Then:

³¹Thus, not even the notion of a free variable is written in stone.

If $\text{free}(\varphi) \subseteq X$ and $s =_X t$, then $\mathcal{M}, s \models \varphi \Leftrightarrow \mathcal{M}, t \models \varphi$.

Here is a key feature of this analysis of dependence language, [16]:

Fact 17

Validity in LFD is decidable.

LFD also has Hilbert- and Gentzen-style axiomatizations supporting Interpolation and Beth definability theorems. As for the above theme of language redesign, the language supports extensions with identity atoms, modalities for independence, and dynamic modalities for model change. In this manner, LFD can analyze notions of dependence in topology, dynamical systems, and games.³² But whatever the merits (and drawbacks) of LFD, for us here, it just illustrates the interest in rethinking of what look like definitive compositional semantics.

Thus once more, the needs of providing a compositional semantics do not fix one unique logic for an important area of reasoning, in this case, for dependence. Further considerations matter, of the sort identified in the above analysis.

8 Perspectives from games

This paper has provided a broad canvas for understanding compositionality as it functions in logical semantics at the interface of linguistics and philosophy. In particular, we identified some explicit design patterns and criteria for judging success that occur across the board, and that deserve attention.

In this section, we connect up with a major feature of Abramsky’s work: its crucial use of *games*. What follows is a light discussion of some key features of the long-standing game perspective in logic, language and computation, and how it might affect received views of compositionality in philosophy and linguistics. In this setting, we will sketch some distinctive features of Abramsky’s game semantics for programs, which is in the tradition of proof theory and categorical logic rather than model theory. However, a full discussion doing justice to this program with its many ramifications would require a separate paper.

8.1 Games in logic, language, and computation

We start with some basic background about interfaces of logic and games.³³

Games in logic. Games have long been used in logic for a variety of tasks. Perhaps the most basic one of these is *semantic evaluation* of formulas in given models. Originally developed for FOL in the 1960s, semantic games now exist for modal logic, fixed point logics (such as the μ -calculus and the logic $\mu(\text{FOL})$ extending FOL with monotonic fixed points), [37, 50], IF logic and dependence

³²Here concrete notions of dependence add additional valid principles. For instance, dependence in vector spaces satisfies the Steinitz Principle $D_{X \cup \{y\}} z \rightarrow (D_X z \vee D_{X \cup \{z\}} y)$.

³³For a much more comprehensive survey, including contacts with game theory, see [78].

logics, resource sensitive formalisms such as linear logic, but also logics with transfinite conjunctions, or infinitely deep alternations of quantifiers.

Evaluation games usually have two players: Proponent (Verifier, Myself) and Opponent (Falsifier, Nature) with opposite goals (establishing the truth or falsehood of the formula). Also, most games are sequential: only one player moves at any stage, as dictated by the syntax of the formula: the leading propositional connective, modality or quantifier determines the turns. Play is typically finite, ending when atomic formulas are reached (with the winner determined by the truth value of the atom), but games with infinite play also occur, e.g., in languages that contain fixed-points or infinitely deep formulas, in which case a winning rule has to be specified for the infinite plays. Evaluation games are usually games of perfect information, though famously, imperfect information is needed in the games that match the Hintikka-Sandu IF logic or Väänänen’s DL, two extensions of classical first-order logic discussed in earlier sections.

Typically, logical evaluation games have the following basic property:

A formula φ is true in a model M iff the Verifier has a winning strategy in the φ -game played w.r.t. M .

Likewise, falsity matches the existence of a winning strategy for the Falsifier. Thus, *strategies* are a key notion here, which may be viewed as more fine-structured ‘reasons’ for the truth or falsity of the formula.

In addition to games for evaluating formulas, there are other important logical games. In particular, *Ehrenfeucht-Fraïssé games* are widely used in model theory for analyzing the expressive power of logical languages, and unlike evaluation games, they analyze two given models in parallel without being guided by the syntax of one particular formula. Winning strategies for the ‘Duplicator’ in these games correspond with structural analogies between models, such as potential isomorphism up to some depth, while winning strategies for ‘Duplicator’ match up with formulas true in one model and false in the other.

More relevant to us here, in addition to semantic scenarios, there are two-player games for analyzing argumentation. Formal *proof games* were initiated by Paul Lorenzen, [53], in a study of the foundations of logic, and the resulting ‘dialogical paradigm’ has had an influence reaching into philosophy and argumentation theory. These games, played between a Proponent of the conclusion, and an Opponent granting the premises, are not played w.r.t. given models: they probe internal structure and links between premises and conclusions. Again, the key notion is that of a strategy. A winning strategy for the Proponent in argumentation corresponds to a proof for the claim from the premises, [29]. Interestingly, winning strategies for the Opponent in this game match up with constructions of *counter-models* for the claim w.r.t. the premises.

While evaluation games and proof games clearly share some formal features, they analyze different logical notions: truth versus consistency or derivability. This coexistence reflects the ever-intriguing interplay of Model Theory and Proof Theory as two fundamental perspectives or working styles in logic.

Games in natural language. Games can also be understood as taking the dynamic semantics of language use mentioned in Section 3 to a multi-agent setting. After all, natural language is typically a medium for communication between different agents who produce and analyze speech or text in turn.

A well-known paradigm here is the ‘game-theoretic semantics’ of [40], [41], introduced in Section 3, whose main features are as in evaluation games for FOL. Such a game may be viewed as a use of the relevant sentence in a concrete situation (context, model), and while ‘winning’ or ‘losing’ may record truth values, the winning strategies are what determines the total meaning.³⁴

Yet, in the study of natural language, too, other types of game exist. For instance, in pragmatics, there is a long tradition of *signaling games* starting with Lewis’ work on conventions and the emergence of meaning, [51], and continuing with further themes in Skyrms [69, 70], van Rooij [80], Zollman [85], and others.³⁵ But the earlier dialogical argumentation games, too, make sense for natural language, since much of discourse is about maintaining consistency rather than direct checking of truth against some independently available reality. For more on comparing and connecting various game-theoretic approaches to the syntax, semantics and pragmatics of natural language, cf. [77].

As said, the use of games for natural language is related to current dynamic semantics. But in a way, games are ‘more dynamic’: unpacking meaning now involves a *process* of evaluation by which speakers and listeners interact with the given sentence and with other agents in a finite or infinite dialogue.

Games in computation. Games also play an important role in the model-theoretic semantics of computation. For instance, evaluation games for fixed-point logics model nested recursions in sequential computational processes. These games go beyond simple logical evaluation games in allowing infinite play. The widely used *parity games* for the modal μ -calculus produce infinite histories with a winning condition based on which fixed-points in the given formula are unfolded infinitely often.³⁶ There is an extensive theory of these games, with striking results such as the Positional Determinacy Theorem stating that memory-free winning strategies suffice for parity games. The use of games and strategies in this area is deeply connected with Automata Theory, [27].

The possible infinity of play in some model-theoretic evaluation games (or in the earlier-mentioned Ehrenfeucht-Fraïssé games) fits with contemporary views of computing as an open-ended process, which may or may not terminate, may or may not compute any function, but whose main output is *behavior* across time. This same behavioral perspective, now also with an emphasis on concurrency and interleaved action, is central to the proof-theoretic tradition in

³⁴Hintikka makes a suggestive distinction here between ‘literal’ and ‘strategic’ meaning.

³⁵There are major differences between competitive logical games and signaling games for natural language. Signaling games are partly cooperative, with coordination as a shared goal (though players may have further conflicting goals: maximum informativity for listeners, minimum effort for speakers). Thus, signaling games typically settle on Nash equilibria, indicating satisfactory mutual understanding rather than brute winning strategies.

³⁶Formally, when play hits a propositional fixed-point variable, no evaluation takes place against the model, but a return to the body of its matching subformula.

the semantics of computation, culminating in Abramsky’s work. To get there, compositionality challenges had to be met on the way.

Abramsky’s game semantics. Lorenzen’s original system had a mixture of ‘logical rules’ determining which moves of defense and attack players could make and ‘procedural rules’ determining how to schedule attacks over time, whether repetitions of defenses were allowed, and so on. The resulting dialogues could be infinite, at least for first-order logic, in case the initial claim does not follow from the premises. By manipulating procedural rules, the analysis could then be made to fit intuitionistic logic, classical logic, or yet other proof systems. While this mixture of logical and procedural rules may be a realistic model of actual argumentation, it makes it hard to see a compositional pattern.

This challenge is addressed in Abramsky’s work, [6], [1], with several innovations inspired by a computational view of the content of logical systems. Formulas now denote games between a System and its Environment. Here, proposition letters now stand for subgames which can themselves have any complexity. Moreover, to avoid imposing classical Excluded Middle by design, non-determined infinite games are the paradigm. This is not just a technical fix for blocking some law, it also means that play can now keep returning to subgames, allowing for a much richer view of processes modeled by logical formulas.

Next, logical operators are interpreted as *game constructions* for new games out of given ones. For instance, as in most logical evaluation games, negation corresponds to dualizing the roles of players. Importantly, here, a significant ‘split’ occurs leading to a richer repertoire of logical notions. One natural disjunction is choice between subgames right at the start, as in most evaluation games, but an equally natural ‘delayed disjunction’ occurs with *parallel play* of two games where one designated player has the right to *switch* from one game to the other, with a suitable winning condition for the resulting infinite plays. This richer game setting absorbs some procedural rules (as far as needed) into the semantics of the logical operators, and transforms the Lorenzen setting. In particular, there is now a notion of an implication game $A \rightarrow B$ with the subgames A, B played in parallel, where winning strategies for the System can be viewed as exemplifying the implication. The resulting logical inferences validated by these games turned out to match up eventually with *linear logic*, [6].

The mathematical setting for all this are *categories* whose objects are games, and whose morphisms are strategies in implication games. In the exact design of these categories, various further compositionality challenges are encountered, as explained in [1]. One approach uses partial strategies and a simple notion of composition of morphisms, but it does not extend to total strategies in the usual game-theoretic sense. This problem is solved by Abramsky’s restriction to *winning strategies* defined in just the right way on parallel disjunction and implication games, whose details need not concern us here.³⁷

³⁷Finding a good notion of composition of morphisms in a category is not exactly the same as solving a compositionality problem in our general sense, but the two tasks are related. Abramsky explains the issues in terms of validating the Cut Rule in the given proof system, and what requirements this puts on composition of total winning strategies in games for $A \rightarrow B$ and $B \rightarrow C$ to winning strategies for the game $A \rightarrow C$, making sure that episodes in

This is the start of a much broader program. Infinite parallel games, with a pattern of moves or calls to subgames are an excellent model of computational processes where input can be consulted, information can be passed internally, and output produced. Surprisingly, powerful memory-free ‘logical’ strategies in such games like ‘Copy-Cat’ turn out to capture essentials of interaction.³⁸ Enrichments and variations of this game semantics handle a wide variety of features of programming in a compositional manner. And beyond program semantics, game semantics is a theory of communicating processes, whose relation to earlier frameworks such as Process Algebra is discussed in [5].

Abramsky feels that this approach, with the above formulated in category-theoretic terms with games as objects and strategies as morphisms, offers an abstraction level that is more finely grained than many model-theoretic semantics, while avoiding the excessive syntactic detail of some proof systems:

“Game Semantics mediates between [...] operational and denotational semantics, combining the good structural properties of one with the ability to model computational fine structure of the other”.

Aside: Model theoretic and proof theoretic semantics. Much more could be said about the proof-theoretic flavor of Abramsky’s semantics in terms of categories of games, whose roots go back to the Curry-Howard Isomorphism in the lambda calculus and the Brouwer-Heyting-Kolmogorov semantics of intuitionistic logic. Here, we just note that the general approach in this paper covers both proof-theoretic and model-theoretic views, as it includes *algebraic semantics*, a pattern common to both. The coexistence of model theory and proof theory in the study of logical systems and natural language raises many intriguing and often unresolved questions (cf. [76] for discussion) that we cannot address here.

8.2 General themes

Our earlier points about compositionality return with games in logic and computation, though with new twists.

For instance, Abramsky’s compositional game semantics for linear logic, IF logic, and their generalizations, [6, 3, 8] is based on parametrizing the context using parallel composition of strategies. At an abstract level, this demonstrates the technique of Currying of Section 4, with contextuality now internalized by a new parameter: the *environment* to which a player’s strategy is responding (or the information of the player about her environment). Strategy composition works by ‘plugging in’ one player’s environment into the other players’ strategies, thus achieving compositionality in a simple and natural way.³⁹

the subgame B remain hidden, and that play does not get stuck forever in these subgames.

³⁸Accordingly, in Abramsky’s view, the matching tensor product in categories is more fundamental than the sequential composition in other logics of computing. For *truly concurrent* game semantics with simultaneous moves by all players, cf. [2, 3].

³⁹Obtaining a yet more abstract compositionality result for game semantics in the style of Hodges’ result in Section 2 seems an open (if not yet very well-defined) problem.

Another key point that returns in game semantics is the legitimate variety of compositional solutions, discussed in Section 5.1. In the model-theoretic realm, this occurs, for instance, with the modal μ -calculus. Infinite parity games are one natural evaluation game, but so are finite second-order games that involve picking objects, but also sets. This variety of games may seem less obvious in a proof-theoretic setting, when the target of the analysis is one specific logical system, but it can arise even there. For instance, [78] discusses different proof games for classical first-order logic, one based on Lorenzen dialogues, another on semantic tableaux. And in principle, one could even have more model-theoretic and more proof-theoretic games for the same language and logic. Indeed, this variety is only to be expected, since games are a mathematical realm with much more structure than the usual sparser semantic denotations.

Finally, the richer ontology of games also illustrates the theme of language redesign of Section 7. A game semantics often suggests redesign of the language one started with. The splitting of disjunction in Abramsky’s game semantics and the resulting focus on parallel game constructions was a good example. An illustration inside classical logic is the alternative syntactic view of quantifiers as denoting atomic games rather than unary logical operators in [4], [78].

Beyond these methodological concerns or matters of detailed game design, here is what we ourselves see as the major new ideas in current game semantics.

First, *multi-agent* interaction seems crucial to how language functions, and the Agent/Environment methodology fits very well with the realities of how we use language to learn about and cope with the world. Also, *concurrency* has not yet been prominent in the semantics of natural language, even though much of how language functions and evolves on a large scale is not in face-to-face sequential interactions, but in concurrent events of language use in large communities. But also, the shift from finite to *infinite* games, and from final output to ongoing *behavior* seems to fit natural language use very well. Linguistic activity as a whole can be viewed as a non-terminating ‘operating system’ that provides the arena where terminating tasks can be performed. However, to us, most food for thought is provided by the following more radical perspective.

Inverting the direction: from construction to analysis. In unfolding a game, game semantics takes the view of an agent who ‘observes’ a sentence (in its context), and gradually analyzes its meaning, by taking it apart, rather than constructing it recursively from the bottom. This inverts the standard semantic perspective that is often associated with compositionality: interpreting expressions outside–in, as Hintikka recommended, rather than inside–out.

All this makes eminent sense for natural languages, especially if one also takes on board pragmatics and understanding up to a point that suffices for action, while acknowledging the role of questions or elucidations to increase depth of understanding as needed. And this is not just a vague metaphor. The outside–in perspective relates to the *co-algebraic view* on semantics, [56], [64], [25]. Coalgebra provides a general mathematical model of dynamical systems that proceeds by co-recursion, rather than recursion. Syntactic operators decompose or unfold an initial state into its components or successors. By grad-

ually parsing a sentence in this top-down manner following its syntactic tree, we gradually ‘observe’ or analyze a semantic model, or even the structure of reality itself – and this process could well be infinite.

Stated a bit provocatively, on this view, sentences do not provide complete meanings or thoughts: they are a tip of a mental iceberg. And in line with this, reality is not something that we construct out of smallest units, but a complex something that we can only observe to a certain extent by interacting with it. And thus the quest for compositionality achieves a certain grandeur: it might be seen as a way of keeping this learning process comprehensible.

9 Further aspects of compositionality

This paper may have made some unusual and occasionally controversial comments, but its agenda and modus operandi fit squarely in the established philosophical and mathematical tradition of studying compositionality by logicians. As a counterpoint, it may be useful to observe that our community is surrounded by a broader world today where compositionality is either conceived very differently, or even as something to be abandoned altogether without giving up on mathematical precision, learnability, or computational efficiency.

Cognitive perspectives. Recent studies of compositionality in cognitive science concern the historical emergence of recursive syntax with matching compositional meanings that support a bootstrapping analysis of how meaningful language arose over time. A good example is found in [71]. This work shows how compositionality can emerge in principle in evolutionary scenarios for language development, and it is backed up by computational simulations of linguistic scenarios for communication.

Distributional semantics. A major current challenge to compositional semantics is the emergence of distributional semantics, where word meanings are computed from co-occurrence frequencies in large corpora, represented in high-dimensional vector spaces, [28]. A guiding motto for distributional semantics is Zellig Harris’s oft-quoted dictum, cited in Section 1, that “If you want to know the meaning of a word, look at the company it keeps”. However, as we have seen in Hodges’ analysis in Section 2, Contextuality can co-exist with Compositionality, so it may be too early to say whether compositional semantics is truly at odds with its distributional rival.⁴⁰

In this connection it may be worth noting that a computational linguistic framework like Data-Oriented Parsing, [21], combines more compositional views of syntax and semantics with a probabilistic memory of connections in a large corpus of earlier experiences.

Even so, it may be a good idea to confront current discussions of compositionality in language and meaning with some of these new realities.

⁴⁰[22] makes a similar point about machine learning systems for natural language processing.

10 Conclusion

Compositionality is a perennial methodological issue that has spawned a huge literature, ranging from particular empirical phenomena to the design of semantics for whole languages, with perhaps the most general abstract perspective to be found in mathematics. One of the reasons why it continues to generate debate is its intriguing relationship with what looks like its antipode: contextuality.

In this paper, we have tried to add some new, or at least under-appreciated considerations to received wisdom. At the most abstract level, we identified a ubiquitous ‘currying pattern’ in setting up a semantics that brings together contextuality and compositionality. Next, looking at more specific case studies, we identified a list of basic concerns that help in understanding what a proposed compositional semantics does, and does not, achieve. In particular, we emphasized the entanglement of compositional semantics with the target set of desired valid inferences, and the dynamic role of language redesign once a semantics has been proposed. We believe that all this results in a more relaxed, but also, a richer attitude toward compositional semantics as an ongoing enterprise, generating new designs, as well as new logical questions and results.

Finally, we have looked at the role of games and game semantics for logic and computation as an area where further intuitions come into play, including radical departures from standard constructive views of meaning to outside-in analyses of complex potentially never-ending behavior. Here computer scientists may well have something to teach to linguists and philosophers.

Acknowledgments Over the years, we have been inspired by our colleague Theo Janssen whose pioneering work on compositionality is still well-worth reading after all these years. The continuing influence of Wilfrid Hodges will also be evident throughout. More concretely for now, we thank the editors of this volume, and the two referees for this chapter. In addition we profited from feedback by Peter Pagin and Jouko Väänänen. But most of all, we thank Samson Abramsky for providing the inspiration and impetus to write this paper.

References

- [1] S. Abramsky. Semantics of interaction. In *Proceedings of the 1996 CLiCS Summer School, Isaac Newton Institute, Cambridge*, pages 1 – 31, Cambridge: UK, 1997. Cambridge University Press.
- [2] S. Abramsky. Sequentiality vs. concurrency in games and logic. *Mathematical Structures in Computer Science*, 13:531–565, 2003.
- [3] S. Abramsky. Socially responsive, environmentally friendly logic. In T. Aho and A.V. Pietarinen, editors, *Truth and Games: Essays in Honour of Gabriel Sandu*, volume 78 of *Acta Philosophica Fennica*, pages 17–45. Societas Philosophica Fennica, Helsinki, 2006.
- [4] S. Abramsky. Information, processes, and games. In P. Adriaans and J. van Benthem, editors, *Handbook of the Philosophy of Information*, pages 483–549. Elsevier Science, Amsterdam, 2008.

- [5] S. Abramsky. From csp to game semantics. In A. W. Roscoe C. Jones and K. Wood, editors, *Reflections on the Work of C.A.R. Hoare*, pages 33–45. Springer Science Publishers, 2010.
- [6] S. Abramsky and J. Jagadeesan. Games and full completeness for multiplicative linear logic. *Journal of Symbolic Logic*, 59(2):543–574, 1994.
- [7] S. Abramsky and G. McCusker. Full abstraction for idealized Algol with passive expressions. *Theoretical Computer Science*, 227:3–42, 1999.
- [8] S. Abramsky and J. Väänänen. From IF to BI: a tale of dependence and separation. *Synthese*, 167:207–230, 2009.
- [9] M. Aiello, I. Pratt, and J. van Benthem, editors. *Handbook of Spatial Logic*. Springer, 2007.
- [10] H. Andréka, I. Németi, and I. Sain. Algebraic logic. In D. Gabbay and F. Guentner, editors, *Handbook of Philosophical Logic, Vol. 2, second edition*. Dordrecht: Springer Science Publishers, 2001.
- [11] H. Andréka, I. Németi, and J. van Benthem. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.
- [12] H. Andréka, J. van Benthem, N. Bezhanishvili, and I. Németi. Changing a semantics: opportunism or courage? In M. Manzano, I. Sain, and E. Alonso, editors, *The Life and Work of Leon Henkin: Essays on his Contributions*, pages 307–338. Birkhäuser, Springer, 2014.
- [13] J. L. Austin. *How to Do Things with Words*. Harvard University Press, 1962.
- [14] R. Ballarín. Modern origins of modal logic. In E. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Stanford, 2017.
- [15] A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In *Proceedings of the Seventh TARK conference (1998)*, pages 43–56. Morgan Kaufmann Publishers, San Francisco, 1998.
- [16] A. Baltag and J. van Benthem. A minimal logic of functional dependence. Technical report, Institute for Logic, Language and Computation, University of Amsterdam, 2020.
- [17] J. Barwise and J. Perry. *Situations and Attitudes*. MIT Press. Reprinted by CSLI Publications, Stanford, 1999, 1983.
- [18] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, Cambridge UK, 2001.
- [19] S. Blamey. Partial logic. In D. Gabbay and F. Guentner, editors, *Handbook of Philosophical Logic, 2nd edition, volume 5*, pages 261–353. Dordrecht: Springer Science Publishers, 2001.
- [20] W. Blok and D. Pigozzi. *Algebraizable Logics*, volume 396. American Mathematical Society. Also as a Classic Reprint, Advanced Reasoning Forum, 2014, Providence, 1989.

- [21] R. Bod. The data-oriented parsing approach: Theory and application. In *Studies in Computational Intelligence*, pages 307–348. Heidelberg: Springer Verlag, 2008.
- [22] D. Bonnay. Compositionality in the light of deep learning. Lecture Presentation, Department of Philosophy, Université Paris Nanterre, 2019.
- [23] P. J. Cameron and W. Hodges. Some combinatorics of imperfect information. *Journal of Symbolic Logic*, 66:673–684, 2001.
- [24] W. Carnielli and I. d’Ottaviano. Translations between logical systems: A manifesto. *Logique et Analyse*, 15:67–82, 1997.
- [25] C. Cirstea, A. Kurz, D. Pattinson, L. Schröder, and Y. Venema. Modal logics are coalgebraic. *The Computer Journal*, 54:31–41, 2011.
- [26] M. Dummett. *Frege. Philosophy of Language*. Duckworth, London, 1973.
- [27] W. Thomas E. Graedel and T. Wilke, editors. *Automata, Logics, and Infinite Games*, volume 2500 of *Lecture Notes in Computer Science*. Springer Verlag, Berlin, 2002.
- [28] S. Evert. Distributional semantics: A practical introduction. In *ESSLI Lecture Notes*. ESSLI, Bolzano, 2016.
- [29] W. Felscher. Dialogues as a foundation for intuitionistic logic. In *Handbook of Philosophical Logic, Vol III*, pages 341–372. Kluwer Academic Publishers, Dordrecht, 2001.
- [30] G. Frege. *Grundlagen der Arithmetik*, translated by J. Austin as *Foundations of Arithmetic*, 1950. Blackwell, Oxford, 1884.
- [31] P. Galliani. Sensible semantics of imperfect information. In M. Banerjee and A. Seth, editors, *Logic and Its Applications, vol. 6521 of LNCS*, pages 79–89. Springer, Berlin/Heidelberg, 2011.
- [32] K. Glüer and P. Pagin. *Switcher Semantics. From Propositions to Ingredient Sense*. Under contract with Oxford University Press, 2021.
- [33] K. Gödel. Zum intuitionistischen Aussagenkalkül. In *Anzeiger der Akademie von Wissenschaften in Wien*, volume 69, pages 65–66. Akademie von Wissenschaften in Wien. Repr. and transl. in Feferman et al. (eds.), *Kurt Gödel. Collected Works, Vol. I*, OUP, 1986, 222–225., 1932.
- [34] R. Goldblatt. Mathematical modal logic. a view of its evolution. In D. Gabbay and J. Woods, editors, *Logic and the Modalities in the Twentieth Century*, volume 7 of *Handbook of Modal Logic*, pages 1–98. Elsevier, 2006.
- [35] V. Goranko. Temporal logic. In *Stanford Electronic Encyclopedia of Philosophy*. CSLI, Stanford University, 1999.
- [36] J. Groenendijk and M. Stokhof. Dynamic predicate logic. *Linguistics & Philosophy*, 14:39–100, 1991.
- [37] Y. Gurevich and S. Shelah. Fixed-point extensions of first-order logic. *Annals of Pure and Applied Logic*, 32:265–280, 1986.

- [38] D. Harel. Dynamic logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic, Vol II*. D. Reidel, Dordrecht, 1984.
- [39] D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. The MIT Press, Cambridge MA, 2000.
- [40] J. Hintikka. *Logic, Language Games and Information*. Clarendon Press, Oxford, 1973.
- [41] J. Hintikka and G. Sandu. Game-theoretical semantics. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language, 2nd edition*, pages 415–466. Elsevier, Amsterdam, 2011.
- [42] W. Hinzen, E. Machery, and M. Werning, editors. *The Oxford Handbook of Compositionality*. Oxford University Press, Oxford, 2012.
- [43] W. Hodges. Compositional semantics for a language of imperfect information. *Logic Journal of the IGPL* 5, pages 539–563, 1997.
- [44] W. Hodges. Formal features of compositionality. *Journal of Logic, Language and Information*, 10:7–28, 2001.
- [45] W. Hodges. From sentence meanings to full semantics. In J. van Benthem et al., editor, *Proof, Computation and Agency*, pages 261–276. Springer, 2011.
- [46] T. Janssen. Frege, contextuality and compositionality. *Journal of Logic, Language and Information*, 10:115–136, 2001.
- [47] T. Janssen. Compositionality. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 495–554. Elsevier, Amsterdam, 2nd edition, 2011.
- [48] T. Janssen and P. van Emde Boas. Some remarks on compositional semantics. In D. Kozen, editor, *Logic of Programs*, number 131 in Lecture Notes in Computer Science, pages 137–149. Springer, Berlin, 1981.
- [49] H. Kamp. A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof, editors, *Formal Methods in the Study of Language*, pages 277–322. Mathematisch Centrum, Amsterdam, 1981.
- [50] P.G. Kolaitis and M.Y. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In A. Scedrov, editor, *Proceedings of the Seventh Annual IEEE Symp. on Logic in Computer Science, LICS*, pages 46–57. IEEE Computer Society Press, June 1992.
- [51] D. Lewis. *Convention. A Philosophical Study*. Cambridge: Harvard University Press, 1969.
- [52] D. Lewis. Index, context, and content. In Stig Kanger and Sven Öhman, editors, *Philosophy and Grammar*. D. Reidel, Dordrecht, 1980.
- [53] P. Lorenzen. *Einführung in die Operative Logik und Mathematik*. Springer Verlag, Berlin, 1955.

- [54] A. Mann, G. Sandu, and M. Sevenster. *Independence-Friendly Logic. A Game-Theoretic Approach*. Cambridge University Press, Cambridge, 2011.
- [55] R. Montague. Universal grammar. *Theoria*, 36:373–398, Reprinted in [?], 222–246, 1970.
- [56] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317 (Erratum published *Ann. Pure Appl. Logic* 99:241–259, 1999), 1999.
- [57] I. Németi. The equational theory of cylindric relativized set algebras is decidable. Preprint 63/85, Mathematical Institute, Hungarian Academy of Sciences, Budapest, 1985.
- [58] R. Nouwen, A. Brasoveanu, J. van Eijck, and A. Visser. Dynamic semantics. *Stanford Encyclopedia of Philosophy*, 2015.
- [59] P. Pagin. Communication and the complexity of semantics. In [42], pages 510–529. Oxford University Press, 2012.
- [60] P. Pagin and D. Westerståhl. Compositionality I. Definitions and variants. *Philosophy Compass*, 5(3):250–264, 2010.
- [61] P. Pagin and D. Westerståhl. Compositionality II. Arguments and problems. *Philosophy Compass*, 5(3):265–282, 2010.
- [62] P. Pagin and D. Westerståhl. Pure quotation and general compositionality. *Linguistics & Philosophy*, 33(5):381–415, 2010.
- [63] Peter Pagin. Compositionality and context. In Gerhard Preyer and Georg Peter, editors, *Contextualism in Philosophy*, pages 303–48. Oxford University Press, Oxford, 2005.
- [64] D. Pattinson. Coalgebraic modal logic: soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309:177–193, 2003.
- [65] S. Peters and D. Westerståhl. *Quantifiers in Language and Logic*. Oxford University Press, Oxford, 2006.
- [66] J. Plaza. Logics of public communications. In *Proceedings 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216. 1989.
- [67] R. Rendsvig. Epistemic logic. In *Stanford Electronic Encyclopedia of Philosophy*. 2016.
- [68] B. Russell. On denoting. *Mind*, 14:479–493, 2005.
- [69] B. Skyrms. *Evolution of the Social Contract*. Cambridge: Cambridge University Press, 1996.
- [70] B. Skyrms. *Signals Evolution, Learning & Information*. New York: Oxford University Press, 2010.
- [71] S. Steinert-Threlkeld. Towards the emergence of non-trivial compositionality. *Philosophy of Science* (forthcoming), 2020.

- [72] C. Stirling. Modal and temporal logics. In *Handbook of Theoretical Computer Science, vol. 2, Semantics*, pages 477–563. Amsterdam: North-Holland, 1993.
- [73] J. Väänänen. *Dependence Logic: A New Approach to Independence Friendly Logic*. Cambridge University Press, Cambridge, 2007.
- [74] J. van Benthem. Tense logic and standard logic. *Logique et Analyse*, 80:395–437, 1977.
- [75] J. van Benthem. The logic of semantics. In *Essays in Logical Semantics*, number 29 in *Studies in Linguistics and Philosophy*, pages 198–214. Springer, Dordrecht, 1986.
- [76] J. van Benthem. Proofs, labels and dynamics in natural language. In U. Reyle and H-J Ohlbach, editors, “*We Will Show Them*”, *Festschrift for Dov Gabbay*, pages 31–41. Kluwer Academic Publishers, Dordrecht, 1998.
- [77] J. van Benthem. Games that make sense: Logic, language and multi-agent interaction. In K. Apt and R. van Rooij, editors, *New Perspectives on Games and Interaction*, pages 197–209. Amsterdam: Amsterdam University Press, 2008.
- [78] J. van Benthem. *Logic in Games*. The MIT Press, Cambridge MA, 2014.
- [79] J. van Benthem. Implicit and explicit stances in logic. *Journal of Philosophical Logic*, 48:541–601, 2019.
- [80] R. van Rooij. Evolution of conventional meaning and conversational principles. *Synthese*, 139:331–366, 2004.
- [81] D. Westerståhl. On mathematical proofs of the vacuity of compositionality. *Linguistics & Philosophy*, 21:635–643, 1998.
- [82] D. Westerståhl. On the compositional extension problem. *Journal of Philosophical Logic*, 33:549–582, 2004.
- [83] D. Westerståhl. Compositionality in Kaplan style semantics. In [42], pages 192–219. Oxford University Press, 2012.
- [84] W. Zadrozny. From compositional to systematic semantics. *Linguistics & Philosophy*, 17:329–342, 1994.
- [85] K. Zollman. Talking to neighbors: The evolution of regional meaning. *Philosophy of Science*, 72(1):69–85, 2005.