Many, but almost Holmes?

1 Introduction

Worlds of fiction are incomplete. While much is specified about the world of, for instance, Sherlock Holmes, very much more is left unspecified. According to Conan Doyle's stories, it is true that Holmes lives in Baker Street. However, it is not true (or false) according to the stories that Holmes's blood type is B-positive.

Possible worlds are complete. For each individual in a possible world, it is true or false whether that individual has a blood type of B-positive.

Reflection on this difference between possible and fictional worlds led Saul Kripke to famously argue that worlds of fiction cannot be possible worlds, and, more specifically, that the referents of fictional names cannot be *possibilia*. His argument has influenced various authors to abandon possibilia in the realm of fiction in favor of other kinds of entities such as abstracta or Meinongian objects.[1]

As it happens, there are similarities between the problem Kripke poses for possibilia and another familiar problem – the problem of the many. Once the analogies between the two problems are appreciated, one solution to the problem of the many seems at first glance capable of defeating Kripke's case against possibilia. In what follows, I shall first explore the similarity between the two problems and then consider the prospects for undermining Kripke's argument.

[1]See, for example, Salmon (1998: 292), Thomasson (2003: 17-18), and van Inwagen (2008: 43) for abstractists who endorse Kripke's argument.

2 Kripke's argument

Kripke disqualifies possibilia as the referents of fictional names in three sentences:

> I hold the metaphysical view that, granted that there is no Sherlock Holmes, one cannot say of any possible person that he would have been Sherlock Holmes, had he existed. Several distinct possible people, and even actual ones such as Darwin or Jack the Ripper, might have performed the exploits of Holmes, but there is none of whom we can say that he would have been Holmes had he performed these exploits. For if so, which one? (Kripke 1980: 158)

Kripke's argument relies on a tension between three claims. First, Kripke assumes the *uniqueness* of Holmes: if Holmes exists, there's only one of him. After all, the stories certainly suggest that there's one, and only one, Holmes. Second, Kripke notes a *multiplicity* of possibilia suited to serve as Holmes: given that any number of distinct possibilia could match the description and details of the Holmes stories, there are many distinct possibilia that have what it takes to qualify as the referent of 'Sherlock Holmes'. For instance, one possibile that matches the description of Holmes has a blood type of B-negative, while another has a blood type of B-positive. Third, there is the claim of *arbitrariness*: none of these possibilia enjoys a right to serve as the referent of 'Sherlock Holmes' to the exclusion of the rest. If each answers to the descriptions provided in the stories, there is no principled reason to prefer one over the other. How can a possibilist about fictional characters like Sherlock Holmes say, as the fiction demands, that there is one Holmes, when there are many possibilia suited to be Holmes, and no principled way of saying which of the many he could be?

3 The problem of the many

Kripke's argument against possibilia as the referents of fictional names appeals to a tension between three claims: *uniqueness*, *multiplicity*, and *arbitariness*. We find the very same principles, and the apparent tension between them, driving another conundrum – the problem of the many. This problem, introduced by Peter Unger (1980), is presented by David Lewis as follows:

> Think of a cloud – just one cloud, and around it clear blue sky. Seen from the ground, the cloud may seem to have a sharp boundary. Not so. The cloud is a swarm of water droplets. At the outskirts of the cloud the density of the droplets falls off. Eventually they are so few and far between that we may hesitate to say that the outlying droplets are still part of the cloud at all; perhaps we might better say only that they are near the cloud. But the transition is gradual. Many surfaces are equally good candidates to be the boundary of the cloud. Therefore many aggregates of droplets, some more inclusive and some less inclusive (and some inclusive in different ways than others), are equally good candidates to be the cloud. Since they have equal claim, how can we say that the cloud is one of these aggregates rather than another? But if all of them count as clouds, then we have many clouds rather than one. And if none of them count, each one being ruled out because of the competition from the others, then we have no cloud. How is it, then, that we have just one cloud? (Lewis 1993: 164)

Here, as in Kripke's argument, there is a claim of *uniqueness*: there is only one cloud. We also find *multiplicity*: there are many equally good candidates for being the cloud. And, finally, *arbitrariness*: it seems arbitrary to pick any particular aggregate of droplets to count as the cloud to the exclusion of the rest. How can we say there's one cloud, when there are many candidates for being the cloud, and no principled way of saying which of the many the

cloud is?

The similarity between the two problems is striking. It is equally striking that theorists of fiction have never considered viewing Kripke's problem as a variant of the problem of the many. For one solution to the problem of the many attempts to resolve the very tension that motivates Kripke's rejection of possibilia as the referents of fictional names.[2] If this solution can be adapted to the case of fiction, we have vindicated possibilia against Kripke's argument.

4 Supervaluations and the problem of the many

David Lewis advertises a solution to the problem of the many that resolves the tension between *uniqueness*, *multiplicity*, and *arbitrariness*. The tension is resolved with a *supervaluationist* move: no matter how we might have made the decision of which of the candidates should count as the cloud, much of what we want to say about the cloud will be true. In particular, "There is but a single cloud" will, as we intuitively desire, come out true since it will be true however we might have made the semantic decision. This allows us at the same time to admit that there are many equally good cloud-candidates while conceding that it would be hopelessly arbitrary to say of any of them that it should be considered the single cloud in the sky. Lewis explains the details as follows (moving from cloud-candidates to cat-candidates):

> Call a sentence *super-true* if and only if it is true under all ways of making the unmade semantic decisions; *super-false* if and only if it is false under all ways of making those decisions; and if it is true under some ways and false under others, then it suffers a truth-value gap... Each intended interpretation of our language puts one of the cat-candidates on the mat into the extension of

[2]Other solutions to the problem of the many simply reject one of the principles in tension. Such solutions aren't happily applied to Kripke's argument, since it's hard to see which of the principles *could* be rejected if thinking about fictional worlds as possible worlds, or fictional characters as possibilia.

> the word 'cat', and excludes all the rest. Likewise each intended interpretation picks out one cat-candidate, the same one, as the referent of 'Tibbles'. Therefore it is super-true that there is just one cat, Tibbles, on the mat. Because it is super-true, you are entitled to affirm it. And so you may say what you want to say: there is one cat. (Lewis 1993: 172-3)

Lewis's supervaluationist solution gives us a way of maintaining *uniqueness*, *multiplicity*, and *arbitrariness* while eliminating the apparent tension between them. Given that what motivates Kripke's argument against possibilia just is the supposed tension between such principles, the availability of an analogous solution in the case of fiction would seem to provide a way to block the argument.

5 Supervaluations and possibilia

The supervaluationist solution is easily adapted to the realm of fiction. In the Holmes case, we have a legion of possibilia, each of which is a candidate for being Sherlock Holmes in virtue of matching the description in the stories. A supervaluationist solution here holds, as with the problem of the many, that a uniqueness sentence like "There is one Sherlock Holmes" is super-true if and only if it is true on every interpretation, or precisification, of the name 'Sherlock Holmes'.

We should expect precisifications of 'Sherlock Holmes' to deliver Holmes-candidates, candidates for the referent of the name, but what *exactly* are Holmes-candidates? Kripke's argument assumes that they are individuals, possible people. However, a Plantingan might suggest that precisifications deliver instead unactualized haecceities. Alternatively, a Lewisian could hold that precisifications deliver not individuals but rather sets (or instead mereological sums) of such individuals, something akin in the fictional case to the set (or sum) of counterparts for an actual individual.

These latter Plantingan and Lewisian suggestions are unamenable to a possibilist account of fiction. For neither unactualized haecceities nor

sets/sums of possible individuals are properly suited to serve as Holmes-candidates. As Kripke brings out and as mentioned above, for the possibilist, Holmes-candidates are supposed to be entities that do what is attributed to Holmes in the fiction.[3] Part of the very appeal of the possibilist position in the realm of fiction is that it allows us to say things that we intuitively want to say about Holmes. We want to say that Holmes smokes pipes and solves mysteries, and construing Holmes as a possible individual would allow us to say just that. Sets/sums do not smoke pipes or solve mysteries. Nor do unactualized haecceities.[4] I accordingly follow Kripke in his assumption that Holmes-candidates are, and thus that precisifications deliver, individuals. I shall leave open, however, whether they are best considered as Lewisian world-bound entities or instead as Kripkean transworld individuals (where transworld identity is a primitive relation).

A precisification or interpretation of the name 'Sherlock Holmes' leaves no facts unaccounted for. It yields a complete individual occupying a complete world. On some interpretations, the candidate delivered will have a B-positive blood type. On others, it will not. Some interpretations will deliver a Holmes-candidate that reads *The Times* with Watson on the first Sunday of each month. Others will deliver a candidate that does not.

Facts extrinsic to the Holmes-candidates themselves will also figure in each precisification of a fictional name. Some Holmes-candidates will occupy a world with an even number of Tory members of Parliament, others will occupy a world with an odd number, and yet others will occupy a world where there are none.

Depending on the details of the fiction, temporal or locational facts might also be needed to distinguish two qualitatively similar candidates for the

[3]Cf. Lewis (1983: 267): "[T]he sense of "Sherlock Holmes" as we use it is such that, for any world w where the Holmes stories are told as known fact rather than fiction, the name denotes at w whichever inhabitant of w it is who there plays the role of Holmes."

[4]Moreover, the unactualzied haecceity view resembles little more than a terminological variant of an abstractist view in the final analysis. On the unactualized haecceity picture, Sherlock Holmes ends up being an abstract entity, since an unactualized haecceity of Holmes is simply the property of being Sherlock Holmes, which is uninstantiated at the actual world. Plantingan possibilism in the realm of fiction, then, doesn't give us a *distinctively possibilist* position.

referent of a fictional name. For instance, descriptions in certain stories could be satisfied by distinct possibilia occupying the *same* world. The Holmes stories appear to take place in the Victorian era, given the details provided of London and its inhabitants. We could imagine, though, that the stories were less detailed such that one couldn't determine that they took place in any particular location or historical era. As such, one possibile could meet the description of the Holmes stories at one point in the history of (or at one place in) that world, while a distinct possibile could do so at a later point in the history of (or location in) the same world. In these cases, an interpretation of the name 'Sherlock Holmes' will deliver one or the other of these candidates depending on how the particular interpretation completes the locational and temporal details of the world sketched in the stories.

The supervaluationist story will also differ depending on one's view of the modal landscape. For instance, those who believe in qualitatively duplicate, yet distinct, Holmes worlds could rely on the distinctness of each such world to generate the distinctness of the various candidates. Or, alternatively, one could simply appeal to haecceitistic differences between the candidates themselves. Those who deny qualitatively duplicate worlds can get the distinctness of the candidates simply in virtue of the distinctness of the sets of properties had by the candidates themselves, as well as those had by the worlds they occupy.

Each interpretation of the name 'Sherlock Holmes' delivers a unique Holmes-candidate to the exclusion of the rest, since each interpretation offers a distinct completion of the Holmes-candidate and his world. As a result, the sentence "There is one Sherlock Holmes" is super-true. Following Lewis, since the sentence is super-true, you are entitled to affirm it. Thus, *uniqueness* holds, in spite of the fact that *multiplicity* and *arbitrariness* also hold. In the case of fiction, then, supervaluations seem initially promising in resolving the tension between the three principles, thereby undercutting the motivation behind Kripke's argument.

6 A paradoxical result?

While the supervaluationist solutions above promise a way to resolve the tension between *uniqueness*, *multiplicity*, and *arbitrariness*, there is a significant worry about how they provide this resolution. When we consider the object language in such solutions, uniqueness sentences like "There is just one cloud" and "There is just one Holmes" happily come out true. Yet when we examine the metatheory for that language, we find, paradoxically, that there are *many* things, each of which qualifies as a cloud (or, in the fictional case, Sherlock Holmes), on some interpretation of the language.

To get a sense of the paradox's structure, it is worth comparing the semantic consequences of these supervalutionist solutions with what some have seen as a paradoxical result of the Löwenheim-Skolem theorem in light of Cantor's theorem:

> Cantor's theorem says that there are uncountable sets – sets which are too big to be put into one-to-one correspondence with the natural numbers. The Löwenheim-Skolem theorem says that if a countable collection of first-order sentences has a model, then it has a model whose domain is only countable. Skolem's Paradox arises when we note that the standard axioms of set theory are themselves a countable collection of first-order sentences. (Bays 2006: 616)

The paradox can be cashed out in terms of the difference between what Cantor's theorem says, and the model in which that theorem is true. Cantor's theorem says that there are uncountably many sets, and yet this can be true in a *countable* model. How exactly could a sentence that posits the existence of uncountably many things be true in a countable model? To bring out even further the paradoxical nature of this result, consider that the model could be the actual world, a world of countably many things. A standard reaction to the paradox is to suggest that a sentence that we take to express there being uncountably many things cannot really do so when interpreted in a countable model. Thus, one might think Cantor's theorem

simply cannot receive its standard interpretation in such a model.[5]

The worry for the supervaluationist solutions discussed above is analogous. How can a sentence like "There is only one cloud" or "There is only one Sherlock Holmes" be true in light of a metatheory that contains *many* distinct cloud-candidates or Holmes-candidates? Of course, such sentences *come out* true given a supervaluationist model. But the results conflict with our ordinary conception of uniqueness sentences. We typically take such sentences to be made true by the existence of a single thing. Yet on the supervaluationist model, uniqueness sentences are ultimately made true by the existence of *many* things, not one.

We are left with a general structural problem for the supervaluationist solutions discussed above. This problem suggests that *no* such solution can satisfactorily resolve the tension between *uniqueness*, *multiplicity*, and *arbitrariness*.

7 Almost-identity

Lewis can be seen as responsive to just these worries. He augments his supervaluationist solution to the problem of the many with a notion of "almost-identity". Though it is true that *many* cloud-candidates or cat-candidates entitle us to affirm that there is one cloud or cat, the various candidates enjoy massive mereological overlap. Each cloud-candidate differs from the rest by just a few particles. Every cat-candidate differs from the rest by just a few hairs. In virtue of this overlap, Lewis holds that, while distinct, the various candidates are "almost" one:

> We have a spectrum of cases. At one end we find the complete identity of a thing with itself: it and itself are entirely identical, not at all distinct. At the opposite end we find the case of two things that are entirely distinct: They have no part in common. In between we find all the cases of partial overlap: things with

[5] As Bays (2006) argues, however, once we understand the subtleties of the background mathematics involved, the paradox in *this* area vanishes.

> parts in common and other parts not in common. (Sometimes one of the overlappers is part of the other, sometimes not.) The things are not entirely identical, not entirely distinct, but some of each. They are partially identical, partially distinct. There may be more overlap or less. Some cases are close to the distinctness end of the spectrum: Siamese twins who share only a finger are almost completely distinct, but not quite. Other cases are close to the identity end. For instance, any two of our cat-candidates overlap almost completely. (Lewis 1993: 177)

This emendation ultimately reveals Lewis's sensitivity to the paradoxical combination of, on one hand, the uniqueness sentences in our object language that come out true and, on the other, the many things appearing in the metatheory for that language. Indeed, Lewis's description of the merits of almost-identity can be seen as demonstrating just such an awareness:

> Sometimes, especially in our offhand and unphilosophical moments, context will favour the second, one-cat sort of interpretation; and then the supervaluation rule, with nothing to defeat it, will entitle us to say that there is only one cat. But sometimes, for instance when we have been explicitly attending to the many candidates and noting that they are equally cat-like, context will favour the first, many-cat sort of interpretation...But even then, we still want some good sense in which there is just one cat (though we may want a way to say the opposite as well). That is what almost-identity offers. (Lewis 1993: 180-1)

We want (as *uniqueness* demands) sentences like "There is one cat on the mat" to come out true. And yet (as *multiplicity* demands), when we notice the number of things that figure in our metatheory – the many cat-candidates – we are pushed to acknowledge that there are *many* things on the mat. According to Lewis, almost-identity gives us a way to say that, in *some* sense, even in the metatheory we have a number of candidates that are one thing. Well, almost.

By stressing the mereological overlap among the various cat-candidates, almost-identity may well help partially explain away the paradoxical results of the supervaluationist solution to the problem of the many. I can't adequately settle this issue here, but it's less clear that it *fully* addresses the worry set forth above. After all, we are still left with an object language that tells us that there is a single thing falling under some sortal, while our metatheory still features *many* (albeit almost-identical) things falling under that sortal on some interpretation of the language. No matter the result for the problem of the many, however, things look much worse when we turn our attention to the fictional case.

8 Almost-identity and possibilia

Almost-identity is a matter of mereological overlap. The various candidates in problem of the many cases are considered almost one because they enjoy near total part-sharing. But one who holds that Holmes is a possibile would be hard pressed to establish that the various Holmes-candidates should likewise be considered almost-identical.

Holmes-candidates are distinct possibilia occupying distinct worlds. Given this fact, it simply is not clear that we can maintain that they enjoy mereological overlap. After all, the stories certainly tolerate Holmes's being composed by any number of different combinations of parts. So we should expect, or at least allow, that Holmes-candidates are composed of various *non-overlapping* classes of parts. To generate almost-identity, one would have to stipulate that each of these candidates is composed of (almost all) the very same parts, i.e., that most of the parts that compose a particular Holmes-candidate in one world exist in and are shared by Holmes-candidates in all their respective worlds. It is not obvious what could successfully motivate such a stipulation, given the tolerance and flexibility the stories display with regard to what composes Holmes.

Moreover, given the vast differences among the Holmes-candidates, we have even further reason to deny that they are almost-identical. The Holmes-candidates will, of course, share *some* properties since they each will meet

the descriptions in the stories. However, there are many more properties, both intrinsic and extrinsic, that they do *not* share. Almost-identity allows for *some* discernibility amongst things that are almost identical. But as Lewis (1993: 178) notes with regard to the many almost-identical cat-candidates, we should "expect almost-identical things to be very similar in a great many ways: size, shape, location, weight, purring, behaviour, not to mention relational properties like location and ownership." As we've seen, given the lack of specificity in the stories, the Holmes-candidates will be incredibly *dissimilar* along many more dimensions than those along which they are very similar, thus violating the expectation of strong similarity that comes with almost-identity. The various Holmes-candidates are overwhelmingly different in a great deal of respects. Indeed, they are *so* different that our intuitive judgment would place them close to the distinctness end of the spectrum of identity, thereby making a plausible claim of almost-identity unavailable.

Finally, the case of fiction is inconsistent with the spirit behind the almost-identity proposal. Almost-identity emerges in part out of respect for the kinds of considerations mentioned above that lead to paradox. Lewis recognizes that we are of two minds when describing the metaphysical scenario in typical problem of the many cases. *Uniqueness* is driven by how things unreflectively seem to be, while *multiplicity* reflects our recognition that there seem to be many things equally suited to fall under a given sortal. The case of fiction certainly recommends *uniqueness*. The fiction suggests that there is but a single Holmes, and it seems to us there would be but one of him were he to exist. But we are drawn to *multiplicity* only as an artifact of considering fiction in the context of possible worlds. Nothing about how we think about the fiction in itself, or about Holmes himself, suggests a plenitudinous framework of Holmes-candidates stretching across the plurality of worlds. Almost-identity gains traction in how it figures in a reconciliation of our tendencies toward both *uniqueness* and *multiplicity*. No such duality of thought appears when we consider fiction. Thus, part of the very motivation for augmenting supervaluationism with almost-identity goes missing in the case of fiction.

When it comes to fiction, almost-identity cannot steer the supervaluationist solution from paradox. Thus, barring some other resolution, supervaluationism cannot satisfactorily resolve the tension between *uniqueness*, *multiplicity*, and *arbitrariness* in the fictional case.

9 Conclusion

Contrasting the case of fiction with the problem of the many as we have above can provide a more precise conception of the particular claims in tension in both cases. Perhaps the similarities and differences between the problem of the many and Kripke's argument are best brought out by pointing to a tension between *uniqueness*, *arbitrariness*, and not simply *multiplicity*, but *multiplicity of* **wholly distinct individuals**. Supervaluations may be utilized in an attempt to resolve this tension, but are only able to do so at the cost of paradoxical semantic results. Almost-identity can then be deployed in the problem of the many case as a method for responding to the paradox, namely by challenging whether and to what extent to which the various candidates for some sortal term that feature in the metatheory should be seen as *wholly* distinct. Framing things in this way brings out the real failing of supervaluationist solutions in the realm of fiction: no notion like almost-identity is available to challenge the complete distinctness of the various possibilia which might serve as a particular fictional character. Thus, no resources in the case of fiction can speak to the paradoxical results of a supervaluationist solution. And as we have seen, with no way to escape the paradox charted above, the semantic situation is rather dire.

Kripke's argument against possibilia as the referents of fictional names has been influential. Yet, given its similarity to the problem of the many, and the apparent adaptability of a solution to that problem, it has been difficult to see *exactly* why we should accept it. Once we set the argument in the context of the problem of the many, we can see the precise limitations one faces in trying to answer Kripke's argument. It thereby becomes clear not only that the argument works, but also *why* we should consider it a success.

References

[1] Bays, T. (2006). The mathematics of Skolem's paradox. In *Philosophy of Logic*, ed. Dale Jacquette, 615-48. Amsterdam: Elsevier.

[2] Kripke, S. (1980). *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.

[3] Lewis, D. (1983). Truth in fiction. In his *Philosophical Papers, vol. 1*, 261-75. Oxford: Oxford University Press.

[4] Lewis, D. (1993). Many, but almost one. In *Ontology, causality and mind*, ed. J. Bacon, K. Campbell, and L. Reinhardt, 23-42. Cambridge: Cambridge University Press. Reprinted in and page references to his *Papers in Metaphysics and Epistemology*, 164-82. Cambridge: Cambridge University Press, 1999.

[5] Salmon, N. (1998). Nonexistence. *Nous* 32: 277-319.

[6] Thomasson, A. (2003). *Fiction and Metaphysics*. Cambridge: Cambridge University Press.

[7] Unger, P. (1980). The problem of the many. *Midwest Studies in Philosophy* 5: 411-67.

[8] Van Inwagen, P. (2008). McGinn on existence. *The Philosophical Quarterly* 58: 36-58.