

Rough draft, not for citation. Comments welcome. In preparation for *Prospects for Meaning* ed. Richard Schantz.

Gabriel M. A. Segal

Five flies in the ointment: some challenges for traditional semantic theory

1. Five challenges for traditional semantic theory

Truth-conditional semantic theory for natural language appears to be flourishing. The work deploys apparatus from the tradition of Frege, Tarski, Carnap, Davidson and Montague.¹ But is all as well as it seems? Work in this tradition relies, obviously, on notions like truth, reference, satisfaction and extension. It assumes that many sentences of natural language, or utterances of them, are true. It assumes that many words of natural language refer to, extend over, or are true of things in the world. If those assumptions are faulty, then all is not well with semantics for natural language. And these core assumptions do face serious challenges. In this chapter, I will consider a few well-known ones that are particularly interesting, mostly emanating from the work of Noam Chomsky. I will list them first, then look at how they might be met.

Challenge 1:

¹ Not all the major contributors to the tradition believed that natural language has a formal semantics. In my view, Tarski and Carnap did not. Montague and Davidson did. And Frege did too, at least in his later years. See Segal (2006) for discussion.

A watermelon is green on the outside and red on the inside. I visit a greengrocer's shop. I want to know if the watermelon I have picked up is one of those that is red on the inside, as opposed to some other colour. In response to my query the greengrocer says: 'Yes it is red'. I buy the watermelon. I visit my friend, the artist. I have the watermelon with me. He asks to borrow it so that he can include it in his still-life. He explains that this work hinges on contrasts between red and other colours. 'I shall place this tomato, which is red, next to the watermelon, which is not red'. Challenge: the watermelon both is and is not in the extension of 'red', so the whole idea of a predicate having an extension is kaput.

Here is a very similar example involving 'water'. I am in a restaurant, I ask the waiter for a glass of water. He brings me a cup of tea. I say 'That isn't water'. Some tea leaves have been dumped into the reservoir that supplies water to my house, as a purifier. What comes out of my tap is chemically indistinguishable from the tea in the restaurant. But I still call it 'water'. Challenge: tea both is and is not in the extension of 'water'. The whole notion of extension is kaput.²


Challenge 2:

² The watermelon example is adapted from Travis, pc. The 'water' example is from Chomsky (1995, 22). Here is another one. In an anthropology class the lecturer says 'Humans drink many varieties of impure water. In poor countries the water is sometimes dirty or contaminated. In more prosperous ones, impurities are deliberately added to enhance the taste or other beneficial properties – tea leaves for example'. I find this example very compelling. But it might be objected that the anthropologist is using 'water' in a technical, scientific sense and that his term is not the same as the natural-language homonym. I am not convinced. The anthropologist could be speaking ordinary language. One type of folk language game involves adopting terms from scientific discourse.

Theseus's ship sails the seas. As its timbers age, they are replaced with new ones. The old timbers are kept, restored and combined to form a ship. Call the original ship 'A', the ship that had all its timbers replaced, 'B', and the ship that was, at the end of the story, made from the original planks, 'C'. The individuation conditions we associate with 'ship' entail that A is the same ship as B and that A is the same ship as C. But B is not the same ship as C. So 'x is the same ship as y' extends over <A, B> and over <A, C>, but not over <B, C>. The whole notion of extension is kaput.

Challenge 3:

Like ships like people. Individuals A and B are in the hands of a mad scientist. She explains that she has an apparatus that will copy all the information in A's brain into B's and vice versa.³ After that, one person will receive a million pounds and the other will be tortured. Call the person whose body is A's body at T1, the "A-body-person-at-T1", the person whose body is B's body at T1, the "B-body-person-at-T1", etc.. Then we have this picture:



³ We can assume that the entire wiring pattern is copied.

The scientist asks B to decide which body-person will be tortured and which will receive the million pounds.⁴ B chooses the A-body-person-at-T3. Upon awakening, the A-body-person-at-T3 recalls having made the decision as the B-body-person-at-T1 and is very pleased. Hence <The-A-Body-Person-at-T3, The-B-Body-Person-at-T1> are in the extension of 'x is the same person as y'.

A person is a human being. The A-body-person-at-T1 is the same human being as the A-body-person-at-T3. Hence the A-body-person-at-T1 is the same person as the A-body-person-at-T3. Hence <The-A-Body-Person-at-T3, The-A-Body-Person-at-T1> are in the extension of 'x is the same person as y'. But of course <The-A-Body-Person-at-T3, The-B-Body-Person-at-T3> are not in the extension of 'x is the same person as y'. The whole idea of extension is kaput.

Let us suppose the A-body-person-at-T1 is called 'Fred'. The name 'Fred' refers to the A-body-person-at-T3. And it refers to the B-body-person-at-T3. But it cannot refer to both. The whole idea of reference is kaput.

Challenge 4:

An utterance of (1) could easily be true.

- (1) John gave a book to Mary, but she already had it, so he read it himself then shredded it.

⁴ The example is from Williams (1970)

But then 'book' extends over objects that are both abstract and concrete and 'it' refers to something that is both abstract and concrete.⁵ But nothing is both abstract and concrete. So the ideas of extension and reference are kaput.

Challenge 5:

An utterance of (2) could be true.

(2) The average American family has 2.3 children.

But traditional semantic theory is committed to treating the sentence as having something like the form partially depicted in (2*):

(2*) $(\exists x)(\text{family}(x) \ \& \ \text{average-for-American}(x) \ \& \ \text{has} \langle 2.3 \text{ children}, x \rangle)$

And utterances of (2*) are, if meaningful at all, then false.

Similarly, an utterance of (3) could be true:

(3) Joe Sixpack's priorities are changing

But according to traditional semantics, all such utterances should be false since the proper name doesn't refer.

Likewise for (4) and (5):

(4) Sherlock Holmes is more famous than any real detective

(5) Sherlock Holmes is a fictional detective

⁵ Chomsky (2003). See Ludlow (2003), to which Chomsky is replying, for discussion of this example and the general topic of this paper.

And (6) too, because traditional semantics treats ‘some’ as an existential quantifier, implying real-world existence:

(6) Some fictional detectives are more famous than any real detective

And finally:

(7) There is a flaw in the argument.

An utterance of (7) could be true. According to the traditionalist, this could only be so if flaws were things. But they are not.⁶

The problems we have with the individuation of the denizens of the manifest image: ships, water, books, people and so on might lead one to think of the manifest image as ‘a dream modulated by sensory input.’⁷ (Llinas 1987, quoted with approval in Chomsky 1995). According to Chomsky, all these problems go away if we refrain from thinking of manifest kinds and particulars as real-world objects. They are not things out there for our commonsense words and concepts to extend over or refer to. Real kinds and particulars are to be found only in the scientific image.⁸ And there is no overlap between the two. It is, according to Chomsky, ok to say things like ‘books exist and so do people, but unicorns don’t’ using ordinary

⁶ Chomsky (1981 324), Hornstein (1984 58).

⁷ The actual quote is ‘perception is a dream modulated by sensory input, from Llinas 1987.

⁸ The expressions ‘manifest image’ and ‘scientific image’ are from Sellars (1962). I don’t think Chomsky uses the terms. But they are apt.

language. In Wittgensteinian spirit: language is alright as it is. But this is a relatively non-committal usage ('without metaphysical import'). It would not do, for example, to translate the claim into (8) and then give it a standard semantics:

(8) $(\exists x)(\exists z)(\text{book}(x) \ \& \ \text{person}(y) \ \& \ \neg(\exists z)(\text{unicorn } z))$.

For Chomsky "the semantic properties of an expression focus attention on selected aspects of the world as it is taken to be by other cognitive systems, and provide intricate and highly specialized perspectives from which to view them, crucially involving human interests and concerns even in the simplest cases." For example if I am inside a house I can clean it, affecting the inside, but I cannot see it, unless an exterior surface is visible. And if I am inside it, I cannot be near it, even though, in the unmarked case, it is surface (like a cube, to which the same point applies). The semantic properties of an expression provide "instructions" to "conceptual-intentional" systems for building up these intricate representations, which can then be used to interpret speech acts. (Chomsky 2005, 20).

Well maybe that is right. It is an appealing picture. But what is the alternative? Certainly, objects and kinds in the manifest image are not reducible to objects and kinds in the scientific image. Even manifest material particulars are not physical objects, in the sense of being objects that could be identified in physics. Ships and cups are not bundles of molecules. Remove a molecule from a healthy ship and you are left with the same ship, but not with the same bundle of molecules. The same goes for orangutans, planets, elm trees and steel girders. But all of these feature in the special sciences: geology, astronomy, botany and material science respectively.

These sciences seem to be perfectly respectable. If orangutans, planets and steel girders can be the subject of serious scientific study without being reducible to objects in physics, then maybe the same could be said for some of those things that appear in the manifest image only. Maybe – if, that is, we can get around the problems of individuation introduced in the Challenges and their kin.⁹

So how should a traditional semantic theorist who is a realist about the manifest image respond to the challenges? The next section outlines what I take to be the most promising package of answers.

1.2. Addressing the Challenges

1.2.1 watermelons and water

Let us begin with the problem of red melons and glasses of impure water. The solution I recommend is offered by Rothschild and Segal (in preparation).

The traditional semanticist has two main options. He could deny the apparent data and hold that the sentences don't have context-dependent truth values after all. He might argue that the greengrocer's utterance of 'It is red' isn't really true. Only objects that are largely red on the outside are really red, he might say. He might argue that tea really is water, but that we are reluctant to call it so in a restaurant because that would be confusing. Or, indeed, he could take the opposite line and argue that in the tap context, I am wrong to call the tea 'water'.

Denying the apparent data is methodologically questionable. Speakers' intuitions pretty consistently opt for context-sensitivity of truth

⁹ Such as sorites paradoxes, about which I have nothing much to say here.

values in these and a wide variety of similar cases. Of course, it is theoretically possible that these intuitions rest on some kind of confusion of semantic and pragmatic factors, or some other confounding factor. But there doesn't appear to be any evidence for this.

The alternative is to claim that there is some sort of indexicality going on. There are three ways to pursue this option. One: one can suppose that there are one or more variables present at logical form, in the underlying syntax of the sentences. Zoltan Szabo, for example, has argued that 'red' associates with two variables at logical form, one designating a comparison class to specify how red the object has to be and one designating the part of the object that has to be red: 'red(c, p)' (Szabo 2001). Two: one can suppose that one or more variables are supplied by the context. These are not present in the object language. But they are present in the meta-language and enter into the interpretation of the target sentences. Third: one can suppose that words like 'red' and 'water' are themselves indexical: they have different extensions in different contexts of utterance. Thus in the greengrocer's context the watermelon is in the extension of 'red' and in the context of the artist's studio it is not.

In my view, only the third option is viable. Let's switch to 'tall'. Suppose that Fred is 6ft tall. Then, in a typical conversational context in the Netherlands, an utterance of 'Fred is tall' would be false. In a typical conversational context in China, an utterance of 'Fred is tall' would be true. Suppose now that 'tall' means something like 'tall[c]', 'c' being a variable that picks out a comparison class.¹⁰ In the each context, 'c' picks out some salient group. Fred is not tall relative to the Dutch. Fred is tall relative to the Chinese. So everything is fine.

¹⁰ As proposed by Ludlow (1989).

Except it isn't. The proposal gets the extensions right. But not the meanings. An utterance of 'tall' in a context that determines a specific comparison class doesn't mean what an utterance of 'tall[c]' with 'c' referring to that class would mean. Let me explain why.

One might know, e.g., that Fred is tall relative to Chinese people without knowing what it is to be tall relative to Chinese people: for example, simply because one doesn't know how tall the Chinese people are. Now, please imagine Fred in China. He is in a supermarket with some local friends. They want to purchase an item that is on a high shelf and they can't quite reach it. One of them says (8):

(9) Fred can reach it. He is tall.

Compare (9) with (10), assuming that the variable refers to the group of friends:

(10) Fred can reach it. He is tall[c]

Someone who didn't know how tall members of the group were could know that the thought expressed by 'He is tall[c]' in (10) is true, without knowing that the thought expressed by 'He is tall' in (9) is true.¹¹ Hence the utterances express different thoughts, hence they mean different things.

Notice also that there is a certain ad hocery about the extra-variable approach.

¹¹ Let me put the point more clumsily but less ambiguously. Let us call the thought expressed by (10) 'T1' and the thought expressed by (9), 'T2': someone could stand an attitude of propositional knowledge to T1 without standing in an attitude or propositional knowledge to T2. So, using the normal Fregean method of distinguishing contents, T1 and T2 are different thoughts.

What is the relevant parameter for counting as water in the restaurant context versus counting as water in the tap context? It is not comparison class. In the restaurant, the water in the water glass might, for all I know, contain a lower percentage of H₂O than the water in my water glass. In the restaurant, the relevant classification system is one for distinguishing different kinds of drink. Chemical composition is only one relevant factor among others.¹² So what would the extra variable pick out? ‘C’ for ‘classification system’ perhaps? Or simply ‘c’ for ‘context’. Again the proposal gets the extension right, but not the meaning. One could know that something is water[c] without knowing what it is to be water[c].¹³ Thus what is expressed by ‘water’ in a particular context of utterance differs from what is expressed by ‘water[c]’, with ‘c’ referring to that context.¹⁴

Let us consider option three, treating the target terms as indexicals. The next subsection simply expounds Rothschild and Segal’s proposal.¹⁵

1.2.1.1

¹² See Malt (1994) for an empirical study of some of the uses of ‘water’. Szabo (2006) denies that ‘water’ shifts its extension in the way Chomsky claims it does. Responding to Chomsky’s example, he writes: “I disagree: I think many of us would be reluctant to stand by both judgments upon learning the chemist’s verdict [that the stuff coming out of the tap and tea in a tea cup are chemically indistinguishable]; we might not know which one to give up, but that does not mean that they must have the same standing.” I disagree: most of my informants are happy with both judgments.

¹³ Could ‘c’ work like a special kind of indexical, such that understanding its use in a specific context entails knowing what the relevant parameter is, under the right description? Could ‘tall[c]’ mean, something like ‘tall for this context’, where one is only in a position to understand ‘this context’ if one knows what it is to be tall for the context? I doubt it very much. Suppose that you are in the supermarket in Beijing and you encounter the group of friends. You haven’t met them before. But you have had a few words with them about the shop and its goods. You are part of the conversation, when one of them says (9). You are blind and you don’t know how tall the friends are. Then you understand (9) in the manner of ‘He is tall for this context’. But you don’t know what it is to be tall for the context. So you don’t understand (9) in the way the speaker and the other participants do. You don’t know which thought the utterance expresses.

¹⁴ Any attempt to extend the variable approach will have to contend with many different kinds of parameter. Consider for example an utterance of ‘the shoes are under the bed’ when they are eight floors below. And, for another, ‘the door must be kept shut at all times’, which sometimes means: *except when you are passing through it*, and sometimes doesn’t. Thanks to Travis for both.

¹⁵ Rothschild and Segal have other arguments against the extra-variable approach.

Rothschild and Segal's account of indexical predicates is an extension of the standard approach taken by T-theoretical semanticists to singular indexicals, particularly 'that'. On this approach, one proves what are sometimes called 'conditionalized T-sentences'. These are items along the lines roughly sketched in (11):

(11) If u is an utterance of 'that is a watermelon', and the speaker uses 'that' in u to refer to x , then u is true iff x satisfies 'is a watermelon'.

The information specified in (11) is context-independent. Now suppose that you are in a particular context of utterance and you know that the speaker used 'that' to refer to watermelon α , then you can move on to (12):

(12) u is true iff α is a watermelon.

The extension of the basic idea to predicates, not an entirely straightforward affair, proceeds as follows.

An utterance of 'Fred is tall' is true in a certain Chinese context, but false in a certain Dutch context. This is because different standards are operative in the different contexts: to be tall by the standards of the Chinese context is easier than to be tall by the standards of the Dutch context. We distinguish tokens of 'tall' in different contexts by subscripting. And we keep contexts and tokens in line using numbers. Thus: if the Chinese context is the k th context, all tokens of 'tall' governed by the standards of that context are subscripted with 'k': "tall_k". "Tall_k" is a syntactic type (and 'k' is

a numeral). All and only utterances of ‘tall’ in the k th context are utterances of “tall $_k$ ” (although the numeral remains unpronounced). An object satisfies “tall $_k$ ” iff the object is tall by the standards of the k th context.

Correspondingly, we distinguish the sentences “Fred is tall $_j$ ”, “Fred is tall $_k$ ” etc.. And we think of these as evaluated relative to their contexts of utterance. Suppose again that the Chinese context is the k th. Then “Fred is tall $_k$ ” is true relative to the context because Fred satisfies “tall $_k$ ”, because Fred is tall by the standards of the context. Finally, we treat utterances as bearers of truth values absolutely: an utterance of a sentence in a context is true, absolutely, iff the sentence uttered is true relative to the context.

When we say “Fred is tall by the standards of the Chinese context, but not by the standards of the Dutch context”, we are using “tall” in a rather abstract and general way. This is how we use “tall” when we are theorizing about height in general and related matters. Let us call this “tall $_g$ ” (the g th context is the general one).

Putting these ideas together we get:

(A1) $(x)(n)(x \text{ satisfies “Fred”, } c_n \text{ iff } x=\text{Fred})$

(A2) $(x)(n)(x \text{ satisfies “is tall”}^n, c_n \text{ iff } x \text{ is tall}_g, c_n)^{16}$

(A3) $(S)(NP)(VP)(\text{If } S=NP^{\wedge}VP, \text{ then } ((n)(S \text{ is true, } c_n \text{ iff } (\exists x)(x \text{ satisfies } NP, c_n \text{ and } x \text{ satisfies } VP, c_n)))^{17}$

(A4) $(u)(n)(S) \text{ (if } u \text{ is an utterance of } S \text{ in } c_n, \text{ then } (u \text{ is true iff } S \text{ is true, } c_n))$

¹⁶ For all x , for all n , x satisfies the expression composed of ‘is tall’ concatenated with the n th index relative to the n th context iff x is tall by the standards of the n th context.

¹⁷ For any sentence S , any noun phrase NP and any verb phrase VP , if S is the concatenation of NP and VP , then, for any n , S is true relative to the n th context iff for some x , x satisfies NP relative to the n th context and x satisfies VP relative to the n th context.

(A1)-(A4) provide the context-independent T-theory. To get beyond this, you need to be in a particular context and to know what the standards of the context are. If you are in the k th context, you need information of this sort:

(A5) u_k is an utterance of $S_k = \text{'Fred is tall}_k\text{'}$ in c_k

(A6) $(x)(x \text{ is tall}_g, c_k \text{ iff } x \text{ is tall}_k)$

You can then combine the T-theory with the contextual information to get you to:

(A7) u_k is true iff $(\exists x)(x = \text{Fred } \wedge x \text{ is tall}_k)$

From which you might infer:

(A8) u_k is true iff Fred is tall $_k$

1.2.2 Ships

Given the way we individuate ships (artifacts generally, and other kinds of thing too), a ship can survive gradual replacement of all of its planks. A ship can also be gradually dismantled, have its planks removed, and then be put back together. So it can be that what began as a single ship can survive as two distinct ships. Given the logic of identity, this causes us a problem. We need to revise our individuation criteria for ships. A ship $_1$ can survive gradual replacement of all of its planks, but cannot be dismantled and put back together. Once taken to bits, it's history. A ship $_2$ can be dismantled and reassembled. But it cannot survive gradual replacement of its parts. Once

you have replaced fifty percent of its parts, it is history, and a new ship₂ has come to be.

It would be convenient if by ‘ship’ we sometimes meant *ship*₁ and other times meant *ship*₂ and the traditional semanticist could treat ‘ship’ like ‘tall’ and ‘red’. Life, though, is rarely convenient. *Ship*₁ and *ship*₂ are artificial concepts, not concepts deployed by Jane Winecooler.

So how do we deal with Jane Winecooler’s utterance of “Aristotle owns lots of ships.”? I suggest a supervenient approach. It is true iff “Aristotle owns lots of ship₁s” and “Aristotle owns lots of ship₂s” are both true, false iff they are both false and indeterminate if exactly one is true.

1.2.3 People

I believe that we are born dualists. We are born with an innate ‘theory of mind’, according to which minds are immaterial entities that could exist disembodied or move from body to body. There is decent (though not 100% conclusive) evidence for this from anthropology and psychology. (A) All natural human groups believe in ghosts (Boyer 2003). (B) Babies don’t expect humans to obey physical laws (Kuhlmeier et. al 2004). (C) Untutored children tend to believe that people survive the death of their bodies (Bering and Bjorklund (2004)). (d) The mind/body problem persists: after all these years of physicalism, we remain baffled at the idea that a physical thing can think and feel.¹⁸

¹⁸ For further discussion see Segal (forthcoming)

Cartesian dualism is false. The folk concept needs to be revised. But a workable concept of mind could retain many of its aspects. Minds cannot exist disembodied. But they can exist embodied. It seems to me that the concept of a mind that can move from one brain to another is not defective. We can coherently conceive of three-dimensional continuants individuated by relations of psychological continuity. In the Williams scenario there are things that fall under this slightly physicalist concept of mind: A's mind, which moves into B's body and B's mind, which moves into A's.

The Williams scenario also, of course, features human beings. Does 'person' pick out the mind, or the human being? Does 'Fred' pick out the mind, or the human being? The lack of agreement amongst philosophers about these issues suggests that the answer is: "it is indeterminate". It is indeterminate whether Fred is the A-body-person-at-T3 or the B-body-person-at-T3.

I used to think that that meant that it was then open to each of us to choose which concept we wished to express by 'person' (and 'self') and which object one would pick out with 'I'. And I want to be a mind. If I were the B-body-person-at-T1, I would want the A-body-person-at-T3 to get the reward.

But I no longer think I can choose to be a mind. If it is indeterminate whether I am a mind or a human being, then I can't resolve the indeterminacy by ditching my old concept of self and replacing it with the slightly physicalist concept of a mind. The human being is a thinking thing, just as the mind is. And it can't choose to be other than it is. Equally the mind already is the mind, and it doesn't need to make a choice in order to be self-identical.

If this is right, then it gives Fred in the Williams scenario a serious problem. When the scientist asks Fred to choose the fates of the later body-people, the human being should choose for the benefit of the B-body-person-at-T3 (or so it seems to me, at least) and the mind should choose for the benefit of the A-body-person-at-T3. And of course the two have not yet divided: the human being and the mind think and want and speak as one. If Fred chooses happily, without a feeling of conflict, then something loses out. If he chooses to benefit the A-body-person, then the human being loses out, if he chooses to benefit the B-body-person, then the mind loses out.¹⁹

There might be a problem here for persons, but that doesn't mean that there is one for semantic theorists. Again, the theorist might reasonably opt for a supervenient approach to 'person' and 'Fred'.

I suggest that the supervenient approach recommended here for 'ship' and 'person' might extend to a wide range of similar cases, cases in which lay concepts don't have fully determinate extensions. Consider, for example, the extensions of 'hammer', 'dog' and 'tooth' in the following scenarios. On a distant planet with no carpentry, there naturally occur objects that are intrinsically physically identical to hammers. On a distant planet, there are animals that are intrinsically physically just like Alsatians, but that share no ancestor with any Earth creature. On a distant planet, there occur naturally things exactly like your teeth. They grow in the ground by themselves and are not associated with eating. An android has objects just like dentures

¹⁹ For extensive discussion see and a different view, see Noonan (1998). For general discussion see Eric Olson's entry 'Personal Identity' in *The Stanford Encyclopaedia of Philosophy*.

permanently in place in its mouth. He uses them to chew food. Arguably, these examples reveal indeterminacy in the extension of the terms: ‘hammer’, ‘dog’ and ‘tooth’. Supervaluationism can make sense of our quotidian use of these terms.

1.2.4 Books

(1) means something like (1’):

(1’) John bought [(a copy of) [a book]_i]_j for Mary. But Mary already had (a copy of) it_i. So he read it_j then shredded it_j

But then nothing in the logical form of (1) needs to extend over anything that is both abstract and concrete. Sometimes ‘book’ extends over abstract things, and sometimes it extends over concrete things. Likewise some occurrences of ‘it’ that are anaphoric on ‘book’ refer to abstract objects and others refer to concrete ones. But no occurrence of ‘book’ or ‘it’ extends over anything that is both abstract and concrete. In (1) ‘book’ extends over book types, which are abstract. (Peter had already bought it for her, which is why she already had it.) The first ‘it’ refers to the type. The second and third refer to a concrete thing, the copy that John bought, read and shredded.

I don’t know how to get the meaning expressed in (1’) out of (1). I don’t know what the logical form of (1) is. I have no argument that it looks like (1’). But, as far as I can see, the problem of how to get the meaning (1’) out of (1) remains whether one is doing truth-theoretic semantics or describing

instructions to the conceptual-intentional systems. And there is no apparent reason to suppose that the problem would be any easier to solve if embedded in the latter enterprise.²⁰

Chomsky (p.c.) has objected to this proposal. Suppose I am holding a book and say (13):

(13) This book won the Pulitzer Prize

The traditional semantic theorist has to treat ‘This book’ as referring both to the type and to the copy. I am not convinced. What won the prize was the type. So it is reasonable to hold that ‘this book’ refers to the type. But it doesn’t also have to refer to the copy. It is true that the speaker draws attention to the copy as part of the pragmatic mechanism that secures reference to the type. But that doesn’t enter the semantics.²¹

1.2.5 Nonexistents

There are three standard strategies that a semantic theorist can deploy in relation to (2)-(7): (i) posit funny objects, such as average American

²⁰ Chomsky says very little about what a theory of instructions to conceptual-intentional systems would look like. He thinks we know very little about the matter. Paul Pietroski (Pietroski forthcoming) has sketched a small proposal about how to do Chomsky semantics. It appears to be fairly widely agreed that if Chomsky right about the nature of semantics, then a lot of what truth-theoretic theorists actually do could be recast in the theory of instructions. This looks plausible in the case of the little T-theory for ‘tall’ presented above.

²¹ The problematic phenomenon is not restricted to books, nor to an abstract/concrete duplex. We also have ‘I painted the door brown then walked through it’ and ‘The baby finished the bottle then broke it’ (both adapted from Chomsky 1995, drawing on Pustejoski 1993). It looks as though when two kinds of things are systematically related –figure and ground, contents and container – we can use the same term to pick out either one. Again, the response is: these phenomena are hard to explain. But as far as one can tell, they can be explained just as well within a traditional framework as within the proposed alternative.

families, existing yet fictional objects and flaws (ii) monkey around with the sentence's logical forms in an effort to show that the apparent ontological commitment disappears under analysis, (iii) adopt a fictionalist approach, according to which utterances of the sentences are not literally true, but they relate in appropriate ways to thoughts that are literally true. I think that (i) offers a plausible account of (7) and (iii) offers a plausible account of (3)-(6). (2), though, requires more extended discussion, as follows.

(2) can appear extremely problematic for a traditional semantic theory because none of (i)-(iii) look promising. The matter deserves some discussion.

Notice first that there is an innocent and relatively unproblematic use of the 'the average American':

(14) The average American is concerned about falling fertility rates.

Here the 'the' is generic 'the', as in 'the tiger likes to hunt at night' and 'average' just means: *typical* (Stanley 2001, Higginbotham 1985). This use of 'the average' appears to be different from 'the average' in (2). Consider (15) a and b:

- (15) a. The average American man's tastes are changing. He used to prefer beer to wine. Now he prefers wine to beer.
- b. *The average American man is concerned about falling fertility rates. He used to have 2.3 children. Now he only has 2.1

One can't combine the numerical constructions together with the *typical* use of 'average'.

The problematic 'the average N' construction appears to appear only in relation to numbers. The most successful account of this construction appears to be that of Stanley (2007).²² According to Stanley's theory, the sentences at issue have the structure: [[NP The average Φ N][VP P]]. [NP The average Φ N] denotes a set of properties, derived as follows. Φ denotes a function from properties to functions from dimensions to measure functions, defined only over the domain of those properties, where a dimension is a way of scaling things along a numbered scale (e.g. height in feet). So Φ denotes a function of type $\langle\langle e, t \rangle\rangle\langle\langle \text{dimension}, \langle e, d \rangle \rangle$. [Φ N], $\langle\langle \text{dimension}, \langle e, d \rangle \rangle$, given a dimension, yields a measure function whose domain is restricted to the extension of N. Given a value of type

²² For an alternative, see Carlson and Pelletier (2002).

<dimension, <e, d>>, 'average' then yields a set of properties, which [The average Φ N] denotes. The set is defined as follows:²³

P is in [The average Φ N] iff (a) where D is some dimension in the domain of [Φ N], P is the property of having the average of [Φ N] (D) along D, or (b) P is the conjunction of Q and R, where Q and R are in [The average Φ N].

One virtue of his account is that it explains the contrast between (16) and (17), since sets of properties are not easily thought of as tall:

(16) The average tall American has 2.3 children

(17) * The tall average American has 2.3 children.

A consequence of the account which is not obviously happy is that if Americans and Canadians are exactly the same in respect of all averages, then (18) comes out true:

(18) The average American is the average Canadian.

Whatever one thinks of (18), there are at least two major problems with the account. First, it predicts that if half of the Americans favourite

²³ We don't bother to avoid use/mention simplification here.

whole integer is 12 and the other half's favourite is 13, then (19) comes out true:

(19) * The average American's favourite whole integer is 12.5.

Secondly, Stanley's account treats 'the' as semantically inert, and more generally makes the construction look utterly unique and bizarre. Why would such an oddity exist?

But a fictionalist approach is highly problematic as well. Suppose that the average French family has 2.3 children, the average German family has 2.3 children etc.. Then, under a fictionalist account, according to which we are pretending that these average families are entities, we ought to be able to say things like (20):

*(20) Some/most/all average European families have 2.3 children.

But we can't.²⁴

Is semantics for natural language doomed, then, by this whacky construction? Of course not. The construction isn't one of natural language. It comes from the science of statistics. I assume 'average' here has to mean

²⁴ The content of most of the preceding discussion of 'the average' is drawn from Stanley's presentation and ensuing discussion at the Arché workshop.

either *median* or *mean*. But these are surely not terms of natural language, but technical terms from mathematics.²⁵ That is precisely why the strange ‘the average’ construction only works with numbers, why it appears so bizarre to the natural-language semantic theorist, and why we natural-language semantic theorists really don’t have to worry about it at all. Phew!

Moving on. I think the traditionalist could do worse than go factionalist about (3) through (6). Utterances of all of them are, strictly speaking, false. But they can be used to get across messages that might be true. I don’t myself believe in non-existent objects. And I am not sanguine about the chances of discovering ontologically innocent logical forms for all of these examples or, indeed, any of them. So I want to suggest that there is nothing wrong with a fictionalist approach to these examples.

We often say things that are literally false in order to communicate something that we take to be true. Sometimes there are Gricean mechanisms that explain how the process works. And I think something roughly like a Gricean account works for (3). When we say (3) (‘Joe Sixpack’s priorities are changing’) isn’t it obvious that we are pretending to refer to a man who is typical of a certain kind of individual? It is only a pretence and so what we say isn’t literally true. (3) is a device we use to convey something we take to be true, something along the lines of: *the priorities of most typical lower-class American men are changing*. I don’t suppose that there is any specific

²⁵ Statisticians also use the expression ‘the median American family’. I suspect that ‘the average American family’ is simply a variant form of that, since the expression is most naturally interpreted as talking about the median. However, a median number of children can’t be 2.3. So perhaps sometimes ‘the average American family’ means: *the mean American family*. I am not sure where the ‘2.3’ example comes from. It might be a linguist’s invention rather than something taken from a real statistical report. Perhaps we should stick with ‘2.5’.

true thought is conveyed by an utterance of (3). But there are a range of very thoughts all of which capture the same gist.

‘Joe Sixpack’ seems to me to have a fairly obvious air of pretence about it. On the other hand, ‘Sherlock Holmes’ in serious uses like (4) and (5) (‘Sherlock Holmes is more famous than any real detective’, ‘Sherlock Holmes is a fictional detective’) and ‘some fictional detectives’ in (6) (‘Some fictional detectives are more famous than any real detective’) don’t. Intuitively (4), (5) and (6) seem to be literally true. I suggest that this is because we intuitively believe that fictional characters exist (real ones, not Joe Sixpack) – not in the natural world but in some realm of fiction. The realm of fiction really exists and so do the characters that occupy it. And we can think about them and refer to them. We believe this intuitively. But on reflection we should give these beliefs up.

So I propose an error theory for (4), (5) and (6). If we say these things, we unintentionally speak falsely. But it doesn’t matter, because there are truths that we are gesturing at. Let us call the idea that the universe contains a realm of fictional characters the ‘meta-fiction’. According the meta-fiction, Sherlock Holmes really exists and really has properties, such as being the creation of Conan Doyle and being very famous. (He fictionally, but not really, has properties such as being a detective and being nearly as smart as Mycroft). Then we can express the truths gestured at by (4), (5) and (6), by (4’), (5’) and (6’) respectively:

(4’) It is meta-fictionally the case that Sherlock Holmes is more famous than any real detective

(5’) It is meta-fictionally the case that Sherlock Holmes is a fictional character

(6') It is meta-fictionally the case that some detectives are more famous than any real detective²⁶

Alternatively, if the meta-fiction ploy doesn't work, we have options such as (4''), (5'') and (6''):

(4'') More people understand 'Sherlock Holmes' than possess any mental representation of a real detective.

(5'') There is a fiction in which Sherlock Holmes is a detective.

(6'') More people understand some fictional names of detectives than grasp possess any mental representation of a real detective.

Of course I am not suggesting that Jane Winecooler thinks the thoughts expressed by (4''), (5'') and (6''). But the sentences express truths that capture the gist of the idea that she tries to get across by acting as if Sherlock Holmes and the like are objects of reference. Suppose Jane utters (4). I might respond by saying: "I am sorry, I cannot accept that. Sherlock Holmes doesn't exist. Therefore he can't have any properties at all, including fame. Would it be ok if I agreed to the proposition that more people understand 'Sherlock Holmes' than possess any mental representation of a real detective?" Jane would say "ok".

When it comes to (7) ('There is a flaw in the argument') I would bite the bullet and claim that flaws really do exist. Chomsky suggests that not even Joe Bloggs treats flaws as real-world objects, offering (21)-(25):

²⁶ My suggestion is that the 'meta-fiction' operator will help us capture a certain range of serious uses of apparent reference to or quantification over fictional and mythical characters. It won't apply to all such uses. Consider: 'Sherlock Holmes exists in the meta-fiction'. This isn't well rendered by 'It is meta-fictionally the case that Sherlock Holmes exists in the meta-fiction'. Thanks to Stacie Friend for this point.

- (21) There is a fly in the bottle.
- (22) There is believed to be a fly in the bottle.
- (23) There is believed to be a flaw in the argument.
- (24) There is a fly believed to be in the bottle.
- (25) There is a flaw believed to be in the argument.

(21)-(24) are acceptable, but (25) is deviant. Constructions like (24) and (25) 'have existential import in some manner beyond' (7), (21), (22) and (23) (Chomsky 2003, 293).

I suggest that (7), (21), (22) and (23) have existential import and that this is harmless. Flaws are just as real as flies. The deviance of (25) has to do with the relation between the flaw and the argument. The fly can be taken out of the bottle and still exist. The flaw cannot be taken out of the argument and still exist. That doesn't mean the flaw in the argument doesn't exist. It just means that its existence is conceptually connected to the argument. And this connection is reflected in our feelings about (25), which somehow suggests that the flaw is only contingently related to the argument. Compare (26):

- (26) There is a hole believed to be in the bucket.

(26) appears deviant, unless one conceptualizes the hole as somehow independent of the bucket. One can do this. There is a large shelf, high up above us. On it there is a bottle, a barrel and a bucket. There is liquid on the floor, below the shelf. There is a hole somewhere. Where is it? Well, it is

believed to be in the bucket. So, indeed, there is a hole believed to be in the bucket.²⁷

3. Conclusion

I conclude that the traditional semantic theorist who takes a realist attitude to the manifest image can respond appropriately to each challenge. In all the cases bar one, he can meet the challenge head on and say something sane and reasonable about how he might account for the relevant data. The only outstanding cases are those of reference-switching, as in ‘book’, ‘bottle’, ‘door’ and so on. But, on the face of it, those phenomena would appear to be equally problematic for everyone. Now, if I just had an account of sorites and an account of conditionals²⁸

²⁷ Higginbotham (1985) suggests that maybe ‘there is a flaw’ could be analyzed as a kind of predicate, rather than a nominal: *the argument is flawed*. Chomsky responds with examples like ‘We fixed three of the flaws you found but the rest of them resisted our efforts’. Ludlow (2003) rejoins with the suggestion of ‘We fixed three of the flawed steps you found but the rest of them resisted our efforts’. But ‘flawed step’ would bother Chomsky just as much as ‘flaw’. A flawed step is no likelier a real-world object than a flaw. Moreover, we have the usual contrast: ‘there is believed to be a flawed step in the argument’ versus ‘there is a flawed step believed to be in the argument’.

²⁸ Many thanks for comments or discussion to: Berit Brogaard, Mat Carmody, Noam Chomsky, Stacie Friend, Mike Gabbay, Jen Hornsby, Ruth Kempson, Shalom Lappin, Richard Larson, Guy Longworth, Peter Ludlow, Harold Noonan, David Papineau, Paul Pietroski, Daniel Rothschild, Richard Samuels, Jason Stanley and Charles Travis.

BIBLIOGRAPHY

- Antony L., and Hornstein, N., (2003) *Chomsky and his Critics*. Blackwell: Oxford.
- Bering, J. M & Bjorklund, D. F (2004) "The Natural Emergence of Reasoning About the Afterlife as a Developmental Regularity" *Developmental Psychology* 40(2), 217-233
- Boyer, P., (2003) "Religious thought and behaviour as by-products of brain function" *Trends in Cognitive Science* 7 3
- Carlson, G., and Pelleteier, F (2002) 'The Average American has 2.3. Children' *Journal of semantics* 19: 73-104.
- Chomsky, N., (1981) *Lectures on Government and Binding*, Foris: Dordrecht.
- (1995) 'Language and Nature', *Mind* 104: 1-61.
- (2003) Reply to Ludlow in Antony and Hornstein (eds) 287-295.
- Higginbotham, J., (1985) 'On Semantics' *Linguistic Inquiry* 16 547-593.
- Hornstein, N., (1984), *Logic as Grammar*, MIT Press, Cambridge MA.
- Kuhlmeier, V. A., Wynn, K. & Bloom, P, (2004) Bloom, "Do 5-Month-Old Infants See Humans as Material Objects?" *Cognition*, 94, 95-103. 'The) E. LePore, E., and B. Smith (eds) (2006) *Oxford Handbook of Philosophy of Language* Oxford University Press: Oxford.
- Llinas, R. (1987) 'Mindness as a functional state of the brain' in Blakemore, C., and Greenfield, S., (eds) *Mindwaves: Thoughts on Intelligence, Identity and Consciousness*. Blackwell: Oxford.
- Ludlow, P., (1989) Implicit Comparison Classes' *Linguistics and Philosophy*, 12: 519-533.
- (2003) 'Referential Semantics for I-Languages' in Antony and Hornstein eds 140-162.
- Malt, B., (1994) 'Water is not H₂O', *Cognitive Psychology* 27, 41-70.
- Noonan, H, (1998) 'Animalism versus Lockeanism A Current Controversy' *The Philosophical Quarterly* 48 302-318.
- Pietroski, P., (forthcoming) Character Before Content' in *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, edited by J. Thomson and A. Byrne, OUP).
- Pustejoski, J., 1993 'Type Coercion and Lexical Selection' in Pustejoski ed. *Semantics and the Lexicon* Kluwer: Dordrecht.
- Rothschild, D., and Segal, G, (in preparation), 'Indexical Predicates'.

Segal, G. (2006) 'Truth and Meaning' in LePore and Smith eds.
Segal G., (forthcoming) 'Poverty of Stimulus Arguments Concerning Language and Folk Psychology' in P. Carruthers and Laurence, S. (eds) *The Innate Mind: Foundations and Future*. Oxford: Oxford University Press.
Sellars, W., (1962) 'Philosophy and the scientific image of man' in *Frontiers of Science and Philosophy*, Robert Colodny, R. (ed.) University of Pittsburgh Press: Pittsburgh 35-78
Stanley, J. (2001) 'Hermeneutic Fictionalism', in *Midwest Studies XXV: Figurative Language*, French and Wettstein (ed.) (Blackwell, 2001): 36-71
(2007) 'The Average American', paper presented to the Arché workshop on adjectives, St Andrew's May 2007.
Szabo, Z., (2001) 'Adjectives in Context', in Kenesei I. and Harnish, R., (eds) *Perspectives on Semantics, Pragmatics and Discourse. John Benjamin*: Amsterdam 119-146
- (2006) 'The Distinction Between Semantics and Pragmatics.' in LePore and Smith (2006)
Williams, B., (1970) 'The Self and the Future', *The Philosophical Review* 79: 161-180.