

Would this paper exist if I hadn't written it?

Samuel Lebens

Published online: 31 January 2015
© Springer Science+Business Media Dordrecht 2015

Abstract This paper wants to know whether it would exist, or could exist, in worlds in which I didn't write it. Before we can answer this question, we first of all have to inquire as to what, exactly, this paper is. After exploring two forms of Platonism (pure and impure), and a theory that defines literary works in terms of events, I shall argue that the term 'this paper' is actually infected with ambiguity. Does this paper need me? It depends upon what you mean by 'this paper'. I lay out the options for what you might mean, and answer the question for each of the options.

Keywords Ontology · Art · Literature

1 Introduction

Would this paper exist if I hadn't written it? Before we can answer this question, in the eighth section, we first of all have to inquire what, exactly, this paper is. After exploring two forms of Platonism (pure and impure) about literary works, and a theory that defines literary works in terms of *events*, I shall argue that the term 'this paper' is infected with an ambiguity ignored by the standard theories about literary (and art) works. Does this paper need me? It depends upon what you mean by 'this paper'.

2 Author essentialism

According to Thomasson (1999, p. 8), George Orwell's *Animal Farm* is essentially an anti-Stalinist satire. The same sequence of words written before the rise of

S. Lebens (✉)
Center for Philosophy of Religion, Rutgers University, 106 Somerset Street,
New Brunswick, NJ 08901, USA
e-mail: samuel_lebens@hotmail.com

Stalin,¹ wouldn't have constituted the same literary work. Thus, a literary work can have properties merely in virtue of its provenance:

For that reason, it seems that a literary work is best conceived not as an abstract sequence of words but as an artifact that had to be created in those original circumstances in which it was created. (Thomasson 1999, p. 9)

One of the features of the circumstances of a literary works' creation is its creator. Of course, we're lacking a conclusive argument for author-essentialism—the thesis that literary works are essentially tied to their particular authors—but we're supposed to be moved in its general direction by the idea that the nature of a literary work can be so tightly tied to the circumstances of its birth.

A second motivation for Thomasson's author-essentialism (Ibid., chapter 3) is her view that author-essentialism is what grounds the following intuitive sort of counterfactual:

- (1) Had it not been for my existence (and my writing of this paper), this paper would not have existed²

It is only *intuitions* that motivate Thomasson's author-essentialism.³ And, there seem to be equally strong intuitions that we could bring to bear *against* author-essentialism. Imagine: soon after this paper's publication, a controversy erupts. My wife claims to have written it. Two camps emerge. One claims that this paper was written by me, the other that it was written by my wife. How should we characterise this controversy?

- (2) It is possible that this paper was written by my wife. It is also possible that it was written by me. The controversy centers upon which of these two possibilities is actual.
- (3) There are two possible papers that happen to be *lexically* identical. One is authored by me and the other is authored by my wife. The controversy centers upon which of these distinct possible papers is actual.

I wager that most people would prefer to characterise the controversy in terms of (2) rather than in terms of (3). The controversy isn't, *intuitively*, about the identity of the literary work, but about the identity of its *author*. Our intuitions here are cutting against the doctrine of author-essentialism. One paper could have been authored by my wife or by me.

A second thought-experiment: my wife has become frustrated with my meager list of philosophical publications. To help me out, she decides to write a paper in my name. Unbeknown to her, at roughly the same time, I sit down to write a

¹ Or, later by an ignoramus.

² Thomasson's actual example concerns Arthur Conan Doyle and the Sherlock Holmes stories.

³ To be fair to her, she is doing more than *merely* appealing to intuition. She argues at length that fictional entities and cultural artifacts are the sort of things about which common-sense everyday talk is particularly authoritative (Thomasson 2003, 2004, 2008). She is appealing to common-sense intuitions against the backdrop of a theory that gives those intuitions particular authority in this area of discourse.

philosophical paper. Lo and behold, we produce indistinguishable papers. Having read hers, do you feel any need to go and read mine? Or, have you already read it? If your answer is that you feel no need to read mine, having read my wife's, doesn't it seem like our one paper was authored twice over?

You might resist this conclusion. You don't need to read both papers because they're identical, but *only* in a certain respect. That doesn't mean that the two works are the same *work*. Let me push matters further. Imagine that the only physical and digital copies of my paper are destroyed. Thankfully, my friend Dustin has a photographic memory.⁴ He's read my paper. Normally, we'd say that if a literary work survives, word for a word, in a person's mind, then it can be recovered. But Dustin has also read my wife's lexically identical paper. Would it be possible for him to intend to copy mine from memory, but not hers?⁵ What would we make of the situation if he attempted, in his one copy, to create a copy of *both* works?⁶ It becomes harder to deny that there is only one literary work here. But, if the same work could have been written independently by either or both of two people, then author-essentialism would be false.

Furthermore, the truth of (1) is quite independent from the truth or falsehood of author-essentialism. A Lewisian account of (1) runs as follows:

(4) In the closest possible world in which I don't exist (and thus never write this paper), this paper doesn't exist.

(4) doesn't entail (5):

(5) In all of the possible worlds in which I don't exist (and thus never write this paper), this paper doesn't exist.

And thus, the truth of (1) is compatible with (6):

(6) In some possible worlds, I don't exist (and don't write this paper); nevertheless, in some of the more remote of those worlds, this paper still exists (written by somebody else)

(1) is compatible with (6). It lends no support to author-essentialism.

Had we been able to adopt author-essentialism, we would have arrived at a quick and easy solution to the puzzle of this paper. If author-essentialism is true, this paper couldn't exist without me and my act of writing it. But, our inconclusive exploration of the issue quickly descended into a war between different sets of intuitions. As we shall see, some of these intuitions are rooted in a Platonist conception of literary works, whereas Thomasson's intuitions are rooted in an *impure* Platonism.

⁴ It was, in actual fact, my friend Dustin Crummett who came up with the second stage of this thought-experiment.

⁵ Imagine that the two papers were visually indistinguishable. Even the most visually discriminating memory would fail to distinguish the 'two' papers.

⁶ Thanks to Jeff Speaks for posing this question to me.

3 What is a work? Platonism

By ‘work’ I mean, in the remainder of this paper, to indicate a broad category, including works of art and works of non-fictional literature.

Platonism in the ontology of art seeks to identify works with abstract structures. Of course, you might want to be Platonist about some mediums of art, and not others. Platonism about musical works suggests that *Beethoven’s Ninth Symphony* is identical to the sonic-structure embodied by its performances.⁷ A Platonist about the visual arts will say that the *Mona Lisa* is identical to the finely grained pictorial form embodied by the relevant piece of canvas in the Louvre.⁸ Platonism regarding literary works suggests that this paper is some sort of linguistic-structure.⁹ Platonism about literary works copes well with the anti-author-essentialist intuitions of the previous section—if the structure is the same, then so is the work. But, a running theme of this paper is that intuitions alone don’t seem to come down decisively on the side of any particular ontology of literary works, and thus, it should come as no surprise that we can marshal a number of powerful counter-intuitions *against* Platonism:

- (a) As Thomasson has already pushed: If *Animal Farm* is just a linguistic-structure, then it exists quite independently of Stalin. It cannot be *essentially* anti-Stalin. A work, on a Platonist account, cannot possess the essential features that works intuitively seem to possess merely as a consequence of their provenance.
- (b) Works, conceived as pure abstracta, cannot be *created*. Philosophers may discover, rather than create, philosophical theories, but we tend to think that they *are* responsible for the *creation* of the literary works that present those theories; just as artists are responsible for the *creation* of art works. The Platonist has to ditch, or re-construe, that way of thinking.
- (c) Platonists are going to struggle to account for forgery. A forged canvas found to be indistinguishable from the original *Mona Lisa*, will be an equally valid instantiation of the work.

In response to (a), the Platonist could follow a suggestion of Wolterstorff (1991).¹⁰ She could argue that the linguistic structure-type that is *Animal Farm* may have the property, *relative to its discovery by Orwell*, of being anti-Stalin, without holding that property intrinsically or essentially.

In response to (b), the Platonist will probably bite the bullet. What they can say, though, is *this*: creativity doesn’t always entail ontological creation. The *discovery* of certain abstracta *can* be an act that embodies creativity and ingenuity, and thus,

⁷ Kivy (1983) has defended such a Platonism regarding musical works.

⁸ Davies (2004, pp. 43–44) observes that Platonists about music and literature tend to eschew Platonism about the visual arts, although the logical space for such a Platonism certainly exists, as does the space for what Davies calls a ‘global structuralism’—the adoption of Platonist analyses of *all* types of work.

⁹ Platonism about literary works has been defended by Goodman (1978, pp. 207–211), Goodman and Elgin (1988, ch. 3), and Wollheim (1980:4–10, 74–83). These examples were gathered by Davies (2004, p. 43). As we shall see, Yagisawa (1999) seems to take the same sort of approach.

¹⁰ As cited and developed by Davies (2004, pp. 47–48).

the Platonist can salvage our talk of artistic and literary creativity even if she can't salvage our talk of actual *creation*.¹¹

Regarding forgery, the Platonist can adopt the attitude of Lessing (1995): forgery is a moral, but not an aesthetic, impairment. I may have plagiarised this paper word-for-word from Phil Swenson, with whom I share an office.¹² That fact wouldn't constitute a *philosophical* or a *literary* impairment in this paper, rather than a dark episode in my biography.

4 Failed-art

Mag Uidhir (2013, ch. 1) attacks Platonism via his argument that works *have* to be created. For any category of things, ϕ , you can divide the world into ϕ s and non- ϕ s. Often, being a non- ϕ is a trivial matter. Mag Uidhir's cheese sandwich and the hunk of pressed metal in the junkyard are both *non-cars*. But, the way in which the hunk of metal in question *came to be* a non-car tells us something interesting about it, namely, that it *used to be a car*, before it was crushed. Another way for a non- ϕ to be non-trivially *not a ϕ* is if it is a *failed- ϕ* . I am, like Mag Uidhir's cheese sandwich, trivially a non-lawyer. But, had I *trained* to be a lawyer, sat the bar exam, and *failed*, then my not being a lawyer would be a much more significant biographical detail. I would be a *failed-lawyer*. According to Mag Uidhir, a thing is a failed-lawyer, if and only if:

- (a) The thing is a non-lawyer
- (b) The thing attempted to be a lawyer in the right sort of way (which he calls a lawyer-attempt)
- (c) The thing's being a non-lawyer is the result of the failure of that thing's lawyer-attempt

To be a failed- ϕ is to be a non- ϕ as the result of a ϕ -attempt; ϕ -attempts are the sort of actions that, when successful, are necessary and sufficient for the creation of a ϕ . Mag Uidhir is convinced that there is such a thing as failed-art—not merely bad art, but things which fail to be art, despite somebody's best efforts. His example:

¹¹ Wolterstorff (1980, p. 89) bites this bullet. It has been suggested to me, by an anonymous reviewer, that the Platonist can do a better job than Wolterstorff's bullet biting, when talking about the creation of works of art. The artist really does create something, in that she creates a token (or a token of a blueprint for use, such as a score or mold) and, via this creation, she gives the rest of us access to the hitherto undiscovered Platonic object. In the words of the reviewer, she 'puts the platonic object (if the work is repeatable) on our ontological map.' But I'm not sure that this really achieves the goal of placing a wedge between artistic creativity and scientific discovery. The scientist also produces a token (or a token of a blueprint for use) when she writes down her discoveries; as does the explorer who draws a map of the lands they discovered, and the route to those lands. The token that the scientist, or explorer, *produces*, and in fact, the *creation* of the very experiment, or voyage, as an historical event, or performance, gives the rest of us epistemic access to the hitherto undiscovered law or land; it is the creation of a token that places the law or the land onto our map (in the second case, quite literally). It seems to me that the Platonist still fails to place a wedge between artistic creation, and mere discovery.

¹² Of course, I would have had to exchange my name for his name in this paragraph, but the Platonist might be able to accommodate that—it all depends upon how finely grained a literary-structure we're talking about.

Marcel Duchamp wrote that he was thinking of signing the Woolworth Building in New York City so as to convert it into a readymade artwork... Most assume that unlike his other readymade works, this would have failed to produce an artwork. Presumably his attempt would have been of the same kind as those behind his other readymades, and those apparently succeeded, so his attempt would have been of the right sort (*an art-attempt*).

The possibility of failed-art directly entails that there is such a thing as an art-attempt. Successful art *has* to be the product of a successful art-attempt, whatever that may be. Works must have creators. This paper can only exist as the product of an attempt. It cannot exist in worlds in which nobody successfully attempts to create it. Platonism, it seems, is false.

Mag Uidhir has tried to present us with a conclusive argument against Platonism. But the argument hangs upon his insistence that there is such a thing as failed-art, rather than mere bad-art and non-art. The only example he presents is the Woolworth Building in the counterfactual situation in which Duchamp attempted to make it into a readymade sculpture. But, who's to say that that art-attempt would have transformed that building into a failed-artwork? Perhaps it would have been such a stunning failure as to have no effect at all. In fact, the Woolworth building is and was already a work of architectural art by Cass Gilbert, its architect, and it would have remained a work of art after Duchamp's failed-attempt. It's not clear that you can turn non-art into failed-art, rather than just bad art. And it's certainly not clear that you can turn actual art, like the Woolworth building, into failed-art, by adding one's signature to it. If you can't find an uncontroversial example of a failed-artwork, you can't oil the wheels of Mag Uidhir's argument for the claim that art-works have to be the products of successful art-attempts. All the anti-Platonist seems to have are anti-Platonist intuitions, left to fight it out with Platonist intuitions.

5 What is a work? Impure Platonism

Those who reject Platonism don't want to say that a work is identical with the physical object/s that embody it. A painting and its canvas have different modal properties. The painting probably wouldn't survive a soapy scrub down. The canvas might.¹³ This paper could survive any given copy of it being burnt or digitally deleted (as long as some copy remained). We're groping for an elusive ontological category: works aren't physical, but, according to anti-Platonist intuitions, they can be *created*, so they're not abstract.

As Thomasson (2004, p. 90) puts it:

¹³ Of course, this makes a mockery of anti-Platonism. As an anonymous reviewer put it: 'Nobody would be tempted to identify the painting with the bare canvas, instead of with the canvas with the paint on it arranged just so.' But, in a sense, that illustrates my point: anti-Platonism simply can't be so crude as to identify an artwork with a canvas; a canvas with paint arranged on it thus and so is a more complicated type of entity. The painting isn't just the canvas. It's the canvas with the paints arranged just so. This is the sort of (more complex) entity to which anti-Platonism might need to appeal.

[W]orks of literature and music seem to fall between the cracks of traditional category systems: accommodating them will require acknowledging intervening categories for temporally determined, dependent abstracta: abstract artifacts created by human intentional activities.

One famous spelling out of an impure abstracta view is put forward by Levinson (1990). Think of a sonata by Beethoven written in 1806. Levinson suggests that we identify it with 'the performed-sound-structure-as-indicated-by-Beethoven-in-March-1806' (p. 224). Not with the structure itself, but with *the structure as indicated*. On this view: this paper was actually created by me. It is an impure abstracta; a literary-structure, as indicated by me in 2014.

6 Problems with impure Platonism

Impure Platonism gives rise to three pressing worries: (1) the view cannot actually account for the creation of works nor for (2) the longevity of works, and (3) it bundles too many essential properties upon works. I will explore these worries in turn.

6.1 Creativity and impure abstracta

What makes a concretum concrete, on some definitions, is that it can enter into causal relations. Abstracta, on such a view, cannot. An impure-abstractum, defined as an abstractum that can be created—*caused* to come into existence—makes no sense; it contradicts the very definition of an abstractum. Dodd (2000, 2002) levels just this complaint against impure Platonism.

Caplan and Matheson (2004) respond: You can't *prove* that abstracta can't enter into causal relations simply by *defining* abstracta as causally inert! The whole point of impure Platonism is to challenge old ontological taxonomies and definitions and, to echo Thomasson, to find cracks between traditional categories. To adopt impure Platonism is precisely to *deny* the definition that forbids abstracta from causal activity. The abstract-concrete distinction is notoriously difficult to draw.¹⁴ Dodd needs an argument for his preferred definition. And even if he maintains his definition, let him have it. We just mean something different; neither concrete, nor by Dodd's uncompromising definition, abstract. Call it 'shabstract' if you have to. You need to give me an argument if you want to convince me that there can be no such category.¹⁵

Another move open to opponents of impure Platonism is to claim that even if abstracta *can* enter into causal relations, they are still eternal; they cannot *come into existence*, and thus, they cannot be *created*. Caplan and Matheson point out that Dodd already accepts, for good reason, that some abstracta *do* come into existence (Dodd 2002, pp. 392, 397)—impure sets can come into existence, such as the

¹⁴ Cf. Lewis (1986, §1.7).

¹⁵ This paragraph benefited greatly from my being present at a reading-group at Rutgers, in which Caplan and Matheson's paper was discussed. I found Jonathan Schaffer and Cameron Domenico Kirk-Giannini's contributions particularly helpful for this paragraph.

singleton set whose sole member is the Eiffel tower, which came into existence in 1889. Even if, for definitional reasons, they cannot be *caused* to come into existence, Dodd is forced to accept that they are non-causally *brought into existence* by our actions. That might be enough. The creation of works could be cashed out in terms of ‘ontological responsibility’ rather than in terms of ‘causation’—I didn’t *cause* this indicated-type to come into existence, but my agency is *responsible* for its creation, as an ontological free lunch.

Mag Uidhir (2013, ch. 4) presents a more powerful argument than Dodd.¹⁶ Mag Uidhir is responsible for there being two doghouses in his backyard. He created them. Furthermore, he (non-causally) brought the set of dog houses in his backyard into existence, and is responsible for its having a prime number of members. But, of course, it doesn’t follow that Mag Uidhir made the number two a prime number! That would be a ridiculous claim. Mag Uidhir thinks that indicated type theory is bound to make an analogously ridiculous claim about the power of artists.

Recall Beethoven’s sonata. Let’s assume that it is an indicated type; a set, whose members include the sonic-structure of the sonata, Beethoven, a time, and perhaps an act of indication. This set was (non-causally) *brought into existence* by Beethoven. He is responsible for its existence. It contains a sonic-structure. That sonic-structure is dazzlingly beautiful. But Beethoven isn’t responsible for the dazzling beauty of the sonic-structure in question, even if he is responsible for the existence of the set that it joins. Beethoven is no more responsible for the beauty of the sonic-structure than Mag Uidhir is responsible for the cardinality of his set of doghouses being prime. Sure, he is responsible for the cardinality of that set, but not for the properties that that cardinality has.

The reason the impure Platonist fought so hard to make sense of artists creating their art works was to make them *responsible* for the aesthetic qualities of their artworks; and to make philosophers responsible for the literary properties of their literary works. Mag Uidhir’s point is this: Indicated-types *can* be (non-causally) created, but they *can’t* help us to do justice to our pre-philosophical notion of artistic and literary creativity.

I would respond to Mag Uidhir as follows. No one ever meant to claim that Beethoven is responsible for making the sonic-structure of his composition beautiful. They only claim that he is responsible for creating a sonata with a beautiful sonic-structure. I created this paper, and I’m responsible for creating a paper with the properties that this one has (even if I’m not responsible for the properties of those properties).

6.2 Longevity and impure abstracta

Caplan and Matheson raise their own concern with impure Platonism. If we think of Beethoven’s sonata as a set that includes Beethoven as a member, then, presumably,

¹⁶ Irvin (2014) understands Mag Uidhir merely to be rehashing Dodd-style arguments. But, as we shall see, Mag Uidhir’s argument goes further.

once Beethoven has died, the set will disappear.¹⁷ We could argue that impure sets exist for as long as they have any members remaining. But, if that is the case, then we're back with the previous problem: works cannot come into existence after all, and consequently cannot be created, because *one* of their constituents—the form or structure constituent—is *eternal*. If one member of the set's existing is enough for the set to exist, then the set in question will *always* exist.

As Caplan and Matheson point out, banishing the Platonic structure of the work from the analysis of its constituents, in addition to claiming that an impure set continues to exist for as long as it has any members remaining, allows you to avoid their problem. Perhaps this paper is actually the set of my act of writing it, along with all of the copies, present and future of it (and will exist until the final copy is lost); but, the linguistic structure exhibited by this paper isn't a constituent of this paper. Banishing the eternally existent element from the analysis of the paper helps them to avoid their problem. As we shall see in the next section, there are other avenues out of their problem. But even Caplan and Matheson's suggestion, of banishing the form/structure from the membership of the set, a suggestion which needs ironing out and further motivation, is enough to suggest that the problem needn't be devastating.

6.3 Modality and impure abstracta

Davies (2004, p. 108) asks us to compare Andy Warhol's *Brillo Boxes* with a hypothetical oil-painting, *Prairie Snowscape*, painted by 'a "naïve" occasional painter':

Barring any unusual features of the context of its creation, it seems that *Prairie Snowscape* could have been executed a few years earlier or later.

The same cannot be said for Warhol's finger-on-the-pulse-artwork, which seems:

much more closely tied to a particular time of execution ... Produced either substantially earlier or substantially later, the "point" of the work would be different.

The *Prairie Snowscape* is such that it could have been painted substantially later or substantially earlier without changing its identity. The indicated type theory, which includes a time frame as a constituent of the indicated type in question, cannot

¹⁷ It's tempting to think that this concern only gets going if you presuppose certain philosophies of time; this concern was pressed upon me by an anonymous reviewer. Only if you're a presentist would you worry that when Beethoven dies, he ceases to exist in an absolute sense and therefore has to take the set down with him. But, I think the worry has purchase whatever your view of time: even if you're an eternalist, and you think that the past and the future, and everything in them, timelessly exist, you still have to make sense of our regular talk of existing and ceasing to exist. The fear is that once Beethoven ceases to exist, in the sense in which even an eternalist and growing block theorist accept that things cease to exist, then any impure set in which they belong will also cease to exist, in the sense in which the eternalist and the growing block theorist think that things can cease to exist. And thus, we can worry, whatever our ontology of time, that works will die as soon as their creators do—again, I'm grateful to the metaphysics reading group at Rutgers for helping me to see that the concern is neutral to one's philosophy of time.

support that fact. And, even though it rightly demands that *Brillo Boxes* could not have been executed substantially earlier or later, it wrongly dictates that *Brillo Boxes* would be the same piece of work:

in a world in which the entire New York art-scene of the 1950s and 1960s doesn't exist as long as Warhol, in that world, executes a piece having the same outward appearance as *Brillo Boxes* at the same time as he executed the piece *Brillo Boxes* in the actual world.

Impure Platonism fails to respect the 'work-relativity of modality'.

It is only because we are viewing works as ordered n -tuples, including specific time frames, creators and structures that we got into this mess. Perhaps we could think of a work as an impure abstractum *without* being an ordered n -tuple. Perhaps works are simple. This might save the work-relativity of modality, and would also get us out of Caplan and Matheson's concern.

Perhaps we could model literary works upon fictional characters. Goodman (2004), following in the footsteps of Searle (1975), van Inwagen (1977) and Thomasson (1999), thinks of fictional characters as *simple* impure abstracta; as the posits of literary theory which finds itself constantly quantifying over fictional characters.

Our social practices and customs allow me to introduce a new fictional character simply by writing a story (or starting to write a story) in which I say, 'There once was a boy called Jack.' I keep attaching new properties to him. I say, 'He owned a very fast car.' And now it is true of him, in the story, that he owned a very fast car. I say, 'He was tall and had a peculiar birth-mark on his left cheek.' Now that becomes true of him, in the story. These properties don't really belong to Jack the fictional character. Instead, they belong to him in the story. But, on the view in question, there are some properties that he really has. He *is* a character in a story, for example. He *was* created to illustrate a philosophical point.

Let's apply the same analysis to works. I opened a new file on *word for windows* with the intention of creating a literary work (namely, this paper). From that moment on, this paper existed. At that stage in its development it was a mere draft, or draft-attempt, or work-in-progress; just as all trees were once saplings, all papers were once works-in-progress. Through the actions that I take, this work-in-progress acquires more and more properties. It acquires a subject matter, a determinate form, a tone, a style, and more. At some point in time this work-in-progress will become a work.¹⁸

One might worry that old drafts that remain on my hard-drive cannot be identical to this paper. In response, we could say that as a paper comes into existence, it sheds drafts just as a snake can shed its skin a number of times. This paper has had two substantially different drafts. The drafts still exist, but the digital files and human

¹⁸ Goodman (2004, footnote 26) concedes that regarding fictional characters there will be some vagueness around the timing of a fictional being's coming into existence. The same will be true for works. We needn't be too worried by this. It's an example of a larger philosophical problem—vagueness—that needs to be discussed elsewhere.

memories and physical copies that embody those drafts, no longer embody this paper, which is now embodied by a different set of files and memories and copies.

As the paper comes into existence, it needn't be an n -tuple. It needn't be complex. Like Goodman's fictional characters, it can be a simple impure abstractum. If we accept that works are simple impure abstracta we can respect the work-relativity of modality. Taking each work on its own terms, looking at the processes by which they were created, we can assess, work by work, what properties must be essential to it, and what properties must be accidental. Impure Platonism is able to escape the three problems we've raised with it, especially if we think of the impure abstractum in question as a simple entity, rather than an indicated type or n -tuple.

Neither Platonism nor impure Platonism seem open to a devastating critique.

7 Art as performance

Leaving both pure and impure forms of Platonism behind, Davies (2004, p. 117) suggests that works should be identified with the generative performances of creators. How do you respect the work-relativity of modality? Even if you come to the table with the assumption that works are simple impure abstracta, when you try to make cross-world identifications of those abstracta, you're going to end up asking yourself whether the generative performance that created that abstracta could have been executed in that situation. Picasso's *Guernica*-creating-performance wouldn't have been the same performance if it hadn't been a response to the bombing of Guernica.

[O]ur intuitions as to the modal properties of works are of the form: "Could this have been *done* under those circumstances?," where "this" refers to a generative performance rather than the product of such a performance.

These intuitions are part of a cumulative case for identifying works with their generative performances. What is it that you appreciate when you behold a work of art or literature, if not the creator's creative activity, her brushwork, her appreciation of light, or her way with words? You are appreciating a generative performance.

Unlike, Currie (1989), Davies doesn't want to associate a work with a *type* of an action. Instead, he wants to associate works with *tokens* of a type of an action. Davies makes the distinction clear via the following question: When John takes a hat-trick in a cricket test-match, what is it that Davies appreciates? Not the action type, *taking a hat-trick*. No. He appreciates *John's taking of a hat-trick*. Davies concedes that it 'is only insofar as I regard the taking of a hat-trick in a test match as a remarkable *type* of achievement, of which John's performance is a *token*, that I can appreciate John's particular action *as* a token of this type. But this doesn't make John's performance a type.'

We would generally think of a particular performance as an *event*. But, events are thought to have time instants as constituents.¹⁹ This is what landed Levinson's

¹⁹ As Davies points out (p. 168), if events don't have time instants as constituents, as in the very different accounts of Kim (1976) and Cleland (1991), then they contain a time-slot relative to a causal chain, as argued by Davidson (1980).

indicated type theory in trouble. If a time instant is a constituent of a work, then it seems that the work couldn't be produced at any other time. This fails to respect the work-relativity of modality. To avoid this problem, Davies postulates a special class of event, which he calls 'doings' or 'happenings' (pp. 116, 167–178).

Davies' argument is that the time of occurrence enters into the individuation of an event only in a limited respect, because events need to have a unique location in the time-line of a world. But, that doesn't mean that the location itself is an individuating factor *across* worlds. Some event-tokens are going to be tightly modally restricted temporally, and some won't be. For instance, if the non-temporal constituents of an event-token don't give rise to a particularly rich analysis, then the event-token will end up essentially tied to a very particular moment in time. Davies' example (p. 172) is John's drinking a cup of coffee at 9 a.m. on July 1, 2002. Given the poverty of the analysis open to this event:

There may be no way of giving content to the idea that it is *this* event-token that is being considered counterfactually, rather than another act of coffee-drinking on John's part [other than by paying attention to the precise time of occurrence]...

'The sinking of the *Titanic*', on the other hand, taken as an elliptical reference to 'the sinking of the *Titanic* by an iceberg on the ship's maiden voyage across the Atlantic,' allows us more latitude in cross-world evaluations:

The event-token can only occur in counterfactual scenarios in which it follows the construction of the *Titanic*, its sailing on its maiden voyage across the Atlantic, and its entering that region of the Atlantic where icebergs are a potential threat to navigation. This, however, allows for the event-token to occur at any moment in an extended temporal interval, given that other general features of early-twentieth-century history are held constant. (p. 171)

And thus, we can distinguish between two varieties of event. There are events which are 'simple' or 'basic', admitting only of an impoverished analysis. And, there are events which admit of richer analyses. For the latter type of event, there will generally be much more cross-world latitude regarding when the event needs to be located. It is the latter type of event that Davies calls a 'happening' or a 'doing'.²⁰

Dilworth (2005) fears a regress:

[A]rguably the reason ordinary objects, unlike events, are able to persist through change and be modally flexible in their spatiotemporal positions is because they are capable of being contingently changed by, and reidentified relative to... *events* ... which in general provide the basis for the most basic scientific descriptions of the universe. Thus the concept of empirical contingency for objects presupposes a two-tier structure, in which object contingency or modal flexibility is explained in terms of their reidentifiability

²⁰ We could attempt to save Levinson's indicated type theory by construing acts of indication as happenings or doings, not tied too tightly to a particular time, but then Levinson's theory just collapses into Davies'—an art work isn't an indicated type but an act of indicating a type.

relative to distinct clusters of modally inflexible events. But clearly a similar framework cannot be applied to events themselves, without embarking on an unacceptable explanatory regress that would have to postulate proto-events to explain the contingency of ordinary events, and so on.

An unacceptable explanatory regress is one that is either infinite or, at least, cumbersome. But this regress is neither infinite nor cumbersome. Some events are simple, some are complex. The events that play Dilworth's basic scientific role—proto-events—can all be simple. That doesn't rule out the existence of a richer variety of events—or *happenings*—that are parasitic on the existence of simpler varieties of event. *Doings* exist, but they are an ontological free lunch. In fact, two distinct *doings* with distinct modal profiles can be embodied in a given possible world by a single set of simpler events. Certain running and throwing events embody John's hat-trick, *and* the first hat-trick of the Cricket season. In another possible world, where another test-match is played prior to John's test-match, those two *happenings* might be embodied by *different* sets of running and throwing events.

I think a much more probing concern is that Davies' distinction between events and doings is based upon a confusion between events and descriptions of events. A single event can admit of more and less fine-grained descriptions. That doesn't mean that we've got more than one thing; events and doings. We've got one type of thing—events—which admit of finer or coarser descriptions. John's drinking his tea could be described with or without reference to the time at which he did it; but those two descriptions shouldn't give rise to more than one thing in our ontology—a thing that is essentially tied to a time and a thing that isn't.²¹ But, perhaps we can offer Davies a defense here in which there is no such confusion, and in which *doings* really justify their claim to being an ontologically free lunch. Perhaps events and doings are the same things but under different counterpart relations.

For Davies, events, whatever their fundamental ontology, have a restricted counterpart relationship that only allows for counterparts that occur at the same time in other worlds. *Doings* might be the same things as events, ontologically, but if so, *doings* are events under a more liberal array of counterpart relations; depending upon the conversation at hand. Just as my counterpart relations are context-sensitive, such that in some conversations 'I could have been a fish' will be true, and in other conversations it would make no sense at all ('No, Sam, you couldn't have been a fish. You could have been a dentist, but not a fish!'); the counterpart relations of *doings* might be likewise context-sensitive, sometimes concerned more with temporal location, and sometimes less. Events are just doings under a much more restricted counterpart relation, and that's why descriptions are so relevant: they are what make certain counterpart relations salient in a conversation. But, 'events' is a term of art. Accordingly, events never allow for counterpart relations that differ in their temporal location, whatever the conversational context, and it is that counterfactual temporal rigidity, perhaps, that gives *events* the scientific role that Dilworth is so anxious to protect.

²¹ Thanks to an anonymous reviewer for pushing this concern.

Some (Kania 2005; Stecker 2005) attack Davies for his highly revisionary analysis of our art and literature talk. We want to say that ‘the *Mona Lisa* can be seen at the Louvre.’ But, you can’t see an event, or hang one on a wall! Davies (p. 187) argues that he has merely uncovered an ambiguity in the way that we talk about processes and products. Take any process, say, sunburn. When we talk of ‘*studying, understanding, or investigating,*’ sunburn, we are ‘characteristically [interested in] *the process generative of the product*’. But, when our interest is some sort of physical operation, our focuses shift to the product of the process. ‘[W]hen I treat John’s sunburn, I treat the product of a particular process, but when I study sunburn, it is the process that concerns me.’

If the performance theory of art is correct, we should expect a similar ambiguity. When you want to understand or study this essay, you are trying to understand my thought processes, the problems I was trying to solve, and the route I attempted to navigate towards their solution. When you’re interested in turning a page, you relate not to the process that generated this paper but to the product of that process, in your hands, or on your screen. Of course, you can’t hang an event on the wall, but our talk about art works is no less ambiguous than our talk about sunburn.

Unfortunately, Davies has revealed an ambiguity that is more endemic than he seems to realise. He thinks there’s a simple disambiguation strategy—if our concern is with a work as an object of appreciation, then ‘the work’ refers to the generative performance; if our concern is with some sort of physical operation, then ‘the work’ refers to the product of that generative process. But, matters seem more complicated. Take the sunburn case as an example. If I’m a diagnostician, I might be supremely interested in studying the product, rather than the process of sunburn. Let the dermatologist worry about the process, while I just chart and marvel at the product!

Imagine that I know nothing about art history, cannot appreciate a deft brush stroke, nor have any knowledge of the sorts of processes that gave rise to the painting before me. Perhaps all I’m appreciating is the form of the thing; the Platonic pictorial-structure; just as the musical ignoramus may only be appreciating the sonic-structure of a musical work, rather than its achievement, or its place in musical history. The reader of this paper, completely disinterested in who I am, and how I wrote this, may be interested only in the propositional content of this paper and not in its production. Other consumers, no doubt, are interested in the generative processes, just as Davies spells out.

8 Taking stock

We started this paper with intuitions for and against author-essentialism. It turned out that the intuitions pulling against author-essentialism were of a Platonist bent, and that the intuitions tugging in the other direction were of an *impure* Platonist bent. Neither theory is open to an insurmountable challenge, but, at times, both theories have to make a number of uncomfortable claims. Perhaps the reason that neither direction satisfies all of our intuitions is that there is no single answer to what a work *is*. Instead, our talk about works is littered with the sort of ambiguity

that Davies discovered, and other ambiguities too, which don't conform to the simple disambiguation strategy that Davies put forward.

The Platonist admits that there is an ambiguity. Sometimes when we talk about an art-work, we're talking about an *instance* of an artwork—the painting hanging in the Louvre—and sometimes we're talking about the Platonic form in question; the *real* art-work. The impure Platonist admits a parallel ambiguity. Sometimes we're talking about a physical embodiment of an artwork—such as the painting hanging in the Louvre—and sometimes we're talking about the *actual* art-work; the impure abstractum. Davies, as we've seen, subscribes to a parallel ambiguity between the processes and the products of art. But, according to all of these theories, there is the *real* artwork (whether it be a Platonic form, an impure abstractum, or a generative performance) and there is the physical thing somehow *associated* with that real artwork (the thing that embodies the artwork, or the product of the artwork). The ambiguity is that 'artwork' sometimes applies to the real thing, and sometimes to the associated physical thing.

In a sense, what all of the standard theories accept is that there is something akin to a type-token ambiguity in the vicinity; an ambiguity analogous to that which infects the word 'word', as pointed out by Peirce (1906, p. 505):

There will ordinarily be about twenty *thes* on a page, and of course they count as twenty words. In another sense of the word "word," however, there is but one word "the" in the English language; and it is impossible that this word should lie visibly on a page or be heard in any voice...

There are tokens of the word 'the' and there is the type. Sometimes the word 'word' refers to tokens of words, and sometimes to a type of a word, and sometimes to tokens of the word 'word', and sometimes to the type of that word.

Everyone accepts that talk about art-works and literary works gives rise to a similar ambiguity, between the more abstract thing (be it the Platonic form, the impure abstracta, or the generative performance) and the more concrete thing (the canvas or the manuscript). But, I want to suggest that the ambiguity runs deeper than any of these standard theories realise. For any work, there are *more than two* things which we might be referring to when we say 'this work is ϕ '; more than one type, and more than one token (or token-type). Given the myriad ways in which most words can be used, both from a semantic and pragmatic point of view,²² the multiple potential significances of a term like 'this work' should hardly surprise us. But, the standard schools that I've been referring to only seem willing to countenance a binary ambiguity: according to them, 'this work' is ambiguous between the more abstract thing, and the more concrete thing, and that's it. But that *isn't* it!

Sometimes we are merely interested in coming to an uneducated or de-contextualised appreciation of the *content* of the work—'Gosh this work is beautiful'. In those situations the Platonist is right. The object of our interest is the pictorial-form, or the sonic-structure, or the linguistic structure type.

²² See, for example, Ziff (1972), who demonstrates the multiple meanings latent in almost any word.

Sometimes we are interested in the work as a cultural phenomenon—‘This work is a national treasure’, and then I think the impure Platonist is right—we are dealing with a simple impure abstractum created by the artist—a posit of social ontology.²³

Sometimes we are concerned with the physical work—‘this work is hanging in the Louvre’. True, when I’m interested in the *Mona Lisa* hanging in the Louvre, I’m not interested in a piece of canvas, but in that canvas qua embodiment/instantiation of the *Mona Lisa*; when I want to know the whereabouts of Michelangelo’s David, I’m not asking after the whereabouts of that lump of stone, but for the whereabouts of that lump of stone qua embodiment/instantiation of Michelangelo’s David. But sometimes my concern might be purely physical: ‘help me lift this statue, I need to move it out of the way!’

Finally, sometimes, when I’m interested in the achievement of the work, or in coming to an educated appreciation of the work, the object of my interest will be the generative performance.

The ambiguity is neither binary nor simple. For any ‘work’, there is a pure Platonic form, a simple impure abstractum, and a generative performance (inter-related in important ways). This explains why some thought experiments can lead you towards Platonism, because they play on your habit of thinking of the work in terms of the Platonic form. Other thought experiments lead us in other directions, by playing on our habit of thinking of the work in terms of the abstract particular, or in terms of the generative performance. We are habituated to think in all of these ways. Different thought-experiments merely tease out different habitual modes of thinking.²⁴

For example, Davies marshaled the intuition that *Brillo Boxes* would not have been the same work had it been created substantially later or earlier. Others might have the intuition that it would have been exactly the same work but would have merely been much less valuable.²⁵ But surely this clash of intuitions comes down to the fact that they are not thinking about the very same thing. Davies is thinking about *Brillo Boxes*, the social phenomenon, which is very tightly associated with a specific time and place in a cultural landscape.²⁶ The clashing intuition is thinking of *Brillo Boxes* as a certain pile of boxes arranged thus-and-so by Andy Warhol. Our intuitions clash because we’re not all thinking about the same thing.

9 Disambiguating ‘this paper’

Given the conclusion of the last section, conversation about this paper is open to an ambiguity. The phrase ‘this paper’ might be used to refer to any of three (or maybe more) distinct things. What are those things?

²³ Perhaps deconstructionists identify the literary work with the impure abstractum *as received by a culture at a time*. These culturally received entities might be an important element of a full ontology of literature. But, they lie beyond the scope of this paper.

²⁴ Williams (1970) notes how, in the different tellings of mind–body thought-experiments, contrary intuitions can be elicited from the same audience.

²⁵ This was the intuition of an anonymous reviewer.

²⁶ He would claim that the social phenomenon was actually a generative performance. It looks more like a simple impure abstractum to me—a posit of the social sciences. But this is beside the point.

9.1 'This paper' as linguistic structure-type

Sometimes when we think of 'this paper', we're thinking of a specific linguistic structure-type. Interesting questions arise here when we want to figure out what a linguistic structure-type might be. Yagisawa (1999) makes the Platonic suggestion that a literary work *is* an ordered set of letter-shape types, and spaces, in a specific sequence. To be fair to him, no sooner does he make this suggestion than he rescinds it. Those line-shapes need to be *interpreted* in a specific way if you want it to give rise to the form/structure of a specific literary work (the same shapes, interpreted differently, would give rise to the structure of a different work). Ultimately, he suggests that a literary work is a sequence of letters under a specific interpretation (p. 168). But, for all I know, you might be reading this paper via a Mandarin translation. There's a very important sense in which you'd still be reading my paper, despite not reading it in English. If this paper, or its content, is a set of syntactic objects under an *English* interpretation, then there's no obvious sense in which you're reading my paper, or its content, if you're reading it in Chinese. That seems wrong.

Perhaps a more promising route is to think of the literary structure as a set of propositions. This is consistent with our intuition that a translation of this paper is still, in some important sense, *this* paper. But, firstly: sometimes we're impressed, not just with what the author said, but *how* the author said it, her choice of words, tone, poetic rhetoric, etc. Secondly: this paper talks about *you*. It imagines you reading this paper in Chinese. One might assume that the referent of 'you' changes every time that somebody else reads this paper; the propositions *you* entertain when you read this paper are *different* to the propositions that somebody else would be invited to entertain.

In order to respect the indexical nature of certain words within literary works, one could suggest that literary-structures are actually sets of Kaplanian characters. The character of a sentence, in Kaplan's (1989) sense of 'character', is a function from utterances to propositions. The proposition expressed by the sentence, 'You're standing on my foot', changes depending upon the circumstances of its utterance. But, there is something constant about its meaning. It always means to assert that the addressee is standing on the foot of the utterer, whoever they may be. The *character* of the sentence is what remains constant; the function from utterance to intended-proposition. The character of an English sentence can be identical to the character of a French sentence: 'You're standing on my foot' and 'Vous êtes debout sur mon pied'²⁷ have the same character; the same function maps utterances of these two sentences to their target propositions.

But, there are two problems with this suggestion. First of all, as it stands, it's under-motivated. It's not clear that the words 'you' and 'I' and 'me', though normally indexical, are really indexical in contexts like this paper—or in many literary works. 'You' probably functions in the second-person plural, addressing all of my readers—its meaning doesn't shift from reader to reader, and the words, 'I'

²⁷ Although, in actual fact, the idiomatic way of saying it would be, 'vous marchez sur mon pied'.

and ‘me’, in the context of this paper probably function in a descriptive sense, picking out the person who wrote this work, whoever that happens to be. Secondly, one of the reasons that we wanted to move away from a propositional account of the structure of this work is that we wanted to respect the appreciation we have for the way in which a work *expresses* propositions. Kaplanian characters could quite easily map two sentences with very different tones onto the very same propositions when uttered in the same circumstances: ‘Oi you’ and ‘Excuse me’ may well have the very same character.

We don’t want to reduce the structure of this work to an ordered set of English words under an English interpretation, because you might be reading it in Mandarin. But, on the other hand, we don’t want to reduce the structure of the work to a set of propositions, because we often appreciate *the way* in which a work says what it says, rather than merely appreciating *what* it says. We could strike this balance by appealing to a restricted *type* of Kaplanian character, of my own invention. Let’s call it a translation-function, or t-function. A t-function, like a character, maps sentence-tokens onto propositions given a circumstance of utterance, but, unlike characters, a single t-function only maps sentences with similar literary qualities. ‘Oi you’ and ‘Hey you’ might have the same t-function, but ‘Oi you’ and ‘Excuse me’ will not.

When a translator translates a work, she tries not merely to swap sentences of the object-language with sentences of the target-language that share a Kaplanian character, she also hopes to find sentences in the target-language that share a *t-function* with the sentences in the object-language. Of course, whether two sentences have the same t-function, will often be vague—but this is as it should be, as it respects the ways in which translation is an art rather than a science.

If we view a literary work qua literary structure as a set of t-functions, we secure three desiderata: (1) we maintain the language neutrality that Yagisawa failed to protect; (2) we make room for the fact that the propositional content expressed by a literary work *might* shift from reader to reader (if certain words are indexical); and (3), we make sure that the literary structure of the work is tied to the literary qualities of the sentences it contains, at least in a loose sense.

But, true to the spirit of this paper, perhaps we should accept that even the phrase ‘the structure of this paper’ is ambiguous. Sometimes a reader may be interested merely in what is said. In which case, the reader might be referring to the set of propositions expressed. Sometimes the reader may be interested in my command of the English language, in which case, she may be referring to an ordered set of English words under an English interpretation. Finally, the translator, or two readers reading this paper in different languages and yet sharing an appreciation for the structure of this work, might be referring to an ordered set of t-functions.

9.2 ‘This paper’ as simple impure abstractum

Sometimes when we think of ‘this paper’ we’re thinking of that simple impure abstractum that I described in Sect. 6.3, which came into existence as a work-in-progress when I opened *Word for Windows* with a specific intent, shedding drafts along the way until it bore the properties that it bears today.

9.3 'This paper' as my generative performance

Sometimes, when we think of 'this paper' we're thinking of my goal-driven work that created the product that you have in your hands or on your screen.

10 Would 'this paper' exist if I hadn't written it?

When we ask whether this paper would exist if I hadn't written it, we need to be clearer. What do we mean by 'this paper'?

10.1 Would this structure exist?

Singular propositions containing fictional characters might be thought not to exist until those characters are brought into existence. On this view, both the set of t-functions *and* the set of propositions that could be thought to constitute *A Study in Scarlet* can only exist once somebody has brought the relevant fictional characters and fictional entities into existence. Similarly, the set of propositions and the set of t-functions, either of which could be said to constitute the structure of *this* paper, can only exist once somebody, anybody, has created *Jack*, the fictional boy who has a fast car and a peculiar birthmark, to whom this paper refers. On such a view, literary structures really are dependent upon certain creative acts when, as in this paper, the work refers to fictional characters. But, it will never matter *who* creates the characters. Once *they* exist, the structure automatically comes into being. The structure isn't authored.

Furthermore, this paper refers to non-fictional philosophers. Until their births, propositions about them might be thought not to exist, nor t-functions yielding such propositions as values. Accordingly, literary structures that refer to temporally bounded beings, might only come into existence once the relevant beings come into existence. But, the structure itself won't need an author. Once again, it will automatically arise. Alternatively, if propositions about future-entities exist timelessly, then so do the relevant t-functions. Either way, the structure of this paper doesn't need my act of authoring it.²⁸

10.2 Would this impure abstractum exist?

This paper qua impure abstractum is the product of a certain generative performance, namely my *motivated manipulation of the linguistic medium through*

²⁸ Interestingly, Bertrand Russell, though a Platonist about universals (such as pictorial-form and sonic-structures), was (for a while) a constructivist about *propositional* content (while he held on to his Multiple Relation Theory of Judgement). On that view, a view which I have defended at length in my PhD thesis (Lebens 2010), literary structures are going to have to be exceptions. Pictorial-structures and sonic-structures generally exist timelessly. *Literary* works, on the other hand, because their structure is cashed out in terms of propositional content, and because propositions are constructed by minds, might be different. I relegate this, my actual view, to a footnote because I don't want my eccentric views about propositional content to distract people from my eccentric views about the ontology of works!

which the structure of this paper was specified as a focus of appreciation (echoing the language of Davies 2004). To be *this* impure abstractum is, it seems, to be the product of *this* generative performance, and of no other. The question is, in another possible world, could somebody else have performed that performance and thus given rise to *this* abstractum? Rohrbaugh (2005) argues that they could not have done.

Some possible works cannot co-exist. Rohrbaugh imagines Maya Ying Lin's Vietnam Memorial in Washington, DC, built in a world in which the Vietnam War never occurred. It would 'commemorate a war that never happened, men who never died. Its significance in that world seems so thoroughly altered that we ought to say that it is not our memorial.' The two possible works—Lin's bizarre memorial to an unfought war and Lin's *actual* memorial to the Vietnam War—are not compossible. But, this impossibility is a function of the 'conflict between the essential dependencies of the works.' This entails that when there is no conflict between the essential dependencies of two possible works then the possible works will be compossible.

Now, imagine that in some possible world my wife creates an impure abstractum that comes to have the linguistic structural properties of this paper. She performs a *motivated manipulation of the linguistic medium through which the structure of this paper was specified as a focus of appreciation*. But, even if she does perform *that* action, we simply can't say that it produces *this* work because her possible creative act shares its essential dependencies with my *actual* creative act. That means that the two creations—my wife's possible creation, and my actual one—are compossible. But, given that we're not talking about Platonic structures, but impure abstract products, in the possible worlds in which my wife's act and my act *coexist*, the resultant works are *distinct*. 'And since what is possibly distinct is actually distinct as well' we know that the work that my wife *could* have produced is not identical to *this* work. This simple abstractum, it seems, could only have been created by me.

10.3 Would this generative performance exist?

Rohrbaugh's argument seems to entail author-essentialism regarding works qua impure abstracta, *and* regarding works qua generative performance. But, actually, since we're viewing generative performances as *doings* rather than *events*, things get murky.

One of the things that we learnt in Sect. 7 is that one set of simple events can embody a number of doings/happenings. The set of head-scratching events, mumbling events, thoughtful walk events, and typing events, that constitute the generative performance of *my first paper dedicated to the ontology of literature*, also constitutes the generative performance of *the best piece of work ever to have the title, 'Would this paper exist if I hadn't written it?'* In other possible worlds, those two happenings come apart. The same set of actual simple events embodies a number of modally distinct *doings*.

If two *doings* are done by different agents in the *same* possible world, then the doings must be distinct. But, *across* possible worlds, I would wager otherwise. Take, for instance, the killing of Kennedy. The people bothered by the question as to

who killed Kennedy aren't bothered about the identity of the *doing*, but by the identity of the *doer*, and thus we have at least generated an intuition in favour of the existence of *doings* that have their *doers* contingently. *Lee Harvey Oswald's killing of Kennedy* exists in fewer worlds than does *the killing of Kennedy simpliciter*. Lee Harvey Oswald is an essential constituent of *his* killing of Kennedy. But, he isn't an essential constituent of *the* killing of Kennedy. I am an essential constituent of *my motivated manipulation of the linguistic medium through which the structure of this paper was specified as a focus of appreciation*, but I am not an essential constituent of *the* motivated manipulation of the linguistic medium through which the structure of this paper was specified as a focus of appreciation. Given the definite article in its description, that doing can exist only *once* in each world. In worlds in which my wife beats me to it, she will be the *doer* of that *doing* (even if, by the lights of Rohrbaugh's argument, it gives rise to a different product).

11 Conclusion

There is no such thing as this paper *simpliciter*. Talk about 'this paper' is actually ambiguous between three (and often many more) referents: a linguistic structure (which is in turn ambiguous between different abstract entities), an impure abstractum, and a generative performance (which is also ambiguous between a number of *doings*).

- Author-essentialism is false for linguistic structures. They exist (or perhaps come into existence) independently of any particular author... nobody authors linguistic-structures.
- Author-essentialism is true for this work qua impure abstractum.²⁹
- Finally, the simple events that gave rise to the impure abstractum in question (typing events, editing events, and so on) embody a number of different *doings*. Some of those *doings* have me as an essential component, some do not.

Acknowledgments Thanks to Carl Mosser, Phil Swenson, Evan Fales, Ross Inman, Natalja Deng, Dustin Crummett, Robin Dembroff, Jeff Speaks, Tzvi Novick, Curtis Franks, and Gaby Lebens for very helpful conversations about these topics, and to all of the other people I've harangued about the nature of literary works. Thanks also to the anonymous reviewers who have helped me to improve this paper immeasurably. Any remaining faults in this paper, are mine alone—More accurately, I take responsibility for the possible referents of 'this paper' that anybody actually authored!

References

- Caplan, B., & Matheson, C. (2004). Can a musical work be created? *British Journal of Aesthetics*, 44(2), 113–134.
- Cleland, C. (1991). On the individuation of events. *Synthese*, 86, 229–254.
- Currie, G. (1989). *An ontology of art*. New York: St Martin's Press.

²⁹ Rohrbaugh (2005) accepts that his argument admits of a small number of exceptions, beyond the scope of this paper. Either way, *this* simple abstractum is no exception to his argument!

- Davidson, D. (1980). *Essays on actions and events*. Oxford: OUP.
- Davies, D. (2004). *Art as performance*. Malden, MA: Blackwell.
- Dilworth, J. (2005). Review of David Davies, *Art as performance*. *The Journal of Aesthetics and Art Criticism*, 63(1), 77–80.
- Dodd, J. (2000). Musical works as eternal types. *British Journal of Aesthetics*, 40(4), 424–440.
- Dodd, J. (2002). Defending musical platonism. *British Journal of Aesthetics*, 42(4), 380–402.
- Goodman, N. (1978). *Ways of world making*. Indianapolis: Hackett.
- Goodman, J. (2004). A defense of creationism in fiction. *Grazer Philosophische Studien*, 67, 131–155.
- Goodman, N., & Elgin, C. (1988). *Reconceptions in philosophy and other arts and sciences*. London: Routledge.
- Irvin, S. (2014). Review of art and art-attempts. *Notre Dame Review of Books*. <http://ndpr.nd.edu/news/46001-art-and-art-attempts/>
- Kania, A. (2005). Review of David Davies, *Art as performance*. *Mind*, 114(453), 137–141.
- Kim, J. (1976). Events as property exemplifications. In M. Brand & D. Walton (Eds.), *Action theory*. Dordrecht: D. Reidel.
- Kivy, P. (1983). Platonism in music: A kind of defence. *Grazer Philosophische Studien*, 19, 109–129.
- Lebens, S. (2010). *On judgement: A multiple relation*. PhD Thesis. Birkbeck College, London.
- Lessing, A. (1995). What is wrong with a forgery? In A. Neil & A. Ridley (Eds.), *Arguing about art*. New York: McGraw-Hill.
- Levinson, J. (1990). *Music, art and metaphysics*. Ithica and London: Cornell University Press.
- Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Mag Uidhir, C. (2013). *Art and art-attempts*. Oxford: OUP.
- Pierce, C. S. (1906). Prolegomena to an apology for pragmatism. *The Monist*, 16(4), 492–546.
- Rohrbaugh, G. (2005). I could have done that. *British Journal of Aesthetics*, 45(1), 209–228.
- Searle, J. (1975). The logical status of fictional discourse. *New Literary History*, 6, 319–332.
- Stecker, R. (2005). Review of David Davies, *Art as performance*. *The Journal of Aesthetics and Art Criticism*, 63(1), 75–77.
- Thomasson, A. (1999). *Fiction and metaphysics*. Cambridge: CUP.
- Thomasson, A. (2003). Fictional characters and literary practices. *British Journal of Aesthetics*, 43(2), 138–157.
- Thomasson, A. (2004). The ontology of art. In P. Kivy (Ed.), *Blackwell guide to aesthetics*. Oxford: Blackwell.
- Thomasson, A. (2008). Existence questions. *Philosophical Studies*, 141, 63–78.
- van Inwagen, P. (1977). Creatures of fiction. *American Philosophical Quarterly*, 24, 299–308.
- Williams, B. (1970). The self and the future. *Philosophical Review*, 79(2), 161–180.
- Wollheim, R. (1980). *Art and its objects* (2nd ed.). Cambridge: CUP.
- Wolterstorff, N. (1980). *Works and worlds of art*. Oxford: Clarendon Press.
- Wolterstorff, N. (1991). Review of An ontology of art. *The Journal of Aesthetics and Art Criticism*, 49(1), 75–77.
- Yagisawa, T. (1999). Against creationism in fiction. *Philosophical Perspectives*, 15, 153–172.
- Ziff, P. (1972). What is said. In D. Davidson & G. Harman (Eds.), *Semantics of natural language*. Dordrecht: D. Reidel.