

## IV. TRUTH IN FICTION

DAVID LEWIS

WE can truly say that Sherlock Holmes lived in Baker Street, and that he liked to show off his mental powers. We cannot truly say that he was a devoted family man, or that he worked in close cooperation with the police.

It would be nice if we could take such descriptions of fictional characters at their face value, ascribing to them the same subject-predicate form as parallel descriptions of real-life characters. Then the sentences "Holmes wears a silk top hat" and "Nixon wears a silk top hat" would both be false because the referent of the subject term—fictional Holmes or real-life Nixon, as the case may be—lacks the property, expressed by the predicate, of wearing a silk top hat. The only difference would be that the subject terms "Holmes" and "Nixon" have referents of radically different sorts: one a fictional character, the other a real-life person of flesh and blood.

I don't question that a treatment along these Meinongian lines could be made to work. Terence Parsons has done it.<sup>1</sup> But it is no simple matter to overcome the difficulties that arise. For one thing, is there not some perfectly good sense in which Holmes, like Nixon, is a real-life person of flesh and blood? There are stories about the exploits of super-heroes from other planets, hobbits, fires and storms, vaporous intelligences, and other non-persons. But what a mistake it would be to class the Holmes stories with these! Unlike Clark Kent *et al.*, Sherlock Holmes is just a person—a person of flesh and blood, a being in the very same category as Nixon.

Consider also the problem of the chorus. We can truly say that Sir Joseph Porter, K.C.B., is attended by a chorus of his sisters and his cousins and his aunts. To make this true, it seems that the domain of fictional characters must contain not only Sir Joseph himself, but also plenty of fictional sisters and cousins and aunts. But how many—five dozen, perhaps? No, for we cannot truly say that the chorus numbers five dozen exactly. We cannot

truly say anything exact about its size. Then do we perhaps have a fictional chorus, but no fictional members of this chorus and hence no number of members? No, for we can truly say some things about the size. We are told that the sisters and cousins, even without the aunts, number in dozens.

The Meinongian should not suppose that the quantifiers in descriptions of fictional characters range over all the things he thinks there are, both fictional and non-fictional; but he may not find it easy to say just how the ranges of quantification are to be restricted. Consider whether we can truly say that Holmes was more intelligent than anyone else, before or since. It is certainly appropriate to compare him with some fictional characters, such as Mycroft and Watson; but not with others, such as Poirot or "Slipstick" Libby. It may be appropriate to compare him with some non-fictional characters, such as Newton and Darwin; but probably not with others, such as Conan Doyle or Frank Ramsey. "More intelligent than anyone else" meant something like "more intelligent than anyone else in the world of Sherlock Holmes." The inhabitants of this "world" are drawn partly from the fictional side of the Meinongian domain and partly from the non-fictional side, exhausting neither.

Finally, the Meinongian must tell us why truths about fictional characters are cut off, sometimes though not always, from the consequences they ought to imply. We can truly say that Holmes lived at 221B Baker Street. I have been told<sup>2</sup> that the only building at 221B Baker Street, then or now, was a bank. It does not follow, and certainly is not true, that Holmes lived in a bank.

The way of the Meinongian is hard, and in this paper I shall explore a simpler alternative. Let us not take our descriptions of fictional characters at face value, but instead let us regard them as abbreviations for longer sentences beginning with an operator "In such-and-such fiction . . .". Such a phrase is an intensional operator that may be pre-

<sup>1</sup> In "A Prolegomenon to Meinongian Semantics," *The Journal of Philosophy*, vol. 71 (1974), pp. 561–580, and in "A Meinongian Analysis of Fictional Objects," *Grazer Philosophische Studien*, vol. 1 (1975), pp. 73–86.

<sup>2</sup> I have also been told that there has never been any building at that address. It doesn't matter which is correct.

fixed to a sentence  $\phi$  to form a new sentence. But then the prefixed operator may be dropped by way of abbreviation, leaving us with what sounds like the original sentence  $\phi$  but differs from it in sense.

Thus if I say that Holmes liked to show off, you will take it that I have asserted an abbreviated version of the true sentence "In the Sherlock Holmes stories, Holmes liked to show off." As for the embedded sentence "Holmes liked to show off," taken by itself with the prefixed operator neither explicitly present nor tacitly understood, we may abandon it to the common fate of subject-predicate sentences with denotationless subject terms: automatic falsity or lack of truth value, according to taste.

Many things we might say about Holmes are potentially ambiguous. They may or may not be taken as abbreviations for sentences carrying the prefix "In the Sherlock Holmes stories . . .". Context, content, and common sense will usually resolve the ambiguity in practice. Consider these sentences:

Holmes lived in Baker Street.

Holmes lived nearer to Paddington Station than to Waterloo Station.

Holmes was just a person—a person of flesh and blood.

Holmes really existed.

Someone lived for many years at 221B Baker Street.

London's greatest detective in 1900 used cocaine.

All of them are false if taken as unprefixes, simply because Holmes did not actually exist. (Or perhaps at least some of them lack truth value.) All are true if taken as abbreviations for prefixed sentences. The first three would probably be taken in the latter way, hence they seem true. The rest would probably be taken in the former way, hence they seem false. The sentence

No detective ever solved almost all his cases.

would probably be taken as unprefixes and hence true, though it would be false if taken as prefixed. The sentence

Holmes and Watson are identical.

is sure to be taken as prefixed and hence false, but that is no refutation of systems of free logic<sup>3</sup> which would count it as true if taken as unprefixes.

(I hasten to concede that some truths about

Holmes are not abbreviations of prefixed sentences, and also are not true just because "Holmes" is denotationless. For instance these:

Holmes is a fictional character.

Holmes was killed off by Conan Doyle, but later resurrected.

Holmes has acquired a cultish following.

Holmes symbolizes mankind's ceaseless striving for truth.

Holmes would not have needed tapes to get the goods on Nixon.

Holmes could have solved the A.B.C. murders sooner than Poirot.

I shall have nothing to say here about the proper treatment of these sentences. If the Meinongian can handle them with no special dodges, that is an advantage of his approach over mine.)

The ambiguity of prefixing explains why truths about fictional characters are sometimes cut off from their seeming consequences. Suppose we have an argument (with zero or more premisses) which is valid in the modal sense that it is impossible for the premisses all to be true and the conclusion false.

$$\frac{\psi_1, \dots, \psi_n}{\therefore \phi}$$

Then it seems clear that we obtain another valid argument if we prefix an operator "In the fiction  $f$  . . ." uniformly to each premiss and to the conclusion of the original argument. Truth in a given fiction is closed under implication.

$$\frac{\text{In } f, \psi_1, \dots, \text{In } f, \psi_n}{\therefore \text{In } f, \phi}$$

But if we prefix the operator "In the fiction  $f$  . . ." to some of the original premisses and not to others, or if we take some but not all of the premisses as tacitly prefixed, then in general neither the original conclusion  $\phi$  nor the prefixed conclusion "In the fiction  $f$ ,  $\phi$ " will follow. In the inference we considered earlier there were two premisses. The premiss that Holmes lived at 221B Baker Street was true only if taken as prefixed. The premiss that the only building at 221B Baker Street was a bank, on the other hand, was true only if taken as unprefixes; for in the stories there was no bank there but rather a rooming house. Taking the premisses as we naturally would in the ways that make them true, nothing follows: neither the unprefixes conclusion that Holmes lived in a bank nor the prefixed con-

<sup>3</sup> For instance, the system given in Dana Scott, "Existence and Description in Formal Logic" in *Bertrand Russell: Philosopher of the Century*, ed. by Ralph Schoenman (London, 1967).

clusion that in the stories he lived in a bank. Taking both premisses as unprefixes, the unprefixes conclusion follows but the first premiss is false. Taking both premisses as prefixed, the prefixed conclusion follows but the second premiss is false.<sup>4</sup>

Our remaining task is to see what may be said about the analysis of the operators "In such-and-such fiction . . .". I have already noted that truth in a given fiction is closed under implication. Such closure is the earmark of an operator of relative necessity, an intensional operator that may be analyzed as a restricted universal quantifier over possible worlds. So we might proceed as follows: a prefixed sentence "In fiction  $f$ ,  $\phi$ " is true (or, as we shall also say,  $\phi$  is true in the fiction  $f$ ) iff  $\phi$  is true at every possible world in a certain set, this set being somehow determined by the fiction  $f$ .

As a first approximation, we might consider exactly those worlds where the plot of the fiction is enacted, where a course of events takes place that matches the story. What is true in the Sherlock Holmes stories would then be what is true at all of those possible worlds where there are characters who have the attributes, stand in the relations, and do the deeds that are ascribed in the stories to Holmes, Watson, and the rest. (Whether these characters would then *be* Holmes, Watson, and the rest is a vexed question that we must soon consider.)

I think this proposal is not quite right. For one thing, there is a threat of circularity. Even the Holmes stories, not to mention fiction written in less explicit styles, are by no means in the form of straightforward chronicles. An intelligent and informed reader can indeed discover the plot, and could write it down in the form of a fully explicit chronicle if he liked. But this extraction of plot from text is no trivial or automatic task. Perhaps the reader accomplishes it only by figuring out what is true in the stories—that is, only by exercising his tacit mastery of the very concept of truth in fiction that we are now investigating. If so, then an analysis that starts by making uncritical use of the concept of the plot of a fiction might be

rather uninformative, even if correct so far as it goes.

A second problem arises out of an observation by Saul Kripke.<sup>5</sup> Let us assume that Conan Doyle indeed wrote the stories as pure fiction. He just made them up. He had no knowledge of anyone who did the deeds he ascribed to Holmes, nor had he even picked up any garbled information originating in any such person. It may nevertheless be, purely by coincidence, that our own world is one of the worlds where the plot of the stories is enacted. Maybe there was a man whom Conan Doyle never heard of whose actual adventures chanced to fit the stories in every detail. Maybe he even was named "Sherlock Holmes." Improbable, incredible, but surely possible! Now consider the name "Sherlock Holmes," as used in the stories. Does the name, so used, refer to the man whom Conan Doyle never heard of? Surely not! It is irrelevant that a homonymous name is used by some people, not including Conan Doyle, to refer to this man. We must distinguish between the homonyms, just as we would distinguish the name of London (England) from the homonymous name of London (Ontario). It is false at our world that the name "Sherlock Holmes," as used in the stories, refers to someone. Yet it is true in the stories that this name, as used in the stories, refers to someone. So we have found something that is true in the stories but false (under our improbable supposition) at one of the worlds where the plot of the stories is enacted.

In order to avoid this difficulty, it will be helpful if we do not think of a fiction in the abstract, as a string of sentences or something of that sort. Rather, a fiction is a story told by a storyteller on a particular occasion. He may tell his tales around the campfire or he may type a manuscript and send it to his publisher, but in either case there is an act of storytelling. Different acts of storytelling, different fictions. When Pierre Menard re-tells *Don Quixote*, that is not the same fiction as Cervantes' *Don Quixote*—not even if they are in the same language and match word for word.<sup>6</sup> (It would have been different if Menard had copied Cervantes' fiction from memory, however; that

<sup>4</sup> Thus far, the account I have given closely follows that of John Heintz, "Reference and Inference in Fiction" (unpublished).

<sup>5</sup> Briefly stated in his addenda to "Naming and Necessity" in *Semantics of Natural Language*, ed. by Gilbert Harman and Donald Davidson (Dordrecht, 1972); and discussed at greater length in an unpublished lecture given at a conference held at the University of Western Ontario in 1973 and on other occasions. My views and Kripke's overlap to some extent. He also stresses what I have called the ambiguity of prefixing and regards the storyteller as engaged in pretence. The conclusions he draws from the present observation, however, differ greatly from mine.

<sup>6</sup> Jorge Luis Borges, "Pierre Menard, Author of the *Quixote*" in *Ficciones* (Buenos Aires, 1944; English translation, New York, 1962).

would not have been what I call an act of storytelling at all.) One act of storytelling might, however, be the telling of two different fictions: one a harmless fantasy told to the children and the censors, the other a subversive allegory simultaneously told to the *cognoscenti*.

Storytelling is pretence. The storyteller purports to be telling the truth about matters whereof he has knowledge. He purports to be talking about characters who are known to him, and whom he refers to, typically, by means of their ordinary proper names. But if his story is fiction, he is not really doing these things. Usually his pretence has not the slightest tendency to deceive anyone, nor has he the slightest intent to deceive. Nevertheless he plays a false part, goes through a form of telling known fact when he is not doing so. This is most apparent when the fiction is told in the first person. Conan Doyle pretended to be a doctor named Watson, engaged in publishing truthful memoirs of events he himself had witnessed. But the case of third-person narrative is not essentially different. The author purports to be telling the truth about matters he has somehow come to know about, though how he has found out about them is left unsaid. That is why there is a pragmatic paradox akin to contradiction in a third-person narrative that ends "... and so none were left to tell the tale."

The worlds we should consider, I suggest, are the worlds where the fiction is told, but as known fact rather than fiction. The act of storytelling occurs, just as it does here at our world; but there it *is* what here it falsely purports to be: truth-telling about matters whereof the teller has knowledge.<sup>7</sup> Our own world cannot be such a world; for if it is really a fiction that we are dealing with, then the act of storytelling at our world was not what it purported to be. It does not matter if, unbeknownst to the author, our world is one where his plot is enacted. The real-life Sherlock Holmes would not have made Conan Doyle any less of a pretender, if Conan Doyle had never heard of him. (This real-

life Holmes might have had his real-life Watson who told true stories about the adventures he had witnessed. But even if his memoirs matched Conan Doyle's fiction word for word they would not be the same stories, any more than Cervantes' *Don Quixote* is the same story as Menard's. So our world would still not be one where the Holmes stories—the *same* Holmes stories that Conan Doyle told as fiction—were told as known fact.) On the other hand, any world where the story is told as known fact rather than fiction must be among the worlds where the plot of the story is enacted. Else its enactment could be neither known nor truly told of.

I rely on a notion of trans-world identity for stories; this is partly a matter of word-for-word match and partly a matter of trans-world identity (or perhaps a counterpart relation) for acts of storytelling. Here at our world we have a fiction *f*, told in an act *a* of storytelling; at some other world we have an act *a'* of telling the truth about known matters of fact; the stories told in *a* and *a'* match word for word, and the words have the same meaning. Does that mean that the other world is one where *f* is told as known fact rather than fiction? Not necessarily, as the case of Menard shows. It is also required that *a* and *a'* be the same act of storytelling (or at least counterparts). How bad is this? Surely you would like to know more about the criteria of trans-world identity (or the counterpart relation) for acts of storytelling, and so indeed would I. But I think we have enough of a grip to make it worthwhile going on. I see no threat of circularity here, since I see no way of using the concept of truth in fiction to help with the analysis of trans-world identity of acts of storytelling.

Suppose a fiction employs such names as "Sherlock Holmes." At those worlds where the same story is told as known fact rather than fiction, those names really are what they here purport to be: ordinary proper names of existing characters known to the storyteller. Here at our world, the storyteller only pretends that "Sherlock Holmes" has

<sup>7</sup> There are exceptions. Sometimes the storyteller purports to be uttering a mixture of truth and lies about matters whereof he has knowledge, or ravings giving a distorted reflection of the events, or the like. Tolkien explicitly purports to be the translator and editor of the Red Book of Westmarch, an ancient book that has somehow come into his possession and that he somehow knows to be a reliable record of the events. He does not purport to be its author, else he would not write in English. (Indeed, the composition of the Red Book by several hobbits is recorded in the Red Book itself.) I should say the same about a first-person historical novel written in English in which the narrator is an ancient Greek. The author does not pretend to be the truthful narrator himself, but rather pretends to be someone of our time who somehow has obtained the Greek narrator's story, knows it to be true, and passes it on to us in translation. In these exceptional cases also, the thing to do is to consider those worlds where the act of storytelling really is whatever it purports to be—ravings, reliable translation of a reliable source, or whatever—here at our world. I shall omit mention of these exceptional cases in the remainder of this paper.

the semantic character of an ordinary proper name. We have no reason at all to suppose that the name, as used here at our world, really does have that character. As we use it, it may be very unlike an ordinary proper name. Indeed, it may have a highly non-rigid sense, governed largely by the descriptions of Holmes and his deeds that are found in the stories. That is what I suggest: the 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. Part of that role, of course, is to bear the ordinary proper name "Sherlock Holmes". But that only goes to show that "Sherlock Holmes" is used at  $w$  as an ordinary proper name, not that it is so used here.<sup>8, 9</sup>

I also suggest, less confidently, that whenever a world  $w$  is not one of the worlds just considered, the sense of "Sherlock Holmes" as we use it is such as to assign it no denotation at  $w$ . That is so even if the plot of the fiction is enacted by inhabitants of  $w$ . If we are right that Conan Doyle told the Holmes stories as fiction, then it follows that "Sherlock Holmes" is denotationless here at our world. It does not denote the real-life Sherlock Holmes whom Conan Doyle never heard of, if such there be.

We have reached a proposal I shall call ANALYSIS 0: *A sentence of the form "In fiction  $f$ ,  $\phi$ " is true iff  $\phi$  is true at every world where  $f$  is told as known fact rather than fiction.*

Is that right? There are some who never tire of telling us not to read anything into a fiction that is not there explicitly, and Analysis 0 will serve to capture the usage of those who hold this view in its most extreme form. I do not believe, however, that such a usage is at all common. Most of us are content to read a fiction against a background of well-known fact, "reading into" the fiction content that is not there explicitly but that comes jointly from the explicit content and the factual background. Analysis 0 disregards the background. Thereby it brings too many possible worlds into

consideration, so not enough comes out true in the fiction.

For example, I claim that in the Holmes stories, Holmes lives nearer to Paddington Station than to Waterloo Station. A glance at the map will show you that his address in Baker Street is much nearer to Paddington. Yet the map is not part of the stories; and so far as I know it is never stated or implied in the stories themselves that Holmes lives nearer to Paddington. There are possible worlds where the Holmes stories are told as known fact rather than fiction which differ in all sorts of ways from ours. Among these are worlds where Holmes lives in a London arranged very differently from the London of our world, a London where Holmes's address in Baker Street is much closer to Waterloo Station than to Paddington.

(I do not suppose that such a distortion of geography need prevent the otherworldly places there called "London," "Paddington Station," . . . from being the same as, or counterparts of, their actual namesakes. But if I am wrong, that still does not challenge my claim that there are worlds where the stories are told as known fact but where it is true that Holmes lives closer to Waterloo than to Paddington. For it is open to us to regard the place-names, as used in the stories, as fictional names with non-rigid senses like the non-rigid sense I have already ascribed to "Sherlock Holmes." That would mean, incidentally, that "Paddington Station," as used in the stories, does not denote the actual station of that name.)

Similarly, I claim that it is true, though not explicit, in the stories that Holmes does not have a third nostril; that he never had a case in which the murderer turned out to be a purple gnome; that he solved his cases without the aid of divine revelation; that he never visited the moons of Saturn; and that he wears underpants. There are bizarre worlds where the Holmes stories are told as known fact but where all of these things are false.

Strictly speaking, it is fallacious to reason from a mixture of truth in fact and truth in fiction to con-

<sup>8</sup> A rather similar treatment of fictional names, different from mine in that it allows the actual and purported meanings of "Sherlock Holmes" to be the same, is given in Robert Stalnaker, "Assertion" (unpublished).

<sup>9</sup> Many of us have never read the stories, could not produce the descriptions that largely govern the non-rigid sense of "Sherlock Holmes," yet use this name in just the same sense as the most expert Baker Street Irregular. There is no problem here. Kripke's causal picture of the contagion of meaning, in "Naming and Necessity" (*op. cit.*), will do as well for non-rigid senses, as for rigid ones. The ignoramus uses "Sherlock Holmes" in its standard non-rigid sense if he has picked it up (in the right way) from someone who knew the governing descriptions, or who picked it up from someone else who knew them, or . . . Kripke's doctrines of rigidity could not be defended without the aid of his doctrine of contagion of meaning; contagion without rigidity, on the other hand, seems unproblematic.

clusions about truth in fiction. From a mixture of prefixed and unprefixd premisses, nothing follows. But in practice the fallacy is often not so bad. The factual premisses in mixed reasoning may be part of the background against which we read the fiction. They may carry over into the fiction, not because there is anything explicit in the fiction to make them true, but rather because there is nothing to make them false. There is nothing in the Holmes stories, for instance, that gives us any reason to bracket our background knowledge of the broad outlines of London geography. Only a few details need changing—principally details having to do with 221B Baker Street. To move the stations around, or even to regard their locations as an open question, would be uncalled for. What's true in fact about their locations is true also in the stories. Then it is no error to reason from such facts to conclusions about what else is true in the stories.

You've heard it all before. Reasoning about truth in fiction is very like counterfactual reasoning. We make a supposition contrary to fact—what if this match had been struck? In reasoning about what would have happened in that counterfactual situation, we use factual premisses. The match was dry, there was oxygen about, and so forth. But we do not use factual premisses altogether freely, since some of them would fall victim to the change that takes us from actuality to the envisaged counterfactual situation. We do not use the factual premiss that the match was inside the matchbox at the time in question, or that it was at room temperature a second later. We depart from actuality as far as we must to reach a possible world where the counterfactual supposition comes true (and that might be quite far if the supposition is a fantastic one). But we do not make gratuitous changes. We hold fixed the features of actuality that do not have to be changed as part of the least disruptive way of making the supposition true. We can safely reason from the part of our factual background that is thus held fixed.

By now, several authors have treated counterfactual conditionals along the lines just sketched. Differences of detail between these treatments are unimportant for our present purposes. My own version<sup>10</sup> runs as follows. A counterfactual of the form "If it were that  $\phi$ , then it would be that  $\psi$ " is non-vacuously true iff some possible world where both  $\phi$  and  $\psi$  are true differs less from our actual world, on balance, than does any world where  $\phi$  is

true but  $\psi$  is not true. It is vacuously true iff  $\phi$  is true at no possible worlds. (I omit accessibility restrictions for simplicity.)

Getting back to truth in fiction, recall that the trouble with Analysis 0 was that it ignored background, and thereby brought into consideration bizarre worlds that differed gratuitously from our actual world. A fiction will in general require some departures from actuality, the more so if it is a fantastic fiction. But we need to keep the departures from actuality under control. It is wrong, or at least eccentric, to read the Holmes stories as if they might for all we know be taking place at a world where three-nostrilled detectives pursue purple gnomes. The remedy is, roughly speaking, to analyze statements of truth in fiction as counterfactuals. What is true in the Sherlock Holmes stories is what would be true if those stories were told as known fact rather than fiction.

Spelling this out according to my treatment of counterfactuals, we have ANALYSIS 1: *A sentence of the form "In the fiction  $f$ ,  $\phi$ " is non-vacuously true iff some world where  $f$  is told as known fact and  $\phi$  is true differs less from our actual world, on balance, than does any world where  $f$  is told as known fact and  $\phi$  is not true. It is vacuously true iff there are no possible worlds where  $f$  is told as known fact.* (I postpone consideration of the vacuous case.)

We sometimes speak of *the* world of a fiction. What is true in the Holmes stories is what is true, as we say, "in the world of Sherlock Holmes." That we speak this way should suggest that it is right to consider less than all the worlds where the plot of the stories is enacted, and less even than all the worlds where the stories are told as known fact. "In the world of Sherlock Holmes," as in actuality, Baker Street is closer to Paddington Station than to Waterloo Station and there are no purple gnomes. But it will not do to follow ordinary language to the extent of supposing that we can somehow single out a single one of the worlds where the stories are told as known fact. Is the world of Sherlock Holmes a world where Holmes has an even or an odd number of hairs on his head at the moment when he first meets Watson? What is Inspector Lestrade's blood type? It is absurd to suppose that these questions about the world of Sherlock Holmes have answers. The best explanation of that is that the worlds of Sherlock Holmes are plural, and the questions have different answers at different ones. If we may assume that

<sup>10</sup> Given in *Counterfactuals* (Oxford, 1973).

some of the worlds where the stories are told as known fact differ least from our world, then these are the worlds of Sherlock Holmes. What is true throughout them is true in the stories; what is false throughout them is false in the stories; what is true at some and false at others is neither true nor false in the stories. Any answer to the silly questions just asked would doubtless fall in the last category. It is for the same reason that the chorus of Sir Joseph Porter's sisters and cousins and aunts has no determinate size: it has different sizes at different ones of the worlds of *H.M.S. Pinafore*.<sup>11</sup>

Under Analysis 1, truth in a given fiction depends on matters of contingent fact. I am not thinking of the remote possibility that accidental properties of the fiction might somehow enter into determining which are the worlds where that fiction is told as known fact. Rather, it is a contingent matter which of those worlds differ more from ours and which less, and which (if any) differ least. That is because it is a contingent fact—indeed it is *the* contingent fact on which all others depend—which possible world is our actual world. To the extent that the character of our world carries over into the worlds of Sherlock Holmes, what is true in the stories depends on what our world is like. If the stations of London had been differently located, it might have been true in the stories (and not because the stories would then have been different) that Holmes lived nearer to Waterloo Station than to Paddington Station.

This contingency is all very well when truth in fiction depends on well-known contingent facts about our world, as it does in the examples I have so far given to motivate Analysis 1. It is more disturbing if truth in fiction turns out to depend on contingent facts that are not well known. In an article setting forth little-known facts about the movement of snakes, Carl Gans has argued as follows:

In "The Adventure of the Speckled Band" Sherlock Holmes solves a murder mystery by showing that the victim has been killed by a Russell's viper that has climbed up a bell rope. What Holmes did not realize was that Russell's viper is not a constrictor. The snake is therefore

incapable of concertina movement and could not have climbed the rope. Either the snake reached its victim some other way or the case remains open.<sup>12</sup>

We may well look askance at this reasoning. But if Analysis 1 is correct then so is Gans's argument. The story never quite says that Holmes was right that the snake climbed the rope. Hence there are worlds where the Holmes stories are told as known fact, where the snake reached the victim some other way, and where Holmes therefore bungled. Presumably some of these worlds differ less from ours than their rivals where Holmes was right and where Russell's viper is capable of concertina movement up a rope. Holmes's infallibility, of course, is not a countervailing resemblance to actuality; our world contains no infallible Holmes.

Psychoanalysis of fictional characters provides a more important example. The critic uses (what he believes to be) little-known facts of human psychology as premisses, and reasons to conclusions that are far from obvious about the childhood or the adult mental state of the fictional character. Under Analysis 1 his procedure is justified. Unless countervailing considerations can be found, to consider worlds where the little-known fact of psychology does not hold would be to depart gratuitously from actuality.

The psychoanalysis of fictional characters has aroused vigorous objections. So would Gans's argument, if anyone cared. I shall keep neutral in these quarrels, and try to provide for the needs of both sides. Analysis 1, or something close to it, should capture the usage of Gans and the literary psychoanalysts. Let us find an alternative analysis to capture the conflicting usage of their opponents. I shall not try to say which usage is more conducive to appreciation of fiction and critical insight.

Suppose we decide, *contra* Gans and the literary psychoanalysts, that little-known or unknown facts about our world are irrelevant to truth in fiction. But let us not fall back to Analysis 0; it is not our only alternative. Let us still recognize that it is perfectly legitimate to reason to truth in fiction from a background of well-known facts.

Must they really be facts? It seems that if little-known or unknown facts are irrelevant, then so are

<sup>11</sup> Heintz (*op. cit.*) disagrees; he supposes that for each fiction there is a single world to be considered, but a world that is in some respects indeterminate. I do not know what to make of an indeterminate world, unless I regard it as a superposition of all possible ways of resolving the indeterminacy—or, in plainer language, as a set of determinate worlds that differ in the respects in question.

<sup>12</sup> Carl Gans, "How Snakes Move," *Scientific American*, vol. 222 (1970), p. 93.

little-known or unknown errors in the body of shared opinion that is generally taken for fact. We think we all know that there are no purple gnomes, but what if there really are a few, unknown to anyone except themselves, living in a secluded cabin near Loch Ness? Once we set aside the usage given by Analysis 1, it seems clear that whatever purple gnomes may be hidden in odd corners of our actual world, there are still none of them in the worlds of Sherlock Holmes. We have shifted to viewing truth in fiction as the joint product of explicit content and a background of generally prevalent beliefs.

Our own beliefs? I think not. That would mean that what is true in a fiction is constantly changing. Gans might not be right yet, but he would eventually become right about Holmes's error if enough people read his article and learned that Russell's viper could not climb a rope. When the map of Victorian London was finally forgotten, it would cease to be true that Holmes lived nearer to Paddington than to Waterloo. Strange to say, the historical scholar would be in no better position to know what was true in the fictions of his period than the ignorant layman. That cannot be right. What was true in a fiction when it was first told is true in it forevermore. It is our knowledge of what is true in the fiction that may wax or wane.

The proper background, then, consists of the beliefs that generally prevailed in the community where the fiction originated: the beliefs of the author and his intended audience. And indeed the factual premisses that seemed to us acceptable in reasoning about Sherlock Holmes were generally believed in the community of origin of the stories. Everyone knew roughly where the principal stations of London were, everyone disbelieved in purple gnomes, and so forth.

One last complication. Suppose Conan Doyle was a secret believer in purple gnomes; thinking that his belief in them was not shared by anyone else he kept it carefully to himself for fear of ridicule. In particular, he left no trace of this belief in his stories. Suppose also that some of his original readers likewise were secret believers in purple gnomes. Suppose, in fact, that everyone alive at the time was a secret believer in purple gnomes, each thinking that his own belief was not shared by anyone else. Then it is clear (to the extent that

anything is clear about such a strange situation) that the belief in purple gnomes does not "generally prevail" in quite the right way, and there are still no purple gnomes in the worlds of Sherlock Holmes. Call a belief *overt* in a community at a time iff more or less everyone shares it, more or less everyone thinks that more or less everyone else shares it, and so on.<sup>13</sup> The proper background, we may conclude, comprises the beliefs that are overt in the community of origin of the fiction.

Assume, by way of idealization, that the beliefs overt in the community are each possible and jointly compossible. Then we can assign to the community a set of possible worlds, called the *collective belief worlds* of the community, comprising exactly those worlds where the overt beliefs all come true. Only if the community is uncommonly lucky will the actual world belong to this set. Indeed, the actual world determines the collective belief worlds of the community of origin of the fiction and then drops out of the analysis. (It is of course a contingent matter what that community is and what is overtly believed there.) We are left with two sets of worlds: the worlds where the fiction is told as known fact, and the collective belief worlds of the community of origin. The first set gives the content of the fiction; the second gives the background of prevalent beliefs.

It would be a mistake simply to consider the worlds that belong to both sets. Fictions usually contravene at least some of the community's overt beliefs. I can certainly tell a story in which there are purple gnomes, though there are none at our collective belief worlds. Further, it will usually be overtly believed in the community of origin of a fiction that the story is not told as known fact—storytellers seldom deceive—so none of the worlds where the fiction is told as known fact can be a collective belief world of the community. Even if the two sets do overlap (the fiction is plausible and the author palms it off as fact) the worlds that belong to both sets are apt to be special in ways having nothing to do with what is true in the fiction. Suppose the story tells of a bungled burglary in recent times, and suppose it ends just as the police reach the scene. Any collective belief world of ours where this story is told as known fact is a world where the burglary was successfully covered up; for it is an overt belief among us that no such

<sup>13</sup> A better definition of overt belief, under the name of "common knowledge" may be found in my *Convention* (Cambridge, Mass., 1969), pp. 52–60. That name was unfortunate, since there is no assurance that it will be knowledge, or even that it will be true. See also the discussion of "mutual knowledge" in Stephen Schiffer, *Meaning* (Oxford, 1972), pp. 30–42.

burglary ever hit the news. That does not make it true in the story that the burglary was covered up.

What we need is something like Analysis 1, but applied from the standpoint of the collective belief worlds rather than the actual world. What is true in the Sherlock Holmes stories is what would be true, according to the overt beliefs of the community of origin, if those stories were told as known fact rather than fiction.

Spelling this out, we have ANALYSIS 2: *A sentence of the form "In the fiction f,  $\phi$ " is non-vacuously true iff, whenever w is one of the collective belief worlds of the community of origin of f, then some world where f is told as known fact and  $\phi$  is true differs less from the world w, on balance, than does any world where f is told as known fact and  $\phi$  is not true. It is vacuously true iff there are no possible worlds where f is told as known fact.* It is Analysis 2, or something close to it, that I offer to opponents of Gans and the literary psychoanalysts.

I shall briefly consider two remaining areas of difficulty and sketch strategies for dealing with them. I shall not propose improved analyses, however; partly because I am not quite sure what changes to make, and partly because Analysis 2 is quite complicated enough already.

I have said that truth in fiction is the joint product of two sources: the explicit content of the fiction, and a background consisting either of the facts about our world (Analysis 1) or of the beliefs overt in the community of origin (Analysis 2). Perhaps there is a third source which also contributes: carry-over from other truth in fiction. There are two cases: intra-fictional and inter-fictional.

In the *Threepenny Opera*, the principal characters are a treacherous crew. They constantly betray one another, for gain or to escape danger. There is also a streetsinger. He shows up, sings the ballad of Mack the Knife, and goes about his business without betraying anyone. Is he also a treacherous fellow? The explicit content does not make him so. Real people are not so very treacherous, and even in Weimar Germany it was not overtly believed that they were, so background does not make him so either. Yet there is a moderately good reason to say that he is treacherous: in the *Threepenny Opera*, that's how people are. In the worlds of the *Threepenny Opera*, everyone put to the test proves treacherous, the streetsinger is there along with the rest, so doubtless he too would turn out to be treacherous if we saw more of him. His treacherous nature is an intra-fictional carry-over from the

treacherous natures in the story of Macheath, Polly, Tiger Brown, and the rest.

Suppose I write a story about the dragon Sculch, a beautiful princess, a bold knight, and what not. It is a perfectly typical instance of its stylized genre, except that I never say that Sculch breathes fire. Does he nevertheless breathe fire in my story? Perhaps so, because dragons in that sort of story do breathe fire. But the explicit content does not make him breathe fire. Neither does background, since in actuality and according to our beliefs there are no animals that breathe fire. (It just might be analytic that nothing is a dragon unless it breathes fire. But suppose I never called Sculch a dragon; I merely endowed him with all the standard dragonly attributes except fire-breathing.) If Sculch does breathe fire in my story, it is by inter-fictional carry-over from what is true of dragons in other stories.

I have spoken of Conan Doyle's Holmes stories; but many other authors also have written Holmes stories. These would have little point without inter-fictional carry-over. Surely many things are true in these satellite stories not because of the explicit content of the satellite story itself, and not because they are part of the background, but rather because they carry over from Conan Doyle's original Holmes stories. Similarly, if instead of asking what is true in the entire corpus of Conan Doyle's Holmes stories we ask what is true in "The Hound of the Baskervilles", we will doubtless find many things that are true in that story only by virtue of carry-over from Conan Doyle's other Holmes stories.

I turn finally to vacuous truth in impossible fictions. Let us call a fiction *impossible* iff there is no world where it is told as known fact rather than fiction. That might happen in either of two ways. First, the plot might be impossible. Second, a possible plot might imply that there could be nobody in a position to know or tell of the events in question. If a fiction is impossible in the second way, then to tell it as known fact would be to know its truth and tell truly something that implies that its truth could not be known; which is impossible.

According to all three of my analyses, anything whatever is vacuously true in an impossible fiction. That seems entirely satisfactory if the impossibility is blatant: if we are dealing with a fantasy about the troubles of the man who squared the circle, or with the worst sort of incoherent time-travel story. We should not expect to have a non-trivial concept of truth in blatantly impossible fiction, or perhaps we should expect to have one only under the

pretence—not to be taken too seriously—that there are impossible possible worlds as well as the possible possible worlds.

But what should we do with a fiction that is not blatantly impossible, but impossible only because the author has been forgetful? I have spoken of truth in the Sherlock Holmes stories. Strictly speaking, these (taken together) are an impossible fiction. Conan Doyle contradicted himself from one story to another about the location of Watson's old war wound. Still, I do not want to say that just anything is true in the Holmes stories!

I suppose that we might proceed in two steps to say what is true in a venially impossible fiction such as the Holmes stories. First, go from the original impossible fiction to the several possible revised versions that stay closest to the original. Then say that what is true in the original is what is true, according to one of our analyses of non-vacuous truth in fiction, in all of these revised versions. Then nothing definite will be true in the Holmes stories about the location of Watson's wound. Since Conan Doyle put it in different places, the different revised versions will differ. But at least it will be true in the stories that Watson was wounded elsewhere than in the left big toe. Conan Doyle put the wound in various places, but never there. So no revised version will put the wound in the left big toe, since that would change the story more than consistency demands.

The revised versions, like the original fiction, will be associated with acts of storytelling. The

revised versions, unlike the original, will not actually be told either as fiction or as known fact. But there are worlds where they are told as fiction, and worlds where they are told as known fact.

Even when the original fiction is not quite impossible, there may be cases in which it would be better to consider not truth in the original fiction but rather truth in all suitably revised versions. We have a three-volume novel set in 1878. We learn in the first volume that the hero had lunch in Glasgow on a certain day. In the third volume, it turns out that he showed up in London that same afternoon. In no other way does this novel purport to be a fantasy of rapid transit. The author was just careless. We could without vacuity apply our analyses directly to the novel as written. Since the closest worlds where it is told as known fact are worlds with remarkable means of travel, the results would astonish anyone—for instance, our forgetful author—who had not troubled to work out a careful timetable of the hero's movements. It would be more charitable to apply the analyses not to the original story but instead to the minimally revised versions that make the hero's movements feasible by the means of travel that were available in 1878. At least, that would be best if there were ways to set the times right without major changes in the plot. There might not be, and in that case perhaps truth in the original version—surprising though some of it may be—is the best we can do.<sup>14</sup>

*Princeton University*

*Received April 30, 1976*

<sup>14</sup> I thank the many friends and colleagues who have given me helpful comments on a previous version of this paper, and I thank the American Council of Learned Societies for research support. Special thanks are due to John G. Bennett and Saul Kripke for valuable discussions.