

Chris Fox

University of Essex

foxcj@essex.ac.uk

The meaning of formal semantics

Abstract: This paper considers some of the foundational issues of formal semantics. These issues include the nature of the data that is to be accounted for, the nature of the formalism used to characterise that data, and the criteria for determining whether a particular formal account is adequate. If the relevant data is merely the arguments and truth conditions that subjects agree with, then the terms in which the theory is expressed may be irrelevant. But a case can be made that it is also legitimate to be concerned with the intuitions that people have about how they reason with language, as well as technical and philosophical issues relating to the chosen formalism. The diversity of issues and objectives is highlighted by considering a selection of different problem domains, and by questioning the sometimes implicit assumptions about what counts as an appropriate formal framework. Ultimately the meaning of formal semantics depends upon assumptions about the nature of the data that is to be accounted for, and in what sense the formal theory itself needs to be sympathetic to that data.

Key words: Formal semantics, formalisation, set theory, model-theoretic analysis, proof-theoretic analysis, ontology

0. Introduction

Those of us engaged in formal semantics are sometimes asked to justify what it is that we do. In particular we may be questioned as to whether what we are engaged in constitutes some form of testable, rigorous scientific endeavour, or whether it is more about telling plausible sounding stories about aspects of language for which we have stories to tell. If the latter, then semanticists may be vulnerable to accusations of ignoring confounding data and issues, or failing to take seriously the need for some objective criteria by which the proposed account can be evaluated, and determined to be better or worse than some alternative account. In this context it is worth considering, at an abstract level, the nature of theories that are proposed by formal semanticists, and the kinds of problem domains in which they might be applied.

This paper develops some of the arguments presented by Fox and Turner (2012).

A typical semantic theory will seek to characterise some certain aspects of meaning for a fragment of the language. Perhaps the most common contemporary approach is to find a way of translating linguistic examples into a set-theoretic representation, for which the rules of set theory itself then mimic the behaviour of the linguistic examples in some sense.¹ We might describe this as the model-theoretic tradition, where any semantic feature of the language is to be characterised as a set. In this setting, even a relation between sets is just another set. The translation might be mediated by a logical representation, but usually the set-theoretic model of this logic is taken to be the “real” semantics.

We may question to what extent such characterisations constitute a semantic theory: what do they tell us that is new or informative? Of course, if we take an extensional view — a view which is consistent with the set-theoretic approach — we could say that such theories themselves characterise what it is to be a semantic theory. But rather than accept this impredicative characterisation, we could break the circle and ask some more reflective questions about this approach. For example, is it right to conflate seemingly distinct notions by collapsing everything to sets? Furthermore, if this is an empirical account, the question remains as to the nature of the data that is being captured. In particular, the question arises as to whether essential aspects of the data, or our intuitive understanding of language, are being lost or overlooked by such characterisations.

Such questions can also be asked of other approaches and paradigms besides the model-theoretic tradition, where for example various notion are captured in terms of expressions in λ -calculus or some form of type theory. In such cases, is it appropriate for distinct notions to be conflated when formulated as λ -terms or as types?

1. Problem Domains

While it may be appropriate to consider formal semantic interpretation independent of any particular problem domain, it can be useful to consider what role a formal semantic analysis might play in various contexts: different issues can come to the fore in different domains. Here we considered

¹A similar methodology can be adopted with frameworks other than set theory, such as constructive type theory (see Ranta, 1994, for example).

text-based computer games, statue law, as well as plurals and mass terms.² The first of these exemplifies a narrow, controlled domain, the second involves the use of rigorous language, and the last exemplifies a general linguistic phenomena that is not constrained to a particular field of use.

One early area for computer-based natural language “understanding” was text-based “adventure games”.³ Such games provide a very clearly defined domain, with a finite number of objects, places and actions, but with a potentially unlimited number of *commands* and *questions* that can be issued by the player. The problem is then how such commands and questions can be interpreted mechanically. In contrast to other problem domains, this provides a clear criteria for correctness, and relatively “clean” data.⁴

In the case of the formal interpretation of statute law, some relevant issues are whether it is possible to determine, mechanically, when the law is being complied with, just by interpreting its language. Intuitively at least, one would expect such laws to have a moderately clear criteria for correctness of interpretation. In contrast to the adventure game problem, the language used is complex, and the data to which the law is applied is quite “messy” and unclear. At the very least, to determine whether a law is being complied with, the relevant circumstances have to be described at an appropriate level of abstraction.

More generally, those working in formal semantics are interested in capturing certain universal aspects of language and its meaning. An example of this is the interpretation of sentences with mass terms and plurals. In such cases, the data is often complex and uncertain. It is not always clear what entailments hold. We give three somewhat arbitrary examples to illustrate the range of fundamental conceptual problems that can arise even with what appears to be a relatively elementary linguistic notion.

²The areas described here illustrate a range of issues, but they are not intended to be paradigm cases, or in any sense exhaustive.

³An early example of this genre is the “Colossal Cave Adventure” (Montfort, 2003, p10).

⁴The view that the domain is limited and straightforward might in part be based on an understanding of how such games actually behave, which may limit expectations about what is appropriate and reasonable given the constraints of the time. Players of contemporary computer games might now expect more sophisticated language comprehension. Certainly, the language processing abilities of early text-based games were rather primitive.

In the case of the sentence “*Three men lifted two pianos*”, we may wonder how many people, pianos and lifting events were involved. For the sentence “*The water is dirty*”, if we take “*the water*” to be defined as referring to some salient collection of water molecules, and nothing else, we may wonder in what sense such a collection of molecules can be considered “dirty”. The statement “*The cards are on top of each other*” can be used correctly to describe a stack of cards, even though the bottom card is not on top of any other card (and there is no card on the top card). There are many other problematic cases, but these are perhaps sufficient to indicate that there are difficulties in giving highly reductive accounts of meaning.⁵

The appropriate ontological structures for human reasoning can be vague. For example, while we may know that material objects composed of atoms, the meaning of language does not depend on such an understanding of the nature of the physical world. In that sense at least, meaning and the nature of our reasoning appear to be independent of scientific theories of the world.

Accounts of plural and mass entities also highlight the question of whether first-order theories are adequate. One structure proposed for analysing such terms is a complete lattice (Link, 1982). Yet the axiom for lattice-theoretic completeness is second-order in nature (requiring quantification over collections of entities).

Such examples highlight the problem that the criteria for correctness are uncertain and the data messy, even for apparently simple, single-sentence examples. For this reason we may be forced to limit the scope of the formal semantic analysis, and leave some aspects of meaning unanalysed.

2. Issues

Some salient issues for a formal account of meaning include the nature of meaning, the nature of the data, the kinds of expressions (“speech acts”) involved, and whether we are aiming for an implementation, a formal characterisation, or an explanation.

We need to reflect on the kind of meaning we are seeking to formalise, and whether it depends on the specific problem, or whether there is a general notion of linguistic meaning common to such problems. If there is a

⁵These examples are discussed in more detail by Fox (2000).

coherent notion of linguistic meaning, we may also wonder to what extent it can be characterised by a formal theory.

The data is more than just the language. For the adventure game, knowledge of the game state is required. For statute law, knowledge of salient parts of a process or procedure are required, together with knowledge of how to formalise descriptive information at an appropriate level of abstraction. In the case of linguistic phenomena such as plurals and mass terms, we seem to require some knowledge of folk ontology, and a way of dealing with vague and incomplete meanings. In all cases, the data to be characterised is some kind of abstraction, of both language and the world.

It is clear that there are a range of linguistic expressions other than propositional assertions. For example, the adventure game requires an analysis of imperatives and questions, and statute law, obligations and permissions. It is possible that ultimately we may decide it appropriate to have some propositional interpretation of these apparently distinct notions.⁶ Even so, ontologically we still need the means to reflect on whether this is an appropriate thing to do.

Semantic theories may have different intended uses. This may effect what counts as an appropriate form of analysis. For the adventure game, it is important to have an effective interpretation that supports an implementation. For statute law, we require an effective description that characterising the key aspects of the law. For linguistic puzzles, we might be more interested in something that could be characterised as an “explanation”. We may question whether these are really distinct (or even appropriate) characterisations, and to what extent they may share some common elements.

The general objectives of the *formal* aspect of formal semantics can be construed as having at least two related aspects. First, there is the *evaluation* of language, and linguistic expressions, in logic. This could be achieved by using using a formal theory to give truth conditions for propositions, answerhood conditions for questions, and compliance conditions for imperatives, for example. Second, formal semantics can capture structural patterns of entailment, seeking to mirror our intuitions about linguis-

⁶For example, Tichý (1978); Hamblin (1973); Karttunen (1977), and others, argue that questions and propositions should be the same type and that any distinction resides in our relationship to them. Similarly in the case if imperatives, some propose reducing them to properties (Portner, 2005, for example).

tic reasoning, presuppositions, connotations, and formalising general rules about when, and in what sense, one proposition, imperative, etc., follows from another.

3. Formal Semantics

Contemporary formal semantic theories assume that a rigorous formal analysis of at least some aspects of meaning is possible. Typically such theories present, or at least assume the existence of a *compositional* translation.

Compositionality guarantees completeness of the coverage of the semantic interpretation for sentences that are generated by the formal syntax, or grammar. It is not necessarily a constraint on the nature of the interpretation itself (Zadrozny, 1994) other than it not being completely arbitrary.⁷ Perhaps the most rigorous version of such an analysis is the “method of fragments” (Montague, 1970a,b, 1973; Partee, 2001) in which the full details of the grammar and semantic analysis are spelt out for a fragment of the language. In more recent work on formal semantics, such an approach is often presupposed but with the details left unstated. There is still an assumption that, in principle at least, the relevant part of language can be given a rigorous grammatical characterisation, and that we can give appropriate compositional rules of interpretation.

In order to make progress on a formal analysis, it has to be assumed that there are some stable, core intuitions about meaning shared by a community of language users. The common canonical core might not include all aspects of meaning (cf. Section 1). There is also the issue of performance *vs.* competence (de Saussure, 1916; Chomsky, 1965) which pervades the analysis of language; our intuitions about what language “ought” to mean may differ from actual practice. There is then a question about whether our normative ideas about language are a genuinely reflection of the intrinsic nature of language itself — assuming that notion is coherent — and whether, and in what sense, the formalisation of such norms corresponds to the formalisation of linguistic meaning.

Concerning the way a theory is expressed, at a minimum the framework in which meaning is to be formalised consists of a system of symbolic notation and deterministic rules whose use and interpretation is understood

⁷See Westerstahl (1998) for a critical appraisal of these arguments.

by an appropriate community. To constitute a logic, as such, the formal system ought to at least support notions of entailment and contradiction. One key question concerns what features a logic must have in order to adequately and faithfully capture appropriate aspects of meaning, and also to express a compositional analysis without resorting to additional, extra-logical machinery.⁸

Typically a logic will be characterised in terms of its *proof theory* (the patterns of entailment that it supports), and a *model theory* (an interpretation of the symbols within some other formal framework). The behaviours of the proof theory and model theory should be consistent with each other (at a minimum, the model theory should not contradict the proof theory). Technically, the role of a model theory is to show that there is a consistent interpretation of the proof theory (where the logic is described as being *sound* with respect to that interpretation), and perhaps that the logic completely characterises the model (where the logic is described as being *complete* with respect to the model).⁹

We may wonder whether there is a “natural” logic. Informally at least we may debate the criteria for what is *appropriate*, or *adequate* for capturing human reasoning. The question of whether there is a natural logic may then be refined into the question of whether there are theory-independent criteria for judging the correctness of some a logic system (or its interpretation). Evidence from axiomatic geometry suggests that such questions are non-trivial, and fall outside the realm of logic itself.¹⁰

Given the difficulty of formulating objective criteria to identify a “natural” logic, it is conventional to fall-back on some standard systems. These

⁸Including the mechanisms for compositionality in the logical framework helps ensure that all relevant aspects of the compositional analysis have a coherent interpretation in the logical framework. This usually requires the logic to support some mechanism for performing substitution, such as the λ -calculus.

⁹There is some potentially confusing use of terminology, where a *model* is sometimes described as providing the *semantics* of the logic, a terminology that perhaps presupposes the model-theoretic view of formal semantics (Section 4).

¹⁰The argument can be summarised as follows. We can formalise rules for geometry. The rules can be considered to be either descriptive or normative for the intended system. They satisfy certain notions of consistency, but there is nothing *within the formal theory* that allows us to say that some formalisation is objectively “right”, and another “wrong”, as we can demonstrate that two different, otherwise incompatible systems of axiomatic geometry are mutually consistent: if one theory is consistent, then so is the other. In the case of axiomatic geometry this follows from the independence of the Euclidean parallel postulate (Beltrami, 1868).

include *propositional logic* (a logic that can be given a straight-forward algebraic definition), *first-order logic* (a logic that is singled out by Lindström’s theorem, and which supports both compactness and the downward Löwenheim–Skolem theorem), *second-order logics* (the weakest logics that are categorial with respect to their models), and *Higher Order Logic* (a logic that can be defined in Church-typed λ -calculus). As can be seen, these “standard” logics can be characterised as formal systems that have, in some sense, mathematically “interesting” properties. But this does not necessarily mean they are “natural”, adequate or appropriate for capturing how we reason with natural language.

Those who attribute a special foundational status to possible worlds might take them to provide a criteria for characterising a “natural” logic; one that is sound and complete with respect to a set-theoretic possible worlds interpretation. On such a view, the logic may be considered to be an eliminable layer of representation. Theories can then be expressed directly in terms of a model theory. This can be seen to characterise the “model theoretic” tradition.

4. Model Theoretic Semantics: a conventional approach

In much of the literature relating to formal semantics, there is a presumption that set-theoretic models are the proper target of formal interpretation. The received story in the Montagovian tradition is that the process of giving a formal interpretation involves finding mappings from natural language to logic, and from logic to a model. Furthermore, this should be done in such a way that composing the two mappings gives a direct translation into the model, allowing the logic to be eliminated (see Thomason, 1974, for example).¹¹

The models typically involve possible worlds (or some alternative, such as context change potential) so that the translation is then from natural language into possible worlds (or context-change potential). In turn, these theories are usually expressed in terms of set theory, allowing us to characterise model-theoretic semantics as being, broadly, the translation of

¹¹There is a point of controversy here about whether this was Montague’s own view. The received view corresponds to the characterisation of Montague semantics presented by Thomason (1974). But this is not uncontested. Montague (1969) hints at the idea that the logical *is* important. This view is also considered by Partee (2013) and Cocchiarella (1981, 1997) — P.C. Partee, July 2013.

natural language into set theory. We might call this *the standard model-theoretic account*.

The model-theoretic approach can be exemplified with the analysis of propositions, questions and imperatives. Essentially, worlds can be taken to be sets of propositions, *or* (equivalently) propositions can be taken to be sets of worlds. Properties are then sets of world–individual pairs, relations are sets of pairs mapped to sets, and quantifiers (at least the simpler cases) are set inclusion and set intersection relationships. Ultimately everything is then mapped to sets or structures over sets, which can all be reduced to a “non-representational”, pure set theory.

Questions can be treated as partitions of worlds (Groenendijk and Stokhof, 1997) corresponding to different answers. Answers then correspond to an identification of a subset of the partition, with appropriate adjustments to account for propositional *vs.* non-propositional answers.¹²

For imperatives, broadly speaking the model-theoretic possible worlds analysis is given in terms of preference relationships over possible worlds (see Lascarides and Asher, 2004; Segerberg, 1990, for example). A similar approach is adopted for possible worlds interpretations of deontic expressions (for example, von Wright, 1951, 1953).

The conventional approach is perhaps motivated in part by the view that set theory has a foundational status, and that any logical representation itself can only have meaning if it too is interpreted in set theory. On such model-theoretic accounts, set theory, and possible worlds, are taken to be independently “given”, perhaps with a more natural or foundational status than a logic. The justification might be that this avoids a reliance on some form of representationalism, for both logic and model theory (as criticised by Wittgenstein and others).

Arguments could be made that semantic accounts that rely on pre-existing, independently motivated frameworks, such as set-theoretic possible worlds, also have some explanatory power. There is a sense in which the theory can be said to “predict” the relevant semantic behaviour in that the behaviour is already contained within the theory: the role of our compositional analysis is then to show how language can be mapped systematically to sets that already have the relevant behaviour.

¹²One alternative is to interpret questions and answers in terms of abstractions and structured propositions. Depending on how this is formalised, related ontological issues can arise, as sketched later in Section 5 (see also Fox, 2013).

In some cases, the translation into set-theory may be somewhat disguised by a narrative that interprets the sets in terms of possible worlds or context-change potential. Semanticists will also talk of “questions”, “partial states”, “answers“, “propositions”, “worlds” and “states” rather than just “sets”. But typically there is an implicit assumption of the paradigm that in their formalisation, such notions can and should ultimately be reduced to pure set theory — even in cases where a logic representation is used.

One downside of this set-theoretic reduction is that the intended interpretation, and the distinctions between different conceptual notions, are not necessarily evident in the formal theory itself, but instead may be provided by a narrative, a narrative that is not strictly part of the formalisation.

5. Is the standard approach adequate?

The arguments that might be made for the standard model-theoretic view can be summarised as follows.

1. It takes a pre-existing “foundational” framework as a starting point.
2. There is no need to “justify” the basic rules of a logic.
3. It avoids representationalism.
4. There is a potential to explain, not just describe

We may question how much weight and credence should be give to such arguments, and whether the standard, set-theoretic possible worlds account is fundamentally different to any other theory or framework. In brief, is it right to credit the standard account as being a non-representational theory with explanatory power?

The standard approach implicitly presumes that all relevant notions can be reduced to set-theoretic characterisations. We may wonder whether there are problems in such reductions (Fox and Turner, 2012). Here it is appropriate to reprise Benacerraf’s dilemma (Benacerraf, 1965). We can purport to define numbers in set theory by identifying those sets that capture the relevant properties and patterns of behaviour. But there are different, equally valid ways of doing this. On the von Neumann account, a number is the set of all its predecessors, with zero being the empty set.

$$\begin{aligned} \text{“0”} &\longrightarrow \{\} \\ \text{“1”} &\longrightarrow \{\{\}\} \end{aligned}$$

$$\begin{aligned} \text{"2"} &\longrightarrow \{\{\}, \{\{\}\}\} \\ \text{"3"} &\longrightarrow \{\{\}, \{\{\}\}, \{\{\}, \{\{\}\}\}\} \end{aligned}$$

That is, "1" is represented by {"0"}, "2" by {"0", "1"}, and "3" by {"0", "1", "2"}, and so on. The notion of "less than" ($<$) is then captured by set membership.

$$\text{"1"} < \text{"3"} \longrightarrow 1 \in 3$$

With the Church-style analysis, a number is a set containing its predecessor, where zero is the empty set.

$$\begin{aligned} \text{"0"} &\longrightarrow \{\} \\ \text{"1"} &\longrightarrow \{\{\}\} \\ \text{"2"} &\longrightarrow \{\{\{\}\}\} \\ \text{"3"} &\longrightarrow \{\{\{\{\}\}\}\} \end{aligned}$$

The notion of "less than" ($<$) is then captured by the transitive closure of the subset relation.

$$\text{"1"} < \text{"3"} \longrightarrow 1 \subset^+ 3$$

Unfortunately, there are crucial ways in which these set-theoretic characterisations are inappropriate. They have behaviours that go beyond number theory. For example, the reduction of numbers to set theory means that statements such as " $1 \in 3$ ", and " $1 \subset 3$ " can be expressed, even though they do not reflect any intuitions about numbers. Their truth conditions are an artefact of the representation rather than numbers as such. Furthermore, these unintended behaviours of "numbers" are different for different characterisations. Thus set theory fails to provide a faithful, canonical interpretation of numbers. We might argue that in some sense set theory cannot provide a "transparent" representations of numbers: there are artefacts of the behaviour of the representation that allow us to distinguish the behaviour of the representation from that of the numbers themselves (cf. Harman, 1990).

The message of Benacerraf's argument is that set-theoretic characterisations cannot be definitive — there is a disparity between the theory and the "data". We may also question whether set-theory alone can ever be said to "explain" or "describe" numbers, given that we appear to require a pre-existing structural understanding if we are to make sense of these set-theoretic characterisations.

Similar concerns about such ontological reductions arise in natural language semantics, where on the standard account:

1. questions *are* partitions of worlds;
2. answers *are* elements of a partition of worlds; and
3. there is no difference between a question and (the set of) its possible answers.

This is in addition to technical issues with the standard account. To pick a couple of examples: propositions have insufficiently fine-grained intensionality; and some of the deontic and imperative “paradoxes”, such as Chisholm’s “contrary-to-duty” obligations (Chisholm, 1963) appear in part to be due to unintended consequences of the conventional possible-worlds interpretations. Effort has to be spent addressing issues that arise as an unintended consequence of reducing everything to a set-theoretic possible worlds analysis.

As noted in the introduction, some of the problematic issues concerning a reductive analysis of meaning are not confined to set-theoretic models. For example, some use λ -abstracts for questions (Ginzburg and Sag, 2000; Hausser, 1983; Hausser and Zaefferer, 1979; von Stechow and Zimmermann, 1984). There is then a potential ontological difficulty in that within the formalisation, questions are indistinguishable from other λ -abstracts. Others have proposed using λ -abstracts and pairs for structured propositions (von Stechow, 1982; Krifka, 2001). Again, there is a potential ontological problem here in that such pairings are not distinguishable from other pairings involving λ -abstracts.

Some of the methodological and ontological concerns about reducing meaning to set-theory, or λ -calculus, may also apply to accounts that adopt other pre-existing and “independently motivated” frameworks, such as constructive type theory, depending on how the interpretation is formulated.

Returning to the question of data, in general formal semanticists are concerned about capturing intuitions about linguistic behaviour, both in isolation and in context. The data is messy, and we are obliged to consider abstractions. Even when resorting to corpus studies of spontaneous language use, or to psycho-linguistic experiments — perhaps to avoid some of the limitations of personal intuitions and introspection — there is the unavoidable issue of abstraction and interpretation. It seems impossible to

consider abstractions without some form of ontological framework with appropriate structural relations and patterns: the nature of our abstractions, and our characterisations of the data, invariably appeals to ontological intuitions (cf. Dummett, 1991).

Ultimately, it seems clear that characterisations of semantic theory within the standard account cannot be taken to *define* the various notions to which we appeal. Any narrative, and an appropriate intermediate representation, ought to be considered an essential part of the theory, providing a representation that is faithful to the intended interpretation. What we could argue is that, given a choice, it would be better if the ontological presumptions and intuitions of the narrative were considered an intrinsic part of our semantic analysis, and were subject to the same formal rigour (cf. Feferman, 1992).

6. An Alternative Approach

An alternative to the reductive model-theoretic approach is to adopt a clean, minimal meta-theory. We can then seek to formalise ontological intuitions non-reductively within this framework, and then go on to formalise meaning in terms of these ontological notions. Essentially this is a proposal to accept and embrace some form of representationalism (and structuralism). An example of a clean, minimal meta-theory is Typed Predicate Logic (TPL, Turner, 2008, 2009). This theory has four basic judgements.

1. T Type (T is a type)
2. $t : T$ (t belongs to type T)
3. p Prop (p is a proposition)
4. p True (p is (a) true (proposition))

These appear in sequents of the form

$$\Gamma \vdash \Phi$$

This represents the claim that judgement Φ follows in the context of judgements Γ . All of the syntactic and logical behaviour of a representation language can be expressed by way of sequent rules of the following form:

$$\frac{\Gamma_1 \vdash \Phi_1 \quad \Gamma_2 \vdash \Phi_2 \quad \dots \quad \Gamma_n \vdash \Phi_n}{\Gamma \vdash \Phi}$$

This says that $\Gamma \vdash \Phi$ follows from the premises $\Gamma_i \vdash \Phi_i$ (for $1 \leq i \leq n$).

In effect we can formulate semantics and grammar in the same system. We also treat the logic as a first class notion that does not require a set-theoretic model to give it meaning above and beyond what is intended in the narrative. We first exemplify the use of TPL by sketching rules for formation (F), introduction (+) and elimination (−) for propositional logic’s conjunction (\wedge) and universal quantification (\forall). The formation rules effectively define the syntax of the expressions, and their logical behaviour is characterised by way of the introduction and elimination rules. Here we highlight some of the salient expressions in the object language as an aid to readers who are not very familiar with these kinds of rules.¹³

Conjunction:

$$\frac{\Gamma \vdash p \text{ Prop} \quad \Gamma \vdash q \text{ Prop}}{\Gamma \vdash \boxed{p \wedge q} \text{ Prop}} \wedge F$$

$$\frac{\Gamma \vdash \boxed{p} \text{ True} \quad \Gamma \vdash \boxed{q} \text{ True}}{\Gamma \vdash \boxed{p \wedge q} \text{ True}} \wedge + \quad \frac{\Gamma \vdash \boxed{p \wedge q} \text{ True}}{\Gamma \vdash \boxed{p} \text{ True}} \wedge -$$

Universal quantification:

$$\frac{\Gamma, x : T \vdash p \text{ Prop}}{\Gamma \vdash \boxed{\forall x \in T \cdot p} \text{ Prop}} \forall F$$

$$\frac{\Gamma, x : T \vdash \boxed{t} \text{ True}}{\Gamma \vdash \boxed{\forall x \in T \cdot t} \text{ True}} \forall + \quad \frac{\Gamma \vdash \boxed{\forall x \in T \cdot t} \text{ True} \quad \Gamma \vdash \boxed{s : T}}{\Gamma \vdash \boxed{t[s/x]} \text{ True}} \forall -$$

Although there is not space to elaborate on this point, it is pertinent to observe that formulating a logic in this framework allows us to develop a fine grained analysis of intensionality without possible worlds (see Turner, 2005, for example). We may still need a model for this logic, which might be formulated in set theory, but its role is only to assist in proofs of its consistency.

We can also sketch a theory of questions in this framework. Questions can be taken to be expressions of the form “[$x \in T \mid \phi$]” where x corresponds to the “missing” component identified by a wh-clause, T is its type, and ϕ is the body of the question, as in the clause “*is running*” in the question

¹³Note that all the formalisations as presented here are intended merely to illustrate the approach; they do not by themselves constitute fully-fledged semantic theories.

“*who is running?*”. Here is a formation rule for question types.

$$\frac{\Gamma \vdash T \text{ Type}}{\boxed{\text{Question}(T)} \text{ Type}}$$

The formation rule for the representation of questions themselves would then be something like:

$$\frac{\Gamma, x : T \vdash \phi \text{ Prop}}{\Gamma \vdash \boxed{[x \in T \mid \phi]} : \text{Question}(T)}$$

Answers can be analysed in terms of structured propositions, represented by “ $\langle f \mid t \rangle$ ”, where f is the part of the term that is in focus, and t is the topic (following the approach of Krifka, 2001). For example, an answer to the question “*who is running*” might be “*John is running*”, where “*John*” is the focus, and “*is running*” is the topic. This would be represented by something akin to “ $\langle \text{John} \mid \lambda x.x \text{ is running} \rangle$ ”. The syntax of an answerhood relation can be governed by the following rule.

$$\frac{\Gamma, x : T \vdash \phi \text{ Prop} \quad \Gamma \vdash a : T}{\Gamma \vdash \boxed{\langle a \mid \lambda x.\phi \rangle \text{ ans } [x \in T \mid \phi]} \text{ Prop}}$$

We have a true answer when “ $\langle a \mid \lambda x.\phi \rangle \text{ ans } [x \in T \mid \phi]$ ” is a well-formed proposition *and* “ $\phi[a/x]$ ” is true. For our example,

$$\langle \text{John} \mid \lambda x.x \text{ is running} \rangle \text{ ans } [x \in T \mid x \text{ is running}]$$

would be a proposition, if “*John*” is of type B , and it would be a true answer if John were indeed running. As can be seen, there is no reduction to set-theoretic possible worlds, and our representation of questions and answers can be specific to those categories. For a more complete presentation along these lines see Fox (2013).

In the case of imperatives, we can state rules governing the judgement that something is an imperative, and that an imperative has been satisfied. The notions of being an imperative, and being satisfied (“fulfilled”, or “complied with”) can be formulated as judgements or types. Here are two candidate rules governing the formation of disjunctive imperatives and their satisfaction conditions (where a and b are imperatives).

$$\frac{\Gamma \vdash a \text{ Imperative} \quad \Gamma \vdash b \text{ Imperative}}{\Gamma \vdash \boxed{a \vee b} \text{ Imperative}} \vee F$$

$$\frac{\Gamma \vdash \boxed{a} \text{ Satisfied} \quad \Gamma \vdash b \text{ unSatisfied}}{\Gamma \vdash \boxed{a \vee b} \text{ Satisfied}} \vee +$$

To give an example, the first rule says that “*Watch television or go to the beach!*” is an imperative if both disjuncts are imperatives. The second rule says that the disjoined imperative is deemed satisfied in the event that the first disjunct is satisfied, and the second is unsatisfied; that is, the subject is watching television and has not gone to the beach.

The framework allows us to concentrate on our intuitions rather than the problem of formalising a set-theoretic possible worlds characterisation, and dealing with any unintended consequences.

The flexible approach to judgements in TPL allows us to consider our intuitions about hybrid expressions, such as so-called “pseudo imperatives” (Franke, 2005a,b). The following rules show how we can type an expression of the form “ $a \wedge p$ ”, where “ a ” is an imperative and “ p ” a proposition (as in “*Take another step and I will shoot*”, or “*Have another beer and you will be happy*”), and “ $a \vee p$ ” (as in “*Stand still or I’ll shoot!*”). Both the conjunctive and disjunctive forms combine a proposition with an imperative. Arguably, in the conjunctive case the result is a proposition expressing a threat or promise. In the disjunctive case there could be said to be both imperative and propositional content, where the latter expresses a salient alternative in the event of non-compliance with the imperative content. Within TPL, the typing for these hybrid expressions can be captured by rules of the following form.

$$\frac{\Gamma \vdash a \text{ Imperative} \quad \Gamma \vdash p \text{ Prop}}{\Gamma \vdash \boxed{a \wedge p} \text{ Prop}}$$

$$\frac{\Gamma \vdash a \text{ Imperative} \quad \Gamma \vdash p \text{ Prop}}{\Gamma \vdash \boxed{a \vee p} \text{ Prop}} \quad \frac{\Gamma \vdash a \text{ Imperative} \quad \Gamma \vdash p \text{ Prop}}{\Gamma \vdash \boxed{a \vee p} \text{ Imperative}}$$

Unlike standard type theory, expressions can belong to more than one type.

Arguably, detaching the formal analysis and characterisation from a reductive model-theoretic framework makes it more straight-forward to consider incoherent commands, as their representation in the framework need not result in an inconsistency within the semantic theory itself. See Fox (2012a) for a more complete formalisation of imperatives along these lines. A theory of deontic logic within TPL is also outlined by Fox (2012b).

7. Explanation and Description

As mentioned in Section 4, one argument that is sometimes made (informally at least) is that model-theoretic approaches that work with a pre-existing, independently motivated theory offer some kind of *explanation*. In contrast, a more bespoke, axiomatic approach might be characterised as merely *describing* the data in a different form. A counter-argument is that reductive model-theoretic accounts require us to identify and circumscribe the relevant behaviour, and that they too must assume that we have some characterisation of the behaviour we wish to circumscribe that is independent of its model theoretic interpretation. There is often an appeal to these intuitions in the description of the intended meaning of the model-theoretic analysis. It is not then clear what any supposed “explanatory power” of the *formal* theory then amounts to if *informal* narrative, and structural insights, are still required to make sense of the theory.¹⁴

It might be argued that a set-theoretic approach to formalising semantics, and the way in which ontological distinctions are removed, allows us to find a unifying analysis when faced with cross-categorial phenomena, such as disjunction. But in principle there is nothing to prevent us considering more abstract characterisations of such notions outside set-theory by identifying and formalising generic, overarching patterns of behaviour.

In any case, given that we need narrative to characterise intended interpretations, and that we are interested in rigorous formal analyses, then it seems appropriate to formalise the insights of that narrative.

It might not be entirely clear what the “data” is, or whether and in what sense a formal analysis of meaning is a scientific theory, given the difficulty in determining an objective criteria for “correctness” of all aspects of the analysis. But such formalisation does at least allow us to check coherence of insights, and attempts to “explain” meaning in terms of ontological assumptions.

¹⁴This is not to say that, for example, a possible worlds perspective is inappropriate; some useful insights can be obtained by considering how a problem might be framed in terms of possible worlds.

8. Conclusions

Semantics is concerned with characterising aspects of meaning. As with any intellectual endeavour, it is natural to presuppose a framework of ontological notions and classifications in which to express abstractions and generalisations about the behaviour of language.

Formal semantics is in essence concerned with characterising behaviour with unambiguous rules. But such formalisation also relies upon a framework of ontological notions and characterisations. In model-theoretic semantics, these may only be evident in the narrative that accompanies the formalisation.

The view being put forward here is that such substantive narrative is part of the semantic theory, and as such it too should be subject to the same rigour. In brief, ontological insights are important. We might even go so far as to say that they may constitute a core theoretical contribution of formal semantics.

Other arguments can be made, of course. Ultimately, it seems there will always be subtle aspects of meaning that are inevitably lost by our clumsy attempts at formalisation. In selecting patterns of behaviour and intuitions that we are able to capture, we cannot help but ignore other aspects of meaning.¹⁵

References

- Beltrami, Eugenio, 1868. Saggio di interpretazione della geometria non-euclidea. *Giornale di Matematiche* VI, 285–315. (Essay on an interpretation of non-euclidean geometry).
- Benacerraf, Paul, 1965. On what numbers could not be. *The Philosophical Review* 74, 47–73.
- Burke, Kenneth Duva, 1966. *Language as Symbolic Action*. University of California Press.
- Chisholm, Roderick M., 1963. Contrary-to-Duty Imperatives and Deontic Logic. *Analysis* 24, 33–36.
- Chomsky, Noam, 1965. *Aspects of the Theory of Syntax*. Cambridge, Massachusetts: MIT Press.
- Cocchiarella, Nino, 1981. Richard Montague and the logical analysis of language. In G. Fløistad (ed.), *Philosophy of Language/Philosophical Logic*, volume 2 of *Contemporary Philosophy: A New Survey*, 113–155. The Hague: Martinus Nijhoff.

¹⁵Cf. Burke's "terministic screens" (Burke, 1966).

- Cocchiarella, Nino, 1997. Formally-oriented work in the philosophy of language. In John V. Canfield (ed.), *Philosophy of Meaning, Knowledge and Value in the 20th Century*, 39–75. London; New York: Routledge.
- Dummett, Michael, 1991. *The Logical Basis of Metaphysics*. London: Duckworth, and Cambridge MA: Harvard University Press.
- Feferman, Sol, 1992. What rests on what? The proof-theoretic analysis of mathematics. Invited lecture, 15th International Wittgenstein Symposium: Philosophy of Mathematics, Kirchberg/Wechsel, Austria.
- Fox, Chris, 2000. *The Ontology of Language*. Stanford: CSLI.
- Fox, Chris, 2012a. Imperatives: a Judgemental Analysis. *Studia Logica* 100(4), 879–905.
- Fox, Chris, 2012b. Obligations and Permissions. *Language and Linguistics Compass* 6(9), 593–610.
- Fox, Chris, 2013. Axiomatising Questions. In Vit Puncocar and Petr Svarny (eds.), *Logica Year Book 2012*, 23–34. College Publications.
- Fox, Chris and Raymond Turner, 2012. In Defense of Axiomatic Semantics. In Piotr Stalmaszczyk (ed.), *Philosophical and Formal Approaches to Linguistic Analysis*, 145–160. Ontos Verlag.
- Franke, Michael, 2005a. How And How Not To Employ Discourse Relations To Account For Pseudo-Imperatives. In Paul Dekker and Michael Franke (eds.), *Proceedings of the Amsterdam Colloquium*, 83–88. ILLC/Department of Philosophy, Amsterdam: University of Amsterdam.
- Franke, Michael, 2005b. *Pseudo-Imperatives*. Master's thesis, Institute for Logic, Language and Computation, University of Amsterdam.
- Ginzburg, Jonathan and Ivan Sag, 2000. *Interrogative Investigations: The Form, Meaning and Use of English Interrogatives*. Stanford: CSLI.
- Groenendijk, Jeroen and Martin Stokhof, 1997. Questions. In Johan van Benthem and Alice ter Meulen (eds.), *Handbook of Logic and Linguistics*, 1055–1124. Amsterdam: North Holland.
- Hamblin, C. L., 1973. Questions in Montague English. *Foundations of Language* 10, 41–53.
- Harman, Gilbert, 1990. The Intrinsic Quality of Experience. *Philosophical Perspectives* 4, 31–52. Issue title: Action Theory and Philosophy of Mind.
- Hausser, Roland, 1983. On questions. In F. Kiefer (ed.), *Questions and answers*, 97–158. Dordrecht: Reidel.
- Hausser, Roland and Dietmar Zaefferer, 1979. Questions and Answers in a Context Dependent Montague Grammar. In Frans Guenther and Max Schmidt (eds.), *Formal Semantics and Pragmatics for Natural Languages*. Dordrecht: Reidel.
- Karttunen, Lauri, 1977. Syntax and Semantics of Questions. *Linguistics and Philosophy* 1, 3–44.
- Krifka, Manfred, 2001. For a Structured Meaning Account of Questions and Answers. In C. Fery and W. Sternefeld (eds.), *Audiatu Vox Sapientia. A Festschrift for Arnim von Stechow*, 287–319. Berlin: Akademie Verlag.
- Lascarides, Alex and Nicholas Asher, 2004. Imperatives in Dialogue. In Peter Kühn-

- lein, Hans Rieser and Henk Zeevat (eds.), *The Semantics and Pragmatics of Dialogue for the New Millennium*, 1–24. Benjamins.
- Link, Godehard, 1982. The Logical Interpretation of Plurals and Mass Terms: a Lattice Theoretic Approach. In Reiner Bauerle, Christophe Schwarze and Arnim Von Stechow (eds.), *Meaning, Use, and the Interpretation of Language*. Berlin: de Gruyter.
- Montague, Richard, 1969. On the nature of certain philosophical entities. *The Monist* 53(2), 159–194. Also in *Formal Philosophy: Selected Papers of Richard Montague*, Yale University Press, 1974, edited by R. H. Thomason.
- Montague, Richard, 1970a. English as a Formal Language. In Bruno Visentini (ed.), *Linguaggi nella società e nella tecnica*, 189–223. Mailand. Also in *Formal Philosophy: Selected Papers of Richard Montague*, Yale University Press, 1974, edited by R. H. Thomason.
- Montague, Richard, 1970b. Universal grammar. *Theoria* 36, 373–398. Also in *Formal Philosophy: Selected Papers of Richard Montague*, Yale University Press, 1974, edited by R. H. Thomason.
- Montague, Richard, 1973. The Proper Treatment of Quantification in Ordinary English. In K. J. J. Hintikka, J. M. E. Moravcsik and P. Suppes (eds.), *Approaches to Natural Language*, 221–242. Dordrecht: D. Reidel. Synthese Library. Also in *Formal Philosophy: Selected Papers of Richard Montague*, Yale University Press, 1974, edited by R. H. Thomason.
- Montfort, Nick, 2003. *Twisty Little Passages: An Approach to Interactive Fiction*. Cambridge, Massachusetts: MIT Press.
- Partee, Barbara H., 2001. Montague Grammar. In Neil J. Smelser and Paul B. Baltes (eds.), *International Encyclopedia of the Social and Behavioral Sciences*. Oxford: Pergamon/Elsevier Science.
- Partee, Barbara H., 2013. Montague’s “linguistic” work: Motivations, trajectory, attitudes. In Emmanuel Chemla, Vincent Homer and Grégoire Winterstein (eds.), *Proceedings of Sinn und Bedeutung 17, 8th–10th September, 2012*, 427–453. Paris: École Normale Supérieure (ENS).
- Portner, Paul, 2005. The Semantics of Imperatives Within a Theory of Clause Types. In Kazuha Watanabe and Robert B. Young (eds.), *Proceedings of Semantics and Linguistic Theory 14*. Ithaca, NY: CLC Publications.
- Ranta, Aarne, 1994. *Type Theoretic Grammar*. Oxford: Oxford University Press.
- de Saussure, Ferdinand, 1916. *Cours de linguistique générale*. Lausanne and Paris: Payot. Edited by C. Bally and A. Sechehaye, with the collaboration of A. Riedlinger.
- Seegerberg, Krister, 1990. Validity and Satisfaction in Imperative. *Notre Dame Journal of Formal Logic* 31(2), 203–211.
- von Stechow, Arnim, 1982. Structured Propositions. Manuscript.
- von Stechow, Arnim and Thomas Ede Zimmermann, 1984. Term answers and contextual change. *Linguistics* 22, 3–40.
- Thomason, Richmond H., 1974. Introduction. In R. H. Thomason (ed.), *Formal Philosophy: Selected Papers of Richard Montague*, 1–69. Yale University Press.

- Tichý, Pavel, 1978. Questions, Answers, and Logic. *American Philosophical Quarterly* 15(4), 275–284.
- Turner, Raymond, 2005. Semantics and Stratification. *Journal of Logic and Computation* 15(2), 145–158.
- Turner, Raymond, 2008. Computable Models. *Journal of Logic and Computation* 18(2), 283–318.
- Turner, Raymond, 2009. *Computable Models*. London: Springer-Verlag.
- Westerståhl, Dag, 1998. On Mathematical Proofs of the Vacuity of Compositionality. *Linguistics and Philosophy* 21(6), 635–643.
- von Wright, Georg Henrik, 1951. Deontic Logic. *Mind* 60, 1–15.
- von Wright, Georg Henrik, 1953. *An Essay in Modal Logic*. New York: Humanities Press.
- Zadrozny, Wlodek, 1994. From Compositional to Systematic Semantics. *Linguistics and Philosophy* 17(4), 329–342.