



Atomism, Concepts, and Polysemy

Kamil Lemanek¹

Received: 30 August 2021 / Accepted: 9 November 2021 /

Published online: 26 November 2021

© The Author(s) 2021

Abstract

The aim of this paper is to examine the theoretical architecture of semantic atomism and its consequences with respect to natural language. In particular, it looks to explore the notion of possible concepts using the fundamental distinction between simple and complex concepts and expressions in Jerry Fodor's atomism. The distinction is exploited to produce an unusual type of concept referred to as a correlate, which effectively mirrors complex concepts while maintaining a distinct underlying structure. Though harmless in and of themselves, their presence in the context of polymorphemic expressions suggests that atomism harbors a tacit and unintuitive form of polysemy that is problematic in its own right and that leads to other complications, some of which may be demonstrated on the example of communication. These issues are tied to the way atomism is structured, and although they seem to have gone largely unnoticed, they appear to bear negatively on the adequacy of atomism where natural language is concerned.

Keywords Theory · Correlate · Meaning · Communication · Possibility · Fodor

Semantic atomism is a striking position; it has lodged convincing and influential arguments against its rival theories while forwarding a promising picture of its own—one that operates on primitives, compositionality, and reference. When put together, it presents a uniquely reliable and stable characterization of meaning. That perceived stability along with its cognitive trappings make for a compelling case, but the position suffers from subtle complications when the structural distinctions unique to atomism are taken seriously and confronted with natural language, as they lead us to an idiosyncratic and undesirable form of polysemy. In order to demonstrate the problem, I sketch a formulation of atomism presented by its most explicit and successful advocate, Jerry Fodor (see, e.g., 1998; 2008; Fodor & Lepore 1992, 2002). Although the analysis below focuses on his atomism in particular, it

✉ Kamil Lemanek
K.Lemanek@uw.edu.pl

¹ Faculty of Philosophy, University of Warsaw, Warsaw, Poland

communicates issues that may be generalized to any position with similar theoretical architecture and commitments.

I begin by outlining and situating semantic atomism in section 1. In section 2, I discuss possibility with respect to concepts and language and present the idea of a correlate concept. I problematize these ideas in section 3, illustrating the abovementioned polysemy and what makes it undesirable, along with a brief discussion of how it affects communication as an example of the consequences it can have for our understanding of natural language in general.

1 Section 1 – Semantic Atomism

The best way to get a grip on semantic atomism is to get a better idea of the landscape it features in. We can start with a simple argument devised by Devitt (1996, 10) to set the stage. It goes as follows:

- (i) Some of a token's inferential properties constitute its meaning.
- (ii) If some of a token's inferential properties constitute its meaning, then they all do.
- (iii) So, all of a token's inferential properties constitute its meaning.

If one accepts the premises through to the conclusion, you get a radical form of holism. If you reject (ii) you get molecularism, and if you reject (i) you get a form of atomism, which is to say that the meaning of a token isn't constituted by *any* of its inferential properties (see Fodor 2008, 143–144; cf. Margolis 1999). Instead of relying on inferential properties, atomism roughly equates meanings with *concepts*, so that the meaning of a token is whatever concept it corresponds to (Fodor 1998, 2), which of course raises the question of what a concept is. Fodor's account has undergone changes over the years, but the core idea is that a concept is a mental particular composed of a content together with a mode of presentation. Fodor accepts a full-blooded referentialism with respect to content (1998, ch. 1; 2008, 16–18; 2015, 1), according to which the content of a given concept is whatever it refers to. Modes of presentation are more complicated; they are what we might consider the forms of mental representations—the *vehicles* of representation or the *way* we represent something (Fodor 1998, 16). They are distinct facets of concepts that are “*ipso facto* available to be proximal causes of mental processes” (Fodor 1998, 19). As such, they are taken to be distinguished functionally with respect to the mental processes they cause. In the context of Fodor's computational theory of mind, modes of presentation effectively communicate the syntactic aspects of concepts in addition to the semantic character of representation. Content and mode of presentation comprise concepts. That being said, there is also a structural distinction built into the latter that plays a uniquely important role in characterizing them—the distinction between simple concepts and complex concepts.

It's a distinction that cuts across atomism at a theoretical level and looms large in the architecture of the position. *Simple* concepts are atomic, which is to say that they have no internal structure, and they're accordingly primitive. *Complex* concepts are

composites, formed of simple concepts which serve as their constituents. In effect, simple concepts are the building blocks, and complex concepts are what you make with them. The idea really comes together in the context of language.

To ease into it, we'll start with Mentalese—Fodor's notion of a language of thought. Simplifying, Mentalese is a lot like our familiar natural languages. It has a set of simple concepts analogous to words, and it can combine those simple concepts into structures analogous to phrases and sentences. The combinations themselves are complex concepts. The relevant difference is that Mentalese doesn't have a surface form: it doesn't have spelling, pronunciation, or any explicit syntactic structure—Mentalese is thought, not spelled out or said. So, if you've got the concept WATER, you can produce a mental token whose meaning is WATER. And if you also have the concept COLD, you can combine it with WATER to produce a token whose meaning is COLD WATER. COLD and WATER are taken to be two simple concepts, and COLD WATER is taken to be a single complex concept.

Where simple concepts are concerned, COLD and WATER are both considered structurally atomic. In fact, Fodor considers SPIGOT, BROWN, and COW atomic too. The idea is that most of the lexical concepts we have are atomic. And insofar as they're atomic, they're also primitive; in the words of an objection Fodor entertains and actually commits to, "if most lexical concepts have no internal structure, then most lexical concepts must be primitive" (Fodor 1998, 123). They can't be broken down into their parts because they have no parts, and because they have no parts, they can't be defined in a traditional sense, where the *definiens* constitutes the *definiendum*.¹ Hence, they're basic. Fodor would ultimately furnish the claim with the narrower and more precise assumption that "concepts expressed by *monomorphemic* names of kinds are basic, including not just natural kinds (STAR, WATER, LEVER) but also everyday kinds of kinds like GRASS and CARBURETOR" (Fodor 2008, 67; emphasis added). While identifying simple concepts relative to monomorphemic expressions obviously ties into natural language, which will prove important later, it also gives us a clear picture of just how many atomic and primitive (and so simple) concepts Mentalese operates with—which is perhaps the most controversial facet of Fodor's position; as unusual as it may sound to suggest that CARBURETOR is a simple concept, Fodor holds that "all the evidence—philosophical, psychological, and linguistic" suggests that these basic lexical concepts are just that: unstructured primitives (1998, 147).

Complex concepts are more intuitive. COLD WATER is a single complex concept. Its content is exhausted by the contribution of COLD and WATER, which effectively define the complex concept as its constituents. The relevant theoretical feature here is compositionality, one of the principles at the heart of atomism. In its simplest form, it requires that "the content of a thought is entirely determined by its structure together with the content of its constituent concepts" (Fodor 2008, 17). And compositionality captures productivity, which can be tied together as follows:

¹ On the notion of definition, see Fodor (1998, 48–49).

[Per productivity,] there are infinitely many concepts that a person can entertain. (*Mutatis mutandis* in the case of natural languages: there are infinitely many expressions of *L* that an *L*-speaker can understand.) Since people's representational capacities are surely finite, this infinity of concepts must itself be finitely representable. In the present case, the demand for finite representation is met if (as, as far as anyone knows, only if) all concepts are individuated by their syntax [read: mode of presentation] and their contents, and the syntax and contents of each complex concept is [per compositionality] finitely reducible to the syntax and contents of its (primitive) constituents. (Fodor 1998, 94-95)

And that's where one of the essential strengths of Fodor's position lies. We effectively assume a large base of simple concepts that can then be used to transparently construct a theoretically infinite set of complex concepts (cf. Pullum & Scholz 2010). That isn't a novel idea, but its liberal use of referentially situated primitivity is. By taking simple concepts to be atomic and therefore primitive, Fodor's semantics provide an exceptionally solid foundation for supporting complex concepts. The resilience of the position is further strengthened by the way concept individuation and concept possession are characterized in semantic atomism—that is, how concepts are distinguished from one another and what it takes for someone to actually possess a simple concept.

As noted in the passage above and built into the way concepts are comprised, a concept is individuated by its mode of presentation and its content. Just how individuation is understood in practice and how it relates to concept possession is perhaps best observed on the basis of some familiar and traditionally problematic cases. Differentiating on the basis of content isn't particularly hard to follow. If two concepts don't share the same reference, then they don't have the same content. And if they don't have the same content, then they aren't the same concepts—they're straightforwardly individuated on the basis of that difference in content. But what about co-extensive concepts, like (1) WATER and H₂O or (2) CICERO and TULLY? Both sets seem to share the same content (where water is H₂O and Cicero is Tully); what differentiates them has to do with their mode of presentation from the standpoint of possession, including its syntactic implications.

- (1) The difference between WATER and H₂O is a matter of complexity and constituency. Notice that where WATER is a simple concept, H₂O is a complex concept consisting of HYDROGEN, TWO, and OXYGEN. So, where someone might have the concept WATER, it doesn't necessarily follow that they need to have the concept HYDROGEN (Fodor 1998, 73-75; 2008, 66). Conversely, if someone has the concept H₂O, they necessarily have the concept HYDROGEN, simply because hydrogen is a constituent of H₂O at the level of presentation. You can't have H₂O if you haven't got H. Despite sharing the same content, the two are distinguished on these structural lines built into their modes of presentation.
- (2) Both CICERO and TULLY are simple concepts; neither of them admits of any relevant constituents at a conceptual level. However, the two representations are effectively type distinct. Where referentialism about content indeed compels

one to see them as semantically identical, one can nonetheless see them as being syntactically distinct (Fodor 2008, 67-71). Given a computational theory of mind, CICERO-type representations have different causal connections than TULLY-type representations, something that can be shown as trivially as that someone with the concept CICERO might believe *Cicero is tall*; something that is made possible by the syntax of the CICERO-type, but which is not made possible by the TULLY-type. This is a roundabout way of saying that having CICERO lets you construct complex concepts that include CICERO, something that having TULLY doesn't afford—though it does afford constructing complex concepts that contain TULLY, something that having CICERO doesn't afford. Again, it's their modes of presentation that make the difference, though in a different way than they did in the preceding case.

Each set illustrates how modes of presentation can play a role here. Cases resembling the first will be of particular interest going forward, as they expressly involve both simple and complex concepts in addition to shared content.

This picture of individuating concepts through content and mode of presentation notably refrains from appealing to or drawing from our knowledge of a given concept or its object—which is to say that atomistic individuation doesn't take recourse to our epistemic situation in a traditional sense. It's a hallmark of the position as a whole, one that's reflected in Fodor's idea of concept possession (and acquisition) too, which is understood in causal terms and may be construed as our quite literally *having* a concept. Roughly, having a concept is a matter of possessing neural-mental structures that are appropriately related to and whose tokenings are (or would be) reliably caused by the instantiation of a given thing. Illustrating that in terms of acquisition, one picks up the concept WATER by being exposed to water, and in turn developing an informational sensitivity to it. The requisite exposure doesn't have to be local; you don't have to have physically encountered a zebra to have a corresponding concept. In fact, Fodor notes a number of ways one might acquire a concept non-locally, including gossip, stories, and theoretical inference (1998, 77-78). Hence one might have the concept of a zebra from hearing about them, or unicorns from stories, and protons from a theory that allows one to find them via theoretical inference. The range of available concepts is interesting in its own right, evidently including natural kinds, imagined creatures, and theoretical posits—as well as a vast collection of so-called mind dependent concepts, like DOORKNOB and SPIGOT. However, we will set the subtle nuances of Fodor's conceptual taxonomy aside. What matters for our purposes is that all of these concepts are captured in nomologically grounded mental representations whose tokenings are caused by something or other, whether locally through interaction and perception or distally through gossip or theoretical inference. In epistemic terms, that basic form of causal exposure (e.g., through physical contact) and sensitivity doesn't require us to know how a given concept relates to other concepts in order to acquire it, nor does it require us to be explicitly aware that we've come to possess it. You can have a concept without knowing anything about it, or even without knowing you have it.

The way concepts are individuated and the sense in which they're possessed within the atomistic framework ensures intersubjective and broadly accessible

concepts. Hence, Homer might have had the same concept WATER that we do, even if we happen to know that water is H_2O while he has no idea what hydrogen is. The reference of his WATER and our WATER could still be the same (i.e., that stuff we drink and that falls from the sky—things we have exposure to), and so our concepts may share the same content. While it is a nomic necessity that water is H_2O , whether or not you *know* about it is irrelevant to whether or not you *have* the concept. The same can be said for the mode of presentation, in that it isn't shaped by what you know about the concept. It has its mode of presentation in virtue of its syntactic and structural properties as a concept in and of itself. Accordingly, Homer can have the same concepts that we do despite being in a radically different epistemic situation. Of course, it goes without saying that what you know about water has important practical consequences, but as far as atomism is concerned, none of that knowledge is required or even relevant to possessing a given concept.

This epistemic insensitivity renders simple concepts remarkably invariant, remaining the same across time and varying contexts. Our invariant simple concepts can then be put together to create complex concepts, which in turn exhibit the same invariance and intersubjectivity. The fairly vague foundational stability offered by a theory in which simple concepts are unstructured and primitive is bolstered by the more tangible and practically applicable idea that they also provide us with concepts that, at least in theory, retain their shape over time and remain accessible (through a fairly broad assortment of means) to people with different backgrounds. Between the characterization of simple concepts, their combination in complex concepts, and the ways they can be acquired, we've arrived at a basic picture of semantic atomism.

We've drawn closer to natural language with the transition from concepts and representations to Mentalese and lexical items together with a sketch of concept individuation and possession. Now we just need to take the final step. Where Mentalese lacks a surface form, natural language requires one. This calls for a small but significant alteration, a move from talk of concepts and their tokens to talk of concepts and the expressions they are associated with. The idea is basically that concepts are regimented under the words of our language, so an English language expression like "water" means WATER, and "cold water" means COLD WATER.² The significance of the difference is that we now have to deal with a layer of surface forms, with an extant lexicon, its history, and its idiosyncrasies—the details and consequences of which we'll see in the next two sections.

However, before we conclude this sketch, it needs to be said that Fodor was ambivalent with regard to natural language and semantic atomism. While he often used English as a model for discussing and developing Mentalese, he would also equally often note the limitations between the comparison. For instance, he accepted that compositionality can't account for idioms (2008, 131), whose explicit constituents are often overridden. He also suggested that English might fail to be compositional altogether (2001; 1998; 2008), and that perhaps English has no semantics at all, inasmuch as the semantics of thought come first (Fodor 2008, 198–199). In

² It is what it looks like—as Fodor puts it, "what I'm selling is a *disquotational* lexicon" (Fodor 1998, 55).

his most recent works, he and Pylyshyn would go as far as to hold that meaning in the traditional sense is “a myth” (2015, 58). Despite all of that, Fodor’s work influenced and often made claims directly concerning natural language and meaning, particularly when criticizing competing semantic positions (Fodor & Lepore 1992, 2002, 2007) and when addressing expressly linguistic phenomena (Fodor 1998, ch. 3; Fodor & Lepore 2002, ch. 5). With that in mind, the question is whether natural languages should be thought of as atomistic, or even within the purview of atomism; while I’m perfectly happy to agree that it accounts quite well for a language like Mentalese, it isn’t as clear that it can do the same for a natural language like English, as I’ll endeavor to illustrate.

2 Section 2 – Possibility

In order to understand how natural language may be problematic, we need to have a better idea of the range of concepts we have to work with over and above our simple criterion relating to monomorphemic expressions; more fundamentally, it isn’t enough to know what concepts are and how they are acquired or individuated, we need to have a better idea of what concepts there are for us to acquire and individuate. They can be neatly defined as just those mental particulars that are theoretically well-formed with respect to atomism (as a framework) and the way our world is (as a matter of fact). Anything that satisfies those two conditions is a concept we could have: a *possible* concept. The set of possible concepts represents every concept there is to acquire and individuate, and it’s just that set that I’d like to explore a bit in order to give us an idea of how large and diverse it might be; sketching possibility in this context and how we can approach it will prove relevant as we move towards problematizing Fodor’s position with exotic concepts. To that end, the best place to start is with concepts we actually seem to have, as possession is *sufficient* for possibility here. That being said, it isn’t *necessary* for possibility, and so it needs to be understood that the concepts we actually have aren’t all there is to the set of possible concepts as such. While that’s very nearly a truism, it’s abstract enough to warrant some explanation.

I can illustrate what I mean using an example. Suppose that there are concepts that only trained carpenters happen to have, like TENON or DADO. There being carpenters that have these concepts requires that the concepts themselves be possible: if the concepts in question were theoretically defective or incoherent, then these carpenters wouldn’t have been able to acquire them in the first place, which is to say that possession is sufficient for possibility. Now suppose all the carpenters in the world suddenly vanish; we’re left without any trained carpenters, and so we’re left without any people that have TENON. Although nobody has the concept, it stands to reason that the concept itself is still possible, in that it isn’t as if the concept is somehow rendered defective because nobody has it. The same can be said for any concept when it’s first acquired; the first carpenter to fashion a tenon must have acquired a concept that nobody else had, if for no other reason than because there were no tenons before they made one. So while we cannot have possession without possibility, we can have possibility without possession (hence, sufficient but not necessary).

It follows that the set of possible concepts isn't limited in any relevant sense to the concepts that we or anyone else actually has, although the concepts we have are an intuitive and worthwhile place to start.

With that in mind, the set of possible concepts features simple concepts with varying contents like WATER, COW, and TENON along with complex concepts like COLD WATER OF BROWN COWS ARE DANGEROUS. Beyond that, there are countless possible concepts for which we have no words, all of which are possible in virtue of their consistency with respect to the framework of atomism and the way the world is. Within that vast field, there is a particularly unusual type of concept unique to atomism that hinges upon that fundamental distinction between simple and complex concepts, and one that I believe to be problematic when considered in the context of natural language. In one of the appendices that close his *Concepts* (1998), Fodor briefly outlines an example of the idea:

[Atomism] allows there to be a *simple* concept of [the property of *being a red square*]; viz. a *primitive* mental representation RED SQUARE (*sic*; this *is* intended to be a structural description) that is locked to *being red and square*. Presumably, one could acquire RED SQUARE ostensively. That is, one could get locked to *being red and square* (not by first getting locked to *being red* and *being square*, but) by learning that redsquares (*sic*) are the things that look like *those*. So Informational Atomism acknowledges the metaphysical possibility of having the concept of a red square without having either the concept RED or the concept SQUARE. (Fodor 1998, 165)

This is a significant observation that has received little attention or development over the years. It effectively substantiates two distinct concepts distinguished on structural grounds: RED SQUARE and RED SQUARE. The former is complex and the latter is simple, with both referring to the same thing in the world. Whereas RED SQUARE is undoubtedly instantiated and in circulation, Fodor's own tone and explanation implies that RED SQUARE isn't. Either way, RED SQUARE appears to qualify as at least possible.

This unusual simple concept seems to essentially be the product of an inference from a structurally complex concept to a structurally simple counterpart. Fodor invents a plausible mode of acquisition to go along with it: ostension. But the mode of acquisition is secondary to that initial inference; it's just a bit of storytelling to help us imagine that there could be a some such concept (and perhaps what its content might be). It's something that we can dispense with for now, given that possession isn't necessary for the possibility of concepts to begin with. Instead, we can just concentrate on the initial inference, which may itself seem a bit unintuitive in this fairly theoretical context. To make it a bit clearer and to better communicate its generality, we can think of it in terms of an analogy. Imagine that we have a vibrant picture of the Eiffel Tower; capturing the verdant green of the surrounding lanes, the subdued brown of its structure, and the blue of the sky ranging over it. Whatever colors comprise the picture, we can pretty easily imagine what it might look like if it were monochrome (e.g., sepia or greyscale). Although the two would differ in color, they'd still be from the same perspective, feature the same measurements, levels, etc. More importantly, they'd share the same content in that they'd both be

pictures *of* the same thing, in the sense that the full color picture of the Eiffel Tower and the monochrome picture of the Eiffel Tower are both still pictures of the Eiffel Tower. On the basis of that train of thought, we can fairly reasonably say that we can infer from almost any given full color picture to the possibility of there being a monochrome version of it. If you think of the colors of the photo as analogous to the parts constituting a complex concept, inferring to a monochrome version of a color photo is like inferring to the structurally simple counterpart of a complex—in both cases inferring to some distinct possibility while preserving content, other representational features, etc. And from color to monochrome and complex to simple, the inferences point us to their corresponding possibilities.

The simplicity of the inference allows it to be applied to practically any complex concept. COLD WATER, RED SQUARE, and BROWN COW can all be inferred to have structurally simple counterparts that mirror them; I'll call these simple concepts *correlates*. Complexes and their correlates share the same referent; in other words, they have the same content.³ Along with the same content, we can also assume they share congruent functional roles, representational features, etc. What distinguishes them is their structure: one is complex and the other is simple. The consequence of that difference in structure can be felt in their collateral commitments. Complex concepts have these commitments and simple concepts don't. That is, we need to have both RED and SQUARE to have the complex RED SQUARE. To have the correlate RED SQUARE, we don't need either. The same goes for COLD WATER and COLD WATER, and any other complex and correlate pair.

It might seem like what I've described is just a quirky and harmless consequence of atomism's theoretical architecture, in that it follows straightforwardly from its reliance on complex and simple structures. However, correlates have significant consequences insofar as we're interested in natural language; the link between complexes and correlates leads to difficulties when they're paired with expressions. The introduction of an additional layer of representation, communication, and possession causes issues that all stem from the otherwise innocuous presence of correlates, but as we'll see, they're only really noticeable in the context of a particular type of expression.

3 Section 3 – Complication

I take it as given that whatever the status of a correlate is, a complex expression of a natural language can only be associated with a complex concept. That is, “red square” corresponds to the complex RED SQUARE and not its correlate RED SQUARE. The explicitly discrete surface structure of the expression only plausibly corresponds to the complex concept, not the correlate—a point Fodor makes in passing that I'm fairly happy to grant (2008, 66; 1998, 42). So RED SQUARE could very well be more than merely possible (perhaps somebody out there actually has

³ This is particularly relevant for Fodor's handling of impossible concepts like ROUND SQUARE (2008, 164). Though impossible concepts are interesting enough, they aren't generally relevant here.

it), but it still wouldn't be something that concerns natural language as long as there isn't a simple expression associated with it.

We've already established that simple expressions like "water" and "spigot" correspond to simple concepts, and so they don't have correlates. And I just conceded that complex expressions like "red square" correspond to complex concepts. Simple expressions correspond to simple concepts, and complex expressions correspond to complex concepts. That might seem to close the book as far as correlates are concerned, but that's not quite right. Recall that Fodor holds that *monomorphemic* lexical items correspond to simple concepts. Of course, not every simple expression is monomorphemic. Consider words composed of multiple free morphemes, expressions I'll refer to as *polymorphemes*: examples include "doorknob," "breakfast," and "laptop." They appear to present the problematic possibility of simple expressions that corresponds to complex concepts.

The intuitive reasoning being that, despite being simple expressions, they contain elements that normally correspond to simple concepts (e.g., door, knob, fast, etc.) without the discrete surface structure that breaks up complex expressions, hence the intuition that they could plausibly correspond to complex concepts. That being said, the case for polymorphemes corresponding to complex concepts can be put more technically and forcefully in terms of Fodor's initial commitment concerning monomorphemes. If monomorphemes contrast with polymorphemes and monomorphemes are singled out as corresponding to simple concepts, then at least some polymorphemes must correspond to complex concepts—that is, if polymorphemes also corresponded to simple concepts through and through, then there wouldn't be any reason to single out monomorphemes to begin with (i.e., it would be a pointless distinction). As such, some simple expressions (e.g., monomorphemes) only correspond to simple concepts and others (e.g., polymorphemes) may correspond to complex concepts. Unfortunately, we have no clear way of telling which polymorphemes correspond to complex concepts and which correspond to simple concepts, nor do we have any reason to think that they couldn't, for instance, do both. In other words, although we have a means for differentiating complex concepts and simple concepts as such, and we can do a bit of work with that (and we will shortly), we don't have a means for antecedently discerning whether a given polymorpheme corresponds to one or the other (or both). Of course, we have to start somewhere, so I'll take it for granted that polymorphemes generally correspond to complex concepts. As indicated above, at least some of them certainly do, so it's something that atomism has to contend with one way or another, and it's an assumption that will allow us to more easily examine the theoretical features surrounding those cases. At any rate, we won't stray very far from the idea of polymorphemes corresponding to simple concepts—it's the connection between the two that will ultimately prove problematic.

With that in mind, we can assume that a polymorpheme like “doorknob” expresses the concept DOOR-KNOB, which is composed of DOOR and KNOB.⁴ Insofar as it’s complex, there is also a correlate: DOORKNOB. Seeing as the primary difference between DOOR-KNOB and DOORKNOB is their structure, and therefore their possession conditions, we can use the constituent concepts of the complex concept to differentiate between them, just like we did with WATER and H₂O. So if someone has a concept that denotes doorknobs without also having both the concept DOOR and KNOB, they must have the correlate. One way to see how that might actually be the case is to consider a child acquiring a language. It isn’t difficult to imagine that they learn about doors and doorknobs before they ever acquire the more abstract notion of a knob. Or putting it somewhat differently, we can go around showing a learner doorknobs so that they’re exposed to the causal source of the concept, while at the same time helping them discern doorknobs without also helping them pick up or isolate the concept KNOB (just like in the case of red squares and REDSQUARE). However we frame it, it strikes me as quite reasonable to think that there are plenty of children that have the correlate concept and not the complex. And I think it’s reasonable to think likewise for plenty of other polymorphemes too, like “breakfast” and “laptop,” which are plausibly complex but which are likely acquired as simple concepts first. If the story holds, then correlates aren’t just one of the many concepts we could have, they’re quite plausibly concepts some of us do have.⁵

Going further, these kids eventually grow up, and they undoubtedly pick up the concept KNOB along the way. Assuming they have both DOOR and KNOB, the complex concept DOOR-KNOB is available to them. But that doesn’t mean that they lose the correlate they originally acquired, which is to say that they may ultimately have both DOOR-KNOB and DOORKNOB.

Unlike in the case of “red square,” where the expression is only plausibly associated with its complex, it’s quite clear that “doorknob” may be associated with either DOOR-KNOB or DOORKNOB, which is to say that one and the same polymorpheme can express either the complex or the correlate. So, for instance, if a child says “doorknob” without ever having acquired the correlate, it must correspond to the complex. And if a child says “doorknob” before acquiring KNOB, it must correspond to the correlate. Moreover, when the child with the correlate eventually gains access to both the complex and the correlate, their utterance may very well correspond to one or the other and may be distinguished with a counterfactual, very much in Fodor’s

⁴ Note, Fodor takes “doorknob” to correspond to a simple concept in his *Concepts* (1998) and uses it as the basis of his doorknob/DOORKNOB problem. I’ve opted to use it here because it serves as an intuitive example, and as Fodor himself notes, “actually, of course, DOORKNOB isn’t a very good example [for the doorknob/DOORKNOB problem], since it’s plausibly a compound composed of the constituent concepts DOOR and KNOB. But let’s ignore that for the sake of the discussion” (122). This is a different discussion, so there is no need to ignore it.

⁵ The same case could be made for adults learning a second language in terms of both acquiring concepts as complexes and as correlates; although that sort of example might ultimately be more relatable and precise, it also requires a fair number of additional assumptions concerning the native language and its conceptual commitments. I haven’t developed that here for the sake of brevity and keeping the basic point as simple as possible. Suffice it to say that these considerations don’t need to involve or rely on the ever-controversial notion first-language acquisition.

style. Putting that all together, we arrive at a plausible example of an expression that may pair with two distinct concepts distinguished by their structures, where the expression is simple and the concepts may be more than merely possible—that is, concepts some of us may actually have.

3.1 Polysemy

Given that concepts serve as meanings, this basic picture of one linguistic form paired with two distinct concepts draws us into questions of ambiguity with respect to the place of polymorphemes in natural language. Where familiar examples of simple and complex concepts with shared content (e.g., WATER and H₂O) tend to come with distinct surface forms (e.g., “water” and “H₂O”), polymorphemes stand apart precisely because they apparently allow us to plausibly regiment complexes and correlates under a single surface form. Given that a polymorpheme may correspond to one or the other concept, it seems to present us with some form of homonymy or polysemy. The difference between them is subtle and somewhat controversial (see, e.g., Tuggy 1993; Taylor 2003; Sennet 2016); as Cuyckens and Zawada (2001) put it, homonymy “refers to the semantic phenomenon where the same linguistic form refers to two separate and unrelated words, each with their own meaning” (xiii) and polysemy “in its simplest form ... refers to the phenomenon in language where one linguistic form has a number of different, yet related meanings” (xi). Putting it schematically and in the atomistic idiom, both homonymy and polysemy are based on a single linguistic form or expression being associated with a number of distinct concepts. The difference between the two comes down to whether the concepts associated with the expression in question are related in some theoretically relevant or salient sense. A basic reconstruction of each will be sufficient for our purposes.

Where standard homonymy is concerned, “bank” is a good example; in atomistic terms, there are two concepts associated with it that are differentiated on the basis of their content—one refers to financial institutions and the other refers to riversides. The two concepts don’t appear to support any relevant relation, etymological, conceptual, or otherwise. Roughly put, the lack of any connection between the two is what renders them so distinct, and so it’s what makes the expression homonymous rather than polysemous. Where standard polysemy is concerned, uncontroversial examples are a bit harder to come by. One is “mink,” and again putting it into the atomistic idiom, there are two concepts associated with it that may be differentiated on the basis of their content—one refers to an animal and the other refers to its fur or coat. However, unlike in the case of “bank,” these two concepts appear to be fundamentally related in that the expression’s use with respect to fur is derivative from or functionally dependent upon its use with respect to the animal—or however your favorite theory might choose to put it. There being some such relation is what renders the expression polysemous rather than merely homonymous.

Turning to polymorphemes, though schematically similar to both examples above, they stand apart here too, presenting us with unusual cases where the concepts in question share the same content, where both the complex and correlate associated with a given polymorpheme refer to one and the same thing; that is, both

the complex and the correlate associated with “doorknob” refer to doorknobs. Obviously, it is rather unusual to put so much emphasis on concepts and referents when considering homonymy and polysemy, but in the context of atomism it’s a natural way of approaching them, one that’s reflected in Fodor’s own work. And in that context, where content plays such an important role, two concepts sharing the same content looks like just the kind of relation that would qualify the expression they correspond to as polysemous rather than merely homonymous—that is, just the kind of relation that would qualify polymorphemes like “doorknob” as polysemous. To better see why and to substantiate this unusual notion of polysemic polymorphemes, we need to take a look at just how Fodor handles polysemy in his atomism.

In discussing *verbs* like “keep,” he writes, “my theory is that there is no such thing as polysemy. The appearance that there is a problem is generated by the assumption that there are definitions; if you take the assumption away, the problem disappears” (1998, 53). The problem he means is that we seem to have the intuition that words like “keep” are univocal but that they also differ in how they feature in expressions like “keep the money” and “keep the crowd happy.” One way of handling the difference that Fodor considers is to hold that “keep” has a single definition, but that the definition is subject to a semantic field feature that adjusts the meaning to the context (1998, 50; see Jackendoff 1992). Or putting it differently (and alluding to a different approach), we might say that “keep” has a “core meaning together with a system of metaphorical extensions” (1998, 50; see Lakoff 1988). Fodor rejects both and instead argues that “keep” is univocal and that it simply means KEEP in both; he maintains that its contribution to complex expressions remains invariant despite differences in linguistic context, and that the seeming difference between keeping money and keeping the crowd happy comes down to the difference between having money and the crowd being happy, rather than a difference appended to definitions or “keep” alone. He and Lepore (2002, ch. 5) pursue a similar point in criticizing Pustejovsky (1995), where they introduce complex lexical entries to ensure that expressions such as “want” remain invariant across grammatically distinct constructions like “wants a beer” and “wants to have a beer,” a theoretical development that allows them to continue to hold that “want” contributes WANT to the complex expressions it features in. There is a bit more nuance there in that the complex entry specifies the expression’s rule of composition, but it still boils down to the idea that “want” means WANT. Despite appearances, they maintain that all of these different uses involve the same concept featuring the same structure and the same content.

Fodor and Lepore do make room for polysemy in standard cases of *nouns* like “lamb,” but they downplay its importance (2002, 109–110; see Rakova 2003, 135–137). They concede that these expressions may really be associated with distinct concepts, where each of the concepts may be thought of as having different contents because they refer to slightly different things—very much like the case sketched earlier for “mink.” Briefly put, they hold that polysemy relies on there being some nomic or logical relation obtaining between the referents of the concepts in question, on that relation affording some sort of overlap or coincidence in their contents, and, crucially, that we see or understand that connection—such that our knowing about the connection is ultimately what makes an expression polysemous rather than

homonymous. Their idea isn't very clearly developed or elaborated, but they provide an example that might help us get a better grip on it:

Suppose it's right that "lamb" is polysemous between the animal and the meat. Surely that's because lamb-the-meat comes from lamb-the-animal. Surely there just *couldn't* be a word that's polysemous between *lamb-the-animal* and (say) *beef-the-meat*? ... That there couldn't may itself sound like a deep fact of lexical semantics. But no; it's just the truism that, the less one can see what the relation between X and Y might be, the more one is likely to think of an expression that is X/Y ambiguous as homonymous rather than polysemous. Opportunities for polysemy arise from *how things are in the world* (or, anyhow, from how we take them to be). (2002, 117)

Now as far as the relations between things in the world are concerned, obviously they generally obtain whether or not we see them. Which is to say that whether or not we realize that lamb-the-meat comes from lamb-the-animal, lamb-the-meat really does come from lamb-the-animal. If we *do not* know about that connection, we'll consider that expression homonymous between the meat and the animal. If we *do* know about it, we'll consider that expression polysemous between the meat and the animal. Notice, however, that the distinction doesn't make any difference to the relations themselves—their objects enter into them whether or not we know about them. At the same time, and perhaps less obviously, these relations don't make any difference to the concepts they seem to concern. While they obtain between the referents of the contents of certain concepts, the content of a concept is just whatever it refers to, not whatever it refers to together with the relations those referents enter into.⁶ As important as the relations between things in the world generally are, they don't figure in contents, and so they don't figure in concepts. So although the whole scenario starts with the conceptually significant assumption that we're dealing with distinct concepts that have different contents, nothing following that assumption has any bearing on the concepts in question as such. It's that line of thinking that stands behind Fodor and Lepore's strong claim earlier in that paper: "We suspect that there is nothing that is both interesting and general to say about such cases; the meanings of words can overlap in all sorts of ways, so there are all sorts of ways in which polysemous terms can differ from mere homonyms. Nothing in the literature convinces us that there are powerful generalizations to state" (2002, 116). For them, it is precisely because polysemy and homonymy have no real bearing on concepts as such that they don't present interesting cases. And there not being any generalizations to state is simply due to the fact that the relations in question come down to the way things in the world are and what we think we know about them, not the way concepts are construed. From their perspective, even though they're prepared to say that we do consider some expressions homonymous and others polysemous, there really isn't any theoretically significant or relevant difference between homonymy

⁶ Of course, those relations may be used for the purposes of individuating concepts (Fodor 1998, 74); the point is that they don't constitute them.

and polysemy where concepts are concerned—in short, for them, nothing really turns on that distinction.

So on the one hand, Fodor (and Lepore) contend that monomorphemic verbs like “want” and “keep” are only associated with a single concept from the outset, and so they aren’t polysemous any more than they’re homonymous (which is to say that they aren’t either). And monomorphemic nouns like “lamb” and “mink” are polysemous, but the polysemy isn’t linguistically interesting because it has more to do with connections between things in the world and our epistemic situation than concepts or meanings themselves. For those concepts, the difference in content is what makes them distinct and that’s that. But polymorphemes are unlike either set of expressions. Obviously they aren’t like the first set, in that the contextual sensitivity that threatens “keep” and “want” due to the seemingly different ways they feature in expressions like “keep the money” and “keep the crowd happy” or “wants a beer” and “wants to have a beer” simply don’t translate to a noun like “doorknob.” There aren’t any intuitively problematic expressions to mediate, and however the connection between complexes and correlates might be construed, it certainly doesn’t have anything to do with definitions, which is what Fodor was primarily concerned with in their case. The comparison to the second set is much more interesting. Where monomorphemic expressions like “lamb” correspond to two concepts that both have the same simple structure and *differ in content*, polymorphemic expressions like “doorknob” correspond to two concepts that both have the same content and *differ in structure*. The connection or similarity between the complex and the correlate is the content, but at the level of content alone they aren’t distinct. It’s only at the level of concepts that we can take into account content together with structure, and so it’s only at the level of concepts that we can appreciate that there is a difference between them—and accordingly it’s only at the level of concepts that we can even discern that we’re dealing with two concepts rather than one. So whereas the concepts corresponding to “lamb” and “mink” are differentiated by their distinct contents, the concepts corresponding to “doorknob” are differentiated by their distinct structures (as evidenced earlier by their possession conditions). It follows that complexes and correlates certainly “overlap” in terms of content, which is to say that, by Fodor’s own criteria, they qualify their polymorphemes as polysemous in virtue of their shared content, and they remain distinct in terms of their structures, which is to say that it doesn’t all come down to the connections between the referents of the concepts but rather our construal of conceptual structure, which is altogether independent of those referents. It isn’t so much about what we think we know about things in the world as it is about what we know about our theory of concepts, of meaning—in a word, what we know about atomism. And having learned about complexes and correlates and sketched a few plausible modes of acquisition, we’ve effectively substantiated this bizarre form of polysemy within the framework.

That’s where the generalization lies, and it’s what makes this polysemy “interesting,” contrary to Fodor and Lepore’s claim that there are no interesting generalizations here. The fact that the complexes and correlates associated with polymorphemes are individuated by possession conditions relocates the usual differentiating factor to the structure of the concept rather than its content. While that speaks to the surprising flexibility of atomism’s theoretical framework, it

also leads us to very particular theoretical commitments and structures, and they all come as a consequence of atomism's theoretical architecture. In this case, what allows us to arrive at an expression with a complex and correlate pair is the very idea of allowing concepts to be complex or simple and allowing those concepts to also serve as meanings. It's tantamount to having complex and simple meanings as such, and that creates the conceptual space we need to have one content represented in two distinct *structures*, both of which can then be regimented under one expression. It's also what renders the polysemy explored above so abstract and effectively tacit—it's something that comes about at an intermediate theoretical level, between the content the concepts represent and the surface form of the expression, making it unusually subtle and theory-bound. Now, of course, none of that is necessarily problematic, but it is when the structures and commitments it secures are sufficiently unintuitive and when they lead to increasingly unintuitive consequences.

The issue is that at bottom atomism commits itself to a polysemy that attaches to a class of expressions that are not even intuitively ambiguous to begin with—whatever your view of polysemy and homonymy might be, expressions like “doorknob” and “breakfast” aren't intuitively ambiguous, and *a fortiori* they aren't intuitively homonymous or polysemous. Not only do polymorphemes present an incredibly large class of expressions, it's an abstract, structural, and unexpected polysemy motivated solely by atomism's theoretical commitments. Granted, the structural possibility of this polysemy is one thing, whether it actually obtains is another. Insofar as a polymorpheme corresponds to a complex concept, and some of them must, given that Fodor characterized monomorphemes as uniquely corresponding to simple concepts (as noted earlier), then they have an available correlate. Insofar as it is available, someone out there might be relying on it—which means that any number of polymorphemes may be plausibly paired with a complex and correlate. While the theoretical picture is quite clear cut and uniform, the actual distribution of complexes and correlates is anything but. For the polymorphemes that correspond to complexes and which therefore have correlates, there is the further matter of actually establishing which complexes and correlates are in play (that is, in use) over and above the theoretical possibility facilitated by the position. I believe this state of affairs is a serious count against atomism as a theory, in the way that any theory suffers when it can be shown to lead to unintuitive consequences and problems. Those consequences include the proliferation of correlates (and perhaps even the very idea of a correlate, for that matter) that can be understood as grounding a polysemy among expressions that intuitively aren't even ambiguous. Those problems include actually determining which of those polymorphemes correspond to complexes and which of those have correlates that are actually in play, the solution to which only further situates an already uninviting and likely unwanted picture. Even if it could be shown to be limited to only a handful of cases, they would still count against the plausibility of the position. And as long as it potentially affects *any number of polymorphemes*, it threatens to multiply meanings and the number of such cases on a large scale.

3.2 Communication

To get a better sense of how this aspect of atomism complicates our understanding of natural language in other respects, we can examine something more tangible, like communication. A simple sketch of the field together with potential issues will suffice for our purposes. An example involving a standard polysemic expression and communication in atomistic terms is the best place to start before moving to an analogous example involving a polymorpheme and the issues that entails.

Imagine that someone says “bank” while only possessing the concept corresponding to the institution, and someone listening hears them say “bank” while only possessing the concept corresponding to riversides, which is to say that one only knows the word to mean the institution and the other only knows the word to mean riversides.⁷ Clearly, they would just be talking past one another due to their conceptual differences; the hearer wouldn’t be able to correctly receive the utterance because they wouldn’t have the concept it expresses for the speaker. Practically speaking, that difference would lead to all sorts of problems that would demonstrate their failure to communicate. Let’s say that you’re the speaker and I’m the listener; if you told me to meet you at the “bank,” I would almost certainly be waiting in the wrong place, next to some river. The miscommunication would be evident in our failing to meet. And of course, even if we did find occasional success, it wouldn’t necessarily translate to our having successfully communicated. For instance, imagine that the local financial bank happens to be located on the bank of a river. Clearly, the fact that we would then succeed in meeting up despite our conceptual differences would just be a happy (and unlikely) accident; we still would have been talking past one another, we just got lucky.

We can consider how the same scenario would play out with a polymorpheme. Let’s use “breakfast.” It plausibly has two concepts associated with it too, a complex and a correlate, but this time the concepts both have the same content. Suppose you ask me about what I had for breakfast. If our concepts match, we can say that we understood one another, and so I’ll tell you about what I ate this morning. If our concepts don’t match, say you only have the complex and I only have the correlate, I would still tell you about what I ate this morning. The practical outcomes are the same whether we share the same concepts or not, because either way our concepts share the same content and so refer to the same thing. But if we don’t share the same concepts with respect to that particular expression, which is to say that it doesn’t have the same meaning to us, it seems that we should say that we’re talking past one another again. And yet, where failure in communication was evidenced by our usually failing to meet in the case of “bank,” there isn’t an analogous sense in which we could fail in the case of “breakfast,” in fact, on the contrary, we would find nothing but practical success in talking about breakfast.

⁷ Of course, though more complicated, the same considerations can be put forward for people who have both concepts associated with “bank” but who simply misinterpret one another—*mutatis mutandis* in what follows.

That brings us to the issue. The clearly accidental nature of the practical success in the bank case squares with the intuition that we don't successfully communicate when our conceptual repertoires don't match with respect to the expressions in play, which is just the uncontroversial idea that we don't successfully communicate if we're talking past one another, even if it occasionally leads to practical success. But the polymorpheme case doesn't present the same accidental picture; despite the difference in conceptual repertoires, our example would seem to lead us to regular practical success. There are two ways to approach this. Either we actually fail to communicate despite our practical success or we actually succeed in communicating despite our differing concepts. The former lines up with a simple decoding picture of communication and the latter with a more flexible practice-oriented picture, neither of which plays out very well.

Going with the former, holding that we fail to communicate in a case like the one sketched above follows from the notion that communication requires arriving at identical concepts. In brief, the idea is that we have some thought or information in our heads, we encode that information in a message, and someone can then receive that message and reconstruct the information it contains, recreating the original thought. *Mutatis mutandis* for atomism, where in our case we're concerned with concepts rather than information per se. Of course, this sort of approach usually relies on some notion of interpretation and access, which don't seem particularly compatible with atomism, but nonetheless it does seem to be a position that Fodor is sympathetic to in his critical works—for instance, claiming that communication requires expressions to “mean the same to all of us” (Fodor & Lepore 2007, 684). Our example leads to our failing to communicate precisely because the expression being used can't “mean the same” to us, as I simply don't have the same concept that you do. The problem is just that this approach paired with atomism requires us to accept that we may technically fail to communicate despite all outward appearances; we might find nothing but regular practical success, and yet it may register as regular failure to communicate due to this abstract difference in concepts. The problem is compounded when we take into account our earlier point concerning just how many polymorphemes there are and that many of them may present the same case as “breakfast,” which means that this unusual form of miscommunication may affect a large proportion of the things we say. The tacit polysemy explored earlier leads to an equally unintuitive and abstract possibility of regular yet seemingly inconsequential miscommunication here.

The latter may seem more appealing to an enterprising atomist weary of the above. The idea would be that we can actually succeed in communicating in cases where we don't have the same concepts, presumably in light of the fact that we can achieve regular practical success (see, e.g., Quine 1992; Brandom 2007). Naturally, it would probably turn on content, as content is what seems to lead us to that success. Whatever shape this sort of position ultimately takes, it would obviously render matching conceptual structures not necessary for communication. Nor would it take matching conceptual structures to be sufficient, as having two concepts with the same structure but radically different contents would undermine practical success in most cases. The focus would have to be on having concepts with the same content, whatever their structure. But that seems to be at odds with the emphasis placed on

that distinction elsewhere. For instance, Fodor's solution to Frege cases like WATER and H₂O take conceptual structure to be precisely what made all the difference, and that was presumably intended to be the case in communication as well as in abstraction. Conceptual structure is also what undergirds our notion of complex concepts as corresponding to phrases and sentences as opposed to simple concepts in the context of natural language; ignoring that would potentially further complicate our picture of communication. More generally, if atomism takes conceptual structure to be a fundamental distinction, and our conception of communication flouts that distinction, then it isn't at all clear that the two are even compatible insofar as they both concern natural language. That is, it isn't clear how our atomistic picture of natural language could maintain its theoretical commitment to distinguishing conceptual structures and at the same time concede that conceptual structure is irrelevant to natural language wherever communication is concerned, which seems to be practically everywhere. Seeing that, the enterprising atomist might try to soften the position and hold that conceptual structure may make a difference in some principled cases or on some condition in order to keep it in play. The problem we would then face is actually accounting for conceptual structure in communication given that it can't simply be read off of the surface form of the expression and that there obviously aren't any linguistic conventions concerning conceptual structure; in effect, the problem is accounting for conceptual structure in our pragmatics, and it's a serious problem at that.⁸ Speakers don't have intentions or make inferences concerning complexes or correlates (see, e.g., Grice 1989; Sperber & Wilson 1986; Bach & Harnish, 1979), there certainly isn't anything in the common ground concerning conceptual structure (see, e.g., Stalnaker 1978, 2002; Garcia-Carpintero 2015), and there is nothing discernible in the context more broadly construed (e.g., Kaplan 1989) that tells us anything about the structure of the concepts being expressed. Even if one of these approaches could somehow be applied here, there are further problems that follow; for instance, how cases in which we don't share the same concepts are actually supposed to be resolved or, if the mechanism involved were, say, inference, what would even prompt someone to make an inference to a concept with a different structure if both concepts in question always work with the initial interpretation and lead to the expected outcome. In short, there doesn't appear to be an approach to pragmatics that could actually account for something as abstract and fine as a complex and correlate concept regimented under a single surface form. Although what I've described above is only a sketch, the issues it illustrates are clear enough. It is difficult to imagine how the nature of conceptual structure in atomism could be reconciled with a workable picture of communication. Whether it's considered from the perspective of decoding or some variation of practice, we face significant obstacles as a consequence of atomism's facilitation of complexes and correlates, further problematizing its relation to natural language.

⁸ Fodor deferred to pragmatics in his (1987), but the considerations he explored in relation to twin-earths through quantification and universes of discourse have no application here.

3.3 Closing Remarks

What this section has illustrated is that atomism struggles with a pervasive form of polysemy. Not only does it affect a broad collection of expressions (i.e., polymorphemes), it seems to straightforwardly follow from atomism's theoretical architecture. As we noted above, it's a consequence of the way concepts are structured and the way they serve as meanings; insofar as we have correlates, the threat of this polysemy is right around the corner, given how Fodor conceives of concept possession. That there is a structural polysemy unique to atomism is, of course, something of a find in and of itself, seeing as Fodor was so hostile to it. That there is potentially so much of it, amidst expressions no one considers polysemous, isn't so much interesting as it is alarming. That concern deepens when that polysemy is only felt on an abstract theoretical level and may also be unevenly distributed among users. The consequences include, but are not limited to, the way communication is rendered in the context of polymorphemes: following one path, we may consistently theoretically fail to communicate, though apparently without consequence or anyone's knowledge, and following the other, we may theoretically succeed in communicating, but at the price of undercutting the application of the position to communication altogether, significantly weakening the case for atomism's place in accounting for natural language.

In closing, I'd like to quickly address two potential replies. The first is the impulse to revise our initial assumption concerning simple expressions and simple concepts: rather than holding that monomorphemic simple expressions correspond to simple concepts (and hence that at least some polymorphemic simple expressions do not), we might assume that all simple expressions correspond to simple concepts. Setting aside the fact that Fodor had reasons for hedging the claim and the seemingly ad hoc nature of the proposed adjustment, it doesn't actually address the core issue of structurally distinct concepts sharing the same content—an arrangement that intentionally holds for WATER and H₂O, among others. Insofar as that theoretical possibility is preserved, we also preserve the possibility of producing an example of a polymorpheme that expresses a complex concept, which would render the assumption false and anything relying on it unsound. Unless someone is prepared to commit themselves to the further claim that no such cases could obtain, it doesn't get us very far (suffice it to say, it wouldn't be very hard to produce an example undercutting that further claim, and thus, again, the assumption). The second is to question the notion of a correlate. I consider it a useful idea here because it gives us a clearer picture of what's going on and provides us with a way of talking about it, but it isn't the source of the problem. If we were to somehow block the theoretical inference from complex concepts to simple counterparts, it wouldn't resolve the intuitive problem posed by congruent and difficult to discern concepts floating around. Whether we accept the idea of a correlate or not, we still need to contend with the fact that expressions like “doorknob” or “breakfast,” within the atomistic framework, plausibly have multiple concepts associated with them that share the same content. If anything, the idea that for every complex concept there is a correlate provides us with some uniform footing that could be used to fashion a systematic solution. Rejecting it just leaves us with a mess of particular cases; it doesn't improve the atomist's position. If there

is a way of cleaning all of this up, I think it will require something more in-depth, or simply withdrawing from natural language altogether.

As far as withdrawing from natural language goes, I'm sure an atomist would be quick to point out that nothing I've described in this section concerns Mentalese. Mentalese isn't something we use to communicate with (it's involved in communication, but isn't the mode of communication); it's essentially private (see, e.g., Fodor 2008, 80, 214; 2009). Furthermore, it's a surfaceless language, which means that concepts have no additional layer over and above the concept itself. So, although we may have both a complex and a correlate, we don't have a single expression or representation under which to regiment them. Without the additional layer, we can't have polysemy (or homonymy, for that matter). As such, there isn't any trouble with polysemy, and there isn't any trouble with communication. It's worth reiterating the point I made earlier: I'm happy to believe that atomism accounts quite well for Mentalese. The question was, however, whether it can do the same for natural language.

4 Conclusion

There is much to admire in semantic atomism, but some of the more complex issues that it creates for itself seem to have gone unnoticed. If atomism is to compete as a foundational approach to natural language, these issues need to be addressed or at least taken into account. The fact that this form of polysemy and the set of terms it affects are unique to atomism makes it a relevant point of contention, in addition to being a concern for atomists that may have thought their position to be mostly polysemy-free even where natural languages are concerned. The consequences for communication are also relevant, as it's an integral function of natural language; given that the difficulties atomism runs into have no bearing on or reflection in use, we have reason to be suspicious of it as effectively creating theoretical problems where there don't seem to be any in practice. Whether one is an atomist or one of its critics, these matters ought to be of interest; for the former as something to consider and for the latter as a new set of tools.

Funding NCN Preludium Grant – Nr. 2019/33/N/HS1/02156, “Meaning, Possibility, and the Non-finite Architecture of Language”.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bach, K., & Harnish, R. (1979). *Linguistic communication and speech acts*. MIT Press.
- Brandom, R. (2007). Inferentialism and some of its challenges. *Philosophy and Phenomenological Research*, 74(3), 651–676.
- Cuyckens, H., & Zawada, B. (Eds.). (2001). *Polysemy in cognitive linguistics*. John Benjamins Publishing Company.
- Devitt, M. (1996). *Coming to our senses: A naturalistic program for semantic localism*. Cambridge University Press.
- Fodor, J. (1982). Cognitive science and the twin-earth problem. *Notre Dame Journal of Formal Logic*, 23, 98–118.
- Fodor, J. (1998). *Concepts: Where cognitive science went wrong*. Clarendon Press.
- Fodor, J. (2001). Language, thought and compositionality. *Mind & Language*, 16, 1–15.
- Fodor, J. (2008). *LOT 2: The language of thought revisited*. Clarendon Press.
- Fodor, J. (2009). Enough with the norms already! In A. Hieke & H. Leitgeb (Eds.), *Between the mind and the brain* (pp. 1–8). De Gruyter.
- Fodor, J., & Lepore, E. (1992). *Holism: A shopper's guide*. Blackwell.
- Fodor, J., & Lepore, E. (2002). *The compositionality papers*. Clarendon Press.
- Fodor, J., & Lepore, E. (2007). Brandom beleaguered. *Philosophy and Phenomenological Research*, LXXIV, 677–691.
- Fodor, J., & Pylyshyn, Z. (2015). *Minds without meanings: An essay on the content of concepts*. MIT Press.
- Garcia-Carpintero, M. (2015). Contexts as shared commitments. *Frontiers in Psychology*, 22. <https://doi.org/10.3389/fpsyg.2015.01932>
- Grice, P. (1989). *Studies in the ways of words*. Harvard University Press.
- Jackendoff, R. (1992). *Languages of the mind: Essays on mental representation*. MIT Press.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan*. Oxford University Press.
- Lakoff, G. (1988). Cognitive semantics. In U. Eco (Ed.), *Meaning and mental representations* (pp. 119–154). Indiana University Press.
- Margolis, E. (1999). How to acquire a concept. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings*. MIT Press.
- Pullum, G., & Scholz, B. (2010). Recursion and the infinitude claim. In H. Hulst (Ed.), *Recursion in human language* (pp. 113–138). De Gruyter.
- Pustejovsky, J. (1995). *The generative lexicon*. MIT Press.
- Quine, W. V. (1992). *Pursuit of truth*. Harvard University Press.
- Rakova, M. (2003). *The extent of the literal: Metaphor, polysemy, and theories of concepts*. Palgrave Macmillan.
- Sennet, A. (2016). Polysemy. Oxford Handbooks Online. Retrieved 29 Jan. 2020, from <https://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199935314.001.0001/oxfordhb-9780199935314-e-32>.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Blackwell.
- Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, 9, 315–332.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25(5-6), 701–721.
- Taylor, J. (2003). Cognitive models of polysemy. In B. Nerlich, Z. Todd, V. Herman, & D. Clarke (Eds.), *Polysemy: Flexible patterns of meaning in mind and language* (pp. 31–48). Mouton de Gruyter.
- Tuggy, D. (1993). Ambiguity, polysemy, and vagueness. *Cognitive Linguistics*, 4(3), 273–290.

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