

October 2021

Worlds without End: A Platonist Theory of Fiction

Patrick Grafton-Cardwell
University of Massachusetts Amherst

Follow this and additional works at: https://scholarworks.umass.edu/dissertations_2



Part of the [Aesthetics Commons](#), and the [Metaphysics Commons](#)

Recommended Citation

Grafton-Cardwell, Patrick, "Worlds without End: A Platonist Theory of Fiction" (2021). *Doctoral Dissertations*. 2329.

<https://doi.org/10.7275/23777855> https://scholarworks.umass.edu/dissertations_2/2329

This Open Access Dissertation is brought to you for free and open access by the Dissertations and Theses at ScholarWorks@UMass Amherst. It has been accepted for inclusion in Doctoral Dissertations by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

**WORLDS WITHOUT END:
A PLATONIST THEORY OF FICTION**

A Dissertation Presented

by

PATRICK ELLIOT GRAFTON-CARDWELL

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

September 2021

Department of Philosophy

© Copyright by Patrick Elliot Grafton-Cardwell 2021

All Rights Reserved

**WORLDS WITHOUT END:
A PLATONIST THEORY OF FICTION**

A Dissertation Presented

by

PATRICK ELLIOT GRAFTON-CARDWELL

Approved as to style and content by:

Ned Markosian, Chair

Phillip Bricker, Member

Amie Thomasson, Member

Alejandro Pérez Carballo, Member

Daniel Altshuler, Member

Phillip Bricker, Department Chair
Department of Philosophy

DEDICATION

To JoAnna. For your steadfast love and encouragement.

The poet's eye, in fine frenzy
rolling,
Doth glance from heaven to earth,
from earth to heaven;
And as imagination bodies forth
The forms of things unknown, the
poet's pen
Turns them to shapes and gives to
airy nothing
A local habitation and a name.

Shakespeare, *Hamlet*

The universe (which others call the Library) is composed of an indefinite and perhaps infinite number of hexagonal galleries, with vast airshafts between, surrounded by very low railings. From any of the hexagons one can see, interminably, the upper and lower floors. The distribution of the galleries is invariable. Twenty shelves, five long shelves per side, cover all the sides except two; their height, which is the distance from floor to ceiling, scarcely exceeds that of a normal bookcase. . .

Jorge Luis Borges, *The Library of Babel*

"Your eyes have not failed you," Master Ultan said after a time. "Do you apprehend any termination to this aisle?"

"No, *sieur*," I said, and in fact I did not. As far as the candlelight flew there was only row upon row of books stretching from the floor to the high ceiling.

Gene Wolfe, *The Shadow of the Torturer*

ACKNOWLEDGMENTS

I want to thank my committee, Ned Markosian, Phillip Bricker, Amie Thomasson, Alejandro Pérez Carballo, and Daniel Altshuler for their many helpful comments and conversations along the way, as well as for their encouragement. I also want to thank Tim Juvshik especially for reading so much of my work, giving incredibly helpful comments on everything, and for friendship and encouragement.

Thanks to Justin Mooney, David Friedell, and the attendees of the UMass philosophy department dissertation seminars at which much of the material was presented, all of whom gave me very helpful detailed feedback.

I also received helpful feedback from commenters and audience members at a number of conferences at which some of this material was presented, including the 2018 SEP Conference, 2018 Metaphysical Mayhem, 2018 ASA Annual Meeting, 2019 Syracuse Graduate Philosophy Conference, CPA Congress 2019, 2020 ASA Annual Meeting, and 2021 APA Eastern Meeting. Thanks to those involved in those events and especially to the commenters and those who followed up with extra feedback and engaging conversations.

I owe many debts to a number of people who, along the way, have encouraged me, taught me, or created opportunities for me when they had no obligation to do so. I want to acknowledge some of those people, knowing that of course I will omit some. Thanks to Matt Jenson, Nathaniel Warne, Tim Pickavance, Dave Horner, Rick Floyd, Jeff Brower, Patricia Curd, and Vanessa de Harven for being mentors and offering training and encouragement to me. Thanks to those I've accidentally omitted as well (of whom I'm sure there are many) who helped me along the way.

I also want to thank my family, especially my parents and my parents-in-law, who have both been sources of deep and continual encouragement to me. And to my wife, who has been my steadfast partner and an encourager in hard times, thank you.

Thanks, once more, to Ned, who encouraged me, through his teaching and through conversation, to think about fiction (and art more generally). It took me a long time to find my philosophical “home”, and that was it. So, thank you.

Finally, I want to thank my God, through whom I have received strength and wisdom to complete this.

ABSTRACT

WORLDS WITHOUT END: A PLATONIST THEORY OF FICTION

SEPTEMBER 2021

PATRICK ELLIOT GRAFTON-CARDWELL

B.S., BIOLA UNIVERSITY

M.A., BIOLA UNIVERSITY

M.A., PURDUE UNIVERSITY

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Ned Markosian

I first ask what it is to make up a story. In order to answer that question, I give existence and identity conditions for stories. I argue that a story exists whenever there is some narrative content that has intentionally been made accessible. I argue that stories are abstract types, individuated by the conditions that must be met by something in order to be a properly formed token of the type. However, I also argue that the truth of our story identity attributions—sentences like, “Peter Jackson’s *Lord of the Rings*” is the same story as JRR Tolkien’s *Lord of the Rings*”—are relative to their context of use, so that the same attribution of identity might be true in one context or false in another because of a change in what’s being said about the works by that attribution in the differing contexts. From those existence and identity conditions, I draw the conclusions that *story* is best understood to be a phase sortal

of some abstract types, and that what it is to make up a story is to identify a type of narrative representation and its content, such that the content is neither based on actual events nor on the content of other stories to an objectionable degree.

In chapter 2, I ask what it is for a story to be complete. I argue for the general meta-thesis I call “Completion Pluralism”, viz. that there are many kinds of artwork completeness, many corresponding senses of ‘complete’, and no kind of artwork completeness is objectively more important than any other. I give a pair of arguments against those who would deny Completion Pluralism, endorsing Completion Monism instead. First, I show how Monistic analyses of the one (or one most important) kind of artwork completeness fail. I also show that each Monistic theory’s failure is contingent on the assumption of Completion Monism: the denial of Completion Monism dissolves each argument against the extensional adequacy of extant theories of artwork completeness. I argue that this is all best explained by the idea that Completion Pluralism is true.

Second, I argue that Completion Pluralism fits best with the variety of interests that we bring to questions about artwork completeness: there is no single or overriding concern we have when we want to know whether an artwork is complete. We shouldn’t expect completeness to be a single sort of thing, then. I show how to put Completion Pluralism to work by identifying different kinds of artwork completeness in artworks. Finally, I introduce a new kind of completeness particular to stories, what I call “narrative completion”, and show how it’s distinct from other kinds of completeness discussed.

In chapter 3, I argue that the distinction between fictional and non-fictional stories is the degree to which stories are made-up. In short, a fictional story is a made-up story. I first provide a couple desiderata for thinking about theories of fiction, then show how the two most popular extant theories of fiction—the prop theory and the genre theory—fall short in meeting one or more of those desiderata to some

extent. I then show how the theory that fictions are made-up stories better meets the desiderata, using my account of what it is to make up a story from chapter 1.

I turn to explicitly ontological matters in chapter 4. Addressing the question of what stories are, I show that both the theses that they are eternal and immutable and that they are created and historical seem plausible for independent reasons. I discuss a schema of an argument some philosophers who hold to the latter view (those I call “Creationists”) attempt to give for their view. However, I show how those who hold to the view that stories are eternal and immutable (“Platonists”) can account for all the same appearances as the Creationists, usually at least as well and sometimes with a better accounting. Similar, deductive arguments for Platonism will give rise to similar responses from the Creationist, however. In the end, then, the style of argument in particular is unlikely to advance the discussion much.

Finally, I attempt to advance the discussion by building an abductive case for Platonism. First, I vindicate a particular metaontological methodology for the debate at hand, a methodology that allows for “revisionary” outcomes in our ontologies of art objects. Second, I build a case for Platonism by arguing that it best explains a number of features of stories discussed in the previous chapters. Moreover, I argue that Creationism is committed to an implausible account of how some abstract objects come into existence. Given the combination of those facts, I argue that we should prefer Platonism to Creationism.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	vi
ABSTRACT	viii
CHAPTER	
INTRODUCTION	1
1. MAKING UP STORIES	6
1.1 Introduction	6
1.2 The Existence of Stories	8
1.3 The Identity of Stories	13
1.4 Making Up Stories	40
2. COMPLETING STORIES	47
2.1 Introduction	47
2.2 Against Completion Monism	51
2.3 In Favor of Completion Pluralism	66
2.4 Putting Pluralism to Work	68
2.5 Conclusion	75
3. FICTIONS ARE FABRICATION	76
3.1 Introduction	76
3.2 Desiderata for a Theory of Fiction	79
3.3 Fiction Is Not a Prop	84
3.4 Fiction Is Not a Genre	95
3.5 Fiction Is a Fabrication	102
3.6 Conclusion	111

4. HOW NOT TO ARGUE FOR CREATIONISM ABOUT STORIES	113
4.1 Introduction	113
4.2 Temporality	119
4.3 Temporal Flexibility	122
4.4 Modal Flexibility	126
4.5 Essentiality of author	128
4.6 Essentiality of Origin	131
4.7 Teleology	133
4.8 Conclusion	136
5. HOW TO ARGUE FOR PLATONISM ABOUT STORIES	140
5.1 Introduction	140
5.2 Vindicating Discovery	141
5.3 In Favor of Platonist Ontology	153
5.4 The No Creation Argument	161
5.5 Causal Abstracta	171
5.6 Conclusion	178
BIBLIOGRAPHY	179

INTRODUCTION

Fictional characters provide us with a number of puzzles. One puzzle is that we tend to form intense emotional attachments to fictional characters, even while having no interaction with them. Another, related puzzle is that we form these attachments even while denying that they exist. Gregory Currie calls this the “Paradox of Caring”.¹ It doesn’t seem coherent to fear or love something one doesn’t think exists. It seems especially bizarre when one considers the time and energy we devote to emotionally engaging with fictional characters.

Another puzzle is just what to say about those assertions of non-existence. Are people really committed to the non-existence of fictional characters? If so, what is the content of their emotions and beliefs apparently about those characters? And when someone denies that, say, Harry Potter exists, what exactly are they referring to?²

Another puzzle, for those who think fictional characters exist, is to say just what kinds of things they are. What sort of ontological category could fill all the artistic, psychological, and metaphysical roles that fictional characters do?³

Given these ways that fictional characters are puzzling, and given the integral role that thinking and caring about them plays in our lives, it’s no wonder that there

¹See Currie (1997). For more discussion of that paradox see Kind (2010), Walton (1978), and Tullman and Buckwalter (2014).

²Negative existentials have been a focal point of much of the debate concerning the existence of fictional characters. For a summary, see Friend (2007). For some realist attempts to make sense of these statements, see Kripke (2013), Thomasson (2003), and von Solodkoff (2014).

³The question of what fictional characters could be, were there any, is raised and discussed extensively by Thomasson (1999).

has been a lot of work devoted to theorizing about them in the last few decades. Philosophers have taken steps towards solving some of the puzzles mentioned above, though in many cases they've been at impasses, with conflicting solutions proposed and no clear way to adjudicate.

Meanwhile, there has been a relative paucity of attention given to a pair of related topics that are relevant for theorizing about fictional characters: the metaphysics of stories and the activity of making up stories. It's typical of the philosopher of fiction to ask a question about fictional characters and to keep her focus only on the characters, taking as evidence for her work certain facts about how we, the audience of various stories, talk about and think about the characters in the stories. There are a few reasons why it's worth taking a step away from the questions about the characters, though, to consider stories instead, thinking as well about the activity of making up stories in order to consider how we came to be related to these things in the first place.

First, fictional characters find their homes, so to speak, in *fictions*. We, as audiences of fictional works, and as philosophers, may be especially interested in fictional characters and in the puzzles they present us with, but it's implausible that very many narrative artists begin (much less complete) their work with the intention solely of making up a character. Much more common, we can assume, is that they begin with the intention of making up a *story*, and reference to characters comes about as the artist carries out that process. Stories, then, are artistically fundamental. It is in virtue of the artist's making up a story that she makes up the characters in it. Even in cases where an artist begins her work with the intention to make up some character, it's unlikely that she could do so at all robustly (i.e. in a way likely to provoke interest and emotional attachment) without portraying her character in some story. We may be interested in portraits of fictional people, where only the suggestion of narrative is at play. But it will be difficult for us to become fully invested in a charac-

ter until we come to know about them the way we do through narrative art. Because of this dependency of characters on stories for emotional significance, it seems like if we really want to understand puzzles like the Paradox of Caring, then we'll need to understand what stories are, and how we are related to them.

Moreover, understanding what stories are should be a topic of philosophical interest in its own right. They present many of their own puzzles, similar to those of fictional characters. We are emotionally invested in the *events* of stories as much as we are attached to their characters. But this is just as paradoxical in the case of events as in the case of characters, since we're just as likely to admit that these are things that have not happened. Why should we care about them, then? Unlike in the case of characters, philosophers are generally not inclined to the view that stories don't exist, but just as with fictional characters it's difficult for the realist about stories to say just what they are, given that they exist.

If we want to get at the truth about stories and about fictional characters, though, it's plausible that a good source of evidence for thinking about both is to construct a theory of the activity by which we come to be related to them, i.e. the activity of making up a story. There are at least a couple reasons why we should be interested in such a theory. First, a theory of what it is to make up a story can give us insight into the ontology of stories and of fictional objects. That is, when we ask ourselves whether there are any such things as stories or fictional objects, and if so what they're like, we are more likely to arrive at the truth if we get clear about what making up a story (and so inventing characters) involves.

Second, the activity of making up stories is of philosophical interest regardless of what the truth about the existence of fictional objects is. Take the following argument:

1. Either fictional objects are created, or they're discovered, or they don't exist.
2. If they're created, then they're created when someone makes up a story.

3. If they're discovered, then they're discovered when someone makes up a story.
4. If they don't exist, then talk putatively about fictional objects is made true (or in some way "good") by someone making up a story.
5. If fictional objects are created when someone makes up a story, or discovered when someone makes up a story, or they don't exist but our talk putatively about them is made true (or in some way "good") by someone making up a story, then we should be interested in a theory of what it is to make up a story.
6. So, we should be interested in a theory of what it is to make up a story.

The point is that, even though we should hope that a theory of what it is to make up a story will settle the question about the ontology of stories, it's also something we should expect to be philosophically interesting and worth pursuing whatever the answer to that question turns out to be.

For these reasons, instead of focusing on fictional characters, I want to direct some attention to the nature of stories and the activity of making them up. I will answer questions like,

- When is someone making up a story?
- What is the difference between making up a story and merely telling a story?
- When is a story finished?
- What is the difference between fictional and non-fictional stories?
- What are the existence and identity conditions for stories?
- Are stories created or discovered?

I will start by providing existence and identity conditions for stories, which will lay the groundwork for a theory of what it is to make up a story. I'll move from there

to give an account of in how many ways a story can be complete. I'll use my theory of what it is to make a up a story to give a theory of what it is for a story to be fiction, an account I call the Fabrication Theory of fiction. In the final two chapters, I provide a cumulative case for Platonism about stories, the view that stories are eternal, abstract types. I begin by showing that the traditional case for Creationism is dialectically ineffective: Platonism can account for all of the features of stories that Creationism can. Next, I show how there are some features of stories that Platonism seems distinctly better suited to account for. Finally, I offer an argument against the plausibility of Creationism from the causal inertness of abstract objects. In the end, I take it that we have reason on balance to favor Platonism to Creationism for stories.

CHAPTER 1

MAKING UP STORIES

1.1 Introduction

I would like to begin by giving an account of what it is to make up a story. This activity (i.e. making up stories) is one that artists of many kinds engage in: novelists, filmmakers, comic book artists, and occasionally even authors and poets make up stories. While stories are ubiquitous in popular art, we don't have a uniform vocabulary for referring to the people who make them up, or for describing what it is to make them up. For clarity's sake, I'll call this activity "authoring"—acknowledging that we commonly reserve this term for the creation of literary works, whether narrative or not—and the people who engage in it "authors". The aim of this chapter, then, is an account of the activity of authoring, or an account of what it is in virtue of which someone counts as the author of a story.¹

An account of authoring, as I'm conceiving of it, is a list of necessary and sufficient conditions for when someone counts as making up a story. For the purposes of this analysis, I'll restrict my attention to acts of authoring that are artistic (as opposed to, for example, lying or scientific theorizing) and to stories that are art objects (as opposed to, for example, lies or theories).² I take artistic stories to be the paradigm

¹Authoring is distinct from another narrative artistic practice for which we already have a name: storytelling. Authoring involves introducing an original story to an audience. Storytelling involves the presentation of a story, regardless of the story's originality. Authors are very often storytellers for their stories. But the relationship is orthogonal: authors don't always tell their stories, and storytellers are not always authors of their stories.

²I'll say more below about what I take to be distinctive about artistic stories. One thing that's worth noting up front is that a mark of a story's being art is that there are norms for how one

for our understanding of the activity of authoring. When we refer to lying or to scientific theorizing as making up a story, the authoring of artistic stories serves as the salient activity similarity to which licenses that description. Clearly, we often rightly describe other activities as making up a story, but my thought is that we do so by extension.

In order to give an account of authoring, I turn first to the existence and identity conditions for stories. In order to know what it is to make up a story, we need to know what conditions must be satisfied for there to be a story, distinct from any other story. In §2 I examine the existence conditions for stories. In §3 I turn to identity conditions for stories. In §4 I show how the results of the previous two sections can be applied to provide a theory of authoring in terms of novel content.

Another question that's intimately related to the question of what it is to make up a story is the question of what stories are. An account of the existence and identity conditions for stories will get us some way towards answering the question of what stories are, as will some accounts of authoring. To give identity and existence conditions for stories is to say just when there is a story, distinct from any other stories. It's a matter of debate, however, whether giving anything beyond the specification of existence and identity questions is required for saying what an object really is. At the risk of leaving out an important issue, I am going to defer any prolonged, explicit discussion of the metaphysics of stories in this chapter. Instead, I aim to first take on the question of what authors are doing when they make up stories, utilizing the answer as evidence for giving a metaphysics of stories in a later chapter.

ought to interact with the story that are a result of its being the work of an author, e.g. norms of canonicity or artistic authority. Artistic stories are such that there is a canonical version of them, and we ought to respect the authority of the artist over such stories at least to a certain extent. A cautionary tale I make up to teach my children may resemble an artistic story in many ways, but the same norms won't apply.

Since I'm not going to say here what I take stories to be (and since that will be a topic of discussion in a future chapter) I'm not going to provide a definition of 'story'. However, here are a few things I take to be true of stories. For one, stories have content. The content of a story is what's true in the story, or what's represented as being the case in the story.³ I take the content of stories to be propositional in nature. My understanding of this is fairly open, not requiring any specific theory of content (e.g. of the ontology of content) or of truth in fiction to hold. Likewise, stories have versions. A version of a story is a telling of the story. We can distinguish between types of versions and token versions. For example, there is the film version of *Star Wars: A New Hope*, and there's the novel version of it, a distinction at the level of types. There's also my copy of the *Star Wars: A New Hope*, despecialized edition on blu-ray, as distinct from the copy of the film I had on VHS as a kid (or the copy you might have had on VHS), a distinction at the level of tokens. We can think of a type of version as a way of telling a story; a token version, on the other hand, is an instance of a particular way of telling a story. I'll refer back to the content of stories, as well as to versions of them, throughout, with these understandings in mind.

1.2 The Existence of Stories

To help illuminate the notion of authoring, it will help first to understand under what conditions a story comes into being. I'll start, then, with the existence conditions for stories, the necessary and jointly sufficient conditions for there to be a story. An account of the existence of stories will make it clearer what an author must do to be a source of original narrative content, i.e. to introduce us to a new story.

One clear necessary condition on the existence of a story S is that there be some content c that is the content of S . In order for there to be a story, there has to

³For accounts of truth in fiction, see Lewis (1983b), Currie (1990), Byrne (1993), Lamarque and Olsen (1994), and Badura and Berto (2018). Truth in non-fiction is just truth.

be a way the story goes. That “way the story goes” makes up at least part of the content of the story, i.e. what’s true in the story or what’s represented as being the case either explicitly or otherwise. When some content c is the content of a story S , I’ll sometimes say that S has c .⁴

Not just any sort of content can be the content of a story. The content in question must be *narrative* content. Narrative content is content that represents that a sequence of events occurs, and that there are some important relations between the events in question.⁵ I will remain agnostic as to what kind of relations between events count as “important” in the sense that makes for narrativity. Carroll (2001), Velleman (2003), and Currie (2006) have each suggested accounts. Each advocates that a narrative must consist of a sequence of events that’s connected in some salient fashion, but there are important differences in the way they develop the notion of salience in question. I will also remain agnostic as to how much of a story’s content must be narrative content. Some components of any story’s content will be non-narrative:

⁴This way of describing things might seem to imply an endorsement of what Irvin (2006) calls Critical Monism, the view that there can be at most one correct attribution of meaning (or content) to a work. I’m not committed to Critical Monism. I speak here as if a story has just one content or meaning because that’s simpler. What I really mean is that there have to be some facts about what’s true in the story, even if it’s also true that there are multiple, inconsistent but correct attributions of meaning for a given story. In such a case, we might say that a given story has more than one content, we might say it’s content is indeterminate on some matters, or we might say that the norms of correct attribution of meaning outstrip what is strictly speaking the content of the story (so that an interpretation can be correct even while in some sense *going beyond* what the story says). I’m not going to settle what we should think about these kinds of cases here. My project is to get clear on what stories are, not to further our understanding of truth in fiction. Suffice to say, the account I develop of the nature of stories is consistent with Critical Pluralism, the view that there can be more than one correct attribution of meaning to a work. The reader who is especially sympathetic to Critical Pluralism can reformulate anything I say in Pluralist terms.

⁵One point of contention here is discussed by Nanay (2009), who advocates for a theory of narrative on which a story may represent only one event. Nanay’s argument focuses on narrative pictures (e.g. Jacques-Louis David’s *The Consecration of the Emperor Napoleon and the Coronation of Empress Joséphine on December 2, 1804*), which Nanay contests have narrative content even while representing only a single event. I’m skeptical of Nanay’s argument, but I will remain neutral on whether the sequence of events that constitutes a narrative can be composed of a single event. The question of whether narrative content may represent only one event goes back much further. Others who have discussed it include, for example, Genette (1980) and Labov (1973). Thanks to Daniel Altshuler for pointing this out.

character identities or setting details, for example. I won't draw any precise threshold for determining how much of a story's content needs to be narrative content.

Importantly, any view that makes narrativity relative to salience also makes it relative to a community in a context. Something is salient only in a context to a community and may fail to be salient outside of that context or that community. This is because salience is a function in part of the interests of a given community, and interests vary between situations and between groups. The result of this view of narrative is that content is never narrative in the abstract, but is narrative relative to the interests of particular people in particular situations. Of course, the correct theory of narrative may privilege an author's interests, or those of the (actual or hypothetical) audience, when accounting for whether some content is narrative content. It may privilege either the context of creation or a context of reception as well.⁶ The point remains that allowing narrative to be a function of salience makes it relative to some community or other (even if just a community of one).

I take this result to be plausible. It seems that cultural context plays an essential role in our judgement of any particular artwork's possession of narrativity. When we want to know of a historical work of literature whether it was a narrative or what distinctly narrative characteristics it has, it's important to look to its context of origin. Likewise, when attempting to understand ancient works of narrative art, it's common to attempt to gain some understanding of the context of reception for the work. In order to understand how the events portrayed in the work are saliently connected, we need to understand the context of the work's reception. Finally, we can imagine how a story that possesses a high degree of narrativity for us (in virtue of ordering its events under a set of relations that are especially important or salient to our minds) might possess a relatively low degree of narrativity to, say, an alien species with a

⁶Thanks to Alejandro Pérez Carballo for pointing this out to me.

radically different form of consciousness (for example as the Heptopods portrayed in Ted Chiang's *The Story of Your Life* (1998), who experience all the events of their life in a "simultaneous" mode of consciousness, aware of their past, present, and future at every moment at which they exist). The fact that some stories have satisfying narrative connections or give satisfying narrative explanations of some of their events in terms of others is at least in part a feature of the way those stories interface with our own psychological makeup.

Making narrativity relative to community and context doesn't rule out the idea that some content might necessarily be narrative. That is, it's possible on this account that for some content c , for any possible context and any possible community (or for any possible context for any community sufficiently like us), c would count as narrative in that context for that community. Thus, while it seems that our psychology plays an important role in the construction of narrativity, we can still make sense of the intuition that some stories would be stories no matter what.

Once we've established that the existence of narrative content is necessary for the existence of a story, we should wonder if it's also sufficient. I don't think it is. For there to be a story (in the sense we're restricting our attention to: stories as art objects) it seems that the content of the story must be intentionally represented in a determinate fashion, or that there must at least be instructions specifying how the author intends that the content be represented. This is to say that for there to be a story, qua art object, there has to be some version of it.

There are two aspects to this condition. One is that the content be made shareable. One way this can be achieved is by representing the content in some determinate form or another (say as a particular novel or film). Another is by providing instructions as to the canonical representation of the content (as we might conceive of screenplays as doing). If this constraint hasn't been met, then it's indeterminate what the object of aesthetic evaluation is. One fairly plausible constraint on art objects (especially

representational art like narrative art) is that for something to be an art object it must be susceptible of aesthetic engagement and evaluation. It seems that no story could be an art object without its content being presented in a way that's accessible to and evaluable by an audience. This isn't requiring that it actually be shared, or that it take any particular form or other. Rather, the idea is that art is in principle sharable and that it can't be entirely formless. Say that when an author represents the content of their story or provides instructions for how to do so they have made that content accessible.

The other aspect of this condition is its intentional component. It's common to say that nothing can be an art object unless it's (in some sense) the result of intentional action. I agree that for there to be a story (in the sense we've restricted ourselves to here), an author must intentionally make some narrative content accessible (either by representing it or by specifying instructions for a canonical representation). There are various ways we might imagine that someone could make narrative content accessible to us by accident. The result would not be an art object, and so not a story in the sense we're concerned with.⁷

That completes my account of the existence conditions of (artistic) stories. A story S exists iff there is some narrative content c such that c has intentionally been made accessible and S has c . These conditions just tell us what it takes for there to

⁷One might imagine another, nearby case that seems problematic, given what I've just said: imagine the waves and wind on a beach erode a group of rocks until (by chance) those rocks are shaped so as to spell out the words of a story. According to my proposed definition, this would not count as an artistic story, although we might imagine the story (once again, by pure chance) being rather good, something literary critics would hail as a great find. My understanding of such cases (if ever any do occur) is that they are cases of appropriation: until someone comes along and reads the "words" in the rock, there really aren't any words in the rock. In coming along and reading (and appreciating) the story, the audience appropriates the shapes in the rock as words, text, narrative, and work simultaneously. In that way, I take it, this kind of case is not all too different from the popularly discussed case of a driftwood winerack, in which someone appropriates a piece of driftwood for use as a winerack. I take cases of appropriation to be sufficient in most cases to satisfy the intentional component necessary to make something an art object. Thanks to Daniel Altshuler for raising this case.

be a story. It will help also to have an account of the conditions under which some stories are the same or distinct.

1.3 The Identity of Stories

Determining the identity conditions for stories will also help us get a sense of what it is to make up a story. The identity conditions for stories are the necessary and jointly sufficient conditions under which some stories are identical. The identity conditions for stories also function as individuation conditions for stories, since two stories are distinct just in case they are not identical.

To get a sense of what we're looking for when we're after the identity conditions for stories, consider the following two very short stories:

Maggie's Wild Ride

Maggie felt the air whip through her hair as she piloted the biplane through the mountains, weaving a path through the open space between peaks. She had never felt this exhilaration and weightlessness before. She didn't even remember learning to fly! As that thought began to dominate her attention, the biplane veered suddenly downwards, plummeting towards a mountain face! Maggie awoke with a start, safe and warm in her bed.

Maggie's Wild Glide

Maggie felt the air whip through her hair as she piloted the hang glider through the mountains, weaving a path through the open space between peaks. She had never felt this exhilaration and weightlessness before. She didn't even remember learning to fly! As that thought began to dominate her attention, the hang glider veered suddenly downwards, plummeting towards a mountain face! Maggie awoke with a start, safe and warm in her bed.

When we're wondering what the identity conditions for stories are, we're interested in questions like whether "Maggie's Wild Ride" is the same story as "Maggie's Wild Glide". What's the right answer, and why?

It's tempting to try to individuate a story by its content. The way a story goes, or what's true in a story, can strike us as essentially just what a story is. After all,

consider “Maggie’s Wild Ride” and, for example, Ted Chiang’s (2007) short story “The Merchant and the Alchemist’s Gate”. They are clearly distinct stories. Plausibly, what makes them distinct is the almost complete difference in setting, character, and event details they present, i.e. in their content.

Smuts (2009) examines several accounts of story identity that identify stories by their content. First is what he calls the COMPLETE ACCOUNT, according to which the story is individuated by “the continuum of events presupposing the total set of all conceivable details, that is, those that can be projected by the normal laws of the physical universe.”⁸ (Smuts 2009, 6) A story’s total actual content isn’t sufficient by itself to individuate the story on this view. What’s needed is an extension of the story’s content that includes “all conceivable details” projectable given the actual laws of nature (modulo, we might assume, cases where the laws of the story are meant to differ from the actual laws).

I can see two ways of disambiguating the COMPLETE ACCOUNT. One is to take seriously the language of conceivability. Allow a maximal conceivable extension c^+ of some content c to be the result of a process by which one adds to c every imaginable detail, constrained only by nomological consistency. Then, we can understand the COMPLETE ACCOUNT as follows.

COMPLETE CONCEIVABILITY ACCOUNT: For any story S with content c , S is individuated by a maximal conceivable extension c^+ of c .

Smuts (2009, 6) points out a fairly obvious problem for the COMPLETE CONCEIVABILITY ACCOUNT: “indefinitely many supplementary details could be compatible with the contents of the story. . . , but not all such sets of details are consistent with each other.” That is, for any given story, there are many conceivable extensions of

⁸This account is taken verbatim from Chatman (1978, 28).

its content. In any particular case, it won't be clear which extension is to play the individuating role after all, on this version of the COMPLETE ACCOUNT.⁹

On the other hand, we might not take the COMPLETE ACCOUNT very seriously when it uses the language of conceivability, instead understanding it in terms of entailment. Allow a maximal logical extension c^+ of some content c to be the result of a process by which one adds to c every detail that is logically entailed by c and the laws of nature. A maximal logical extension of some content c will differ in what it represents from at least some maximal conceivable extensions of c . For example, some maximal conceivable extension of Disney's *Frozen* represents that Anna and Kristoff have a child with brown hair, but the maximal logical extension of *Frozen* doesn't. Then, we can understand the COMPLETE ACCOUNT as follows.

COMPLETE ENTAILMENT ACCOUNT: For any story S with content c , S is individuated by a maximal logical extension c^+ of c .

The COMPLETE ENTAILMENT ACCOUNT seems to avoid the problem raised for the COMPLETE CONCEIVABILITY ACCOUNT. There will only be one complete logical extension of the content of any given story.

However, there doesn't seem to be any good reason to accept the COMPLETE ENTAILMENT ACCOUNT. If we think a story is individuated by its content, why would we look to an *extension* of that content to individuate it, and not at the content of the story alone?¹⁰

⁹Cf. Lorand (2001) for additional reasons to reject any attempt to reach a conception of stories as maximal in their portrayal of reality.

¹⁰Phil Bricker points out to me that one might take a story's content to be closed under logical entailment, so that what I've called the maximal logical extension of a story's content would just be its content. In that case, the COMPLETE ENTAILMENT ACCOUNT will turn out no different from the view that says a story is individuated by its total actual content. Whether we should treat story content as closed under logical entailment is a matter of debate, however. Those who say we should not will point to the various ways in which the content of stories seem to be incomplete, representing nothing (either explicitly or implicitly) sometimes even concerning matters about which their explicit content would entail some facts.

Moreover, there are at least two reasons not to accept the COMPLETE ENTAILMENT ACCOUNT. First, if the logic used to generate the maximal logical extension of a story's content is classical, then the COMPLETE ENTAILMENT ACCOUNT will be unable to individuate impossible fictions. At least, any fiction that contains a contradiction (of which there are many) will generate an explosive maximal logical extension, identical to the extension of any other fiction that contains a contradiction. Intuitively, however, there are many distinct stories with contradictory content.¹¹ Wildman and Folde (2017) provide compelling evidence that there are, in fact, distinct fictions with the same maximal logical extensions: in such fictions classical logic is explicitly true and a contradiction holds. Thus, the maximal logical extension of any such fiction includes all propositions.

A second problem is that it seems like the COMPLETE ENTAILMENT ACCOUNT will have a hard time making sense of what has been called the “transposability of the story.” In discussing some constraints on an acceptable account of the identity conditions for stories, Smuts (2009, 6) argues, “the primary assumption of much narratology is the ‘transposability of the story’ thesis.” The thesis is that one story can be told using disparate media, in different artistic styles, even when formal differences between versions introduce differences in content. For example, the story of *The Watchmen* has been told both as a graphic novel (published in 1986-87) and as a film (released in 2009) (and continued as a television show in 2019). The transposability of the story allows for difference in form (e.g. in medium, focal perspective, or narrative structure) as well as content between versions of a story.¹²

Whatever one says about the identity of stories, one desideratum for any account is to explain the transposability of the story, i.e. to explain how the same story can

¹¹See Nolan (forthcoming) for a discussion of impossible fictions. See also Priest (1997) for a putative example of an impossible fiction.

¹²For more on the transposability of the story, see Chatman (1978) and Bal (2017).

have many versions, each of which can differ in form and content. It turns out this is a difficult constraint to meet. Different versions will have different contents since they'll vary with respect to the (sometimes trivial but sometimes seemingly important) details of the story. Different contents will have different maximal logical extensions. It's not clear then how two versions could be versions of the same story on the COMPLETE ENTAILMENT ACCOUNT.

As an alternative way of individuating stories by their content, Smuts (2009, 6) introduces the STRICT ACCOUNT, according to which a story is individuated by the explicit (or what Wildman and Folde 2017 refer to as the “primary”) content of the story. Like the COMPLETE ACCOUNT, the STRICT ACCOUNT individuates stories according to their content, but it doesn't require anything beyond the primary content of the work and so doesn't fall prey to some of the initial difficulties of the COMPLETE ACCOUNT. For example, as Smuts points out, one can't point to any contradictions in the secondary (i.e. imported, imagined, or entailed) content of the story to conjure up questions as to which story is under consideration. Likewise, the STRICT ACCOUNT can distinguish between the distinct “unlimited” fictions proposed by Wildman and Folde (2017) because it individuates stories by means of their primary content, which can differ between such stories.

The STRICT ACCOUNT also provides a clear answer to the question with which we began this section: according to this account, “Maggie's Wild Ride” and “Maggie's Wild Glide” are distinct stories. Their primary content differs in that one represents Maggie as flying a biplane while the other represents her as flying a hang glider. Since stories are individuated by their primary content on the STRICT ACCOUNT, it tells us the stories must be distinct.

This is not obviously an advantage of for the STRICT ACCOUNT, however. Like the COMPLETE ENTAILMENT ACCOUNT, the STRICT ACCOUNT doesn't seem to be

able to account for the transposability of the story. Take Smuts's own example to illustrate this:

Imagine two versions of *I Am Legend* where the main character has a different eye, hair, or skin color; perhaps a translator made a slight mistake and wrote that Neville has a brown instead of blonde beard. If the story is all of the event, character, and setting details relayed by the discourse, then the two hypothetical versions of *I Am Legend* do not present the same story.... So, according to the strict theory these two instances, with just a minor difference in beard color, could not be the same story. (Smuts 2009, 6)

If the only difference between two versions is what they represent of a character's beard, we wouldn't usually judge them to be two different stories. However, the STRICT ACCOUNT has the consequence that they are. The STRICT ACCOUNT doesn't seem to fit with the idea that there can be two versions of a story that differ only in minor details of what they represent.

What we've found in examining these various attempts so far is that individuating a story merely by means of its content is not a promising strategy. Individuating the story by means of its total content leads us to uncertainty about what the story is that we're focused on, since versions differ in content. Individuating the story merely by means of primary content makes clear what story we're focused on, but it leaves us with the wrong results about which stories are distinct.

What we need is a theory that allows for distinct versions of the same story (to make sense of transposability). One potentially helpful observation is that story versions seem like types rather than tokens. The same version can be tokened in more than one copy or instance, for example, and we don't think versions are located where their tokens are. Interestingly, if stories have multiple versions, then it also seems as if stories are types to which the versions belong. Smuts (2009, 9) introduces another account, the HIGHER-ORDER TYPE ACCOUNT, that seems to make sense of this. According to the HIGHER-ORDER TYPE ACCOUNT, the story is a higher-order type, i.e. a type of type. On this account, stories are individuated by a set of important

attributes, viz. the attributes something must possess to be a token of the type.¹³ Two story types are the same if and only if they have the same important attributes. Two token stories are the same if and only if they are tokens of the same type.

The HIGHER-ORDER TYPE ACCOUNT attempts to provide an answer to the question of what the identity conditions for stories are by saying more about what stories are. Because it identifies a story with a type, it can allow that only some of the attributes of the story are important for determining whether something is a token of it. Thus, whether two works are instances of the same story (even if they differ in certain aspects of their form and content) depends on which attributes of the story are important.

This account bears important similarities to the account of musical works argued for in García-Carril Puy (2019). On García-Carril Puy's account, musical works are higher-order types with versions as typical instances. This is meant to account for the way that musical versions relate to their parent works. For example, Tchaikovsky's *Symphony No. 2* is a higher order type, according to this view, that has both its 1872 and 1879 versions as instances, which in turn have as instances their various performances. The HIGHER-ORDER TYPE ACCOUNT says likewise that stories are types of types with versions as their instances and copies as instances of the versions.

An important question for the HIGHER-ORDER TYPE ACCOUNT is which attributes of any given story are the essential (or normative) attributes of it. The proponent of this account needs a story about just which attributes of a story matter and which don't. This question comes into focus when we consider once again the

¹³While Smuts doesn't make use of this idea, it is also open to the HIGHER-ORDER TYPE theorist to employ the notion of a norm-type, an idea pioneered by Wolterstorff (1980) and used by Dodd (2007). On Dodd's understanding of the notion, a norm-type is a type such that some properties are "normative within" it. It's in virtue of this fact that norm-types (and not other types) admit of improperly formed tokens. A Higher-Order Type theorist about stories might take stories to be norm types and give an account of which properties of a story are normative within them, still allowing that a discourse could tell a story *S* even if it lacked some property that was normative within *S*.

question of whether “Maggie’s Wild Ride” and “Maggie’s Wild Glide” are the same story or not. For the HIGHER-ORDER TYPE ACCOUNT, it seems to depend on what the core features of each story are.

Smuts (2009, 10) points out that “the elements that are essential for any given story will be highly idiosyncratic.” This is to say that no general account can be given by the advocate of the HIGHER-ORDER TYPE ACCOUNT of which attributes of stories are important for individuation. In some cases, the essential attributes will be that some event-types occur. In some cases, we’ll judge that some character details (like race, gender, or background) are essential. In other cases we might even judge that particular formal features are essential. (Perhaps, for example, we might judge that *A New Hope* wouldn’t be the story it is without the battered, space frontier aesthetic).

Smuts points out that we need an explanation for the variability in which features seem essential to a given story. He thinks the best explanation is that the elements we judge as essential for a particular story are those that are most salient to our understanding, impression, and memory of the story. This gives rise to what is a significant problem for the HIGHER-ORDER TYPE ACCOUNT. Smuts puts the problem this way:

It seems that the elements that we think are essential for a particular story are those that are salient—those that stand out in our memories, those that grab our attention. The problem is that the salience of any given character, event, or setting detail is largely a product of the presentation. If salience is sometimes relevant, retelling a story would not simply be a matter of what was told, but of how it was told. (Smuts 2009, 10)

It seems as if we judge a characteristic of a story as essential to the story because we experience that characteristic as salient in a larger whole, viz. in a version of the story we’re acquainted with. But salience is a function (at least in part) of how the particular version of a story interacts with our psychologies on a given occasion.

The problem is that if what's salient is a function in part of what version is being presented, then it's not so clear that the HIGHER-ORDER TYPE ACCOUNT can make sense of the transposability of the story. This is because what's salient in one version often won't be salient in another. Changing the medium of presentation or the ordering of events in the narrative, for example, can change what stands out to us (and thus what's salient). But if the identity of the story depends on which attributes of the story are salient, and what's salient depends on which version of the story is under consideration, then to tell the same story will often require sticking to the same version (and perhaps even presenting that version in the same context). Transposability tells us that any story can be told in multiple versions, each varying in its form or content. The HIGHER-ORDER TYPE ACCOUNT seems to have the consequence that at least some stories can only be told one way. Making sense of transposability was supposed to be the virtue of the HIGHER-ORDER TYPE ACCOUNT, but if it can't divest itself of the need to rely on the notion of salience (and thus on the means of presentation), there doesn't seem to be any special reason to accept it.

Worse for the higher-order type theorist: the HIGHER-ORDER TYPE ACCOUNT collapses an important distinction between story identity and story type. We want to be able to capture the difference between story identity claims (e.g. someone might ask, "Is the movie *To Kill a Mockingbird* the same story as Harper Lee's novel?" and might sensibly be answered in the affirmative) and story type claims (e.g. "*The Epic of Gilgamesh* and Homer's *Odyssey* are the same type of story"). As Smuts (2009, 11) points out, the HIGHER-ORDER TYPE ACCOUNT seems to collapse this distinction. Take for example the Cinderella story, a story-type in which a persecuted heroine receives a quick reversal in fortune. On the HIGHER-ORDER TYPE ACCOUNT, to tell the same story is to tell a version of a type of story, so it seems on that account that we should take all the cases of folks (often in different cultures and at different

times) making up Cinderella stories to be cases of their telling the same one story, even given all the variation in character, setting, and event details. The HIGHER-ORDER TYPE ACCOUNT can make sense of our talk about story identity, but only by reducing it to talk about typehood, and that's theoretically costly.

This brings to light another desideratum for an account of the identity conditions of stories: whatever account we give should be able to make sense of the distinction between story identity and story type. When people in different cultures and at different times make up Cinderella stories, they're each making up the same type of story, but certainly not in every case are they making up the same story. We want an account that makes sense of both the type/identity distinction and the transposability of the story.

It's time to take stock of where we've come thus far. The COMPLETE ACCOUNT was rejected for a variety of reasons. The STRICT ACCOUNT was rejected because it failed to account for the transposability of the story. The HIGHER-ORDER TYPE ACCOUNT was rejected because it seemed unable to account for either the transposability of the story or the story identity/type distinction.

It might seem at this point that individuating stories according to their content or form is the problem, since either is subject to change or variation across versions. If we can avoid looking to content or form at all in our identity conditions, then perhaps we have a better chance at making sense of transposability. In searching for an alternative, we can look to another literature for guidance. On the question of the identity conditions for words, Hawthorne and Lepore (2011, 31) discuss the proposal that words are individuated by something about their history, viz. their originating event.¹⁴ Irmak (2019, 1150) points to a problem for such a view: it seems that words can undergo fission and fusion, so that distinct words can have the same originating

¹⁴Thomasson (1999) seems to embrace a view like this for literary works.

event (because of fission) or that one word can have multiple, distinct originating events (because of fusion). The same seems like it's true of stories. An author can begin the work of authoring a story, and the result can be two very different stories. Likewise, one story might be the end result of two authorial projects that were at an earlier time distinct.

Irmak (2019, 1150-1) does take up the suggestion that words are individuated by their history. However, instead of looking only to a word's originating event, Irmak suggests we look to the whole of its history. That is, words w_1 and w_2 are identical iff they have the same history. The same account might seem plausible for stories. Stories s_1 and s_2 are identical iff they have the same history. Call that the HISTORICAL ACCOUNT. The benefit of the HISTORICAL ACCOUNT is that it seems it can meet both of the constraints above. Since the HISTORICAL ACCOUNT doesn't individuate a story by either its form or content, there's no reason to think it can't account for the transposability of the story: the same single generic entity may vary in its features across time and throughout its instances. The HISTORICAL ACCOUNT can also distinguish in principle between story identity and story type. What's needed is a supplementary account of story types that's given in terms of something besides history, perhaps similar to what the HIGHER-ORDER TYPE ACCOUNT says about core or essential features: stories with the same core features are of the same type. The account of story type could go a number of different ways. What's important is that the HISTORICAL ACCOUNT doesn't make the type/identity distinction impossible to draw.

The HISTORICAL ACCOUNT isn't without problems, though. Irmak (2019, 1151) points to one problem for the analogous account of words: "Perhaps the most important problem is determining how much similarity and difference two word histories must share to establish their identity..." The problem is that if a word's history is what individuates it, we seem to need an account of the individuation of word

histories to know how much change a word can undergo before there is a change in word history (along with an attendant change in word identity). The same is true for stories. At what point does change in content or form result in the genesis of a new story (because a new story history)? Irmak (2019, 1151) defers to “the scientific study of word histories” to provide particular answers in particular cases. We might opt likewise to defer to the narratologists or social scientists who study particular narrative works to tell us how much change the works have undergone and in which cases some changes resulted in a new work. But the lack of a principled answer as to how much change a single story can undergo leaves the account incomplete.

Moreover, there seems to be something deficient about the HISTORICAL ACCOUNT’s way of describing what’s happening in cases where a story changes. In cases where a story changes significantly and we deem the changes sufficient to declare a change in story identity, it will generally be most natural to describe what has happened as a change in the story’s content or form bringing about a change in story identity. That is, introducing enough difference in how the story goes or how it’s told can introduce a change in what story is being told. What the HISTORICAL ACCOUNT requires us to say, however, is that change in a story’s content or form can introduce a change in that story’s history, which in turn results in a change in story identity. It’s entirely unclear how a change in a story that isn’t itself a substantial change could introduce a change in what history belongs to the work, however. In fact, it’s unclear how a change in *any* object that isn’t itself a substantial change could introduce a change in what history belongs to the object. Any way that I change, for example, is a part of *my* history, unless that change is itself a substantial change, resulting in my non-existence. Thus, the HISTORICAL ACCOUNT seems to require us to posit an unmotivated variety of historical changes.

The Historical Theorist may respond with an alternative way of understanding their account: the way to understand the HISTORICAL ACCOUNT is merely as iden-

tifying the fact that identical stories share a story history. It doesn't require us to posit that non-substantial changes in a story result in a change in what story history belongs to what story. If the HISTORICAL ACCOUNT *does* allow that changes in a story's form and content are what are doing the work of determining story identity, then it may escape the objection at hand, but it will be utterly uninformative. In that case, the account will merely tell us that stories are identical iff they share a story history. But what is it for stories to share a story history? Intuitively, stories s_1 and s_2 share a story history h iff every state of s_1 is a state of s_2 . That is, on this account, two stories are identical iff any state of one is a state of another. That seems true, but it doesn't seem to really tell us anything about the identity conditions for stories. It seems true of stories, but it seems true of any kind of object whatever, and it doesn't tell us much about what distinguishes one story from another, in principle.

If the HISTORICAL ACCOUNT offers this trivial platitude, that would place more weight on the problem raised by Irmak, viz. that it's not clear how much change a story history can contain while being a history of just one story. The HISTORICAL ACCOUNT doesn't seem able to give any principled account of when change gives rise to something new, and this problem becomes starker when one realizes the view doesn't tell us anything about stories that doesn't seem true of objects in general.

Before I introduce my own account of story identity, let me introduce one more phenomenon a theorist ought to account for when giving an account of story identity: our attributions of story identity are highly sensitive to context. In many contexts, people will be inclined to say that Robert Mulligan's *To Kill a Mockingbird* is the same story as Harper Lee's *To Kill a Mockingbird*. However, in a context, for example, in which some folks are discussing how the centrality of the protagonist determines the moral perspective of a story, those same discussants will be much more likely to judge the film and novel to be different stories. This is because while Scout Finch is clearly the protagonist in the novel, her father Atticus Finch is put into that role by the

movie. The perspectival shift results in a marked difference in the moral perspective between the book and the movie. When this difference in moral perspective becomes salient, it will tend to elicit judgments of distinctness rather than identity. Likewise, in many contexts people will be inclined to say that Tolkien's *Lord of the Rings* is the same story as Jackson's *Lord of the Rings*. However, in a context in which the great number of differences in their contents is salient, many people will judge that they are not the same story.

This is important to the question of whether "Maggie's Wild Ride" and "Maggie's Wild Glide" are the same story. Suppose that, before presenting those stories, I had included a discourse on the important mechanical and historical differences between biplanes and hang gliders, only afterwards including the two stories. It's plausible that if then asked whether the stories are the same or distinct, you'd be more inclined to think they're distinct. The reason is that in such a context, you'd have been primed to consider the details about what kind of vehicle Maggie is flying (biplane or hang glider) to be core.

Of course, one might take some of these judgments to be "loose talk" or attribute confusion to those making them. I think there are a few reasons not to do so. First, the context-sensitivity in question is pervasive. These kinds of discrepancies in judgments occur frequently in our non-academic discussions of stories. But they arise in academia no less often. For example, one often finds a theorist endorsing an implicit account of story identity that suits their broader point, as when Alter (1978, 355) claims that "in biblical narrative more or less the same story often seems to be told two or three or more times about different characters, or sometimes even about the same character in different sets of circumstances." The examples Alter gives involve variation in character, setting, and event details, making it quite easy for many stories to count as the same according to the criteria being employed. In the context Alter

has introduced, not that much is required to say of two stories that they're the same story.¹⁵

A second reason to take the context-sensitivity of our talk about story identity seriously is that doing so allows us to resolve apparently deeply held disagreements. One person judges that Francis Ford Coppola's *Bram Stoker's Dracula* (1992) tells the same story as Tod Browning's *Dracula* (1931). Another person judges that Browning's story is better than Coppola's. It doesn't seem that both can be right. However, if the truth of our talk about story identity really is sensitive to context, they can be, so long as the utterances occur in separate contexts.

A final reason to take the context-sensitivity of our talk about story identity seriously is that, as will be seen shortly, doing so turns out to be theoretically fruitful. When we posit that the truth of story identity attributions really are sensitive to context, we discover a way of understanding transposability while retaining the distinction between story type and story identity. I suggest, then, that we take the context-sensitivity of story-identity attributions seriously.

What we need—if the foregoing attempts at outlining desiderata for an account of story identity have hit the mark—is an account that has the following features:

1. It allows for versions of the same story to differ in content and form,
2. It preserves the story identity/story type distinction, and
3. It takes the context-sensitivity of our story identity attributions seriously.

I just gave reasons for (3). (2) was established in the discussion of the HIGHER ORDER TYPE ACCOUNT, which collapses the relevant distinction. (1) is just the transposability of the story, a feature that's been hard to preserve for most any account thus far.

¹⁵Another theorist, asking the question of whether one can tell the same story twice, answers “yes and no. It all depends on what we might possibly mean by ‘story’.” (Polanyi 1981, 333)

In order to satisfy all three desiderata, I suggest we do two things. First, I suggest we adopt a type-theoretic account of stories, on which stories are abstract types that have tokens. Second, I suggest we take the truth of story identity attributions to be relative to context. Call this the SENSITIVE TYPE ACCOUNT. Positing that stories are types will allow us to satisfy desideratum (2), the preservation of the story identity/story type distinction, since story types can be instances of higher-order types (an important lesson of the HIGHER-ORDER TYPE ACCOUNT). It will also allow us to make sense of the earlier observation that stories have both content and a canonical version: theories of types give us a way of understanding what it is for a version of a story to be canonical while still allowing for variation. Positing context-sensitivity of our story identity attributions will allow us to satisfy desiderata (1) and (3), the transposability of the story and the apparent context-sensitivity of our story identity attributions.

On the SENSITIVE TYPE ACCOUNT, stories are individuated the way we commonly understand types to be individuated when those types admit of properly and improperly formed tokens. These are what Wolterstorff (1980, 54-8) and Dodd (2007, 32) call “norm types”, and they are individuated by the condition something must meet in order to be one of its properly formed tokens. According to the SENSITIVE TYPE ACCOUNT, then, the question of whether “Maggie’s Wild Ride” and “Maggie’s Wild Glide” are the same story comes down to the question of what conditions something must meet to be a properly formed token of each. If the conditions are the same in each case, then they’re the same story. If not, then they’re distinct. Plausibly, since properly formed tokens of each differ in what they say about what Maggie is flying, the stories themselves are distinct on this account.

However, it’s also the case according to the SENSITIVE TYPE ACCOUNT that the truth of our story identity attributions is sensitive to context, so that whether a sentence like “Tolkien’s *Lord of the Rings* is the same story as Jackson’s *Lord of the*

Rings” is true is relative to the context of use. How exactly one should understand this context-sensitivity? How can the same sentence attributing identity to some stories be true in one context and false in another? In explaining this, I want to point first to a common way of understanding this phenomenon in contextualist theories of knowledge attributions. Such theories commonly hold that what proposition a sentence attributing knowledge expresses is relative to a context of use. For example, Richard (2004, 215) says, “On standard formulations of contextualism about knowledge, the relation expressed by a use of the verb ‘knows’ — and thus what is said by a use of a sentence such as “George knows that he has hands” — depends upon the “standards of knowledge” supplied by the context of use.” On this way of understanding context-sensitivity, the context of use of a given sentence plays a role in determining what proposition is expressed by that sentence in that context.

Adapting this kind of strategy to the case of story identity, I think we ought to say that the relation expressed by a use of the predicate ‘same story as’ — and thus what is said by a use of a sentence such as “Tolkien’s *Lord of the Rings* is the same story as Jackson’s *Lord of the Rings*” — depends upon the standards of story identity supplied by the context of use. Thus, the truth-value of an attribution of identity to stories can vary according to the context of use, but this is because the content of such attributions can vary according to context.

To clarify the account, I want to say more about standards of story identity. To do so, let me first distinguish the concept of *work* from Chatman’s concept of *kernel*. (Chatman 1978, 53) A work is the total, artistic whole resulting from an artist’s activity of authoring. Works are things like Spenser’s *The Faerie Queene*, Kurosawa’s *Rashomon*, and George Eliot’s *Middlemarch*. The work can be helpfully contrasted with the kernel. Chatman describes a kernel as a major event in the work, a “node” or “hinge” in the structure that forces “movement into one of two (or more) possible paths.” (Chatman 1978, 53) On Chatman’s understanding, kernels

are essential events in the “narrative logic” of a work. Chatman also contrasts this concept with that of *satellite*, a minor plot event that is not at all crucial in the way that a kernel is. On Chatman’s understanding of kernels and satellites, then, they are components of the work’s content, some of which are essential to the narrative logic and others of which are merely ornamental. A work’s content can be construed as a collection of kernels and satellites, all of which stand in something like a part-whole relation to the total content of the work, but such that only kernels are *essential* parts of the work.

I want to appropriate and alter the kernel/satellite dichotomy. First, instead of taking a kernel to be an essential component of a work, take it to be an *important* component of a work. Contextual salience is key to determining what components of a work’s content are important. Thus, if a proposition p is part of the content of some work x , then p is a kernel of x in context c if and only if p is salient in c . For any proposition p , such that p is part of the content of work x , if p is not a kernel of x in context c , then p is a satellite of x in c .

I will also extend the kernel/satellite distinction to cover all of the features of works, not just their content. For example, I will take components of the form of a work to be important or unimportant relative to context, and thus as either kernels or satellites of a work. As an example, suppose a rustic verbiage and a battered old west aesthetic are components of the form of some work. In a context in which we’re comparing that work to the films of John Wayne, those formal features will likely function as kernels of it, given the well-known occurrence of such formal elements in films in which John Wayne appears. Likewise, historical and even obviously extrinsic features of a work, including facts about its reception or how it has influenced (or is the result of particular influences from) other authors can count as kernels in a context in which they’re especially salient.

The general idea is that we can reduce a narrative work to the elemental components of form and content that make it up. In a given context, some of these components are its kernels and some are its satellites. Unlike Chatman, I allow all of the features of a work into this landscape, and I hold that whether something is a kernel or a satellite of the work is a function at least in part of context (rather than merely of “narrative logic”). I allow these alterations because I take it that this is how our actual judgment and discourse about the important components of works functions. Sometimes our attention is on formal components, sometimes on components of the content, sometimes both (and sometimes even on what may seem in other contexts to be merely extrinsic features of the work). When discussing the features of a work in what follows, I will generally refer merely to its kernels or satellites, leaving open whether they are components of the work’s content, form, or something else. I’ll make the distinction explicit in cases where it matters.

Here is what it means, then, to say there are contextually sensitive standards of story identity: a context of use of the term ‘story’ introduces a counterpart relation (a relation of similarity) between works, by which they are ordered from maximally to minimally similar according to their possession (or not) of contextually salient kernels. Two works are counterparts in a context if and only if they share the most kernels in that context. For any works x and y , and for any context c , an assertion that “ x is the same story as y ” is true in c iff x is a counterpart of y in c . The standards of story identity change in the sense that which components of a work are kernels change from context to context. While two works x and y may share all their contextually salient kernels in one context, they may fail to do so in another. Thus, the truth about story identity can change from context to context.¹⁶

¹⁶I also allow that the term ‘story’ is ambiguous, so that while in most cases our claims about stories refer to works, the things that are the types I’ve identified with stories, in some cases they refer instead to kernels of the works, either to some content or to some formal elements of the works. In some contexts it may also be that we refer to a token work rather than to works on the level of

We can find similar kinds of theories in other domains. For example, the SENSITIVE TYPE ACCOUNT gives story identity conditions in terms of the conditions something must meet to be a token of the type that is the story, while simultaneously positing that our talk about story identity is made true by certain types functioning as counterparts. Likewise, the stage theory of personal identity gives personal identity conditions in terms of the conditions something must meet to be the person-stage that is the person, while simultaneously positing that our talk about personal identity is made true by certain person-stages functioning as counterparts. One way of understanding the stage theorist is as relativizing the truth of attributions of personal identity to a context of use, just as the SENSITIVE TYPE ACCOUNT does for attributions of story identity. A context of use introduces a counterpart relation between person-stages, by which they are ordered from most to least similar to a contextually salient person-stage. If y is in the class of person-stages that are most saliently similar to x in context c , then the assertion that “ x is the same person as y ” is true in c .¹⁷ One objection that’s likely to be leveled against the SENSITIVE TYPE ACCOUNT is that it speaks out of both sides of its mouth: it says that stories are individuated strictly—a story is individuated by the conditions that must be met to be a properly formed token of it—but the account also allows the truth about story identity to be sensitive to context. Notice, however, that this is a feature of any view that takes some domain of truths to be a function of objective reality and contextual

types. Context is what settles whether we’re comparing total works or kernels thereof. Likewise, context determines whether names like *Lord of the Rings* or *Middlemarch* refer to the works or to kernels of them. In a context in which reference to stories is to kernels of works, it’s true that two stories are the same if and only if the kernels are counterparts in that context. This gives rise to a second way in which the truth of story identity attributions can vary across contexts: whether we refer to a total work or merely to kernel(s) can vary across contexts.

¹⁷For more on stage theory, see Sider (1996, 2001). Another similar theory is counterpart theory about de re modality. It says the truth of our de re modal discourse is context-sensitive. The counterpart theorist here posits that a context introduces a counterpart relation between objects. When we speak of what an object might, must, or could never have been like, the truth of such statements is relative to context. For more on counterpart theory, see Lewis (1968, 1971, 1973, 1986).

input. Thus, the stage theorist similarly speaks out of both sides of their mouth: they say that people are individuated strictly—as momentary time-slices—but allow the truth about personal identity to be sensitive to context. This may be a cost of the SENSITIVE TYPE ACCOUNT, then, but it’s not a cost it bears alone.

Recall now how the STRICT ACCOUNT didn’t seem able to account for the transposability of the story. The STRICT ACCOUNT individuates stories by their primary content, and different versions of a single story can differ in primary content, so it wasn’t clear how the STRICT ACCOUNT could account for multiple versions of a story each having the content of the story while also differing from each other in content. The SENSITIVE TYPE ACCOUNT can succeed where the STRICT ACCOUNT failed. It tells us that the truth of our story identity attributions is determined contextually by the stories’ salient components, their kernels. On the SENSITIVE TYPE ACCOUNT, the claim that two stories x and y are versions of the same story can be understood as true in a context in which x and y are counterparts. There are plenty of contexts in which it’s true to say of two works that differ in form and content that they’re the same story (or versions of the same story) because they share most or all of their kernels in that context. Likewise, the SENSITIVE TYPE ACCOUNT helps explain how there could be other contexts where it would be right to say of those very same works that they aren’t the same story: a change in context can introduce a change in which components of the works are salient.

The SENSITIVE TYPE ACCOUNT also reconciles the contextual sensitivity in our talk about story identity with a way of retaining the identity/type distinction. We’re sometimes inclined to make distinctions in story identity on account of very slight differences (e.g. recall the case of Jackson’s vs. Tolkien’s *Lord of the Rings* presented above). On the other hand, sometimes people make the claim that there are only so many stories in all of human history (I’ve seen 6, 7, and 10 suggested as the right numbers). The latter kind of claim often seems to be a claim about story type

(and indeed people often speak of types of stories or of plots rather than of stories simpliciter), but on the SENSITIVE TYPE ACCOUNT we can make sense of the claim that there are only seven stories as itself introducing an extremely course-grained context, one that effectively collapses the difference between story identity and story type. This explains the ease with which critics can switch between talk of there only being a few stories and talk of there only being a few types of stories, since in a very course-grained context the truth-conditions for these claims will be the same. The story identity/story type distinction is a contextually relative distinction. In most contexts it applies, but there are contexts where it doesn't.

Another advantage of the SENSITIVE TYPE ACCOUNT is that it's able to account for the way that stories seem temporally and modally flexible. It seems like stories can differ in their content over time and over possible scenarios without change in identity. For example, George Lucas's *A New Hope* (1977) arguably did not represent the proposition that Darth Vader is Luke's father at the time of its release. Arguably, now it does. It seems the content of *A New Hope* has changed over time.¹⁸ Any account that merely individuates works strictly according to their form or content will have trouble making sense of these kinds of changes. Even the HISTORICAL ACCOUNT has a hard time saying just which changes a work can or cannot undergo without fission or fusion occurring (or just how a work could have been different without being a different work). Applying the SENSITIVE TYPE ACCOUNT diachronically allows us to make sense of these kinds of cases. The proponent can point out that what we'll say about the original *A New Hope* versus the *A New Hope* of (1980) and after depends in part on the context. Most contexts are lax enough that small changes in the form or content of a work won't matter for our talk about what story it is, or whether it's the same story as another work. Only in a very stringent context would we say that,

¹⁸This example is due to Caplan (2014), who provides an extensive discussion of change in the content of serial works.

e.g. *A New Hope* became a new story when Lucas made it the case that Vader was Luke's father.

The SENSITIVE TYPE ACCOUNT will of course face objections. One potential objection is that it seems to divorce numerical identity from story identity. For example, as I've pointed out previously, in some contexts the statement "Jackson's *Lord of the Rings* is the same story as Tolkein's *Lord of the Rings*" will be true while in others it will be false. But for this to be true the following must hold as well. Suppose there to be some objects x and y , such that x and y are stories and x and y are not numerically identical. A consequence of the SENSITIVE TYPE ACCOUNT is that in some contexts the assertion " x and y are the same story" is true.¹⁹ This is because in some contexts every contextually salient kernel $k_1 \dots k_n$ of x will be a kernel of y .

This will strike some as intuitively bizarre. A natural way to think about the claim that some objects are the same story, after all, is that they are each a story and they are numerically identical. What's going on in cases of story identity is not that there's a unique relation, *same story as*, distinct from the relation of identity, but just that there are stories and some of them are identical. The SENSITIVE TYPE ACCOUNT seems to entail that this natural way of thinking about story identity is false. The challenge is for its proponent to explain how this could be.

¹⁹Suppose instead that there are some objects x and y , such that x and y are stories and x and y are numerically identical. Is a consequence of the SENSITIVE TYPE ACCOUNT that in some contexts the assertion that " x is not the same story as y " is true? It depends on whether the standards of story identity can be applied irreflexively. One reason we might think they could be is if we think of those standards as being applied by means of something like a Fregean sense, so that even though the last movie I watched in the theater is numerically identical to Abrams' *The Rise of Skywalker*, I might still sensibly (though falsely) claim something like, "*The Rise of Skywalker* wasn't as good as the last movie I watched in the theater." If the standards for story identity are applied similarly, those standards would allow for cases where an object that is a story can truly be said not to be the same story as itself. I don't think we should allow the standards of story identity to be applied this way, however. Doing so would allow the question of whether a story even contains a given kernel to be relative to context, and that doesn't seem at all right. Whether a given kernel counts as important may be contextually relative, but surely not whether or not its a component of the story.

My explanation of the way numerical identity and story identity come apart in the SENSITIVE TYPE ACCOUNT is that *story* is a phase sortal. A phase sortal is a particular kind of property, viz. one that's typically exemplified by an object only for a portion of the time that object exists. Phase sortals are typically expressed by sortal predicates such as "is a child" or "is a teacher". Identifying something by one of its phase sortals is a way of saying what it is, but, we might say, not of saying what it *really* is. We can give existence and identity conditions for teachers, for example, but to do so would be to say under what conditions there is an object that exemplifies the property of *being a teacher*, distinct from other exemplifications of that property. It wouldn't be to give existence or identity conditions of the objects themselves that exemplify the property of *being a teacher*, since those objects can go on existing even after they stop being a teacher.

Markosian (2010) points to a way of formulating questions of identity in cases where numerical identity and sortal identity come apart. Discussing the case of personal identity, Markosian suggests we formulate the problem of personal identity in terms of instances, instantiations, and episodes of personhood. An instance of a property ϕ is the having of ϕ by an object at a time. Instances are complex, momentary entities. An instantiation of a property ϕ is the having of ϕ by an object, whether at a single moment or in a temporally extended fashion, where temporally extended instantiations are composed of multiple instances. An episode of ϕ -ness, for some property ϕ , is a maximal instantiation of ϕ . Maximality is understood such that "instantiation E of property ϕ is maximal iff the fusion of E with any further instance of ϕ -ness (i.e., one that is not a part of E) is not a single instantiation of ϕ -ness". (Markosian 2010, 137) Markosian's suggestion for how to formulate the problem of personal identity is as follows:

The problem of personal identity consists of trying to provide an answer to the following question: What are the conditions under which an instance of personhood at t_1 is part of the same episode of personhood as an instance of personhood at t_2 ? (Markosian 2010, 138)

What Markosian's formulation provides (among other things) is a way of formulating the problem of personal identity that allows for an answer on which *person* is a phase sortal and on which personal identity can come apart from numerical identity.

Markosian is not the only one to allow personal identity to come apart from numerical identity. Consider the theory of personal identity defended by Sider (2018), on which the relation of personal identity is asymmetric. On Sider's view, personal identity is grounded by a variety of psychological identification relations, each of which is asymmetric. On this view, it's possible for one person-stage, p_1 , to be the same person as another person-stage, p_2 , even though p_2 is not the same person as p_1 , nor are they numerically identical person stages. On this view, personal identity is certainly not reducible to numerical identity.

This way of framing the problem of personal identity points towards a similar way of framing the problem of story identity on which story identity and numerical identity can come apart. If we can entertain the relation of personal identity being something separate from the relation of numerical identity as applied to persons, it should be possible to do the same with the relation of story identity.

One might doubt that the analogy with personal identity is apt, however, since the kinds of questions Markosian and Sider's views are meant to answer generally are questions of the diachronic identity of persons—i.e. identity of persons across time. My claim is that two works can either be or fail to be the same story in different contexts; but contexts may be located at the same time, and so the solution to the question of what makes two works the same story must be given an answer that can apply synchronically—i.e. at a single time. This is no real problem, though. Just as Markosian defines the problem of personal identity in terms of instances of personhood and episodes of personhood, we can define the problem of story identity in terms of instances and episodes of storyhood. Following Markosian, take an instance of storyhood to be a work that has the property of storyhood at a time. Likewise,

allow an instantiation of storyhood to be a work having the property of storyhood either at a moment or in a temporally extended fashion. Finally, take an episode of storyhood to be a maximal instantiation of storyhood by a work. (Maximality can be understood, as before, so that an instantiation E of storyhood by a work is maximal iff the fusion of E with any further instance of storyhood is not a single instantiation of storyhood.) The question of story identity is the question of under what conditions an instance of storyhood at t_1 is part of the same episode of storyhood as an instance of storyhood at t_2 . These times may be identical or distinct. The SENSITIVE TYPE ACCOUNT answers that this is the case, relative to a context, whenever the works are counterparts in the relevant context.

One might grant that the context sensitivity of the the SENSITIVE TYPE ACCOUNT is important but level the objection that it's not obvious why we should combine that feature with the claim that stories are types. Why not combine context-sensitivity with the STRICT ACCOUNT, for example? This objection is not hard to answer. Take a moment to consider the STRICT ACCOUNT as it would be applied to token stories, on the one hand. It would tell us that token stories are individuated by what they represent explicitly. Even when combined with context-sensitive account of story identity attributions, this seems false. Two token stories can be instances of the same story even though they fail to represent the same things. Say, for example, someone drops a line in a performance of *Hamilton*. It's still a performance of *Hamilton*, not some other story.

Token stories seem to be individuated by the types they're tokens of, then. If, on the other hand, we apply the STRICT ACCOUNT to story types, then it already agrees with the SENSITIVE TYPE ACCOUNT about what stories are (i.e. types) but disagrees about how they're individuated. The SENSITIVE TYPE ACCOUNT says they're individuated by the conditions that must be met by something to be a properly formed token of the story. The STRICT ACCOUNT would say they're individuated by

what they represent. But the SENSITIVE TYPE ACCOUNT seems to give the clear right answer as to how types are individuated. Something doesn't get to be a token of *Wonder Woman* (2017) merely by explicitly representing everything the movie does. Someone could write the list of propositions that exhausts what the movie explicitly represents, after all, but they would have failed to reproduce the movie in doing so. In order to be an instance of *Wonder Woman*, something has to have more than mere representational adequacy.

We might also consider combining the context-sensitivity of the SENSITIVE TYPE ACCOUNT with the individuation conditions of the HIGHER-ORDER TYPE ACCOUNT. The problem here is that this would be redundant. Positing that stories are higher-order types was meant to explain the transposability of the story, but that's explained well enough by taking context-sensitivity seriously. We don't need to say that every story is a higher-order type.

The HIGHER-ORDER TYPE ACCOUNT did highlight an important distinction for us: we sometimes talk about two stories being the same (or about two works telling the same story) while, other times we talk about one story being a version of another. We might wonder if there's an important difference in the relationships we're tracking with that talk.

I think there is a subtle difference in the kind of relationship we tend to track with the different terms. When we talk of versions, it tends to be in cases where what's salient to us is the historical relationship between works. For example, when we know that two stories are similar but (saliently) they're also the result of the same project, then we seem more inclined to say they're versions of a single story. For example, it's common knowledge that there are many versions of the film *Blade Runner* (1982). On the SENSITIVE TYPE ACCOUNT, each "version" is a distinct story, a type in and of itself. However, in almost any context we're inclined to call these stories versions

of the same one story because what remains salient across most contexts is their historical relationship and the features they share.

I think there is on balance compelling reason in favor of adopting the SENSITIVE TYPE ACCOUNT of story identity, then. It makes sense of our talk about stories, the transposability of the story, and story identity/type distinction without falling prey to the problems of the HISTORICAL ACCOUNT. Now, with a firmer grasp on the existence and identity conditions for stories, we should be able to say what it is authors are doing when they make up stories.

1.4 Making Up Stories

Existence and identity conditions for stories help make clearer what must be going on in the activity of authoring (i.e. of making up a story). Authoring is the activity by which we are introduced to new stories. The existence conditions for stories tell us that for there to be a story we're introduced to, there needs to be a way the story goes and some canonical version of the story. The identity conditions for stories tell us stories are individuated by the conditions required for being a properly formed token of them, but also that claims about whether a given story's content is novel are in large part relative to context. I'll elaborate on these points here and use them to guide my account of authoring.

Authoring must obviously involve, minimally, the identification of new stories. For an author to introduce us to a new story, there must be a story for them to introduce us to, distinct from any other. Given the existence conditions for stories we've settled on, it seems that authors must identify both a way for their story to go and a canonical version of the story. An author introduces both novel narrative content and a novel way of telling or presenting it.

Given the identity conditions for stories, so I've argued, an author introduces us to a new story when the story they introduce us to is unique in the conditions required

for properly tokening it. The way a story is tokened, importantly, is the way the story is told. Thus, an author introduces us to a unique story when the way their story is properly told is unique.

Of course, the way a story is told, as well as the content that's related in a telling, can be more or less novel. As I've argued above, our judgments about story identity are sensitive to context. So, whether we judge two stories to be distinct (and thus whether we judge a particular story to be made up) will be likewise sensitive to context.

Let's put these observations to work in an initial, first pass account of authoring.

FIRST PASS ACCOUNT: For any subject S , S makes up a story iff

1. there is a narrative representation N , and narrative content c , such that N has c ;
2. S identifies N and c ;
3. c does not non-accidentally resemble some other narrative content c' of a different story; and
4. c does not non-accidentally represent actual events.

Take condition (1) first. By a 'narrative representation' I have in mind the type of representation—visual, linguistic, or in some way mixed—that is used to portray the content of a story. A narrative representation is an abstract type, naturally taken to be a way of telling a story. If the canonical version of a story is a novel, then the narrative representation in question is a linguistic type. If the canonical version is a film, then the narrative representation is a cinematic type. When an author settles on a canonical form of representation for the content of their story, they are (as I'm conceiving of it) identifying a narrative representation for their story and making it the case that that representation has that content.²⁰ By doing so, an author meets the existence conditions for stories.

²⁰Of course, the author doesn't establish by fiat what representations have what content. There are also facts about linguistic and pictorial conventions and other contributions to meaning from an author's community.

That brings us to condition (2), the requirement that an author identify both the narrative content and representation of their story. Different authors do this differently. Some focus more on the narrative representation, concerning themselves more with the form their work is taking and less with the content of the story. Others concern themselves primarily with the content and attempt to create a work with a form to fit that. This cognitive diversity of authors is not a threat to the claim that someone must identify both the form and content of a story in order to count as authoring it. By identification I have in mind a process of intentional mental selection, which may be mediated (so that one selects the content of a story, e.g. by selecting its canonical form, or vice versa). By satisfying condition (2), an author makes a story their own.

Condition (3) is an anti-plagiarism constraint. Someone doesn't count as authoring a story if the content of their story comes from another story, whether this is intentional or otherwise. This is obvious to us. When we find out that someone purports to have authored a story but they really got all the ideas for its content from another story they've read or heard somewhere, we reject their claim to authorship.

Condition (4) is an anti-journalistic constraint. An author doesn't count as making up a story if the content of their story is a report of actual events. This constraint is also somewhat obvious to us. If someone purports that their story is fiction, at least, we expect them to say something if the story was in some important sense based on real events. Likewise, if someone claims to have made up a story as told in their new novel and we find out the story is in fact autobiographical, we're liable to contest the claim that it was made up (or, on the other hand, the work's classification as fiction).

It's important to note that someone can author a story that unintentionally non-accidentally represents actual events. Matravers (2014, 22-3) and Currie (1990, 45) each consider a particular case of this sort. Currie spells the case out in detail.

Smith has real-life experiences of so horrible a kind that he represses them. He then invents, so he supposes, a story, and this story exactly retells these

events. This is no coincidence; Smith's subconscious somehow provides him with the information for his story. (Currie 1990, 45)

The content of Smith's story is (unbeknownst to him) an accurate representation of the actual events of his life that occurred many years prior. The fact that the content is a representation of actual events is no accident. Had the events not occurred, Smith's story would not have had the content it does. Because they occurred, Smith's story does. The lesson to be learned from such cases is that people can make up (or fail to make up) stories on accident. This is to say that (perhaps surprisingly) intention plays no essential role in determining whether a story is made up. In normal cases, authors will succeed at doing what they mean to do. But authors, being human being, are susceptible to mistakes.

A related, important question is what the distinction between accidental and non-accidental representation at play in conditions (3) and (4) amounts to. There are a couple of different ways we can disambiguate the claim that some content represents some events. First, it might be that the content in question is *about* the events, no matter whether what it represents of them is accurate. Second, it might be that the content in question *accurately* represents the events. We'll often only be inclined to say some content accurately represents some events if it's about those events. However, if we're going to make sense of the notion of accidental representation, it seems we need to allow for cases in which someone identifies some content that, by chance, accurately represents some events that it isn't actually about. I'm not going to settle here what it takes for some content to be about some events. Instead, I'll focus on the accidental/non-accidental distinction as it applies to accuracy.

One natural thought is that non-accidental representation of actual events obtains when there's counterfactual dependency between the actual events occurring just as they do and the story's having the content it does. On this way of thinking about non-accidentality, a story S non-accidentally represents some actual events $e_1 \dots e_n$

iff the following is true: were $e_1 \dots e_n$ not to have occurred as they did, S would not have the content it does. In honor of Nozick (1981), call this the ‘Sensitivity Account’ of non-accidentality. If the content of a story is Sensitive to the actual events it represents, then the story is not made up.

While the First Pass Account should be clear now, there’s an obvious problem with it that will lead us to a final account. Multiple components of the the First Pass Account are only met by degree. For example, resemblance is a degreed matter, so that all content resembles all other content to some degree. Condition (3), which prohibits non-accidental resemblance of the content of a story to the content of other stories is clearly too strong. What’s important is that resemblances can’t be too great, or that if they are it must be by chance.

Likewise, it seems clearly to be a matter of degree the extent to which a given story non-accidentally represents actual events. Every story is at least inspired by many actual events insofar as an author must draw from their actual experiences to inform their choices about what events will occur in their stories. Many stories we would say are made up are more directly inspired by actual events, however. Historical fictions often include many actual historical events as a backdrop for a fabricated plot, for example. Often, the line between what’s made up and what’s taken from reality is even more blurry. For example *The Book Itch*, by Vaunda Michaeux Nelson depicts a real person—Lewis Michaux, owner of the National Memorial African Bookstore in Harlem. The story is told from the perspective of Lewis’s son, Lewis Jr. It is based on actual conversations between Lewis Jr. and the author, but Michaeux Nelson thought a more captivating story would be one told by Lewis Jr (rather than merely narrating his responses to questions). Speaking about this decision, Michaeux Nelson says, “I knew this would mean expanding on his words and crossing the line between nonfiction and fiction. So be it. What matters most is what works best for the storytelling. However, I think it’s important to inform readers about what’s been

done with the nonfiction content, to let them know what they are being served.”²¹ I take it what’s going on here is that Micheaux Nelson wanted to partially represent some actual events. In order to do so, she wrote about some non-actual events, representing them as having some similar features to the actual events. Micheaux Nelson’s conviction was to tell a story that wasn’t about any actual events but that was importantly based on and resembled some actual events. It’s not totally clear to me (or to the author, apparently) whether this story is better classified as having been made up or not. It’s obviously not the archetype of an original story.

In light of these considerations, I take it that when someone counts as making up a story is a matter of degree. A story isn’t just totally, wholly made up or totally, wholly not. Some stories depict mostly actual events, but along with a few non-actual events. Some impart important truths about reality (including about real events) by depicting non-actual events. There’s a lot of room for variety.

I submit the following, then, as my considered account of what it is to make up a story:

FINAL ACCOUNT: For any subject S , and any story x , such that S identifies the narrative representation and content (N, c) of x , S makes up x to the degree that

1. c does not non-accidentally resemble some other narrative content c' of a different story; and
2. c does not non-accidentally represent actual events.

This account fits best with the existence and identity conditions for stories settled on above. By making authorship a degreed endeavor, it also fits well with our understanding of narrative creativity as sitting on a spectrum.

This account of authorship also accords well with the SENSITIVE TYPE ACCOUNT of the identity conditions of stories, which says that the truth of our story identity

²¹From Hinz (2018); <https://lernerbooks.blog/2018/12/blending-nonfiction-and-fiction-4-recent-picture-books.html>, visited on 04/30/2019.

attributions is relative to context. We can posit that the same contextually salient similarity relations that determine the truth of our story identity attributions are what determine the standards for how much non-accidental resemblance is enough to disqualify a story as original. Clearly no story is perfectly novel in the sense given by conditions (1) and (2) of the FINAL ACCOUNT. It also seems that our judgments of originality are contextually-sensitive in the same kind of way as our judgements of story identity. The proponent of the SENSITIVE TYPE ACCOUNT of story identity can explain this.

I've argued that we ought to understand the existence of stories in terms of both form and content. For there to be a story, there has to be a way it goes and a way of telling it. The identity of stories should be understood in terms of what it takes to properly token them, while the truth of our assertions about the identity of stories is a contextually sensitive question. The account of authorship given is one on which authors make narrative content accessible to us, and they count as authors to the degree that the content of their works is original (i.e. both un plagiarized and not based on actual events).

CHAPTER 2

COMPLETING STORIES

2.1 Introduction

In the previous chapter, I set out conditions under which an author counts as making up a story, but when is this activity finished? When is a story complete? The more general question of artwork completeness has been a topic of recent interest to philosophers and other theorists.¹ They've already rehearsed the importance of the question. Some sense of artwork completeness is important for making sense of legislation about the moral rights of artists.² A sense of artwork completeness also plays an important role in the practice of displaying artworks in galleries and other communal spaces (e.g. concert halls, homes, gardens, or corporate environments). In many cases, if a work is incomplete in some relevant sense, it's a violation of artistic norms (if not moral or legal norms) to display it. Moreover, as Trogdon and Livingston (2014) point out, many artists seem driven by the goal of completing

¹For a volume containing some recent sociological discussion, see Becker et al. (2006). For the philosophical literature, see Beardsley (1982), Cray (2018), Gover (2011, 2015, 2018), Hick (2008a,b), Livingston (1999, 2005, 2008), Livingston and Archer (2010), Rohrbaugh (2017, 2018), and Trogdon and Livingston (2014, 2015).

²Trogdon and Livingston (2014, 225) claim the "legal category of derivative works depends on the category of previously completed works, those on which derivative works are based." Gover (2018, 45-46) corrects this though: "US copyright law says that the work must be fixed in a tangible medium, but it does not say that the work has to be complete in any absolute sense." See Gover (2011, 2018) for discussion of legal cases where artists cite lack of completeness of a work as sufficient reason for suit against a museum for displaying the work. While Gover is clear to point out that the legal and moral rights of the artist in these cases don't depend on the completeness of the work, the completion status of the work is cited by the artists and others as morally and legally relevant. Moreover, it's not implausible for a Completion Pluralist (a view to be introduced below) to understand the notion of fixation as itself a kind of completeness. Since nothing in my arguments turn on this, I won't emphasize these points in what follows.

their works. We should want to know just what it is artists are trying to achieve. Trogdon and Livingston (2014) and Rohrbaugh (2017) also discuss the fact that for some sense of completeness, if an artwork is incomplete then this restricts the manner in which its acceptable to evaluate or appreciate it.³ Since philosophers of art are already motivated to understand artistic evaluation and appreciation, we should want to understand what sense of completeness is relevant to those activities. Finally, philosophers of art are also motivated to understand what works of art are, what ontological categories particular works fall under and what important properties they have. But it often seems one cannot answer these questions for a particular artwork unless one knows whether it's complete, for some sense of completeness.

There are thus concepts of artwork completeness that we employ in our reasoning about the moral and legal rights of artists; the practices of displaying, appreciating, and evaluating art; the artistic process; and the work itself. An important question is whether the concept of artwork completeness is the same in each case. Philosophers writing on artwork completeness have been ambivalent about this.

On the one hand, they have sometimes recognized that there is more than one way for an artwork to be complete. Consider the opening page of Livingston's inaugurating piece on the topic:

To say that a work of art is finished can mean at least two very different things, depending on whether we are focusing primarily on some item's artistic or aesthetic features, or whether attention is drawn to an aspect of the creative process. (Livingston 1999, 14)

If we're focusing on an object's artistic or aesthetic features, we may (like Beardsley 1982, 55) come to the conclusion that for an artwork to be finished is for it to be "worthy of standing by itself, as an object of aesthetic enjoyment." This notion of completeness, understood as a normative condition on a work's aesthetic value,

³Of course, we can evaluate and appreciate unfinished works. The point is that there is a sense in which we must hold something back in these activities until the artist finishes their work and puts it out for the full scrutiny of the public. Thanks to an anonymous reviewer for pointing this out.

Livingston calls “aesthetic completion”. He calls the latter kind, the kind having to do less with the work’s aesthetic features and more with the creative process, “genetic completion”. Livingston posits that both kinds can be present in a work, that neither might be, and that they can come apart, so that a work might be complete in one way but not another.

While much of the subsequent literature has paid obligatory homage to the aesthetic/genetic distinction, the focus and rhetoric of that literature has continually undercut the impact of this observation. The decision of Livingston and others to refer to this second kind of completeness as “genetic completion” (or as Gover (2015, 457) calls it “ontic completion”) paints the picture that this is the kind of completeness that matters for determining whether a work is *really complete*, or for what the work *really is*. This picture is confirmed by the complete lack of interest shown in aesthetic completion by those writing on the general topic of artwork completeness. Recently, mention of the aesthetic/genetic distinction has altogether dropped out of the literature on the topic.⁴

In this chapter, I argue that the tendency of authors to either situate genetic completion as the “true topic” or to leave aesthetic completion out of the discussion altogether is a mistake. This is because the following thesis is true:

COMPLETION PLURALISM: There are many kinds of artwork completeness, many corresponding senses of ‘complete’, and no kind of artwork completeness is objectively more important than any other.

Completion Pluralism is a meta-thesis. It has the consequence that the question “When is an artwork complete?” is ambiguous. The conditions for completion will vary depending on the kind of completion under consideration. The upshot is that instead of arguing about when a work is really complete, we should list the kinds

⁴The later Livingston (in Trogon and Livingston (2014)) leaves the distinction out. Rohrbaugh (2017, 132) says that setting aside aesthetic completion “allows us to focus on our true topic.” Gover (2018) discusses artwork completion at length, treating it as a single monadic (if vague) property of artworks and never mentioning aesthetic completion.

of completion and analyze each kind in turn. Call the denial of Completion Pluralism “Completion Monism”. Completion Monism can be understood in at least two different ways.

NUMERICAL COMPLETION MONISM: There is only one way for an artwork to be complete.

PRIORITY COMPLETION MONISM: There is more than one way for an artwork to be complete, but one kind of artwork completeness is objectively more important than the others.

A natural way to read the literature on artwork completeness so far is as implicitly assuming Priority Completion Monism. They acknowledge other ways for an artwork to be complete besides whatever way they are trying to identify, but they take there to be some kind of completion that is most important for what an artwork is, and they’re engaged in the project of trying to analyze that way of completing a work. Regardless, it won’t matter exactly what version of Completion Monism others have held, as an argument for Completion Pluralism will count equally against both versions of the view.

In what follows, I provide two arguments in favor of Completion Pluralism. The first argument is negative: even if Completion Monism were true, the extant views of completeness wouldn’t be good candidates for the one (or most important) kind of completeness. These views fail to capture any satisfactory monistic conception of completeness, and the best explanation of this fact is that Completion Pluralism is true. The second argument is positive. We have disparate reasons on different occasions for caring about whether an artwork is complete (in some sense of complete). Our concerns are moral, legal, social, theoretic, artistic, aesthetic, and ontological. These concerns can come apart, however, and so can the various properties of artworks they track. A work can be intuitively complete according to a sense of ‘complete’ that tracks one property but not another, and in these cases there’s no reason to say that one sense of ‘complete’ is tracking an objectively more important property. The only good explanation for this is that Completion Pluralism is true. After providing

these two arguments for Completion Pluralism, I apply the thesis by identifying four distinct kinds of artwork completeness, showing how each addresses a distinct concern we sometimes have when wondering if a work is complete, and arguing that each can be present in the absence of the others. I go on to identify a distinct kind of completeness I think applies only to stories: narrative completion.

2.2 Against Completion Monism

Here I give a negative argument for Completion Pluralism: it is the best explanation for the failure of recent theories of artwork completeness. Extant philosophical theories of artwork completeness fall into two basic camps, each of which places a psychological constraint on the completion of a work. The first, I'll follow Rohrbaugh (2017) in calling "Psychologism." This is the view that whether an artwork is complete is, in some way to be specified, merely a matter of the artist's psychology. Livingston has defended various versions of Psychologism (along with co-authors Archer and Trogdon). Gover (2015, 2018) also arguably endorses Psychologism about completeness. She distances herself from Livingston by emphasizing the variation in the means by which artworks come to be completed, the complexity of the phenomenon in question, and the indeterminacy of it. But she can be found speaking psychologically in characterizing what she's after. For example, in her considered theory of authorship, Gover (2018, 82) says that we must posit an intentional moment in which the author indexes the work as complete. This moment marks the completion of the first intention necessary to authorship, the intention to make a work of art. The main points of debate for proponents of Psychologism involve whether completeness is always determinate; whether it's terminal (i.e. whether a work must remain complete once it has become so); and whether it must involve a decision of the artist, a disposition of the artist, or whether the determination of completeness can admit of different kinds of psychological states.

The virtue of Psychologism is the way in which it captures the intuitive sense in which a work can be complete (or not) that's up to the artist(s) who authors the work. Consider the case of the *non-finito*, a genre the name of which tags each token work as incomplete or unfinished. *Non-finiti* are cases of artworks that are intentionally designed to appear as incomplete parts of other wholes, such as other artworks or bodies. A typical *non-finito* might be a sculpture of an arm alone, or a sculpture of a person that appears to have been abandoned.⁵ A novel might be considered a *non-finito* if it ends abruptly, in the middle of the action, with no resolution to the main narrative. We can distinguish a completed instance of this genre from an intrinsic duplicate for which the work of the sculptor or author was merely interrupted. There seem to be complete *non-finiti*: complete incompletes. The hypothesis of Psychologism is that the way in which these works are complete is that the author is done with them. The debates between proponents of the view center on the best way to make this precise.

The other extant camp has a sole (published) occupant. Rohrbaugh (2017) holds what I call "Satisfactionism," the view that an artwork is complete iff it is intrinsically such as to have satisfied the artist's creative plan, their intentions in making the work. It thus places a psychological constraint on the completion of the work (viz. the plan of the artist) but stops short of analyzing work completeness in merely terms of artist psychology, as Psychologism does. Satisfactionism offers two promises: it can explain the same phenomena as Psychologism, and it can do so without running into the same problems. It's worth taking a moment to consider whether Satisfactionism delivers on these promises.

Since part of Satisfactionism's promise is that it can avoid the problems suffered by Psychologism, I'll examine the problems of Psychologism first. Rohrbaugh (2017)

⁵A number of historical cases of *non-finiti* are discussed in Becker et al. (2006).

raises a significant objection for Psychologism. I'll follow Cray (2018) in calling it "Rohrbaugh's Regress". The threat of Rohrbaugh's Regress is that Psychologism gets wrong how artists reason about completion. To see this, first notice that the thesis of Psychologism is generally formed either in terms of judgments or dispositions.⁶ Take the judgment version first: this says (roughly) that a work is complete iff the artist(s) of the work make an uncoerced judgment that the work is complete (Cf. Livingston and Archer (2010, 446)). Now, consider the artist who wants to know whether their work is complete. Judgment Psychologism tells them that whether the work is complete is a question of whether they've judged it to be complete. So, what they ought to do is introspect and ask, "Have I judged the work to be complete?" But the notion of completeness appears here again, and so what they're really asking is whether they've judged that they've judged the work to be complete. And we're off to the races.

Of course, the proponent of Psychologism might opt for the disposition version. This says (roughly) that a work is complete iff the artist(s) of the work have acquired a completion disposition with respect to their work (Cf. Trogdon and Livingston (2014)). In this case, the artist isn't so much forming a judgment concerning the question, "Is this complete?" Instead, she's asking herself, "Should I stop now?" The problem with Disposition Psychologism is that it has to be able to distinguish between cases in which artists decide to stop because they've abandoned their projects and those in which they really do complete them. To make this distinction, Trogdon and Livingston (2014, 228) say that "when an artist has a completion disposition [as opposed to a mere refrainment disposition], she is disposed to refrain from making further changes to the work in virtue of having exercised her capacity to reason about art and its production in a certain way." As has been identified by both Gover (2015,

⁶Cray (2018) also distinguishes between cognitive and non-cognitive versions of Judgment Psychologism. I don't think this complication changes anything that follows, so I omit it.

458) and Rohrbaugh (2017, 135), talk of “exercising a capacity to reason about art” is just a way of implicitly smuggling in a completion decision. As Gover puts it, “this qualification is just another way of saying that the artist has effectively decided that the work is done.” (Gover 2015, 458) In that case, Disposition Psychologism faces Rohrbaugh’s Regress all over again.⁷

It’s important to note when we’re evaluating Rohrbaugh’s Regress that it only succeeds on the assumption that Completion Monism is true. That is, the reason Rohrbaugh’s Regress is effective is because both Rohrbaugh and his opponents assume Psychologism is supposed to be giving an analysis of the notion of real completeness, whatever that turns out to be. What Rohrbaugh’s Regress shows is that if there were only one (or one most important) kind of artwork completeness, Psychologism wouldn’t be a good candidate for that one kind. The Completion Pluralist can understand Psychologism differently, though. A work is Psychologically complete, they might say, just in case the artist of the work makes an uncoerced judgment that the work is complete in some non-Psychologic way, for example in the way described by Satisfactionism or in a manner relative to the artistic kind of the work (a proposal I’ll discuss further below). Rohrbaugh’s Regress would never get off the ground.

Another reason Psychologism can’t be the only (or most important) kind of completeness is that it can’t account for certain important cases of incompleteness. One historical example is Schubert’s Symphony No. 8 in B minor, often referred to merely as the *Unfinished*. Schubert stopped working on the *Unfinished*, and it seems undeniable that he formed a refrainment disposition with respect to the work. For all we know, this refrainment disposition was based in the exercise of Schubert’s rational

⁷It’s worth pointing out here that regardless of whether this argument targeting the disposition version of Psychologism succeeds, the point that inspires it is surely correct. Namely, if an artist wants to know whether their work is done, in at least some cases they must do more than merely introspect.

capacities, so that it also counts as a completion disposition. However, we count the *Unfinished* as an archetypical case of an incomplete work. The reason we are so certain it's incomplete is not because we feel sure about the question of whether Schubert's disposition towards the piece is a completion disposition or a mere refrainment disposition (or whether he decided it was done). The reason is that we know the *Unfinished* is a symphony, and something must have a certain structure to be a complete symphony. The *Unfinished* thus satisfies the definition of Psychologism for all we know, but it's an incomplete symphony. The kind of completeness we're sure it lacks is not Psychologistic.

Let's return to Satisfactionism. It promised to explain what Psychologism does without suffering the same problems. Recall that Psychologism is meant to explain the intuitive sense in which a work can be complete (or not) that is up to the artist(s) who authors the work. Satisfactionism makes sense of this not by allowing artists to decide when a work is done, but by allowing them to set the parameters for completeness according to their plan.

Does Satisfactionism really explain what Psychologism does? Here's one reason to think not: it seems artists can decide to be done with works that don't satisfy their creative intentions. As Trogdon and Livingston (2014, 228) point out, some of the considerations salient to an artist when deciding whether to continue working on an artwork include how much time has passed, how much money has been spent, and what genre one is working in. Whether one is bored of a project, how meaningful it makes one's life, whether there's a due date, whether there's a paycheck to be had, whether there are other more exciting projects on the horizon, and whether an artist is working with people who expect certain results also seem like relevant factors for an artist to consider when deciding whether to stop working on something. Given all these factors, it should be quite a usual occurrence for an artist to stop working on an artwork even though it doesn't perfectly satisfy their plan.

On the whole, I would hazard the guess that cases where an artwork completely satisfies the plan of the artist are in the minority. The threat to Satisfactionism from this point is double. First, it seems Satisfactionism and Psychologism will disagree about what works are complete, so Satisfactionism won't actually explain the same intuitive cases of completion that Psychologism does, and Satisfactionism doesn't live up to its promise of successfully accounting for all the intuitive cases of completion that Psychologism does. Second, cases where an artist decides a work is done even though it doesn't satisfy their plan look like counterexamples to Satisfactionism's claim of capturing the one (or the most important) kind of completion. In such cases, the work is intuitively complete, but Satisfactionism seems to have the result that it's incomplete.

The Satisfactionist might claim that what's actually going on when an artist appears to "settle" by publishing a work that doesn't satisfy their plan is that the artist is really altering their plan to fit the present state of the work. In effect, the creative process is almost always one where intentions and work both change until they fit each other. In cases where it appears an artist has completed a work that didn't satisfy their plans, they must have actually altered their plan. Surely, it's true that creative intentions are dynamic and subject to change throughout the creative process. However, if this response is correct then it's in principle impossible to distinguish between the phenomenology of the following two cases:

1. Esperanza plans to write an epic, young adult novel about a little girl with magical powers. While still in the process of writing, she decides to publish a short fragment of the longer story she wants to write due to pressure from her publisher to finish.
2. Mirabel plans to write an epic, young adult novel about a little girl with magical powers. In the process of writing she realizes that a short story would better fit the mood and some of the initial content of her story than the longer novel she

had originally planned. She publishes a fragment of the story she had originally planned to write as a way of satisfying her altered plan.

Mirabel will experience her artistic plan as changing, and Esperanza won't. In Esperanza's case, she stops working not because of convergence between work and plan, but because of external pressures. The kind of reply being considered on behalf of the Satisfactionist can't capture the important phenomenological difference between these two cases. In particular it can't capture the phenomenon of what it's like to decide an artwork is done even though it doesn't do what you want it to. Now, of course there may be a sense of 'plan' in which Esperanza will experience her plan as changing. After all, she didn't plan on writing a short story, and she did. But the kind of plan that's important for Satisfactionism is an artistic plan, the thing the artist is trying to realize by means of their artistic activity. The fact that Esperanza decides not to work more on her story doesn't mean she didn't have an artistic plan that called for a longer story.

Rohrbaugh may have a story to tell about cases where artists finish works that don't satisfy their plans.⁸ Rohrbaugh's account is that a work is complete just in case there is a kind of fit between work and plan, where the artist's plan is constituted by their intentions to bring about a work with certain features. But Rohrbaugh also mentions a "broader" notion, viz. that of a "practical state" guiding an artist in their creative endeavors. (Rohrbaugh 2017, 136) Thus, Rohrbaugh actually entertains a pair of questions as central to the Satisfactionist theory of completion:

How well does the work in progress reflect [the artist's] plan? Or, more generally, does the work in progress satisfy the condition imposed by the content of the guiding practical state? (Rohrbaugh 2017, 136)

One question here is what the guiding practical state is. Depending on the answer, it might be that we have two versions of Satisfactionism lurking. It's clear that

⁸Thanks to an anonymous reviewer for pointing this out to me.

Rohrbaugh doesn't take that to be the case, as in a later discussion note he says that the version of Satisfactionism given in terms of the artist's plan is just a more colloquial way of putting what he says about a guiding practical state. (Rohrbaugh 2018, 105) Unfortunately, Rohrbaugh leaves the notion of a guiding practical state (intentionally, I take it) unclear. Here's what I think the view amounts to. Since in reality plans are messy, we can conceive of them as consisting of intentions, desires, fears, impulses, obsessions, and a variety of other (cognitive and noncognitive) states, some of which will count as guiding the work, perhaps in virtue of their possessing a majority (or even a trumping) influence on the outcome of the artist's process.

In that case, Rohrbaugh might respond to the case presented above by claiming that if Esperanza caved to her publisher then her guiding practical state must have been something that led her to do so (e.g. a desire to please the publisher or a fear that she might not find further work). But then there is a kind of fit between the work and the guiding practical state when she alters her work to fit it, and the work is complete, even if there is a sense in which it doesn't do what Esperanza wants it to. In that case, the Satisfactionist can account for cases where a work is finished even though the author is not satisfied with it, as long as it satisfies the guiding practical state. Call this version of Satisfactionism GPS-Satisfactionism for "guiding practical state Satisfactionism."

I see two problems with GPS-Satisfactionism. First, it threatens to trivialize the phenomenology of dissatisfaction. GPS-Satisfactionism analyzes any case where an artist decides a work is done as a case where the work satisfies the guiding practical state (or at least where the artist thinks it does). But then the only room remaining for the artist to feel dissatisfied with their work is if there is internal conflict between the content of the guiding practical state and some other practical states of the artist. Surely, though, it's possible for an artist to be internally unconflicted, to end their

work on an artwork, and for that artwork not to do what they wanted it to. GPS-Satisfactionism has the result that this is impossible.

The other problem with GPS-Satisfactionism as we've been thinking about it is that it doesn't account for the difference between artistic and non-artistic practical states. Rohrbaugh introduces the twin notions of an artistic plan and a guiding practical state with two questions: "What is the artist trying to do? Or, more broadly, what practical state is guiding the agent in his or her creative efforts and what is the content of that state?" Rohrbaugh (2017, 136) The plan, as an artistic plan, is supposed to be the result the artist is trying to bring about *qua artist*. The guiding practical state is whatever is principally guiding the creative efforts of the artist. It seems Rohrbaugh wants to run these questions together as pursuing a single target, but the problem is that their targets can diverge in principle. For example, consider a case where an artist is trying to paint caricatures on a crowded street to make money. The artistic plan of each caricature will have something to do with simultaneous replication and exaggeration of the facial features (and some other significant features) of the subject. The guiding practical state of the artist, we can stipulate, is the desire to make money. Which one sets the standard for determining whether a work is complete, the plan or the guiding practical state? It seems there really ought to be two versions of Satisfactionism after all, then, each with a different answer to the question at hand. Call them "GPS-Satisfactionism" and "Plan-Satisfactionism".

The Satisfactionist faces a dilemma, then. Either they embrace GPS-Satisfactionism or Plan-Satisfactionism. GPS-Satisfactionism trivializes the phenomenology of artist dissatisfaction, and it makes artwork completion relative to non-artistic ends. GPS-Satisfactionism has a way of accounting for the Esperanza/Mirabel case, but at a very high theoretical cost. Plan-Satisfactionism, on the other hand, has no way of accounting for the Esperanza/Mirabel case. It seems that Satisfactionism doesn't

really explain what Psychologism does, then. This isn't too surprising. Psychologism aims to give us an account of when an artist has stopped working on something. Satisfactionism aims to give us an account of when a work does what an artist wants it to. Those are just different notions.

Satisfactionism's promise was also premised partly on the claim that it doesn't suffer the same problems as Psychologism. Cray (2018) has challenged that idea. Cray's challenge starts from the observation (made by Rohrbaugh) that plans change and develop. If an artist working on a particular artwork entertains a succession of plans, which plan is the one the work must satisfy to be complete (and why that one)? Suppose, for example, an artist happens to satisfy the plan p they are entertaining at a given moment t . But suppose the artist also realizes at t that they don't actually prefer p and so they continue to work, altering their plan and altering the artwork accordingly. Will we judge that the work was completed at t ? Unlikely. On account of scenarios like this, Cray thinks it necessary to give conditions for plan completeness. However, Cray takes it that any such conditions will be psychologistic, and thus leave the account of plan completion open to Rohrbaugh's Regress. In that case, an essential component of Satisfactionism (the notion of plan) would open the view up to the same kind of problem besetting Psychologism.

Rohrbaugh agrees this would be a bad consequence for his view, but he doesn't think Satisfactionism must suffer it. He says he's "tempted to say there is no such case as what [Cray] describes." (Rohrbaugh 2018, 107) Instead, Rohrbaugh guesses we are conceiving loosely of a case where someone satisfies a large part of their plan but not all of it (so there is no convergence) and they don't like the result. Rohrbaugh does acknowledge the possibility of cases in which there really is convergence between plan and work. For such cases, if the artist decides to change their plan and alter the work, Rohrbaugh prefers to describe them as cases of "change in an already completed work." (Rohrbaugh 2018, 106)

Of course, it's natural (although not uncontested) to think that sometimes an artist completes a work (for some sense of 'complete') and later changes it. However, Rohrbaugh seems to bite a bullet in leaving no room for cases like those described by Cray, those in which an artist really does satisfy their plan, immediately regrets it, and opts to alter the plan and keep going.

One option that's open to Rohrbaugh here is to deny that completion is terminal, i.e. to allow that a work can go from being complete to incomplete.⁹ This version of Satisfactionism would define completeness for both plans and works relative to each other. For any plan p and any artwork a , and for any time t_1 , if p and a converge at t_1 , then p and a are complete at t_1 ; and for any time t_2 , if p and a fail to converge at t_2 , then p and a are incomplete at t_2 . The problem with such a view is that it would have troubling practical and theoretical ramifications. For example, a decision about whether to display a work that turned on the work's completion status would be almost impossible to make. Likewise, if the completion-status of the work determines what it is, then a work's nature would be malleable and highly unstable. This is, of course, assuming Completion Monism. On Completion Pluralism, it's not particularly problematic for Satisfactionism to posit a non-terminal kind of completion, so long as some other kind of completion is terminal and applies naturally to cases where completion must be terminal.

Assuming Rohrbaugh doesn't opt for the non-terminal variety of Satisfactionism, I think Cray's argument sticks. There's no good reason to deny the possibility that an artist's plan and work converge and they immediately regret the convergence. One way the artist might respond to this is by throwing the work in the trash heap. Gover (2018, 42-48) considers such cases, in which an artist completes an artwork but de-

⁹Hick (2008a, 71), Gover (2015, 459) and Trogdon and Livingston (2015, 461) discuss reasons for and against adopting a terminal view of completion. Hick and Trogdon and Livingston opt for a terminal view. Gover argues that completion needn't be terminal on Psychologism.

clines to ratify it, thereby neglecting to admit the work to their corpus. Another way for the artist to respond, of course is to change the plan, and change the work. In a case where the artist's effort is continuous, it would be a strained interpretation of their activity to hold that the artist is now only making changes to their previously finished work. Instead, it's much more natural to acknowledge that the work is incomplete as long as the plan is subject to change and the work is ongoing. But in that case, an account of plan completeness is needed that doesn't relativize it to convergence with the work, and Rohrbaugh's Regress raises its head again. Once again, this problem presents itself only because of the assumed meta-thesis of Completion Monism. There is nothing obviously objectionable (from the Pluralist point of view) about the idea that a work is complete (in the sense Satisfactionism identifies) at any time at which there is a convergence between work and plan.

A final problem for the Satisfactionist has to do with what are commonly referred to as completions of artworks. Here is a somewhat common phenomenon: an artist begins an artwork, progresses it to a fairly advanced state, and is suddenly prevented from working on it anymore. Call these cases of unintended incompleteness. Maybe the artist dies, maybe the work is stolen from them, or maybe they are rendered incapable of working altogether. Sometimes in cases of unintended incompleteness, another artist completes the work at a later time. Christopher Tolkien and Guy Gavriel Kay completed J.R.R. Tolkien's *The Silmarillion*. Max Brod completed a number of Franz Kafka's incomplete works. Elsa Respighi completed Ottorino Respighi's opera *Lucrezia*.¹⁰

It's surprising, from the point of view of Satisfactionism, that one artist could complete another artist's work. That's because if the original artist is dead then the

¹⁰For more on these cases, see <http://www.tolkienlibrary.com/reviews/silmarillion.htm>, Kafka (1988), Diamant (2003), and <http://www.musicweb-international.com/respighi/lucrezia.htm>.

completing artist often won't have any access to the creative intentions of the original artist besides through the unfinished work itself.¹¹ (This is especially obvious in cases where the completing artist has no personal relationship with the original artist, as when Barry Cooper completed Beethoven's Symphony No. 10.) But it seems that later completions are rather common in cases of unintended incompleteness. This is hard to explain if Satisfactionism is our only account of completion.

Even worse for Satisfactionism: sometimes in cases of unintended incompleteness the incomplete work receives multiple completions. For example, Modest Mussorgsky's opera *Khovanshchina* was completed separately by Nikolai Rimsky-Korsakov, Maurice Ravel and Igor Stravinsky working together, and Dimitri Shostakovich. Anton Bruckner's *Symphony No. 9* has received no fewer than 10 completions.

What should the Satisfactionist say about such cases? Rimsky-Korsakov's completion of *Khovanshchina* is aesthetically distinct from Shostakovich's completion; they have different scores, they sound differently, and proper performances of each will leave audiences with divergent impressions as to their meaning. It seems thus that at most one of them could really be a completion of *Khovanshchina*, assuming the truth of Satisfactionism.

There are a number of options for the Satisfactionist. One option is to say that none of the so-called completions is *really* a completion of the work in question. The work remains incomplete. The so-called completions are distinct works based on the original, incomplete work. However, this way of describing things undercuts the sense that completions somehow fulfill the promise of the original work, that they really finish what was started, even if they do so in disparate ways.

¹¹Although there are notable cases in which completions proceed on the basis of notes, explicitly for the purpose of trying to complete the work in a way that satisfies the artist's plan. For example, Christopher Tolkien has published and is still publishing unfinished works of J.R.R. Tolkien using the latter's notes to guide his writing decision. Similarly, Robert Jordan left meticulous notes for the completion of his Wheel of Time series when he learned of his impending death, which were used to great effect by Brandon Sanderson. Thanks to Tim Juvshik for these examples.

Another option is to claim that all the so-called completions genuinely complete the work. To make sense of this, the Satisfactionist will have to distinguish each completion as a distinct, co-authored work. Each of the co-authored works would share an author and, perhaps, a proper part (depending on how one is conceiving of the metaphysics of the work-type in question). While this way of describing things captures the idea that each completion is legitimate, it undercuts the idea that these are completions of the original work, rather than merely co-authored works. In co-authored works, all the authors have a genuine say. In the case of completions, the original author's plan seems to be the only one that matters, and the completer works only to serve that plan as best they can.

The last option for the Satisfactionist is to say that only one at most (and perhaps none) of a set of so-called completions is actually a completion of the original work. Since completers do seem to work to serve the plan of the original author, it seems only one person at most could've achieved that (assuming two people don't independently produce identical completions). The obvious problem with this is that we tend to count many adequate attempts as legitimate completions. Of course (once again) this decision only arises for the Satisfactionist assuming they endorse Completion Monism. The Completion Pluralist can understand cases of posthumous completion using a distinct notion of completion: all so-called completions genuinely complete the work as versions of it that are complete according to the *art-kind* to which it belongs. The fact that the completers often don't have access to the original artistic plan for the work doesn't prevent them from making a version of it that's complete according to the salient kind.

In looking closely at both Psychologism and Satisfactionism, we've uncovered problems for both views. Assuming Completion Monism, neither view seems like a promising candidate for the one kind (or the most important kind) of artwork completion. Importantly, the problems suffered by each view hold in every case only on

the assumption that Completion Monism is true. If Completion Pluralism is adopted instead, the objections to both Psychologism and Satisfactionism as legitimate kinds of completeness lose their power.

Moreover, both views seem to be getting at something intuitive about artwork completeness. Sometimes, when we ask whether a work is complete, we're likely to be interested in one or both of the questions, "Is the artist done with it?" and "Does it do what the artist wants it to?" Thus, it's not just that both Psychologism and Satisfactionism both seem to fail as accounts of a monistic account of completeness. They do seem to identify something important about what we care about when we want to know whether works are complete. There can be better or worse accounts of when a work actually satisfies these intuitive criteria (as there are, for example, competing accounts of Psychologism), but adopting Completion Pluralism allows us to treat both questions as legitimate without eliminating one or prioritizing one as objectively special. The best explanation for the failure of these views on the assumption of Completion Monism thus seems to be that Completion Pluralism is true. We ought to drop the monistic assumption.

One might wonder whether, in giving this brief account of the unity of different kinds of completeness (and senses of 'complete'), I have just undermined my credibility as a Completion Pluralist. Completion Pluralism says that there are many kinds of artwork completeness and many corresponding senses of 'complete'. I've just suggested that all these senses are unified by a common feature. All the kinds of completeness are kinds of *completeness* insofar as they are conditions for equipoise and stability in the features of a work. But then it may sound as if I'm a Completion Monist, according to whom achieving equipoise and stability is fundamentally the only way for an artwork to be complete.

The important thing to notice is that on the view I've developed, everything Completion Pluralism says is true: there are many kinds of artwork completeness, many

corresponding senses of ‘complete’, and no kind of artwork completeness is objectively more important than any other. Adding the claim that the kinds of artwork completeness are importantly related by a common higher order feature doesn’t undercut any of these claims.

2.3 In Favor of Completion Pluralism

The foregoing negative argument had the primary aim of undercutting support for Completion Monism. Here I’ll provide a distinct, positive argument for thinking Completion Pluralism is true. As I’ve already stressed, we have a variety of reasons for caring about artwork completeness. We want to make sense of moral, legal, social, artistic, aesthetic, and ontological features of artworks and of our social practices of art.

These phenomena we want to explain are not always compresent, however. A work can be intuitively complete according to a sense of completeness that tracks one property but not another. Take an example. A curator for a museum has been waiting for months for an artist to complete a particular installation. The curator stops by the installment during her lunch break one day and asks an assistant, “Is it finished yet?” In this case, what the curator is interested in, we can presume, is whether the artist is done working on the installation. It doesn’t particularly matter for the curator’s interests whether the piece satisfies the artist’s plan, whether the installment is as aesthetically revelatory as it could be, or even what the thing itself is. On the other hand, suppose the artist takes a step back and asks herself, “Is it finished?” Rohrbaugh’s *Regress* should have us convinced the artist is not interested in the question of whether she herself is done working. Rather, she is asking herself something about the work. Suppose the artist decides it’s finished and ratifies it as her own, allowing the plans for its display to move forward. Imagine now an audience member, viewing the work after its eventual date of revelation, unsure of

the norms appropriate to the genre in which the artist is working and looking at the seemingly randomly placed materials that make up the installation. This person asks his friend, “Is it finished?” The audience member, we can imagine, is wondering whether what he’s looking at is a complete instance of the kind of thing the artist was attempting to make. It may look to him like an incomplete statue or some kind of half-finished architecture, and his question is in part an attempt to navigate the question of whether it’s an incomplete instance of one art-kind or a complete instance of another.

There isn’t a single, overriding concern we have when we’re inquiring into the nature of artwork completeness or the question of whether some particular work is complete. Likewise, there isn’t a single property of artworks that all our completion-concerns track. Moreover, the properties our completion-concerns track can be instantiated either together or separately, so that depending on the context people will be inclined to call a work complete (or incomplete) if it satisfies (or fails to satisfy) only some of the many properties we associate with completeness (as for example with the *non-finiti* discussed earlier).

What seems to be the case, then, is that there are different senses of ‘complete’ that can be applied separately. ‘Complete’ is ambiguous between a number of different, related conditions, none of which is objectively more important than any other. This should be the case only if one of two alternatives obtains: Either Completion Monism is true and our concerns about artwork completeness don’t track the truth about completion, or Completion Pluralism is true. We have no reason to think that our completion-concerns fail to be truth-tracking, though. We clearly do have some handle on when a work is complete, for some senses of ‘complete’. We ought to reject the first disjunct above, then, and conclude that Completion Pluralism is true.

One potential objection to this line of argument: there’s clearly *something* the curator, artist, and museum-goer are all talking about when they’re talking about

completeness. Completion Pluralism seems to have the consequence that they're just talking past each other. I agree this would be bad. We don't want to say that 'complete' is ambiguous in the way that words like 'bank' or 'bat' are. We need an explanation of why all these kinds of completeness count as kinds of *completeness*. A natural thing to say here is that 'complete' is a polysemous term. Polysemous terms are those with multiple, distinct but closely related meanings. I won't a detailed theory of how the multiple meanings of 'complete' are closely related. However, here's a suggestion: every kind of completion involves a way in which an artwork attains a certain equipoise, a balance in its features that results in greater stability and resistance to change. Psychologism identifies as a condition equipoise in the artist's effort's. Satisfactionism identifies as a condition equipoise between the intrinsic state of the work and the artist's plan. Both kinds of completeness have the result that a work is less likely to undergo change once it's attained the sort of equipoise in question. In each case, I posit that a kind of completeness will strike us as intuitively so only if it has identified a condition for some form of equipoise (and resultant stability) in the work. As there are many dimensions along which a work can attain this, there are many ways it can be complete.

2.4 Putting Pluralism to Work

In the remainder of this chapter, I apply the thesis of Completion Pluralism. To do so, I first identify what I take to be four intuitive kinds of artwork completion that apply to artworks generally. I show how each addresses a distinct concern we sometimes have when we ask whether a work is complete. Additionally, I show how each kind of completion can be present in the absence of some other intuitive kinds of completion. Finally, I introduce a kind of artwork completion I take to apply to stories only, narrative completion.

We've already identified aesthetic completion above as an intuitive kind of completeness an artwork can have when it's "worthy of standing by itself, as an object of aesthetic enjoyment." (Beardsley 1982, 55) Others have interpreted Beardsley as claiming that all artists are trying to make aesthetically complete works. While that seems a tad strong, surely it seems right that at least some artists are going for this. Thus, when an artist is trying to complete their work, a good candidate for what at least some have in mind is aesthetic completion. After all, we know from our discussion of Rohrbaugh's *Regress* that psychologistic completion is not what artists are after. Satisfactionism may seem to identify what artists are after. Naturally, we might think, artists want their works to satisfy their plans. We should be circumspect about this line of thought, though. Artists very rightly are sometimes worried about whether their plan will be the best thing for the work. Rather than getting what one wants, then, a mature artist may wish merely that the work be a worthy object of appreciation and engagement, whether or not it ends up satisfying their plan.

Sure enough, this just may come about even while the work remains incomplete in other ways. For example, suppose the work attains aesthetic completion even as the artist is still working on it. It's not psychologically complete, since the artist isn't done. There isn't convergence between work and plan, we can assume, as evidenced by the artist's continuing to work.

Psychologistic completion is, intuitively, the kind of completeness a work has when the artist is done with it. While it doesn't seem to be a primary concern of artists, it certainly seems to be a concern of art curators and audiences. Psychologistic completion can clearly obtain in the absence of aesthetic completion. I've argued above that psychologistic completion can obtain without convergence between work

and plan (or vice versa) as well, in the Esperanza/Mirabel case and case of the artist who changed their plan upon achieving it.¹²

Satisfactionism identifies another kind of completeness, “satisfactory completion,” that certainly seems like a relevant concern for those appreciating and evaluating art. Part of the evaluation of art is the evaluation of the artist, and one wants to know that the work is a proper window into the artist’s intentions. This is the case only if it does what the artist wants it to, i.e. if there’s convergence between work and plan. This can come about, however, in works that are aesthetically incomplete, and even as the artist is still laboring. Thus, satisfactory completion can obtain in absence of aesthetic or psychologistic completion.

It’s worth identifying another kind of completeness that applies to artworks in general to show that the extant theories don’t exhaust the kinds of completeness a Completion Pluralist might posit. Return to the case of Schubert’s *Symphony No. 8 in B minor*, i.e. the *Unfinished*. Recall that, in fact, Schubert never finished the *Unfinished*. Let’s stipulate that Schubert formed a completion disposition with respect to the *Unfinished*, based on the exercise of his rational faculties. Moreover, let’s stipulate (so that it doesn’t matter whether our scenario is true or merely imagined) that Schubert’s plan for the *Unfinished* was satisfied. Clearly the *Unfinished* is also worthy of standing by itself, as an object of aesthetic enjoyment. In the scenario we’re considering, the *Unfinished* satisfies the definitions for artwork completion for each of the kinds explored thus far. In such a scenario it would be aesthetically, psychologically, and satisfactorily complete. But it still seems that people would think it was incomplete in an important sense. What sense is that?

¹²Importantly, even though I’ve argued for Completion Pluralism above, there is still a question about the right account of Psychologism, or the kind of completeness that results from an author’s being done with a work. Several accounts exist in print. The truth of Completion Pluralism doesn’t entail that all arguments about artwork completion are moot.

I think the relevant sense in which the *Unfinished* would be incomplete in this scenario is relative to kind. The *Unfinished* is a symphony. Symphonies typically have a particular structure, usually involving four movements. The *Unfinished* was left with only two movements. Consider, then, a fourth kind of artwork completeness, “kind-relative completion.” For any object x and any artistic kind K , x is complete relative to kind K iff x is a K , but x is lacking some property F that is normative within K . The notion of a property’s being “normative within” a kind I take from Dodd (2007, 32-3), who makes a distinction between types that do and types that do not admit of improperly formed tokens. The former kind Dodd follows Wolterstorff (1980) in calling “norm-types”. Some examples of norm-types include stories, novels, and words. What it is to be a properly formed token is determined by the properties that are normative within the type, none of which a properly formed token lacks. My notion of kind-relative completion thus relies on the idea that art-kinds are like norm-types: there can be artworks that are members of the kind even though they lack some property that is constitutive of what it is to be a properly formed member of the kind. Just as with other kinds of artwork completeness, I take it my account of kind-relative completion identifies a condition for equipoise in a work, viz. equipoise between the intrinsic state of the work and the range of intrinsic states that are normative for works of that kind. If a symphony lacks a final movement, then there is something out of balance between that work and the norms of the kind. When the symphony is completed, kind-relative equipoise is achieved. My suggestion is that the *Unfinished* is a kind-relative incomplete symphony.

We began with the question of what a complete story is. My answer should be clear now: there is more than one way in which a story can be complete. Here at the end of this discussion, however, we should wonder if there are any kinds of completeness that apply particularly to stories. I think there’s at least one, a kind of completeness I want to call “narrative completion”. A common way of understanding the function of

stories is that a good story provides understanding by way of its narrative. Velleman (2003, 19) argues for an account of narrative understanding according to which a good story provides emotional understanding. A story does this when some beginning events set up an emotional ‘tick’ and the concluding events provide the answering ‘tock’, with intervening events delaying the resolution. “[t]he audience of a story understands the narrated events, first, because it *knows how they feel*, in the sense that it experiences them as leading through a natural emotional sequence; and second, because it *knows how it feels about them*, in the sense that it arrives at a stable attitude toward them overall.” Narrative understanding doesn’t come about by happenstance, of course. Velleman theorizes that a story leads an audience to this understanding when its events are presented (beginning, middle, and end) in a way that leads the audience through the relevant emotional experience. Thus, a telling of a story that provides us with no opportunity for emotional engagement will not be a candidate for providing us with narrative understanding either, on this view. Whether a story (or a telling of one) can provide us with narrative understanding is a matter of fit between the features of the story and the way we are emotionally constituted.¹³

On my understanding of narrative completion, a story is narratively complete if it provides the audience with the kind of narrative understanding identified by Velleman, i.e. an emotional understanding of the events of the story taken as a whole, so that the audience both knows how they (the events) feel, and how it (the audience) feels about them. I’ll defer to Velleman’s account of how it is that a good story does this. The details aren’t important for the truth of the claim that a story is narratively complete when it provides the audience with narrative understanding.

¹³It’s important to note that Velleman offers his account of narrative understanding as a condition for what it is to be a story. That’s not what I’m doing. I take it that bad stories and incomplete stories are still stories. On my view, narrative understanding is a way of understanding what it is for a story to be more or less complete (and so, in one way, better or worse) as a story.

In order to see that narrative completion is a distinct kind of completeness, it's worth identifying a couple cases, one in which narrative completion is present in the absence of other kinds of completion, and one in which it's absent in the presence of other kinds of completion. Take the former case first. Imagine Edmund Spenser in the late 1580s, working fastidiously on a draft of Book I of *The The Faerie Queene*. In the imagined scenario, Spenser's draft reaches a point where the events certainly do provide any potential audience with the emotional tick, tock necessary for narrative understanding. The plot is largely as we know it. However, the verse needs quite a bit of work. There are typos and problems with the structure of a significant number of stanzas. The work fails to be aesthetically complete, then. While it's a compelling story, it's not quite worthy of standing by itself as an object of aesthetic enjoyment, at this point. Likewise, it's not psychologically complete. Spenser isn't done with it. Neither is it satisfactorily complete, naturally. The work isn't intrinsically such as to have satisfied Spenser's artistic plan. Given that many of the stanzas are structurally incomplete, the book is not kind-relative complete either. A poem written in Spenserian stanza must meet certain structural conditions to be a complete instance of its kind. At the imagined point in time we're considering, then, the version of Book I of *The The Faerie Queene* would lack all four kinds of completion surveyed above, but it would still be narratively complete.

Take, on the other hand, Daniil Kharms's short story, "Sonnet":

A peculiar thing happened to me: I suddenly forgot what comes first—
7 or 8?

I set off to ask my neighbors what their thoughts were on the matter.

How great was their surprise—and mine, too—when they suddenly realized that they also could not recall the counting order. 1, 2, 3, 4, 5, and 6 they remember, but what comes next they've forgotten. . . .

We all went down to the commercial store called "Gastronom" that's on the corner of Znamenskaya and Basseynaya Streets, and asked the cashier there about our incomprehension. Smiling a sad smile, the cashier extracted a small hammer from her mouth, and twitched her nose slightly. She said: "In my opinion seven comes after eight, but only when eight comes after seven."

We thanked the cashier and in utter joy ran out of the store. But after pondering deeply the cashier's words grief came over us again, for it seemed that not a word of hers made any sense to us.

What was there to do? We went to the Summer Garden and began counting the trees there. But when we reached the number 6 we stopped counting and began to argue: some thought seven was next in the order, others—8.

We would have argued very long, but, luckily, just then somebody's child toppled off a park bench and broke both of its jaws. This distracted us from the argument.

After that, everyone went home. Kharms (2007)

“Sonnet” seems to be a good example of a story that's distinctly lacking in narrative completion. In fact, that's kind of the point. The story sets up the tick (“Why can't these people remember that 7 comes before 8?!”), it delays the resolution, building the tension in a conventional manner, and then it refuses to provide the accompanying tock, sending everyone home without a resolution. How should we feel about these events? What do they mean? Who knows.

“Sonnet” appears to be complete in every other way we've considered here, however. It's aesthetically complete: the story certainly has no significant aesthetic flaws. Reading it provides one with an experience of surreality. It causes one to reflect on potential meanings, the kind of puzzle-box experience common to much contemporary art. “Sonnet” is also psychologically complete, or we have no reason to think it isn't. Kharms seemed to have formed a completion disposition towards the work, at least, not making any changes to it. These are also good reasons to think it's satisfactorily complete as well, that the work does what Kharms wanted it to do. Finally, it's complete relative to its kind. Unlike symphonies and epic poems, the normative constraints on short stories are limited: it has to be a story and it can't be too long. This work clearly checks those boxes. There's reason to think Kharms's work is complete in every way we identified above, then. But there still seems to be a way in which it's incomplete, a way in which it leaves one with a feeling that it lacks

something (albeit intentionally so). My posit is that the incompleteness is narrative. “Sonnet” is an intentionally incomplete story.

By identifying kind-relative completion and narrative completion as distinct kinds of artwork completion I take myself to have further demonstrated the promise of Completion Pluralism for opening up new horizons in research on artwork completion. The cases examined in this section show that these kinds of completion are of legitimate interest and that a work can be have or lack kind-relative or narrative completion regardless of what other kinds of completion it has or lacks.

2.5 Conclusion

The recent thread of literature on artwork completeness began with an observation about two kinds of completeness. I’ve argued that the significance of that observation has been entirely overlooked. Proceeding on the (either explicit or operational) assumption of Completion Monism has led to debates that are leading the wrong direction. In reality, there isn’t just one (or one most important) kind of completeness. A work might be worthy of standing on its own as an object of aesthetic enjoyment. It might have satisfied the artist’s plan. The artist might be done with it. It might or might not be a properly formed member of its kind. These (and perhaps other) notions can use theoretical refining, so that there are certainly best versions of Psychologism, Satisfactionism and of other notions of completeness. The way forward for understanding artwork completeness lies at least in part in taking on the project of discovering the many kinds of completeness and the ways they can be instantiated either together or apart.

CHAPTER 3

FICTIONS ARE FABRICATION

3.1 Introduction

Having argued for accounts concerning what it is to make up a story and what it is to finish a story, I want now to take up the question of what the distinction between fictional and non-fictional stories is.¹

It seems to me that fictions are fabrications. What I mean is that a story is a fiction to the extent that it's made up.² This strikes many as an unremarkable claim, a platitude unworthy of defense. Nevertheless, in this chapter I will offer both an articulation and defense. I offer a defense in part because this claim has taken abuse from others writing on the nature of fiction, and in part because I think it fits naturally with the views for which I've argued in chapters 1 and 2. Thus, a vindication of the view that fictional stories are those that are made up serves as a confirmation of the picture I've argued for thus far.

¹Throughout this chapter I often refer to fictional stories as fictions. It's important to note that this is a shorthand, one I drop when the distinction between non-narrative fictions and fictional stories becomes relevant. On my view, the fiction/non-fiction distinction and the narrative/non-narrative (or story/non-story) distinction overlap. This is to say that some stories are fiction, some non-stories are fiction, and some of each are non-fiction as well. What I'm interested here is to say (keeping the focus on stories) just when a story is fiction or not. The account I give should be directly applicable in the non-narrative case as well, but I won't endeavor to show that.

²Importantly, as will become clearer later, when I say that a fictional story is one that's made up, I don't mean it's one that's created. I mean that the story was identified in the way discussed in Chapter 1. This theory of fiction is logically consistent with a Creationist metaphysics of stories, but it isn't the same thing.

The claim that fictions are made up is certainly confirmed by appealing to the dictionary and common usage.³ Dictionary definitions aren't decisive evidence to settle conceptual or metaphysical matters. However, in this case, I think the dictionaries have hit the conceptual nail on the head. Our concept of *fiction* just is the concept of a made up story. The complimentary concept of *non-fiction* is the concept of a story that hasn't been made up.

Currie (1990, 1) claims, "There can hardly be a more important question about a piece of writing or speech than this: Is it fiction or nonfiction?" I take it he's right about this, but it's worth dwelling on the point for a moment to be clear on why. In a cultural environment in which mistrust in the news and misinformation are among the chief public concerns, it can be helpful to gain a clearer understanding of what features of a story make it fiction, as well as whether those features can be unintentionally bestowed on a story and whether they might be exemplified by things like news reports. As a matter of more perennial interest, philosophers of art should want to understand the nature of the fiction/nonfiction distinction. It seems analytic that they stand in an exclusive relation. Understanding the distinction should also add to our understanding of the differences in the way we engage with instances of the two categories.

Call the theory that a fiction is a made up story the Fabrication Theory. Despite the intuitive appeal of the Fabrication Theory, it has gone largely undiscussed. I

³For example, Merriam-Webster gives this as the primary definition of 'fiction':

1a : something invented by the imagination or feigned
specifically: an invented story
"...I'd found out that the story of the ailing son was pure fiction."
— Andrew A. Rooney

The OED, in two separate definitions, mentions literature "that describes imaginary events and people," and things that are "invented or untrue." <https://www.merriam-webster.com/dictionary/fiction>. <https://en.oxforddictionaries.com/definition/fiction>. Both visited 12/03/2018.

know of only two philosophers who have defended it.⁴ When the Fabrication Theory is mentioned, it's usually quickly dismissed.⁵ Since a number of philosophers have adopted theories of fiction in conflict with the Fabrication Theory without much engaging with the reasons that favor it, and since the Fabrication Theory seems like it ought to be the theoretical starting point, I set out here to offer a case in favor of it.

In order to defend the Fabrication Theory, in §2 I present two important desiderata for a theory of fiction: a theory of fiction is good insofar as it exhibits both intensional adequacy and explanatory power. Next, I outline two competing theories of fiction and show both how are lacking with respect to at least one relevant desideratum. The first theory takes fictions to be props, devices of varying types the function of which is to produce imaginings. The second theory takes fiction to be a genre, and individual fictions to be works within that genre. Both theories suffer from failures of intensional adequacy. After arguing against them, I flesh out the Fabrication Theory by showing how its central claim can be understood in terms of my earlier analysis (in chapter 1) of what it is to make up a story. In providing a full account of the Fabrication Theory that avoids the problems of competitors, I take myself to be making a case for the theory, since the idea that fictions are fabrications is the natural starting point, and to date there seem to be no good arguments given against it.

⁴Both Deutsch (2000) and Young (2002) embrace the Fabrication Theory, giving an analysis of fiction in terms of making something up. Both of their understandings of what it is to make something up diverges considerably from mine, but in different ways. I'll discuss the details of their views below.

⁵For example, Friend (2012) says "this position is neither plausible nor popular," but the claim about its plausibility goes unexplained. Stock (2017, 156-7) mentions Deutsch's theory but her only argument against it is to argue more generally against "patchwork" views of fiction, on which a work of fiction often consists of a mixture of fictional and non-fictional statements. I'll discuss patchwork views in more detail below. Matravers (2014, 31) mentions a theory on which fictions are "narratives that deal with situations believed to be non-actual." He says that such a theory is "patently inadequate", referencing what he takes to be a discussion of the theory in Currie (1990, 4-9). However, Currie's discussion of the semantic properties of fiction there only goes so far as to show that fictions can be entirely true (by virtue of the fact that someone can accidentally make up a story that is in fact true).

3.2 Desiderata for a Theory of Fiction

A theory of fiction is an answer to the question, “What are fictions?” An ideal theory of fiction would give an answer that both accurately predicts which things are fictions and makes sense of the way we interact with fictions. If a theory failed to meet either of these constraints, then it might be true but it wouldn’t be very illuminating. If it didn’t accurately predict which things are fictions, we couldn’t trust it to be telling us about what fictions actually are. If it made our practices seem like nonsense, we would need a very good argument to believe it, since our fiction-directed practices are our evidential starting point for our thinking about fictions. By “fiction-directed practices” I have in mind the variety of ways we interact with fictions that are characteristic of how we interact with fictions. Some examples include the following: we classify certain things (and not others) as fictions; we evaluate fictional stories according to the presence of features like narrative explanation, originality, surreality, creativity, and other aspects (most of which have nothing to do with truth, except in certain genres of historical fiction); we use fictions to stimulate the imagination, as occasions for pretense, and as means of escapism; and we treat authors of fictions as authoritative concerning what’s true in their fiction, which is different than how we treat authors of non-fictions. I don’t take this list to be exhaustive, but merely illustrative of some of our fiction-directed practices.

To be an ideal theory of fiction, then, a theory \mathcal{T} should have at least the following two qualities:

1. **INTENSIONAL ADEQUACY:** Necessarily, for every object x , \mathcal{T} says that x is a fiction if and only if x is a fiction.
2. **EXPLANATORY POWER:** \mathcal{T} explains why we engage in the fiction-directed practices we do.

Intensional adequacy is just the basic condition that an ideal theory of fiction should not misclassify things. It shouldn't say of fictions that they aren't fictions, and it shouldn't say of things that aren't fictions that they're fictions. Moreover, this desideratum holds not only with respect to what's actual, but also with respect to the merely possible. A theory \mathcal{T} might get lucky and correctly classify all and only actual fictions as fictions. But say it's a consequence of \mathcal{T} that some possible object o would be a fiction, even though o would not be a fiction. In such a case, \mathcal{T} would fail to be fully intensionally adequate. It's natural to think some theories are closer than others to meeting the desideratum of intensional adequacy, even if none meets it perfectly. As a simple way of regimenting this, allow there to be two theories \mathcal{T} and \mathcal{T}' , such that every case \mathcal{T}' classifies correctly, \mathcal{T} classifies correctly, but such that there are cases \mathcal{T} classifies correctly that \mathcal{T}' doesn't. In this case, we can say that \mathcal{T} is intensionally adequate to a higher degree than \mathcal{T}' .⁶ It's also natural to think some ways of failing to meet this desideratum are worse than others (some misclassifications count as worse mistakes than others). I won't attempt to regiment this idea, however. Usually, we apply the criterion of intensional adequacy by appeal to cases: if there's a clear case of a (possible or actual) fiction that a theory classifies as a non-fiction (or vice versa), then that counts against the intensional adequacy of the theory.

Explanatory power is a virtue a theory of fiction has when it helps us better understand why people say or do the kinds of fiction-directed things mentioned above. A theory of fiction \mathcal{T} has explanatory power to the degree that it helps explain why it makes sense to participate in these kinds of practices, given what \mathcal{T} says fictions are. The more a theory's account of what fiction is coheres with our fiction-directed practices, so that one should expect us to treat fictions the way we do given what

⁶Thanks to Alejandro Pérez Carballo for suggesting this way of putting things.

the theory says they are, the more explanatory power it has. A theory \mathcal{T} lacks explanatory power to the degree it does not explain our fiction-directed practices. A theory \mathcal{T} incurs an explanatory debt if \mathcal{T} suggests our practices are incoherent given what \mathcal{T} says fictions are.

One might worry about whether the appeal to explanatory power here is felicitous. For example, someone might worry about whether the explanations for our practices might be better sought not in the natures of the objects themselves but in cultural history, sociology, or psychology.⁷ My response to this worry is to point out that these kinds of explanations don't rule each other out. Understanding why we engage in the fiction-directed practices we do will plausibly take understanding *both* what it is for something to be a fiction *and* the cultural history of our engagement with stories (along with the relevant sociological and psychological information that may explain some of that history).

An example in a different case might illustrate how a theory can help explain a social practice. As an example, take the painting-directed practice we have of privileging the state of the work upon its publication. If a painting is damaged or faded over time we employ experts to restore it to its original (i.e. original published) condition. If a painting is damaged, we would not accept an intrinsic duplicate of the painting in its original published condition as an appropriate substitute for aesthetic appreciation. Intrinsic duplicates are mere copies, or worse, forgeries. Perhaps they are worth some amount of curiosity or study, but they certainly don't have the same artistic value. Contrast this with our attitudes towards novels: except in cases where a token novel has special significance as a physical artifact (e.g. it's a first edition, signed copy, or has significant sentimental value), we are much more ready to replace a worn copy with a new one.

⁷Thanks to Amie Thomasson for this point.

A theory of paintings needs to take these practices and attitudes into account. Now consider two theories of paintings. One is a type-token theory on which a painter selects pre-existing image types and tokens them on canvases. Let this theory identify paintings with the image type selected by the painter. A second theory is a physical object theory. On that theory the painting is just the physical object the painter creates, i.e. the paint on canvas. The type-token theory will have a much harder time, *prima facie*, explaining the social practices mentioned above than the physical object theory. For example, if a painting is a pre-existing image type, then it won't seem to matter much which token of it we display for aesthetic appreciation, so our disdain for copies or forgeries doesn't make a lot of sense. Importantly, the explanation of our painting-directed practices that's provided by a theory of what paintings are still leaves room for historical, sociological, and psychological explanations of those same practices. After all, we might wonder *why* we care more about token paintings than we do about image types (when the same doesn't seem to be true of music, for example). A historical or psychological explanation can probably help provide the answer.

In introducing these desiderata for a theory of fiction, I'm not claiming that a theory must satisfy any particular degree of intensional adequacy or explanatory power in order to be considered worthy of consideration. Every theory exhibits both to some greater or lesser degree, and it may be that a very great degree of explanatory power is worth some loss of intensional adequacy (as measured by a failure to match intuitions in a number of cases) or vice versa. What I *am* claiming is that a theory of fiction is better to the degree that it exhibits these qualities. Thus, we should prefer a theory that exhibits both to a high degree over a theory that exhibits only one to a high degree.

My notions of intensional adequacy and explanatory power bear similarities to what Kathleen Stock characterizes as opposing methodologies for theories: concep-

tual analysis vs. explanation.⁸ On Stock’s characterization, conceptual analysis aims merely at matching folk intuitions, while explanatory theories aim at explanatory power (while trying not to result in too great a mismatch with the folk). Stock and I agree that the intensional adequacy of a theory can be measured by its resistance to counterexamples. However, Stock believes takes this methodological goal to be opposed to that of explanation. I don’t see any reason to characterize accuracy and explanation as opposed in the way Stock does. Rather, when weighing theories against each other, both qualities should be seen as virtues.⁹ If we can identify a theory that does both to a high degree, there’s no need to choose between desiderata.

One other constraint worth mentioning here: a theory of fiction should illuminate something about the nature of non-fiction and our “non-fiction-directed practices”. The relevant contrast to the category of *fiction* is the category of *non-fiction*. Our fiction-directed practices are linked in important ways to contrasting non-fiction-directed practices, so that how we evaluate a work, what we think about the author’s authority over the work, and how we experience a work while reading or viewing it for the first time often depends in large part on whether we think the work is a work of fiction or non-fiction. Whether we use a story primarily to stimulate the imagination or as a means to acquire information is a matter of whether story is fiction or non-fiction. Likewise, whether we take an author to be authoritative of what happens in their work or to be subject to norms of accuracy depends on whether we take that author to be presenting a work of fiction or non-fiction. What we really want then is a full account of the fiction/non-fiction distinction, one that explains the social practices aimed towards both kinds.

⁸Stock (2016, 204-5)

⁹Stock (2016, 205) herself tacitly acknowledges this when she admits that an explanatory theory should “at least fit with some uncontroversial cases of the original folk concept, in order to avoid the impression of talking past the original topic.”

My basic argument in what follows is that the Fabrication Theory of fiction meets both the above desiderata to a higher degree than any other proposed theory. The Fabrication Theory is ideal, as measured by these two desiderata. Perhaps there are other virtues to consider, but as it stands the Fabrication Theory is a cut above.

3.3 Fiction Is Not a Prop

Most theories of fiction on offer are some version of what I label the ‘Prop Theory’ of fiction. The Prop Theory says that fictions are objects the function of which is to prompt imaginings in an audience. It’s worth spending significant time on the Prop Theory because of the way it has dominated the philosophical conversation about the nature of fiction in recent years. The originator of the Prop Theory is Walton (1990).¹⁰ On Walton’s view, the prop metaphor is taken seriously. Works of fiction literally function as pieces in games of make-believe. They gain their function by playing a role in a certain society that uses them.

There are two separate claims that constitute the heart of Walton’s view, then:

SUFFICIENCY CLAIM: Possessing the function of prompting imaginings is sufficient for being a fiction.

NECESSITY CLAIM: Possessing the function of prompting imaginings is necessary for being a fiction.

Not all Prop Theorists endorse both claims, but all agree with Walton in endorsing at least one. Walton’s view is distinctive in its account of how something possesses the function of prompting imaginings. While it’s commonplace among many more recent versions of the Prop Theory to posit that only the artist’s intentions play the role of determining whether a work has the function of prompting imaginings in an audience, Walton posits instead that a work has this function if and only if it plays a certain role in a society.

¹⁰Other proponents of the view include Currie (1990), Davies (1996), García-Carpintero (2007), Lamarque and Olsen (1994), and Stock (2017).

For Walton, a prop is an object that generates a fictional truth within the context of a particular game of make-believe, according to certain generation principles.

Props generate fictional truths independently of what anyone does or does not imagine. But they do not do so entirely on their own, apart from any (actual or potential) imaginers. Props function only in a social, or at least human, setting. The stump in the thicket makes it fictional that a bear is there only because there is a certain convention, understanding, agreement in the game of make-believe, one to the effect that wherever there is a stump, fictionally there is a bear. [This is] a *principle of generation*. (Walton 1990, 38)

Principles of generation are psychosocial principles by which prop-facts get generated in the context of particular games of make-believe, games that are taken on by particular social groups. On Walton's view, works of fiction are props in these games.

Walton's theory does not have a high degree intensionally adequacy. When we talk about fictions, we generally have in mind narrative works, like novels, poems, films, or comics. Sometimes we have in mind non-narrative representations where the content of the representation has been made up. What we don't have in mind is a mere prop in a game of make-believe, something that in itself has no content and by itself represents nothing. Perhaps a novel does function as a prop when I use it to imagine some scenes in a story, but the same novel can also be a shield in a make-believe game of "slay the dragon" I play with my daughter. It's not a fiction in virtue of it's ability to do the latter sort of thing. Walton's theory lets in too many things as fictions by allowing anything that plays the role of prop to be a fiction.

The over-broadness of Walton's Prop Theory appears to be an intended feature, however. Walton (1990, 72) worries that the everyday concept of fiction will result in cases of some works that count as both fiction and non-fiction. His goal is thus to introduce a theoretical alternative to the common concept that explains our engagement with fiction without running fiction and non-fiction together. We can understand Walton, then, as providing a theory that has significant explanatory power, while at the expense of giving up significant intensional adequacy.

Currie (1990) is another Prop Theorist. Currie differs from Walton in his view of how fictions gain their function as props. For Walton, whether something has the function of a prop is a contextual matter, something that's relative to societies. (Walton 1990, 52-53) Currie's view, on the other hand, is that something has the function of prompting make-believe in virtue of the intentions of its maker. Currie construes fictions as utterances of a certain kind, viz. fictive utterances, and something can be a fictive utterance only if the person from whom the utterance originates intends it that way.¹¹

Another important difference between Currie and Walton is that Currie denies the Sufficiency Claim. He includes an additional, necessary condition for something to be fiction. He takes his pair of necessary conditions to provide an account of what's sufficient for fiction. Cases that push Currie away from the Sufficiency Claim include works that are entirely (or even just mostly) true, in which the truth of the work is non-accidental, but such that the function of the work is to produce imaginings in the audience. Someone might produce a purely historical report but with the reflexive intention that the audience make-believe it, for example. (Friend (2012, 184) offers Truman Capote's *In Cold Blood* (1965) as a historical example of such a work.) Currie concludes that in order to be fiction an utterance must not be non-accidentally true.

Currie's considered version of the Prop Theory is thus that, for any x , x is a fictive utterance if and only if

- (i) x was made with a reflexive intention to produce make-believe, and

¹¹Fictive communication in general is distinct in being a kind of communication wherein the communicator has a reflexive intention to produce make-belief in the audience. A reflexive intention is a complex intention to produce some mental state in an audience, and furthermore that the mental state be produced *by means of* the recognition by the audience that the communicator has that intention. (The idea of reflexive intentions is due to Grice (1957).) A reflexive intention to produce make-belief is thus the intention to produce make-belief in an audience by means of the audience's recognition of one's intention to produce make-belief in them. A fictive utterance is just an utterance produced with a reflexive intention to produce make-belief.

(ii) x is not non-accidentally true.

Currie says that “a work as a whole is fiction if it contains statements that satisfy the conditions of fictionality I have presented, conditions we can sum up briefly by saying that a statement is fiction if and only if it is the product of an act of fiction-making. . . and is no more than accidentally true.” (Currie 1990, 49) The theory is that a fiction will generally turn out to be a kind of “patchwork” of statements that are fictional and statements that are non-fictional, i.e. a patchwork of fictional utterances and non-fictional utterances. The work as a whole will be fictional in virtue of facts about the statements that author it.¹² Call the view that fictions are made up of both fictional and non-fictional utterances the “Patchwork View” of fiction.

Currie’s Prop Theory has been extensively criticized because of its commitment to the Patchwork View. Matravers (2014, 39), Stock (2017), and Walton (1990, 79) all argue against the Patchwork View, and by extension against Currie’s Prop Theory. I’ll defer a discussion of the objections to the Patchwork View for now, but this leads to a final candidate version of the Prop Theory, offered by Stock (2017). Stock, like Currie, analyzes fiction in terms of fictive utterances. She holds that, for any x , x is a fiction if and only if it consists entirely of fictive utterances. Stock’s theory is thus explicitly anti-patchwork. A fictive utterance, for Stock, is an utterance made with the reflexive intention to produce imaginings. This is not any different from Currie’s understanding of what makes an utterance fictive. However, Stock combines this understanding with the explicitly anti-patchwork condition that a fiction can’t contain any non-fictive utterances. A non-fiction, on the other hand, is “any set of utterances *not* wholly reflexively intended by their utterer to be imagined; though some of them might be.” (Stock 2017, 160) That is, although Stock’s Prop Theory

¹²Although Currie (1990, 49) thinks it’s a “bad question” to ask what ratio of fictional-to-non-fictional statements one needs to get a fictional work. The view seems to be something like this: a work authored entirely of fictional utterances is clearly fictional; a work authored entirely of non-fictional utterances is clearly non-fictional, and there are harder cases in between.

rejects a Patchwork View of fiction, it must embrace a Patchwork View of non-fiction as a direct consequence.

What Walton (1990), Currie (1990), and Stock (2017) have in common as Prop Theorists is a commitment to the view that a fiction is something the function of which is to get an audience to make-believe or imagine some content. There are important points of disagreement: Whose intentions matter for determining the function of the representation? Is having an imagination-producing function sufficient or merely necessary for being a fiction? What is the relationship between the fictionality of the work and the fictionality of the utterances that make it up? However, while versions of the Prop Theory differ on how they answer these and some other questions, they draw a similar picture of what fictions are in terms of their function. That picture constitutes a shared theoretical framework.

One way to critically discuss the Prop Theory would be to focus on each iteration of it, providing different arguments that fit each iteration in turn. I think that would be ineffective, since it would allow for the development of even more versions of the theory that escape piecemeal criticism. Instead, I'll offer a more global diagnosis of where the Prop Theory goes wrong. I'll support this diagnosis by arguing that both the Sufficiency Claim and the Necessity Claim are false, which is just to say that every version of Prop Theory is appreciably lacking in intensional adequacy (whether it accepts one of those claims or both).

What I think is going wrong with the approach taken by Prop Theorists is that they've identified a common function of fictions and attempted to apply that to give an analysis of the nature of fiction. They've identified the fact that fictions are standardly used to prompt imaginings and make-believe. Fictions that don't function that way must certainly be rare. However, that we generally use fictions this way doesn't entail that this is essentially what it is to be a fiction. Being a prop to produce imaginings is neither necessary nor sufficient for being a fiction.

I'm not the only philosopher to diagnose the problem with the Prop Theory this way. For example, Friend (2012, 193) suggests that Prop Theorists “mistake features that are merely standard in our current practice for necessary and sufficient conditions that define fictionality for all time.” In giving her argument against the Prop Theory, Friend offers cases where a non-fictional text seems to be aimed at producing imagining. For example, she claims that “[a]nyone who reads Ernest Shackleton’s *South* (1920), an account of his failed expedition to Antarctica, without imagining the terrible odyssey that unfolded after his ship was crushed by ice has simply not engaged properly with the story.” (Friend 2012, 183) *South* is an example of a work of non-fiction the function of which is to produce imagining. There are many other such works. These examples show the Sufficiency Claim is false.¹³

One possible response to this line of argument that Walton (1990) seems at times to entertain is to claim that fiction and non-fiction are compatible categories. (Friend (2012, 205-6) also entertains the possibility.) That is, a work could be both fiction and non-fiction. That way, even if a work was a prop in a game of make-believe, it might also fulfill whatever the criteria for non-fiction are and be both at the same time.¹⁴

The reason it would be wrong to endorse this response is that the concepts of *fiction* and *non-fiction* should illuminate each other. As discussed above, they seem to stand in a contradictory relation. Our fiction-directed practices are different in quality from our non-fiction-directed practices. To deny the exclusivity of the distinction would be to ignore the way in which we treat them as importantly different. Of course, this is not to deny that the distinction may be vague, with borderline cases. It’s to deny that we should posit any cases that are clearly both fiction and non-

¹³This point against the Prop Theorist is raised in Friend (2008, 2011, 2012) and Matravers (2014).

¹⁴This can’t be the case according to the explicit definition of non-fiction Walton (1990, 72) adopts but it’s at least open as a line of response for the Prop Theorist.

fiction. So, while I acknowledge that this is something a Prop Theorist *can* say, I think this kind of response carries with it an obvious and serious cost in explanatory power.

Stock's Prop Theory is designed to resist counterexamples to the Sufficiency Claim like *South*. Stock (2016) gives an account of exactly what kind of imagining is prompted by fictions, a kind she stipulatively refers to as 'F-imagining'. According to Stock, when someone F-imagines that p they "propositionally imagine" that p . One characteristic feature of propositional imagining, according to Stock is that if someone is propositionally imagining that p , they may also believe that p . However, if someone propositionally imagines that p at time t , either they don't believe that p at t , or they do believe that p at t , but they either occurrently or dispositionally conjoin in thought p and some further proposition q , where q is not the content of any belief of theirs at t . The upshot of Stock's account is that while it's possible to propositionally imagine something one believes, one must at least be disposed to connect that imagining to something one doesn't believe.

I won't give a full analysis of Stock's theory, but there are a couple things to say about whether this can save the Sufficiency Claim. Stock rests her support for her account of the interaction between belief and propositional imagination on a dispositionalist account of belief that says "to have a belief is to stand in actual or dispositional relations to certain 'outputs'." (Stock 2016, 147) There are reasons we might doubt this support. First, Stock seems to think that if we favor a dispositionalist account of belief, then we must favor a corresponding account of imagination, one that individuates imaginings by their behavioral outputs. She is looking for "an inevitable difference in behaviors characteristically dispositionally associated with, respectively, believing that p on its own and imagining-whilst-believing that p ." (Stock 2016, 147-8) But why favor a theory of imagination that individuates imaginings according to behavioral outputs, as opposed to, say, content and phenomenology? We

don't typically take imaginings to terminate in behavior the way beliefs do. Consider an alternative account, on which believing that p and believing and imagining that p could have the same behavioral outputs. What would distinguish belief and imagination, on this account, would be a phenomenological rather than a behavioral difference.

Stock doesn't provide us with any reason to favor a dispositionalist account of belief. A priori considerations may tell against it. Similarly, empirical considerations might tell against Stock's portrayal of propositional imagining. Indeed, when reflecting on the possibility under consideration, it seems to me that I can hold an empty cup in my hand and imagine that it's empty even as I believe the same. It seems to me that I can do this without being disposed to conjoin my imagining to any other proposition I don't believe (since the imagining isn't part of any larger game or story). Individual reflection, empirical research, and a priori reflection about the nature of belief and imagination all have the potential to defeat Stock's claim about what's possible (or not) for imagination.

The other point to make here is that even if Stock is right about what's possible for our imagination, she still doesn't seem to have saved the Sufficiency Claim. The Sufficiency Claim says that possessing the function of prompting imaginings is sufficient for something's being a fiction. Counterexamples (like Shackleton's *South*) show that there are non-fictions the function of which is to prompt imaginings. Stock's theory attempts to save the Sufficiency Claim by holding that, actually, it's just the function of prompting *F-imaginings* that's sufficient for something's being a fiction. But then it seems Stock has changed the subject (from imaginings in general to F-imaginings in particular). We might rightly object: the question was not whether prompting F-imaginings was sufficient for being fiction, but whether prompting imaginings (in some more general sense) was. In reducing the scope of what's sufficient for fiction, Stock's theory is weaker than the Sufficiency Claim. This isn't to say that Stock's

theory falls short of being a version of the Prop Theory. It's just to say that she seems (despite her efforts to save it) to have abandoned the Sufficiency Claim. Her Prop Theory requires something more than merely possessing the function of prompting imaginings for something to be a fiction.

Stock is not the only Prop Theorist to go beyond the Sufficiency Claim. As we saw earlier, Currie adds the requirement that a fictional utterance also must not be non-accidentally true. Currie, Stock, and others, despite going beyond the Sufficiency Claim, are committed to the Necessity Claim. To see if there's anything going for the Prop Theory, we should also examine the plausibility of the Necessity Claim, then. I think it's false. We generally do use fictions as props to produce imagining, so it's natural to think that fictions must do this. But it's perfectly coherent to imagine a fiction that doesn't function this way. For example, imagine an author who makes up a story from whole cloth and sells it as a memoir, not unlike James Frey's *A Million Little Pieces* (2003) or Janet Cooke's "Jimmy's World" (1980). The function of the piece, what it's *for*, we can imagine, is to get people to generate beliefs (perhaps in a way that elicits emotional engagement). If folks found out what this author had done, many would (as many people did with both Frey's and Cooke's works) say the author had produced a fiction and passed it off as non-fiction. This is in spite of the fact that the author had no intention for the function of her work to be the prompting of imagining in their audience (at least not in any way distinct from how normal memoirs are used to do this).

The Prop Theorist faces a choice here: either they can say that all memoirs and news articles (and any other kinds that could be used to represent a narrative) have the function of prompting imaginings, regardless of what their author intends, or they can deny that an entirely fabricated work would be fiction. The former response implausibly divorces the function of a work from the intentions of its author. The intentions of an author may not determine the function of their work in every case,

but they are also not functionally irrelevant. This response would render them so. Moreover, if the function of all memoirs, news articles (and any other narrative art kind) is to produce imaginings, then a question remains for the Prop Theorist: what is the difference between a memoir (which is non-fiction) and a fictional novel? The answer can't be the function of producing imaginings, since that's a function they share. The most plausible difference is that the latter is made up. If instead the Prop Theorist denies that works like those under consideration would be fictions, they will have failed the test of intensional adequacy, insisting on the claim of their theory without attending to counterexamples.

To drive home this point, imagine a world in which our social practices were radically different (or, if not our practices, take those of an alien people). In this world, people make up stories for different purposes than we do. The goal of those who make up stories is to present them in as beautiful and surreal a form as possible. For example, the aesthetic effect of the particular words in a novel, the color scheme or use of camera effects in a film, or any other salient formal techniques that could be used to capture the attention of an audience are all considered to be of paramount importance in the fiction-making and story-telling practices of this world. The authors of these fictions don't want their audiences to get distracted by the content of the story. Their goal is to give them an unforgettable aesthetic experience, and the story merely serves as a vehicle for the formal complexity, unity, diversity, and so forth. They aim not for their audiences to imagine the content of their stories, but instead design their stories so that the story's form (the words, the visual fields, etc.) occupies the audience's attention principally. The function of these works, then, is aesthetic engagement with the work's formal features.

The status of the work as a fictional story still matters in this alternate scenario. Authors who work in fiction do so because making the story up frees them from the responsibility of accurate representation. Audiences, knowing that the story is made

up (rather than representing actual events) *expect* a work that engages them through its formal features, and they appreciate it accordingly. The knowledge that a work is fiction directs the audience’s attention and appreciation. What seems true to me is that what it is for a work to *be* a fiction is no different in this imagined world than it is in our own. The difference between their world and ours is that fiction functions differently in each. This scenario demonstrates a possible case in which fictions don’t have the function of producing imaginings, and thus this function is not necessary for something to be a fiction.¹⁵

In sum, I think the Prop Theory identifies something important about our fiction-directed social practices: we use fictions to prompt imaginings. The Prop Theory is able to explain why this is so, but it wields this insight like a sledgehammer. While it can explain one thing about our fiction-directed practices (viz. the connection between fiction and imagination), it doesn’t explain other things. For example, it doesn’t seem able to explain how we really go about classifying fiction and non-fiction, since we’re characteristically concerned with these categories’ connection to truth. Moreover, given that it doesn’t seem to have given us accurate necessary or sufficient conditions for something’s being a fiction, it’s lacking in intensional adequacy.

¹⁵Matravers (2014, chap. 3) also argues extensively for the falsity of the Necessity Claim. His argument begins with the claim that there’s no difference in “what goes on in our heads” when we engage with a work of fiction than when we engage with a work of non-fiction. In each case the mental state in question is some kind of representation or simulation. If that’s the case, then either (a) functioning as a prop to produce imaginings is not necessary to be a fiction or (b) it’s only trivially necessary. Whether (a) or (b) is true depends on what we take imaginings to be. If they’re what goes on in the head whenever anyone interacts with a representation, then a work of fiction *does* have to function as a prop to produce imaginings, but this doesn’t reveal anything about the nature of fiction. Non-fiction represents things too, and so will produce imaginings in that sense. If imaginings are something more substantive, then producing them is not necessary to be a fiction (for reasons like those I’ve given in this section).

3.4 Fiction Is Not a Genre

As an alternative to the Prop Theory, Friend (2012) offers the Genre Theory of fiction, arguing that it allows us to make better sense of how categorizing works as fiction or non-fiction affects our appreciation of works. On the Genre Theory, there are no necessary and sufficient conditions for something's being fiction. The whole project of giving necessary and sufficient conditions is a mistake. We're better off understanding fiction in terms of a cluster of non-necessary properties, taking contextual and historical facts into account.

A genre is “a way of classifying representations that guides appreciation, so that knowledge of the classification plays a role in a work's correct interpretation and evaluation.” (Friend 2012, 181) Membership in a genre is determined “by a cluster of non-essential criteria that include not only features internal to the work... but also facts about the work's origins, in particular the category in which the artist intended the work to be appreciated, or in which the artist's contemporaries would have placed it.” (Friend 2012, 187) A genre, then, according to Friend, is like what Walton (1970) calls a ‘category of art.’ It is a kind of conceptual label under which works can be sorted, such that the sorting can be subject both to the intentions of the authors (in cases where authors choose to assert themselves) and to the features of the work. How a work is sorted subsequently affects how we properly appreciate it.

Membership in a genre generates expectations about the features of a work. Knowing a work belongs to the genre of fiction allows us to appreciate it in terms of its possession (or lack thereof) of features that are standard, contra-standard, or variable for that genre. Standard features of a genre tend to count in favor of a work belonging to that genre. Contra-standard features of a genre tend to count against a work belonging to that genre. Variable features tend to differ between works of a genre and don't count one way or the other for a work's belonging to the genre. The Genre Theory says that while something like fabrication is standard for fiction and

truth-telling or reporting contra-standard for fiction, these are not essential features of a work of fiction, so it's possible to have works of fiction that are not made up, works of non-fiction that are made up, and works that are in other ways aberrant. The features of the work aren't sufficient for determining what genre a work belongs to. The history of the work or its context of origin can play a role in contravening how one would expect a work to be classified, based solely on its features. (Friend 2012, 193) Authorial intentions for a story to belong to a particular genre (or, in the absence of clear such intentions, contextual facts about how the original artistic audience would have classified the work) are especially important for correct classification of a work in many instances.

It's important to distinguish two distinct theses of the Genre Theory.

ANTI-ESSENTIALISM: There are no necessary and sufficient conditions by which one can appropriately characterize the category of *fiction*.

APPRECIATIVE RELATIVISM: The way we appreciate a particular work is relative to expectations generated by historical, contextual, and social facts.

Friend says at the outset of her paper that she's focused on "the nature of fiction and non-fiction and our appreciation of works in these categories." (Friend 2012, 179) That is, her project is concerned both with the nature of fiction and with how we appreciate it. Friend's commitment to Anti-Essentialism relates to the first concern, viz. the nature of fiction. It's her answer to the ubiquitous Prop Theory.

Appreciative Relativism, on the other hand, is Friend's way of addressing our appreciation of works of fiction and non-fiction. Our appreciation of a work is always in response to our expectations of what features the work will have. Our expectations are a function of the classification of the work, though. If we don't know what a work is, we don't have the grounds to expect anything in particular of it. Much of Friend's discussion of genres serves as an in-depth examination and explanation of the psychological facts surrounding categorization and appreciation of works of fiction and non-fiction, culminating with the conclusion

... we do not classify works as fiction or non-fiction based on necessary and sufficient conditions. . . . Instead, as with other genres and categories of art, classification turns on a cluster of non-essential criteria: in particular the possession of standard features. . . , the intention of the author that the work be read in a particular category and the conventions associated with contemporary categorization practices. (Friend 2012, 195)

Friend's view is that our application of the *fiction* concept proceeds according to a cluster of properties, rather than by necessary and sufficient conditions. That categorization in turn affects how we interact with and appreciate the work, since it alters our expectations of the work.

Friend presents Anti-Essentialism and Appreciative Relativism as a theoretical package. It's easy to see, though, that they're completely extricable. The argument for Anti-Essentialism is an abductive argument. The Genre Theory is offered as the best explanation of a number of things, including the failure of the Prop Theory and the truth of Appreciative Relativism. But we might agree that the Prop Theory fails and that Appreciative Relativism is true without agreeing with Friend's conclusion that the Genre Theory best explains these facts. This is precisely my view.

It seems clear to me that Appreciative Relativism is true: whether a work is classified as fiction or non-fiction will affect our appreciation of the work. This should lead us to think there can be good artistic reasons (at least in some circumstances) to take a story that was entirely made up and publish it, market it, and sell it as non-fiction. In making up a story, an author might have artistic goals that are better suited if their audience interacts with the work as non-fiction. For example, an author might (like James Frey did with his *A Million Little Pieces*) write a largely made up story and publish it as memoir. Our imagined author might want to bring about some social or psychological change in the audience that would not plausibly be wrought

without the pretense that their book was non-fiction. There are a number of potential artistic reasons to publicly classify a made-up work as non-fiction.¹⁶

It's clear then, that a made up story might be publicly classified as non-fiction in a way that's justified by the author's artistic goals. Would such a story *be* non-fiction, though? I don't think so. What a work is can diverge from how we classify it (personally or publicly). Moreover, the latter must be explained in terms of the former. That is, facts about how we should classify a work have to be explained by facts about what the work is, not the other way around. The nature of the work plays an important explanatory role in our publishing practices. Both the nature of the work and how it's published likewise play a subsequent role in our application of the concept of *fiction*.

Another significant part of Friend's view already discussed is that the way we apply the concept of *fiction* is according to a cluster of properties, rather than according to necessary and sufficient conditions. That may well be true. But if it is, it's only obviously a truth about our psychology, not about the category of *fiction*. Compare another example that may prove illustrative: it's perfectly consistent with an essentialist theory of the category of race (on which there are necessary and sufficient conditions for belonging to a given race) that we apply the concept of *race* by using a cluster of properties rather than the necessary and sufficient conditions that correspond to the category. This could be so even on a view according to which race is a mind-dependent social kind. Our conceptual machinery needn't be as finely-tuned as the categorial structure of the world.

In the kinds of cases I've considered above, where it seems permissible to classify a work that's been made up as non-fiction, it's natural to think we're identifying

¹⁶Similarly, an author might publish a book as fantasy, as biography, as narrative non-fiction, or under any number of descriptions for a number of different reasons, artistic or pragmatic. Thanks to Tim Juvshik for pointing out to me that publishing fiction as non-fiction (or vice versa) is simply one case of a more general kind of action an author might take for a variety of reasons.

artistic reasons to publish something under a false description. In a case where we publish something that's made-up as non-fiction (or something that isn't made up as fiction) it's because we want people to interact with the work under the mistaken belief or the knowing pretense (or some other propositional attitude) that the work is not what it is. James Frey admitted to initially shopping his *A Million Little Pieces* as a novel (i.e. as fiction), but publishers weren't interested until he "recast" it as a memoir. (Mnookin 2006) It mattered that the book turned out to be fiction. It lost what inspirational power it might've had once people discovered Frey made most of it up. A cursory review of news articles and internet message-boards from the time show that many people were inclined to accuse Frey of attempting to pass a work of fiction as non-fiction, and in this case they were angry because the justification was profit, not art. Many libraries and bookstores (and even the publisher) went to the trouble of reclassifying the work as fiction.

Now, the Genre Theory can allow for cases of works that are non-fiction being classified as fiction. For one thing, it's not committed to the claim that classification as fiction is necessary and sufficient for being fiction. For another, the Genre Theorist can posit that fiction is a genre and still allow for the phenomenon of mis-classification. If a genre is a way of classifying representations that guides appreciation, one could presumably misclassify a particular representation, so that appreciation is misguided.

There is still a problem here for the Genre Theory, though. Friend (2012, 204) says a work is only likely to lack a definite genre when "an author's intention sits uneasily with contemporary expectations." She also mentions cases when "an author's intention or contemporary practices allow for works that do not fit clearly into either category." (Friend 2012, 204) What's suggested by Friend's remarks is that the genre a work belongs to is determined most intimately by the author's intentions and the audience's expectations, not in a way that can be specified by necessary and sufficient conditions perhaps, but determined nonetheless. Return to the kinds of cases I've

been considering, then, in which an author attempts to publish a made-up story as non-fiction. If the author intends for it to be classified as non-fiction, and the audience accepts it as such, there doesn't seem to be any good reason to deny that the work would *be* fiction, according to the Genre Theory. Intuitively, however, the work would not be fiction. The takeaway is that while the Genre Theory is not committed to the idea that every story that is published as fiction automatically counts as fiction, it doesn't rule out the possibility of a fiction that consists entirely of truth-telling, or of a non-fiction that's entirely made up.¹⁷ The classification of the work on the Genre Theory is thus divorced in principle from the features of the work, from its content, and from any other condition one might attempt to place on it besides the will of the artist and the acceptance of the audience. This leads to a clear failure of intentional adequacy.

Here's what seems like another bad consequence of the theory. There's no reason a genre's features can't change over time, since nothing about them is essential to a genre. So, there is a possible history for us, according to the Genre Theory, in which fiction is not a genre in which it's standard to make things up. Likewise, it seems possible according to this theory for non-fiction to become a genre in which it *is* standard to make things up. All I can say about the people in this alternate history is that whatever they mean by 'fiction', it's not what we do. Of course, it might be possible for reference of the word 'fiction' to shift to some other concept than ours so that in this alternate history the later people refer to something different by their use

¹⁷Friend (2012, 193) thinks "[w]e can safely say that a work that lacks any standard features of a category... will be excluded from the category." It's not clear why we can safely say this, though. After all, the history of art from at least the 20th century on is a history of artists attempting to disprove such bold claims as the claim that no work of non-fiction will be completely made up, for example. And given that Friend posits the artist's intentions as playing an especially important role in determining how a work is correctly classified, *especially* in cases where the work doesn't possess a significant number of features standard for the intended genre, it's not at all clear why an artist couldn't publish a work of non-fiction or a work of fiction either of which was completely contra-standard in its features.

of our same term ‘fiction.’ But the Genre Theory seems to entail the possibility not just of reference shift but of a radical change in the polarity of the fiction/non-fiction distinction.

A proponent of the Genre Theory might respond to this objection by insisting that ‘fiction’ is a rigid designator, picking out a classificatory practice indexed to the actual present, so that if our practice were to change in any significant way it would not be fiction that we were sorting works into anymore but some similar seeming genre, call it ‘schmiction.’¹⁸ While I think this response is open to the Genre Theorist, I don’t think it’s very plausible. It makes the meaning of fiction seem more stable than it actually is. Friend herself illustrates this point well:

It was only in the late sixteenth century that historians began to eschew the representations of inner thoughts. . . . We *could* say that historical writing prior to the seventeenth or eighteenth century counts as fiction rather than non-fiction. But surely it is more plausible to say that conventions for writing non-fiction history have changed over time. (Friend 2012, 185)

The standards for classifying fiction and non-fiction have changed over time. If the Genre Theory posits a rigid meaning of ‘fiction’, then it has a harder time accounting for this change in conventions over time. On the other hand the Genre Theory can forgo rigidification, but then it posits a meaning of ‘fiction’ that is *too* subject to change. Of course, part of how to evaluate this version of the Genre Theory depends on what it would take ‘fiction’ to rigidly designate. We need a theory of fiction that posits an essential core so that the nature of fiction can’t change completely over time while still allowing for change in convention. I think the Fabrication Theory does just that. Any version of Genre Theory developed to accommodate more changes in our practices over time threatens to be subject to more counterexamples as well.

¹⁸I don’t think Friend would make this response, since she places importance on the fact that the features associated with fiction and non-fiction seem in fact to have changed throughout history.

The Genre Theory of fiction does purport to have significant explanatory power with respect to our practices of classification and appreciation of works fiction and non-fiction. We should believe the theory only if it best explains these facts about classification and appreciation. One of my worries is that it lacks the explanatory power it lays claim to. It seems like whether a work is fiction or not plays some explanatory role in whether we classify it as such. But the Genre Theory allows for cases where this distinction (between what a work is and how we classify it) seems to collapse. On the other hand, I think the Fabrication Theory can accommodate Appreciative Relativism and offers just as much explanatory power, without the difficult counterexamples.

3.5 Fiction Is a Fabrication

The previous theories each fall short with respect to intensional adequacy. My argument is that the Fabrication Theory can meet both desiderata set out in §2 to a fuller extent. For that reason, and the fact that it seems to be the common sense starting point, I think we ought to prefer it. To get a handle on the Fabrication Theory, we can apply the picture of story identity we settled on in chapter 1, as well as the account of what it is to make up a story.

In chapter 1, I argued that a story S exists iff there is some narrative content c such that c has intentionally been made accessible and S has c . I also argued for the SENSITIVE TYPE ACCOUNT of story identity, according to which stories are abstract types that are individuated by the conditions that must be met to properly token them, but the truth-values of story identity assertions are relative to their context of use. The picture endorsed there, roughly, is that stories are abstract, representational types with narrative content.

Recall the account I gave of authorship (the term I assigned to the activity of making up a story) was one on which authorship was a degreed matter. A story isn't

just totally, wholly made up or totally, wholly not. Some stories represent a mixture of actual and non-actual events. Some impart important truths about actual events by depicting non-actual events. In light of those considerations, this was the account of authorship I settled on:

For any subject S , and any story x , such that S identifies the narrative representation and content (N, c) of x , S makes up x to the degree that

1. c does not non-accidentally resemble some other narrative content c' of a different story; and
2. c does not non-accidentally represent actual events.

The picture is one on which an author makes up a story by identifying the content of the story and at least one canonical way of representing that content, and on which the content of the story must be neither plagiarized nor a report of actual events. To the extent that either of those conditions fails to be met, that counts against the thought that the author is making something up.

We can reframe the account to answer the question not of when someone is making up a story, but rather when a story (taken as a whole work, psychologically complete at the least) is made up or not. Here's that matching account, removing reference to the author.

For any story, x , and any content, c , such that x has c : x is made up to the degree that (1) c does not non-accidentally resemble any other narrative content c' such that c' is had by a distinct story y , and (2) c does not non-accidentally, accurately represent any actual events.

The fact that a story is made up only to a certain degree once again gives us the intuitively correct result that there are stories that are clearly made-up, cases of stories that are clearly not made-up, and borderline cases where it's naturally harder to say.

Representation can be *de dicto*, *de re*, and *de se*. A story can represent (*de dicto*) that there is a girl running in circles, or it can represent (*de re*) Lucy running in circles, and likewise Lucy can tell a story that represents (*de se*) of herself that she

runs. Representation can be accurate or fail to be accurate in each of these respects. A story can represent (as many do) *de dicto* that a person loves another person, for example. A story might on the other hand represent *de re* of a person, Lucy, that she loves her sister, Margaret. There is thus more than one way for a story to fail to be made up according to condition (2). In particular cases, it might be helpful to determine in just which way a story represents actual events non-accidentally, but on the view being presented accurate representation in any way counts against a story being made up if it is non-accidental.

The pair of theorists who have argued for a version of the Fabrication Theory likely won't favor this account of authorship. Deutsch (2000) and Young (2002) each require the existence of a plenitude of possible stories (a fictional plenitude), such that "to make something up is to record attributive descriptions of aspects of the fictional plenitude." Deutsch (2000, 156) While I'm sympathetic to the notion that for every distinct narrative content, there is a possible fictional story with that content, I don't endorse Deutsch's and Young's views. First, a theory of what it is to make up a story doesn't need to posit a fictional plenitude. Such a posit will certainly be controversial in a debate about the ontology of stories. We should prefer a theory that can make sense of the phenomenon in a way that's amenable to the nominalist, the platonist, and the abstract creationist alike. Working merely in terms of a story, its content, and the relations they bear to other stories, content, and events, should do the job.

Second, both Deutsch's and Young's accounts make a particular kind of imaginative process necessary for the production of fiction. In Deutsch's view, a story is made up only if, in generating it, "the author's imagination is not a slave to (our epistemic access to) how things actually are. . . . fiction may be defined as the product of the author's or maker's imaginative activity." (Deutsch 2013, 366) Young (2002, 31) likewise proposes that authors are "making up a story when the only constraints on what they write are self-imposed, or imposed by their imagination." The common

theme in both views is that the activity of making up a story can be understood as a process limited only by the author's imagination. The more constrained a process is, the less likely the result will count as fiction.

I acknowledge that authors who make up their works are typically less constrained, but I don't agree that this is what it is for a work to be made up, nor is it true that they are always only (or even typically only) constrained by their imaginations. First, authors of fiction usually face many pragmatic and sociopolitical constraints. Some authors face more such constraints than others, but this doesn't diminish the status of their works as fiction. Second, as my discussion of unintentionally made up works shows, I think it's possible to make up works without intending to do so. But in cases in which an author unintentionally makes up a work, it will be normal for them to work under constraints common to works of non-fiction. Alternatively, it's possible on my view to fail to make up a work when one intends to (as in the Smith case discussed by Currie 1990). But any such case will fit the description of a made up work given by Deutsch and Young, since the only constraint on Smith's process will be self-imposed. Thus, while I'm in agreement with Deutsch and Young that a fiction is a made up story (the Fabrication Theory), I differ on what making up a story amounts to. Insofar as I've given reason to think a story could unintentionally be made up, that's reason to favor my version of the Fabrication Theory over Deutsch's or Young's.¹⁹

My version of the Fabrication Theory takes its inspiration from Currie (1990). Currie construes fiction in terms of fictive utterances, requiring that they not be non-accidentally true. Where Currie goes wrong, I contend, is in requiring that a fiction be made with a reflexive intention to produce imagining (or make-believe).

¹⁹For those who remain unconvinced, my version could be brought into consonance with theirs by adding two stipulations. First, add the assumption that for every distinct narrative content there exists a distinct possible story. Second, add a third condition onto my account of what it is to make up a story, "... (3) and the production of *S* is not limited by anything but the author of *S*."

My version of the Fabrication Theory is a way of embracing and refining Currie's non-accidentality condition after rejecting the approach of the Prop Theory.

As mentioned above, Currie's theory has been criticized because of his embracing what we can call a "Patchwork View" of fiction, on which works of fiction turn out to be constructed out of both fictional utterances and non-fictional utterances. Matravers (2014, 39), Stock (2017), and Walton (1990, 79) all argue against the Patchwork View. One objection is that we ought to prefer a theory of fiction that can give a perfectly precise answer to the question of when a work is a fiction or not, while the patchwork view will result in many cases where it is difficult to say (or perhaps there is not even an answer to the question) whether a given work is a fiction. As should be clear, I don't think much of this objection. In reality it can be hard to say whether a work is fiction or non-fiction. Authors sometimes purposefully attempt to ride the line between the distinction, making up their story only in part.

The other objection that gets leveled against the Patchwork View is that it entails that factual statements in a work of fiction are interruptions, disrupting the flow of the reader's experience. That is, it's Currie's view that when reading a work of fiction one will often go from reading a fictional utterance to reading a non-fictional utterance and back again. But it's also Currie's view that the proper way to engage with a fictional utterance on the one hand is to make-believe its content (since a fictional utterance is made with a reflexive intention to produce make-believe), while the proper way to engage with non-fictional utterances on the other hand is to believe their content (since a non-fictional utterance is made with a reflexive intention to produce belief). But Currie takes belief and make-belief to be mutually exclusive, given that he analyzes them in terms of reflexive intentions to produce distinct mental states. On Currie's account, then, non-fictional utterances will constitute interruptions in the experience of the reader who is properly engaging with a work of fiction. But of course that's not our actual experience of works of fiction. When an author includes

a factual statement (about when some event occurred in history, say) our mode of engagement with the work doesn't change. There's no experience of interruption elicited by non-fictional additions. Thus, we ought to reject the patchwork view. Call this the "phenomenological objection."

It seems that the phenomenological objection succeeds against Currie's Prop Theory. But this is because of Currie's combination of (a) the Patchwork View with (b) the claim that the proper means of engagement with fictional utterances differs from that of non-fictional utterances. A view that denies either of (a) or (b) escapes the phenomenological objection. Importantly, the Fabrication Theory is not committed to (b). Thus, a version of the Fabrication Theory that accepts the Patchwork View is not thereby susceptible to the phenomenological objection. On the Fabrication Theory, one can distinguish the property of being made up at the level of particular utterances, allowing whole works to be patchworks of fictional and non-fictional utterances. Such a view wouldn't thereby be committed to saying that non-fictional utterances interrupt our experience of a work. This is all to say that the Fabrication Theory could be developed in a way that incorporates the Patchwork View.

I don't favor this way of developing the Fabrication Theory. There's another important objection to the Patchwork View besides the precision and phenomenological objections that has gone undiscussed. The Patchwork View treats stories as wholes composed of utterances, and it takes the fictionality of a story to be a function of the fictionality of the simple utterances that compose it. Objectors to Currie haven't generally been opposed to this account. Stock (2017), for example, explicitly embraces a view on which fictional utterances in a work are "F-units", and the important point of disagreement between her and Currie is whether all or only some of the utterances that compose a fiction must be F-units. There are two reasons we might question this picture, though: (1) not just literary works are fictions (films and comic books are too, for example) so the claim that works are composed of "utterances" doesn't

seem quite right; (2) even if we grant that works are composed of utterances, there doesn't seem to be any simple way in which the fictionality of the work is a function of the fictionality of its parts.

Let me focus for a moment on the second worry, the idea that the fictionality of works is not simply determined by the fictionality of its parts. One problem with the Patchwork View is that it seems to imply that determining whether a work is fictional is a simple matter of stacking up the number of simple utterances of the work that are made up against those that are not. If the right ratio is reached, the work is made up; if not, it isn't. But the question of whether a work is made up or not (and thus of whether it's fictional) doesn't seem to be reducible to a ratio of utterance types. Certain features of a work will count more in favor of the whole's being made up than others. For example, I might write a story that's a mundane report of my day, but include as the conclusion, "But then I awoke to find it was all a dream!" Arguably, the inclusion of this last sentence has an outsized effect on the content of the entire work. Moreover, even the sense in which a simple utterance of a work can be made up or not should be taken as degreed. For example, an allegory about World War II could represent features of D-Day by depicting a playground brawl. Alternatively, a given utterance might represent an actual event but with some non-actual feature substituted in, such as representing that an important historical military and political leader crossed the Delaware but representing it as Napoleon instead of George Washington. In such cases, we should be inclined to say the particular F-unit (sentence, picture, or whatever it is) is made-up, but only to some degree.

There are a number of lessons here. One is that the fictionality of an utterance is not independent of the fictionality of other utterances in the same work (otherwise, adding an "all a dream" sentence would make no difference). Another is that the fictionality of an utterance itself admits of degrees. Additionally, even if we were to suppose that the fictionality of a work was a function of the fictionality of its

utterances, the function would have to be rather complex. For example, it couldn't be that one merely checks all the utterances to see if they're fictional or not and then calculates the ratio of fictional to non-fictional utterances in the work. Rather, it actually seems to be the case that features of the work as a whole—such as its genre, how the author intends it to be received, or its role in a broader corpus of work— affect the representational features of the specific utterances that might be thought to make up the work.

Add to this the first problem mentioned, viz. that some works (e.g. films) are certainly not reducible to atomic utterances and don't seem easily reducible to any other atomic element that would determine the fictionality of the whole. Films are composed of frames, of course, but it's an obviously implausible picture to think that frames are each either fictional or not and that the fictionality of the whole film is determined by some function of the fictionality of its frames.

All of this tells against the Patchwork View, then, and in favor of a kind of holism that allows the fictionality of the work and its parts to be taken as separate but importantly related questions. I've said above that a work is made up to the degree that its content is novel (i.e. un plagiarized and not reporting actual events). Some of a story's content might fail these tests to a higher degree than other portions, though, and whether a story's content does this might have to do in part with features that are primarily features of the work. The "holist" view of fictionality that I'm espousing, then, doesn't deny that the content of portions of the work plays a role in determining the fictionality of the whole work. It merely denies that the fictionality of the work is simply a function of the fictionality of its parts.

It should be clear now that the Fabrication Theory meets the desiderata set out in §2. Stories that are entirely made up are definitely fictions. Stories that are entirely not made up are definitely not fictions. There are some harder cases in between, and this is as it should be. In actuality, we aren't always sure whether a given story is

fiction or non-fiction. The theory seems intensionally adequate, then. Whether a story is made up corresponds with its status as fiction.

The Fabrication Theory also has significant explanatory power. It explains our artistic practices with respect to fiction well. Endorsing the Fabrication Theory of fiction helps explain why people talk about fictions as being made-up and non-fictions as being based on reality or as being true. Intentionally basing a story on reality will be required in normal cases for non-accidentally depicting actual events. If fictions are stories that are (to some large extent) made-up, then it's natural to talk about them that way.

The Fabrication Theory also explains why we treat authors as rarely (if ever) wrong about what is true in their fictional stories. Since the stories are made-up, the authors can't go wrong about what's true in them. It also explains the primary phenomena its competitors are meant to capture. In general, fictions *are* meant to be imagined, but that's incidental to what a fiction is (at least, if I'm right about the Prop Theory). It just happens to be that if you want a story that's good for the imagination, it is usually expedient to be able to make it up. And how we classify stories (as fiction or non-fiction) *does* affect our appreciation and evaluation of them, but we need an account of the nature of fiction and non-fiction to make sense of our classificatory practices. The Fabrication Theory seems to have at least as much (if not more) explanatory power as its competitors, then. Meanwhile, both competitors suffer a failure of intensional adequacy. Unless there is further, compelling reason to reject the Fabrication Theory, then, it ought to be preferred to the prop or genre theories.

Here's one final reason to favor the Fabrication Theory I think worth mentioning. The focus here has been on works of fiction as artistic artifacts. However, the Fabrication Theory explains the way we often extend the use of 'fiction' to apply to non-artistic stories or to non-narrative statements, theories, lies, and other represen-

tational entities that are fabricated or false. A natural thing to think is that ‘fiction’ is a polysemous term with a focal meaning. As discussed in the previous chapter, polysemous terms are those with multiple, closely related meanings. A polysemous term has a focal meaning if and only if each of its meanings is to be understood by way of one, central meaning (the focal meaning of the term).²⁰ The focal meaning of ‘fiction’ is plausibly the meaning that applies to made-up stories. Other uses of ‘fiction’ should be understood in terms of that meaning, which is to say that we use the term ‘fiction’ to apply to other things (things that aren’t stories, say, or that aren’t wholly made-up) when they appropriately resemble the objects that fall under the domain of the focal meaning. We call a lie, a false scientific theory, or an untruth (unintentionally propagated) fictions because of the ways those things resemble central, paradigmatic cases of fiction. The prop theory, at least, can’t give a similar explanation of these sorts of uses of the term ‘fiction’, since it takes fictions to be props meant to stimulate imaginings. But false theories and lies, for example, are not meant to stimulate imaginings. They’re meant to stimulate belief.

3.6 Conclusion

In this chapter, I’ve claimed that we need a theory of fiction that gets right which stories are fictions and which non-fictions while simultaneously illuminating as much as possible about our practices surrounding those categories. We want to both understand the nature of fiction and non-fiction and to gain an understanding that makes sense of how we interact with those entities.

I surveyed two theories of fiction that seem to fall short of doing both those. The Prop Theory, which says that fictions are narrative works with the function of serving as a prop to produce imaginings, fails to get right which stories are fiction and which

²⁰Cf. McDaniel (2009, 293-295) for a discussion of polysemy, focal meaning, and related phenomena.

non-fiction. It mistakes a common feature of our fictional stories for their essence. The Genre Theory, which says that fiction and non-fiction are genres of which individual works are instances, also fails to get right which stories are fiction and which non-fiction, at least in principle. The Anti-Essentialism writ into the Genre Theory leaves open the possibility of bizarre instances of each category, or gerrymandered categorial shifts.

The Fabrication Theory of fiction is a natural candidate given how ordinary people talk and reason about fiction. It seems to get the right answers about which stories are fictional and which aren't. It explains a lot about our use of the term 'fiction,' our attitudes towards authors of fiction, the fact that fictions are in general meant to be imagined, and the way that knowing a story is fiction affects our appreciation of it. The Fabrication Theory seems to do better on these counts than its extant rivals. Given that it's a natural starting point, we ought to prefer it. Sometimes the true answers to philosophical questions turn out to be stranger than fiction, but in this case the truth is that a fiction is made up story.

CHAPTER 4

HOW NOT TO ARGUE FOR CREATIONISM ABOUT STORIES

4.1 Introduction

To this point we've considered what authors are doing when they make up stories, what it is for a story to be complete, and what the distinction between fictional and non-fictional stories consists in. It's time to turn our attention to the question of what a story is. To do so, first I want to focus on the question of whether, in authoring a story, an author creates it.

Consider some work of fiction, say Spenser's *The Faerie Queene*. *The Faerie Queene* presents us with a puzzle. First, *The Faerie Queene* has some features that theorists say point in favor of its being an eternally existing, immutable abstract object. For example, *The Faerie Queene* is repeatable: there are many copies of it, all of which are instances of it. Moreover, each copy of the story allows us an opportunity to experience the story in full. Let me elaborate a bit on how each of these features have been presented in the relevant literature.

Repeatability: Dodd (2007) and Kivy (1983, 1987) focus on the repeatability of stories. Repeatability is something theorists generally claim can only be a feature of abstract objects like types or universals, which they then argue must be eternal and immutable.

Experiencability: Dodd (2007) emphasizes the audibility of musical works as evidence of their being eternal and immutable, since certain versions of that view (e.g. the view that musical works are types of which their performances are tokens) seem to make better sense of talk about audibility than certain versions of the view that musical works are historical individuals (e.g. the musical perdurantism explored in Caplan and

Matheson 2006). One can point to the fact that a story is capable of being experienced in full in its instances as a similar sort of evidence: it's better explained by the view that stories are eternal and immutable than it is by some versions of the view that stories are created and historical.

On the other hand, *The Faerie Queene* has several features that many contemporary ontologists would say point in favor of its being a created, historical entity. *The Faerie Queene* originated with Spenser's authorial activities. It was susceptible to change after Spenser authored it, according to Spenser's decree. *The Faerie Queene* could have gone differently than it does. It was authored by a particular person in a particular literary-historical context. Finally, *The Faerie Queene* is meant to be understood by an audience. These ideas have also been developed in recent literature.

Temporality: Works come into and go out of existence. Rohrbaugh (2003) stresses this point. His claim is that a work's existence depends on there being what he calls "embodiments" of it, objects (and events, I take it, for artworks that can be embodied by being performed like stories or plays) that ground the facts about what the work itself is like. An embodiment of a particular story is anything that "preserves what it is like and leads to new [instances]." (Rohrbaugh 2003, 191) Rohrbaugh's discussion of embodiments arises in the context of talking about photographs. Dodd (2007, 107-108) considers its application to musical ontology.

Temporal Flexibility: Works can undergo change. Friedell (2020b) discusses this issue extensively, including why it is that in most cases only a work's author (and not I, for example) can change the work. We want to explain both the ways works can change and the ways they can't change.

Modal Flexibility: Works could have had other properties than they do. Rohrbaugh (2003) also emphasizes this point as one reason we should take some artworks to be historical individuals rather than eternal types.

Essentiality of Author and Origin: Works' authors are essential to them. That is, for any story, it is impossible for that story to exist while having been authored by someone other than the person (or people) who in fact authored it. Levinson (1990) and Fisher (1991) endorse and emphasize this claim with regard to stories. Works' contexts of origin are also essential to them. That is, for any story, it is impossible for that story to exist while having been authored in a different literary-historical context or by a different author. Levinson especially emphasizes this point by the use of modal arguments and the invocation of Leibniz's Law. I'll

address the essentiality of authors and contexts of origin separately below.

Teleology: Works exist in order to be understood. This point belongs to Morris (2007), who argues that since works have *teloi*—they are for something—it must be the case that they are created.

The fact that *The Faerie Queene* is repeatable and experienceable tells in favor of its being eternal and immutable, but it has several features that tell in favor of its being created and historical. It can't be both. So which is it, and what features should we prioritize when theorizing about stories?

One way to resolve the puzzle is to take only some of the relevant features of *The Faerie Queene* to be important (or of primary importance) for theorizing about the nature of the work. This sort of resolution seems fairly popular amongst ontologists of art. For example, Dodd (2007) takes the repeatability and audibility of musical works to be of primary importance. A large portion of his book is spent explaining away other apparent features of musical works. Levinson (1990) considers primarily repeatability, temporality, and essentiality of composer and origin. Rohrbaugh (2003) takes temporality, temporal flexibility, and modal flexibility as the primary features in his argument. In the context of arguing that fictional characters and fictional works (rather than stories) are created abstract artifacts, Thomasson (1999) relies heavily on the temporality of the works to make the case.

It's not an uncommon dialectical strategy, then for an ontologist of art to make a case for some position about what stories (or some other repeatable artwork) must be by directing their focus entirely or primarily on a subset of the set of (arguably) ontologically significant features of the work. One might not think this strategy misguided. After all, if one can construct a sound deductive argument that will resolve the puzzle with which we started using only some of the features of stories, that would seem to be a success.

And lo and behold, for each feature F that tells in favor of stories being created and historical, one *can* construct an argument (and in some cases theorists have done so, for stories or for some other repeatable works) that runs as follows:

1. For any story s , s has feature F only if s is created.
2. For any story s , s has feature F .
3. Therefore, for any story s , s is created.

The fact that there are multiple such features builds a seemingly strong cumulative case in favor of the idea that stories are created, historical objects. Call this view “Creationism.”

Creationism as it’s normally understood is really the combination of (at least) three theses:

1. Stories exist.
2. Stories are objects.¹
3. Stories are created by their authors.²

¹There’s an obvious way we sometimes use the term “object” on which anything that exists is an object. Everything that’s anything is a thing, so to speak. But sometimes we use the term “object” to distinguish a more limited class of things in our ontology in order to contrast them with things that lack a certain particularity or thisness. When I say Creationists hold that stories are objects, I use the term “object” in this second way. Creationists typically hold that stories are particular things, not properties, events, states of affairs, norms, or other things we might think of as non-objects when drawing the sort of contrast just mentioned. See Rettler and Bailey (2017) for more discussion on the category of object and more ways of drawing object/non-object distinctions.

²When I speak of works being created by authors here, I mean to use ‘author’ in the broad sense on which anyone who makes up an artistic story is an author. Folk tales, pulp fiction, and soap operas count as authored stories just like the works of Shakespeare (even in instances of works that are co-written by cultural groups, or by collections of writers, artists, and producers). I’ll stipulate here that in order to create something, one must act in such a way as to directly bring that thing into being, and not merely as a logical or ontological consequence of something else one brings into being. I want to rule out this kind of case: Lucy directly brings into being a statue of a ladybug, and as a consequence there comes to be the singleton set of the the ladybug statue. Lucy created the ladybug statue, but she didn’t create the singleton of the ladybug statue. I don’t want to rule out the possibility of unintended creation, however. It isn’t necessary that one intend to bring an

Creationism is a view both about the nature of stories and of the artistic activity of authoring. It takes authoring to be a kind of creation, and it takes stories to be particular things that authors directly bring into being.

On the other hand, one might think that when someone authors a story, they thereby discover some new object.³ Call the view that by composing music one discovers some new object(s) “Discoverism.” Discoverism, as I’ve just defined it, is not incompatible with Creationism. It’s possible, for all we know, that authors both discover objects and create objects while authoring stories. (Indeed, Levinson (1990) apparently holds this view of composers.) Call a view “Platonism” if it entails that Discoverism is true and Creationism is false.⁴ According to Platonism, in the course

object into existence to create that object, but objects that merely follow logically or ontologically (rather than etiologically or in some other more substantive way) from one’s actions won’t count as being created in the sense in which I’m talking about creation. In the next chapter, I’ll address this understanding of creation and those that diverge from it, providing reasons to prefer this way of understanding creation.

The Creationist holds that when someone authors a story, they thereby create some new object, and that object is the story. There are varying possible species of Creationism. Some Creationists like Levinson (1990), Thomasson (1999), Evnine (2016), and Friedell (2020b) take created works to be abstract. There are differences amongst their views about the ontological complexity of the works. Other Creationists, like Caplan and Matheson (2006, 2008) take created works to be concrete.

³I’ll say that to discover some object o is to become acquainted with o when (1) one was not acquainted with o at any prior time, (2) o ’s existence predates one’s acquaintance with o , and (3) one’s acquaintance with o is not causally explained by the fact that anyone else is acquainted with o .

⁴This is, of course, a stipulative definition of ‘Platonism’, one that contrasts with, for example, Kania (2017, §2.1), who defines Platonism as “the view that [stories] are abstract objects,” and contrasts Simple Platonism (“[stories] are eternal existents, existing in neither space nor time”) with Complex Platonism (“[stories] come to exist in time as the result of human action”). While on my definitional schema, Levinson comes out as a Discoverist but not a Platonist, on Kania’s, Levinson comes out as a Complex Platonist, not a Simple Platonist. These are clearly merely terminological differences, but it’s important to be clear on them. This importance comes to light, for example, when we notice that Kania claims that Platonism “respects more of our pre-theoretic intuitions about stories than any of the other theories.” (Kania 2017, 2.1) This claim may seem to be especially germane to the topic at hand, but since Kania’s notion of Platonism doesn’t rule out Creationism, the relevant contrast class of theories doesn’t include all and only the Creationist theories. His point seems to be that identifying stories with abstracta respects more of our pre-theoretic intuitions. It’s not clear exactly why Kania makes this claim in absence of a discussion of what kind of abstracta (e.g. created or eternal) he has in mind, though. Thanks to an anonymous reviewer for helping me to identify the differences between my discussion and Kania’s and to relate my terminology to Kania’s.

of composing a story, an author discovers some new object that is the story, and they do not create any object that is the story.⁵ Platonism agrees with Creationism about the first two of its theses:

1. Stories exist.
2. Stories are objects.

Platonism endorses a different third thesis, however:

4. Stories are discovered by their authors.

Since discovery requires the pre-existence of the discovered object, it is impossible for one object to be discovered and created in the same activity.⁶ Thus, the Creationist and Platonist positions differ over what occurs in the activity of authoring. The Platonist identifies stories with eternally existing, immutable abstract objects, suitable candidates for discovery by authors. The puzzle that began this section can best be understood then as a trial between Creationism and Platonism. Some features of stories seem to entail that Platonism is true. Some seem to entail that Creationism is true. Platonism and Creationism can't both be true, so something has to give.

My argument in this chapter is that the cumulative case for Creationism does not settle the matter decisively. Every case of putative evidence in favor of Creationism can be explained by the Platonist. In what follows, I parse the case for Creationism into its component arguments and show how the Platonist can easily respond to each. In §§2-6, I show how the Platonist can respond to the Creationist's arguments derived

⁵It's worth noting that Creationism and Platonism aren't the only possible views to hold in this neighborhood. They are complex ontological theses that hold contradictory answers to questions about what stories are and what authors are doing. However, they are natural views to hold and the most popular views represented in the literature, so they are the only views I discuss here.

⁶There are certainly senses of the word 'discover' according to which discovery does not require the pre-existence of an object. I'm using the term stiputively, as mentioned above, to single out a sense that does require pre-existence of the discovered object. The dialectic between the Platonist and Creationist could be reformulated without this stipulation, but it helps simplify things.

from each feature of stories that are supposed to tell in favor of Creationism. Each section discusses a particular feature and the argument for Creationism stemming from that feature, as well as the way(s) a Platonist can respond. I conclude by discussing lessons to be learned. The main lesson is that the stalemate between Creationism and Platonism cannot be broken merely by focusing on a few features of stories and attempting to argue for one view or the other by way of a deductive argument that makes use only of those features. Some methodological clarity is needed, and progress is not likely to be so neat.

4.2 Temporality

Rohrbaugh (2003) argues that just as paintings come into existence at a time (viz. with the painter's act of painting), so do photographs.⁷ Likewise, just as a painting can't be identified with any particular collection of matter (since it can degenerate and be restored), so a photograph can't be identified with any particular bits of matter. Rohrbaugh takes it that part of the temporality of a photograph is that it depends for its existence on the existence of embodiments of it. Embodiments, according to Rohrbaugh, are "those physical objects [and presumably events] which ground the facts about what [a work of art] is like." (Rohrbaugh 2003, 191) An embodiment of a work needn't be an occurrence of it, but it does need to encode information about what the work (or occurrences of it) must be like. To be an embodiment of a work, something must preserve what the work is like and lead to the work's perpetuation. Rohrbaugh's suggestion is that the painting case and the photograph case are analogous because the constitution relation and the embodiment relation are both species of the generic relation of ontological dependence. So paintings (which are

⁷Levinson (1990, 2012) stresses creation when it comes to musical works. Thomasson (1999) gives similar arguments addressing fictional characters and works of fiction. See Rossberg (2012) for a discussion of the destruction of artworks that deals with music but is focused more on computer art.

constituted by their matter) and photographs (which are embodied by various objects) both are art objects that ontologically depend on other more mundane objects.

Rohrbaugh's argument translates easily to the case of stories. It seems like stories originate with the authorial activity of their authors, and it seems like they depend on their embodiments (objects and events such as books, performances, recordings, and memories). Stories are temporal in just that sense. But, the Creationist will tell us, stories are temporal only if they're created. So, stories must be created.⁸

The Platonist needn't accept this argument. They must deny the premise that stories are temporal. They identify stories with eternally existing, immutable abstract objects. Those are not the kinds of things that come into existence or that depend for their existence on mundane, concrete objects or events.⁹ However, the Platonist can explain why it seems that stories are temporal. To do so, the Platonist makes a distinction between the intrinsic and extrinsic properties of a story. The intrinsic properties are the properties a story has "in itself," those that any qualitative duplicate of the work must possess. The extrinsic properties a story has are those the work has that are not intrinsic, properties it has in virtue of the way the work is related to other things.¹⁰ Intrinsic properties of my shirt, for example, include its *being-green* and *having-mass*. Extrinsic properties of my shirt include *being-owned* and *being-worn*.

⁸Rohrbaugh himself doesn't press the point about creation. He's interested in arguing for the thesis that certain artworks are historical individuals, particular objects that come into existence at a point in time, depend for their existence on other objects or events, and can cease to exist. It's easy to see how his argument can be extended to the case for Creationism, however, and the idea that stories are historical in Rohrbaugh's sense is still inconsistent with Platonism about stories. For an extension of Rohrbaugh's thesis about artworks as historical individuals, see Magnus (2012).

⁹By asserting that eternal objects don't come into existence, all I mean to rule out is that they have a beginning. I take it that both temporal and eternal objects exist in time, but eternal objects exist at every time while temporal objects do not (having a beginning and possibly an ending).

¹⁰For more discussion of the intrinsic/extrinsic distinction, see Lewis (1983a, 1986).

On the Platonist's view, a work cannot come into existence or depend ontologically on mundane objects or events. However, it's natural for the Platonist to acknowledge that prior to an author's act of authoring, the story was not a story. This is to say that the kind *story* is a phase sortal. A phase sortal is any kind (like *parent*, *professor*, or *town hall*) something can come to fall under and later fail to fall under without ceasing to exist. Phase sortal kinds are accidental to the things that have them, as opposed to essential. Now, some phase sortals require an intrinsic change in an object for the object to fall under them. For example, everyone undergoes some intrinsic changes as a physical necessity in order to become an adult, even though *adult* is a phase sortal, a kind that isn't essential to any person. But not every phase sortal requires this kind of change. Some are such that an object can fall under them without any intrinsic change in itself, for example the kinds *uncle* or *secret crush*. The Platonist should posit that the kind *story* is just like that. Whatever a story is fundamentally (whether type, universal, or something else besides) it comes to fall under the kind *story* in virtue of relations it bears to things distinct from it, and its being a story depends on the existence of embodiments of it. If there are no more embodiments, then the story ceases to be a story (even though it continues existing). In this way, the Platonist can make good sense of our intuitions about temporality without acknowledging any of the ontological theses of the Creationist.¹¹

¹¹This view of stories is endorsed by (at least) Wolterstorff (1980, 88). In discussing temporal flexibility, Dodd (2007, 91) agrees that the type/token version of Platonism "has no choice but to say that Bruckner's [unfinished] Ninth Symphony could not have differed with respect to its intrinsic properties: had he done what we would describe as 'completing the work', he would, in fact, have authored a work distinct from the work that exists actually." It seems that Dodd's view is like the view described, then, at least insofar as Dodd denies the temporality of stories. However, it's not clear from this passage whether Dodd feels compelled to say (like Wolterstorff) that *story* is a phase sortal.

4.3 Temporal Flexibility

Friedell (2020b) argues for a particular answer to the question, “Why can’t I change Bruckner’s *Eighth Symphony*.” The answer, according to Friedell, is that while musical works are created abstract objects with no parts, they *have* sound structures. *Having-a-sound-structure* is an extrinsic property of a musical work, and our contingent social practices make it so that only a work’s composer can determine what sound structure a work has. If our social practices were radically different, though, then I would be able to change Bruckner’s *Eighth Symphony* (i.e. I and anyone else would be able to determine what sound structure it has). Friedell thus gives a satisfying account not only of the fact that musical works can change, but that not just anyone can change them. This is what their temporal flexibility consists in.

While Friedell’s interest lies explicitly only with musical works, his thesis can be adapted easily to deal with the alleged temporal flexibility of stories. It seems that stories can change, but that not just anyone can change them. One way to account for this is with Friedell’s strategy: stories are created abstract objects with no parts, but they *have* plots (or word-structures, or whatever it is about a story we might consider an author changing). *Having-a-plot* is an extrinsic property of a story, and our contingent social practices make it so that only a story’s author can determine what plot it has. If our social practices were radically different, though, then I would be able to change Spenser’s *The Faerie Queene*.

Friedell’s account highlights an important feature of social practices. Social practices constrain what’s socially possible. Given that it’s a social practice in your group to drive on the right side of the road or to grant ownership of property to certain people (and not others) who have satisfied certain conditions, these practices constrain where one can drive and what kind of claims one can feasibly make on property. Likewise, if Friedell is right, our social practices surrounding music constrain facts

about who can work on what stories. We can change the social facts, though, to change what's socially possible. People can work to change practices about driving or property. They can change practices about how people dress or when they eat. It's open to us to change the practices surrounding music or stories, too, so that it would be possible for anyone to change what features a given musical work or story has.

What should the Platonist say about temporal flexibility? The Platonist says stories are eternal, immutable abstract objects. Stories' immutability just consists in them not undergoing any intrinsic changes. Clearly, the Platonist must deny that stories change intrinsically. Most versions of Platonism can't quite say the same thing as Friedell, either. This is because Platonists about stories will tend to identify them with something like a linguistic or pictorial type. The thought is natural: if you think authors are discovering something, you probably think what they're discovering is either a kind of collection of words or pictures that express the content of the story, or with the content itself. So any normal Platonist can't say (like Friedell) that a work's undergoing change is its coming to have a new linguistic structure or its coming to have new content, things it has only extrinsically. Platonism about stories will likely make stories identical either to their linguistic or pictorial structure or to their content, or to make it the case that *having-a-linguistic-structure* or *having-a-content* is intrinsic to the work.

The ideal strategy for the Platonist is instead to provide truth-conditions for claims people make about artwork change. The key observation for the Platonist is that when an someone authors a story, they establish certain social practices with respect to the abstract object they discover. By creatively discovering some abstract object, in combination with the powers granted them by the social practices of their artistic community, they make it the case that when anyone wants to perform, listen to, watch, read, evaluate, or otherwise interact with the work, they must interact in

those ways with an embodiment of the object identified by the author. They might read a book that is an instance of the work. They might watch an instance of it on a streaming service on their television. They might attend a performance of the work. They might instead listen to a recording of the performance they attended. But the fact that it is a social practice for members of an artistic community to interact with these embodiments of the abstract object in a certain way is an extrinsic feature of that object, something brought about by the author, in concert with the pre-existing social practices of the community.

What the Platonist should say in cases of apparent change is that it's not the work but actually the social practice that's changing. For example, if someone claims that Spenser revised *The Faerie Queene*, the Platonist should say this is true iff Spenser altered the social practice he had established. Instead of requiring people to interact with embodiments of the original abstract object, *o*, the social practice now required them to interact with embodiments of a different abstract object, *p* (different manuscripts, published works, and recordings thereof) in order to interact with his work. The change is not an intrinsic change in any work (i.e. not in any eternal type), but an extrinsic change in which object we refer to by means of the term "Spenser's *The Faerie Queene*" (first *o* then *p*) and in what our beliefs and practices centered on Spenser's work are about. On Platonism, stories don't undergo intrinsic change. However, *something* undergoes intrinsic change that makes it right to say the works change. In particular, what changes are the social practices instituted by the author. Spenser originally instituted a social practice of interacting with embodiments of one abstract object. But in "changing the work" he changed the social practice so that it required one to do these things with embodiments of a different object instead.

An obvious drawback to the Platonist's account of artwork change is the one it wears on its face: it says, "stories don't actually change; it's just right to say they do sometimes." One might think Friedell's Creationist picture is on better footing here,

since it seems to allow stories actually to change. It's not clear this is right, though. For one thing, Friedell's Creationism only allows story to undergo extrinsic change, and the Platonist allows this too. The Platonist says works go from being works to not being works and that their aesthetic properties change based on their relations to artistic communities, but they understand these changes as purely extrinsic.

Some might not think Friedell's Creationism even captures what we are after when we think stories are temporally flexible. After all, it denies that stories change in themselves. Rohrbaugh (2003) and Eynine (2009, 2016) stand apart from Friedell here. Their versions of Creationism do allow stories to undergo intrinsic change and so obviously best account for the feature of temporal flexibility if it's meant to be the ability of stories to undergo intrinsic change. Friedell (2020b, 12) points out though that Rohrbaugh and Eynine are left with a mystery when attempting to explain why only some people (and not others) can change a work. To show this, Friedell points out that social practices do not prevent anyone from changing the intrinsic properties of artworks like paintings or sculptures. The mystery, on Rohrbaugh and Eynine's views, is why it should be possible for some people and not others to change the intrinsic properties of stories. Both Friedell and the Platonist have a way of making sense of change for stories without having to face up to that mystery.

Platonism actually appears to be even better off than Friedell's Creationism here though, since the latter says all stories are intrinsically just the same. While Friedell's Creationism can say that stories change, it can't say that any of them is intrinsically any different from any other. In itself, Spenser's *The Faerie Queene* is just the same as E.L. James's *Fifty Shades of Grey*. That's a surprising result, and the Platonist can avoid it, since they can identify the works with linguistic structures (or at least allow having such a structure to be an intrinsic feature of a work). The Platonist can say a lot of different things about how eternal, immutable abstract objects are intrinsically, and most versions of Platonism will make stories come out as intrinsically distinct.

So Platonism is on at least equal footing with Creationism here, if not actually better off.

4.4 Modal Flexibility

To say that a work is modally flexible is to say that it could have had different properties than it actually does. Rohrbaugh (2003, 181) points out that what is at stake is *de re* modality. People think of a work like *The Faerie Queene* that *it* could have been different, not just that some sentence featuring the term ‘*The Faerie Queene*’ could have been true that isn’t actually true.

It seems that *The Faerie Queene* could have had different properties than it actually does. But if *The Faerie Queene* is an eternally existing, immutable abstract object, then *The Faerie Queene* couldn’t have had different intrinsic properties than it actually does. Say, for example, *The Faerie Queene* is an abstract linguistic structure. Then the identity conditions for *The Faerie Queene* are determined by what a particular collection of words must be like in order to be a token of *The Faerie Queene* (or what a particular collection of words must be like in order to be an ideal token, if *The Faerie Queene* is what Dodd (2007) calls a ‘norm-type’). Consider another possible world *w* in which we suppose that normal (or ideal) instances of *The Faerie Queene* are phrased differently than they actually are. Then the work that’s called ‘*The Faerie Queene*’ in *w* is not identical to *The Faerie Queene* in the actual world, since properly formed tokens of *The Faerie Queene* in *w* differ from properly formed tokens of *The Faerie Queene* in the actual world. What we’re imagining in considering world *w* is a world where Spenser discovered a different object and called it ‘*The Faerie Queene*.’ In such a scenario, it wouldn’t have been the case that *The Faerie*

Queene had different properties than it actually does. It just would have been that Spenser discovered some other object and we called *that* thing *The Faerie Queene*.¹²

The Platonist has at least a couple options for what to say about the purported modal flexibility of works. Which option they take will likely depend on their broader commitments in the metaphysics of modality. The first option is to agree with the paragraph above and deny that stories are modally flexible. What's really true, the Platonist might say, is the broader possibility that things in general could have been different than they are. For example, we can imagine a possible world *w* in which everything is just as it actually is, with the exception that the poem called '*The Faerie Queene*' is slightly intrinsically different from how *The Faerie Queene* actually is. In *w*, Canto I of Book I begins with the line "A Gentle Knight was riding ore the grass" (as opposed to the actual "A Gentle Knight was pricking on the plaine"), completing the rhyme by rhyming 'grass' with 'mass'. People may then mistakenly (and understandably) move from that imagined scenario to the claim that *The Faerie Queene* could have had different intrinsic properties than it actually does. But it's not clear what else about our experience of the work would make an evidential difference between the thesis that things could have been different in general and the thesis that *The Faerie Queene* could have been different in itself. The Platonist has no need to abandon their theory since the fact that it seems like works could be different is just as well explained by their theory as by the Creationist's. The claim that stories are modally flexible seems like just one interpretation of what we're conceiving of in conceiving of a world with a different version of *The Faerie Queene*.

¹²This argument is very similar to the main argument of Hazlett (2012). He gives a condensed version as follows,

[I]f there are repeatable artworks, they are abstract objects; no abstract object has any accidental intrinsic properties; would-be repeatable artworks have at least one accidental intrinsic property; therefore, there are no repeatable artworks. (Hazlett 2012, 162)

The other option for the Platonist is to embrace a counterpart theoretic analysis of *de re* modal properties of stories. On counterpart theory, *de re* modal properties are analyzed in terms of a counterpart relation, a relation of similarity. Simply put, *The Faerie Queene* could have been *F* iff there exists a possible object *a*, such that *a* is a counterpart of *The Faerie Queene* and *a* is *F*. What possible objects are counterparts of *The Faerie Queene* depends in large part on how intrinsically similar those objects are to *The Faerie Queene* and in part on the context of evaluation (since a normal feature of counterpart theory has it that the counterpart relation is context-relative).¹³ A Platonist who adopts a counterpart theoretic approach to *de re* properties of stories can agree that *The Faerie Queene* really could have been intrinsically different. This is just because *The Faerie Queene* is very similar to another possible work that is intrinsically different. Once the Platonist has taken this approach, they no longer have to agree with the Creationist's claim that a work is temporally flexible only if it's created.

4.5 Essentiality of author

Some theorists take it that a particular story could only have been authored by the person(s) who actually authored it. What's supposed to tell in favor of this thesis is the fact that, for example, were John Donne to have authored *The Faerie Queene* it would have certain aesthetic or artistic properties it doesn't actually have. Compare Fisher (1991, 135), who argues against the idea that musical works are discovered:

in the typical case of composition there is nothing to be discovered that is independent of the process of composition, filled, as it is, with the composer's taste, values, interests, competence, favorite motives, etc. . . . These are "discoveries" that only make sense within each particular artist's body of work. Hence they are not discoveries of pre-existing independent objects that others could have discovered.

¹³For more on counterpart theory, see Lewis (1968, 1971, 1973, 1986).

The claim that a musical work only “makes sense” within a particular artist’s body of work seems to be the claim that it has certain important aesthetic and artistic properties in virtue of its being composed by that artist. An analogous argument could be made with respect to stories: Fisher or others might be inclined to claim that a story only makes sense within a particular artist’s body of work, which is just to say that it has certain aesthetic and artistic properties that it wouldn’t have were it to have been authored by someone else. Works that differ significantly in their aesthetic and artistic properties must be distinct. But the claim here is that what aesthetic and artistic properties a work has depend in part on who authors it.

The Creationist argument in this case really has two steps. The first is to establish the essentiality of authors to works. The Creationist does this by imagining a possible scenario in which two different authors author two intrinsically linguistically identical works, *a* and *b*. The Creationist insists that *a* and *b* will vary in some of their aesthetic and artistic properties in virtue of their having been authored by different authors. By Leibniz’s Law, *a* and *b* are distinct, since they have different properties. Since the example was arbitrary, it seems to generalize that intrinsic linguistic duplicates with distinct authors are necessarily distinct. Thus, a work’s author is essential to it. The second step of the argument is to establish the created status of stories. stories have their authors essentially only if they’re created, the Creationist will say. After all, if a story is an eternally existing, immutable abstract object then it couldn’t have any essential tie to a contingently existing being like an author. Since stories do have their authors essentially, it can be concluded that stories are created.

Once again, the Platonist can say at least a couple different things in response. The first option for the Platonist is to protest that the Creationist’s argument from Leibniz’s Law is unsound. It depends upon the premise that possible intrinsic linguistic duplicates authored by distinct people must have distinct aesthetic and artistic properties. But this begs the question against the Platonist. After all, on (at least

some versions of) Platonism, if possible work *a* is an intrinsic linguistic duplicate of possible work *b*, then *a* is identical to *b*. In that case *a* has all of the aesthetic and artistic properties *b* has. What the Creationist is forgetting, the Platonist will point out, is that a single work can be authored by more than one author (however unlikely that is in actuality). In the event that this occurs, the one single work will have a number of aesthetic and artistic properties we'll only be inclined to attribute to it in one intensional context or another (e.g. thought of as *a*). But that's a matter of what's going on in our heads, not a matter of how the work is in itself.

The Creationist might push back here by pointing out that in acknowledging that a work could have different aesthetic and artistic properties were it to be (or to have been) authored by another author, the Platonist seems to be committed to both the temporal and modal flexibility of stories. After all, if works can acquire new aesthetic and artistic properties along with the advent of being newly authored, then it's possible for any given work either at another time or another possible world that someone else authors it and it acquires new such properties. But this is not actually a problem for the Platonist. The aesthetic and artistic properties in question are clear cases of extrinsic properties, and the Platonist has already been ready to acknowledge that a work can change extrinsically (or that it could have been different extrinsically). *Of course* a story possesses (at least some of) the aesthetic and artistic properties it does only by being part of an author's larger oeuvre. (For example, Langston Hughes' "Let America Be America Again" is rightfully read as a mournful but hopeful plea at least in part because of Hughes' Blackness and his personal history struggling as a poet, playwright, and novelist. Taking another example from the world of music, it's only because of Stravinsky's prior work composing *The Firebird* and *Petrushka* that White (1961, 59) can say that in *The Rite of Spring* Stravinsky "pushed [his technique] to its logical conclusion." Likewise, James Joyce's *Ulysses* is hailed as one of the most important works of modernist literature in part because Joyce employed

literary techniques foregrounding the process of thinking that no author had ever used prior, not even Joyce, and to great effect.) That only goes to show that those properties are properties the work has by its relation to things outside it, namely to other works (of the author's or of other authors). But then such properties are clearly extrinsic, so that even the Creationist shouldn't hold that they're intrinsic, and it's no problem for the Platonist to allow works to undergo change with respect to such properties.

The second option to respond to the case for Creationism here is open for the Platonist who endorses counterpart theory for *de re* modal properties of stories. The counterpart theorist has no need to deny the essentiality of authors to works. On their view, it's a contextual matter what *de re* modal properties an object has. If our modal intuitions strongly favor the importance of authors to works, then it will come out as true that necessarily, Spenser authored *The Faerie Queene*, since in that context all of *The Faerie Queene*'s counterparts will have been authored by Spenser (or by intrinsic duplicates, or counterparts of Spenser). Thus, the counterpart theorist can agree with the Creationist that stories have their authors essentially, but deny the Creationist's premise that this is so only if stories are created.

4.6 Essentiality of Origin

Similar to the argument for essentiality of author, Levinson (1990) argues that the musico-historical context in which a work was composed is essential to it.¹⁴

... a musical work, stripped of its contextual coordinates in musico-historical space—thus yielding, roughly, a pure sound sequence—is incapable of bearing many of the determinate aesthetic properties that we ascribe to it... the way to show this is to observe that for all we know there *are*, and at any rate easily *could be*, works containing (incorporating) *identical sound sequences* and yet presenting nontrivial aesthetic

¹⁴For a more recent and more circumspect discussion of the role of historical origins in the individuation of a story, see Moore (2012).

differences; the conclusion is inescapable that such works, if they truly differ aesthetically *cannot* be identified with the lone sound sequence itself. (Levinson 1990, 222)

The style of argument is the same as in the author/composer case. Levinson considers an alternate possibility in which two people compose works with the same sound sequence in different contexts. But their works would have different properties were they composed in different contexts. By Leibniz's Law, the two works are distinct and we must conclude that the context is essential to the work. Once again, the Creationist will want to apply this argument to stories, saying that this can be the case only if the story is created, and thus conclude that stories are created.

It's worth pointing out that Levinson's point in this section is more idiosyncratic than Fisher's. He hopes to show not just that musical works (or stories) are created, but that a pure sound sequence is "incapable of bearing many of the determinate aesthetic properties that we ascribe to" a musical work. The analogous argument with respect to stories would be that an abstract linguistic or pictorial structure is incapable of bearing many of the aesthetic properties we ascribe to stories. It's unclear how Levinson's thought experiment is supposed to support that claim, however.

To see this, examine the situation from the Platonist's point of view. On their view, a story is perhaps an abstract linguistic structure, or some narrative content. In being authored, the work gains some properties it didn't previously have, but the Platonist will understand these properties as extrinsic to the work, properties it has in virtue of the relations it bears to its author, the historical context of its origin, the community of people in which it is appreciated, and so forth. Of course some content or a linguistic structure that hasn't been authored would lack some of the aesthetic properties it has in virtue of being authored, on the Platonist's view. But the aesthetic properties it gains on being authored are all extrinsic to the work itself, thinks the Platonist. Levinson's thought experiment doesn't seem to support at all the idea that a pure sound structure (or a linguistic structure, or some content)

can't bear the aesthetic properties we attribute to stories. Rather, it supports the idea that a story would be appreciably aesthetically different were it authored in a different context. The Platonist has a story to tell about this.

Returning to the case for Creationism, a similar reply to the case of the essentiality of author is in order. The Platonist can object that Levinson's modal argument with the invocation of Leibniz's Law is unsound. The Platonist's way of understanding the imagined scenario is such that the two authors have each authored the same work. Perhaps it has some of its aesthetic properties in virtue of one act of authoring and others in virtue of the other act, but all the aesthetic properties in question belong to the one work.

The Platonist who endorses counterpart theory for *de re* modal properties of stories can give a similar response here as well. The context of origin *is* essential to a story. This is given the usual counterpart theoretic analysis: all of a story's counterparts were authored in intrinsically identical or counterpart contexts of origin. The Creationist's claim that context of origin is essential only if a work is created can be rejected again.

4.7 Teleology

Morris (2007) argues that musical works are teleological entities, by which he means that they have the property of being *in-order-to-be-understood*. Morris's argument first proceeds from the observation that musical works are meaningful. Indeed, it seems that necessarily, all artworks are meaningful. Morris's argument ought to apply as much to stories as it does to musical works, then. Morris (2007, 60) then suggests a necessary condition on the property of being meaningful: "Something is meaningful [only] if it is, in some sense, *there to be understood*." The idea is that an object that's meaningful has a telos; it's "for" something, and "being understood is what it is for." (Morris 2007, 60)

According to Morris, though, something can exhibit this kind of teleology only if it is created. “Nothing can have that specific teleological purpose [of being meaningful, and therefore requiring understanding] without having been created.” (Morris 2007, 64) If that’s right, then stories must be created.

The Platonist will deny the premise that stories have this telos (and are meaningful) only if they’re created. Instead, the Platonist will adopt the view that what matters for an object’s being meaningful in this sense is its standing in the right intentional relation to someone. Someone might create an object and imbue it with meaning from the get go, but another salient possibility is that someone selects a pre-existing object and presents it in such a way as to imbue it with the same kind of meaning.

Morris (2007, 67) argues that Platonism cannot take this route, since he has “argued that works of art are not only meaningful, but *essentially* meaningful. We could not have the very same things without their being meaningful. Consequently, we could not have had the very same things before the intervention of an artist.” Morris’s argument that works of art are essentially meaningful relies on the claim that the other features of an artwork we don’t think could vary all seem to be tied to its meaning. Some of us don’t think a particular novel could have had entirely different words. Likewise, some of us don’t think a particular symphony could have had totally different themes and instrumentation. Morris thinks the explanation for these intuitions is that in the supposed scenarios the meaning of those works wouldn’t be retained. A natural Platonist explanation of the modal intuitions Morris discusses is that we don’t think a story could have a different meaning and be the same *artwork*, even if it would still be the same object. On Platonism, an author selects a pre-existing object and imbues it with meaning. If it didn’t have that meaning, the Platonist can very well admit it wouldn’t be the very artwork it is, but that’s because *artwork* is

a phase sortal, not an essential kind. Morris doesn't provide a reason to favor his explanation of these intuitions over a Platonist explanation.

Here's a case to press the point against Morris. Imagine an alien civilization whose language, Formic, consists of the practice of taking found objects in their environment and presenting them, just as they are, for various amounts of time, in various relations to each other in space in order to create units of meaning. The lexicon of the language then just consists of the domain of physical objects available to be picked up by the speakers of the language. An object's color, shape, and size can all make contributions to the overall meaning of a sentence, as do the ways in which speakers of Formic arrange the objects they use to author sentences in time and space. Since the objects Formic-speakers pick up are the basic units of communication in the language, it seems that they have meaning. They are to be understood. However, it's stipulated in the thought experiment that those objects are found and not created. This seems like a counterexample to Morris's claim that for something to be meaningful it must be created.

Morris anticipates this kind of case and says "so-called 'ready-mades' are not properly so-called; but that is hardly a counter-intuitive claim: it is quite natural to think that taking an already-made object and placing it in a gallery creates a work of art which did not exist before." (Morris 2007, 67) Morris insists on his principle then. In the case of Formic speakers he would reject the stipulation and posit that what we're imagining is a group that routinely creates new objects that didn't exist prior. Presumably, these objects will share their spatial location with the mundane objects the Formic speakers select. This posit seems ad hoc, though. There's no obvious reason to dig in one's heels and refuse to imagine the scenario as described.

It's worth noting that Morris's claim about what it's natural to think in the case of ready-mades is contested by other theorists writing on the subject. Evnine (2013) considers the case of ready-mades in detail. When asking just the question at hand

(viz. whether it's right to say that a ready-made is an object created by an artist out of another object or whether it is a pre-existing object merely selected by an artist), he says the following:

The point of ready-mades is precisely that they confront us with these two options between which there is no definitive basis for decision. Is the artist effecting an imposture, passing off as a sculpture something that is not one? Is he, in the terms I used above, simply using a urinal as a work of art although it is not one? Or is he demonstrating something about the nature of sculpture and about the creative powers of the artist, who can indeed make something new merely by selecting another, already existing object? (Evnine 2013, 420)

Evnine's point is not to deny that there is a fact of the matter as to whether ready-mades are created or mere appropriated objects. Rather, the point is that the purpose of historical ready-mades is to raise just that question. It shouldn't be obvious to anyone, then, that a ready-made is a creation. Moreover, ruling out the possibility of their being appropriated objects does not look like an advantage of a view.

Morris gives a further theory of under just what conditions a work of art is meaningful, according to which, "[i]t is only when an artist intentionally exploits those independent properties of the medium [viz. those properties of the medium that the artist uses to make the meaning of the work] that we have something which is meaningful at all." (Morris 2007, 70) It's not at all clear why the Platonist can't agree with this claim, however. The Platonist thinks that artists intentionally exploit the properties of their media to communicate meaning through pre-existing works. It isn't at all clear why an object has to exist prior to the kind of activity Morris is describing for the activity to be possible.

4.8 Conclusion

The case for Creationism might be sound. It presents us with a number of features of stories, and for each feature there is a deductive argument the conclusion of which is that stories are created and historical. It's clear that this kind of case is a dialectical

failure, though. None of the deductive arguments the Creationist can construct give the Platonist a compelling reason to accept their view. Indeed, more than one version of Platonism has an easy time responding to the case for Creationism. For every individual argument the Creationist constructs, the Platonist can either justifiably deny that stories have the feature in question and provide an explanation for why it seems that way, or the Platonist can agree that stories have the feature in question and deny the inference to creation. There may be an associated theoretical cost for the Platonist of denying that stories have some of the features associated with the Case for Creationism, but the Platonist presumably has broader reasons for their Platonism that they see as worth paying the cost. In at least some cases (e.g. temporal flexibility), Platonism is arguably *better* situated to explain the feature in question than the best Creationist candidate. The point is that the Case for Creationism does not (and cannot) settle the question of whether stories are created or eternal.

This does not by itself establish an upper hand for Platonism in the ontology of art. The Creationist of course has things to say about how to account for the repeatability of stories, and will not be moved by a corresponding case for Platonism. The stalemate leaves us with the puzzle with which we began: some features of stories seem to entail that they're eternal and immutable; some seem to entail that they're created and historical. Both can't be true. How do we resolve the puzzle? Some philosophers have continued to press points on one side or the other of the first-order debate. A number of papers in Mag Uidhir (2012) (many of which are cited here) point to the repeatability, temporality, temporal flexibility, modal flexibility, or origins of works as important evidence for solving the question of what they are. Even more recently there have been a number of new attempts to explore Creationist theories of artworks and at least one attempt at breaking new Platonist ground.¹⁵

¹⁵Much of the work here is in the ontology of music. For new work on developing and defending Creationism about musical works, see Aliyev (2017), Cray and Matheson (2017), Evnine (2016),

What this chapter should have shown, however, is that so long as any such attempt fits the mold of a deductive argument that takes a feature of stories and attempts to use that feature to show that they are either created or eternal, the attempt is unlikely to resolve the debate.

A fair amount of discussion in ontology of art these days has been about methodology and metaontology. A number of authors have recently debated the extent to which how artworks seem to us must be a guiding feature of our theorizing in the ontology of art.¹⁶ Rohrbaugh (2012) calls the principle that one’s methodology in theorizing about the ontology of art must be constrained by the practice of art “Ontological Pragmatism.” The question at issue amongst these (and other) theorists is the nature, scope, and validity of this constraint. For example, when approaching the question of what stories are, just which apparent features of stories (or practices surrounding stories) must constrain our theorizing? Should we prioritize the beliefs and practices of experts (e.g. authors, musicians, and music critics) over others, or should we take a more democratic approach? If we adopt a strong pragmatist attitude, will this lead to incoherence (as the puzzle with which we began seems perhaps to suggest) and perhaps even to nihilism about certain art kinds? Without some further resolution on these questions we are liable to end up in dialectical stalemate. If the constraint is too strong, however, the puzzle with which we began threatens to escape resolution.

In order for theorists to make progress as a community, then, it’s likely that some measure of agreement about methodology must first be achieved. Progress could come about in different ways. If theorists agreed on a particular principle that settled the question of which features of stories (or other artworks) were important for theorizing

Friedell (2017, 2020b), and Moruzzi (2018). For a new Platonist theory of musical works, see Letts (2018).

¹⁶ For example, see Davies (2009), Dodd (2013), Irvin (2008), Kania (2008), Predelli (2009), Rohrbaugh (2012), Stecker (2009), and Thomasson (2006).

about them, that would help settle the standoff. On the other hand, theorists might settle on a metaontology that allowed one to resolve standoffs by weighing broader theoretical concerns.

In the next chapter, I'll take up recent discussions about the metaontology of art and show how they point to a different approach to settling this debate. The temptation is to appeal to a metaontology that reads the answer off of our practices. I'll argue instead for a metaontological approach that allows for a more radical theory, but one that best explains our experience of the works in question.

CHAPTER 5

HOW TO ARGUE FOR PLATONISM ABOUT STORIES

5.1 Introduction

Progress in ontology of art is a tough row to hoe. The problem is not a lack of suggestions about what various art-kinds are. The problem is, rather, an overproliferation of conflicting views. As we saw in the previous chapter, conflicting accounts arise naturally out of the varying features we tend to attribute to stories given our artistic practices of authorship, appreciation, criticism and so forth. Stories seem to present us with a metaphysical puzzle. Some of their features tell in favor of their being created and historical, while others tell in favor of their being eternal and immutable.

Some theorists have tried to resolve that puzzle with what I called the case for Creationism, but that case turns out to be a dialectical failure. Every case of putative evidence in favor of Creationism can be explained by the Platonist. In each case, the Platonist either denies the claim that stories have the feature the Creationist attributes to them, accounting for appearances with an error theory, or they deny that stories have to be created to have the feature in question. Either way, the case for Creationism doesn't budge the Platonist from their perch. Of course, an analogous case for Platonism shouldn't budge the Creationist either.

The question is how we can make any kind of progress in discussions about the ontology of stories (or of art in general). It's tempting to think that debate can coalesce around agreement if we first agree on some metaontological ground rules. The thought is that there is a difference in quality between the case for Creationism

and the case for Platonism that's revealed by settling on principles for adjudicating between the two.

In what follows, I first examine the main consensus about what our metaontology of art should be, viz. the idea that our artistic practices should constrain what we say in our theories of art ontology. There are different ways to understand this constraint. I show what the important distinctions in different versions of the constraint are and then argue for a version of the constraint that allows for substantive discoveries in the ontology of art. This "Discovery View" is opposed to those who argue that the ontology of art must (for any of a number of reasons) not "go beyond" our practices. I defend the Discovery View from the objection that it can't make sense of how our art-kind terms (or names for artworks) successfully refer. With the Discovery View vindicated, I turn to a new kind of case for Platonism. Instead of presenting a deductive argument that turns on the claim that stories must be eternal, abstract objects given some feature they have, I present a two-stage argument in favor of Platonism. First, I argue that Platonism best explains a number of features of stories already discussed in the previous chapters. Second, I argue that Creationism is committed to an implausible account of how some abstract objects come into being. Given these two facts, we should on balance prefer Platonism to Creationism. As will be seen, this way of arguing for Platonism escapes the pitfalls of the deductive cases examined in the previous chapter. Instead of pretending the question of creation can be settled simply, I show how to argue for Platonism by way of an abductive case that compares the theoretical advantages of the two views.

5.2 Vindicating Discovery

Thomasson (2005) issues a call to metaontology of art as a way to resolve stalemates like that between the Creationist and the Platonist. It's a tempting thought that in this debate, some methodological ground rules are what will help resolve the

disagreement. If we can agree about what exactly an ontological theory must take into account (and how), then we should be able to achieve some convergence in at least the terms of the debate, if not in the outcome. While new suggestions about the ontology of art have continued to proliferate, Thomasson's call has successfully stimulated discussion of the proper method for art ontology.

Agreement in this discussion has coalesced around the principle that, in some way or to some degree, our artistic practices constrain what we should say in our theories of art ontology. Following Davies (2004), let's call this principle the Pragmatic Constraint.

Pragmatic Constraint: Our artistic practices constrain what we should say in our theories of art ontology.

The Pragmatic Constraint says that our artistic practices place some kind of limit on what we can say about the metaphysics of art. Different theorists articulate and employ the Pragmatic Constraint differently, however. One difference can arise when we consider the kinds of practices in question. Some theorists distinguish between what we say or believe about artworks from how we act towards them, for example. We can also distinguish between different kinds of practices, for example authorial, critical, evaluative, or consumptive practices. These distinctions have a downstream effect on what a version of the Pragmatic Constraint prescribes for our ontological theorizing.

Likewise (and probably more importantly) there are differences amongst theorists over how exactly practices constrain what we should say. As a way to introduce the differences here, let's distinguish (using broad strokes) between what has been called "Descriptivism" and what's been called the "Discovery View".¹

Descriptivism: The correct ontology of art must make explicit the conception of works inherent in our practices; an ontology of art must reflect the conception inherent in our practices.

¹See, e.g. Thomasson (2005), Dodd (2012, 2013), and Davies (2017).

Discovery View: The correct ontology of art must explain our artistic practices; an ontology of art may fail to reflect (to some extent) the conception inherent in our practices.

The Descriptivist, generally speaking, says ontologies of art have to posit that artworks have the properties we act like they do. In effect, Descriptivism requires that we “read off” our ontology from our practices. The Discovery View, on the other hand, just says an ontology of art has to give an explanation of our practices. That is, it must make sense of our practices, given what it says the art objects in question are. Another way to put this: On the Discovery View, what’s important is that an ontology of art make it likely (to some degree) that we would have the artistic practices we do, given what the ontology says about what the art objects in question are. I take this to be a natural way of understanding explanation that we use in a variety of contexts: E explains H (to some degree) just in case H is more likely if E is true than if E is false. I won’t say exactly how we should judge the strength of explanations like this; that’s beyond the scope of this work. The basic idea here is that what the Discovery View is looking to do is give an account of what art objects are that makes it very likely that we would have the practices we do (more likely than it would be if the ontology in question were false).

But one way to do this is to give an error theory of some aspect(s) of our practices, a theory on which we are in some way mistaken to speak about or treat the objects in question as we do, but on which there is a natural explanation for why we speak or act this way. It would be quite a tall order, even on the Discovery View, to argue that we’re *entirely* mistaken to treat some artworks of a given kind as we do. But the Discovery View leaves open the possibility that we can find we were wrong in some ways about what some art kinds are really like.

There are different ways of developing these views, of course. It should be obvious that Descriptivism, as stated, needs supplementation. We can’t just read off our ontology from our practices. One lesson from the previous chapter was that our prac-

tices sometimes seem to imply that stories are eternal and immutable, and sometimes seem to imply that stories are created and historical. If we just read an ontology off of that puzzle, we'll end up with a contradictory ontology.

Two noted descriptivists have developed the view differently. I'll start by discussing Davies' version of Descriptivism and then move on to critiquing Thomasson's. In part to avoid worries about inconsistent practices, Davies (2004, 2017) holds that the practices that constrain our ontological theorizing are the critical and appreciative practices *that we would hold upon reflection*.²

A proposed ontology of art must cohere with a theoretical representation of the norms that *should* govern the judgments that critics make concerning works. . . . A defense of any such theoretical representation of "right artistic practice," or of any unwillingness to revise our actual practice in the light of such a representation, will necessarily appeal, implicitly or explicitly, to conceptions of the goals of our critical practice, the values sought in such practice, and the values offered by the works that such practice takes as its objects. It is in the context of such debates, where we weigh the merits of alternative theoretical reconstructions of our actual practice and its objects, that ontological proposals find themselves accountable to artistic practice. (Davies 2004, 20)

The rational reconstruction or theoretical representation of right artistic practice that constrains ontology of art according to Davies must be arrived at in a spirit consistent with that of our actual practice. It would be illicit, for example, for one to come to the broader ontological conclusion that a one-category ontology is to be preferred, and then to apply that conclusion by construing all artworks as being of the same ontological category, independent of any considerations about our practices. Rather, the kinds of considerations sanctioned by Davies for reconstructing our artistic practices are those that come from within the practice itself, considerations that invoke

²Davies may not be quick to embrace the descriptivist label. This doesn't strike me as particularly important. It will become clear shortly the ways in which Davies' particular views fit the characterization of Descriptivism offered above. Moreover, Davies (2017) doesn't fight the label when discussing Dodd's characterization of him as a "local descriptivist", i.e. someone who advocates a descriptivist approach to art ontology but not to ontology in general. It's not clear to me why Davies limits the Pragmatic Constraint to our critical and appreciative practices.

the goals or values that are the currency of the practice. Those considerations are given pride of place in determining what right artistic practice must be, and right artistic practice is what constrains ontological theorizing about art.

Davies' rationale for adopting this version of Descriptivism is that anything short of this methodology amounts to changing the subject.

our philosophical interest in “art” and in “artworks” is grounded precisely in [artistic] practice. It is because certain features of that practice puzzle us, or because the entities that enter into that practice fascinate us, that we are driven to philosophical reflection about art in the first place. To offer an “ontology of art” not subject to the pragmatic constraint would be to change the subject. . . there is no alternative but to start from critical reflection on our actual artistic practice because the very notions of “art” and “artwork” are parasitic upon that practice—artworks just are the things that play a particular kind of role in a particular kind of practice. (Davies 2004, 21)

Davies' point is twofold. First, he points out that our interest in art-ontological questions begins with, and is rooted in artistic practice. Given that our point of connection to artworks is via our artistic practices, those practices provide the only natural starting point. His second claim is stronger, though. To go beyond or revise the conception of artworks inherent in this starting point would be to “change the subject.” In Davies' view, we establish the subject of art-ontological inquiry by reference to our practices of appreciating and criticizing certain things. Artworks are just whatever are the things we're criticizing and appreciating. Right artistic practice (i.e. the artistic practices we would hold upon reflection) thus establishes a kind of theoretical role, and Davies' Descriptivism prescribes that ontology of art posit something that exactly fits the role in question.

It's not at all obvious that Davies' version of Descriptivism is inconsistent with the Discovery View. In fact, in his most recent defense and articulation of his view, Davies (2017) argues that they are consistent. The apparent inconsistency between Descriptivism and the Discovery View (in their more general versions) arises from the question of whether an ontology of art may falsify the conception of art inherent in

our practices. Descriptivism says it may not, while the Discovery View says it may. The inconsistency is resolved in this case, however, when we realize the equivocation occurring between Davies' use of 'practices' and the way that a proponent of the Discovery View (e.g. Dodd) uses the same term. While Davies doesn't allow an ontology of art to falsify the conception of art inherent in our *reflective* practices, Dodd allows an ontology of art to falsify the conception of art inherent in our *actual* practices. Since actual practices may not be (and likely are usually not) reflective practices, there's no problem here. While Davies and Dodd may differ over to what extent they allow an ontology to diverge from actual practices, they don't disagree in principle about whether an ontology of art may falsify the conception inherent in our actual practices.

Where Dodd and Davies *do* seem to disagree is over the question of how many of our practices must be taken into consideration before ontological theorizing begins. Davies' complaint with Dodd is that Dodd's case for musical Platonism depends (much in the way the case for Creationism above did) on taking just some features of musical works and building the case for Platonism solely from those features. Instead of taking the whole of our practices into account to arrive at his view, Dodd favors the aspects of our practices that seem to indicate that musical works are audible and repeatable, infers that musical works must be abstract types, and explains away the features of musical works that tell in favor of their being created and historical after the type-theory has been established as the default.³ Davies insists that the Pragmatic Constraint requires us to take the whole of musical practice into account at the start of the enterprise.

It's hard to see how much of a difference this will make, however. The explanation for the form of Dodd's argument is almost certainly that Dodd finds Platonist meta-

³Cf. the critique in Davies (2009) of the overall argument found in Dodd (2007).

physics in general more compelling and so gives a greater weight to certain practices (viz. those that emphasize repeatability or audibility for musical works) over others. Thus, even if Dodd were to take all of our practices into account before building his ontology of musical works, it's unlikely that he would end up with any different view. Unless Davies also insists that one must weight all aspects of every practice equally in applying the pragmatic constraint, it's not clear that taking all the practices into account will make a difference for someone with strong metaphysical inclinations towards one or another view already rooted.

There is good reason for Davies not to insist on this, however. If Davies were to insist upon such an equal weighting, that very insistence would seem to leave his version of Descriptivism unable to resolve cases in which our practices seem to be inconsistent. We *need* to be able to bring in considerations that help us to emphasize one aspect of practice over another in cases where the practice is puzzling. Davies' own view of artworks, on which all artworks are performances, seems to be a result of a process on which some of our practices have been weighted more heavily than others. There's no reason, from Davies' own perspective, then, to insist on an argument that takes all of our practices into account before giving an account of a certain art-kind.

Thomasson (2005, 2006, 2007, 2010, 2015) has developed a deflationary version of Descriptivism that's much more clearly at odds with the Discovery View. On Thomasson's view, which follows from her general "Easy Approach" to ontological questions, ontology of art is properly accomplished by determining via conceptual analysis what the application and co-application conditions for our artkind terms are and checking to see whether those conditions are fulfilled. Application conditions for a term establish the existence conditions for the kind in question. Co-application conditions establish the identity conditions for the kind. Conceptual analysis that uncovers these constitutive semantic rules for our art-kind terms unveils the structure

of the concept in question, and objects can't rightly be said to have existence or identity conditions distinct from those built into the concept.

I won't weigh in on Thomasson's broader meta-ontological project here. Instead, I'll focus on her case that Easy Ontology is the correct brand of Descriptivism for ontology of art. Thomasson's main argument stems from what has been called the qua problem. The qua problem (so termed by Devitt and Sterelny (1987)) is the problem that when establishing the reference of a name or kind term, there are inevitably many candidate referents in the neighborhood. It seems that some disambiguating concept must be used to pick out one candidate from the many. But then the disambiguating concept plays an important role in linking art-kind terms and names of art objects to the objects they pick out. According to Thomasson (2005, 223), disambiguating concepts "determine the ontological kind, if any, picked out by the term, and so the ontology of the work of art must be something we learn about through conceptual analysis of the associated concepts." In order to find out about what works of art are, then, Thomasson argues that we need merely to analyze our concepts of the works, the concepts that must have been used to ground the terms in the first place.

Dodd (2013) points to a problem with this argument. We can acknowledge, with Thomasson, that a disambiguating concept is required to pick out one candidate from the many. However, this acknowledgement by itself does nothing to settle *how* the disambiguating concept does the picking. Thomasson holds that the disambiguating concept picks out a candidate by establishing which modal or ontological features the things in question can truly be said to have. Thus, an ontology of a particular art-kind cannot truly posit that that kind has an ontological profile distinct from that given by the application and co-application conditions constitutive of the disambiguating concept. But Dodd argues that the theory of reference espoused by Evans (1985) allows for a different way of resolving the qua problem, one that doesn't require a kind to perfectly match the ontological profile provided by the disambiguating concept.

On Evans' theory, "the item determined as the referent of a singular term is that thing that is the *source*—or, more precisely, the *dominant* source—of the relevant body of information." (Dodd 2012, 84) Thus, it's not that the disambiguating concept provides an ontological profile something must meet to be the object (or kind of object) in question. Rather, disambiguation happens by means of causal influence: the term refers to the object that's the dominant source of information associated with it. On the view Dodd articulates, a disambiguating concept is required to establish reference, but there's no reason to think the disambiguating concept must function as an exact guide to the ontological profile of the kind in question. Thus, as the Discovery View holds, an ontology of art may falsify (to some extent) the conception inherent in our practices.

Thomasson (forthcoming) notes what seems to be a clear problem with applying Evans' theory of reference to the qua problem:

Evans' suggestion was not designed to be a solution to the *qua* problem, and could not succeed as a solution to the problem at issue here... if we are asking whether our musical work terms (say) refer to pure structures, indicated structures, abstract artifacts, etc., the candidate referents differ only in their *modal* properties, not their empirical and causal properties—they are not separate sources of causal information.

The problem here is not the familiar question of whether abstracta can be causal. Thomasson's worry can be elucidated by first supposing there to be such varied ontological kinds as pure structures, indicated structures, and abstract artifacts. Now, ask whether "Beethoven's *Eroica* symphony" refers to a pure structure, an indicated structure, or an abstract artifact? Thomasson's suggestion is that (granting our supposition about the existence of these three kinds) there will be three different objects that are candidates for the name "*Eroica*" and there's no reason to think any one of the candidates has played a distinct causal role from any of the others. Just as in the case of the statue and lump (supposing there to be two colocated objects),

causal efficacy by itself doesn't seem to be enough to distinguish abstracta that are tokened in all of the same instances.

Dodd and Thomasson both discuss a potential response: one might appeal to salience rather than to causal dominance. Instead of worrying about what entities cause what, we might just focus on how our interests determine what the referents of our terms are, with our interests often being in what objects are the sources of the information associated with the disambiguating concept. Along this line, Dodd suggests that given "given the kinds of creatures we are and the kinds of interests we have, objects of a certain kind will, so to speak, leap out at us from a sample and attract our interest more readily than others; and it is items of this more salient kind that will tend to be the dominant source of the information associated with a term when its reference is being grounded." (Dodd 2012, 93) Now, if "dominant source" here is understood in terms of causation, Dodd's suggestion is still susceptible to the worry under consideration: abstracta (at least certain groups of them) don't seem to be susceptible to being distinguished merely in terms of being dominant causal sources of information.

An alternate way of understanding the talk about being a "source" of information is in terms of explanation. On this understanding, we need an explanation for why our disambiguating concept has the content it does and not some other content. One object can be a better explanation of that kind of conceptual fact than another even in a case where there's no difference in their causal influence on our concepts. The relevant difference is instead located in the way the features of one object align with our interests more than the features of another, so that the one object is more salient. For example, two mathematical structures may both be adequately described by the content of a certain person's concept of *set*. Given that person's interests in using sets, however, it may be that the structure described by ZFC (for example) is more salient to them than that described by naive set theory. In that case, on the view

in question, we could posit that their term ‘set’ refers to the structure described by ZFC, because it’s the structure described by ZFC that’s the dominant *explanatory* source of the information associated with their term ‘set’. And this holds not because of any causal story but because of a more general story about the way the features of the structure in question align with the person’s interests.

It’s important to note that causal explanation is a kind of explanation, so that the kinds of cases that were the focus of Evans’ original theory should still count as dominant sources of information on this theory. What switching from causation to a broader notion of explanation allows is a way of making sense of how something can be a dominant source of information in lieu of making any causal difference to our conceptual content. The answer is that even without causing us to have the conceptual content we do, some items are more salient than others in virtue of the way their features align with our interests.

Thomasson (forthcoming) argues that since the idea of salience comes from perceptual fields, and since there is no difference in perceptual properties between things like different set-theoretic structures or other abstracta, “it isn’t clear how we can understand this appeal to salience...in any way other than as metaphorical.” She allows that we may say it is in virtue of our interests that our term refers to what it does, but responds by asking, “How do we come to discover our interests, though? Presumably by appealing to a careful study of our practices: to what we care about when we use the relevant term... But this brings us back to the descriptivist camp—thinking that the best way of telling what ontological and modal features the referents of our art-kind terms (and names of works of art) have is to study the practices we are engaged in.”

What Thomasson doesn’t seem to allow for is the idea just discussed: it may be that our interests make it the case that a term refers to what it does, but not because our interests set out an inviolable description of what the thing must be. Rather, it

may be (as Dodd seems to have in mind) that in virtue of our interests a term refers to one kind of object over another simply because that kind better fits our interests, *even though our concept doesn't perfectly describe it*. Take for example, the earlier case of *set*. Imagine a world in which (similarly to the previously described case) there is a community of people for whom the structure described by ZFC is most salient to them given their interests in using sets and so forth, so that the structure in question is the dominant explanatory source of their concept of *set*. However, suppose that it's part of this community's concept of *set* that sets are semi-divine beings, worthy of worship for their work of "holding" objects. We can take this to be an example of a case where the community's interests (how they want to use sets or set theory) make it the case that their term ('set') refers to what it does, but not because their interests set out an inviolable description of what the thing must be. Sets aren't divine, after all. Rather, it's in virtue of their interests that their term refers to the structure it does rather than, say, that described by naive set theory, because the structure described by ZFC better fits their interests, even though their concept doesn't perfectly describe it.

Once one has accepted the broader, explanation-based version of Evans' theory, there's no need to accept Thomasson's solution to the qua problem. So, even though our interests *are* important in determining the referents of our terms, they aren't important in the way Thomasson says (by establishing rules of use for our terms, which in turn establish the modal features of the objects to which those terms refer). In placing emphasis on our interests, we may, as Thomasson asserts, have "arrived back at the descriptivist camp", but we certainly aren't in the Deflationist's tent. That is, contrary to Thomasson's claim, appealing to our interests to make sense of the meaning of our art-kind terms does not automatically bring us to a deflationary Descriptivism. If we do arrive back at Descriptivism, I take it to be the sort endorsed by Davies that is not principally at odds with the Discovery View.

What's most important to take away from this discussion of the qua problem is not that the proposed solution is the right one, but that it's at least one way to respond to the qua problem that doesn't rely on the notion of a dominant causal contributor. That shows that the argument from the qua problem to Easy Ontology is resistible.

If we don't have good reason to accept the deflationary version of Descriptivism represented by Thomasson's Easy Approach to ontology and there doesn't seem to be a substantive difference between Davies' Descriptivism and the Discovery View, then there doesn't seem to be any reason not to continue to utilize the Discovery View in our approach to the metaontology of art. Our practices constrain art ontology insofar as the correct ontology of art must offer an explanation of why we have the practices we do, but the explanation needn't be purely descriptive. It may falsify to some extent the conception inherent in the practices. Part of the justification for adopting such an ontology may come from an appeal to aspects of the practices of themselves, but it may also involve appeals to broader considerations, e.g. the theoretical utility of positing the entities in question or the instability of competing views. It's to such a case for Platonism that we'll now turn.

5.3 In Favor of Platonist Ontology

If the Discovery View has been vindicated, there's nothing illicit about making a case for Platonism that proceeds from some important features of our practices and shows how Platonism functions as the best theoretical explanation of those features. The kind of case I'll develop here is importantly different than the kind of case considered in the previous chapter, however. I won't be offering a deductive argument for the conclusion that Platonism is true. That is, I won't be claiming that if our practices have the features in question, then Platonism must be true. Instead, I'll offer an abductive case for Platonism: it's the best theoretical explanation of a number of

important features of stories. Given that it can make sense of the features Creationism can—as was the upshot of the previous chapter—Platonism should be preferred to Creationism on balance.

This abductive case for Platonism about stories starts by recognizing features from our practices that were uncovered in earlier chapters. Let me summarize some of the relevant features. In chapter 1, I argued for the following existence and identity conditions for stories:

Existence: A story x exists iff there is some narrative content c such that c has intentionally been made accessible and x has c .⁴

Identity: For any stories x and y , x is the same story as y iff the conditions for being a properly formed token of x are the same as the conditions for being a properly formed token of y .

For an author to make some narrative content accessible, recall, is for the author to represent the content or to lay down instructions for a canonical means of representation. Simply put, it's for the author to bring it about that there is a version of the story. Recall also my claim that our story identity attributions are context sensitive. A context of use of the term 'story' introduces a counterpart relation between stories, by which they are ordered from more to less similar according to their possession (or not) of story "kernels", where a story kernel is a feature of a story that is important or salient relative to a context. This is to say that story identity attributions are true relative to a context of utterance.

⁴A worry that one might have at this point is whether this account of the existence conditions for stories assumes Platonism, or at least whether it could only be true were Platonism true. If that were the case, then it wouldn't help us in a (non-question-begging) argument for Platonism. I don't take this to be the case, however. A Creationist about stories can make sense of Existence in a way friendly to their account: a story comes into existence whenever an author makes some narrative content accessible to an audience, namely by bringing into existence the thing that has that content. Thanks to Amie Thomasson for raising this question.

I argued earlier that in order to explain these existence and identity conditions, we need to posit that *story* is a phase sortal. I gave the following account of what it is to make up a story:

For any subject S , and any story x , such that S identifies the narrative representation and content (N, c) of x , S makes up x to the degree that

1. c does not non-accidentally resemble some other narrative content c' of a different story; and
2. c does not non-accidentally represent actual events.

The picture developed is one on which an author makes up a story by identifying a way for the story to go and a way of representing it, such that way the story goes is not merely borrowed from another story or from actual events.

Three important considerations that have come out of these findings are that stories seem to be abstract, that *story* seems to be a phase sortal, and that making up a story seems to involve an intentional activity by an author of identifying some content and a means of representation for the story. Stories seem to be abstract because they are repeatable types, all that's required for their existence is the identification of some content, and what distinguishes them is the conditions for their tokening. *Story* seems to be a phase sortal because, as I argued in chapter 1, numerical identity and story identity can come apart: A claim of the type " x and y are the same story" may be true even when x and y are not numerically identical. This is best explained by positing that *story* is a phase sortal of narrative representation types, things like linguistic or pictorial types that can be used to represent content. Recognizing that authoring a story involves intentionally identifying the content and means of representation of the story is just the result of my analysis of that activity.

The first claim in my case for Platonism is that Platonism about stories best explains each of these findings. I will compare Platonism with Concretism (quickly) and Abstract Creationism (more substantively) to argue that this is the case.

Consider first Concretism, the view that stories are concrete objects. Concretism denies the thesis that stories are abstract, and it will have a difficult time explaining the phenomena that suggest the abstractness of stories, viz. their repeatability and experiencability in their instances and what seem to be the identity conditions for stories. This by itself is not enough to simply rule Concretism out, since the Concretist will certainly be able to give some account of why stories *seem* to be repeatable or experiencable in their instances. They will likewise be able to offer an alternative account of the identity of stories. It certainly counts as a cost of Concretism that it says that stories are not abstract, however.

Likewise, Concretism seems to offer a poor explanation of our earlier findings about what it is for an author to make up a story. If stories are concrete objects, then authoring must involve the identification or manipulation of features of the concrete objects that are the stories. However, as we've noted, making up a story can involve merely identifying a way for the story to go and a specification of a version of it, without reference to any particular concrete object at the time of its story's genesis. That is, an author can make up a story merely by thinking about it. But then it seems like an author can make up a story without there being any concrete object that is the story.

To be sure, making up stories often involves the creation of concrete objects, for example the writing of manuscripts, the illustration of pictures, the filming of scenes, and more. However, these objects themselves are not repeatable. Nor are we inclined to identify them with the stories that seem to depend on their creation. And, more to the point at hand, the creation of such concrete objects doesn't seem essential for making up a story. Concretism seems poorly positioned to explain these findings, then.

One way of responding for the Concretist is to suggest that in making up a story by thinking about it, an author *has* created a concrete object: namely, they've created

a brain state that is the story. There are a couple things worth saying about this kind of response. First, Concretism was supposed to be a version of Creationism, but this actually seems like a departure from that view. Previously, I defined Creationism about stories as the combination of three theses, one of which was the thesis that stories are objects. Since this way of understanding Concretism takes stories to be properties of our brains, or states of them at a time, it seems to be a departure from that standard version of Creationism. This doesn't make the view wrong, but it's an important change to track.

Second, I take it that stories as artworks must be shareable. Part of what it is to be art is to be susceptible of engagement by an audience.⁵ However, the Concretist faces a bit of a dilemma here. If they identify stories with token brain states of authors, then they don't seem sharable. Token brain states aren't multiply instantiable and not the kind of thing that can be had by more than one person. If, on the other hand, they identify them with brain state types, then they may be shareable but it's not at all obvious that they were created by the author. A plausible view of types has it that they exist prior to and independently of any particular tokening of them.⁶ However, not too much hinges on what the correct theory of types is. It's independently implausible that an author creates *either* their token or type brain states, even if they bring it about that those brain states exist. Creation is an intentional relation. When an author makes a story up merely by thinking about it, however, the attendant brain state (type or token) seems to be a by-product of that intentional activity, not its target. (I'll offer a more substantial argument for this claim below.)

Concretism is simply one way of developing a Creationist view of stories, and not at all the most popular. A much more common approach is the one adopted

⁵For support of this claim, see Grafton-Cardwell (forthcoming).

⁶See the theories of types laid out in Dodd (2007) and Wetzel (2009) for example.

by Thomasson (1999), on which works of fiction (along with fictional characters, and presumably many other repeatable artworks) are abstract artifacts, intentionally created abstract objects.⁷ Call this view Abstract Creationism.⁸

Abstract Creationism obviously accounts for the idea that stories are abstract objects. It also seems to accord with the idea that making up a story involves identifying content and a way of representing that content. On at least one typical way of developing Abstract Creationism, an abstract artifact is generated when an artist indicates certain eternally existing abstract objects that they want to be constitutive of their artifact. For example, Levinson (1990) posits that a musical work is an “indicated type”, a kind of object that is generated whenever a composer indicates a pure sound structure at a time and that is individuated by the composer’s act of indication and the context in which the act of indication occurred. The Abstract Creationist may agree with the account of authoring I give then, positing that the conditions are those required to generate a new indicated type of fiction.

Another advantage of Abstract Creationism is that it does all of the above while vindicating the idea that stories are created, historical entities. Abstract Creationism can thus seemingly explain why it seems that stories are temporal, temporally flexible, modally flexible, and teleologically oriented.

⁷Among those who have followed Thomasson’s lead are Friedell (2016, 2017, 2020b,a), Irmak (2012, 2020), Trivedi (2002), and Walters (2013).

⁸Sometimes the term ‘artifactualism’ has been used, and the proponents of this view often refer to things like stories and musical works as abstract artifacts. I think it would be misleading for me to use this terminology, since I also take these objects to be artifacts. On a theory of artifacts on which an object can be “artifactualized” by appropriation (for example, a piece of driftwood being appropriated as a winerack), *artifact* is a phase sortal, and making new artifacts doesn’t require the creation of new objects. On the version of Platonism I prefer, we ought to hold such a view of artifacts and posit that, in making up a story, an author appropriates a type of narrative representation as a kind of artifact. What’s really at issue between the Abstract Creationist and the Platonist about stories, then, is the question of whether stories are created or discovered by their artists.

The person I’m labeling “Creationist” may not appreciate the term, as some take it to be a slur. Let me cancel that implicature here.

However, there are a couple important ways in which Abstract Creationism doesn't seem to fit with the earlier findings about the nature of authoring of stories. First, it's not actually so obvious that Abstract Creationism can actually posit that authoring stories involves the intentional identification of content and form of the story, unless they also allow for vague existence. This is because the identification of content and form comes in degressed stages. When an artist is early in their process of working on the story, it may be in many respects indeterminate what the content or form of the story is. If there's indeterminacy in the identification of content or form, though, and the conditions for generation of a story are the identification of those things, then there's indeterminacy in the extent to which the conditions for generation of the story have been met. Either the Abstract Creationist needs to posit a sharp cutoff in what degree of identification is required for the generation of an abstract artifact or they need to deny that identification comes in degrees. If neither move is taken, then the Abstract Creationist will be committed to the vague existence of stories.

The problem is that both moves seem objectionable. I've just said why it seems that identification is a degressed notion. There are plenty of times when an author is working on a project where it seems that the story exists but it's indeterminate what the content of the story is or what the canonical means of representation is (or will be). The most palatable option for the Abstract Creationist who wants to avoid vague existence, then, may be for the Abstract Creationist to posit a sharp cutoff in the degree of authoring required for a new abstract artifact to be generated.

But the sharp cutoff view should strike us as implausible. It has the consequence that in such cases, when an author is working on a story (perhaps even for years), trying to settle some important decisions about, say, whether it will be a novel or a graphic novel, or whether the main character will be young or old, that there is no story they're working on. Only when there is a determinate content and form does the story come into existence, susceptible to change. Not only does that seem ad

hoc, as a story about the metaphysics of such situations, it also fails to accord with what the artist's phenomenology is likely to be. For Abstract Creationists that don't want to be committed to vague existence, it seems necessary then to come up with an alternate account about what is going on when someone makes up a story.⁹

Abstract Creationism also must deny that *story* is a phase sortal. The reason is that if *story* is a phase sortal, there can come to be a new story without any new object coming into existence. However, Creationists posit that the way a story is brought about just is by bringing a new object into existence. They cannot countenance the possibility of something's merely becoming a story (while already having existed prior to its being a story). I argued in chapter 1 that *story*'s being a phase sortal follows from the identity conditions for stories. The Abstract Creationist must therefore reject those identity conditions as well.

It's likely that they will offer up some version of a historical account of the identity conditions for stories instead, on which stories are individuated by their histories. In chapter 1 I examined some difficulties for historical accounts, including the challenge such accounts face to give a principled answer about how much change stories can undergo, and the way such accounts seem unable to escape a reliance on content and form to determine the identity of stories. Those arguments count against Abstract Creationism then, insofar as that view is committed to a historical account of story individuation.

I take these considerations to offer significant explanatory reason to favor Platonism over Abstract Creationism. As we saw in Chapter 4, Platonism seems at least as well situated (if not better situated in some cases) to explain the features of stories that Creationism is meant to explain. But there are some important considerations that Platonism can offer a better account of, including the identity conditions for

⁹Of course, this argument only functions against the Abstract Creationist who wants to deny vague existence. Importantly, Thomasson (2007) allows for vague existence.

stories (and truth conditions for story identity attributions), what it is to author a story, and the plausibility of a phasalist understanding of stories. On balance, I think we should prefer a Platonist theory.

To supplement these considerations, however, I want offer up a more direct argument against Abstract Creationism. First, I'll argue that if the relation posited between authors and stories by Abstract Creationism is non-causal, then it is not properly understood as creation. Second, I'll argue that if it is causal, then it can't apply to abstract objects. The result is a dilemma for Abstract Creationism: either artists generate (but don't create) the objects in question, or there are no such objects.

5.4 The No Creation Argument

Take the following argument from Irmak (2020), slightly adapted:

- (A1) Stories are abstract objects.
- (A2) Abstract objects are causally inert entities.
- (A3) If one creates something, one causes it to exist.
- (A4) Therefore, stories are not created.

Given that we've framed this as an argument aimed at the Abstract Creationist, there are two ways for the Abstract Creationist to respond: they can reject (A2) or (A3).

Irmak rejects (A3), providing an account of non-causal creation in terms of intentional generation. He acknowledges that (A3) is a popular thesis. Some authors merely put their foot down and insist that creation of an object necessarily involves causing the object to exist. For example, Brock (2010, 343) holds that

... causal creation is the process of bringing something into being or causing it to exist. My view is that [non-causal] creation isn't a kind of creation at all; causal creation is the only variety of genuine creation.

Likewise, Dodd (2000, 431) (citing the support of Predelli 1995) claims that “The creation of an abstract object would have to be a kind of causal interaction between a person and an abstract object or objects,” and rejects it on that account.

Other authors allow that one might offer up a stipulative definition of “create” according to which creation is non-causal, but these authors hold that such an account of creation will necessarily leave unexplained *how* abstract artifacts are generated. It’s just not clear how abstract objects come into being if they aren’t caused to do so, since otherwise their genesis appears to be mysterious, akin to a sort of magic.

In defending the view that creation doesn’t require causation, Irmak takes on the burden of defending non-causal creation from the second kind of charge, resting his case on an attempt to make non-causal creation unmysterious. Irmak attempts to remove the mystery by giving an account that explains just how an abstract object is generated in non-causal creation.

On Irmak’s account, a relation (or an instance thereof) is “generative” “iff its instantiation brings something into existence.” (Irmak 2020, 7) More precisely, a relation is generative just in case

- (i) a numerically distinct entity comes into existence as a result of its instantiation and
- (ii) the generative relatum, i.e. the parts, the causes, the members, and the instantiation of the relation guarantee the existence of a new entity, event, or a property. (Irmak 2020, 8)

A relation is generative just in case the instantiation of the relation by some objects guarantees that a new object will come into existence. Examples of generative relations for Irmak include causation, composition, constitution, and set-formation.

To formulate his account of creation, Irmak distinguishes between species of ontological dependence that are generative and those that are not, as well as between intentional and unintentional instantiations of generation relations. Relations of ontological dependence track what objects depend on what other objects for their existence. They’re generative when they meet the conditions specified above, and inten-

tional whenever some agent intentionally brings it about that the conditions specified in the ontological dependence base are satisfied.

On Irmak's view, the relation between a composer and their musical work is a relation of rigid existential dependence. This is to say, things like musical works depend for their existence on the sound and performance-means structures that are specified by composers, on the particular composers who bring about the satisfaction of their dependence bases, and on the very creative acts those composers take. Adapting the theory to fit stories, the view is that a story depends for its existence on the existence of the content and form indicated by the author, the author, and the author's act of indication of that content and form (at the time they in fact indicate these things). This is a generative relation, since the instantiation of ontological dependence in this case is what brings into existence a new entity, the story. At least in normal cases, it's also an intentional relation, since the author indicates what they do with the intent of generating a story.¹⁰ In total, the hope is to have provided a non-mysterious account of non-causal creation.

I think both accusations against non-causal creation remain intact. That is, I think we still have reason to think there's no such thing as non-causal creation, and we should still take the account on offer to be objectionably mysterious. First, take the claim that only causal creation is genuine creation. We can add support for this claim by considering other cases of generation that don't seem to be creation. For example, assuming a set theory that allows for impure sets, if I bring a new concrete object (say a table) into existence, I thereby bring it about that a set is generated (singleton table). However, there doesn't seem to be any natural sense in which I've created the set. I created the table. The set was a byproduct.

¹⁰We might actually question how often that's the case, but we can spot Irmak this assumption here.

Recall, however, that Irmak distinguishes between intentional and unintentional generation:

There is a wider sense of creation in which any object that comes into existence—partly in virtue of some human (or animal) activity—constitutes a created object. Therefore, we need to make a distinction between mere production or generation and intentional creation of objects. Although different kinds of objects come into existence by the same human activity, it simply is not true that they are all necessarily artifacts. Artifacts are intentional creations. (Irmak 2020, 9)

In Irmak’s view, any product of a generation relation is a creation, but *artifacts* only come about as a result of intentional creation.¹¹ In the case of the set we’ve just considered, Irmak will point out that it’s a case of unintentional creation. This is a substantial difference, on his view, since only cases of intentional creation generate an artifact. The intuition we have that I didn’t create the set can be explained by the fact that I didn’t create an artifact by creating the set.

I don’t think this response saves Irmak from the objection. After all, we can imagine a scenario in which I, being a philosophically informed carpenter, or a boundary-pushing artist, create a table with the intent of generating an impure set. (“The art isn’t the object, it’s the singleton set of the object!”) Perhaps I do this on a regular basis. This intention doesn’t seem to make any difference with respect to my having created the set. It may be an intended byproduct, but it’s still just a byproduct.

I think Irmak will respond by digging in and insisting that if a set is an intentional byproduct then it’s an artifact, and so an abstract artifact. After all, Irmak endorses the claim that “to create an abstract artifact is to intentionally bring about the conditions that are specified in its dependence base.” (Irmak 2020, 9) Likewise, in explaining why the singleton of the Eiffel Tower is *not* an artifact, Irmak says “the

¹¹Irmak acknowledges that some (e.g. Cray (2017)) take “generation to be a broader notion and reserve creation for intentional generation of artifacts,” but claims that “[n]othing in our discussion here depends on this terminological choice.” (Irmak 2020, 9, fn 15) This claim is odd given that part of what’s at issues is what genuine creation amounts to.

singleton of the Eiffel Tower is an unintended product of the people who were involved in the creation of the tower.” (Irmak 2020, 11) That is, the reason the singleton of the Tower is not itself an artifact is because it was an unintended (rather than an intended) byproduct. Intention is the sole difference-maker, then, for Irmak, between generated objects that are and those that are not artifacts.

This commitment to intention as the sole difference-maker comes into conflict, it turns out, with another thesis about artifacts to which Irmak commits himself: “Artifact creation minimally requires the intention to bring an object of a kind K into existence, where K is an artifactual kind.” (Irmak 2020, 9) This commitment seems innocuous at first. It’s connected to the idea that artifacts are intentional creations, so it would be surprising if, as I’m claiming, it turned out to conflict with the claim that intention is all one needs to get an artifact. The important thing to recognize is that, as seems plausible, someone brings an artifact into existence only if they have an intention to bring something of a certain artifactual kind into existence. This brings up the question of whether *set* is an artifactual kind. It seems quite clear that it isn’t. After all, most singleton sets aren’t artifacts. Irmak will agree: they weren’t intentionally generated.

Perhaps Irmak would respond by drawing a distinction between two kinds of sets: sets that are intentionally generated and sets that are unintentionally generated. On this line of response, the intentionally generated sets are an artifactual kind while the unintentionally generated sets aren’t. This distinction would be clearly ad hoc. There would be no difference between the existence or identity conditions of intentionally generated and unintentionally generated sets. Likewise, unintentionally generated sets could still come to be thought about by us and regarded in the same ways as unintentionally generated sets. It doesn’t seem that there could be any independent reason to treat them as intrinsically different kinds of sets.

If *set* is not an artifactual kind, though, as I've just argued, then it doesn't seem that intentionally generating a set is sufficient to produce an artifact, given Irmak's own commitments. Intention isn't enough by itself to get one from a case of unintentional generation to the creation of an abstract artifact. We have good reason, then, even on Irmak's own understanding of artifacts, to think that "non-causal creation" is not genuine creation. Even in cases where it's intentional, the non-causal generation of an object doesn't seem to be creation of that object.

But there's also good reason to think the notion of non-causal creation is mysterious, as others have alleged. Irmak hopes to dispel this worry by giving an account of the conditions under which a new work is generated, but the picture that follows seems to commit the Creationist to the rejection of a plausible anti-idealist principle.

The principle I have in mind is the following: identifying an object does not make an ontologically significant difference. An ontologically significant difference, as I'm conceiving it, is a difference in either what objects exist or how some object is intrinsically. If something makes an ontologically significant difference, then that thing will make a difference in what objects exist or how some object is intrinsically. For the purposes the anti-idealist principle I'm endorsing, identifying some content (in the sense under consideration) amounts to just thinking about it. It may be more than this, as when an author forms some intentions towards the content of their story (for example the intention to represent the content in question). But the thrust of the anti-idealist principle is that merely thinking about an object does not by itself result in a difference in what objects exist or in how some object is intrinsically.

(Of course, strictly speaking that's obviously false. For example, it's subject to the counterexample that what I think about makes a difference in what thoughts exist and in how I am intrinsically. It's even the case that our having thoughts can be the causal impetus for things coming into existence or for intrinsic change. If it occurs to me that I'm hungry, I may bring a sandwich into existence, in the process

changing the intrinsic character of the ingredients by slicing, spreading and so forth. So what am I really trying to say? The idea is that our thoughts, *by themselves*, make no ontologically significant difference to things *external to us*. When I make the sandwich, my thoughts may figure in the causal story, but my thoughts are not by themselves sufficient to bring it into existence (sadly). And while my thoughts are sufficient by themselves to determine some facts about how I am intrinsically, they can't by themselves determine anything about what exists (external to me) or how things (external to me) are intrinsically.¹² Something x makes an ontologically significant difference if x makes a difference in what objects exist or in how some object is intrinsically, external to x . I will bracket the talk of externality except when it's relevant from here on out.)

But why might we think Abstract Creationism is committed to the ontological significance of thought.¹³ The kinds of cases that seem to show this are those in which an author makes up a story without creating any concrete objects: they may do this, for example, just by thinking through the plot of the story and then forming plans to write it (or tell it in some other way). Before the writing has even begun, it seems the story has been made up: it exists. But there doesn't seem to be any

¹²This way of putting things may still seem to be false on the assumption that composition is unrestricted. For if composition is unrestricted then any time I have a thought there are many objects whose intrinsic natures change, solely in virtue of the fact that those objects have me as a part. My response here is to say that the sense of "external" I had in mind when I said that our thoughts don't make any ontologically significant difference to things external to us would not necessarily count such things as external to us. Thus, the fact that my thoughts could change the intrinsic natures of those objects provides no counterexample to my claim. Maybe the best gloss on what I have in mind here by "external to x " is "any object that is not identical to x , that is not a part of x , and that x is not a part of." I'm not sure if that perfectly captures it, but I think it should do well enough here.

¹³Some Abstract Creationists will be happy to commit to this. For example, Thomasson (2015, chapter 2) endorses the idea that there are no conditions on existence that apply universally. In order to understand when objects of any given kind come into existence, we must understand the existence conditions of that kind taken individually. This leaves her view in principle committed to the possibility that thought is ontologically significant, since it leaves open that for any particular kind, it may be that what's required for the existence of members of that kind is that someone think some thoughts.

concrete object the creation of which contributes to the existence of the story. Instead, it's merely the author's activity of identifying the story's content and canonical means of representation that bring it into existence. Call cases in which an author makes up the story by thinking "Brainstorm Cases." In Brainstorm Cases, it's merely the author's thought that brings the story into existence, and if Abstract Creationism is right, then it seems as if thought must be ontologically significant in such cases.

One way the Abstract Creationist might respond to Brainstorm Cases is actually to deny that they are genuine cases of story-creation. That is, the Abstract Creationist might hold that an author has not actually made up a story until they have created some concrete object that can serve as what Rohrbaugh (2003) calls an embodiment of the work, i.e. something that preserves what the work is like and leads to the work's perpetuation. This response seems ad hoc, though. We can all imagine (and many of us have experienced) Brainstorm Cases. They are not phenomenologically all that different from cases in which the process of identification that's required for making up a story is supplemented by the creation of concrete objects like scripts, story boards, or sketches. To deny that the author can make up a story via Brainstorming seems to arbitrarily deny the author power that she seems in fact to have.

If, on the other hand, the Abstract Creationist agrees that Brainstorm Cases are genuine cases of story-creation, then they seem to be committed to the ontological significance of thought. They are committed to the idea that, merely by thinking, an author can bring a new object into being.

Another potential response available here is to posit that an author does create a concrete "object" in Brainstorm Cases, namely a state-of-affairs. Take some content, c and an author, S who has made up a story by thinking about c . Then, the suggestion might be that S , by thinking about c , has created a state-of-affairs, *c-as-thought-about-by-S*. But, since this state of affairs is not *external* to S , S 's creation of that object by thought alone doesn't violate the principle that thought is not ontologically

significant. Likewise, the response might continue, that since the abstract object that is the story depends for its existence on this state-of-affairs that's internal, in the relevant sense, to the author's mental life, we can posit that the story is itself internal. The author brings it into existence merely by thinking, but once again that's no violation of the principle that thought is not ontologically significant. The story is "in" the author, so to speak.

The problem with this line of response is that it threatens to make stories private. As I've already stressed, stories, as representational art objects, must be in principle sharable. If a story is merely internal to its creator, then it doesn't seem to be something that can be shared with others. It's a part of the author's mental life, and not an entity external to the author that can be experienced and enjoyed by others. So, once again, the Abstract Creationist faces a choice: they must allow that stories are sharable but admit that thought is ontologically significant, or if they deny this they seem to make stories into a kind of objectionably private object.

Brainstorm Cases seem to show that the Abstract Creationist is committed to the ontological significance of thought, then. At least, either the Abstract Creationist is committed to the ontological significance of thought, or any attempt to avoid this commitment results in another objectionable claim, like that stories aren't sharable or that Brainstorm Cases aren't genuine cases of story-creation.

Now, one can only say so much directly against the idea that thought is ontologically significant. Perhaps you agree with me that it isn't, or perhaps you don't see the problem. If you find yourself in the latter situation, you might hope to hear whether Platonism holds any theoretical advantages over Abstract Creationism. I think it does, and I've already discussed some reasons why in the previous section. I want to discuss one other advantage of Platonism, one that was touched on briefly in chapter 4. Abstract Creationism, as we've understood it, says there is a plenitude of content from which authors identify their story's content when they make it up.

The Abstract Creationist claims that stories are distinct objects from this content, objects that are created by the author and that have content. An important question to ask about the Abstract Creationist's stories is what they are like in themselves. That is, what are their intrinsic natures?

The answer, it seems, is that they are intrinsically very simple, to the point that there are no real intrinsic differences between stories.¹⁴ Friedell (2020b, 12) puts this point well with respect to musical works: “musical works have no (or hardly any) important intrinsic properties, other than those that follow from their being created abstract partless objects.” An objection one might have to such a view (one Friedell himself is sensitive to) is that if this is the case then there is no intrinsic difference between any such works, say between two stories, or even (on Friedell's view) between a musical work and a story. That is, there is no difference in how one story is in itself from how any other story is in itself, or in how a story is in itself from how a musical work is in itself. But most of us would've been inclined to say there were many intrinsic differences between stories (or between stories and musical works). If two stories have different content, that should be understood as a difference in the intrinsic features of the stories, for example. Likewise, we would've said that musical works are all in some important way intrinsically different from things like novels and movies. Platonism can better account for this. It allows for more intrinsic difference between abstract artifacts, and so can account naturally for the thought that stories (and other repeatable, tokenable artworks) are intrinsically distinct.

There's good reason to side with the Platonist then, whether or not one takes thought to be ontologically significant. If one does accept the principle that thought is not ontologically significant, then one should also accept the conclusion that there

¹⁴I don't think this is actually the only way to develop Abstract Creationism. I brought up Evnine (2016) as an example of Creationism earlier. However, I think Friedell (2020b) offers compelling arguments that his version of Creationism has important advantages over Evnine's, and Evnine doesn't address stories, so I suppress any discussion of his work here.

is a sense in which Abstract Creationism remains mysterious. Irmak gives an account of the conditions under which an abstract artifact comes into existence, but given that those conditions allow for abstract artifacts to be generated by the mind alone, we should be suspicious.

There are a number of reasons to be skeptical of Abstract Creationism, then. Platonism better explains the existence and identity conditions for stories, the fact that *story* seems to be a phase sortal, and the apparent differences between the intrinsic natures of stories. It's also the case that we have good reason to think only causal creation is genuine creation, and we have good reason to think thought is not ontologically significant. Thus, non-causal creation seems impossible, and the account of it is mysterious. In the following section, I'll turn to the argument over (A2) of the No Creation argument, the premise that abstract objects are causally inert.

5.5 Causal Abstracta

In the previous section, I argued that creation must be understood causally. That is, if the relation posited to hold between authors and their stories by Abstract Creationists is non-causal, then it isn't creation. In this section, I will argue that if the relation between authors and their stories *is* causal, then stories are not abstract objects. The conclusion is that abstract objects are causally inert, and so no one could cause them to exist. This is to secure premise (A2) of the previously discussed No Creation Argument:

- (A1) Stories are abstract objects.
- (A2) Abstract objects are causally inert entities.
- (A3) If one creates something, one causes it to exist.
- (A4) Therefore, stories are not created.

Since the Abstract Creationist is committed to (A1), and given my argument for (A3) in the prior section, this will conclude my support for the premises of the argument.

In fact, what I'll argue is that abstract objects are unable to participate in a certain species of causal relation, what I call a causal-production relation. On the other hand, I'll argue that it's exactly this kind of causation that's relevant to the question of creation. Thus, while we have reason to see abstract objects as causal in some other (e.g. difference-making) sense of "causal", they are causally inert in the way that matters for supporting (A3).

My argument here is largely a response to Friedell (2020a), who argues that we should take abstracta to be causally efficacious, contrary to the standard view that takes them to be causally inert. Friedell proceeds from the observation that sentences like

(1) *Uncle Tom's Cabin* caused many Americans to support abolition.

seem true to argue that abstracta are causal.

Friedell's "straightforward" argument is as follows (Friedell 2020a, 135):

P1: (1) is true.

P2: If (1) is true, then some abstracta are causally efficacious.

C: Some abstracta are causally efficacious.

Those who resist P1, Friedell calls "error theorists", since they claim that we're in error when we assert sentences like (1). Those who resist P2, Friedell calls "paraphrasers", since they attempt to provide a way of understanding sentences like (1) on which C doesn't follow.

My response can probably best be described as a version of Friedell's paraphrasing strategy. I don't argue that C is false, however, nor do I contest any of the premises. Instead, I argue that the argument is only sound if terms like "caused"

and “causally efficacious” are understood as denoting a difference-making species of causation. However, the sense of ‘cause’ that’s relevant to the No Creation Argument is not a difference-making sense, but the sense of causation as production or interaction. Thus, while Friedell may offer a sound argument that abstract objects are causally efficacious on one understanding of that term, his conclusion does nothing to contradict premise (A2) of the No Creation Argument.

To begin analyzing Friedell’s argument, notice first that (unlike Irmak 2020) Friedell accepts premise (A3) of the No Creation Argument. “After all, if there are abstract artifacts, then people presumably *cause* them to exist.” (Friedell 2020a, 134) Friedell agrees with me about premise (A3), then, and his response to the No Creation argument must stand or fall with premise (A2).

In order to build a case against premise (A2), Friedell focuses on propositions like (1) above, in which we seem to truthfully attribute causal efficacy to abstract artifacts like Harriet Beecher Stowe’s *Uncle Tom’s Cabin*. Since we seem apt to assert such things, and since they seem to entail that some abstract objects are causally efficacious, we should allow that this is so.

Friedell considers two strategies for responding to his argument. The first, error theories, take propositions like (1) to be false. Friedell thinks the way to motivate such error theories is by appeal to a general view about causation. For example, he considers what he calls the “pushing-touching” view and the “energy transfer” view, on which (respectively) to cause an effect an object must either push or touch an object that’s part of the effect or must transfer energy to an object that’s part of the effect. Friedell rejects both views because of their inability to handle cases of “causation by disconnection,” cases in which an object causes an effect by preventing another effect. For example, I can cause my floor to get wet by opening the window when it’s about to rain.

The second kind of strategy for resisting Friedell’s argument is paraphrasing. Paraphrasers accept the truth of propositions like (1), but deny that such sentences entail that any abstract objects are causally efficacious. The only version of a paraphrasing strategy Friedell considers is the proposal that “*Uncle Tom’s Cabin*” is a metonym, so that the term denotes token copies of the novel instead of the novel itself. Friedell provides compelling reasons to doubt this paraphrasing strategy.

Another, nearby approach remains, however. Instead of positing an ambiguity in our use of “*Uncle Tom’s Cabin*”, it’s open to Friedell’s opponent to posit that there’s an important ambiguity in our use of “cause” and related causal terminology. This is already a common view in the literature on causation.¹⁵ A simple version of such a view takes it that we have “two distinct *concepts* of cause that have become tangled together in our language.” (Godfrey-Smith 2010, 330) I’ll assume that *at least* this much is true. That is, we have at least two distinct concepts of cause (but could have more). One concept (or species of concept, if things divide up further), I’ll refer to as “interactive causation.” Interactive causation is the kind of causation picked out by those who propose pushing-touching or energy-transfer views about causation. Some things we think of as causes because (whether by pushing, touching, transferring energy, or in some other salient way) they directly interact with the objects that are part of the effects they cause.

The other concept (or species thereof) I’ll refer to as “difference-making causation.” This is the kind of causation picked out by those who propose counterfactual analyses of causation. Some things we think of as causes because (even if they don’t directly interact with the objects that are part of their effects) they make a difference as to what happens.

¹⁵For articulations and defenses see Hall (2004) and Godfrey-Smith (2010).

Interactive causation and difference-making causation are not the same thing. Cases of causal overdetermination seem to show this, as it's natural to say that the first bullet with which Rasputin was shot was an interactive cause of his death even if it was not a difference-making cause (given that he was shot two other times, poisoned, wrapped up, and thrown into a freezing river). Likewise, something can be a difference-making cause without being an interactive cause, as our earlier example of opening a window shows: in opening a window I cause the floor to get wet without interacting with it in any way. My opening the window makes a difference as to whether the floor gets wet, even though I don't interact with the water or the floor.

My suggestion is not that premise P2 of Friedell's argument is false, but that the only sense in which it can be true is if "causally efficacious" is understood in terms of difference-making causation. Friedell's own examples seem to lend credence to this. *Uncle Tom's Cabin* made a difference to how many Americans supported abolition because if it hadn't been written (and if Americans hadn't read it) then many fewer Americans would've supported abolition. Other examples provided by Friedell can be understood similarly:

- *Songs in the Key of Life* caused Mary J. Blige and George Michael to record a duet.
- π caused Taylor (the mathematician) to lose sleep.

It seems natural to understand these kinds of claims in a difference-making sense. If *Songs in the Key of Life* hadn't been written, then Mary J. Blige and George Michael wouldn't have recorded a duet (or they wouldn't have recorded the duet they did). If Taylor hadn't known about π , then Taylor wouldn't have lost that sleep.

We have further reason to hold to the idea that abstract objects are only causally efficacious in a difference-making fashion: it seems like an important part of understanding the abstract-concrete divide that abstract objects be unable to directly

interact with concrete objects. After all, if abstract objects *can* directly interact with concrete objects, so that they are even interactively causally efficacious, we might wonder what separates them in any substantive, categorial way from concrete objects.

One potential suggestion (one that's commonly understood to be true about abstract objects) is that they are non-spatiotemporal. However, this option doesn't seem open to the Abstract Creationist. If they posit that abstract artifacts come into being at a particular time, then it seems that they must at least be temporally located. Moreover, if abstract artifacts aren't temporally located, then it's implausible that they can serve as interactive causes: when would a timeless entity interact with a temporal one to produce an effect, after all? It seems that abstract objects must exist in time, then, on Abstract Creationism.

An option many Abstract Creationists have taken (following Thomasson 1999) is to allow abstract artifacts to be temporally located but to distinguish them from concrete objects by denying them any spatial location. Juvshik (2020) has provided a serious reason to abandon this option, however. Juvshik's argument proceeds as follows (Juvshik 2020, 270):

1. If x is causally efficacious, then x has a temporal location.
2. If x has a temporal location, then x has a spatial location.
3. If x has a temporal location, then x is not abstract.
4. So, if x is causally efficacious, then x is not abstract.

The relevant premise for our purposes is (2), which bans the Abstract Creationist from separating spatial and temporal locatedness.

Juvshik's argument for premise (2) is that Lorentzian spacetime—"the spacetime of our best physical theories"—posits "only a single, *spatiotemporal* metric, which

yields *spatiotemporal* distances (or the spacetime interval) between things or events.” (Juvshik 2020, 274) In Lorentzian spacetime, anything with a temporal location has a spatial location, because the only such locations are spatiotemporal. The Abstract Creationist who goes this way thus faces the cost of disagreement with our best physical theories.¹⁶

It’s open to the Abstract Creationist to posit that the abstract-concrete distinction is a primitive one, so that there’s no reduction of the distinction and not even any powers distinctive of one category that members of the other category necessarily lack. That way seems to provide us with no reason to posit that there even is any such distinction, though. If abstract objects are not distinctly different from concrete objects, being located in spacetime and acting as interactive causes, what reason is there to even posit a distinction between them?¹⁷

It seems, then, that causal inertness is thus essential for understanding how abstract objects can be relevantly different from concrete objects. Given that this is so, there doesn’t seem to be any good reason to deny the interactive causality of abstract objects. And given that denial, premise (A2) of the No Creation Argument stands. We seem to have strong reason to take stories not to be created.

¹⁶Alternatively, the Abstract Creationist may simply be engaged in conceptual analysis and not be attempting to say anything that could be construed as able to disagree with our best physical theories.

¹⁷One further way of drawing the divide, suggested by Cowling (2017, 88), who suggests a view on which abstract objects are distinct from concrete in that there can be no object distinct from but qualitatively indiscernible with an abstract object (while there may be for any concrete object). This way may be worth exploring, but it’s not open to an Abstract Creationist like Friedell, who (as previously discussed) takes abstract objects to be simple, having “no (or hardly any) important intrinsic properties, other than those that follow from their being created abstract partless objects.” (Friedell 2020b, 12) For Friedell (and at least some other Abstract Creationists) abstract artifacts are all indiscernible one from another.

5.6 Conclusion

In this chapter I've paved the way for a multi-faceted case for Platonism. After first arguing for the legitimacy of a Discovery View in the metaontology of art, I showed how Platonism best explains a number of features of stories that have been made salient by the discussion of stories in the foregoing chapters. Given that Platonism seems able to explain all the features of stories that Creationism can, while also serving as a superior explanation of some of these other features, I take it that we have reason on balance to prefer a Platonist theory of stories.

I added to that abductive case an argument against Creationism in the form of a defense of both contested premises of the No Creation Argument outlined in Irmak (2020). First, I argued that we have reason to think that if someone creates something, they cause it to exist. Intentional non-causal generation does not by itself seem sufficient to count as creation. Moreover, holding to the existence of non-causal creation turns out to run afoul of what I take to be a plausible anti-idealist principle. I also argued that abstracta are causally inert, in the sense that if stories are caused to exist (in any interactive sense of causation), then they are not abstract.

It follows from these arguments that stories are not created. We have then an abductive case in favor of Platonism supplemented by a case against Creationism. This provides a compelling narrative in favor of the idea that stories are discovered, not created.

BIBLIOGRAPHY

- Aliyev, Alexey. 2017. "Musical Perdurantism and the Problem of Intermittent Existence." *Grazer Philosophische Studien* 94:83–100.
- Alter, Robert. 1978. "Biblical Type-Scenes and the Uses of Convention." *Critical Inquiry* 5:355–368.
- Badura, Christopher and Berto, Francesco. 2018. "Truth in Fiction, Impossible Worlds, and Belief Revision." *Australasian Journal of Philosophy* 97:178–193.
- Bal, Mieke. 2017. *Narratology: Introduction to the Theory of Narrative*. University of Toronto Press, fourth edition.
- Beardsley, C., Monroe. 1982. *The Aesthetic Point of View*, chapter "The Creation of Art". Cornell University Press.
- Becker, Howard, Faulkner, Robert, and Kirshenblatt-Gimblett, Barbara. 2006. *Art from Start to Finish*. University of Chicago Press.
- Brock, Stuart. 2010. "The Creationist Fiction: The Case Against Creationism about Fictional Characters." *Philosophical Review* 119:337–364.
- Byrne, Alex. 1993. "Truth in Fiction: The Story Continues." *Australasian Journal of Philosophy* 71:24–35.
- Caplan, Ben. 2014. "Serial Fiction, Continued." *British Journal of Aesthetics* 54:65–76.
- Caplan, Ben and Matheson, Carl. 2006. "Defending Musical Perdurantism." *British Journal of Aesthetics* 46:59–69.
- . 2008. "Defending 'Defending Musical Perdurantism'." *British Journal of Aesthetics* 48:80–85.
- Carroll, Noël. 2001. *New Perspectives on Narrative Experience*, chapter "On the Narrative Connection". SUNY Press.
- Chatman, Seymoure. 1978. *Story and Discourse: Narrative Structure in Fiction and Film*. Cornell University Press.
- Cowling, Sam. 2017. *Abstract Entities*. Routledge.

- Cray, Wesley D. 2017. "Abstract Generationism: A Reply to Friedell." *The Journal of Aesthetics and Art Criticism* 75:289–292.
- . 2018. "Psychologism about Artistic Plans: A Response to Rohrbaugh." *The Journal of Aesthetics and Art Criticism* 76:101–104.
- Cray, Wesley D. and Matheson, Carl. 2017. "A Return to Musical Idealism." *Australasian Journal of Philosophy* 95:702–715.
- Currie, Gregory. 1990. *The Nature of Fiction*. Cambridge University Press.
- . 1997. *Emotion and the Arts*, chapter "The Paradox of Caring", 63–77. Oxford University Press.
- . 2006. "Narrative Representation of Causes." *The Journal of Aesthetics and Art Criticism* 64:309–316.
- Davies, David. 1996. "Fictional Truth and Fictional Authors." *British Journal of Aesthetics* 36:43–55.
- . 2004. *Art as Performance*. Blackwell Publishing.
- . 2009. "The Primacy of Practice in the Ontology of Art." *Journal of Aesthetics and Art Criticism* 67:159–171.
- . 2017. "Descriptivism and its Discontents." *The Journal of Aesthetics and Art Criticism* 75:117–129.
- Deutsch, Harry. 2000. *Making Up Stories*, 149–182. CSLI Publications.
- . 2013. "Friend on Making Up Stories." *The Aristotelian Society* 113:365–370.
- Devitt, Michael and Sterelny, Kim. 1987. *Language and Reality*. Basil Blackwell.
- Diamant, Kathi. 2003. *Kafka's Last Love: The Mystery of Dora Diamant*. Basic Books.
- Dodd, Julian. 2000. "Musical Works as Eternal Types." *British Journal of Aesthetics* 40:424–440.
- . 2007. *Works of Music: An Essay in Ontology*. Oxford University Press.
- . 2012. "Defending the Discovery Model in the Ontology of Art: A Reply to Amie Thomasson on the *Qua* Problem." *British Journal of Aesthetics* 52:74–95.
- . 2013. "Adventures in the Metaontology of Art: Local Descriptivism, Artefacts and Dreamcatchers." *Philosophical Studies* 165:1047–1068.
- Evans, Gareth. 1985. *Collected Papers*, chapter "The Causal Theory of Names", 1–24. Oxford University Press.

- Evnine, Simon J. 2009. "Constitution and Qua Objects in the Ontology of Music." *British Journal of Aesthetics* 49:203–217.
- . 2013. "Ready-Mades: Ontology and Aesthetics." *British Journal of Aesthetics* 53:407–423.
- . 2016. *Making Objects and Events*. Oxford University Press.
- Fisher, J. A. 1991. "Discovery, Creation, and Musical Works." *Journal of Aesthetics and Art Criticism* 49:129–136.
- Friedell, David. 2016. "Abstract Creationism and Authorial Intention." *Journal of Aesthetics and Art Criticism* 74:129–137.
- . 2017. "Music and Vague Existence." *Res Philosophica* 94:437–449.
- . 2020a. "Abstracta are Causal." *Philosophia* 48:133–142.
- . 2020b. "Why Can't I Change Bruckner's *Eighth Symphony*?" *Philosophical Studies* 177:805–824.
- Friend, Stacie. 2007. "Fictional Characters." *Philosophy Compass* 2:141–156.
- . 2008. *Imagining Fact and Fiction*, 150–69. Basingstoke: Palgrave Macmillan.
- . 2011. "Fictive Utterance and Imagining II." *Proceedings of the Aristotelian Society Supplementary Volume* 85:163–80.
- . 2012. "Fiction as a Genre." *Proceedings of the Aristotelian Society, New Series* 112:179–209.
- García-Carpintero, Manuel. 2007. "Fiction-Making as an Illocutionary Act." *Journal of Aesthetics and Art Criticism* 65:203–216.
- García-Carril Puy, Nemesio. 2019. "The Ontology of Musical Versions: Introducing the Hypothesis of Nested Types." *The Journal of Aesthetics and Art Criticism* 77:241–254.
- Genette, Gérard. 1980. *Narrative Discourse: An Essay in Method*. Cornell University Press.
- Godfrey-Smith, Peter. 2010. *Oxford Handbook of Causation*, chapter "Causal Pluralism", 326–337. Oxford University Press.
- Gover, K. E. 2011. "Artistic Freedom and Moral Rights in Contemporary Art: The Mass MoCA Controversy." *The Journal of Aesthetics and Art Criticism* 69:355–365.
- . 2015. "Ambivalent Agency: A Response to Trogdon and Livingston on Artwork Completion." *The Journal of Aesthetics and Art Criticism* 73:457–460.

- . 2018. *Art and Authority*. Oxford University Press.
- Grafton-Cardwell, Patrick. forthcoming. “The Aesthetic Engagement Theory of Art.” *Ergo* .
- Grice, Herbert Paul. 1957. “Meaning.” *Philosophical Review* 66:377–88.
- Hall, Ned. 2004. *Causation and Counterfactuals*, chapter “Two Concepts of Causation”, 225–276. MIT Press.
- Hawthorne, John and Lepore, Ernie. 2011. “On Words.” *The Journal of Philosophy* 108:447–485.
- Hazlett, Allan. 2012. *Art & Abstract Objects*, chapter “Against Repeatable Artworks”, 161–178. Oxford University Press.
- Hick, Daniel Hudson. 2008a. “When Is a Work of Art Finished?” *The Journal of Aesthetics and Art Criticism* 66:67–76.
- . 2008b. “When Is a Work of Art Finished?” *The Journal of Aesthetics and Art Criticism* 66:395–398.
- Hinz, Carol. 2018. “Blending Nonfiction and Fiction: 4 Recent Picture Books.”
- Irmak, Nurbay. 2012. “Software is an Abstract Artifact.” *Grazer Philosophische Studien* 86:55–72.
- . 2019. “An Ontology of Words.” *Erkenntnis* 84:1139–1158.
- . 2020. “The Problem of Creation and Abstract Artifacts.” *Synthese* 1–14.
- Irvin, Sherri. 2006. “Authors, Intentions and Literary Meaning.” *Philosophy Compass* 1:114–128.
- . 2008. *New Waves in Aesthetics*, chapter “The Ontological Diversity of Visual Artworks”, 1–19. Palgrave Macmillan.
- Juvshik, Tim. 2020. “Relativity and the Causal Efficacy of Abstract Objects.” *American Philosophical Quarterly* 57:269–282.
- Kafka, Franz. 1988. *The Castle*. Schocken Books.
- Kania, Andrew. 2008. “The Methodology of Musical Ontology: Descriptivism and its Implications.” *British Journal of Aesthetics* 48:426–444.
- . 2017. “The Philosophy of Music.” *Stanford Encyclopedia of Philosophy* .
- Kharmis, Daniil. 2007. *Sonnet*. The Overlook Press, Peter Mayer Publishers, Inc.
- Kind, Amuy. 2010. “The Puzzle of Imaginative Desire.” *Australasian Journal of Philosophy* 89:421–439.

- Kivy, Peter. 1983. "Platonism in Music: A Kind of Defense." *Grazer Philosophische Studien* 19:109–129.
- . 1987. "Platonism in Music: Another Kind of Defense." *American Philosophical Quarterly* 24:245–252.
- Kripke, Saul A. 2013. *Reference and Existence*. Oxford: Oxford University Press.
- Labov, William. 1973. *Language in the Inner City: Studies in the Black English Vernacular*. University of Pennsylvania Press.
- Lamarque, Peter and Olsen, Stein Haugom. 1994. *Truth, Fiction, and Literature*. Oxford University Press.
- Letts, Philip. 2018. "The Property Theory of Musical Works." *Journal of Aesthetics and Art Criticism* 76:57–69.
- Levinson, Jerrold. 1990. *Music, Art, and Metaphysics*. Cornell University Press.
- . 2012. *Art & Abstract Objects*, chapter "Indication, Abstraction, and Individuation", 49–61. Oxford University Press.
- Lewis, David. 1968. "Counterpart Theory and Quantified Modal Logic." *The Journal of Philosophy* 65:113–126.
- . 1971. "Counterparts of Persons and Their Bodies." *The Journal of Philosophy* 68:2.
- . 1973. *Counterfactuals*. Blackwell Publishers.
- . 1983a. "Extrinsic Properties." *Philosophical Studies* 44:197–200.
- . 1983b. *Philosophical Papers Volume I*, chapter "Truth in Fiction", 261–275. Oxford: Oxford University Press.
- . 1986. *On the Plurality of Worlds*. Blackwell Publishers.
- Livingston, Paisley. 1999. "Counting Fragments, and Frenhofer's Paradox." *British Journal of Aesthetics* 39:14–23.
- . 2005. *Art and Intention: A Philosophical Study*. Clarendon Press.
- . 2008. "When a Work Is Finished: A Response to Darren Hudson Hick." *The Journal of Aesthetics and Art Criticism* 66:393–395.
- Livingston, Paisley and Archer, Carol. 2010. "Artistic Collaboration and the Completion of Works of Art." *British Journal of Aesthetics* 50:439–455.
- Lorand, Ruth. 2001. "Telling a Story or Telling a World?" *British Journal of Aesthetics* 41:425–443.

- Mag Uidhir, Christy (ed.). 2012. *Art & Abstract Objects*. Oxford University Press.
- Magnus, P. D. 2012. *Art & Abstract Objects*, chapter “Historical Individuals like *Anas Platyrhynchos* and ‘Classical Gas’”, 108–124. Oxford University Press.
- Markosian, Ned. 2010. *Time and Identity*, chapter “Identifying the Problem of Personal Identity”. MIT Press.
- Matravers, Derek. 2014. *Fiction and Narrative*. Oxford University Press.
- McDaniel, Kris. 2009. *Ways of Being*, 290–319. Oxford University Press.
- Mnookin, Seth. 2006. “Picking Up the Pieces: How James Frey flunked rehab, and why his fakery matters.”
- Moore, Joseph G. 2012. *Art & Abstract Objects*, chapter “Musical Works: A Mash-Up”, 284–306. Oxford University Press.
- Morris, Michael. 2007. *Doing Justice to Musical Works*, chapter 2, 52–78. Oxford University Press.
- Moruzzi, Caterina. 2018. “Every Performance Is a Stage: Musical Stage Theory as a Novel Account for the Ontology of Musical Works.” *Journal of Aesthetics and Art Criticism* 76:341–351.
- Nanay, Bence. 2009. “Narrative Pictures.” *The Journal of Aesthetics and Art Criticism* 67:119–129.
- Nolan, Daniel. forthcoming. “Impossible Fictions Part I: Lessons for Fiction.” *Philosophy Compass* .
- Nozick, Robert. 1981. *Philosophical Explanations*. Harvard University Press.
- Polanyi, Livia. 1981. “Telling the Same Story Twice.” *Text* 1:3.
- Predelli, Stefano. 1995. “Against Musical Platonism.” *British Journal of Aesthetics* 35:338–350.
- . 2009. “Ontologese and Musical Nihilism: A Reply to Cameron.” *British Journal of Aesthetics* 49:179–183.
- Priest, Graham. 1997. “Sylvan’s Box: a Short Story and Ten Morals.” *Notre Dame Journal of Formal Logic* 38:573–582.
- Rettler, Bradley and Bailey, Andrew M. 2017. “Object.” *Stanford Encyclopedia of Philosophy* .
- Richard, Mark. 2004. “Contextualism and Relativism.” *Philosophical Studies* 119:215–242.

- Rohrbaugh, Guy. 2003. "Artworks as Historical Individuals." *European Journal of Philosophy* 11:177–205.
- . 2012. *Art and Abstract Objects*, chapter "Must Ontological Pragmatism Be Self-Defeating", 29–48. Oxford University Press.
- . 2017. "Psychologism and Completeness in the Arts." *The Journal of Aesthetics and Art Criticism* 75:131–141.
- . 2018. "Psychologism about Artistic Plans: Reply to Cray." *The Journal of Aesthetics and Art Criticism* 76:105–107.
- Rossberg, Marcus. 2012. *Art & Abstract Objects*, chapter "Destroying Artworks", 62–83. Oxford University Press.
- Sider, Theodore. 1996. "All the World's a Stage." *Australasian Journal of Philosophy* 74:433–453.
- . 2001. *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press.
- . 2018. "Asymmetric Personal Identity." *Journal of the American Philosophical Association* 127–46.
- Smuts, Aaron. 2009. "Story Identity and Story Type." *Journal of Aesthetics and Art Criticism* 67:5–14.
- Stecker, Robert. 2009. "Methodological Questions about the Ontology of Music." *Journal of Aesthetics and Art Criticism* 67:375–386.
- Stock, Kathleen. 2016. *Imagination and Fiction*, chapter 15, 204–216. Routledge.
- . 2017. *Only Imagine: Fiction, Interpretation, and Imagination*. Oxford University Press.
- Thomasson, Amie L. 1999. *Fiction and Metaphysics*. Cambridge University Press.
- . 2003. "Speaking of Fictional Characters." *Dialectica* 57:205–223.
- . 2005. "The Ontology of Art and Knowledge in Aesthetics." *The Journal of Aesthetics and Art Criticism* 63:221–229.
- . 2006. "Debates about the Ontology of Art: What Are We Doing Here?" *Philosophy Compass* 1:245–255.
- . 2007. *Ordinary Objects*. Oxford University Press.
- . 2010. "Ontological Innovation in Art." *The Journal of Aesthetics and Art Criticism* 68:199–130.
- . 2015. *Ontology Made Easy*. Oxford: Oxford University Press.

- . forthcoming. chapter “The Ontology of Art: Keeping it Easy and Keeping it Real”, —. Oxford University Press.
- Trivedi, Saam. 2002. “Against Musical Works as Eternal Types.” *British Journal of Aesthetics* 42:73–82.
- Trogdon, Kelly and Livingston, Paisley. 2014. “The Complete Work.” *The Journal of Aesthetics and Art Criticism* 72:225–233.
- . 2015. “Artwork Completion: A Response to Gover.” *The Journal of Aesthetics and Art Criticism* 73:460–462.
- Tullman, Katherine and Buckwalter, Wesley. 2014. “Does the Paradox of Fiction Exist?” *Erkenntnis* 79:779–796.
- Velleman, J. David. 2003. “Narrative Explanation.” *Philosophical Review* 112:1–25.
- von Solodkoff, Tatjana. 2014. *Empty Representations: Reference and Non-Existence*, chapter “Fictional Realism and Negative Existentials”, 333–352. Oxford: Oxford University Press.
- Walters, Lee. 2013. “Repeatable Artworks as Created Types.” *British Journal of Aesthetics* 53:461–477.
- Walton, Kendall. 1970. “Categories of Art.” *The Philosophical Review* 79:334–367.
- . 1978. “Fearing Fictions.” *Journal of Philosophy* 75:5–27.
- . 1990. *Mimesis as Make-Believe*. Harvard University Press.
- Wetzel, Linda. 2009. *Types & Tokens*. MIT press.
- White, Eric Walter. 1961. *Stravinsky*. Pelican Books.
- Wildman, Nathan and Folde, Christian. 2017. “Fiction Unlimited.” *The Journal of Aesthetics and Art Criticism* 75:73–80.
- Wolterstorff, Nicholas. 1980. *Works and Worlds of Art*. Clarendon Press.
- Young, James O. 2002. “Making it Up: A Definition of Fiction.” *Iyyun: The Jerusalem Philosophical Quarterly* 51:23–40.