Fictional Reports

Paal Fjeldvig Antonsen

Abstract This paper outlines a bicontextual account of fictional reports. A fictional report is a report on something that happens in a fiction, and a bicontextual account is an account that relativizes truth to two contexts. The proposal is motivated by two considerations. First, it explains the intuitive truth conditions of fictional reports without postulating hidden fiction operators. Second, it handles the problem of indexicals in fictional reports better than the standard accounts.

Keywords Fictional reports · Truth in fiction · Indexicality · Bicontextualism · Context of imagination

1 Introduction

I'll begin with a simple observation: the sentence

(1) Alexander the Great was poisoned.

has at least two readings, one false and one true. If (1) is used to make a claim about ancient history, it is false. The consensus among historians, at least, is that Alexander died of typhoid fever or some other natural cause. If (1) is stated during a conversation about Oliver Stone's movie *Alexander*, however, those of us familiar with the story will accept it as true. At a critical point in the narrative, Alexander's closest companions successfully carry out a conspiracy to poison him to death. But how can we really accept the statement as *true*, since Alexander did not in fact get poisoned? It's obvious enough that we're assessing a fictional report, that is a report on some event taking place in the fiction. What's not so obvious is why that makes a difference when it comes to the sentence's truth value.

I want to suggest that the difference comes down to which context we are assessing the sentence from. Broadly speaking, a context of assessment is any situation from which one might potentially assess a statement. When we treat (1) as a historical report, the appropriate

Paal Fjeldvig Antonsen
Department of Philosophy, University of Bergen
E-mail: pal.antonsen@uib.no

context of assessment is the given situation of utterance. It's just a straight up claim about how Alexander met his demise, one that happens to get things wrong. When we treat (1) as a fictional report, on the other hand, the appropriate context of assessment is an imagined situation prescribed by the fiction. Now it becomes a claim about what the movie has invited us to imagine, one that happens to get things right.

The account I'm outlining here is *bicontextual*, in the sense that it makes use of a semantic framework that relativizes truth to both a context of use and a context of assessment. This kind of framework is familiar from MacFarlane's [16] analysis of assessment-sensitive expressions. According to my proposal, the same technical resources can be used to handle fictional reports as well. In what follows, I'll be spending most of my time spelling out two motivations for the bicontextual account. First, unlike the standard accounts, the bicontextual account captures the intuitive truth conditions of fictional reports without postulating hidden fiction operators. Second, it gives an explanation of the behaviour we perceive of indexials in fictional reports, while it is at best unclear how such data can be explained by some prominent rival accounts.

2 Standard account

According to what I'm going to call 'the standard account', the distinguishing mark of fictional reports is that they contain a hidden element.[1] When (1) is used to talk about Stone's movie, the thought goes, it is merely a compressed way of saying what would have been expressed by uttering (2).

(2) In the fiction, Alexander the Great was poisoned.

The reason we willingly give our assent to the fictional report (1) is that our intuitions about its accuracy are really directed at (2). The sentence (2), in turn, decomposes into a fiction operator and an embedded sentence. In order to explain our intuitive assessment of fictional reports, then, we need only come up with a plausible semantics for this hidden operator.[2]

It is at this point that Lewis' famous proposal enters the picture: the fiction operator can be understood as 'a restricted universal quantifier over possible worlds' [15, 35]. On this proposal, we treat the semantic value of a sentence embedded under the fiction operator as something that can take different truth values at different possible worlds. The meaning of 'in the fiction' can then be modelled as a function from these semantic values to truth values, mapping the semantic value of a sentence ϕ to truth just in case ϕ is true at every world compatible with the information provided by the topical fiction. On this line of thought, the following seems like an attractive semantics:

$$[\![\text{In the fiction } \phi]\!]^c_{w,s,g} = 1 \text{ iff } \forall w' \in s : [\![\phi]\!]^c_{w',s,g} = 1.$$

Here c is a context of use, w is a possible world, g is an assignment function that maps variables to objects in the domain of the world of evaluation (used to handle quantifiers), and s is the set of worlds compatible with the topical fiction.[3] Taking this semantics on

---

[1] This paper is exclusively concerned with the tradition that treats fictional reports as genuine *statements* — something that may be true or false. By 'the standard account', I mean the position that most discussions within this tradition use as their starting point. Another tradition, inspired by Walton [32], treats fictional reports as a kind of utterances under pretense that are neither true nor false.

[2] See e.g. [9, 471], [15, 37-38], [8, 174], [6, 184], [4, 24], [13, 83], [22, 287], [12, 241], [10, 62].

[3] Borrowing a term from Yalcin [33], I have chosen a domain semantics for the fiction operator. That is not essential, I just happen to think it makes for a smoother presentation than representing 'in the fiction' as a relational operator [25], or using fiction-specifying operators [15].

board, the sentence (2) has the following truth conditions: $[\![(2)]\!]^c_{w,s,g} = 1$ if and only if, for every $w' \in s$, Alexander was poisoned in w'.

I think there is a lot to be said for a Lewisian-style semantics for 'in the fiction'.[4] In particular, it's not difficult to see how it can be made to fit with the view that fictions invite the audience to engage in imaginings about people, situations and events.[5] It seems plausible, for example, that when we watch Stone's movie, one of the things we're required to imagine is that (1). The propositional content of this imagining can be represented by the set of worlds $\{w : \text{Alexander was poisoned in } w\}$. Now, if we take the intersection of all the propositional contents that Stone's movie invites us to imagine ($p_1 \cap \ldots \cap p_n$), what we get is the set of all the worlds compatible with the information conveyed by the fiction. I'm going call this 'the fictional worlds'. So when someone makes a fictional report, they are in effect reporting on the goings-on in the fictional worlds. The truth conditions associated with (2) becomes more plausible if we add all this to the background. The phrase 'for every $w \in s$, Alexander was poisoned in w', for example, becomes just another way of saying that Alexander was poisoned in all the worlds the movie invites us to imagine.

The last part of the standard account is a description of how we go about evaluating particular statements. The idea is that when we judge a statement's veracity, what we're really assessing is whether it is true relative to its contexts of use, defined as follows:

Truth at a context of use
A sentence ϕ is true at a context of use c iff for any assignment g : $[\![\phi]\!]^c_{w_c,s,g} = 1$, where w_c is the world of c.

To make things more readable, I'll drop the qualifier 'at the context of use' for the most part. The standard account delivers the desired result. The historical report (1) comes out as true just in case Alexander was poisoned in w_c (which I'm going to assume is the actual world). Since Alexander was in fact not poisoned, (1) is false on that reading. The fictional report (1), on the other hand, is true if and only if the sentence (2) is true, and (2) is true just in case Alexander was poisoned in all the fictional worlds. Since that is indeed what the fiction asks us to imagine, (1) comes out as true on that reading.

A recurring objection to the standard account is that it's hard to see how 'in the fiction' gets to be part of a sentence's logical form.[6] For one thing, there doesn't seem to be much syntactic evidence to speak of backing it up. Here's how Bertolet expresses the worry:

> Now it is true that linguists who support the so-called performative hypothesis have postulated such a relation between e.g. 'John loves Mary' and 'I assert that John loves Mary'... But this analysis is motivated by *syntactic* considerations, without which it would be manifestly implausible, and there is not one whit of syntactic evidence of which I am aware that would motivate a similar relation between 'Santa Claus flies reindeer' and 'Some person/story said Santa Claus flies reindeer'. [2, 150]

It's also far from obvious what other mechanism would be in play. I suppose one could try to appeal to some process of free enrichment. One might say, for example, that although 'in the fiction' is left out for reasons having to do with communicative efficiency, it can

---

[4] Lamarque & Olsen go even further: 'The notion of 'in the fiction...' as an intensional operator governing reports of fictive content is now so common in theoretical writing that appeal to it does not need special defence.' [13, 83]

[5] I am using 'fictions invite the audience' as shorthand for 'when an author performs an act of fiction-making, they invite the audience to imagine the information conveyed'. Cf. [7], [32], [29].

[6] See e.g. [1], [2], [23] [24], [3], [26], [5].

be introduced through extraction from information available in the immediate utterance environment. It is often transparent when we are talking about Stone's movie, and in those circumstances we can reliably count on our interlocutors' ability to understand that by *saying* (1), what we're trying to convey is the qualified claim (2). I am, admittedly, skeptical to this line of thought, but let's just grant it for the sake of argument. It's enough for my purposes to register that the standard account comes with a theoretical cost.

Note that I am not trying to cast doubt on the Lewisian semantics for the fiction operator as such, nor the idea that there are fiction operators in English. There are indeed many instances where we do use phrases such as 'in the fiction' and 'according to the fiction', and a full-fledged account of fiction discourse would have to include a description of what these phrases contribute to the semantic value of sentences in which they can occur. The comment about the standard account's theoretical cost concerns the postulation of hidden fiction operators, and the need for an accompanying mechanism that explains how sentences that do not have a fiction operator in their surface structure still contain such elements in their logical form.

3 Indexicals

There is, I think, a basic flaw with the standard account: it can't explain the behaviour we perceive of indexicals in fictional reports. The standard account's shortcomings in this respect have previously been noted by Predelli [25, 283], using (3) as an example.

(3) All Japanese planes could have survived the raid.

To a good approximation, the logical form of (3) is (4).

(4) Possibly $\forall x$ (Actually (x is a japanese plane) $\rightarrow$ x survive)

Here Possibly is the familiar modal operator and Actually is an indexical operator that shifts to the actual world, no matter how deeply embedded.

$$[\![\text{Possibly } \phi]\!]^c_{w,s,g} = 1 \text{ iff } \exists w' \text{ accessible from } w : [\![\phi]\!]^c_{w',s,g} = 1.$$
$$[\![\text{Actually } \phi]\!]^c_{w,s,g} = 1 \text{ iff } [\![\phi]\!]^c_{w_c,s,g} = 1.$$

Like our first example, (4) has two relevant readings. It can be read as a claim about the history of the Second World War, in which case the standard account rightly predicts that $[\![(4)]\!]^c_{w,s,g} = 1$ just in case there is at least one world w' accessible from w, such that all the Japanese planes in w_c survived in w'. But (4) can also be read as a claim about Michael Bay's movie *Pearl Harbour*, where it intuitively says something true just in case all the planes in the fictional worlds could have survived the attack. This is a problem for the standard account, because it requires that the semantic value of 'is a japanese plane' is drawn from the domain of the actual world.

Let's flesh out the problem a bit. According to the standard account, when (3) is used as a fictional report, the sentence to be evaluated is (5).

(5) In the fiction Possibly $\forall x$ (Actually (x is a japanese plane) $\rightarrow$ x suvive)

This sentence has the following truth conditions: $[\![(5)]\!]^c_{w,s,g} = 1$ just in case, for every fictional world w', there is a world w'' accessible from w', such that if x is a Japanese plane in w_c then x survived in w''. That is, after the fiction operator has shifted to the fictional worlds w', Possibly shifts to an accessible world w''. But then Actually takes us back to the actual world w_c for the part 'is a japanese plane', and we end up picking out actual planes.

The planes in w_c, however, are not the ones supposed to be surviving in w'', on the intended reading — it is supposed to be the fictional planes in w'. So the standard account doesn't adequately depict the fictional report's intuitive truth conditions.

The problem with indexicals becomes even more transparent when we turn our attention to "pure indexicals" in reports about interactive fictions, such as video games, adventure books and roleplaying games. On the standard semantics, the first person pronoun picks out the speaker:

$$\llbracket I \rrbracket^c_{w,s,g} = \text{the agent of } c.$$

It is a commonplace observation that people typically talk about their avatars (the characters they play as) in a first person way. When we report on the goings-on in interactive fictions, then, we are free to use the first-person pronoun to refer to our avatars, as well as temporal indexicals to pick out points in the fictional time-line, and locational indexicals to pick out fictional places. Consider, for example, the following scenario.

Scenario 1
Teddy is playing the space opera video game *Mass Effect*. His avatar, Jane Shepard, is trying to romance Kaidan, another character in the story. After they've enjoyed an intimate moment, they embrace in a kiss. Teddy says to his friend: Hey look,

(6) I kissed Kaidan.

The pronoun in (6) is associated with two readings:[7]

– I_c kissed Kaidan
– I_i kissed Kaidan

This first means just what it would if it had been used in an ordinary conversation: it is true just in case the speaker, Teddy, kissed Kaidan. On the second it means what it is used in Scenario 1 to express: something that is true just in case Jane, the person who according to the story, kissed Kaidan. The problem, of course, is that the second reading is unavailable if fictional reports only differ from non-fictional reports by virtue of containing a hidden 'in the fiction' element. Whether we prefix (6) with a fiction operator or not, the context of use remains the same. The operator shifts to the fictional worlds, but the semantics for the first person pronoun is insensitive to this alteration. The standard account of the fictional report (6), then, leads to the odd prediction that Teddy has said something true just in case the fiction invites the audience to imagine that *Teddy* kissed Kaidan. That can't be right, though, since Teddy himself is never even mentioned in the story of *Mass Effect*.

In their discussion of interactive fictions, Robson & Meskin say that such fictions are 'in virtue of their interactive nature... about those who consume them.' [27, 165] In some sense, then, Teddy should be engaged in an act of self-involving imagining when he is playing *Mass Effect* in an appreciative manner. I find their suggestion persuasive, but the self-involving character of the imaginings associated with interactive fictions does not present an obstacle to the view that 'I' refers to Jane in (6). As an analogy, Robson & Meskin ask us to consider the difference between imagining that one is flying over Metropolis, like Superman, and imagining being Superman flying over Metropolis. It is the second kind of "imagination from the inside" that Teddy is engaged in when he plays *Mass Effect*. Teddy, then, is not imagining that his actual self is kissing Kaidan. Rather, he is imagining the possibility of being Jane, engaged in an act of intimate embrace. So the referent of 'I' should still be Jane.

---

[7] Roughly, 'I_c' picks out the speaker, and 'I_i' picks out the avatar.

It is just that Teddy also imagines himself being the heroic protagonist. If Robson & Meskin have identified the proper sense in which interactive fictions are about those who consume them, the self-involving nature of video games does not mean that the first person pronoun should pick out the speaker in the actual world.

4 Context of Imagination

A straightforward diagnosis of what goes wrong in Scenario 1 is that we are using the wrong context. As long as we maintain the assumption that the first person pronoun picks out the agent of the context, Jane has to be that agent. Since Jane doesn't report on the fiction herself, we might be better off evaluating the fictional report (6) relative to another context than the context of use.

Predellli [23] [24], Reimer [26] and Voltolini [30] [31] have all made suggestions along these lines. Predelli, for example, says that when we make a fictional report we are 'talking "from the point of view of the story" … to be represented by pairs containing a context other than the context of utterance' [23, 73-4]. Following the same line of thought, Voltolini writes:

> [W]henever the context relevant for the semantic interpretation of an indexical sentence is not its proper context of utterance, that context is… a fictional context i.e., a context which has at least one fictional parameter: a *pretended* agent, or a *pretended* space, or a *pretended* time, or a *pretended* world [30, 27].

What does such a context look like? A context representing the point of view of some fiction obviously won't include the actual world as one of its parameters. There also aren't many fictions, no matter how elaborate, that contain enough information to uniquely pick out another world that deserves to be called *the* world of the fiction either.[8] Fortunately, it's not too hard to see how a context representing a fictional point of view ends up with an appropriate set of worlds. As mentioned, we're working on the assumption that fictions invite the audience to engage in imaginings. We called the collective content of these imaginings for the fictional worlds. When someone makes a fictional report, we can treat them as engaging in some kind of counterfactual speech; they talk as if the fictional worlds were actual. On this line of thought, it seems natural to say that the world parameter of a context representing the point of a view of some fiction should be that fiction's associated fictional worlds. Let's use the term 'context of imagination' for contexts of this kind. So the idea we're running with for the moment is that fictional reports are evaluated relative to a context of imagination, instead of their contexts of use.

An attractive consequence of introducing contexts of imagination is that we can dispense with fiction operators. This is illustrated by how the 'context shift account' handles our original example. Recall that we started out with the observation that (1) could be used to make a fictional report or a historical report. So on the context shift account, (1) is either to be evaluated relative to the context of use or relative to the context of imagination. A semantic theory underwriting this suggestion must therefore define 'truth at a context' without specifying which context is in play:

Truth at a context

---

[8]   This is just another way of saying that fictions are indeterminate. If we happen to live in an indeterminate universe as well, the context of use won't determine a unique world either, only the set of worlds compatible with past events. Cf. (MacFarlane [17])

A sentence ϕ is true at a context c iff for any assignment g and for every world $w \in W(c) : [\![\phi]\!]_{w,g}^{c} = 1$, where $W(c)$ is the set of worlds overlapping at c.

We arrive at the following explanation. If someone utters (1) as a historical report, c is the context of use, and the set of worlds picked out is the singleton set consisting only of the world of utterance ($W(c)$ = {the actual world}). The historical report (1) will then be true just in case Alexander was poisoned in the actual world. If we utter (1) as a fictional report, on the other hand, c is a context of imagination, a context that picks out the set of worlds that Stone's movie asks us to imagine. The fictional report (1) will then be true just in case Alexander was poisoned in all the fictional worlds. Just as the standard account, the context shift account is able to handle garden variety fictional reports.

Introducing contexts of imagination seems to be a step in the right direction, but I'm not too happy with abandoning contexts of use. One reason is that there are fictional reports where the first person pronoun *does* pick out the speaker. The context shift strategy for dealing with Scenario 1 is obviously to evaluate 'I kissed Kaidan' relative to a context of imagination where Jane is the agent. Even if we suppose that it is successful for that particular example, the same strategy doesn't seem to work for the following scenario.

Scenario 2
Tom Cruise is playing the action game *Blasted*, where the story takes place in a universe not so different from our own. His avatar, Bully, is a ruthless thug who ends up killing a lot of people inside a fancy L.A. nightclub. Looking at the corpses on the floor he notices a familiar face. It's Tom Cruise! He yells out: Oh no,

(7) I killed myself!

On the intended reading, Tom Cruise says something true just in case Bully killed Tom Cruise in the fictional worlds. Now, if we evaluate (7) relative to the context of imagination, we get the mistaken result that both pronouns pick out Bully: (7) is true just in case the fiction has asked us to imagine that Bully committed suicide. Of course, things don't fare better if the context of use is the operative one, because then (7) will say that Tom Cruise committed suicide. The intended reading is only available if two first person indexicals occurring in the same sentence can take different semantic values.

Another reason discarding contexts of use might be too radical is that the actual world sometimes plays an important role. Consider the sentence (8).

(8) Had admiral Yamamoto attacked Pearl Harbour at night, he would have lost fewer planes than he actually did.

For the point I want to make, it doesn't really matter what semantics we use for the counterfactual conditional. I'll settle for a familiar one:

$[\![\phi \mathbin{\Box\!\!\rightarrow} \psi]\!]_{w,g}^{c} = 1$ iff at the worlds w', where $[\![\phi]\!]_{w',g}^{c} = 1$, the world w'' most similar to w is one such that $[\![\psi]\!]_{w'',g}^{c} = 1$.

Let's only consider (8) as a fictional report. The thing note is that the sentence is still ambiguous. On the one hand, it can be used to compare the number of planes the Japanese would have lost, had Yamamoto attacked at night, with the number of planes he lost in the historical battle. Maybe it was voiced in an attempt to show off military knowledge. We arrive at this reading if 'actually' shifts to the actual world: (8) is true if and only if, in the counterfactual worlds w' where Yamamoto attacked at night, the number of planes he lost in the world w'' most similar to the actual world w''' is less than the number of planes he lost in w'''. On the other hand, we can also use (8) to compare the counterfactual losses with

the number of planes he lost in the fiction. You may want to express this if you are simply interested in telling someone about an alternative path the movie could have taken. We arrive at this reading if 'actually' shifts to the fictional worlds, instead of the actual world. Notice that if we relativize (8) to the context of imagination exclusively, we only have access to this second reading. If the context of imagination replaces the context of use, and 'actually' shifts to the worlds of the context, then it must shift us to the fictional worlds. There is simply no actual world available for 'actually' to shift to. This is an unhappy result, because it seems perfectly intelligible to compare what would have happened in Bay's movie, had Yamamoto attacked at night, with what really did happen.

5 Bicontextualism

Let's briefly review the discussion so far. I've complained that even if we grant that there are hidden fiction operators, the standard account is unable to handle straightforward examples of indexical fictional reports. As a potential remedy, we considered letting contexts of imagination play the role of contexts of use. The objection against this move was that doing so leaves us with too few resources. In particular, the context shift account couldn't handle examples where the semantic value of an indexical depended on parameters usually reserved for the context of use.

I now want to try a different approach. As a starting point, let's retain the notion that contexts of use represent possible situations in which a sentence can be uttered. So whether someone is using (1) as a historical report or as a fictional report, the sentence is uttered at the same context of use. There is still a role for contexts of imagination to play, however, if we say that the difference between the two types of utterances is a matter of which context the sentence is assessed from. I'll say that the historical report is assessed from the same context in which it was uttered, since we are only considering whether (1) says something that is true in our actual situation. The fictional report, on the other hand, is assessed from the context of imagination, reflecting that our concern in this case is whether (1) says something that is true in the hypothetical situation the movie invites us to imagine. On this suggestion, the fictional report comes out as true because Alexander was poisoned in the imagined scenario.

It seems appropriate to call this suggestion 'the bicontextual account', as it makes use of a semantic framework that relativizes truth to two context. Appropriating some of MacFarlane's [19] terminology, we'll say that the target notion for a semantic theory for fictional reports is not 'truth at a context of use', but 'truth at a context of use and a context of assessment'. The 'context of assessment' is here an imagined situation from which a sentence can be assessed.[9] This doubly relativizied truth predicate is defined as follows:

> **Truth at a context of use and a context of assessment**
> A sentence ϕ is true at a context of utterance c_1 and a context of assessment c_2 iff for any assignment g and for every world $w \in W(c_2) : [\![\phi]\!]_{w,g}^{c_1,c_2} = 1$, where $W(c_2)$ is the set of worlds overlapping at c_2.

How does this apply to fictional reports? It's easiest to start with (1) as an example. When we evaluate (1) as a historical report, the context of utterance and the context of assessment are in alignment ($c_1 = c_2$). So (1) is true just in case Alexander was poisoned in the actual

---

[9] Although I am using MacFarlane's machinery, the context of assessment is playing a different role here. First, a context of assessment is always imagined, so we can never be in one. Second, the context of imagination is tied to the speaker. When another person evaluate a fictional report, they treat it as the speaker's context of assessment, not their own.

world. When we evaluate (1) as a fictional report, however, the two contexts come apart ($c_1 \neq c_2$). The appropriate context of assessment is not merely removed in time and place from the context of use, but wholly imagined; the context of imagination representing the point of view of Stone's movie. Allowing that a context of imagination plays the role of c_2, we get the intuitively right prediction that (1) is true just in case Alexander was poisoned in the fictional worlds. So far the bicontextual account is on a par with the other two accounts we've considered.

Parsons [21] and MacFarlane [19, Ch. 3] have both pointed out that a bicontextual framework lends itself to an assessment sensitive treatment of indexicals. As described above, the context of assessment is relevant for how we evaluate the accuracy of particular statements, as it determines the world parameter. But the context of assessment can also be used in a parallel way to the context of use in the semantics for indexicals and other context-dependent expressions. Parsons, for instance, argues that there are two distinct conventions regulating the temporal indexical 'now'; one according to which 'now' picks out the time of utterance and another according to which 'now' picks out the time of assessment. I'm not taking a stand on whether he is right about our practices of talking about the present time. What concerns us is the general idea that any context dependent expression could in principle have an assessment sensitive use in addition to its normal one.

The idea that indexicals might be ambiguous is not novel[10], but the bicontextual framework opens up an interesting possibility for fiction discourse. For example, it might be that when we talk about a fiction, 'actually' has one use according to which it shifts to the actual world, and another use according to which it shifts to the worlds of assessment. This idea has some intuitive plausibility. The process of evaluating a fictional report requires something like two mental files, one for the actual situation, the other for the imagined situation. When we engage in fiction discourse we routinely keep track of both by shifting between which one we consider actual. The rules for these two uses can be described as follows:

$$\llbracket \text{Actually}_1 \phi \rrbracket_{w,g}^{c_1,c_2} = 1 \text{ iff } \forall w' \in W(c_1) : \llbracket \phi \rrbracket_{w',g}^{c} = 1.$$
$$\llbracket \text{Actually}_2 \phi \rrbracket_{w,g}^{c_1,c_2} = 1 \text{ iff } \forall w' \in W(c_2) : \llbracket \phi \rrbracket_{w',g}^{c} = 1.$$

By the same token, the framework allows for the possibility that there are two uses of the first person pronoun; one according to which it picks out the agent of the context of use, and another according to which it picks out the agent of the context of imagination. Again, this has some immediate appeal. When we play an interactive fiction, we play as one of the characters in the story. This is reflected in our practices of talking about them in a first person way. At the same time, it is always open for us to talk about our actual selves in the first person mode as well, as Tom Cruise does in Scenario 2.

$$\llbracket I_1 \rrbracket_{w,g}^{c_1,c_2} = \text{the agent of } c_1.$$
$$\llbracket I_2 \rrbracket_{w,g}^{c_1,c_2} = \text{the agent of } c_2.$$

The bicontextual account offers a conservative extension of the classical Kaplan machinery. It preserves the standard semantics in the ordinary cases where $c_1 = c_2$, but adds a local extension in the special cases where $c_1 \neq c_2$. The latter cases are isolated to instances where we are reporting on imaginings induced by some fiction. Of course, the fact that we use a framework that allows for there being two uses of 'actually' and 'I' does not settle the question of whether the indexical expressions *are* ambiguous in this way. So I'll spend the rest of the paper motivating this idea by showing how the bicontextual account handles the examples we've encountered so far.

---

[10] [14, 83], [11, 451n12], [28, 87], [20, 337-8].

When we engage in a conversation about some fiction we carry on as if the story was real, shifting our perspective from the actual world to the fictional worlds. Once our perspective has shifted in this manner, 'actually' seems to be able to take us to the fictional worlds, as we witnessed in (4).

(4) Possibly $\forall x$ (Actually (x is a japanese plane) $\rightarrow$ x survive)

The fictional report (4) says, roughly, that all the fictional planes could have survived the attack. The problem for the standard account was that if Actually shifts to the actual world, we end up picking out actual, instead of fictional planes. But if the operator in (4) is Actually$_2$, we get the desired result: (4) is true (relative to a context of use c_1 and a context of imagination c_2) if and only if there is a world w, such that all the Japanese planes in the worlds of c_2 survive in w. Since the worlds of c_2 are the fictional worlds, we correctly pick out the fictional planes.

The hypothesis that there are two uses of 'actually' can also explain the different truth conditions associated with (8).

(8) Had admiral Yamamoto attacked Pearl Harbour at night, he would have lost fewer planes than he actually did.

The fictional report (8) could either be used to compare the counterfactual number of lost planes with the number lost in the actual world or with number lost in the fictional worlds. The bicontextual account is a straightforward explanation for why (8) is ambiguous between these two readings: it is ambiguous because 'actually' is ambiguous. When 'actually' is read as Actually$_1$, we get the first reading: $[\![4]\!]^{c_1,c_2}_{w,g} = 1$ just in case, in the worlds w' where Yamamoto attacked at night, the number of planes he lost in the world w'' most similar to the actual world w''' is less than the number of planes he lost in w'''. We get the second when 'actually' is read as Actually$_2$, since the actual world is just replaced by the fictional worlds in the truth conditions.

The bicontextual account is also able to make sense of Scenario 1 and 2. To recapitulate, in Scenario 1 Teddy was playing as the character Jane in the videogame *Mass Effect*. After Jane kissed Kaidan in the fiction, Teddy reported on the event by saying 'I kissed Kaidan'. Here's a way of understanding what's going on. One thing Teddy imagines while playing the game is the sequence of events portrayed by the fiction (such as Jane kissing Kaidan). But in addition to that, the fiction asks him to imagine being the character he plays. This is, after all, the salient difference between a regular fiction, like a novel, and an interactive fiction where you take on the role as one of the characters. So the context of imagination induced by *Mass Effect* will be one where Jane is the agent, representing the imagined scenario located on Jane. Teddy conveys this information by talking about Jane in a first-person way.[11] If someone was unaware that Teddy was reporting on the fiction, they would judge his utterance as true just in case Teddy kissed Kaidan in the actual world. If, however, they were aware that Teddy was making a fictional report, they would interpret the sentence as 'I$_2$ kissed Kaidan', and judge it to be true just in case Jane kissed Kaidan in the fictional worlds.

Scenario 2 is interesting because it involves both uses of the first person. In this scenario, Tom Cruise said 'I killed myself' to convey the information that his avatar Bully killed Tom Cruise in the fiction. The bicontextual account gets the truth conditions right when

---

[11] I should briefly add two clarificatory comments. First, I don't think it's metaphysically possible for Teddy to be Jane. I'm only saying that he talks about the fiction from Jane's perspective in a first person way. Second, I am not taking a stance on whether fictional worlds ultimately have to be understood in terms of metaphysically possible worlds either.

the sentence is read as 'I_2 killed myself$_1$': it expresses something that is true just in case Bully killed Tom Cruise in the fictional worlds. But the bicontextual account has a further advantage, in that it is able to explain why the sentence (7) could in principle be read in all the following ways:

– I_1 killed myself$_1$
– I_2 killed myself$_2$
– I_1 killed myself$_2$
– I_2 killed myself$_1$

On the first reading, (7) says something that is trust just in case that Tom Cruise committed suicide. This would be the natural reading had Tom Cruise, for example, written a suicide-note explaining how he died. The other readings are all cases where (7) is a fictional report. On the second, (7) says something that is true just in case the fiction has asked us to imagine that Bully committed suicide. Although that is not intended reading in Scenario 2, it is something Tom Cruise could have said had the story taken that turn. On the third reading, Tom Cruise would be saying something that is true just in case Tom Cruise killed Bully in the fictional worlds. For example, if Bully started shooting in the night club, but Tom Cruise (in the game) turned around and killed him. On the fourth, and intended reading in Scenario 2, Tom Cruise says something true just in case Bully killed Tom Cruise in the fictional worlds. The easiest explanation for why these are all possible readings of (7) is that there are two uses associated with the first person pronoun.

As a final remark, it should be noted that the view presented here comes with a theoretical cost of its own. While the standard account postulates hidden operators, the bicontextual account complicates the semantics by introducing an additional contextual parameter. I do think, however, that the cost is somewhat mitigated by three factors. First, the bicontextual account is more suited in accommodating the data than its rivals. Second, since the context of assessment becomes inert in non-fictional discourse, it amounts to a conservative extension of a plain vanilla intensional semantics. And, third, there is independent motivation for the inclusion of contexts of assessment, as they might be needed to handle other linguistic phenomena. On balance, then, it seems like a bicontextual analyis has more going for it than either the standard account or the context shifting alternative.

6 Conclusion

In this paper, I've outlined a bicontextual account of fictional reports. The central idea was that when we use a sentence to make a fictional report, it is evaluated relative to a context of use and a context of imagination. A context of imagination was a representation of the counterfactual situation that the topical fiction invites the audience to imagine. I gave two motivations for this proposal. First, it gave predictions that corresponded to the intuitive truth conditions without postulating hidden operators. Second, it could accommodate the observation that indexicals are sometimes sensitive to whether they occur in fictional or non-fictional reports.

References

1. Bach, K., Thought and reference. Clarendon Press, Oxford (1987)
2. Bertolet, R., On a fictional ellipsis, Ekentnnis, 21, 185 - 154 numbers (1984)

3. Borg, E., Saying what you mean: Unarticulated constituents and communication. In Ellipsis and nonsentential speech. Springer, Dordrecht, 237 - 242 (2005)
4. Byrne, A., Truth in fiction: the story continued, Australasian Journal of Philosophy, 71, 24 - 33 (1993)
5. Cappelen, H & Hawthorne, J., Relativism and monadic truth. Oxford UP, Oxford (2009)
6. Currie, G., Fictional truth, Philosophical Studies, 30, 153 - 212 (1986)
7. Currie, G., The nature of fiction. Cambridge UP, Cambridge (1990)
8. Devitt, M., Designation. Harvard UP, Cambridge, MA (1981)
9. Field, H., Theory change and the indeterminacy of reference, The Journal of Philosophy, 70, 442 - 481 (1973)
10. Hanley, R., As good as it gets: Lewis on truth in fiction, Australasian Journal of Philosophy, 82, 62 - 128 (2004)
11. Kaplan, D., Demonstratives. In J. P. J. Almog & H. Wettstein (Eds.) Themes from Kaplan. Oxford UP, New York. 481 - 343 (1989)
12. Krasner, D.A., Semantics and fiction, Erkenntnis, 37, 235 - 273 (2002)
13. Lamarque, P & Olsen, S.H., Truth, fiction and literature. Clarendon Press, Oxford (1994)
14. Lewis, D., Anselm and actuality, Nous, 4, 173 - 188 (1970)
15. Lewis, D., Truth in fiction, American Philosophical Quarterly, 7, 37 - 44 (1978)
16. MacFarlane, J., Future contingents and relative truth, The Philosophical Quarterly, 33, 321 - 334 (2003)
17. MacFarlane, J., Truth in the garden of forking paths. In Relative truth. Oxford UP, Oxford. 81 - 102 (2008)
18. MacFarlane, J., Epistemic modals are assessment-sensitive. In Egan, A & Weatherson, B. (eds.). Epistemic modality. Oxford UP, Oxford, 84 - 178 (2011)
19. MacFarlane, J., Assessment sensitivity: Relative truth and its applications. Oxford UP, Oxford (2014)
20. Ninan, D., De se attitudes: Ascription and communication, Philosophy Compass, 3, 331 - 347 (2010)
21. Parsons, J., Assessment-contextual indexicals, Australasian Journal of Philosophy, 85, 1 - 17 (2011)
22. Phillips, J., Truth and inference in fiction, Philosophical Studies, 54, 273 - 253 (1999)
23. Predelli, S., Talk about fiction, Erkenntnis, 44, 45 - 77 (1997)
24. Predelli, S., Contexts: Meaning, truth and the use of language. Clarendon Press, Oxford (2005)
25. Predelli, S., Modal monster and talk about fiction, Journal of Philosophical Logic, 37, 277 - 257 (2008)
26. Reimer, M., The ellipsis account of fiction-talk. In Ellipsis and nonsentential speech. Springer, Dordrecht, 203 - 27 (2005)
27. Robson, J. & Meskin, A., Video games as self-involving interactive fictions, The Journal of Aesthetics and Art Criticism, 74, 83 - 177 (2016)
28. Smith, Q., The multiple uses of indexicals, Synthese, 78, 87 - 151 (1989)
29. Stock, K., Fictive utterance and imagining, Aristotelian society supplementary volume, 83, 83 - 154 (2011)
30. Voltolini, A., Fiction as a base of interpretation contexts, Synthese, 73, 23 - 47 (2006)
31. Voltolini, A., Fiction and indexinames, Journal of Literary Theory, 8, 253 - 322 (2014)
32. Walton, K., Mimesis as make-believe. Harvard UP, Cambridge, MA (1990)
33. Yalcin, S., Epistemic modals, Mind, 64, 583 - 1024 (2007)