

ABSTRACT

Title of dissertation: NOVELS AND THEIR INSTANCES:
 A METAPHYSICAL EXPLORATION

Alexey Aliev
Doctor of Philosophy, 2018

Dissertation directed by: Professor Jerrold Levinson
 Department of Philosophy

What is the ontological status of novels? Are they inscriptions (i.e., concrete texts typically written or printed on something or displayed on the screen of some electronic device)? Sets of inscriptions? Mental representations of some semantic content? Structures of meanings? Syntactic sequences? Or something else? Furthermore, what is the ontological status of instances of a novel (i.e., entities that manifest all the primary properties that must be experienced to fully appreciate this novel)? Are they readings (i.e., sequences of sounds generated as a result of reading aloud)? Inscriptions? Both readings and inscriptions? Or some other entities?

My goal in this dissertation is to answer these questions.

The dissertation is structured as follows. In Part I, I provide some terminological clarifications that must be made before addressing the issues concerning the ontological status of novels and their instances. In particular, in Chapter 1 (“Defining ‘a Novel’”), I define “a novel,” and in Chapter 2 (“Defining ‘an Instance of an Artwork’”), I define “an instance of an artwork.”

Part II is aimed at clarifying the ontological status of instances of novels. I begin, in Chapter 3 (“Against Inscriptions as Instances of Novels”), by arguing against the most widely endorsed ontology of instances of novels—the ontology according to which the paradigmatic, or most typical, entities that serve as such instances are inscriptions. Next, in Chapter 4 (“An Ontology of Instances of Novels”), I put forward and defend an alternative ontology—the one according to which instances of novels are *readings* and *mereological sums of readings and graphic elements*. Finally, in Chapter 5 (“The Novel as a Performing Art”), I examine a peculiar consequence of the foregoing ontology—that the novel is a performing art.

The purpose of Part III is to clarify the ontological status of novels. I begin, in Chapter 6 (“What a Novel Is Not”), with a critical overview of the most promising existing ontologies of novels, arguing that none of these ontologies stands up completely to criticism. Then, in Chapter 7 (“An Ontology of Novels”), I expound and defend a new ontology of novels. According to this ontology, novels are a peculiar kind of concreta—namely, *concrete types composed of certain sonic, semantic, syntactic, contextualist, and visual elements*.

NOVELS AND THEIR INSTANCES: A METAPHYSICAL
EXPLORATION

by

Alexey Aliev

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2018

Advisory Committee:
Professor Jerrold Levinson, Chair/Advisor
Professor Peter Carruthers
Professor Robert Howell
Professor Brian Richardson
Professor Allen Stairs

© Copyright by
Alexey Aliev
2018

Acknowledgements

I want to thank everyone who has been involved with this project.

I am extremely grateful to Jerrold Levinson. I was fortunate to find in him an advisor who is not just a great philosopher but also a kind and generous person. Without his help and support, my experience at Maryland would not have been as successful as it has turned out to be. I am deeply indebted to him.

I am also very grateful to Robert Howell, who familiarized me with analytic philosophy and, in particular, analytic aesthetics. Had not taken his excellent courses at Moscow State University back in 2008, this dissertation would not have been written.

Thanks to the many folks who provided helpful comments on various parts of this dissertation. Among others, my thanks to Wesley Cray, Gregory Currie, David Davies, Erin Eaker, James Hamilton, Andrew Kania, Robbie Kubala, Jerrold Levinson, Robert Howell, Aidan Lyon, Anna Christina Ribeiro, Lee Walters, and audiences at the 2017 British Society of Aesthetics Annual Conference, the 2017 Eastern Division Meeting of the American Philosophical Association, the 2016 Annual Meeting of the American Society for Aesthetics, the 2016 Annual Meeting of the Canadian Society for Aesthetics, the 2015 British Society of Aesthetics Annual Conference, the 2015 Eastern Division Meeting of the American Society for Aesthetics, the 2015 Pacific Division Meeting of the American Society for Aesthetics, the 2015 Central Division Meeting of the American Philosophical Association, and the 2014 Annual Meeting of the Canadian Society for Aesthetics.

Thanks to the members of my dissertation committee: Peter Carruthers, Robert Howell, Jerrold Levinson, Brian Richardson, and Allen Stairs.

Thanks to the Department of Philosophy at Maryland. I am especially grateful to Louise Gilman for her help with various organizational matters.

Finally, many thanks to my parents, Arif and Natalia, for all the support they have given me.

Table of Contents

Acknowledgements	ii
PART I	1
1 Defining “a Novel”	3
1.1 Is It Possible to Define “a Novel”?	4
1.2 Defining “a Novel _r ”	13
1.3 Extending the Definition of “a Novel _r ”	28
1.4 Objections	40
1.4.1 Objections of Group (a)	41
1.4.2 Objections of Group (b)	43
1.5 Final Remarks	49
2 Defining “an Instance of an Artwork”	55
2.1 Defining “an Instance of an Artwork”	56
2.2 Evaluating the Definition	64
2.3 Defining “a Well-Formed Instance of an Artwork”	68
2.4 Defining “a Non-Well-Formed Instance of an Artwork”	70
2.5 Additional Remarks	76
PART II	81
3 Against Inscriptions as Instances of Novels	83
3.1 Against the Strong Orthodox View	85
3.1.1 Why Inscriptions Cannot Manifest Sonic Properties	86
3.1.2 Why the Sonic Properties of a Novel Must Be Experienced to Fully Appreciate This Novel	89
3.1.3 Why the Sonic Properties of a Novel Are Primary	110
3.1.4 Conclusion	110
3.2 Against the Strong Orthodox View: An Alternative Argument	111
3.2.1 Kivy’s Response	113

3.2.2	Other Responses	127
3.2.3	Conclusion	133
3.3	Against the Weak Orthodox View	135
3.4	Against the Orthodox View	137
3.5	Appendix	137
4	An Ontology of Instances of Novels	141
4.1	Primary Appreciation-Relevant Experienceable Properties of Novels	142
4.2	The Ontological Status of Well-Formed Instances of Novels	146
4.2.1	Readings as Well-formed Instances of Non-Visual Novels	147
4.2.2	Mereological Sums of Readings and Graphic Elements as Well-Formed Instances of Visual Novels	164
4.2.3	Conclusion	171
4.3	The Ontological Status of Non-Well-Formed Instances of Novels	172
4.4	Objections	173
4.5	Some Consequences	187
4.6	Final Remarks	195
5	The Novel as a Performing Art	197
5.1	Urmson's and Kivy's Arguments: A Critical Analysis	198
5.1.1	Urmson's Argument	198
5.1.2	Kivy's Argument	203
5.2	The Novel as a Performing Art	207
5.2.1	What Is a Performing Art?	207
5.2.2	What Are Performances of an Artwork?	215
5.2.3	Readings as Performances	224
5.2.4	Conclusion	228
5.3	The Novel as a Performing Art and the History of Literature	230
PART III		234
6	What a Novel Is Not	236
6.1	Abstractionist Views	239
6.1.1	A Novel as a Set of Embodiments	239
6.1.2	A Novel as a Property	243
6.1.3	A Novel as a Pure Type	244
6.1.4	A Novel as an Initiated Type	257
6.1.5	A Novel as a Historical Individual/an Abstract Artifact	291
6.2	Concretist Views	293
6.2.1	A Novel as a Particular Embodiment	294
6.2.2	A Novel as a Mereological Sum of Embodiments	297
6.3	Eliminativism	300
6.4	Toward a Satisfactory Ontology of Novels	306

7	An Ontology of Novels	308
7.1	A Novel as a Concrete Type	308
7.2	Some Objections	319
7.3	The Essential Elements of the Concrete Type to Which a Novel Is Identical	335
7.4	The Ontological Status of a Novel	342
	Bibliography	343

PART I

Introduction to PART I

The main focus of this dissertation is the question “What is the ontological status of novels and their instances?” Before tackling this question, however, it is necessary to clarify what is meant by the expressions “a novel” and “an instance of an artwork.” Given this, my goal in PART I is to define these expressions.

Chapter 1: Defining “a Novel”

Introduction

In this chapter, I aim to define “a novel.” I begin by arguing that there is no real reason to consider “a novel” indefinable (Section 1.1). Next, I define “a novel” in the restricted sense, or, in other words, “a novel” that refers to any novels except the so-called nonfiction novels, novels in verse, and graphic novels (Section 1.2). Then I provide definitions of “a nonfiction novel,” “a novel in verse,” and “a graphic novel”—and, with the help of these definitions as well as the definition of “a novel” in the restricted sense, define “a novel” (Section 1.3). After that, I examine potential objections to the definition of “a novel” provided in the previous section, arguing that none of them stands up completely to criticism (Section 1.4). Finally, I make some remarks concerning the foregoing definition and suggest that “a novella” and “a short story” can be defined in a similar way (Section 1.5).

1.1 Is It Possible to Define “a Novel”?

As is clear from what has been said, the goal of this chapter is to define “a novel.”¹ Can this goal be achieved? There are a number of potential reasons to answer this question in the negative. Let us have a look at these reasons and determine whether any of them is successful.

Here is one potential reason against the possibility of defining “a novel.” There have been a considerable number of serious attempts to define “a novel.” Yet none of these attempts can be considered successful. Meanwhile, if there have been a considerable number of serious attempts to do x , all of which have failed, then doing x is likely to be impossible.

Another potential reason against the definability of “a novel” can be formulated in the following way. Defining “a novel” is possible only if the concept expressed by “a novel” is structured according to conditions satisfied by all entities that actually fall under “a novel” and only by such entities (hereafter: “necessary and sufficient conditions”). However, there are at least two strong considerations against the idea that this concept is structured that way. First, if this idea were

¹As is generally agreed, to define x (where x is some expression) is to explicate the meaning of x by specifying a set of conditions that are (a) satisfied by *all* entities that actually fall under x and *only* by such entities and (b) sufficiently informative—in particular, (1) are not enumerative, or, in other words, do not amount to a (disjunctive) list of objects that fall under x ; (2) do not involve a vicious circle, or, in other words, do not explicitly or implicitly contain the concept expressed by x ; (3) do not contain meaningless expressions or expressions whose meanings cannot be understood; etc.

true, then it would be relatively easy to come up with a satisfactory definition of “a novel.” However, given that, as mentioned in the previous paragraph, none of the attempts to define “a novel” has been successful, defining “a novel” is hard, if possible at all.²

Second, if the concept of “a novel” were structured according to necessary and sufficient conditions, then we would use such conditions when categorizing objects as “novels.” But, in fact, when we categorize objects as “novels,” we do not use necessary and sufficient conditions; we use other, non-definitional classificatory means—such as prototypes³ (mental representations of a novel in general) and

²It is worth noting that in this respect, “a novel” is not unique. Defining other artistic expressions—such as “art,” “poetry,” “music,” and “literature”—seems very hard, if not impossible. Moreover, it is often hard to define even “ordinary” expressions. Consider, for example, “a bachelor.” According to a common definition—call it D_{b1} —a bachelor is an unmarried man of a marriageable age. But this definition is problematic. If it is true, then the Pope is a bachelor. But he is not a bachelor—or so it seems.

According to another possible definition—call it D_{b2} —a bachelor is an unmarried man of a marriageable age who has a right to marry. This definition is also problematic. Although D_{b2} avoids the problem of D_{b1} (the Pope does not have a right to marry and so, according to D_{b2} , cannot be a bachelor), it faces another problem. Consider an unmarried man who has lived with his girlfriend for 20 years. Is he a bachelor? D_{b2} entails that he is. But this result is counterintuitive. It seems wrong to call someone who has lived with his or her sexual partner for 20 years “a bachelor.”

Perhaps there is a way to modify D_{b1} and D_{b2} so that the foregoing problems could not arise. But it is not easy to figure out what this way is. And there seems to be no other potentially acceptable definition of “a bachelor.” As a result, it is unclear how this expression can be defined.

³See Rosch (1978).

exemplars⁴ (mental representations of some particular novel or novels considered paradigmatic or most typical).⁵

Here is yet another potential reason against the possibility to define “a novel.” As Weitz (1956) points out, if a concept is open—that is, if the application conditions of this concept are emendable and corrigible—then the expression corresponding to this concept is indefinable. For suppose an expression that corresponds to a concept is definable. Then the application conditions of this concept are necessary and sufficient. But such conditions are neither emendable nor corrigible, since there can be no entity x that would legitimately require us to modify the concept being discussed by adding a condition that either prevents x from falling under this concept or ensures that x is covered by this concept. Thus, in the case of a definable expression, the application conditions of the corresponding concept must not be emendable or corrigible. Meanwhile, as already mentioned, in order for a concept to be open, its conditions of application must be emendable and corrigible. So the concept corresponding to a definable expression must be closed.

Thus, the openness of a concept entails that the expression corresponding to this concept cannot be defined. But if that is the case, then “a novel” is indefinable. For suppose that the openness of a concept does, in fact, entail that the expression corresponding to this concept cannot be defined. Then “a novel” can be defined only if the concept of “a novel” is not open. But this concept is open. Consider

⁴See Smith and Medin (2002).

⁵Dean (2003) offers a similar reason to show that the concept of “art” is not structured according to necessary and sufficient conditions.

the art of the novel. It is aimed at generating something substantially new and, hence, is essentially creative. Meanwhile, if a practice is essentially creative, then the application conditions of the concept corresponding to a product of this practice are emendable and corrigible.⁶ Thus, the application conditions of the concept of “a novel” are emendable and corrigible. But if this is so, then this concept is open.⁷

Are any of the foregoing reasons against the possibility of defining “a novel” successful? Consider the first reason. It assumes that if there have been a considerable number of serious attempts to do x , none of which has been successful, then doing x is likely to be impossible. Yet this assumption is questionable. Furthermore, according to the first reason, there have been a considerable number of serious attempts to define “a novel.” However, so far as I am aware, there have actually been no such attempts.⁸ (This is not to say, of course, that no one has tried to

⁶Suppose a practice is essentially creative. Then “new cases can always be envisaged or created... which would call for a decision on someone’s part to extend or to close” (Weitz, 1956, 413) the concept corresponding to a product of this practice. Put another way, it is always possible to create or merely imagine an entity x that would legitimately require us to modify the concept corresponding to a product of this practice by adding a condition that either prevents x from falling under this concept or ensures that x is covered by this concept. Meanwhile, if creating or imagining such an entity is possible, then the application conditions of the concept corresponding to a product of the practice under consideration are emendable and corrigible.

⁷This objection to the thesis that “a novel” is definable is due to Weitz (1956).

⁸The lack of interest concerning the definition of “a novel” is puzzling, especially given that there is considerable interest with regard to some closely related expressions such as “literature” and “poetry” (for definitions of “literature,” see, e.g., Beardsley (2004), Lamarque (2009), Ohmann (1971), and Stecker (2004); for a definition of “poetry,” see Ribeiro (2007)).

specify some characteristic, though non-defining, features of novels: A number of attempts to do this have been made by philosophers, literary theorists, historians of literature, literary critics, and authors of dictionary articles concerning the novel. The point is that no one has tried to *define* (in the sense specified above) “a novel.”)

In light of what has been said, the first reason against the possibility of defining “a novel” can hardly be considered successful. What about the second reason? It can be successful *only if* at least one of the considerations advanced to support the thesis that the concept of “a novel” is not structured according to necessary and sufficient conditions is satisfactory. Is any of these considerations, in fact, satisfactory?

According to the first consideration, the fact that defining “a novel” is hard shows that the concept of “a novel” is not structured according to necessary and sufficient conditions. But does this fact really show that? Consider “gold.” Prima facie, if this expression were easy to define, it would be definable with the help of “ordinary,” easily available information—that gold is a precious, yellow metal, is used to create things like rings and necklaces, serves as the monetary standard, etc. But, as Rey (1999)—following Kripke (1980) and Putnam (1970, 1975)—points out, such information “is not in any way necessarily tied to [gold]” (291) and, hence, cannot be used to define “gold.” So defining “gold” is not easy. However, this does not entail that the concept of “gold” is not structured according to necessary and sufficient conditions. Consider the following description: “the basic chemical element with atomic number 79.” According to Rey (1999),

this... description, if true, would appear to provide us with nomologi-

cally necessary and sufficient conditions for the proper use of the term [“gold”], those facts in virtue of which something qualifies as being (made of) gold. And it would also seem to capture our *metaphysically* modal intuitions: again, if the description is *true*, then only something having atomic number 79 *would* be gold; even something yellow, mined in the Sierras, used in wedding bands and as a monetary standard, but lacking that atomic number, *would not* be. In short, the description would appear to provide much that has been standardly asked of a definition.

(291–292)

Thus, if Rey (1999) is right—and there seems to be nothing that would suggest otherwise—there is good reason to hold that the description “the chemical element with atomic number 79,” if true, can be used to define “gold.” Meanwhile, if “gold” is definable, then it is reasonable to maintain that the concept of gold is structured according to certain necessary and sufficient conditions.

Thus, it is highly questionable that the fact that it is hard to define “gold” shows that the concept of gold is not structured according to necessary and sufficient conditions. But if this is so, then there is good reason to question the thesis that in the case of “a novel,” the analogous fact shows that this expression is not structured according to certain necessary and sufficient conditions.

Let us now turn to the second consideration against the thesis that the concept of “a novel” is structured according to necessary and sufficient conditions. According to this consideration, the fact that we do not categorize novels with the help of nec-

essary and sufficient conditions shows that the concept of “a novel” is not structured in accordance with such conditions. But does this fact really show that? As Rey (1999) points out, a concept has two functions: the metaphysical and the epistemological. The metaphysical function consists in specifying the facts by virtue of which something is *correctly* categorized by the concept. The epistemological function, on the other hand, consists, mainly, in providing “the means by which an agent categorizes things, decides whether or not something is of a certain kind” (Rey, 1999, 282).⁹ Taking this into account, suppose that the fact that we do not categorize novels by means of necessary and sufficient conditions shows that the concept of “a novel” is not structured according to such conditions. Then the metaphysical

⁹It is worth noting that the distinction between the metaphysical and the epistemological functions

... corresponds to a crucial, if very battered, distinction in philosophy between *metaphysics* and *epistemology*, or between issues surrounding *how the world is* (what exists, what is true) and issues surrounding *how we know, believe, infer, how the world is*. Although this [latter] distinction is not everywhere perfectly sharp (e.g., in describing our own cognitions), and despite the fact that some people (e.g., various sorts of relativists and idealists) are inclined to *argue* that the distinction is ultimately only apparent, it should seem on its face pretty plausible: there is, after all, all the difference in the world between the issue of *whether there actually is a cow on the road* and the issue of *whether anyone knows, believes, has inferred, or even cares whether there is*. Similarly, then, there would seem to be all the difference in the world between *something being a cow* and *someone knowing, believing, or inferring that it is*. (Rey, 1999, 284)

function of the concept of “a novel” must coincide with the epistemological function of this concept. For suppose that is not the case. Then the metaphysical function can employ the means other than the facts about the way we actually categorize novels, in particular it can employ certain necessary and sufficient conditions. But if this is so, then, since this function is performed by the concept of “a novel,” this concept can be structured according to necessary and sufficient conditions.

Thus, if the fact that we do not categorize novels by means of necessary and sufficient conditions shows that the concept of “a novel” is not structured according to such conditions, then the metaphysical function of this concept must coincide with the epistemological function of this concept. But does the metaphysical function, in fact, coincide with the epistemological one? If it does, then the concept of “a novel” is determined by how we actually categorize novels. However, according to Rey (1999, 288), there is good reason to think that this concept is not determined by that. The ways people categorize novels (as well as other things) differ from person to person. One might categorize novels with the help of a particular set of properties, whereas someone else might do that using a different set of properties or perhaps something entirely different—say, a mental representation of an “exemplar” novel. So if the concept of “a novel” is determined by our actual categorization practices, then each of us has a concept of “a novel” that, in an overwhelming majority of cases, is different, and sometimes *very* different, from a concept of “a novel” possessed by another person. Meanwhile, a genuinely meaningful conversation about x is possible only if the participants of this conversation have the same, or at least very similar, concepts of “ x .” As a result, if the concept of “a novel” is determined by how we

categorize novels, then meaningful conversations about novels must be very rare. But, of course, such conversations are not rare. So the concept of “a novel” is not determined by how we categorize novels.

Thus, the metaphysical function does not coincide with the epistemological one. But then, given what has been said above, the fact that we do not categorize novels using necessary and sufficient conditions does not show that the concept of “a novel” is not structured according to such conditions.

So neither consideration against the thesis that the concept of “a novel” is structured according to necessary and sufficient conditions is satisfactory. And there seems to be no other potentially satisfactory considerations against this thesis. Given this, the second reason against the possibility of defining “a novel” cannot be considered successful.

Let us now turn to the third reason against this possibility. This reason assumes that if a practice is creative, then the concept corresponding to a product of this practice is open. Is this assumption true? Suppose the answer is “Yes.” Then the game of chess is creative only if the concept of “a game of chess” is open. But this concept is closed. Given the fact that the game of chess is strictly determined by certain rules, there can be no x that would legitimately require us to modify the concept of “a game of chess” by adding a condition that either prevents x from falling under this concept or ensures that x is covered by this concept. So the application conditions of the concept of “a game of chess” are neither emendable nor corrigible. Meanwhile, any concept whose application conditions are neither emendable nor corrigible is closed.

Thus, if the assumption being discussed is true, then the game of chess is not creative. But this consequence is clearly false, as games of chess can be creative—for example, by virtue of involving particular creative moves, series of moves, or strategies.¹⁰ So the foregoing assumption is false. But then the third reason against the possibility of defining “a novel” fails.

So none of the foregoing reasons shows that “a novel” is indefinable. Meanwhile, so far as I am aware, there are no other potentially plausible reasons against the possibility of defining “a novel.” Given this, it is reasonable to hold that “a novel” can be defined.

1.2 Defining “a Novel_r”

Let us now turn to the main task of this chapter—defining “a novel.”

Call any novel that is not a nonfiction novel, a novel in verse, or a graphic novel “a novel in the restricted sense” (hereafter: “novel_r”). How can “a novel_r” be defined?

To answer this question, let us consider the characteristic features of a novel_r. One of these features is that a novel_r is a verbal object—that is, an object composed, for the most part, of linguistic elements, such as words and punctuation marks. Note that being a verbal object does not require being composed *solely* of linguistic elements;¹¹ so the feature being discussed does not presuppose that a novel_r cannot

¹⁰This objection is due to Davies (1991).

¹¹Thus, W. G. Sebald’s *Vertigo* is a verbal object, although some of its essential elements—in particular, drawings and photographs—are non-linguistic.

involve non-linguistic entities—photographs, maps, diagrams, drawings, and the like.

Another characteristic feature of a novel_r is that such a novel is written, for the most part, in prose. Note that this feature does not imply that a novel_r cannot contain verse or some other non-prosaic textual elements—for being written, for the most part, in verse is compatible with not being *completely* prosaic.¹²

Yet another characteristic feature of a novel_r is that such a novel has an appropriate length. Here, of course, a natural question arises: What exactly is this length? To answer this question, it is sufficient to answer the following questions: (a) “What is the *maximal* number of words a novel_r can have?” and (b) “What is the *minimal* number of words a novel_r can have?”¹³ Let us first answer the former question.

A lot of novels_r have about 100,000 words. At the same time, there are novels_r that are much longer. Samuel Richardson’s *Clarissa* has over 950,000 words, Marcel Proust’s *À la recherche du temps perdu* exceeds a million words, Madeleine de Scudéry’s *Artamène ou le Grand Cyrus* contains about 2.1 million words, and Mark Leach’s *Marienbad My Love* has approximately 17 million (!) words. Can a novel_r be even longer than that? There seems no reason to think otherwise. It is even possible

¹²An example of an object that is written, for the most part, in prose but contains verse and, hence, is not completely prosaic is R. L. Stevenson’s *Treasure Island*.

¹³Here and in what follows, a novel’s having x words is understood as (a) the possession of x words by this novel—if it is written in English, *or* (b) the possession of x words by its (correct) English translation—if it is written in a language other than English.

to conceive of a novel_r that contains an infinite loop and, hence, is never-ending.¹⁴ Given this, the length of a novel_r does not have an upper word limit.

What about the lower word limit? What is the minimal number of words a novel_r can have? Answering this question precisely—by specifying an exact number of words—is hardly possible. The reason for this is that any attempt to do this faces the paradox similar to the paradoxes of the heap and the bald man. Suppose we find out that the shortest novel_r ever written¹⁵ has 34,381 words (Julie Otsuka’s novel_r *When the Emperor Was Divine* is about that long). It seems bizarre to say that anything that is just one word shorter than this novel_r is not a novel_r. So a novel_r can contain 34,380 words. But taking one word out of this 34,380 word novel will not turn it into a non-novel_r. So it is possible for a novel_r to have 34,379 words. But, again, if we remove just one word from a 34,379 word novel_r, we will not turn this novel into a non-novel_r. Thus, a novel_r can have 34,378 words. But, again... We can continue this reasoning until we reach an absurd claim—that a novel_r can

¹⁴One might object that a never-ending novel is not really a novel, since (a) it is unreadable and (b) being a novel presupposes being readable. But this objection fails. It assumes that a never-ending novel cannot be read. But this assumption is false: Although a never-ending novel cannot be read *completely*, it can be read.

One might also object that a never-ending novel is not really a novel, since (a) it cannot be read in its entirety and (b) something can be a novel only if it can be read in its entirety. But this objection also fails. It assumes that every novel can be read in its entirety. But there seems no real reason to accept this assumption. *Prima facie*, there is nothing wrong with there being novels that cannot be fully read.

¹⁵In this chapter, the word “to write” is used in a broad sense—the sense according to which to write is to generate a text.

have just one word. The absurdity of this claim suggests that there must exist a limit to how short a novel_r can be. But it seems that, regardless of what we take this limit to be, we face the paradox just sketched. So how should we proceed?

I think we should admit that since the limit to how short novels_r can be is (at least, epistemically) essentially vague, it cannot be precisely identified. At the same time, we can try to specify it approximately. Consider an object that is identical to a 10 word sentence. Clearly, such an object cannot be a novel_r. Furthermore, it seems odd to hold that a novel_r can have 5000 or fewer words. Likewise, we, most likely, will not call something that has about 10,000 words “a novel_r.” Now, what about an object that contains 15,000 words? Can it be a novel_r? The answer to this question is less obvious than the answer to the analogous question about a 10,000 or 5000 word object. But it seems that in this case, saying “No” is closer to the truth than saying “Yes.” An object having approximately 15,000 words can be a short story or a novella. But it can hardly be “a novel_r.” Given what has been said, the border that separates a novel_r from a non-novel_r is, I think, in the 30,000 word range. There are objects with a word count in this range that are widely recognized as novels_r—for instance, Albert Camus’s *The Stranger* (c. 36,500 words), E. L. Konigsburg’s *From the Mixed Up Files of Mrs. Basil E. Frankweiler* (c. 31,300 words), Julie Otsuka’s *When the Emperor Was Divine* (c. 34,300 words), and E. B. White’s *Charlotte’s Web* (c. 33,500 words). At the same time, so far as I am aware, there is no novel_r whose word count is in the 20,000 or some lower word range.¹⁶

¹⁶The foregoing substantiation of the thesis that the border separating a novel_r from a non-

Thus, taking into account what has been said, if something is a novel_r, it must (a) be a verbal object, (b) be written, for the most part, in prose, and (c) have at least n words, where n is a number in the 30,000 word range. The converse, however, is not true. Being an n or $n+$ word verbal object written, for the most part, in prose is not sufficient for being a novel and *a fortiori* a novel_r. Consider, for instance, historical monographs. They are typically n or $n+$ word verbal objects written in prose, but they are not novels. And the same can be said about an overwhelming majority of dictionaries, collections of short stories, textbooks, and philosophical treatises. So what property distinguishes novels_r from the foregoing non-novelistic works?

One possible answer draws upon Monroe Beardsley's account of "a literary work." On this account, "a literary work is a discourse in which an important part of the meaning is implicit" (Beardsley, 1981, 126), where "implicit meaning" can be understood as referring to any of the following types of meaning:

- meaning suggested by the use of a word (for example, the idea that a speaker is a member of a religious group suggested by the fact that she uses the word "thee" in her speech);
- a connotation of a word, that is, a meaning that is "commonly associated in a novel_r is in the 30,000 word range is doubtless limited (in order for a substantiation of this thesis to be complete, it must be based on the results of a comprehensive empirical investigation into the issue of what novels are considered the shortest ones and how many words they have). Because of this, this thesis should be treated as *defeasibly* true (not as true *simpliciter* or *absolutely* true) and, hence, as open in principle to revision.

particular society with the object referred to by the word in its primary use” (Lyas, 1969, 86) (for example, the connotation “promise of adventure” of the word “sea”);

- meaning suggested by the context of an utterance (for example, the idea that a person is very tall suggested by calling someone who is very tall “little”);
- meaning suggested by the word order of a sentence (for example, the idea that a person has been sought for suggested by his saying “Here I am,” rather than “I am here”);
- meaning that depends upon sentential ambiguity (for example, the idea expressed by the final line of Wilfred Owen’s “Anthem for Doomed Youth”—“And each slow dusk a drawing down of blinds”—which is ambiguous between “each dusk householders draw down their blinds in memory of the fallen” and “each dusk is itself a vast natural drawing down of blinds in mourning for the fallen” (Lyas, 1969, 86));
- meaning suggested by the content of an utterance (for example, the idea that a person “believes the door to be open and that he wishes to have the door shut” (Lyas, 1969, 89) suggested by his saying “Shut the door!”);
- meaning suggested by the way in which a work is written (for example, the idea that the narrator is concerned about “the plight of the cotton pickers in the Deep South” (Lyas, 1969, 89) suggested by the fact that his description of

this plight is very detailed).¹⁷

Now, in light of Beardsley's account of "a literary work," one could answer the question posed above by saying that what distinguishes novels_r from non-novelistic works like historical monographs, dictionaries, collections of short stories, textbooks, and philosophical treatises (hereafter: "non-novelistic works") is that unlike such works, novels_r have a high level of implicit meaning.

Is the foregoing answer satisfactory? In order for it to be satisfactory, two conditions must be satisfied. First, novels_r must have a high level of implicit meaning. Second, non-novelistic works must have a moderate or low level of such meaning. Are these conditions satisfied? Consider E. B. White's novel *Charlotte's Web*. It is quite semantically transparent—at least, the level of its semantic transparency does not differ much from the level of the semantic transparency of ordinary (day-to-day) discourse. Given this, it seems wrong to say that *Charlotte's Web* has a *high* level of implicit meaning. And the same, I think, can be said about many (though, of course, not all) other children's novels_r, as well as some novels intended for adults. Thus, the first condition is not satisfied.¹⁸ What about the second condition? There

is good reason to think that it is not satisfied either. Consider Kant's *Critique of*

¹⁷See Lyas (1969).

¹⁸Beardsley (2004) seems to, at least partially, agree with this. In his view, the criterion of having a high level of implicit meaning:

...will cover a great deal of prose fiction, including those novels and short stories that are most worthy of attention from a literary point of view. But it will not cover all prose fiction.(Beardsley, 2004, 54)

Pure Reason. It is not a novel. But since it involves a considerable number of passages that can have different, often mutually incompatible, interpretations, it is semantically dense and so may well be characterized as having a high level of implicit meaning.

Thus, the property of having a high level of implicit meaning cannot be used to distinguish novels, from non-novelistic works. So the answer suggested above is unsatisfactory.

Another possible answer to the question of what distinguishes novels, from non-novelistic works is based on the account of “a literary work” offered by Ohmann (1971). According to Ohmann (1971), in the case of a literary work, the conditions for performing felicitous (successful) illocutionary speech acts^{19,20} are not satisfied,

¹⁹An illocutionary speech act is, roughly, an act of stating (asserting), asking a question, giving an order, promising, etc. For a detailed account of such an act, see Austin (1962), Ohmann (1971), and Searle (2012).

²⁰The nature of the conditions for performing a felicitous speech act depends on the kind of this act. Here are the conditions that must be satisfied in order for an act of stating to be performed felicitously:

1. A declarative sentence expressing the statement must be uttered.
2. The utterer must be the right person to make the statement.
3. The statement must be made correctly.
4. The statement must be complete.
5. If the statement is designed for use by persons having certain thoughts or feelings, then it must reflect the beliefs of the utterer.
6. If the statement is designed for the inauguration of certain consequential conduct on the

and, as a result:

A literary work is a discourse whose sentences lack the illocutionary forces that would normally attach to them. Its illocutionary force is mimetic. . . . A literary work purportedly imitates (or reports) a series of speech acts, which in fact have no other existence. (Ohmann, 1971, 14)

Put another way, in Ohmann (1971)'s view, a literary work is a discourse in which none of the sentences expresses illocutionary speech acts that are made felicitously and, hence, all the sentences have purely mimetic, or imitative, illocutionary forces.

It is worth stressing that Ohmann (1971)'s account does not entail that a literary work does not contain sentences expressing statements or other illocutionary speech acts. In his view, literary works can—and in most cases, do—contain such sentences. His point is that in a literary work, no sentence expresses a *felicitous* illocutionary speech act, or a speech act that has any real illocutionary force. Thus, the first sentence of *Pride and Prejudice*—"It is a truth universally acknowledged, that a single man in possession of a good fortune must be in want of a wife"—does not felicitously state that it is a truth universally acknowledged, that a single man in possession of a good fortune must be in want of a wife; rather, it merely imitates (or reports) stating that felicitously.

Now, in light of Ohmann (1971)'s account of a literary work, one could answer the question being discussed by saying that what distinguishes novels_r from non-

part of any participant (the utterer or the addressee), then the utterer must intend to act according to this statement. (Ohmann, 1971, 11)

novelistic works is that unlike the latter, novels_r are composed of sentences that (a) do not express felicitous speech acts and, hence, (b) lack genuine illocutionary forces.

Like the previous answer, this answer cannot be accepted. First of all, it assumes that non-novelistic works are composed of sentences with real illocutionary forces. However, this assumption is questionable. Recall that non-novelistic works include collections of short stories. So if non-novelistic works are composed of sentences with real illocutionary forces, then collections of short stories must be composed of such sentences. But this result is false. If novels_r do not involve sentences that have actual illocutionary forces, then it is reasonable to hold that the same must be true with regard to collections of short stories. Furthermore, if Ohmann (1971)'s account is correct, then such collections do not contain sentences with real illocutionary forces. (Of course, a proponent of the answer being discussed could respond by restricting this account to novels_r. However, for this response to be acceptable, it is necessary to explain why this restriction is justified, and there seems no satisfactory way to do that.)

In addition, contrary to what the answer based on Ohmann (1971)'s account implies, some novels_r do contain sentences with real illocutionary forces. Consider, for instance, *War and Peace*. It involves the following sentences:

Man lives consciously for himself, but serves as an unconscious instrument for the achievement of historical, universally human goals. An action once committed is irrevocable, and its effect, coinciding in time

with millions of actions of other people, acquires historical significance.

(Tolstoy, 2007, 605)

These sentences have real illocutionary forces if the illocutionary act corresponding to them—the act of stating—was performed felicitously. Was this act, in fact, performed felicitously? The answer to this question is “Yes” if the following conditions are satisfied:

1. A declarative sentence expressing the statement must be uttered.
2. The utterer must be the right person to make the statement.
3. The statement must be made correctly.
4. The statement must be complete.
5. If the statement is designed for use by persons having certain thoughts or feelings, then it must reflect the beliefs of the utterer.
6. If the statement is designed for the inauguration of certain consequential conduct on the part of any participant (the utterer or the addressee), then the utterer must intend to act according to this statement.²¹

Are these conditions satisfied? The sentences quoted above are declarative, express a particular statement, and are uttered (by the narrator²²). Hence, condition (1)

²¹See Ohmann (1971, 11).

²²A narrator—one who tells a story in a literary work—is not necessarily the author of this work. Thus, the thesis that the quoted sentences are uttered by the narrator does not necessarily imply that they are uttered by Tolstoy.

is satisfied. Condition (2) is also satisfied, as there is no real reason to think that Tolstoy is not the right person to make the statement being discussed. Furthermore, since this statement is made correctly and complete, conditions (3) and (4) are satisfied as well.

Does the statement being discussed satisfy condition (5)? This condition states that if a statement is designed for use by persons having certain thoughts or feelings, then it must reflect the beliefs of the utterer. The statement being discussed is clearly designed for use by persons having certain thoughts. So condition (5) can be satisfied just in case this statement reflects the beliefs of the utterer. Does it, in fact, reflect these beliefs? It is made by the narrator, and not one of the characters. Meanwhile, the consensus is that in *War and Peace*, the narrator is the author himself. Given this, there is good reason to think that the statement being discussed is made by Tolstoy, which, in its turn, suggests that this statement reflects his beliefs. As a result, it is reasonable to hold that condition (5) is satisfied.

Finally, what about condition (6)? Is it satisfied? There is nothing to suggest that by making the statement being discussed, Tolstoy intended to cause anyone to act in some way. Meanwhile, if he did not intend that, then the foregoing condition is satisfied.

Thus, conditions (1)–(6) are satisfied. As a result, the act of stating corresponding to the sentences quoted above was performed felicitously. But then, given what has been said above, these sentences have real illocutionary forces.

Yet another possible answer to the question being discussed is this: What distinguishes novels, from non-novelistic works is that unlike the latter, novels, were

seriously intended, by their authors, to tell a fictional story. This answer seems more promising than the previous two answers. Yet it is not unproblematic. In order for it to be satisfactory, there must be no work that (a) is a verbal object, (b) is written in prose, (c) has at least n words, and (e) was seriously intended to tell a fictional story, but (f) is not a novel_r. However, there can be such a work. Suppose Mary has written a 50,000 word textbook in biology. Suppose also that, as a result of some psychological aberration, when working on this textbook, she seriously intended to tell a fictional story. Now, Mary's textbook is a verbal object that is written in prose, has at least n words, and was seriously intended to tell a fictional story. But this textbook is not a novel and *a fortiori* not a novel_r.

Or suppose John has written a 100,000 word prosaic text that is *completely* meaningless (it is just a collection of arbitrarily chosen words). Suppose also that when working on this text, he seriously intended to tell a fictional story. Is John's text a novel_r? Prima facie, the answer is "No." A completely meaningless text is not a novel_r. At the same time, John's text is a verbal object that is written in prose, has at least n words, and was seriously intended to tell a fictional story.

Thus, the foregoing answer to the question "What distinguishes novels_r from non-novelistic works?" cannot be accepted. However, there is a way to modify this answer so that it would be acceptable. The reason why Mary's and John's works are not novels_r is that these works do not, in fact, tell any fictional stories. If Mary's and John's works did tell such stories at least to some extent, these works could be considered novels_r. Given this, we can make the foregoing answer acceptable by adding to it the claim that the author's intention to tell a fictional story is realized

at least to some extent. Thus, the acceptable answer to the question being discussed is as follows: What distinguishes novels_r from non-novelistic works is that unlike the latter, novels_r (a) were seriously intended (hereafter: “intended”), by their authors, to tell fictional stories and (b) realize this intention at least to some extent.²³

Regarding the foregoing answer, three remarks are worth making. First, it should be underlined that the mentioned intention is the intention to tell a fictional *story* (a unified sequence of events), not *stories* (several relatively autonomous sequences of events). Second, the story that the author intends to tell does not have to be *completely* fictional; it may involve real people and objects and describe events that have actually occurred.²⁴ Finally, the intention to tell a fictional story does not necessarily involve the intention to present this story as fictional. Put otherwise, when one intends to tell a fictional story, one does not necessarily intend to make it explicit that the story is fictional.²⁵

²³One might ask: What exactly is meant by “fictional”? Answering this question is beyond the scope of this dissertation. For possible answers, see, e.g., Currie (1985), Searle (1975), and Walton (1990).

²⁴This reflects the fact that not all stories of novels_r were intended, by their authors, to be *completely* fictional. To see that this fact actually holds, consider, for instance, the story of *War and Peace*. Given that this story contains descriptions of real people (Napoleon, Kutuzov, Alexander I), objects (Moscow, Saint Petersburg, Borodino), and events (the Patriotic War of 1812), it clearly was not intended, by Tolstoy, to be completely fictional.

²⁵This corresponds to the fact that not all stories of novels_r were intended, by their authors, to be presented as fictional. That this fact actually holds can be shown as follows. Consider the story of *Robinson Crusoe*. When Daniel Defoe was working on this story, he doubtless had the intention to tell a particular story. And there is little doubt that he intended to tell a *fictional* story, for, of

Given what has been said, a novel_r has the following characteristic features: (a) being a verbal object, (b) being written, for the most part, in prose, (c) having at least n words, where n is a number in the 30,000 word range, (d) being intended, by its author(s), to tell a fictional story, and (e) realizing this intention at least to some extent. Yet having these features is still not sufficient for being a novel_r. For consider screenplays and theatrical scripts. Of course, they normally have fewer than 20,000 words and, hence, cannot be novels. But nothing stops us from imagining a screenplay or a theatrical script that has 30,000 or more words. Furthermore, we can imagine that this screenplay/theatrical script (a) is a verbal object, (b) is written for the most part, in prose, (c) was intended, by its author, to tell a fictional story, and (d) realizes this intention to some extent. So if having the mentioned features were sufficient for being a novel_r, then some screenplays/theatrical scripts would be novels_r. But, surely, neither screenplays nor theatrical scripts are novels and *a fortiori* novels_r.

Is there a feature that can be used to distinguish novels_r from screenplays and theatrical scripts? The function of screenplays and theatrical scripts is to provide instructions or guidelines on how to create certain works of art (namely, films—in the case of screenplays; and theatrical performances—in the case of theatrical

course, he realized that the story about Robinson Crusoe does not describe any real events. (This is not to say, of course, that this story is completely unrelated to real events. As is well known, it was inspired by what actually happened to Alexander Selkirk.) But in his description, he did not intend to present the story as a fictional one; he “sought to give the impression of telling the truth” (Lamarque, 2009, 15), which is evidenced by the absence of an explicit mention that the story is fictional and the overt realism of the story itself.

scripts). Thus, neither screenplays nor theatrical scripts are what might be called “self-standing,” that is, intended to serve as *primary* foci of appreciation, and not merely as means to create such foci. But that is not how things are in the case of novels_r. In light of this, we can say that what distinguishes novels_r from screenplays and theatrical scripts is that unlike the latter, novels_r are *self-standing* (in the sense defined above).

Now, taking into account what has been said, “a novel_r” can be defined as follows:

Novel_r: For all x , x is a novel_r if and only if x is a verbal object that (a) is written, for the most part, in prose, (b) has at least n words, where n is a number in the 30,000 word range, (c) was intended, by its author(s), to tell a fictional story, (d) realizes this intention at least to some extent, and (e) is self-standing.

1.3 Extending the Definition of “a Novel_r”

The foregoing definition of “a novel_r” is rather powerful—it covers an overwhelming majority of novels. Yet it does not cover all novels. In particular, it does not cover graphic novels, novels in verse, and nonfiction novels (hereafter: “non-standard novels”). So it cannot serve as a definition of “a novel” *simpliciter*. But can’t it be modified so that it could serve as such a definition?

Let us begin to answer this question by examining non-standard novels. Consider first novels in verse. Like novels_r, novels in verse are verbal objects. Furthermore, there is good reason to think that novels in verse, similar to novels_r, have at

least n words. Next, like novels_{*r*}, novels in verse are intended, by their authors, to tell a story. Moreover, like in the case of novels_{*r*}, in the case of novels in verse, the authors' intention to tell a story is realized at least to some extent. Finally, like novels_{*r*}, novels in verse are self-standing. At the same time, novels in verse differ from novels_{*r*} in that unlike the latter, novels in verse are written in verse, not in prose. Furthermore, in the case of novels_{*r*}, the authors intend to tell a *fictional* story, whereas in the case of novels in verse, the authors intend to tell *some* (not necessarily fictional) story.²⁶

Consider now nonfiction novels. Like novels_{*r*}, nonfiction novels are verbal objects. Also, similar to novels_{*r*}, nonfiction novels have at least n words. Next, like novels_{*r*}, nonfiction novels are intended, by their authors, to tell a story. Furthermore, like in the case of novels_{*r*}, in the case of nonfiction novels, the authors' intention to tell a story is realized at least to some extent. Finally, like novels_{*r*}, nonfiction novels are self-standing. At the same time, unlike novels_{*r*}, nonfiction novels are intended, by their authors, to tell a *real* story. Also, unlike novels_{*r*}, nonfiction novels can be written not only in prose but also in verse.²⁷

Finally, let us have a look at graphic novels. Like novels_{*r*}, graphic novels are intended, by their authors, to tell a story. Also, like in the case of novels_{*r*}, in the case of graphic novels, the authors' intention to tell a story is realized at least to some

²⁶So far as I am aware, there are currently no novels in verse intended to tell a nonfiction story. Yet, from an intuitive viewpoint, such novels are possible.

²⁷I am not aware of any nonfiction novels written in verse. However, there seems no reason to hold that such novels are impossible.

extent. Furthermore, graphic novels are similar to novels_{*r*}, with regard to how long they are: Like novels_{*r*}, graphic novels have an appropriate length. Finally, similar to novels_{*r*}, graphic novels are self-standing. At the same time, in the case of novels_{*r*}, the authors intend to tell a *fictional* story, whereas in the case of graphic novels, the authors intend to tell *some* (not necessarily fictional) story.²⁸ Also, unlike novels_{*r*}, graphic novels are not verbal objects; rather, they are objects that are composed of pictures and/or words organized in a comic-strip format. Finally, graphic novels differ with regard to the minimal word count: Unlike novels _{*r*}, graphic novels may have less than *n* words.

Given the latter fact, as well as the abovementioned fact that graphic novels must have an appropriate length, a natural question arises: What exactly is the appropriate length of a graphic novel? Presumably, the best answer to this question is as follows: The length of a graphic novel is that of a book (under normal printing conditions). Of course, such characterization of the length of a graphic novel is imprecise. But this should not, I think, be considered a shortcoming. It would be a shortcoming if a precise characterization of the length of a graphic novel (that is, a characterization of the form “the length of a graphic novel is at least *x* words/pictures/pages,” where *x* is a particular number) could be given. But such a characterization cannot be given, since there is no fact of the matter regarding the exact minimal length of graphic novels. (This, of course, entails that the concept of “a graphic novel” (and, hence, the concept of “a novel” *simpliciter*) is vague. But

²⁸I am not aware of any graphic novels intended to tell a nonfiction story. But there seems nothing to suggest that such novels cannot exist.

this entailment is unproblematic. As is generally agreed, most, if not all, concepts that were not introduced by means of strict definitions (for example: “cow,” “chair,” “bald,” “perception,” “to paint”) are essentially vague.)

Taking into account what has been said, non-standard novels can be characterized as (1) objects that:

- are verbal
- have at least n words
- were intended, by their authors, to tell a story
- realize this intention at least to some extent
- are written in verse *or* were intended, by their authors, to tell a real story
- are self-standing

—or (2) objects that:

- are composed of pictures and/or words organized in a comic-strip format
- are book-length
- were intended, by their authors, to tell a story
- realize this intention at least to some extent
- are self-standing

However, not only non-standard novels can be characterized this way. Consider a typical biography. Since it is not a novel, it is not a non-standard novel. But it

falls under the first disjunct: It is a verbal object that has at least n words, was intended, by its author, to tell a real story—the story of someone’s life—does, in fact, tell such a story, and is self-standing. Or imagine a historical monograph that is book-length, is composed of pictures and/or words organized in a comic-strip format, was intended, by its author, to tell a story, does, in fact, tell a story, and is self-standing. Such a monograph is not a non-standard novel, as it is not a novel at all. But it falls under the second disjunct.

So what distinguishes non-standard novels from biographies, historical monographs, and similar non-novelistic works? To answer this question, let us first consider Levinson (1989)’s account of “an artwork.” On this account, an artwork is “a thing (item, object, entity) that has been seriously intended for regard-as-a-work-of-art, i.e., regard in any way preexisting artworks are or were correctly regarded” (Levinson, 1989, 21), where “regard” refers to “any mode of interaction with an object” (Levinson, 1989, 31).^{29,30} Thus, according to Levinson (1989), an artwork has the property of having been seriously intended for regard in a way preexisting artworks were correctly regarded. A similar property, I think, is possessed by a non-standard novel. Such a novel was seriously intended (hereafter: “intended”), by its author(s), to be regarded in a way similar to the way novels_r were correctly regarded. Thus, Alexander Pushkin intended *Eugene Onegin* to be regarded in a

²⁹Correspondingly, “regard in any way preexisting artworks are or were correctly regarded” refers to any mode of interaction with an object which was or is proper to some work of art.

³⁰Examples of regards are a regard with close attention to form, a regard with openness to emotional suggestion, and a regard with awareness of symbolism.

way similar to the way novels_r were correctly regarded. And an analogous intention with regard to their non-standard novels was possessed by Truman Capote, Art Spiegelman, Will Eisner, Rodolfo Walsh, and other “non-standard” novelists.

Here, one might ask: What exactly is a way in which a nonstandard novel must be intended, by its author(s), to be regarded? Before answering this question, let us first consider the way novels_r are correctly regarded. This way can be characterized as a set of regards (approaches, attitudes) that includes regards that presuppose:

- knowledge of the historical context (provided that any such context is relevant to the story being told)
- sensitivity to the formal structure
- ability to understand the story being told
- sensitivity to the stylistic features
- sensitivity to the sonic techniques (rhythm, alliteration, consonance, dissonance, etc.)
- ability to empathize with the characters
- willingness to read (or listen) with proper attention
- willingness to attend to the graphic elements (if there are any)
- awareness of the tradition of novel_r writing (that is, the tradition of writing novels like *Don Quixote*, *Pride and Prejudice*, and *War and Peace*)

- knowledge of the fact that the story being told is not intended, by the author(s), to be real
- knowledge of the fact that the main function of a novel_r is not *merely* to inform

—as well as perhaps some other regards. Now, as mentioned above, a way in which a nonstandard novel must be intended, by its author(s), to be regarded is *similar* to the foregoing way. In particular, such a way involves a considerable number of the regards of the way novels_r are correctly regarded, including the regards that presuppose at least some knowledge of the historical context (if any such context is relevant to the story), sensitivity to the formal structure, ability to understand the story, willingness to read (or listen) with proper attention, awareness of the tradition of novel_r writing, and sensitivity to the stylistic features. At the same time, a way in which a nonstandard novel must be intended, by its author(s), to be regarded is not necessarily identical to the way novels_r are correctly regarded. Depending on the kind of non-standard novel, such a way (a) may not involve some of the regards of the way novels_r are correctly regarded or (b) can contain some additional regards. Thus, in the case of a nonfiction novel, the way in which such a novel must be intended, by its author(s), to be regarded does not presuppose regarding this novel with awareness of the fact that the story being told is not intended, by the author(s), to be real. And in the case of a graphic novel, the way in which such a novel must be intended, by its author(s), to be regarded involves a regard that is not involved in the way novels_r are correctly regarded—namely, the regard that presupposes willingness to attend to the comic-strip format as well as at least some

knowledge of the tradition of comic book writing.

It is important to underline that the fact that x was intended, by its author, to be regarded in a way similar to the way novels _{r} were correctly regarded does not necessarily imply that the author had a thought like “I want x to be regarded with at least some knowledge of the historical context (if any such context is relevant to the story), sensitivity to the formal structure, ability to understand the story, willingness to read (or listen) with proper attention, knowledge of the tradition of novel _{r} writing, and sensitivity to the stylistic features, and so on.” Of course, if the author had such a thought, then x was intended by her to be regarded in a way similar to the way novels _{r} were correctly regarded. But x can be considered as having been intended, by x ’s author, to be regarded in such a way even if x ’s author had a different thought—for instance, the thought “I want x to be regarded in a way similar to the way novels _{r} are correctly regarded” or “I want x to be regarded like novels _{r} are correctly regarded.”³¹

Thus, a non-standard novel was intended, by its author(s), to be regarded in a way similar to the way novels _{r} were correctly regarded. Now, what about biographies and other non-novelistic works? Were any of them intended, by their authors, to be regarded in a way similar to the way novels _{r} were correctly regarded? There is no doubt that the answer to this question is “No.” Thus, what distinguishes non-standard novels from biographies and other non-novelistic works is that unlike the latter, non-standard novels were intended, by their authors, to be regarded in a way similar to the way novels _{r} were correctly regarded.

³¹Cf. Levinson (1989) (in particular, his idea of extrinsic and intrinsic modes of artmaking).

Given what has been said, “a non-standard novel” can be defined as follows:

Non-Standard Novel: For all x , x is a non-standard novel just in case x is:

- a novel in verse—a verbal object that (a) is written in verse, (b) has at least n words, where n is a number in the 30,000 word range, (c) was intended, by its author(s), to tell a story, (d) realizes this intention at least to some extent, (e) was intended, by its author(s), to be regarded in a way similar to the way novels_r were correctly regarded, and (f) is self-standing;
- a nonfiction novel—a verbal object that (a) has at least n words, where n is a number in the 30,000 word range, (b) was intended, by its author(s), to tell a nonfiction story, (c) realizes this intention at least to some extent, (d) was intended, by its author(s), to be regarded in a way similar to the way novels_r were correctly regarded, and (e) is self-standing; or
- a graphic novel—an object that (a) is book-length, (b) is composed of pictures and/or words that are organized in a comic-strip format, (c) was intended, by its author(s), to tell a story, (d) realizes this intention at least to some extent, (e) was intended, by its author(s), to be regarded in a way similar to the way novels_r were correctly regarded, and (f) is self-standing.

We are now in a position to answer the question posed at the beginning of this section—namely, the question “Can’t the definition of ‘a novel_r’ be modified so that it could serve as a definition of ‘a novel’ *simpliciter*?” The answer to this question

is “Yes”: The definition of “a novel_r” can indeed be modified so that it could serve as a definition of “a novel” *simpliciter*. To modify it that way, the definiens of “a non-standard novel” should be added, as a disjunct, to the definiens of the definition of “a novel_r.” The result of this modification is the following definition:

Novel (DN): For all x , x is a novel if and only if x is:

- (1) a novel_r—a verbal object that (a) is written, for the most part, in prose, (b) has at least n words, where n is a number in the 30,000 word range, (c) was intended, by its author(s), to tell a fictional story, (d) realizes this intention at least to some extent, and (e) is self-standing; or
- (2) a non-standard novel, that is:
 - a novel in verse—a verbal object that (a) is written in verse, (b) has at least n words, where n is a number in the 30,000 word range, (c) was intended, by its author(s), to tell a story, (d) realizes this intention at least to some extent, (e) was intended, by its author(s), to be regarded in a way similar to the way novels_r were correctly regarded, and (f) is self-standing;
 - a nonfiction novel—a verbal object that (a) has at least n words, where n is a number in the 30,000 word range, (b) was intended, by its author(s), to tell a nonfiction story, (c) realizes this intention at least to some extent, (d) was intended, by its author(s), to be regarded in a way similar to the way novels_r were correctly regarded, and (e) is self-standing; or

- a graphic novel—an object that (a) is book-length, (b) is composed of pictures and/or words that are organized in a comic-strip format, (c) was intended, by its author(s), to tell a story, (d) realizes this intention at least to some extent, (e) was intended, by its author(s), to be regarded in a way similar to the way novels_r were correctly regarded, and (f) is self-standing.

Before proceeding further, it is worth making two remarks concerning the foregoing definition (hereafter: “DN”). First, there is good reason to hold that DN reflects the actual structure of the concept expressed by “a novel.” One of our intuitions about novels is that some entities called “novels” are doubtless novels, whereas other such entities do not seem entirely like novels. Consider novels_r: *Pride and Prejudice*, *War and Peace*, *Moby-Dick*, etc. There is no doubt that each of them is a novel. The same, however, cannot be said about non-standard novels. Consider *Eugene Onegin*. Although it is categorized as a novel, it does not seem like a novel (“How can it be a novel, given that it is written in verse? It’s a poem!”). Or consider *Maus*. The consensus is that it is a novel. But, intuitively, it is not (“How can it be a novel, given that it has a comic-strip format? It’s a comic book!”). What has been said about *Eugene Onegin* and *Maus* can be said about other non-standard novels as well. Thus, *from an intuitive viewpoint*, novels_r are clearly novels, whereas the status of non-standard novels qua novels is dubious. In light of this, there is good reason to hold that the concept expressed by “a novel” is composed of two sub-concepts: (a) the (core) concept that covers novels like *Pride*

and *Pride and Prejudice*, *War and Peace*, *Moby-Dick*, etc. and (b) the (peripheral) concept that covers novels of a less traditional sort, such as *Eugene Onegin* and *Maus*. But if this is so, then a definition of “a novel” reflects the structure of the concept of “a novel” just in case the definiens of this definition is composed of two disjuncts that express, respectively, the concept that covers novels like *Pride and Prejudice*, *War and Peace*, *Moby-Dick*, etc. and the concept that covers novels of a less traditional sort, such as *Eugene Onegin* and *Maus*. Meanwhile, the definiens of **DN** is, in fact, composed of these disjuncts.

Second, **DN** has certain historical implications regarding the novel. One of these implications is that the first novel came into existence in or before the I century AD, long before the time when a considerable number of first paradigmatic novels were written (the XVII–XVIII centuries). Consider *Callirhoe*—one of the so-called “ancient Greek novels” that was written in the I century AD. It is a verbal object that is written in prose, has more than 30,000 words, tells a fictional story, and, clearly, was intended by its author, Chariton of Aphrodisias, to tell such a story. So, according to **DN**, *Callirhoe* is a novel. As a result, if no entity created before the I century falls under **DN**, then **DN** implies that the first novel was written in the I century; if, on the other hand, there is an entity that was created before the I century and falls under **DN**, then **DN** implies that the first novel was written before the I century.

Another historical implication of **DN** is that non-standard novels—nonfiction novels, novels in verse, and graphic novels—were not created before the tradition of novel_r writing came into existence. As mentioned above, to regard *x* in a way

similar to the way novels_r were correctly regarded, it is necessary to regard x with at least some awareness of the tradition of novel_r writing. Suppose now that this tradition does not exist. Then x cannot be regarded in a way similar to the way novels_r were correctly regarded—and, hence, cannot be intended to be regarded in such a way.³² Meanwhile, according to **DN**, having been intended to be regarded in a way similar to the way novels_r were correctly regarded is a necessary property of a non-standard novel. Thus, if **DN** is true, then the creation of the first non-standard novel took place after the tradition of novel_r writing came into existence.

1.4 Objections

Let us now examine potential objections to **DN**.

As pointed out in Footnote 1, to define x (where x is some expression) is to provide a sufficiently informative set of conditions that are satisfied by all entities that fall under x and only by such entities. So a definition of x is satisfactory just in case it is sufficiently informative and covers all and only those entities that fall under x . In light of this, possible objections to **DN** can be divided into two groups: (a) objections aimed at showing that **DN** is insufficiently informative and (b) objections whose purpose is to demonstrate that **DN** is too broad (that is, covers entities that are not novels) or too narrow (that is, does not cover some novels). Let us first consider objections of group (a).

³²The expression “can” here is used in the probabilistic (not the absolute) sense.

1.4.1 Objections of Group (a)

Objection 1. The definiens of **DN** involves expressions containing the term “novel”—namely, the expressions “a novel_r,” “a non-standard novel,” “a novel in verse,” “a nonfiction novel,” and “a graphic novel”—and, hence, implicitly contains the concept of “a novel.” Meanwhile, if the definiens of a definition contains the concept of the expression being defined, then this definition is insufficiently informative.

Response. Objection 1 assumes that the fact that the definiens of **DN** involves the expressions “a novel_r,” “a non-standard novel,” “a novel in verse,” “a nonfiction novel,” and “a graphic novel” implies that this definiens contains the concept of “a novel.” But this assumption is false. The foregoing fact implies that the definiens of **DN** contains the concept of “a novel” only if at least one of the abovementioned expressions is defined using the concept of “a novel.” But none of them is, in fact, defined using this concept.

Objection 2. **DN** does not define the expressions “a story,” “fictional,” “verse,” “book-length,” and “a comic-strip format.” Meanwhile, in order for this definition to be sufficiently informative, it must define these expressions.

Response. Objection 2 assumes that to be sufficiently informative, **DN** must define “a story,” “fictional,” “verse,” “book-length,” and “a comic-strip format.” Why think that this assumption is true? One possible answer is that any sufficiently informative definition must define all the expressions it involves. However, this answer is unsatisfactory. Consider the following definition of “water”: For all x , x

is water just in case x has the molecular structure H_2O . This definition is doubtless acceptable and, hence, is sufficiently informative. However, it does not provide the definitions of “a molecular structure,” “H,” and “O.”

Furthermore, if a definition defines all the expressions it involves, then there must be infinite chains of (non-circular) definitions. But such chains do not exist, as some expressions are basic and, hence, indefinable. Thus, there is no definition that defines all the expressions it involves. So if the answer being discussed is correct, no definition is sufficiently informative, which is, of course, absurd.

Another potential answer is that a sufficiently informative definition must define all definable expressions—that is, expressions that can be defined using solely basic expressions—and the expressions “a story,” “fictional,” “verse,” “book-length,” and “a comic-strip format” are definable. But this answer also fails. Consider, once again, the definition of “water”—“Water is H_2O .” As already mentioned, it is sufficiently informative. But it does not define the definable expressions that it contains—for instance, the expressions “ H_2O ” and “a molecular structure.”

Thus, neither answer is satisfactory. Meanwhile, there seems no other potentially satisfactory answer to the question being discussed. Therefore, the assumption involved in Objection 2—that to be sufficiently informative, **DN** must define “a story,” “fictional,” “verse,” “book-length,” and “a comic-strip format”—can be rejected.³³

³³What has been said above does not imply, of course, that the expressions “a story,” “fictional,” “verse,” “book-length,” and “a comic-strip format” are not worth examining. An examination of these expressions is doubtless worthwhile. Yet such an examination is beyond the scope of the

1.4.2 Objections of Group (b)

Having examined the objections of group (a) (the objections aimed at showing that **DN** is not sufficiently informative), let us now examine the objections of group (b) (the objections aimed at showing that **DN** is too broad or too narrow).

Objection 3. Suppose there is an object O that (a) is verbal, (b) is written in prose, (c) was intended, by its author, to tell a fictional story, (d) does, in fact, tell such a story, and (e) is self-standing. Suppose next that O has 27,700 words. Is it a novel_r? An answer to this question depends on whether O 's word count is in the 30,000 word range. Is O 's word count, in fact, in this range? Intuition cannot help us answer this question. It is not obvious that the number of words O has—27,700—*is not* in the 30,000 word range; likewise, it is not obvious that this number *is* in this range. At the same time, the answer to the foregoing question cannot be found using some principle that precisely determines the lower bound of the 30,000 word range, as, given the essential vagueness of this bound, no such principle exists. In light of what has been said, there is good reason to hold that using **DN**, it is impossible to establish whether O is a novel_r. Meanwhile, if a definition cannot be used to establish whether an object falls under the concept being defined, then according to this definition, it is indeterminate whether this object falls under this concept. Thus, **DN** entails that for some objects, their status as novels is indeterminate. However, in fact, any entity is either clearly a novel or clearly a non-novel. So **DN** either excludes from the extension of “a novel” some novels or fails to exclude from

current project.

this extension some objects that are not novels.

Response. Objection 3 assumes that no entities are indeterminate qua novels. However, this assumption is false. In fact, besides entities that are clearly novels and entities that are clearly non-novels, there are also borderline cases—entities that are neither clearly novels nor clearly non-novels.³⁴

Objection 4. Consider the *Iliad*. It is a verbal object that has at least n words, is written in verse, was intended by its author—Homer—to tell a story, and does, in fact, tell a story. Furthermore, most likely, Homer intended it to be regarded in a way similar to the way novels_r were correctly regarded. Thus, the *Iliad* seems to fall under the definition of “a novel in verse” and, hence, under **DN**. But the *Iliad* is not a novel.

Response. Objection 4 assumes that the *Iliad* satisfies all the conditions of the definition of “a novel in verse.” Is this assumption true? The *Iliad* is doubtless a verbal object that is written in verse, has at least n words, was intended by its author to tell a story, and does, in fact, tell a story. But was the *Iliad* intended by its author, Homer, to be regarded in a way similar to the way novels_r were correctly regarded? Given what has been said in the previous section, in order for the *Iliad* to have been intended, by Homer, to be regarded in such a way, he had to intend it to

³⁴The same, by the way, can be said about other kinds of entities. Consider, for instance, artworks in general. Surely, there are entities that are clearly artworks (the *Mona Lisa*, *David*, Beethoven’s Symphony No. 5, *Moby-Dick*) and entities that are clearly non-art (humans, trees, planets). But there are also entities that are neither clearly artworks nor clearly non-art—for example, computer games.

be regarded in a way that included the regard presupposing at least some awareness of the tradition of novel_r writing. But he could not have intended the *Iliad* to be regarded in a way that included this regard—since the tradition of novel_r writing did not exist at that time. Thus, the *Iliad* was not intended, by Homer, to be regarded in a way similar to the way novels_r were correctly regarded. But if this is so, then, contrary to what Objection 4 assumes, the *Iliad* does not satisfy all the conditions of the definition of “a novel in verse.”

Objection 5. According to **DN**, some works written before the XVII century might be novels. However, the consensus is that the novel qua a genre—and, hence, works of this genre—did not come into existence until the XVII century.

Response. The objection assumes that there is a generally accepted view as to when the novel came into existence. Is this assumption true? There are a considerable number of theorists who believe that the first *European* novel was written in the XVII century.³⁵ But there is no consensus as to when the first novel *simpliciter* was written. Some theorists argue that it was written in the XVII century (Cervantes’s *Don Quixote*). Others claim that it was written in the XI century (Murasaki Shikibu’s *The Tale of Genji*). There are also theorists arguing that the creation of the first novel dates back to the I–III centuries AD (the so-called “ancient Greek novels”: Chariton’s *Callirhoe*, Achilles Tatius’s *Leucippe and Clitophon*, Longus’s *Daphnis and Chloe*, Xenophon’s *Ephesian Tale*, and Heliodorus

³⁵The view that the European novel was invented in the XVII century is not universally accepted, however. Thus, Watt (1967) argues that the European novel was invented later—in the XVIII century.

of Emesa's *Aethiopica*.) Thus, the foregoing assumption of Objection 5 is false.

Objection 6. Consider the Harry Potter series of novels. It is composed of seven books, each of which is an $n+$ word verbal object written, for the most part, in prose. Meanwhile, anything that is composed of $n+$ word verbal objects written, for the most part, in prose is an $n+$ word verbal object written, for the most part, in prose. So the Harry Potter series is such an object. Furthermore, this series was intended by its author, J. K. Rowling, to tell a fictional story—the story about Harry Potter—and this intention was successfully realized. Finally, the Harry Potter series is doubtless a self-standing object. Thus, according to **DN**, the Harry Potter series of novels is a novel. But this series is not a novel; it is a *collection* of novels.

Response. Objection 6 assumes that the Harry Potter series of novels is not a novel. But this assumption, I think, can be rejected. Indeed, the Harry Potter series is not normally categorized as a novel. Yet there seems no real reason against categorizing this series as such.³⁶ Furthermore, it is common for some analogous series to be categorized as novels. Consider, for instance, Marcel Proust's *À la recherche du temps perdu*. Like the Harry Potter series, it is a series of novels. At the same time, it is categorized as a novel.

Here, one might ask: Can *any* series of novels be categorized as a novel? The answer to this question is “No.” To be categorized as a novel, a series of novels must satisfy the conditions of **DN**. Meanwhile, not any series of novels satisfies these conditions (consider, for instance, a series of novels that was not intended, by

³⁶Note that categorizing the Harry Potter series of novels as a novel does not entail that this series cannot be characterized as a collection of novels.

its author(s), to tell a story or a series of novels that tells several stories, and not a (unified) story).

Objection 7. If **DN** is true, then a non-standard novel must be intended, by its author(s), to be regarded in a way similar to the way novels_r are correctly regarded. However, there is no real reason to hold that, say, *In Cold Blood* was, in fact, intended, by its author, Truman Capote, to be regarded in such a way. So if **DN** is true, then *In Cold Blood* is not a novel. But *In Cold Blood* is doubtless a novel.

Response. Objection 7 is based on the idea that there is no real reason to hold that *In Cold Blood* was, in fact, intended, by Capote, to be regarded in a way similar to the way novels_r were correctly regarded. But this idea is highly questionable. According to Capote, “a nonfiction novel” “employ[s] all the techniques of fictional art but [is] nevertheless immaculately factual” (Plimpton, 1966). Now, there is little doubt that by “fictional art” here, he means primarily the art of the traditional novel, or novel_r. So it can be said that according to Capote, a nonfiction novel is to be regarded qua an entity that possesses the relevant features of a novel_r except the feature of having been intended, by the author(s), to tell a fictional story. But if this is the case, then it is reasonable to suppose that in his view, a nonfiction novel should be regarded in a way similar to the way novels_r are correctly regarded. Meanwhile, if this supposition is true, then, given the fact that *In Cold Blood* is explicitly characterized by Capote as a nonfiction novel,³⁷ there is, in fact, a good reason to think that he intended *In Cold Blood* to be regarded in a way similar to

³⁷See Plimpton (1966).

the way novels_r were correctly regarded.

Here, one could ask: Can't Objection 7 be made successful by replacing *In Cold Blood* with some other non-standard novel—in particular, some non-standard novel that was not intended, by its author, to be regarded in a way similar to the way novels_r were correctly regarded? I am not aware of any evidence that can be used to show that there is an entity that (a) is clearly a non-standard novel but (b) was not intended, by its author, to be regarded in a way similar to the way novels_r were correctly regarded. At the same time, in an overwhelming majority of cases, there is evidence suggesting that the author of a non-standard novel did, in fact, intend this novel to be regarded in such a way.³⁸ Taking this into account, the foregoing question, I think, should be answered in the negative.

Objection 8. According to **DN**, a novel must be written. But there can be unwritten novels—in particular, novels created by pronouncing certain words or in one's mind.

Response. As pointed out in Footnote 15, the expression “to write” is used here in a broad sense—the sense according to which to write is to generate a text. So, contrary to what Objection 8 implies, any novel created by pronouncing certain words or in one's mind is a written novel.

Objection 9. Suppose *N* is identical in semantic content to some novel—say, E. B. White's *Charlotte's Web* (c. 33,500 words). Suppose next that *N* is written in a language other than English. Finally, suppose that since this language makes it

³⁸The most common piece of such evidence is the fact that the author *explicitly* characterized his work as “a novel.”

possible to communicate meanings in considerably fewer words than English, N has c. 20,000 words. Then, since N 's word count is not in the 30,000 word range, **DN** entails that N is not a novel. However, given the fact that the semantic content of N is the same as the semantic content of E. B. White's *Charlotte's Web*, there is good reason to consider N a novel.

Response. Objection 9 assumes that N has c. 20,000 words. But this assumption is false. As is clear from Footnote 13, a novel has x words just in case this novel has x English words (it is written in English) or its English translation has x words (it is written in a language other than English). So if N has c. 20,000 words, then, given that N is not written in English, N 's English translation must have c. 20,000 words. But, taking into account the fact that the language in which N is written makes it possible to communicate meanings in considerably fewer words than the English language does, as well as the fact that the semantic content of N is the same as the semantic content of a 33,500 word English novel (E. B. White's *Charlotte's Web*), N 's translation must have c. 33,500 English words.

1.5 Final Remarks

Thus, none of the objections discussed in the previous section stands up to criticism. Meanwhile, there seem to be no other potentially acceptable objections to **DN**. In light of this, there is good reason to think that **DN** is both sufficiently informative and covers all and only those entities that are, in fact, novels. But if that is the case, then **DN** is a satisfactory definition of "a novel."

In closing, it is worth noting that the expressions “a novella” and “a short story” can be defined in a way similar to the way “a novel” has been defined. Consider first the expression “a novella.” It can be understood as referring to novellas_r (that is, novellas that are not nonfiction novellas, novellas in verse, or graphic novellas) and non-standard novellas (that is, nonfiction novellas, novellas in verse, and graphic novellas). Meanwhile, novellas_r and non-standard novellas are not that different from novels_r and non-standard novels, respectively. Novellas_r differ from novels_r in length: While the minimal word count of a novel_r is in the 30,000 word range, the minimal word count of a novella_r is, most likely, in the 15,000 word range.³⁹ Also, while a novel_r does not have a maximal word count (as mentioned above, a novel_r can be infinitely long), a novella_r does have such a word count—prima facie, it is in the 25,000 word range. As regards non-standard novellas, they differ from non-standard novels in that unlike the latter, they are intended, by their authors, to be regarded in a way similar to the way novellas_r are correctly regarded. Furthermore, non-standard novellas differ from non-standard novels in length. Unlike the word count of nonfiction novels and novels in verse, the word count of nonfiction novellas and novellas in verse seems to be in the 15,000–25,000 word range. And unlike the page count of graphic novels, the page count of graphic novellas is, most likely, in the 40–70 page range.

Given what have been said, “a novella” could be defined as follows:

Novella: For all x , x is a novella if and only if x is:

³⁹The given word range as well as the word and page ranges mentioned in what follows are guesstimates and, hence, may well turn out to be inaccurate.

- (1) a novella_r—a verbal object that (a) is written, for the most part, in prose, (b) has n words, where n is a number in the 15,000–25,000 word range, (c) was intended, by its author(s), to tell a fictional story, (d) realizes this intention at least to some extent, and (e) is self-standing; or
- (2) a non-standard novella, that is:
 - a novella in verse—a verbal object that (a) is written in verse, (b) has n words, where n is a number in the 15,000–25,000 word range, (c) was intended, by its author(s), to tell a story, (d) realizes this intention at least to some extent, (e) was intended, by its author(s), to be regarded in a way similar to the way novellas_r were correctly regarded, and (f) is self-standing;
 - a nonfiction novella—a verbal object that (a) has n words, where n is a number in the 15,000–25,000 word range, (b) was intended, by its author(s), to tell a nonfiction story, (c) realizes this intention at least to some extent, (d) was intended, by its author(s), to be regarded in a way similar to the way novellas_r were correctly regarded, and (e) is self-standing; or
 - a graphic novella—an object that (a) has at least m pages, where m is a number in the 40–70 page range, (b) is composed of pictures and/or words organized in a comic-strip format, (c) was intended, by its author(s), to tell a story, (d) realizes this intention at least to some extent, (e) was intended, by its author(s), to be regarded in a

way similar to the way novellas_r were correctly regarded, and (f) is self-standing.

Consider now the expression “a short story.” It can be understood as referring to short stories_r (that is, short stories that are not nonfiction short stories, short stories in verse, or graphic short stories) and non-standard short stories (that is, nonfiction short stories, short stories in verse, and graphic short stories). Like novellas_r and non-standard novellas, short stories_r and non-standard short stories do not differ considerably from novels_r and non-standard novels, respectively. Short stories_r differ from novels_r in length: Unlike the minimal word count of a novel_r, the minimal word count of a short story_r amounts to a few words. Also, while a novel_r does not have a maximal word count, a short story_r does have such a word count—prima facie, it is in the 15,000 word range. As regards non-standard short stories, they differ from non-standard novels in the way they are intended, by their authors, to be regarded: Unlike non-standard novels, non-standard short stories are intended to be regarded, by their authors, in a way similar to the way short stories_r are correctly regarded. Furthermore, non-standard short stories differ from non-standard novels in length. Unlike the word count of nonfiction novels and novels in verse, the word count of nonfiction short stories and short stories in verse seems to be in the 2–15,000 word range. And unlike the page count of graphic novels, the page count of graphic short stories is, most likely, in the 1–40 page range.

In light of what has been said, “a short story” can be defined as follows:

Short story: For all x , x is a short story if and only if x is:

- (1) a short story_r—a verbal object that (a) is written, for the most part, in prose, (b) has n words, where n is a number in the 2–15,000 word range, (c) was intended, by its author(s), to tell a fictional story, (d) realizes this intention at least to some extent, and (e) is self-standing; or
- (2) a non-standard short story, that is:
 - a short story in verse—a verbal object that (a) is written in verse, (b) has n words, where n is a number in the 2–15,000 word range, (c) was intended, by its author(s), to tell a story, (d) realizes this intention at least to some extent, (e) was intended, by its author(s), to be regarded in a way similar to the way short stories_r were correctly regarded, and (f) is self-standing;
 - a nonfiction short story—a verbal object that (a) has n words, where n is a number in the 2–15,000 word range, (b) was intended, by its author(s), to tell a nonfiction story, (c) realizes this intention at least to some extent, (d) was intended, by its author(s), to be regarded in a way similar to the way short stories_r were correctly regarded, and (e) is self-standing; or
 - a graphic short story—an object that (a) has at least m pages, where m is a number in the 1–40 page range, (b) is composed of pictures and/or words organized in a comic-strip format, (c) was intended, by its author(s), to tell a story, (d) realizes this intention at least to some extent, (e) was intended, by its author(s), to be regarded in a

way similar to the way short stories, were correctly regarded, and (f) is self-standing.

Now, a natural question arises: Are the foregoing definitions of “a novella” and “a short story” satisfactory? Answering this question requires a substantial analysis of these definitions. Such an analysis, however, goes beyond the scope of this dissertation.

Chapter 2: Defining “an Instance of an Artwork”

Introduction

My goal in this chapter is to define “an instance of an artwork” as well as some derivative expressions—in particular, “a well-formed instance of an artwork” and “a non-well-formed instance of an artwork.”¹ I begin with an exposition and defense of Davies (2010)’s definition of “an instance of an artwork” (Sections 2.1 and 2.2). Next, I elaborate on this definition by defining “a well-formed instance of an artwork” and “a non-well-formed instance of an artwork” (Sections 2.3 and 2.4). Finally, I make a few additional remarks concerning the expression “an instance of an artwork.” In particular, I define “a token of an artwork,” which is closely related to “an instance of an artwork,” examine certain ontological implications of the definition of “an instance of an artwork,” and provide an alternative formulation of this definition (Section 2.5).

¹For the purposes of this dissertation, there is no need to define “an instance of an artwork” and its derivatives in all contexts; it is sufficient to define them only in the context of the ontology of artworks. Given this, in what follows, the expression “an instance of an artwork” and its derivatives are assumed to be located solely within the mentioned context.

2.1 Defining “an Instance of an Artwork”

How can “an instance of an artwork” be defined? To my knowledge, the most detailed, and the only explicit, answer to this question has been given by Davies (2010). In his view, “an instance of an artwork” can be used in two distinct senses—*the purely epistemic* and *the proveniential*—and so the question posed above amounts to two questions: “How can ‘an instance of an artwork’ used in the purely epistemic sense (hereafter: ‘an instance_e of an artwork’) be defined?” and “How can ‘an instance of an artwork’ used in the proveniential sense (hereafter: ‘an instance_p of an artwork’) be defined?” The former question, according to Davies (2010), can be answered as follows: For all x , x is an instance_e of an artwork just in case x “makes manifest to receivers certain properties that bear experientially upon the appreciation of the work” (Davies, 2010, 415). Or, in other words: For all x , x is an instance_e of an artwork if and only if x manifests certain properties that must be experienced to fully appreciate this work.²

Regarding this definition, a natural question arises: What is meant by the expression “to manifest a property”? Given what Davies (2010) says, and taking into account the standard linguistic practices present in the philosophical literature, this expression can be defined as follows: For all x , x manifests a property just in case this property is apprehensible by directly perceiving x —that is, by perceiving x with the help of one or more of our sensory faculties (such as sight (vision), hearing (audition), taste (gustation), smell (olfaction), and touch (somatosensation)).

²The question of what these “certain” properties are is addressed later in the chapter.

Note that the only thing that determines whether an entity manifests a property is whether this property can be apprehended by perceiving this entity with the help of at least one sensory faculty. Thus, a red apple manifests the property of being red—since this property can be apprehended by perceiving this apple with the help of the faculty of sight. At the same time, an apple produced in Florida does not manifest the property of being produced in Florida—for, this property cannot be apprehended by perceiving this apple with the help of any sensory faculties.

Note also that *manifesting* a property is not equivalent to *having* this property. For, an object can have a property without manifesting it. Thus, an apple produced in Florida *has* the property of being produced in Florida but, as mentioned above, *does not manifest* this property. Similarly, a musical score *has* the property of sounding a particular way but *does not manifest* this property. (If it did, then at least some sonic properties could be apprehended by perceiving it with the help of a sensory faculty. However, no sonic property can be apprehended that way. In order for this to be possible, it must be possible to hear a musical score. But no musical score can be heard, since (a) (strictly speaking) only sounds can be heard, and (b) a musical score is not a sound (rather, it is a concrete sequence of notes and other symbols).³)

Having clarified the expression “to manifest a property,” let us return to Davies (2010)’s definition of “an instance_e of an artwork”—the definition according to which

³Of course, we can meaningfully say that a musical score can be heard. But when we say this, we do not mean that this score can *literally* be heard; what we mean is that the sounds generated with its help (or perhaps the sounds it encodes) can be heard.

an instance_e of an artwork is whatever manifests certain properties that must be experienced to fully appreciate this work. It is important to underline that the set of properties that, according to this definition, an instance_e of an artwork must manifest does not necessarily involve *all* the properties that must be experienced to fully appreciate this work. Thus, consider Alexander Ivanov's painting *The Apparition of Christ Before the People*. It is reasonable to assume that to fully appreciate this painting, it is necessary to learn about the process of creating this painting. Meanwhile, learning about this process is impossible without engaging experientially with at least some of the properties of Ivanov's preparatory sketches.⁴ Thus, experiencing these properties is requisite for a full appreciation of *The Apparition of Christ Before the People*. But they cannot be possessed by any instance of *The Apparition of Christ Before the People*—the original canvas (and perhaps certain very good copies of this canvas)—and, hence, cannot be manifested by an instance of this painting.

Or consider Beethoven's Symphony No. 5. It can be plausibly argued that a complete appreciation of this symphony requires an experiential engagement with those experienceable properties that enable to grasp the history of its composition. At the same time, these properties cannot be possessed by any of the instances of Symphony No. 5—particular musical performances—and, hence, cannot be manifested by an instance of this symphony.

Thus, the set of properties manifested by an instance_e of an artwork does not

⁴This is not to say, of course, that learning about the creative process leading up to an artwork—
all the stops and starts, all the false turns, etc.—is always requisite to full appreciation.

necessarily involve all the properties that must be experienced to fully appreciate this work. But what, then, determines whether a property that must be experienced to fully appreciate an artwork must be manifested by an instance_e of this work? According to Davies (2010), what determines that is whether such a property falls under the category of properties through which the primary content⁵—that is, the set of “those contentful properties that may be the ground of other contentful properties but which are not themselves grounded in contentful properties” (Davies, 2010, 411)⁶—of the work is articulated: If the property falls under this category, then it must be manifested by an instance_e; otherwise, the property does not have to be manifested by this instance_e.

Given what has been said—and assuming that the expression “primary properties” denotes properties through which the primary content of an artwork is articulated—Davies (2010)’s definition of “an instance_e of an artwork” can be formulated more precisely as follows:

Instance_e: For all x , x is an instance_e of an artwork if and only if x manifests certain primary properties that must be experienced to fully appreciate this work.

Note that the only thing that matters for being an instance_e of an artwork is manifesting the relevant primary properties that must be experienced to fully appreciate this work: As long as an entity manifests such properties, it is an instance_e

⁵The term “content” is used here in a broad sense—to refer to the overall artistic content (and not just to the semantic content).

⁶Alternatively, the primary content of an artwork can be characterized as the set of the basic properties that determine the content of this work.

of the corresponding work. Thus, consider, for example, the original canvas of Kazimir Malevich's *Black Square* and its indiscernible counterpart⁷ that was created by someone completely unfamiliar with Malevich's works. Since both the canvas and the counterpart manifest the same properties, they are both instances_e of *Black Square*. Likewise, both a correct recitation of R. L. Stevenson's "To Friends at Home" and its indiscernible counterpart produced by someone who has never encountered Stevenson's poetry manifest the same properties and so are instances_e of this poem.⁸

Let us now turn to the question "How can 'an instance_p of an artwork' be defined?" This question, according to Davies (2010), can be answered in the following way: For all x , x is an instance_p of an artwork if and only if x is an instance_e that is related, in an appropriate historical-intentional respect, to this work. Or, in other words:

Instance_p: For all x , x is an instance_p of an artwork if and only if x (a) manifests certain primary properties that must be experienced to fully appreciate this work and (b) stands in an appropriate historical-intentional relation to it.

It is worth stressing that unlike being an instance_e, being an instance_p of an artwork requires not only manifesting certain primary properties that must be expe-

⁷Following Fisher (1995), I define "an indiscernible counterpart" as follows: For all x and for all y , x is an indiscernible counterpart of y if and only if x and y share all of their manifest properties (where a property is manifest just in case it is manifested by something).

⁸It is assumed that the original canvas of *Black Square* and the recitation of the poem "To Friends at Home" manifest the relevant primary properties that must be experienced to fully appreciate the corresponding works.

rienced to fully appreciate this work but also standing in an appropriate historical-intentional relation to this work. Thus, an indiscernible counterpart of Malevich's *Black Square* that was created by someone completely unfamiliar with Malevich's works is not an instance_p of *Black Square*, since this counterpart does not stand in any appropriate historical-intentional relation to this painting. Similarly, an indiscernible counterpart of a recitation of Stevenson's "To Friends at Home" produced by someone who has never encountered Stevenson's poetry is not an instance_p of "To Friends at Home" because this counterpart does not stand in any appropriate historical-intentional relation to this poem.

Regarding the definition of "an instance_p of an artwork," a natural question arises: What exactly is the appropriate historical-intentional relation to an artwork in which an instance_p of this work must stand?⁹ An answer to this question can vary depending on what kind of art is under consideration. In the case of classical music, the historical-intentional relation to an artwork in which an instance_p of this work stands is usually understood as the relation of being identical to a performance of this work generated with the help of either the original score (i.e., the score directly created—say, written or typed—by the composer) or an entity that stands in the "copy" relation to this score (where x stands in the "copy" relation to y just in case x is a copy of y , or x is a copy of a copy of y , or x is a copy of a copy of a copy of y , and so on).¹⁰ In the case of photographic art, the historical-intentional

⁹Davies (2010) leaves this question open. So the following is not part of his account of "an instance of an artwork."

¹⁰If this treatment is correct, then "an instance_p of a work of classical music" can be defined as follows: For all x , x is an instance_p of a work of classical music if and only if x (a) manifests certain

relation to an artwork in which an instance_p of this work stands is usually taken to be the relation of being identical to a print derived from particular photographic film created by the author.¹¹ In the case of painting, the historical-intentional relation to an artwork in which an instance_p of this work stands is typically treated as the relation of being identical to the original canvas.^{12,13} And in the case of literature, the consensus is that the historical-intentional relation to an artwork in which an instance_p of this work stands is the relation of being identical to the work's original manuscript or an entity that stands in the "copy" relation to this manuscript.¹⁴

primary properties that must be experienced to fully appreciate this work and (b) stands in the relation of being identical to the work's performance generated with the help of a score that is either the original score (i.e., the score directly created by the composer) or an entity that stands in the "copy" relation to this original score.

¹¹If this treatment is correct, then "an instance_p of a photographic work" can be defined as follows: For all x , x is an instance_p of a photographic work if and only if x (a) manifests certain primary experienceable properties that must be experienced to fully appreciate this work and (b) stands in the relation of being identical to a print derived from particular photographic film created by the work's author.

¹²If this treatment is correct, the definition of "an instance_p of a painting" can be formulated as follows: For all x , x is an instance_p of a painting if and only if x (a) manifests certain primary properties that must be experienced to fully appreciate this painting and (b) stands in the relation of being identical to the canvas.

¹³According to a widely endorsed view, a painting is identical to its canvas. If this view is true, then the relation here is that of being identical to the painting.

¹⁴If this treatment is correct, then the definition of "an instance_p of a painting" can be formulated as follows: For all x , x is an instance_p of a literary work if and only if x (a) manifests certain primary properties that must be experienced to fully appreciate this work and (b) stands in the relation of being identical to the work's original manuscript or an entity that stands in the "copy" relation to

It should be noted that the foregoing interpretations of the appropriate historical-intentional relation to an artwork in which an instance_p of this work must stand are not claimed to be correct. In fact, although these interpretations are widely accepted, they may well be misguided. The interpretation provided in the case of classical music implies that an instance_p of a musical work must be a performance. But why can't an instance of a musical work be something other than a performance—say, a playing of a recording of a performance or a reproduction of a performance (that is, a particular sequence of sounds generated by some electronic device with the help of a musical score)? Similar questions arise with regard to the interpretations given in the case of painting and photographic art. According to the interpretation given in the case of painting, a painting has only one instance, namely the canvas. But why can't a painting be instanced not only by its canvas but also by something else—say, certain copies (for example, molecule-for-molecule duplicates) of this canvas? The interpretation provided in the case of photographic art assumes that the only instances of photographic artworks are prints. But why can't photographic artworks be properly instanced by things other than prints—say, copies of prints? Finally, it can be questioned whether the interpretation given in the case of literary works is right in identifying instances of literary works with either original manuscripts or their copies.¹⁵ (Note that what has been said here is not intended to show that the mentioned interpretations are, in fact, misguided.

this manuscript.

¹⁵In the following chapter, I provide an argument showing that this interpretation is actually wrong in doing that.

The goal is to suggest that these interpretations *could* be misguided.)

2.2 Evaluating the Definition

As pointed out in Chapter 1,¹⁶ a definition of an expression is satisfactory just in case this definition is sufficiently informative and covers all and only those entities that fall under this expression. In light of this, Davies (2010)'s definition can be rejected on the grounds that it is insufficiently informative or on the grounds that it does not cover all and only those entities that fall under "an instance of an artwork." There seems no real reason to question the sufficiency of the informativeness of Davies (2010)'s definition. However, one could question whether this definition covers all and only those entities that fall under "an instance of an artwork." In particular, one could argue as follows. Suppose some entity E makes available for experience the primary properties that must be experienced to fully appreciate some artwork A (where y makes x available for experience just in case y makes it possible to experientially engage with x either by directly perceiving x or by perceiving x as a result of applying a special skill, or, in other words, a skill that is not acquired in a natural way (such as the skill of reading or the skill of playing a musical instrument)). Suppose next that E does not manifest these properties. Then Davies (2010)'s definition entails that E is not an instance of A . Is this entailment true? Given the actual use of the expression "an instance of an artwork," (a) to be an instance_e of an artwork, it is sufficient to make available for experience the primary

¹⁶See Chapter 1, Footnote 1.

properties that must be experienced to fully appreciate this work, and (b) to be an instance_{*p*} of an artwork, it is sufficient (i) to make available for experience the primary properties that must be experienced to fully appreciate this work and (ii) to stand in an appropriate historical-intentional relation to this work. In light of this, *E* is an instance_{*e*} of *A*. Moreover, assuming that *E* stands in an appropriate historical-intentional relation to *A*, *E* is an instance_{*p*} of *A*. Thus, the foregoing entailment is false—and, as a result, Davies (2010)’s definition fails to cover all of those entities that fall under “an instance of an artwork.”

This objection is based on the thesis that the actual use of the expression “an instance of an artwork” supports the account of “an instance of an artwork” according to which (a) if *x* makes available for experience the primary properties that must be experienced to fully appreciate an artwork, then *x* is an instance_{*e*} of this work, and (b) if *x* (i) makes available for experience the primary properties that must be experienced to fully appreciate an artwork and (ii) stands in an appropriate historical-intentional relation to it, then *x* is an instance_{*p*} of this work. Is this thesis true? Suppose there is a musical score *S* such that by applying to it a particular special skill—namely, the skill of silent score reading—one can imagine, and, hence, experientially engage with, a performance of *S* that manifests the primary properties that must be experienced to fully appreciate the corresponding musical work *M*.¹⁷ Then *S* makes available for experience the primary properties that must be experienced to fully appreciate *M*. Furthermore, it can be assumed that *S* stands in an appropriate historical-intentional relation to *M*. Given this, if the foregoing ac-

¹⁷Many musical scores are too complex for score reading. So *S* is not *any* musical score.

count of “an instance of an artwork” is true, then S is an instance _{p/e} of M . However, this result does not correspond to the actual use of the expression “an instance of an artwork,” since according to this use, scores of musical works are not instances of these works (the consensus is that instances of musical works are solely musical performances (as well as perhaps their surrogates—playings of recordings of musical performances)).

Alternatively, that the foregoing account of “an instance of an artwork” does not accord with the actual use of the expression “an instance of an artwork” can be shown as follows. Suppose there is a verbal description V of some very simple drawing D —say, a drawing of a black square. Suppose also that by applying a special skill—namely, the skill of reading—to V , one can mentally form, and, hence, experientially engage with, an accurate image of D —an image that makes it possible to perceptually grasp the primary properties that must be experienced to fully appreciate D . Then V makes available for experience the primary properties that must be experienced to fully appreciate D . Furthermore, it can be assumed that V stands in an appropriate historical-intentional relation to B . Thus, if the definition of “an instance of an artwork” being discussed is true, then V is an instance _{p/e} of B . However, this does not accord with the actual use of the expression “an instance of an artwork,” since according to this use, no verbal description of a drawing is an instance of this drawing.

Thus, the foregoing alternative account of “an instance of an artwork” contradicts the actual use of the expression “an instance of an artwork.” As a result, the thesis that the actual use of this expression supports this account is false. Mean-

while, if this thesis is false, then, since, as already mentioned, the objection being discussed is based on this thesis, this objection fails.

Are there any other potentially successful objections to the idea that Davies (2010)'s definition covers all and only those entities that fall under "an instance of an artwork"? No—or so it seems. As a result, given that, as already mentioned, there are no potentially successful objections to the idea that Davies (2010)'s definition is sufficiently informative, there seems no reason to reject this definition. At the same time, there is reason to consider it satisfactory. First, as already mentioned, it seems to satisfy one of the criteria of a successful definition—that of being sufficiently informative. Furthermore, there is reason to hold that it also satisfies the second criterion—that of covering all and only those entities that fall under the expression being defined. As pointed out in Footnote 1, the expression "an instance of an artwork" is assumed to be located within a particular context—namely, the context of the ontology of artworks. So Davies (2010)'s definition reflects the actual use of this expression if it reflects this use *by ontologists of art*. Does Davies (2010)'s definition reflect the latter use? The answer to this question, I think, is "Yes." Many ontologists of art—including, Currie (1989), Danto (1981), Davies (2003b), Levinson (1980), Nannicelli (2013), and Wollheim (1980)—use the expression "an instance of an artwork" according to the proveniential version of the definition, with regard to all artworks. At the same time, a number of ontologists of art use this expression according to the purely epistemic version of the definition, with regard to at least some artworks. For example, Dodd (2000) uses it that way when he talks about instances of musical works, and Goodman and Elgin (1987) use it that

way with regard to instances of notational artworks (such as literary and musical works).¹⁸

Thus, Davies (2010)'s definition seems to reflect the actual use of the expression "instance of an artwork." Meanwhile, if this is so, then there is reason to think that this definition covers all and only those entities that fall under "an instance of an artwork."

2.3 Defining "a Well-Formed Instance of an Artwork"

Davies (2010)'s account of "an instance of an artwork" can be elaborated further. Note that this account does not specify whether an instance_{p/e} of an artwork is capable of manifesting *all* the primary properties that must be experienced to fully appreciate this work. Can such an instance_{p/e} (in principle) manifest all such properties? Apparently, the only plausible reason to answer "No" is that there are no entities capable of manifesting all the primary properties that must be experienced to fully appreciate artworks. But this reason is unsatisfactory. If there are no entities capable of manifesting all the primary properties that must be experienced to fully appreciate artworks, then no artwork can be fully appreciated. But this consequence is doubtless false. A complete appreciation of an artwork may be hard, but in an overwhelming majority of cases, it is, at least in principle, possible.¹⁹

So there seems no real reason to think that an instance_{p/e} of an artwork cannot

¹⁸At the same time, according to Goodman and Elgin (1987), non-notational (analog) artworks, such as etchings and paintings, have instances_p, not instances_e.

¹⁹This is not to say, of course, that *every* artwork can be fully appreciated.

manifest all the primary properties that must be experienced to fully appreciate this work. At the same time, there is a good reason to uphold the opposite thesis—that an instance_{p/e} of an artwork can, in principle, manifest all of these properties. The reason is that according to the consensus among ontologists of art, for most (though not all) artworks, there, in fact, existed, exist now, or will exist instances_{p/e} that manifest all the primary properties that must be experienced to fully appreciate these works. Thus, taking into account what has been said, it is reasonable to conclude that there can be instances_{p/e} of an artwork that are capable of manifesting all the primary properties that must be experienced to fully appreciate this work.

Given the foregoing result, we are justified in adding to Davies (2010)’s account the definition of “an instance_{p/e} that can manifest all the primary properties that must be experienced to fully appreciate the corresponding artwork”—or, in other words, the definition of “a *well-formed* instance_{p/e} of an artwork”:²⁰

Well-formed instance_e: For all x , x is a well-formed instance_e of an artwork if and only if x manifests all the primary properties that must be experienced to fully appreciate this work.

Well-formed instance_p: For all x , x is a well-formed instance_p of an artwork if and only if x (a) manifests all the primary properties that must be experienced to fully appreciate this work and (b) stands in an appropriate historical-intentional relation to it.

²⁰This kind of instance_{p/e} could also be characterized as “strict,” or “genuine,” or “perfect,” or “ideal.”

2.4 Defining “a Non-Well-Formed Instance of an Artwork”

The only difference between well-formed instances_{p/e} of an artwork and instances_{p/e} of an artwork is that well-formed instances_{p/e} manifest *all* the primary properties that must be experienced to fully appreciate the work, whereas instances_{p/e} manifest *certain* primary properties that must be experienced to fully appreciate the work. In light of this, one might ask: Are the class of instances_{p/e} and the class of well-formed instances_{p/e} coextensive? Put otherwise, are all instances_{p/e} well-formed? Consider a slightly damaged print of a photograph or a musical performance that contains one incorrect note. Clearly, neither the performance nor the print provide access to *all* the primary properties that must be experienced to fully appreciate the corresponding works. But, at the same time, both the print and the performance provide access to a significant set of such properties. As a result, it seems reasonable to think that both of them can be (a) non-well-formed instances_e and—assuming that each of them can stand in an appropriate historical-intentional relation to the corresponding work—(b) non-well-formed instances_p. And, in fact, most ontologists of art do think so. Given this, the above question, I think, should be answered in the negative.

In light of the fact that the class of instances_{p/e} is not exhausted by well-formed instances_{p/e}, a natural question arises: How can “a non-well-formed instance_{p/e} of an artwork” be defined? As is clear from what has been said in the previous paragraph, a non-well-formed instance_{p/e} of an artwork manifests only some of the primary properties that must be experienced to fully appreciate this work. Can this fact

alone be used to define “a non-well-formed instance_{p/e} of an artwork”? No—for there are entities that (a) manifest only some primary properties that must be experienced to fully appreciate an artwork but (b) are not instances_{p/e} of this work. Consider, for example, a black and white image of a color painting. Such an image cannot manifest any color properties. Meanwhile, such properties are doubtless among the primary properties that must be experienced to fully appreciate any color artwork. So a black and white image cannot manifest *all* the primary properties that must be experienced to fully appreciate a color painting. But such an image doubtless can manifest some of these properties—for example, those that are concerned with the shapes of what is depicted in this painting. Thus, a black and white image can manifest some of the properties that must be experienced to fully appreciate a color painting. At the same time, as is generally agreed, no such image can be an instance_{p/e} of such a painting.

So what distinguishes non-well-formed instances_{p/e} of artworks from entities that are not such instances? Taking into account the art-ontological context, this question, I think, can be answered as follows: Unlike entities that are not instances_{p/e} of artworks, non-well-formed instances_{p/e} of artworks (a) manifest sufficiently many, though not all, of the primary properties that must be experienced to fully appreciate these works and (b) can, in principle, manifest all sensory kinds of such properties (by “a sensory kind of properties” is meant any kind of properties that is relevantly concerned with a sensory modality—for example, visual properties (i.e., properties related to vision), auditory properties (i.e., properties related to hearing), and olfactory properties (i.e., properties related to olfaction)). Given this, “a

non-well-formed instance of an artwork” can be defined as follows:

Non-well-formed instance_e: For all x , x is a non-well-formed instance_e of an artwork just in case x (a) manifests sufficiently many, though not all, of the primary properties that must be experienced to fully appreciate this work and (b) could, in principle, manifest all sensory kinds of such properties.

Non-well-formed instance_p: For all x , x is a non-well-formed instance_p of an artwork just in case x (a) manifests sufficiently many, though not all, of the primary properties that must be experienced to fully appreciate this work, (b) could, in principle, manifest all sensory kinds of such properties, and (c) stands in an appropriate historical-intentional relation to this work.

Note that one of the features that any non-well-formed instance_{p/e} must possess is the feature of being, in principle, capable of manifesting all sensory kinds of the primary properties that must be experienced to fully appreciate the corresponding artwork. Thus, a performance of a musical work that contains an incorrect note but can, in principle, manifest all sensory kinds of these properties can be a non-well-formed instance_{p/e} of this work. At the same time, given some plausible assumptions, a playing of an audio recording of a live performance of a work of classical music cannot be a non-well-formed instance_{p/e} of at least some classical musical works. Such a playing, being non-visual, cannot, in principle, manifest any visual properties.²¹ Meanwhile, the primary properties that must be experienced to

²¹Note that what is said here applies only to playings of “audio only” (non-video) recordings. Perhaps playings of audio-video recordings—recordings that capture both the sonic and the visual

fully appreciate a work of classical music often include certain visual properties.²² So in some cases, a playing of an audio recording of a live performance of a work of classical music cannot manifest all sensory kinds of the primary properties that must be experienced to fully appreciate this work.

Likewise, a soundless screening cannot be a non-well-formed instance_{p/e} of a sound film. The primary properties that must be experienced to fully appreciate such a film include particular sonic properties. But no soundless screening can manifest any such properties. So no such screening can manifest all sensory kinds of the primary properties that must be experienced to fully appreciate a sound film.

Here, one could object as follows. The foregoing definition implies that being, in principle, capable of manifesting all sensory kinds of the primary properties that must be experienced to fully appreciate an artwork is necessary to be a non-well-formed instance of this work. However, this implication is false. There can be non-well-formed instances that can manifest only some of the sensory kinds of the mentioned properties.

Is this objection successful? It assumes that it is possible for a non-well-formed instance of an artwork to be incapable of manifesting all sensory kinds of the primary properties that must be experienced to fully appreciate this work. This assumption, however, is problematic. According to a widely accepted view, non-

aspects of a performance—can manifest all the primary properties that must be experienced to fully appreciate the corresponding works. For a discussion of this possibility, see Mag Uidhir (2007).

²²For evidence that can be used to support this claim, see Bergeron and Lopes (2009), S. Davies (2001), Kivy (2002), Mag Uidhir (2007), and Nanay (2012).

well-formed instances of artworks are slightly incorrect well-formed instances of these works. Meanwhile, a *slightly* incorrect well-formed instance of an artwork is doubtless capable of manifesting all sensory kinds of the primary properties that must be experienced to fully appreciate this work. So if the abovementioned view is true—and there seems no reason to think otherwise—non-well-formed instances of artworks must be, in principle, capable of manifesting all sensory kinds of the primary properties that must be experienced to fully appreciate these works.

Here is another consideration against the assumption being discussed. If this assumption is true, then there must be non-well-formed instances that are, in principle, incapable of manifesting all the sensory kinds of the primary properties that must be experienced to fully appreciate the corresponding artworks. What entities could serve as such instances? Here are some possible candidates:

- (a) a silent screening of a sound film
- (b) a playing of the sound of a film
- (c) a silent performance of a sound play
- (d) a purely sonic performance of a play
- (e) a purely sonic performance of a musical
- (f) a silent performance of a musical

But can (a)–(f), in fact, serve as non-well-formed instances that are, in principle, incapable of manifesting all the sensory kinds of the primary properties that must be experienced to fully appreciate the corresponding artworks? To be such instances,

(a)–(f) must satisfy two conditions. First, they must be, in principle, incapable of manifesting at least one sensory kind of the primary properties that must be experienced to fully appreciate the corresponding artworks. Second, they must manifest sufficiently many, but not all, of these properties. There is no doubt that (a)–(f) satisfy the first condition. But do they satisfy the second one? Prima facie, (a)–(f) can manifest sufficiently many of the primary properties that must be experienced to fully appreciate the corresponding artworks only if perceiving (a)–(f) enables us to adequately (though, of course, not fully) appreciate these works. However, we cannot adequately appreciate a film solely by watching its silent screening or by listening to its sound; likewise, we are unable to adequately appreciate an (audible) play or a musical just by watching their silent performances or just by listening to the sound of their performances. Thus, there is good reason to hold that (a)–(f) do not manifest sufficiently many primary properties that must be experienced to fully appreciate the corresponding artworks. Meanwhile, if this is so, then (a)–(f) cannot be non-well-formed instances.

Now, what has been said about (a)–(f) can, I think, be said about any other potential candidates for the role of non-well-formed instances that are, in principle, incapable of manifesting all the sensory kinds of the primary properties that must be experienced to fully appreciate the corresponding artworks. So there are no non-well-formed instances that are, in principle, incapable of that—and, hence, the assumption being discussed is false.

2.5 Additional Remarks

In closing, a few additional remarks.

Remark 1. In the philosophical literature, there is an expression that is closely related to “an instance of an artwork”—namely, “a token of an artwork.” This latter expression can be defined as follows: For all x , x is a token of an artwork just in case (a) this work is a type (where a type can be understood as an entity that can have multiple instances _{p/e}) and (b) x is its instance _{p/e} .

Remark 2. The account presented in Sections 2.1–2.3 has certain implications with regard to the existence and identity conditions of artworks. Suppose this account is true. Then:

- (1) The existence of an instance _{p/e} of an artwork entails that this work exists.

Substantiation. Suppose there is an instance _{p/e} of some artwork A . Then A can be adequately appreciated. But if this is so, then there is no real reason to deny the existence of A .

- (2) The fact that an instance _{p/e} of an artwork does not exist does not entail that this work does not exist.²³

Substantiation. Suppose there is no performance, reproduction of a performance, or playing of a recording of a performance of some musical work M .

²³This does not imply, of course, that *any* artwork can exist if none of its instances _{p/e} exist. According to a widely accepted view, a painting is identical to its only instance _{p/e} —the canvas. If this view is correct, then a painting cannot exist if no instances _{p/e} of this painting exist.

Then there are no instances _{p/e} of M . Suppose next that there is an encoding of M —say, a copy of M 's score or a recording of a performance of M . Does M exist in this case? Prima facie, the answer is “Yes.” M exists qua an entity that is, in some sense, contained in that encoding. Thus, a musical work can exist even if there are no instances _{p/e} of this work.

- (3) If (a) there is an instance _{p} of some artwork A and an instance _{p} of some artwork B and (b) these instances (i) manifest the same primary properties that must be experienced to fully appreciate A and B and (ii) stand in the same historical-intentional relation to A , then A is identical to B .

Substantiation. Suppose there is an instance _{p} of some artwork A and an instance _{p} of some artwork B . Suppose next that these instances (i) manifest the same primary properties that must be experienced to fully appreciate A and B and (ii) stand in the same appropriate historical-intentional relation to A . Then, since nothing can stand in an appropriate historical-intentional relation to *more than one* artwork, both of them must be instances _{p} of one and the same work. But if this is so, then A must be identical to B .

- (4) The fact that (a) there is an instance _{e} of some work A and an instance _{e} of some work B and that (b) these instances manifest the same primary properties that must be experienced to fully appreciate A and B does not entail that A and B are identical.

Substantiation. Consider Brahms's Piano Sonata Opus 2 (1852) and an (imaginary) “work identical with it in sound structure, but written by Beethoven”

(Levinson, 1980, 12).

Brahms's Piano Sonata Opus 2 (1852), an early work, is strongly *Liszt-influenced*, as any perceptive listener can discern. However, [the] work identical with it in sound structure, but written by Beethoven, could hardly have had the property of being Liszt-influenced. And it would have had a visionary quality that Brahms's piece does not have. (Levinson, 1980, 12)

Given what has been said, the foregoing works are not identical. Suppose now that there are some instances, I_1 and I_2 , that manifest all the primary properties that must be experienced to fully appreciate Brahms's work. Then, since, by assumption, Beethoven's work is identical in its sound structure to Brahms's work, I_1 and I_2 also manifest all the primary properties that must be experienced to fully appreciate Beethoven's work. Thus, it is possible (a) for an instance_e of an artwork A and an instance_e of an artwork B to manifest the same primary properties that must be experienced to fully appreciate A and B and (b) for A and B to be nonidentical.

- (5) If (a) there is an instance_{p/e} of some work A and an instance_{p/e} of some work B and (b) these instances manifest nonidentical sets of primary properties, then it cannot be inferred that A and B are nonidentical.

Substantiation. Suppose there is an instance_{p/e} of A and an instance_{p/e} of B . Suppose next that the primary properties manifested by the instance_{p/e} of A and the primary properties manifested by the instance_{p/e} of B are not the same.

Then, of course, these instances are not identical. But they can nevertheless be instances_{*p/e*} of the same artwork, since each of them can manifest different sufficient sets of the primary properties of the same work and stand in an appropriate historical-intentional relation to this work. Thus, in this case, *A* and *B* are not necessarily nonidentical.

Remark 3. Given what has been said in Sections 2.2 and 2.3, Davies (2010)’s definition of “an instance of an artwork” can be formulated in a way other than the way it is formulated in Section 2.1. As shown in Sections 2.2 and 2.3, an instance_{*p/e*} of an artwork can be well-formed or non-well-formed. At the same time, there can be no instances_{*p/e*} other than well-formed and non-well-formed ones. In light of this, as well as the definitions of “a well-formed instance_{*p/e*} of an artwork” and “a non-well-formed instance_{*p/e*} of an artwork,” “an instance of an artwork” can alternatively be defined as follows:

Instance_{*e*} (ALT): For all *x*, *x* an instance_{*e*} of some artwork *A* if and only if *x* is either:

- a well-formed instance_{*e*} of *A*—an entity that manifests all the primary properties that must be experienced to fully appreciate *A*; or
- a non-well-formed instance_{*e*} of *A*—an entity that (a) manifests sufficiently many, but not all, of the primary properties that must be experienced to fully appreciate *A* and (b) could, in principle, manifest all sensory kinds of these properties.

Instance_p (ALT): For all x , x an instance_e of some artwork A if and only if x is either

- a well-formed instance_e of A —an entity that (a) manifests all the primary properties that must be experienced to fully appreciate A and (b) stands in an appropriate historical-intentional relation to A ; or
- a non-well-formed instance_e of A —an entity that (a) manifests sufficiently many, but not all, of the primary properties that must be experienced to fully appreciate A , (b) could, in principle, manifest all sensory kinds of these properties, and (c) stands in an appropriate historical-intentional relation to A .

Clearly, the foregoing formulation provides a more detailed account of “an instance of an artwork” than the formulation given in Section 2.1. It should be underlined, however, that the former formulation does not differ extensionally from the latter one; that is, these formulations cover exactly the same set of entities.

PART II

Introduction to PART II

Now that “a novel” and “an instance of an artwork” have been defined, the main question of this dissertation—“What is the ontological status of novels and their instances?”—can be addressed. Clearly, to answer this question, it is sufficient to answer the following questions: “What is the ontological status of novels?” and “What is the ontological status of instances of novels?” The main goal in PART II is to answer the second of these questions.

Chapter 3: Against Inscriptions as Instances of Novels

Introduction

What is the ontological status of instances of novels¹? Put otherwise, what sort of entities can play the role of such instances? According to the view accepted by an overwhelming majority of theorists, including Noel Carroll, Arthur Danto, David Davies, Stephen Davies, John Dilworth, Nelson Goodman, Peter Lamarque, Jerrold Levinson, Christy Mag Uidhir, Aaron Meskin, Richard Wollheim, and Lee Walters, the paradigmatic, or most typical (though not the only), entities that play this role are *inscriptions*—concrete (usually physical) texts written or printed on something (say, paper, papyrus, or parchment) or displayed on the screen of some device (such as a computer or an e-reader).^{2,3} My goal in this chapter is to show that this view, which I will call “Orthodox,” is misguided. I begin with a formulation and defense of an argument against the strong version of the Orthodox View—the version according to which inscriptions are *well-formed* instances of novels⁴ (hereafter: “the Strong

¹Here and in what follows, by “novels,” I mean novels that can, in principle, be read aloud.

²[See the appendix to this chapter (Section 3.5).]

³Although this view is accepted by an overwhelming majority of theorists, it is not accepted by all of them (thus, it is rejected by Kivy (2006)).

⁴As shown in Chapter 2, “a well-formed instance of an artwork” can be defined as follows:

Orthodox View”) (Section 3.1). Next, I provide a critical analysis of an alternative argument against this view and of potential ways to defend this argument, including those based on the ideas championed by Kivy (2006) and Urmson (2004) (Section 3.2). I then turn to an examination of a weaker version of the Orthodox View—the version according to which inscriptions are *non-well-formed* instances of novels⁵—arguing that this version (hereafter: “the Weak Orthodox View”) does not stand up to criticism (Section 3.3). Finally, I provide an argument against the Orthodox View *simpliciter* (Section 3.4).

Well-formed instance_e: For all x , x is a well-formed instance_e of an artwork if and only if x manifests all the primary properties that must be experienced to fully appreciate this work.

Well-formed instance_p: For all x , x is a well-formed instance_p of an artwork if and only if x (a) manifests all the primary properties that must be experienced to fully appreciate this work and (b) stands in an appropriate historical-intentional relation to it.

⁵As shown in Chapter 2, “a non-well-formed instance of an artwork” can be defined as follows:

Non-well-formed instance_e: For all x , x is a non-well-formed instance_e of an artwork if and only if x (a) manifests sufficiently many, but not all, of the primary properties that must be experienced to fully appreciate this work and (b) could, in principle, manifest all sensory kinds of these properties.

Non-well-formed instance_p: For all x , x is a non-well-formed instance_p of an artwork if and only if x (a) manifests sufficiently many, but not all, of the primary properties that must be experienced to fully appreciate this work, (b) could, in principle, manifest all sensory kinds of these properties, and (c) stands in an appropriate historical-intentional relation to this work.

3.1 Against the Strong Orthodox View

According to the definition of “a well-formed instance of an artwork,” a *well-formed instance_e* of an artwork is whatever manifests all the primary properties that must be experienced to fully appreciate this work, and a *well-formed instance_p* of an artwork is whatever (a) manifests all the primary properties that must be experienced to fully appreciate this work and (b) stands in an appropriate historical-intentional relation to it. Thus, regardless of whether the expression “a well-formed instance of an artwork” is used in the epistemic or the proveniential sense, to be a well-formed instance⁶ of some artwork, it is necessary to manifest all the primary properties that must be experienced to fully appreciate this work. Recall now that according to the Strong Orthodox View, inscriptions are well-formed instances of novels. So if this view is true, then inscriptions must be capable of manifesting all the primary properties that must be experienced to fully appreciate the corresponding novels. This consequence, however, is problematic. For, in fact, no inscription can manifest *all* of the properties of the mentioned kind. The reason for this is that there is (at least) one class of primary properties that (a) must be experienced to fully appreciate a novel but (b) cannot be made manifest by any inscription: the sonic properties of the novel, or, in other words, the properties related to how the

⁶*Terminological note:* If it is not specified whether the expression “instance” is used in the purely epistemic or the proveniential sense, then this expression can be used in either of these senses. Also, it is assumed that, regardless of whether “instance” is used in the purely epistemic or the proveniential sense, it is used in one and the same sense throughout the relevant passage.

novel sounds.⁷ Clearly, the thesis that the sonic properties of a novel are primary properties that (a) must be experienced to fully appreciate this novel but (b) cannot be manifested by any inscription is far from obvious and, therefore, requires substantiation. To substantiate this thesis, it is sufficient to show that these properties (a) cannot be manifested by any inscription, (b) must be experienced to fully appreciate the novel, and (c) are primary.

3.1.1 Why Inscriptions Cannot Manifest Sonic Properties

Let us first show that the sonic properties of a novel cannot be manifested by inscriptions. According to the definition of the expression “to manifest a property,”⁸ an entity manifests a property only if this property is apprehensible by directly perceiving this entity. So if sonic properties can be manifested by an inscription, they must be apprehensible by directly perceiving it. But can they, in fact, be apprehended this way? For a property to be apprehensible by directly perceiving an inscription, this property must be apprehensible by means of directly applying *some* sensory faculty to this inscription. What is this faculty, in the case of sonic properties? Presumably, it can only be the faculty of hearing, for no other faculty can be used to adequately grasp sonic properties when it is directly applied to an inscription. So sonic properties can be apprehended by directly perceiving an inscription only if they can be apprehended through *hearing* it. However, an inscription (as

⁷Examples of sonic properties are “being sonorous,” “being mellifluous,” “having a particular rhythm,” “having a particular sounding,” etc.

⁸See Chapter 2, Section 2.1.

opposed to a reading of it) cannot be heard. The reason for this is that (a) (strictly speaking) only sounds can be heard, but (b) an inscription is not a sound (according to our definition,⁹ it is a concrete text written or printed on something or displayed on the screen of some electronic device). Thus, sonic properties cannot be apprehended by directly perceiving an inscription and so cannot be manifested by it.

One might object as follows. The foregoing argument assumes that the only faculty that can be used to adequately grasp sonic properties by means of a direct application of this faculty to an inscription is the faculty of hearing. But this assumption is false. For, in fact, there is another faculty that can be used this way—the faculty of aural imagination.

The foregoing objection, however, fails. It assumes that the faculty of aural imagination can be directly applied to inscriptions. But this assumption is false. Recall that inscriptions are, by definition, concrete. So, since any concrete entity is either physical (spatiotemporal) or mental, every inscription is either physical or mental.¹⁰ Can the faculty of aural imagination be directly applied to a physical inscription? No. In order for this faculty to be directly applicable to a physical inscription, it must be possible to mentally “hear” such an inscription. However, (strictly speaking) no *physical* object (as opposed to a mental object—such as an imaginary sound or a sequence of imaginary sounds) can be *mentally* “heard.” Can

⁹See the introduction to this chapter.

¹⁰It is safe to assume that all existent inscriptions of novels are physical. The argument, however, does not depend on this assumption.

the faculty of aural imagination be directly applied to a mental inscription (that is, an imaginary text)? Again, the answer is “No.” The faculty of aural imagination is directly applicable to a mental inscription only if the latter can be mentally “heard.” However, a mental inscription cannot be mentally “heard,” since (a) such an inscription cannot be a mental *sound* (according to our definition, it can only be a mental *text*—a sequence of written or printed symbols), but (b) (strictly speaking) only mental sounds can be mentally “heard.” Thus, the faculty of aural imagination cannot be directly applied to physical inscriptions; nor can it be directly applied to those inscriptions that are mental. But then, given that any inscription is either physical or mental, this faculty cannot be directly applied to any inscriptions whatsoever.

So we have established that the sonic properties of a novel cannot be manifested by inscriptions. Before proceeding, it is worth emphasizing two things. First, although an inscription does not manifest sonic properties, it can be characterized as having such properties in some non-manifesting way—for example, as *encoding* them (similar to how a musical score encodes the sonic properties of a musical work). Second, although sonic properties are not manifestable by inscriptions, some other properties are. Consider, for instance, visual properties. For these properties to be apprehensible by directly perceiving an inscription, they must be apprehensible by seeing it. And they doubtless can be apprehended in this latter manner. So, since they are apprehensible by directly perceiving an inscription, they can be manifested by it.

3.1.2 Why the Sonic Properties of a Novel Must Be Experienced to Fully Appreciate This Novel

Let us now show that the sonic properties of a novel must be experienced to fully appreciate it. Clearly, among the factors that determine the aesthetic value of a novel, the factor of how this novel sounds is not the most important. (Presumably, the most important factors pertain to the novel's content and structure.) But this should not lead us to think that the sonic dimension is completely irrelevant to the aesthetic value of a novel. Although this dimension is not the main determinant of a novel's aesthetic significance, it is, nevertheless, a determinant. One reason to think so comes from an observation of our professional literary community. Consider novelists. Many of them use sound techniques—such as alliteration, assonance, consonance, rhythm, etc.—in their writing, and presumably intentionally. Here are some examples of such use:

- **Alliteration:** “So we *beat* on, *boats* against the current, *borne back* ceaselessly into the past” (F. Scott Fitzgerald, *The Great Gatsby*); “You may say that *dear diligent Dexter* gets carried away in his job. . .” (Jeff Lindsay, *Darkly Dreaming Dexter*); “. . .neither of these can feel *stranger* and *stronger* emotions than the man does, who for the first time finds himself pulling into the *charmed, churned* circle of the hunted sperm whale” (Melville, *Moby-Dick*).
- **Assonance:** “*Lolita, light of my life, fire of my loins*” (Vladimir Nabokov, *Lolita*); “And stepping softly with her air of blooded ruin about the glade

in a frail agony of grace she trailed her rags through dust and ashes, circling the dead fire, the charred billets and chalk bones, the little calcined ribcage” (Cormac McCarthy, *Outer Dark*); “Perhaps tonight—after a month of waiting—would be the night“ (J. K. Rowling, *Harry Potter and the Order of the Phoenix*).

- **Consonance:** “An oblong puddle inset in the coarse asphalt...” (Vladimir Nabokov, *Bend Sinister*); “When he woke in the woods in the dark and the cold of the night...” (Cormac McCarthy, *The Road*); “A loud, echoing crack broke the sleepy silence like a gunshot...” (J. K. Rowling, *Harry Potter and the Order of the Phoenix*).
- **Rhythm:** “Becky was always good to him, always amused, never angry“ (William Thackeray, *Vanity Fair*); “With the first gray light he rose and left the boy sleeping and walked out to the road and squatted and studied the country to the south. Barren, silent, godless. He thought the month was October but he wasn't sure. He hadn't kept a calendar for years. They were moving south. There'd be no surviving another winter here” (Cormac McCarthy, *The Road*); “Then she took off the hank and looked me straight in the face, and very pleasant, and says, ‘Come on, now, what's your real name?’” (Mark Twain, *The Adventures of Huckleberry Finn*); “TJ's brother is Rex, age forty-four, my second son, currently married to a stripper. Amber is her name, a poor creature without a brain but with a large fake chest, who, I think, is his third wife. Second or third, but who am I to condemn?” (John Grisham, *The*

Testament).

Consider next literary critics. Although their primary focus is usually the plot of a novel, they sometimes attend to how the novel sounds. Here are a few quotes that illustrate the point:

- “[Stevenson’s] writing...remains true to its musical principles. It is the result of trained ear and recognition of language as a conscious instrument...” (Swinerton, 1915, 87).
- “In Twain’s own terms, Cooper actually *is* a word-musician; he does create a sound in his prose that we cannot ignore...” (Kowalewski, 1993, 72).
- “... the opening pages of *A Farewell to Arms* do merge a realist aesthetic with a valorization of form and the musical aspects of language that one might expect of a Mallarmé prose poem” (Eby, 2013, 177).
- “All readers have been affected by Fitzgerald’s style, for Fitzgerald was marvelously sensitive to the sounds and cadences of language” (Lewis, 1985, 89).
- “From youth, Truman Capote was a master of literary style, writing with a true sense of the sound, rhythm and texture of language, with a feel for trenchant detail and metaphor” (Hicks, 2009, 94).
- “The rhythm of Cervantes is incredible. Not always, obviously, because what I call pulsation, or rhythm should have, as in music, moments of climax, moments of extreme tension that cannot be sustained, otherwise they would end up being monotonous” (Yúrkievich, 2003, 227).

- “Balzac was as devoted to style for its own sake as Malherbe, and had the same narrow oratorical ideal of correctness, the same devotion of order, dignity, and sonorous rhythm” (Grierson, 1906, 273).

Finally, it is not uncommon for literary theorists to stress the importance of sonic elements in prose. Thus, according to Aristotle (2010), “the form of a prose composition” should not be “destitute of rhythm” (Aristotle, 2010, 139). Gustave Flaubert notes that “a good prose sentence must be like a good line of verse, unchangeable, as rhythmic, as sonorous” (Llosa, 1987, 219). And according to R. L. Stevenson (2011):

In all ideal and material points, literature, being a representative art, must look for analogies to painting and the like; but in what is technical and executive, being a temporal art, it must seek for them in music. Each phrase of each sentence, like an air or a recitative in music, should be so artfully compounded out of long and short, out of accented and unaccented, as to gratify the sensual ear. And of this the ear is the sole judge. (Stevenson, 2011, 9–10)

So our professional literary community emphasizes the sonic aspect of the novel: Novelists endow their texts with certain sonic properties; literary critics take into account the sound of a novel in their aesthetic evaluation; finally, literary theorists stress the importance of sound in literary prose. But why would our professional literary community emphasize this aspect if the latter were irrelevant to the aesthetic value of novels? The very fact of such emphasis suggests that the

aesthetic value of a novel depends, in part, on how the novel sounds.

Here, one might object that there is an analogous reason against the view that the sound of a novel is aesthetically important, namely this: Most ordinary readers do not, in fact, pay attention to the sonic aspect of the novels they read.¹¹ This reason, however, has much less credibility than the abovementioned reason in favor of this view. As is generally agreed, what is done/said by specialists is a lot more likely to be right than what is done/said by non-specialists.¹² Novelists, literary critics, and theorists are professionally involved in literary practices and, hence, may well be considered specialists in literature. At the same time, it is clear that ordinary readers are non-specialists in literature. So novelists, literary critics, and theorists are, most likely, right in stressing the importance of the sonic properties of a novel, whereas ordinary readers are, most likely, wrong in disregarding these properties.

Thus, that the professional literary community emphasizes the sonic aspect of novels provides a reason to believe that the sound of a novel is a determinant of the aesthetic value of this novel. This is, however, not the only reason to believe so. Another such reason is concerned with certain syntactic, semantic, and pragmatic features of the text of a novel and their relation to the aesthetic value of this text. Before stating this reason, we first need to clarify in what respects sentences¹³ can

¹¹It is worth noting that although, when reading silently, readers usually do not attend to the sound of what they read, they seem to generate this sound in their minds. For recent empirical evidence confirming this, see, e.g., Petkov and Berlin (2013).

¹²I assume here that the specialists and non-specialists are trustworthy.

¹³Hereafter, by “a sentence,” I mean a sentence-token, not a sentence-type.

differ from each other.

So in what respects can sentences do that? First of all, sentences can differ in their *appearance*. Whether they differ in this respect depends on what syntactic structures they have and what font they employ. If sentences have different syntactic structures or use different fonts, then they differ in their appearance; otherwise, they are identical with regard to their appearance.¹⁴ Below are examples of sentences that differ in their appearance because of the difference in their syntactic structure ((1a) and (1b)) and because of the difference in their fonts ((2a) and (2b)), as well as an example of sentences that do not differ in their appearance ((3a) and (3b):

(1a) The idea is obvious.

(1b) The idea is evident.

(2a) The problem can be solved.

(2b) *The problem can be solved.*

(3a) The problem can be solved.

(3b) The problem can be solved.

Second, sentences can differ in their *semantic meaning*. An example of sentences that differ in this respect is sentences (2a) and (1a)/(1b). And examples of sentences that do not differ in their semantic meaning are sentences (2a) and (2b)

¹⁴It is worth noting that sentences that are identical in their appearance may not be identical in certain other respects (for example, they may not be identical in their meaning).

and (3a) and (3b).¹⁵ Note that the fact that sentences differ in their semantic meaning does not necessarily entail that they differ in their appearance. Thus, consider the sentences:

(4a) We approached one of the banks.

(4b) We approached one of the banks.

These sentences have the same appearance. But they can have different semantic meanings. Suppose, for instance, that the first occurrence of the word “bank” is intended to mean “a river bank,” whereas the second occurrence of this word is intended to mean something else—say, “a particular kind of financial institution” or “a receptacle where something may be deposited for recycling.” Then (4a) and (4b) have different semantic meanings.

Third, sentences can differ in their phonology, or, in other words, in their *sonic properties*. Whether sentences differ in this respect depends on what phonemic structures they have. If sentences have different phonemic structures, then they differ in their sonic properties; otherwise, they are identical with regard to their sonic properties. Examples of sentences that differ in their phonemic structures—and, hence, in their sonic properties—are sentences (1a), (1b), (2a)/(2b), (3a)/(3b), and (4a)/(4b). Examples of sentences that do not differ in their phonemic structures—and, hence, in their sonic properties—are sentences (2a) and (2b); (3a) and (3b); and (4a) and (4b). Note that the fact that some sentences differ in their sonic

¹⁵Assuming that the expressions “obvious” and “evident” are exact synonyms, (1a) and (1b) can also serve as an example of sentences that mean the same.

properties does not always entail that these sentences differ in their appearance or meaning. Thus, both (1a) and (1b) as well as the following sentences:

(5a) There is no sewer [in the sense of a person who sews] outside.

(5b) There is no sewer [in the sense of an artificial conduit for carrying out waste water] outside.

—differ in their sonic properties. However, (1a) and (1b) do not differ in their meaning (assuming that the expressions “evident” and “obvious” are exact synonyms), and (5a) and (5b) do not differ in their appearance.

Fourth, sentences can differ in *how natural (or “normal”) they appear to a competent speaker*. Examples of sentences that do not differ in this respect are sentences (1a) and (1b); (2a) and (2b); (3a) and (3b); (4a) and (4b); and (5a) and (5b). Here is an example of sentences that differ in the foregoing respect:

(6a) John picked up his wallet.

(6b) John picked up John’s wallet. [It is assumed that all occurrences of “John” refer to the same individual.]

(6a) does not sound odd to us. On the contrary, we find it quite natural. When faced with a necessity to express the fact that John picked up his wallet, most of us would say “John picked up his wallet.” At the same time, (6b) appears to us considerably less natural than (6a). We would not normally say “John picked up John’s wallet” if we had to describe John’s picking up his wallet.¹⁶

¹⁶Here, of course, the following question arises: Why do we find (6b) less natural than (6a)?

Fifth, sentences can differ in *the structure-based difficulty of comprehension*, or, in other words, in the difficulty of understanding them that arises from the way they are structured. Examples of sentences that do not differ in this respect are sentences (1a) and (1b); (2a) and (2b); (3a) and (3b); (4a) and (4b); (5a) and (5b); and (6a) and (6b). Here is an example of sentences that differ in the foregoing respect:

(7a) “My first kiss will always be recalled by me as how my romance with Shayna was begun” (King, 2000, 117).

(7b) “My romance with Shayna began with our first kiss. I’ll never forget it” (King, 2000, 117).

Clearly, (7a) and (7b) differ in how difficult it is to comprehend them: Comprehending (7b) is doubtless easier than comprehending (7a). The reason for this is not concerned with the meaning of (7a) and (7b), for they express roughly the same idea. Rather, the reason is related to the way they are structured (in particular, the fact that (7a) involves the use of passive voice, whereas (7b) employs active voice). Thus, the difference in how difficult it is to comprehend (7b) and (7a) is structure-based.

Sixth, sentences can differ in their *conversational implicatures*, where by “a conversational implicature” is meant an implicature¹⁷ that depends, for what it is,

Although this question is worth investigating, it is tangential to our main purposes, and so I will leave it unanswered.

¹⁷An implicature *simpliciter* is what is implied by a sentence within a particular context. Thus, suppose someone says:

not on the meaning of the sentence uttered but *solely* on the conversational context in which this sentence is placed.^{18,19} Thus, consider the sentences:

(9a) I am not feeling well.

(9b) I am not feeling well.

Suppose (9a) is Jack's answer to the question "Are you going to the party?" and (9b) is Jill's answer to the question "Are you running a marathon tomorrow?" Then the implicature of (8a) is that Jack is not going to the party, whereas the implicature of (9b) is that Jill is not running a marathon tomorrow. Thus, given the context,

(8) John "is an Englishman; he is, therefore, brave" (Grice, 1975, 44).

Clearly, this sentence implies that John's being an Englishman entails that he is brave. Hence, the implicature here is the idea that "it follows from [John's] being an Englishman that he is brave" (Grice, 1975, 75).

¹⁸An example of a conversational implicature is the implicature "Jill is not going to the colloquium" contained in the following dialog:

Jack: Are you going to the colloquium?

Jill: I am not feeling well.

The reason why this implicature is conversational is that it is generated *solely* by the features of the conversational context of the dialog (note that it is not part of the meaning of Jill's answer, for, strictly speaking, the latter does not contain the idea that Jill is not going to the colloquium)

¹⁹It is worth mentioning that a conversational implicature should be distinguished from what is known as a *conventional* implicature. The latter kind of implicature is an implicature that is part of the linguistic meaning of a sentence. An example of such an implicature is the implicature "his being an Englishman entails that he is brave" contained in (8). Another example is the implicature "some books are not interesting" contained in the sentence "Not all books are interesting."

(9a) and (9b) differ with regard to their conversational implicatures. (It is worth noting that (9a) and (9b) do not differ with regard to their appearance and semantic meaning (or so we can assume). Furthermore, since (9a) and (9b) semantically mean the same, they do not differ with regard to their conventional implicatures either.)

In addition to the mentioned respects, sentences can differ in their *aesthetic value*. For instance, a sentence can provide a more colorful description of a character, or be more pleasing to the ear, or be more visually pleasing (say, as a result of the fact that it has been calligraphed) than some other sentence. It should be stressed, however, that the difference in aesthetic value supervenes on the differences in one of the abovementioned respects.²⁰ If some sentences differ in their aesthetic value, then they differ in (a) their appearance, (b) their semantic meaning, (c) their sonic properties, (d) how natural they appear to a competent speaker, (e) their structure-based difficulty of comprehension, (f) their conversational implicatures, or (g) some combination of the mentioned respects.²¹ For instance, if some sentences

²⁰Note that this does not imply that the difference in aesthetic value is *reducible* to the differences in one of those respects.

²¹One could object as follows. Suppose there are sentences S_1 and S_2 that do not differ in (a) their appearance, (b) their semantic meaning, (c) their sonic properties, (d) how natural they appear to a competent speaker, (e) their structure-based difficulty of comprehension, and (f) their conversational implicatures. Then S_1 and S_2 can still differ aesthetically—if they differ in their rhythmic features.

However, this objection fails. It assumes that S_1 and S_2 can differ in their rhythmic features. But this assumption is false. By assumption, S_1 and S_2 do not differ in their appearance and sonic properties. Meanwhile, if sentences that do not differ in their appearance and sonic properties, then these sentences have the same rhythmic features.

are different because one of them is more visually pleasing than the other, then these sentences differ in their appearance; if some sentences differ in that one of them can provide a more colorful description of a character than the other, then these sentences differ in their semantic meaning; and if the difference between some sentences is that one of them is more pleasing to the ear than the other, then these sentences differ in their sonic properties.

Having characterized the respects in which sentences can differ from each other, we can now formulate the second reason in favor of the thesis that the sonic properties of a novel are relevant to the aesthetic value of this novel. Take any standard novel—a novel that has a text and perhaps some (non-textual) graphic elements. It seems plausible to suppose that this novel contains a sentence/set of sentences S_1 with aesthetic value V_1 and that there is a paraphrase of S_1 — S_2 —such that:

- (I) S_2 means the same or almost the same as S_1 .
- (II) S_2 does not differ from S_1 with regard to the structure-based difficulty of comprehension.
- (III) From a competent speaker's perspective, S_2 seems as natural (or "normal") as S_1 . (That is, when S_2 is encountered by a competent speaker, the latter does not have thoughts like "While we say S_1 , we do not normally say S_2 " or "It would be odd to say S_2 , as opposed to S_1 .")
- (IV) If S_1 has any conversational implicatures, then S_2 has the same conversational implicatures.

(V) S_2 is typed in the same aesthetically neutral black font (such as Times New Roman) as S_1 .

(VI) S_2 has an aesthetic value, V_2 , that is noticeably different from V_1 .

As mentioned above, if there is an aesthetic difference between sentences, then this difference can be explained in terms of differences in (a) their appearance, (b) their semantic meaning, (c) their sonic properties, (d) how natural they appear to a competent speaker, (e) their structure-based difficulty of comprehension, (f) their conversational implicatures, or (g) some combination of the mentioned respects. So the fact that S_1 and S_2 have different aesthetic values (V_1 and V_2 , respectively) should be explainable with the help of one or more of the foregoing kinds of differences. Which of these kinds is/are involved in the explanation? It is not the difference in meaning, since, by assumption, both S_1 and S_2 mean the same or almost the same. It is not the difference in the structure-based difficulty of comprehension because, by assumption, comprehending S_1 and S_2 is equally difficult/easy. It is not the difference in how natural S_1 and S_2 appear to a competent speaker, for, by assumption, they both appear to her equally natural. And it is not the difference in the conversational implicatures, since, by assumption, if S_1 and S_2 have conversational implicatures, these implicatures are identical.

Now, what about the difference in appearance? Can this difference explain the aesthetic difference between S_1 and S_2 ? Perhaps it could—if the visual properties of S_1 and S_2 were intrinsically relevant to the aesthetic values of S_1 and S_2 , respectively. However, in fact, neither the visual properties of S_1 nor the visual properties of S_2

are intrinsically relevant to those values. The visual properties related to the color, size, and shape of S_1/S_2 are not relevant to the aesthetic value of S_1/S_2 at all and, *a fortiori*, are not intrinsically relevant to this value. And the visual properties related to the meaning and sound of S_1/S_2 are relevant to the aesthetic value of S_1/S_2 only instrumentally—as something on which those properties that are intrinsically relevant to the aesthetic value of S_1/S_2 depend (or supervene). Of course, if S_1/S_2 were calligraphed or typed in a fancy, aesthetically interesting font, then S_1/S_2 would possess certain visual properties relevant to the aesthetic value of S_1/S_2 . However, according to (V), S_1 and S_2 are neither calligraphed nor typed in such a font.

Before proceeding, two remarks are worth making. First, while the thesis about the intrinsic irrelevance of the visual properties of the text of a novel to the aesthetic value of this novel applies to *typical* novels (which are typed in “ordinary” fonts), it does not apply to *all* novels. The aesthetic value of the text of some (non-standard) novel can intrinsically depend on the visual properties of this text (consider, for instance, a Japanese calligraphic novel). Second, the thesis that the visual properties of the text of a typical novel are not intrinsically relevant to the aesthetic value of this novel should not be confused with a stronger thesis—that the visual properties of a typical novel (as opposed to just the text of this novel) are not intrinsically relevant to the aesthetic value of this novel. This stronger thesis is false: A typical novel may well contain graphic elements (drawings, photographs, schemes, diagrams, etc.), and in this case, some of its visual properties (namely, the properties related to these graphic elements) *are* intrinsically relevant to its aesthetic value.

Thus, the aesthetic difference between S_1 and S_2 cannot be explained by appealing to the difference in their appearance. Nor can it be explained by appealing to the difference in their meanings, structure-based difficulty of comprehension, conversational implicatures, or in how natural they appear to a competent speaker. So how can the aesthetic difference between S_1 and S_2 be explained? It can be explained, I think, by appealing to the difference in their sonic properties. The reason why S_1 and S_2 have different aesthetic values is that S_1 and S_2 have different sonic properties. Meanwhile, if this is so, then the sonic properties of S_1 are relevant to its aesthetic value (for, otherwise, the aesthetic value of S_1 would be independent of the sonic properties of S_1 , and, hence, no difference in sonic properties could explain why S_1 is aesthetically different from S_2 or some other set of sentences). As a result, since S_1 is an essential part of a novel, these properties are also relevant to the aesthetic value of this novel.

Let us consider some possible applications of the foregoing reasoning, which might be called “the Paraphrasability Argument.” Consider Charles Dickens’s *A Tale of Two Cities*. This novel begins with the sentence:

S_t: “It was the best of times, it was the worst of times, it was the age of wisdom, it was the age of foolishness, it was the epoch of belief, it was the epoch of incredulity, it was the season of Light, it was the season of Darkness, it was the spring of hope, it was the winter of despair. . .”

(Dickens, 1902, 3).

Compare this sentence with the following paraphrase:

P_t: It was the best of times, the worst of times, the age of wisdom, the age of foolishness, the epoch of belief, the epoch of incredulity, the season of Light, the season of Darkness, the spring of hope, the winter of despair. . . .

From an intuitive viewpoint, P_t differs aesthetically from S_t. How can we account for this difference? The visual properties of S_t and P_t are not intrinsically relevant to the aesthetic values of S_t and P_t. So the aesthetic difference between S_t and P_t cannot be explained by saying that they have different visual properties. Can we explain this difference by appealing to the difference in meaning between S_t and P_t? If S_t and P_t differ in meaning, then, obviously, this difference can only result from the fact that in S_t, the expression “it was” occurs several times, whereas in P_t, this expression occurs only once. However, the mentioned fact does not really affect the meaning expressed in S_t/P_t (it does not matter, with regard to the meaning, whether S_t or P_t is used). Thus, the meanings of S_t and P_t are the same. But if this is so, then we cannot explain the aesthetic difference between S_t and P_t by appealing to the difference in their meanings.

Can we explain this difference by saying that the structure-based difficulty of comprehending S_t and the structure-based difficulty of comprehending P_t are different? No. Although S_t and P_t have different syntax, the structure-based difficulty of comprehending each of them is similar: With regard to their structures, comprehending S_t is not more difficult than comprehending P_t, and comprehending P_t is not more difficult than comprehending S_t.

Also, we cannot explain the aesthetic difference between S_t and P_t by saying that S_t and P_t have different conversational implicatures. First of all, even if S_t has a conversational implicature, nothing stops us from assuming that P_t has the same conversational implicature. Furthermore, there is no reason to assume that S_t or P_t has conversational implicatures. Recall that a conversational implicature is an implicature that depends, for what it is, not on the meaning of the sentence but solely on the conversational context in which this sentence is located. So in order for a conversational implicature of S_t to exist, there must be an implicature generated by the context in which S_t was placed by Dickens. However, there seems to be no such implicature. (If it exists, what is it?) Furthermore, a conversational implicature of P_t exists only if there is an implicature generated by the context in which S_t was placed. But there is no such implicature—or so we can assume.

Finally, we cannot explain the aesthetic difference between S_t and P_t by saying that from a competent speaker's viewpoint, P_t/S_t seems less natural than S_t/P_t . From a competent speaker's viewpoint, P_t does not seem less natural than S_t ; nor does S_t seem less natural than P_t . In other words, from such a viewpoint, the “naturalness” of P_t and the “naturalness” of S_t are the same (or, at least, very similar).

Thus, the aesthetic difference between S_t and P_t cannot be explained by appealing to the difference in their appearance, meanings, or conversational implicatures. Nor can it be explained by appealing to the difference in their structure-based difficulty of comprehension or to the difference in how natural they seem to a competent speaker. How can the aesthetic difference between S_t and P_t then be explained?

The only satisfactory way to explain it is by appealing to the difference in the sonic properties of S_t and P_t . Indeed, because of the lack of a particular rhythm, P_t does not sound as good as S_t ; P_t , one might say, is less gratifying to the ear than S_t . But if this is so—if the aesthetic difference between S_t and P_t is to be explained in terms of their sonic properties—then the aesthetic value of S_t and, hence, of *A Tale of Two Cities* depends on these properties.

Here is another application of the Paraphrasability Argument. Compare the following excerpt from John Grisham’s novel *Bleachers*:

S_f: “The bleachers were silent now, waiting” (Grisham, 2011, 8).

with its possible paraphrase:

P_f: The bleachers were silent now. They were waiting.

There seems to be an aesthetic difference between S_f and P_f . But what exactly is responsible for this difference? As before, and by the same reasoning, it can be shown that the aesthetic difference between S_f and P_f cannot be explained by appealing to the difference in their meanings, or visual properties, or conversational implicatures, or how natural they appear to a competent speaker, or their structure-based difficulty of comprehension. The only acceptable way to explain the aesthetic difference between S_f and P_f is by appealing to the difference in their sonic properties. Like in the previous case, in this case, it can be argued that, due to a particular rhythm, S_f sounds better than P_f and, hence, is more aesthetically pleasing. Meanwhile, if the aesthetic difference between S_f and P_f is to be explained in terms of the

sonic properties, then these properties are relevant to the aesthetic value of S_f and, consequently, to the aesthetic value of *Bleachers*.

Let us consider one more application of the Paraphrasability Argument. J. K. Rowling's *Harry Potter and the Philosopher's Stone* contains the following text:

S_h : "Mr. Dursley stopped dead. Fear flooded him. He looked back at the whisperers as if he wanted to say something to them. . . ." (Rowling, 2004, 9).

Compare this text with its possible paraphrase:

P_h : Mr. Dursley stopped dead, and fear flooded him, and he looked back at the whisperers as if he wanted to say something to the whisperers. . . .

Intuitively, there is an aesthetic difference between S_h and P_h . How can we explain this difference? Like in the previous case, in this case, it can be demonstrated, with the help of the reasoning used in the Dickens example, that the aesthetic difference between S_h and P_h cannot be explained by appealing to the difference in their meanings, or visual properties, or conversational implicatures, or how natural they appear to a competent speaker, or how difficult it is to comprehend them because of their structural complexity. The only satisfactory way to explain the aesthetic difference between S_h and P_h is by appealing to the difference in their sonic properties (primarily, the properties related to the rhythms of S_h and P_h). But if this is so—if S_h is aesthetically different from P_h by virtue of having certain sonic properties—then these properties are relevant to the aesthetic value of S_h and, therefore, to the aesthetic value of *Harry Potter and the Philosopher's Stone*.

In a similar way, the Paraphrasability Argument can be applied to any other standard novel (i.e., a novel that contains a text and perhaps some graphic elements). This, of course, requires us to assume that *any* such novel contains a sentence or set of sentences for which there is a paraphrase that (a) is not relevantly different in its appearance, (b) means the same or almost the same, (c) has a similar syntax-based difficulty of comprehension, (d) has the same conversational implicatures (provided that the original sentence or set of sentences has any conversational implicatures), (e) seems equally natural, from a competent speaker's perspective, but (f) possesses a different aesthetic value. But this assumption does not strike me as implausible. On the contrary, I think it is correct. I am not aware of any novel that does not contain a sentence or set of sentences for which there could be no such paraphrase.

Thus, the thesis that the sonic aspect of novels is relevant to their aesthetic value is supported by the Paraphrasability Argument as well as the fact that our professional literary community emphasizes this aspect. This, I believe, is sufficient to show that this thesis is true.

Now, if that is the case, then a complete appreciation of a novel is impossible without an experiential engagement with at least some of the sonic properties of this novel.²² For suppose the thesis that the sonic aspect of novels is relevant to their aesthetic value is true. Then the aesthetic value of a novel is, in part, determined by the sonic properties of this novel—and so we cannot fully appreciate

²²It is worth noting that the foregoing thesis can be used to derive another peculiar consequence—that the structure of a novel involves a sonic dimension. An account of how this can be done is provided in Chapter 7.

this value without appreciating these properties. Meanwhile, to fully appreciate a novel, we must fully appreciate the aesthetic value of this novel. Thus, we cannot fully appreciate a novel without appreciating its sonic properties. Now, can we fully appreciate the sonic properties of a novel without *experiencing* them? If we can, then we must be able to do this with the help of a mere description. But can we fully appreciate the sonic properties of a novel using a mere description? No. Appreciating the sonic properties of a novel solely by means of a description cannot enable us to fully appreciate them—similar to how appreciating the sonic properties of a musical work through a mere description is insufficient for a full appreciation of these properties. Thus, to fully appreciate the sonic properties of a novel, we must experience them. As a result, taking into account what has been said, we cannot fully appreciate a novel without experiencing its sonic properties.

Note that, in and of itself, our inability to fully appreciate a novel without experiencing the sonic properties of this novel does not entail that those who do not hear the physical sound of this novel—such as those who read silently—cannot experience these properties and, hence, cannot fully appreciate the novel. For suppose that to fully appreciate a novel, it is necessary to appreciate the sonic properties of this novel. Suppose also that these properties can be manifested not only by physical sound but also by mental sound (i.e., sound produced by imagination). In this case, those who do not hear the physical sound of a novel could experience the sonic properties of this novel by imagining and attending to the novel’s mental “sounding” and, hence, could fully appreciate this novel.²³

²³That the thesis that we must experience the sonic properties of a novel to fully appreciate this

3.1.3 Why the Sonic Properties of a Novel Are Primary

Having shown that inscriptions cannot manifest sonic properties and that the sonic properties of a novel must be experienced to fully appreciate this novel, let us show that the sonic properties of a novel are primary. Given what has been said above, it is uncontroversial that the sonic properties of a novel articulate a part of the artistic content of this novel—namely, the sonic dimension. And this dimension is primary, since it is not grounded in any other contentful properties. Meanwhile, according to the definition of “a primary property,” if properties articulate the primary artistic content—a particular set of the primary contentful properties (i.e., the contentful properties that are not grounded in any other contentful properties)—of an artwork, then these properties are primary.

3.1.4 Conclusion

So we have established that the sonic properties of a novel (a) are primary properties that must be experienced to fully appreciate this novel but (b) cannot be manifested by any inscription. Meanwhile, as mentioned in the introduction to this section, to be a well-formed instance of a novel, an inscription must be capable of manifesting all the primary properties that must be experienced to fully appreciate novel does not, in and of itself, entail that we cannot fully appreciate this novel without hearing certain physical sounds does not mean that we can fully appreciate a novel without hearing such sounds. In fact, in the following chapter, I argue that a full appreciation of novels requires listening to their physical sound.

this novel. So inscriptions cannot be well-formed instances of novels. But then the Strong Orthodox View—the view that inscriptions are well-formed instances of novels—fails.

3.2 Against the Strong Orthodox View: An Alternative Argument

The argument presented in Section 3.1 is not the only possible argument against the Strong Orthodox View. One could argue against this view as follows:

- (10) Given our ordinary reading and performing practices (that is, practices common among “ordinary” readers and musicians), inscriptions of novels (hereafter: “inscriptions”) and musical scores are relevantly analogous.
- (11) Musical scores are not well-formed instances of musical works.
- (12) If (10) and (11) are true, then inscriptions are not well-formed instances of novels.
- (13) Therefore, inscriptions are not well-formed instances of novels.

Call this “the Analogical Argument.”

Is the Analogical Argument acceptable? Its most controversial premise is premise (10). We do not normally treat musical scores as the ultimate perceptible objects of artistic appreciation; rather, we treat them as instructions (or prescriptions) for generating such objects—musical performances. Given this, in order for (10) to be true, it must be true that (i) we do not normally treat inscriptions as the ultimate perceptible objects of artistic appreciation; (ii) we normally use inscriptions

as instructions (or prescriptions) for generating such objects; and (iii) according to our ordinary reading practices, the ultimate perceptible objects of artistic appreciation that are generated by means of inscriptions are relevantly analogous to the ultimate perceptible objects of artistic appreciation that are generated by means of musical scores. Are (i), (ii), and (iii) true?

There is little doubt that (i), (ii), and (iii) are true in the case of reading aloud. For, in this case, we do not consider inscriptions to be final perceptible products of artistic appreciation. Rather, we treat them as instructions (or prescriptions) to generate such products—namely, readings. Furthermore, readings are relevantly analogous to the ultimate perceptible objects of artistic appreciation that are generated by means of musical scores—musical performances. First, like musical performances, readings are generated with the help of instructions (or prescriptions). Second, it seems plausible to hold that, similar to musical performances, readings are artistic performances.²⁴ Third, like musical performances, readings possess their own artistic properties. Fourth, like musical performances, readings provide experiential access to certain other artistic objects (namely, novels).²⁵

However, in the case of silent reading, the truth of (i), (ii), and (iii) is far from obvious. Do we actually generate any objects of artistic appreciation when we read

²⁴An argument in favor of the claim that readings are artistic performances is provided in Chapter 5.

²⁵The thesis that musical performances provide experiential access to other artistic objects (musical works) does not apply to all kinds of musical works. In those cases in which a musical performance is identical to a musical work (e.g., jazz), the performance does not provide experiential access to any artistic object other than itself.

silently? If we do, what are these objects? Furthermore, why think that they are relevantly analogous to musical performances? Unless a proponent of the Analogical Argument provides satisfactory answers to these questions, premise (10) and, hence, the Analogical Argument cannot be accepted.

3.2.1 Kivy's Response

The foregoing objection is based on the thesis that in the case of silent reading, (i), (ii), and (iii) are questionable. This thesis, however, is rejected by Kivy (2006). In his view, when we read novels silently, we do not treat the inscriptions as the ultimate perceptible objects of artistic appreciation; rather, we treat them as instructions for generating particular final perceptible objects of artistic appreciation—silent readings (that is, mental soundings, or “voicings”)—that are relevantly analogous to the ultimate perceptible objects of artistic appreciation that are generated by means of musical scores—musical performances. Is Kivy (2006)'s view acceptable? Before addressing this question, let us first clarify what he means by “a reading of a novel.”

In characterizing “a reading of a novel,” Kivy (2006) first notes that a reading of a novel is an event that, in most (though not all) cases, takes up “a certain non-continuous period of time” (Kivy, 2006, 5).²⁶ He then specifies what kind of event a reading of a novel is. In his view, such a reading:

²⁶In those cases in which a reading takes up such a period of time, this reading amounts to a complex event—“the sum total of a number of reading events, separated by various, sometimes protracted periods” (Kivy, 2006, 5).

is the kind of event we would describe as an act or an activity: it is an action performed by a reader. And the most important aspect of this act is that it... results in an “experience.” The point of an act of reading *Pride and Prejudice* is to have an experience of it for the usual reasons people have for experiencing works of art of that kind. Some people might say that such a reading act has as its purpose the experiencing of the work “aesthetically.” But I will not say that. I will say rather that its purpose (usually) is the experiencing of it qua art work of that kind: all the art-relevant ways of experiencing it, of which the aesthetic way is one. (Kivy, 2006, 5)

Thus, according to Kivy (2006), “a reading of a novel” is a particular event—namely, an act—that typically results in an art-relevant experience of the novel being read. What sort of act is this act? In Kivy (2006)’s view, it is a *sounding* of the text of a novel, where “a sounding” is understood as a sequence of sounds generated by a reader (such as a human being or a computer). Thus, given what has been said, “a reading of a novel,” according to Kivy (2006), can be defined as follows: For all x , x is a reading of a novel just in case x is a particular act—namely, a sounding of the text of this novel—that usually results in an artistic experience of this novel.

Before proceeding further, it is worth making a few additional remarks concerning Kivy (2006)’s account of “a reading of a novel.” First, on this account, a reading of a novel can be either *audible* or *silent*, where an audible reading is one that is generated as a result of reading the text aloud, and a silent reading is one

that is generated as a result of sounding out, or “voicing,” the text mentally (“in one’s head”).

Second, on Kivy (2006)’s account, a reading of a novel should be distinguished from an interpretation of this novel. Surely, when one reads a novel, one must interpret it at least to some extent. But this does not mean that this reading is identical to, or includes, or is contained in the interpretation of the novel. The two acts, although related to each other, are different. Reading a novel is an act of sounding out the text of this novel; interpreting a novel is an act of understanding what this text means.

Third, on Kivy (2006)’s account, a reading of a novel should be distinguished from an experience of this novel. Although a reading of a novel normally *results* in an experience, it is not an experience; rather, it is an act—namely, the act of sounding out the novel’s text.

Having clarified the sense in which Kivy (2006) uses the expression “a reading of a novel,” let us return to an examination of his view. Recall that according to Kivy (2006), in the case of silent reading, (i), (ii), and (iii) are true: When we read novels silently, we do not treat the inscriptions as the final perceptible objects of appreciation; rather, we use these inscriptions to generate such objects—silent readings—that are relevantly analogous to the final perceptible products generated by means of musical scores—musical performances. Is Kivy (2006) right about this?

Let us first consider the thesis that when we read novels silently, we generate silent readings—that is, silent soundings that usually result in some artistic experiences of what we read. Is this thesis true?

The answer to this question is “Yes” only if in the case of reading silently, we generate silent readings. Do we, in fact, generate such readings in this case? Although Kivy (2006) does not provide any reason in favor of answering this question in the affirmative,²⁷ such a reason can be provided. That we generate silent soundings when we read silently can be supported by the fact that there is psychological evidence suggesting that when we read silently, we mentally “voice” what we read and, hence, generate something that may well be called “a silent sounding.”²⁸ In

²⁷At the same time, Kivy (2006) provides a reason to think that when we read novels silently, we *can* generate silent soundings. The reason is as follows.

...It seems... consistent with what we know about thought and consciousness that reading a story might be experienced as “hearing” a story told in your head, in the way that reading a score, we know, is experienced, by those few who can, as “hearing” music played in your head. For at least there is nothing... implausible about the underlying premise, that one hears a voice in one’s head, with what we know about consciousness: first-person reports seem to confirm the experience. (Kivy, 2010, 111–112)

Thus, the idea that reading a story silently can be experienced as “hearing” this story in our mind reflects what we know about consciousness. So there is reason in favor of the thesis that reading a story silently can be experienced as “hearing” this story in our mind. Meanwhile, if this thesis is true, then, of course, we can mentally sound out this story.

It should be noted, however, that the foregoing consideration does not support the thesis that when we read silently, we *actually* generate silent soundings. For it is one thing to be capable of generating silent soundings, and it is another thing to actually generate such soundings when reading silently.

²⁸See, e.g., Petkov and Berlin (2013).

light of this, the question posed above can, I think, be answered in the affirmative.

So when we read silently, we generate silent soundings. Are these soundings *readings*? It is uncontroversial that silent soundings we generate in the process of reading silently usually result in artistic experiences of what we read. Meanwhile, to be a silent reading, it is sufficient to be a silent sounding that usually results in artistic experiences of what is being read. So the answer to the foregoing question is “Yes.”

Thus, the thesis that when reading novels silently, we generate silent readings can be considered true. Meanwhile, if this is so, then there is good reason to accept the first part of Kivy (2006)’s view—the part according to which when we read novels silently, we do not treat the inscriptions as the ultimate perceptible objects of artistic appreciation; rather, we treat them as instructions for generating particular final perceptible objects of artistic appreciation—silent readings. Now, what about the second part—that our silent readings of novels are relevantly analogous to the final perceptible products of playing music—musical performances. Is this part acceptable?

The answer to this question is “Yes” only if there are no serious disanalogies between our silent readings of novels and musical performances. Are there such disanalogies? Consider musical performances. They belong to the class of artistic performances. Given this, if our silent readings of novels do not belong to this class, then, *prima facie*, there is a serious disanalogy between them and musical performances. So do our silent readings of novels belong to this class?

According to Kivy (2006), they do. To support his view, he appeals to a

particular account of reading. In presenting this account, Kivy (2006) first invites us to consider the case of Ion—an Ancient Greek rhapsode famous for telling the stories of the *Iliad* and the *Odyssey*. According to Kivy (2006), when telling these stories, “Ion is ‘playing the role’ of Homer: he is impersonating the storyteller—not of course in the sense of an imposter, passing himself off as someone else for purposes of deception, but in the sense of an actor or actress impersonating a character in a play, playing a part” (Kivy, 2006, 44–45). In Kivy (2006)’s mind, when we read novels silently, we are like Ion: “*We* impersonate the storyteller silently, as Ion does out loud” (Kivy, 2006, 59). For example, when we read *Pamela* silently, we impersonate “the storyteller telling his story through the representation of letters” (Kivy, 2006, 60).

Regarding Kivy (2006)’s account of how we read novels, two remarks are worth making. First, the fact that x impersonates the storyteller does not entail that x make-believes that he/she is the storyteller. For suppose x impersonates the storyteller. Then, since x *impersonates* y just in case x *plays the part of* y ,^{29,30} x ’s impersonating the storyteller = x ’s playing the part of the storyteller. But x ’s playing the part of y does not entail that x make-believes himself/herself to be y . If someone plays the part of, say, Hamlet, then this does not necessarily mean that he make-believes that he is Hamlet.³¹ So x does not necessarily make-believes himself/herself to be the storyteller.

²⁹See Kivy (2006, 44–45).

³⁰Thus, when Ion impersonates Homer, Ion plays the part of Homer. Conversely, Ion’s playing the part of Homer entails Ion’s impersonating Homer.

³¹See Kivy (2006, 59).

Second, given the fact that when we read a story we impersonate a storyteller, it may seem natural to suppose that when we read characters' words, we impersonate the characters, and not the storyteller. However, according to Kivy (2006), this supposition is mistaken: In fact, when we read characters' words, we impersonate the storyteller impersonating the characters. Similarly, Kivy (2006) notes, "when [Ion] recites the characters' speeches, he is impersonating the storyteller impersonating the characters by reciting their speeches, much as in *Hamlet*, in the play within the play, the actor who plays 'the actor' impersonates an actor impersonating a king: an impersonator of an impersonator" (Kivy, 2006, 46).

Using the foregoing account of reading, Kivy (2006) explains why our silent readings of novels are, in fact, artistic performances. In his view, when we read a novel silently, we read *while impersonating the storyteller*. Meanwhile, whenever we read that way, we read *con espressione*—as a result of which the silent reading we generate acquires certain artistic properties, in particular the property of being "expressive." Of course, *how* "expressive" this reading is may vary. If it is generated by a professional reader, like Julie Harris, then, most likely, it is quite "expressive"; if, on the other hand, it is generated by a relatively unprepared, average reader, then its "expressiveness" is usually less noticeable. However, at least some "expressiveness" is always present, to a certain noticeable extent, in every proper silent reading. Meanwhile, if that is so, then, assuming that being "expressive" is sufficient for being artistic, a silent reading is an artistic entity. Moreover, there is good reason to treat a silent reading as a performance. *Prima facie*, to be a performance, it is sufficient to be an action that is intended to be presented to an audience. A

silent reading is doubtless an action—an activity that is generated by an agent (the reader) and directed at a goal (to generate a silent reading). Furthermore, a silent reading is intended to be presented to an audience (this audience consists of just one person—the reader herself). Thus, taking into account what has been said, our silent readings of novels are artistic performances.

The foregoing argument depends on Kivy (2006)’s account of how we read novels. Is this account satisfactory? Before addressing this question it is important to underline that Kivy (2006)’s account is not normative: It is not aimed at clarifying “how we *should* read [novels]” (Kivy, 2006, 2). Rather, this account is *descriptive*: Its goal is to clarify “how we, at least some of us, *do* read [novels]” (Kivy, 2006, 2). Thus, whether Kivy (2006)’s account is satisfactory depends on whether it reflects how we *actually* read novels, and not how we *should* read them.

So is Kivy (2006)’s account of how we read satisfactory? As D. Davies (2008) points out, according to this account, when reading a novel—say, *David Copperfield*—silently,

the reader impersonates Charles Dickens who is himself impersonating David Copperfield.... [And] in the case of direct quotation in such a fiction, the reader must presumably impersonate the author impersonating the narrator impersonating one of the characters. (Davies, 2008, 90)

However, that is not how we normally silently read *David Copperfield* or other novels. Perhaps, when reading a novel silently, we impersonate the storyteller. But we do not engage in a multi-level impersonation of the kind entailed by Kivy (2006)’s

account.

While D. Davies (2008)'s objection is successful, it fails to essentially undermine Kivy (2006)'s account. For Kivy (2006) can easily defuse this objection by adopting an account of impersonation that does not presuppose any multi-level impersonation—for example, an account according to which impersonation has only one level—say, the level of a storyteller.

A more pressing objection concerns the core of Kivy (2006)'s account—the thesis that when we read novels silently, we impersonate the storyteller. In a critical note on Kivy (2006)'s *The Performance of Reading*, Feagin (2008) writes:

I do not think that I do... what Kivy says he and other ordinary readers do when he writes that we “impersonate the storyteller silently, as Ion does out loud (p. 59).” I may, occasionally, silently impersonate a storyteller when I read, but I do so rarely.... (Feagin, 2008, 93)

Like Feagin (2008), I do not normally impersonate the storyteller when reading a novel silently. And most other readers, I am sure, do not do so either. Given this, the thesis that we impersonate the storyteller in our silent reading seems false. Meanwhile, if it is false, then Kivy (2006)'s account of reading cannot be accepted.

One might suggest that Kivy (2006)'s account can be modified by replacing the thesis that we impersonate the storyteller when reading silently, with the thesis that we impersonate the characters when reading silently. However, such modification will not rectify this account, since the latter thesis also seems false. Perhaps, when reading dialogs, we occasionally impersonate the characters participating in these

dialogs. However, most of the time, we do not impersonate any of the characters of a novel when we read it silently.³² (Here, an interesting question arises: Why don't we impersonate the characters in the process of silent reading? Probably, the main reason for this is that, as Feagin (2008, 93) notes, the act of impersonation takes attention away from the meaning of the text and thereby makes it harder for us to understand the story. Another reason is that we find impersonating characters useless, since, in our view, doing this does not help us appreciate the story or any other relevant aspect of the novel.)

Thus, Kivy (2006)'s account of how we read is unsatisfactory. But if this is so, then, since this account is used in his argument in favor of the thesis that silent readings we generate in the process of reading novels are artistic performances, this argument fails. Meanwhile, to my knowledge, there is no other argument that could support this thesis. So there seems to be no real reason to hold that our silent readings of novels are artistic performances.

At the same time, there is reason in favor of the opposite thesis—that our silent readings of novels are not artistic performances. Consider what Godlovitch (1998) says about a paradigmatic kind of artistic performance—musical performances:

[Musical] performances are deliberate, intentionally caused sound sequences. They are never involuntary like sneezes, nor accidental or inadvertent. A person, unaware of a certain piece, who plays something

³²Note that this does not imply that we *should* not perform this act. Perhaps our current practice of silent reading without impersonation is seriously flawed and, hence, should be abandoned in favor of a practice of silent reading that involves impersonation.

sounding just like it by casually running a bow across a cello could only generously be said to have performed that piece. Such a person could not claim the credit that is normally due to one who has given a performance. The intention to perform and beliefs about the immediate context are integral to performance. (Godlovitch, 1998, 16)

I agree. Whenever one creates a musical performance, one has a particular intention—the intention to perform. And the same, I think, can be said with regard to other paradigmatic kinds of artistic performance, such as performances of dance and theatrical performances. An artistic performance is not created without the intention to perform.³³ But when we read a novel silently, we do not intend to give a performance. Moreover, in this case, we normally do not have any specific intentions with regard to our reading at all.

Another consideration in favor of the thesis that our silent readings of novels are not artistic performances draws upon a particular view on arthood. According to this view, which is accepted by an overwhelming majority of theorists, when one creates art, one must, at least at some moment, have a particular intention with regard to this object—such as the intention to give the object “the capacity to satisfy the aesthetic interest” (Beardsley, 1983, 58), or the intention that the object be regarded “in some overall way that some earlier or preexisting artwork or artworks are or were correctly regarded or treated” (Levinson, 2006, 27–28), or the intention to endow the object with certain aesthetic properties in accordance

³³For an argument in favor of the claim that this intention is necessary for an artistic performance, see Davies (2011, 4–7).

with a particular creative insight (Zangwill, 1995, 307). However, as mentioned earlier, when we generate silent readings, we usually do not have any intentions with regard to them; in most cases, we generate them unintentionally. Consequently, if the foregoing view on arthood is correct, our silent readings of novels cannot be considered artistic entities and, hence, artistic performances.

In light of what has been said, we are entitled to conclude that our silent readings of novels are not artistic performances. Meanwhile, if this conclusion is correct, then there is a serious disanalogy between our silent readings of novels and musical performances: The latter are artistic performances, whereas the former are not.

Are there any other serious disanalogies between them? Consider the practice of generating (silent) readings by reading inscriptions silently and the practice of generating (silent) musical performances by reading musical scores silently. The former practice is doubtless very common. This is not to say, of course, that no one generates readings by reading inscriptions aloud. Parents often generate readings that way when reading to their kids. Also, readings are generated by reading inscriptions aloud when audio recordings (“audiobooks”) of the corresponding novels are produced.³⁴ So readings are sometimes generated by reading inscriptions aloud. But the practice of generating readings by reading inscriptions silently is considerably

³⁴It is also worth noting that in the past, the practice of generating readings by reading inscriptions aloud was not uncommon. Thus, as Cliff-Hodges (2015) notes, this practice “was for a long time common among the wealthy middle classes, from the novel’s inception onwards” (Cliff-Hodges, 2015, 96).

more widespread. However, the situation with the practice of generating musical performances by reading musical scores silently is different. In most cases, we do not generate musical performances by reading musical scores silently. Rather, we do it by actually performing the scores. Given what has been said, there is another striking disanalogy between our silent readings of novels and musical performances: In an overwhelming majority of cases, the former are generated as a result of silent reading, whereas the latter are not.

Can't the foregoing disanalogy be strengthened by saying that generating musical performances by reading musical scores silently is not just something unusual but is something that, in virtually all cases, simply cannot be done? According to Kivy (2006), the answer is "No." He admits that "reading scores... and realizing the sounds of musical works in one's mind is decidedly *not* the customary way of experiencing music nor is it anything but *very unusual*" (Kivy, 2006, 36). But, in his view, some highly gifted musicians are capable of generating musical performances by silent score reading:

The ability to read scores and thereby to successfully experience musical works is part of the aura that surrounds only the most gifted, the account of Beethoven "reading" the scores of Handel's and Schubert's works on his deathbed being a case in point. Of course I am not suggesting that one need be a musical genius to accomplish the feat. Nevertheless, you need to have a musical mind and musical training far beyond even that of most accomplished professional musicians. It is, in other words, a

feat of considerable mental power, considerably exceeding that of the average and even above-average musician, *a fortiori*, beyond that of the most avid and devoted music-lover. (Kivy, 2006, 36)

Is Kivy (2006) right about this?

As is clear from the foregoing quote, to support his view, Kivy (2006) gives the example of Beethoven, who could generate musical performances by silently reading Handel's and Schubert's scores. However, the available evidence confirming this example is anecdotal. At the same time, Kivy (2006) does not offer any other support for the view that highly gifted musicians are capable of generating musical performances by reading musical scores silently. Given this, it is reasonable to consider this view dubious. Perhaps one who has the appropriate skills could generate musical performances by silently reading very simple, "elementary" musical scores. However, that someone—even Beethoven himself—could silently read an average musical score (not to say a score as complex as the score of, say, Beethoven's Symphony No. 5) and generate a musical performance of this score in his mind is hard to believe. Taking this into account, there seems nothing that could stop us from adopting the strong version of the foregoing disanalogy—the version according to which virtually no musical performance can be generated by reading its score silently.

So there are at least two serious disanalogies between our silent readings of novels and musical performances. First, unlike musical performances, our silent readings of novels are not artistic performances. Second, our silent readings of

novels are generated as a result of silent reading, whereas musical performances are not. Thus, there are good grounds to hold that our silent readings of novels are not relevantly analogous to musical performances—and, hence, to reject the second part of Kivy (2006)’s view.

Given this, Kivy (2006)’s view cannot be accepted. As a result, the defense of (i), (ii), and (iii) based on this view, in the case of the silent reading of novels, cannot be accepted either.

3.2.2 Other Responses

Another potential way to defend (i), (ii), and (iii) in the case of silent reading draws upon Urmson (2004)’s account of our silent reading of novels. According to Urmson (2004), when we read novels silently, we do not treat inscriptions as the final perceptible objects of appreciation; rather, we treat them as instructions for generating such objects—*silent reading experiences of novels*—that are relevantly analogous to those final perceptible objects of appreciation that we generate in the case of reading musical scores—silent reading experiences of musical works. Is Urmson (2004) right about this? Before addressing this question, let us first clarify the expressions “a silent reading experience of a novel” and “a silent reading experience of a musical work.”

Consider first the expression “a silent reading experience of a novel.” According to Urmson (2004), it refers to an experience that (a) is acquired as a result of silently reading a particular “instruction”—namely, an inscription—and (b) en-

ables one to realize what one would hear if one heard a reading of this novel. Now, the expression “a silent reading experience of a musical work,” in Urmson (2004)’s view, has a similar meaning. It denotes an experience that (a) is acquired as a result of silently reading a particular “instruction”—namely, a musical score—and (b) enables one to realize what one would hear if one heard a performance of this work.

Given the foregoing account, it might seem natural to assume that both a silent reading experience of a novel and a silent reading experience of a musical work are artistic performances: A silent reading experience of a novel is a particular “voicing”-performance, whereas a silent reading experience of a musical work is a particular musical performance. However, Urmson (2004) rejects this assumption. In his view,

It would be implausible to say that musical score readers are giving a performance to themselves.... Apart from the fact that they need hear no sound (they may or may not hum to themselves), considered, absurdly, as performances, what the best score readers normally do would be intolerably bad. They habitually read through the slower bits far faster than they perfectly well know that the music should go, and, for many reasons, nobody can read a fast complex piece at a speed that he recognizes to be that of the music. Score reading is something quite distinct from... performance. Urmson (2004, 91)

Thus, according to Urmson (2004), what musical score readers generate—a silent

reading experience of a musical work—cannot be characterized as a performance. And the same, in his view, is true of what literary inscription readers generate—a silent reading experience of a novel. (Urmson (2004) does not explicitly provide any reasons why this latter experience is not a performance. However, it seems plausible to suppose that in this case, he could give reasons similar to the ones he gives in the quote above.)

Having clarified the expressions “a silent reading experience of a novel” and “a silent reading experience of a musical work,” let us now turn to an examination of Urmson (2004)’s view. Recall that on this view, when we read novels silently, we do not treat inscriptions as the final perceptible objects of appreciation; rather, we treat them as instructions for generating such objects—silent reading experiences of novels—that are relevantly analogous to those final perceptible objects of appreciation that we generate in the case of reading musical scores—silent reading experiences of musical works. Is Urmson (2004)’s view acceptable?

Let us first consider the thesis that when we read a novel silently, we generate a silent reading experience of this novel, or, in other words, an experience that (a) is acquired as a result of silently reading an inscription and (b) enables us to realize what we would hear if we heard an audible reading of this novel. Is this thesis true?

According to Urmson (2004), it is:

somewhat confirmed by some of the critical remarks we make about literary style. Even in the case of works which would not normally be read aloud it is a commonplace to speak of assonance, dissonance,

sonority, rhythm; we reject as unstylish conjunctions of consonants which would be awkward to say aloud, though we easily read them. We criticize the writing in terms of how it would sound, if it were spoken. Contrast the case of logical notation which is not literature and for which we have only a makeshift oral rendering: who would think of criticizing a piece of writing in formal logic as unstylish because our conventional reading of it was awkward in sound? (Urmson, 2004, 92)

Put otherwise, according to Urmson (2004), the thesis that when we read novels silently, we generate silent reading experiences of these novels is supported by the fact our critical practice emphasizes the sonic aspect of the novel. Is Urmson (2004) right about this? I doubt that. There is no real reason to hold that the foregoing thesis is actually supported by the fact that our critical literary practice emphasizes the sonic aspect of a novel. (It is worth noting, however, that this fact supports, to a certain extent, a *normative* version of Urmson (2004)'s thesis—the version according to which when we read silently, we *should* generate silent readings of novels. That our critical literary practice emphasizes the sonic aspect of a novel corroborates the thesis that the sonic aspect is relevant to the aesthetic appreciation of a novel. And this thesis, in its turn, provides support for the thesis that when reading novels silently, we (that is, ordinary readers) should pay attention to the sound of what we read—which may well require a generation of a silent reading experience of the kind suggested by Urmson (2004).)

Thus, Urmson (2004)'s substantiation of the thesis that when we read novels

silently, we generate silent reading experiences of these novels is unpersuasive. Despite this, however, this thesis might still be true. Is it true? When I read a novel silently, I usually do not have an experience that enables me to realize what I would hear if I heard an audible reading of the novel. And the same, I believe, can be said about most other readers. Given this, it is natural to hold that the process of reading a novel silently usually does not involve any silent reading experience of this novel. Meanwhile, if this process does not, in fact, involve any such experience, then the foregoing question must be answered in the negative.

So the part of Urmson (2004)'s view according to which our silent reading of a novel results in a silent reading experience of this novel (that is, an experience that enables us to realize what we would hear if we heard the novel's audible reading) is, at least, highly questionable. This provides us with a strong reason to reject this view. But it is not the only such reason. Urmson (2004)'s view states that our silent reading experiences of novels and our silent reading experiences of musical works are relevantly analogous. However, there is good reason to doubt that this analogy, in fact, holds. When typical literature readers silently read a novel, they are not interested in an experience that enables them to realize what they would hear if they heard an audible reading of the novel (what they are interested in is an experience of the story told by the novel). However, the situation is different in the case of musical score reading. When musical score readers read a musical score (let us assume, for the sake of argument, that doing this is possible), their primary aim is to get an experience that enables them to realize what they would hear if they heard an actual performance of this work. (Why would they read the score—in

Urmson (2004)'s sense of the word "read"—if they did not aim at that?) As a result, there is an important disanalogy between our silent reading experiences of novels and our silent reading experiences of musical works. Meanwhile, if this is so, then our silent reading experiences of novels and our silent reading experiences of musical works cannot be considered relevantly analogous.

Thus, there are at least two good reasons against Urmson (2004)'s view. Given that, this view cannot be accepted. Meanwhile, if this is so, then defending (i), (ii), and (iii), in the case of silent reading, with the help of this view cannot be accepted either.

Is there any other potentially acceptable way to defend (i), (ii), and (iii) in the case of silent reading? One could try to do this by saying that when we read novels silently, we do not treat inscriptions as the final perceptible objects of appreciation; rather, we treat them as instructions for generating such objects—namely, imaginary visual and sonic "actualizations" of the events portrayed in the novels—that are relevantly analogous to those final perceptible objects of appreciation that we generate in the case of reading musical scores—imaginary sonic "actualizations" (soundings) of musical works. However, this way of defending (i), (ii), and (iii) in the case of silent reading also fails. It entails that our silent reading of novels results in a generation of imaginary visual and sonic "actualizations" of the events portrayed in the novels. But this entailment is false. When we read novels silently, we rarely imaginatively "actualize" the events described in these novels. This is not to say, of course, that we never do this in the process of silent reading. When reading a description of some scene, we may sometimes imagine this scene; when reading

a dialog, we may sometimes mentally “voice” it. However, it would be mistaken to say that our silent reading of a novel is always accompanied by our imaginative “actualizing” of the events described in this novel.

3.2.3 Conclusion

So none of the foregoing ways to substantiate (i), (ii), (iii) in the case of silent reading is successful. Meanwhile, there seem to be no other potentially satisfactory ways to do that. Given this, there is good reason to hold that premise (10) (“Given our ordinary reading and performing practices, inscriptions and musical scores are relevantly analogous”) cannot be substantiated. Meanwhile, as mentioned in the introduction to this section, if this premise cannot be substantiated, then it cannot be accepted—and, hence, the Analogical Argument cannot be accepted either.

Despite the fact that the Analogical Argument seems unsatisfactory, there is a related argument against the Strong Orthodox View which, I believe, is satisfactory. The argument is as follows:

14. Given our best theory of musical and literary practice, inscriptions and musical scores are relevantly analogous.
15. Musical scores are not well-formed instances of musical works.
16. If (14) and (15) are true, then inscriptions are not well-formed instances of novels.
17. So inscriptions are not well-formed instances of novels.

The argument is valid. (15) is uncontroversial. A musical score is not an instance of a musical work, let alone a well-formed instance of such a work. The truth of (16) is beyond doubt. If inscriptions and musical scores are relevantly analogous and musical scores are not well-formed instances of novels, then inscriptions are not well-formed instances of novels. The crux of the argument is clearly premise (14). Is this premise true?

Unlike the truth value of (10), the truth value of (14) does not depend on our *ordinary* practices of reading novels and performing musical works (that is, practices that are common among “ordinary” readers and musicians). The truth value of (14) depends on our best theory of musical and literary practice—in particular, (14) is true if and only if this theory implies that inscriptions and musical works are relevantly analogous. So is (14) true? According to our best theory of musical practice, musical scores should not be treated as the final perceptible products of appreciation; rather, they should be treated as instructions for generating such products—musical performances. Given this, in order for (14) to be true, our best theory of literary practice must imply that (i*) inscriptions should not be treated as the final perceptible objects of appreciation; (ii*) inscriptions should be treated as instructions (or prescriptions) for generating such objects; and (iii*) these objects (whatever they are) should be treated as relevantly analogous to musical performances. Does our best theory of literary practice, in fact, imply that? As shown in Section 3.1, our best theory of literary practice implies that inscriptions cannot be used to fully appreciate novels—and, hence, that we should not treat inscriptions as the final perceptible objects of appreciation. Furthermore, as will become

clear from Chapters 4 and 5, our best theory of literary practice implies (a) that inscriptions should be treated as instructions (or prescriptions) for generating such objects and (b) that the latter should be treated as relevantly analogous to musical performances. Thus, the question about the truth of (14) should be answered in the affirmative.

Given the fact that the argument (14)–(17) is, at least *prima facie*, satisfactory, a natural question arises: Which of the arguments—this argument or the argument offered in Section 3.1—is preferable? The argument (14)–(17) involves the premise that according to our best theory of musical and literary practice, inscriptions and musical works are relevantly analogous. Since no analogy is perfect, it is impossible to *conclusively* show that this premise is true. As a result, the soundness of the argument (14)–(17) can always be questioned. The argument offered in Section 3.1, however, does not have this shortcoming. And, at the same time, this argument is at least as good as the argument (14)–(17) in all other respects. Thus, given what has been said, there is good reason to hold that the argument offered in Section 3.1 is preferable to the argument (14)–(17).

3.3 Against the Weak Orthodox View

As has been shown in Section 3.1, the Strong Orthodox View is false: Inscriptions cannot be *well-formed* instances of novels. But what about the Weak Orthodox View—the view that inscriptions are *non-well-formed* instances of novels? Is it also false?

As shown in the previous chapter, to be a non-well-formed instance of an artwork, it is necessary to manifest sufficiently many of the primary properties that must be experienced to fully appreciate this work. So an inscription cannot be a non-well-formed instance of a novel unless this inscription can manifest sufficiently many of the primary properties that must be experienced to fully appreciate this novel. Can an inscription manifest sufficiently many of these properties? Perhaps inscriptions can manifest *some* primary properties that must be experienced to fully appreciate novels. However, as has been shown in Section 3.1, inscriptions cannot manifest a significant portion of such properties—namely, the sonic properties of novels. Given this, it is reasonable to answer the foregoing question in the negative. Meanwhile, if this answer is correct, then, given what has been said, the Weak Orthodox View is false.

The foregoing argument is based on the thesis that the sonic properties of a novel are a significant portion of the primary properties that must be experienced to fully appreciate novels. What if this thesis is false? Should we then accept the Weak Orthodox View? No. For there is another way to show that this view is false. As demonstrated in the previous chapter, an entity cannot be a non-well-formed instance of an artwork unless this entity is, in principle, capable of manifesting all sensory kinds of the primary properties that must be experienced to fully appreciate this work. So if an inscription is, in principle, incapable of manifesting all sensory kinds of the primary properties that must be experienced to fully appreciate the corresponding novel, then inscriptions cannot be non-well-formed instances of novels. Can an inscription, in principle, manifest all such kinds? As has been established

in Section 3.1, no inscription can manifest certain primary properties that must be experienced to fully appreciate a novel—namely, the sonic properties of this novel. Meanwhile, these properties are a sensory kind. Thus, the answer to the foregoing question is “No.” As a result, taking into account what has been said, the Weak Orthodox View is false.

3.4 Against the Orthodox View

Thus, both the Strong and the Weak Orthodox Views fail: Inscriptions can be neither well-formed nor non-well-formed instances of novels. Meanwhile, the Orthodox View *simpliciter*—that is, the view that novels are instanced by inscriptions—is true *only if* inscriptions are either well-formed instances of novels, or non-well-formed instances of novels, or both well-formed and non-well-formed instances of novels. Consequently, this view is false.

3.5 Appendix

That the mentioned theorists endorse the idea that novels are instanced by inscriptions is evidenced by the following quotes:

- Noël Carroll: “When a play, like the dramatic text of *Strange Interlude*, is considered as a literary work, then my copy of *Strange Interlude* is a token [= an instance] of the art-type *Strange Interlude*...” (Carroll, 1998, 213).
- Arthur Danto: “I can... burn up a copy of the book in which a poem is printed, but it is far from clear that in so doing I have burned up the *poem*,

since it seems plain that though the page was destroyed, the poem was not; and though it exists elsewhere, say in another copy, the poem cannot merely be identical with that copy. For the same reason, it cannot be identified with the pages just burned. This immediately suggests that the poem stands to the class of its copies in something like the relationship in which a Platonic form stands to its instances, where it would have been acknowledged by Plato that the destruction of the instances leaves the form unaffected...; and by parallel reasoning the Poem Itself appears to be logically incombustible” (Danto, 1981, 33).

- David Davies: “Generally, we accord the status of an e-instance only to those events or objects for which we can obtain some guarantee that they have that status, as is the case with copies of novels...” (Davies, 2010, 415).
- Stephen Davies: “If I copy out Austen’s *Persuasion*, I produce another instance of her novel...” (Davies, 2003a, 158).
- John Dilworth: “Type theories are one popular way in which to explain how a particular novel, musical composition, etc. could have multiple copies or performances, yet still be such that all of its instances or tokens are purely concrete items” (Dilworth, 2008, 346).
- Nelson Goodman: “All and only inscriptions and utterances of the text are instances of the [literary] work...” (Goodman, 1968, 209).
- Peter Lamarque: “The view that copies of literary works are tokens [= in-

stances] and the works themselves types is widely endorsed, although the question of what these types are is more controversial” (Lamarque, 2009, 72).

- Jerrold Levinson: “My handwritten copy, I would maintain, is an instance of Black’s poem...” (Levinson, 2011a, 99).
- Christy Mag Uidhir: “According to this standard account, our interaction with such art-abstracta (e.g., *Moby-Dick*, *The Eroica Symphony*) must be mediated by their associated concrete instances or specifications (e.g., copies of *Moby-Dick* on library shelves, performances of *The Eroica* in symphony halls)” (Mag Uidhir, 2013, 167).
- Aaron Meskin: “Multiple instance works (for example, the novel *The Name of the Rose*) are types. Particular instances of such works (e.g., the copy of *The Name of the Rose* in the British Library) are tokens, specific spatiotemporally located instance of the novel” (Meskin and Robson, 2011, 554).
- Richard Wollheim: “*Ulysses* and *Der Rosenkavalier* are types, my copy of *Ulysses* and tonight’s performance of *Rosenkavalier* are tokens [= instances] of those types” (Wollheim, 1971, 50).
- Lee Walters: “It is extremely plausible to think of repeatable works of art as types, of which their particular copies and performances are tokens [= instances of the corresponding types]” (Walters, 2013, 461–462).

The foregoing quotes do not explicitly support the idea that the abovementioned philosophers share the conviction that inscriptions are *well-formed* instances

of novels. In my view, however, this idea is uncontroversial and can be assumed. If it were false, at least some of the proponents of treating inscriptions as instances of novels would point out that inscriptions cannot be *well-formed* instances of novels. However, to my knowledge, no one points that out.

Chapter 4: An Ontology of Instances of Novels

Introduction

As shown in the previous chapter, the most widely held view on the nature of instances of novels—the Orthodox View—fails. But what is then the correct view? What are the entities that serve as instances of novels?

My goal in this chapter is to answer this question. I begin by addressing the question of what entities serve as *well-formed* instances of novels, arguing that these entities are readings (that is, particular sequences of sounds generated as a result of reading aloud) and mereological sums of readings and graphic elements (Sections 4.1 and 4.2). Then I turn to the question of what entities serve as *non-well-formed* instances of novels. I argue that these entities are the same as the entities that serve as well-formed instances of novels—namely, readings and mereological sums of readings and graphic elements (Section 4.3). Next, I respond to possible objections to the view that novels are instanced by readings and mereological sums of readings and graphic elements (Section 4.4). Finally, I examine potential consequences of this view and make some remarks concerning its possible extension as well as the question of what makes a reading a reading of a particular novel (Sections 4.5 and 4.6).

4.1 Primary Appreciation-Relevant Experienceable Properties of Novels

What entities serve as well-formed instances of novels? Before addressing this question, let us first clarify what kinds of properties are contained in the set of the primary properties that must be experienced to fully appreciate a novel.

As shown in the previous chapter,¹ one kind of properties contained in this set is *the sonic properties of a novel*. Are any other kinds of properties contained in it? Consider experienceable properties that provide experiential access to the semantic content of a novel (or, in other words, to what the novel tells the reader). Clearly, the semantic content of a novel is essential to the aesthetic value of this novel. And no novel can be fully appreciated without appreciating its aesthetic value. So a full appreciation of a novel is impossible without appreciating the semantic content of this novel. Meanwhile, to appreciate the semantic content of a novel, it is necessary to experience certain properties that provide experiential access to this content. Thus, the experienceable properties that provide experiential access to the semantic content of a novel must be experienced to fully appreciate this novel. Furthermore, these properties are primary. They provide experiential access to the semantic content of a novel. And this content is clearly part of the novel's primary artistic content—the set of the contentful properties of the novel that are not grounded in any other contentful properties. So the experienceable properties that provide experiential access to the semantic content of a novel provide experiential

¹See Chapter 3, Section 3.1.

access to the primary artistic content of this novel. But then they are properties through which this content is articulated. Meanwhile, any property through which the primary artistic content of an artwork is articulated is primary.²

So another kind of properties contained in the set of the primary properties that must be experienced to fully appreciate a novel is *the experienceable properties that provide experiential access to the semantic content*. Now, what about the visual properties of a novel? Are they contained in this set? Clearly, sonic properties are non-visual. So the fact that such properties are contained in the set of the primary properties that must be experienced to fully appreciate a novel does not entail that any visual properties are contained in this set. What about the experienceable properties that provide experiential access to the semantic content? Does the fact that these properties are contained in the set of the primary properties that must be experienced to fully appreciate a novel entail that any visual properties are contained in this set? No. Of course, the experienceable properties that provide experiential access to the semantic content *can* be visual, since in some cases, they can be apprehended by merely looking at some entity—say, an inscription. But they do not have to be visual. Thus, consider the properties of a (correct) recitation of a novel. These properties provide experiential access to the semantic content of this novel. But they are not visual (they are sonic).

So the fact that the sonic properties and certain experienceable properties that provide experiential access to the semantic content are contained in the set of the primary properties that must be experienced to fully appreciate a novel does not

²See Chapter 2, Section 2.1.

entail that any visual properties are contained in this set. But, of course, this fact does not entail the opposite—that no visual properties are contained in this set. So the question remains: Are any visual properties contained in the set of the primary properties that must be experienced to fully appreciate a novel?

An answer to this question depends on whether novels under consideration have aesthetically relevant graphic elements—such as drawings, diagrams, maps, photographs, etc. Suppose these novels do have such elements.³ Then the answer is “Yes”: Certain visual properties are contained in the set of the primary properties that must be experienced to fully appreciate a novel. For consider novels with aesthetically relevant graphic elements. Must any visual properties be experienced to fully appreciate such novels? As mentioned earlier, no novel can be fully appreciated without appreciating what is essential to its aesthetic value. Clearly, in the case of novels that have aesthetically relevant graphic elements, these elements are essential to the aesthetic value of these novels. So none of these novels can be properly appreciated without appreciating its graphic elements. Meanwhile, a correct appreciation of a graphic element requires an experiential engagement with

³Examples of novels that have aesthetically relevant graphic elements are Dan Brown’s *The Da Vinci Code*, which involves pictures of certain unreadable symbols; Mark Z. Danielewski’s *House of Leaves*, which includes words printed in blue and purple; Kurt Vonnegut’s *Breakfast of Champions*, which contains numerous drawings of things like the yin-yang symbol, female underpants, and Eliot Rosewater’s handwriting; R. L. Stevenson’s *Treasure Island*, which includes a map of Treasure Island drawn by the author; and W. G. Sebald’s *Schwindel. Gefühl*, which incorporates photographs, maps, and drawings.

certain visual properties.⁴ Thus, to fully appreciate novels that involve aesthetically relevant graphic elements, certain visual properties must be experienced. Furthermore, these properties are primary. They provide experiential access to a part of the primary artistic content of a novel—namely, certain graphic elements—and, hence, are among the properties through which this content is articulated. Meanwhile, as already mentioned, any property through which the primary artistic content is articulated is primary.

Thus, in the case of novels that involve aesthetically relevant graphic elements, *certain visual properties* are contained in the set of the primary properties that must be experienced to fully appreciate these novels. At the same time, that is not how things are in the case of novels that do not involve aesthetically relevant graphic elements.⁵ In this case, it is not necessary to experience any visual properties to fully appreciate a novel—and so such properties are not contained in the set of the primary properties that must be experienced to fully appreciate these novels.

Given what has been said, the question of what kinds of properties are contained in the set of the primary properties that must be experienced to fully appreciate a novel can be answered as follows: The kinds of properties contained in the set of the primary properties that must be experienced to fully appreciate *any* novel are certain sonic properties and experienceable properties that provide expe-

⁴This is not to say, of course, that we cannot, in principle, correctly appreciate a graphic element without *actually* seeing it. For, in at least some cases, we can correctly appreciate such a property by “seeing” it with the help our imagination.

⁵It seems safe to maintain that such novels constitute the majority of all novels.

riential access to the semantic content; the kinds of properties contained in the set of the primary properties that must be experienced to fully appreciate a novel that involves aesthetically relevant graphic elements are certain sonic properties, experienceable properties that provide experiential access to the semantic content, and certain visual properties.

4.2 The Ontological Status of Well-Formed Instances of Novels

Having clarified what kinds of properties are contained in the set of the primary properties that must be experienced to fully appreciate a novel, we can now turn to the question of what entities serve as well-formed instances of novels.

In light of what has been said in the previous section, all novels can be divided into two groups:

- **Non-visual novels**, or novels whose primary properties that must be experienced to fully appreciate these novels are certain (a) sonic properties and (b) experienceable properties that provide experiential access to the semantic content; and
- **Visual novels**, or novels whose primary properties that must be experienced to fully appreciate these novels are certain (a) sonic properties, (b) experienceable properties that provide experiential access to the semantic content, and (c) visual properties.

Given this distinction, it is clear that to answer the question “What entities serve as well-formed instances of novels?” it is sufficient to answer two questions: “What

entities serve as well-formed instance of visual novels?” and “What entities serve as well-formed instances of non-visual novels?” Let us answer these questions.

4.2.1 Readings as Well-formed Instances of Non-Visual Novels

In my view, the best candidates for the role of the entities that serve as well-formed instances of non-visual novels are *readings*. Before explaining why this is so, let us first clarify what exactly is meant here by “a reading” and what sort of thing a reading is, ontologically speaking.

Here is a possible definition of “a reading”:

Reading: For all x , x is a reading just in case x is a sequence of particular external (spatiotemporal) sounds generated as a result of reading aloud.

Note that a reading, thus understood, does not have to be directly produced by a human being. Thus, both a sequence of sounds generated as a result of sounding out a text by a computer and a sequence of sounds generated as a result of playing an audiobook can be readings.

Note also that the sense in which the expression “a reading” is used here is different from two common senses of this expression: the sense according to which a reading is an act of extracting meanings from syntactic sequences and the sense according to which “a reading” is synonymous with “an interpretation.” (An example of the use of “a reading” in the former sense is the use of “a reading” in the sentence “Your reading of the instructions is too slow.” In the latter sense, “a reading” is used, for example, in the sentence “Two literary critics have different

readings of James Joyce's *Ulysses*.”)

Having clarified the meaning of “a reading,” let us now clarify what sort of thing a reading is, ontologically speaking. According to the foregoing definition, readings are *sequences of sounds generated as a result of reading aloud* (hereafter: “sequences of sounds”). What is the ontological status of these sequences? To answer this question, it is necessary to answer the following questions:

- (1) Are sequences of sounds concrete or abstract?;
- (2) If sequences of sounds are concrete, then are they objects, events, or something else?;
- (3) What sort of things are the constituents of these sequences—sounds—ontologically speaking?

Let us first answer question (1). Sounds are doubtless concrete. So sequences of sounds are sequences of concreta. Meanwhile, a sequence—at least, as it is understood here—is a mereological sum. So sequences of sounds are mereological sums of concreta. But a mereological sum of concreta is a concretum. So sequences of sounds are concreta.⁶

Let us now turn to question (2) (“If sequences of sounds are concrete, then are they objects, events, or something else?”). To answer it, we first need to highlight some major standard (though not universally accepted) differences between objects

⁶Note that here, by “a sequence” is meant a mereological sum. If “a sequence” were used in a different sense—say, in the sense of a particular set—then, of course, the conclusion of the given argument—“Sequences of sounds are concreta”—would not follow.

and events. As Casati and Dokic (2014) point out, there are five such differences.

First, there is a difference in mode of being: [1] material objects such as stones and chairs are said to *exist*; events are said to *occur* or *happen* or *take place*.... Second, there are differences in the way objects and events relate to space and time. [2] Ordinary objects are supposed to have relatively crisp spatial boundaries and vague temporal boundaries; events, by contrast, would have relatively vague spatial boundaries and crisp temporal boundaries. [3] Objects are said to be invidiously located in space—they *occupy* their spatial location; events tolerate co-location much more easily.... [4] Objects can move; events cannot.... Finally, [5] objects are standardly construed as continuants—they *are in* time and persist through time by being wholly present at every time at which they exist; events are occurrents—they *take up* time and persist by having different parts (or “stages”) at different times.... (2014)

Taking this into account, let us now answer question (2). While sequences of sounds can be said to *exist*, it seems more natural to say that they *occur* or *take place*. Furthermore, since it is unclear where exactly sequences of sounds are spatially located but, at the same time, relatively clear where they are temporally located, sequences of sounds have vague spatial boundaries and crisp temporal boundaries. Next, sequences of sounds appear to tolerate co-location: It seems plausible to say that two different sequences of sounds can be located at the same spatial region at the same time. Also, sequences of sounds do not seem to be capable of moving

from one place to another. Finally, sequences of sounds take up time and persist by having different parts (sounds) at different times. Thus, sequences of sounds seem to possess the characteristics that, on Casati and Dokic (2014)'s account, are peculiar to events. Given this, such sequences should be treated as events, not objects.

Finally, let us address question (3) (“What sort of things are sounds, ontologically speaking?”). There are a number of possible answers to this question. According to the Property View, a sound is a particular property. Depending on the kind of this property, there can be different versions of this view. According to one version, a sound is a dispositional property, or “power,” to produce auditory experiences in a perceiver. According to another version, a sound is a physical property that is capable of producing auditory experiences in a perceiver. According to yet another version, a sound is a mental property of a perceiver’s auditory experiences.⁷

Can the Property View, in any of its versions, be accepted? Here is one reason to say “No.” If a sound is a property, then it is a universal, and not a particular. But is a sound, in fact, a universal, and not a particular? As is generally agreed, a particular has two features that distinguish it from a universal.⁸ First, it is not instantiated by other entities; rather, it itself instantiates other entities—for example, various properties. Second, it can survive change to its intrinsic properties. There is no doubt that a sound has the first of these features. A sound is not instantiated by other entities; rather, it itself instantiates other entities—for example, proper-

⁷A notable proponent of the Property View is Locke (1975). Also, this view has recently been advocated by Pasnau (1999).

⁸It is assumed here that an entity is either a particular or a universal.

ties such as “being loud,” “being high-pitched,” “being mellifluous,” etc. Does a sound have the second feature? Suppose that at a particular moment t_1 , we hear some sound—say, the sound of a siren—and this sound is perceived by us as quiet and low-pitched. Suppose next that what we hear gradually becomes louder and louder—so that at t_2 , it becomes very loud and high-pitched. It would be odd to say that the t_2 sound and the t_1 sound are completely different. The natural view is that the t_2 sound is the t_1 sound that has undergone a particular change—namely, has changed from being quiet and low-pitched to being loud and high-pitched. But if this is so, then, given the fact that being quiet and low-pitched and being loud and high-pitched are intrinsic properties, there is good reason to hold that a sound is capable of surviving change in its intrinsic properties, which means that a sound has the second characteristic feature of a paradigmatic particular.

Thus, a sound possesses both characteristic features of a particular. Meanwhile, if this is so, then a sound should be regarded as a particular, not a universal.

Another consideration against the Property View concerns the possibility of attributing certain properties to sounds. Consider the property of being red. This property is not itself red; what is red is an object that manifests this property. Analogously, the property of being loud is not loud; what is loud is an object that manifests this property. But if this is so, then if sounds are properties, then sounds cannot be literally characterized as “loud.” This consequence, however, is unacceptable. Surely, any potentially plausible account of sounds must be compatible with the idea that sounds can literally be characterized as “loud.” (Obviously, the scope of this objection is not restricted to the case of attributing to sounds the property

of being loud; the objection can be extended to other similar cases.)

In light of the above considerations against the Property View, it is reasonable to search for a less controversial account of the nature of sounds. A potential candidate for the role of such an account is the Sensation View—the view that a sound is a non-physical, mental entity (say, a sensation produced in the listener when the sound waves reach her ear).⁹ To support this view, the following argument can be offered. There is a good reason to hold that a sound can be generated even in the absence of any external (physical) cause. The reason is that there are numerous reports of people literally hearing sounds—such as ringing, buzzing, and even voices—*when there are no external (physical) causes for these sounds*. Meanwhile, a physical entity x can be generated only if there is another physical entity y that generates x —or, in other words, only if there is a physical cause of x . Thus, a sound is not essentially physical. Meanwhile, if an entity is not essentially physical, then it is natural to assume that this entity is not physical at all. As a result, there is good reason to regard a sound as a non-physical entity. But if a sound is such an entity, then it must be either mental or abstract. Surely, it is not abstract, for it can be perceived and, hence, is (directly) perceptually accessible, whereas no abstractum is (directly) perceptually accessible. So a sound is a mental entity.

The foregoing argument in favor of the Sensation View implies that the fact that someone hears a sound when there is no external (physical) cause of this sound is sufficient to demonstrate that this sound can be generated without such a cause.

⁹A notable proponent of the Sensation View is Maclachlan (1989).

Is this implication true? Suppose John claims that he is hearing someone's voice. Suppose also that in fact, there is no external (physical) source that could produce this voice. What would our reaction to John's claim be? Most likely, we would say: "In fact, no one is talking here. You are having auditory hallucinations." If our reaction to John's claim is correct—and there seems to be no reason to think otherwise—then the fact that someone hears a sound when there is no external (physical) cause of this sound is insufficient to demonstrate that this sound actually exists (similarly, the fact that someone sees a pink elephant dancing the Argentine tango does not mean that this elephant actually exists). Meanwhile, if this fact is insufficient to demonstrate that, then it is also insufficient to demonstrate that the foregoing sound can, in fact, be generated without an external (physical) cause.

Thus, the above argument in favor of the Sensation View cannot be accepted. At the same time, there are at least two reasons against this view. One of these reasons is as follows. If the Sensation View is true, then sounds, being mental entities, are not publicly available—that is, they cannot be experienced by anyone other than the person who has them in his/her mind. But this consequence is problematic. Suppose sounds are not publicly available. Then when listening to a musical performance (understood as a particular sequence of sounds) at a concert, each of us is listening to his/her private performance, not one and the same publicly available performance. However, this result seems false. *Prima facie*, when we listen to a musical performance at a concert, each of us listens to one and the same publicly available performance, not his/her own private performance.

Similarly, if sounds are not publicly available, then when listening to a (public)

lecture (understood as a particular sequence of sounds), each of us is listening to his/her private lecture, not one and the same publicly available lecture. But this result does not appear correct. It seems odd to claim that when we listen to a lecture, each of us listens to his/her private lecture, rather than to one and the same publicly available lecture.

Here is another reason against the Sensation View. If this view is true, then, since mental entities have no spatial location, sounds have no such location. However, there is a good reason to hold that sounds do have spatial location. The reason is that according to our auditory experience, many (though, of course, not all) sounds have a nonmental, spatial origin. Thus, when we perceive that somebody is talking to us, we do not think that the sound of her voice originates in our mind; rather, we think that this sound originates outside our mind—somewhere in space (where exactly it originates is another question). Likewise, when we perceive that someone is ringing a bell, we think that the sound generated as a result of this ringing originates in space, not inside our mind.

One could respond that if the mental is identical or can be completely reduced to physical brain states, then, contrary to what the foregoing argument states, sounds have spatial location—they are located where the relevant parts of the brain are located. This response, however, is unpersuasive. First of all, it assumes that the mental is reducible to the physical, but this assumption is controversial. More importantly, the key idea of the response being discussed—that sounds are located where the relevant parts of the brain are located—contradicts our auditory experience. Consider again the sounds of talking and ringing a bell. We do not perceive

these sounds as originating where the relevant parts of one's brain are located. According to our auditory experience, the sound of ringing a bell originates where the bell is, while the sound of talking originates where the one who is talking is.

Thus, like the Property View, the Sensation View is problematic. A more promising account of the nature of a sound is provided by the Wave View. On this view, a sound is a wave “that is generated by a disturbance and that moves through a surrounding medium [such as water or air]” (O’Callaghan, 2009, 28). The Wave View has been rather popular and may well be characterized as the standard account of the nature of a sound. Advocates of the Wave View include most contemporary physicists and a considerable number of contemporary philosophers. Furthermore, this view (or, at least, something very similar to it) was endorsed by Aristotle, Galileo, Descartes, and some other notable thinkers of the past.

The popularity of the Wave View can be explained by the fact that “many perceptual properties of sounds are neatly explained by the presence of strong correlations with properties of waves, in particular pitch and intensity (i.e. volume)” (Casati and Dokic, 2014). Thus:

The felt quality of high pitch is correlated with high frequencies; low pitch is correlated with low frequencies; high volumes are correlated with high, low volumes with low, amplitudes. The directionality of sounds (the fact that they appear to be ‘in a direction’) is related to the fact that the hearer is located on a propagation line from the source. Even more accomplished is the explanation of particular auditory effects, such

as the Doppler effect, whereby the speed of a sounding object in motion contributes to a change in the sound's heard pitch (so that the whistle of an engine passing by is heard to drop in pitch as it travels past us). (Casati and Dokic, 2014)

However, as O'Callaghan (2009) points out, the Wave View is not unproblematic. Suppose sounds are waves. Then, since waves travel through air (or some other medium), sounds must travel through air (or some other medium). However:

...This is not how things seem.... Sounds are not perceived... to travel through the air [or some other medium] as waves do. They are heard to be roughly where the events that cause them take place.... When we hear a clock ticking, the sound seems to be "over there" by the clock; voices are heard to be in the neighborhood of speakers' heads and torsos; when a door slams in another part of the house, we know at least roughly where the accompanying racket takes place. (O'Callaghan, 2009, 28–29)

So, assuming that our auditory experience is not systematically illusory—and there seem to be no persuasive considerations to think otherwise—there is a powerful reason against treating sounds as waves.

Given the foregoing problem, even the Wave View cannot be considered completely satisfactory (although it seems right to characterize this view as less problematic than the Property and the Sensation Views). Is there a satisfactory view on the nature of a sound?

A highly plausible candidate for the role of such a view is the Event View.¹⁰ On this view, a sound is a particular *disturbance event*—an event “in which a medium is disturbed or set into wave-like motion by the movement of a body or interacting bodies” (O’Callaghan, 2009, 36). Thus, the sound that comes into existence as a result of the striking of a tuning fork in air is the event of disturbing the air by the oscillation of the fork (where the oscillation is generated as a result of striking the fork). And the sound of reading is the event of disturbing the air as a result of the reader’s moving her lips, tongue, and certain other parts of her body.

Why think that the Event View is a highly plausible candidate for the role of a satisfactory view on the nature of a sound? One reason to think so is that this view does not have the shortcomings of the views examined above. A disturbance event can take place where we perceive the corresponding sound to be. For example, in the case of the tuning fork, the disturbance event takes place where the oscillation of the fork disturbs the air—and this accords with our perceptual experience of the location of the sound that comes into existence as a result of striking this fork. So if a sound is a disturbance event, then a sound can take place where we perceive it to be. Consequently, the Event View avoids the problem of the Wave View, as well as one of the problems of the Sensation View.

Next, a disturbance event, being a physical entity, is publicly available. Thus, if a sound is a disturbance event, then it is publicly available. Therefore, the Event View avoids the second of the problems of the Sensation View.

Finally, a disturbance event is a particular and can be literally characterized as

¹⁰A notable proponent of the Event View is O’Callaghan (2009, 2010).

having the property of being loud as well as other properties we normally attribute to sounds. So if a sound is a disturbance event, then it is a particular and can be literally characterized as having the mentioned properties. As a result, the Event View avoids the problems of the Property View.

Thus, the Event View does not face the problems of the Property, Sensation, and Wave Views. At the same time, as far as I am aware, this view does not face any other problems.¹¹ In light of this, the Event View can, I think, be regarded as a highly plausible candidate for the role of a satisfactory account of the nature of a sound. Given this, there is, I think, good reason to adopt this view—at least, as a *working* account of the nature of a sound. (It is worth noting, however, that neither the arguments of this chapter nor the arguments presented elsewhere depend on the truth of the Event View.)

Thus, given what has been said, a reading can be characterized, from an ontological viewpoint, as:

- a concretum;
- an event;
- a sequence (a mereological sum) of certain sounds, where the latter are particular disturbance events, namely events in which the air is disturbed by certain movements performed by a reader.¹²

¹¹For a powerful defense of the Event View against potential objections, see O’Callaghan (2009, 2010).

¹²Here, “a reader” can refer not only to humans but also to things like computers.

Having clarified what is meant by “a reading” and what sort of thing a reading is, ontologically speaking, let us return to our main thesis—that the entities that serve as well-formed instances of non-visual novels are readings. Why think that this thesis is true?

Surely, readings, being essentially sonic entities, can manifest the relevant sonic properties. What about the experienceable properties that provide experiential access to the semantic content? Can they be manifested by readings? Clearly, we can grasp the semantic content of a novel solely by listening to a reading of this novel. But if this is so, then the experienceable properties that provide experiential access to this content can be apprehended by means of directly experiencing this reading. Meanwhile, if a property can be apprehended as a result of directly experiencing some entity, then this entity manifests this property. So the foregoing question should be answered in the affirmative.

Thus, readings can manifest both the relevant sonic properties and experienceable properties that provide experiential access to the semantic content. Meanwhile, as already mentioned, the primary properties that must be experienced to fully appreciate a non-visual novel include only the relevant sonic properties and experienceable properties that provide experiential access to the semantic content. As a result, readings are capable of manifesting all the primary properties that must be experienced to fully appreciate non-visual novels. But if this is so, then, since manifesting all the primary properties that must be experienced to fully appreciate an artwork is sufficient for being a well-formed instance_e of this work, readings are capable of being well-formed instances_e of non-visual novels.

Now, what about well-formed instances_p of non-visual novels? Can readings serve as such instances? To be a well-formed instance_p of an artwork, it is sufficient to (a) manifest all the primary properties that must be experienced to fully appreciate this work and (b) stand in an appropriate historical-intentional relation to it. As has been shown above, readings can manifest all the primary properties that must be experienced to fully appreciate non-visual novels. Furthermore, there is nothing to suggest that readings of such novels cannot stand in appropriate historical-intentional relations to these novels. Given this, readings can be well-formed instances_p of non-visual novels.

One might suggest that, besides readings, there is another class of entities that can serve as well-formed instances of non-visual novels—namely, *mental “voicings,”* or, in other words, sequences of particular sounds produced in one’s mind with the help of imagination. Is this suggestion acceptable? To answer this question we first need to answer the following question: Are mental “voicings” capable of manifesting the sonic properties that must be experienced to fully appreciate non-visual novels?

The first thing to note is that there is little doubt that mental “voicings” can manifest *some* sonic properties. Take any mental “voicing.” At least some of its sonic properties can be apprehended by directly perceiving it. Meanwhile, if a property is apprehensible by directly perceiving some entity, then this entity manifests this property.

But can mental “voicings” manifest *the sonic properties that must be experienced to fully appreciate the corresponding novels?* An essential characteristic of at least some of these properties is the possession of a particular degree of sonic

vididness—the degree similar to the one possessed by sonic properties that are manifested by means of real sound. However, properties manifested through mental “voicings” lack this characteristic. The degree of sonic vividness of these properties is much lower than the degree of such vividness of sonic properties manifested through reading aloud.

Thus, there is good reason to regard mental “voicings” as incapable of manifesting all the sonic properties that must be experienced to fully appreciate novels. Meanwhile, as shown in the previous chapter, to be a well-formed instance of a novel, it is necessary to manifest these properties. So mental “voicings” cannot serve as well-formed instances of novels and *a fortiori* non-visual novels.

One might object by saying that the foregoing argument fails because it falsely assumes that sonic properties manifested by mental “voicings” are not as sonically vivid as sonic properties manifested through reading aloud. Is this objection successful? It implies that sonic properties manifested by mental “voicings” are at least as sonically vivid as sonic properties manifested through reading aloud. But this implication is false. Reading a text aloud generates sounds that are considerably more sonically vivid than the “sounds” that result from mentally “voicing” this text. Put otherwise, the sonic vividness of a reading of a text is considerably greater than the sonic vividness of a mental “voicing” of this text. Similarly, the sonic vividness of actual music (such as an actual musical performance or a playing of a recording of such a performance) considerably surpasses the sonic vividness of any music “performed,” with the help of imagination, in one’s mind.¹³ Meanwhile,

¹³This is not to say, of course, that a mental “voicing” of a text/music “performed” in one’s

the sonic vividness of sonic properties is directly proportional to the sonic vividness of the sounds on which these properties supervene. As a result, the sonic vividness of the sonic properties manifested through reading aloud is greater than the sonic vividness of the sonic properties manifested by mental “voicings.”

There seems no other potentially successful objection to the foregoing argument against treating mental “voicings” as well-formed instances of non-visual novels. Given this, this argument can, I think, be accepted. But what if it is, in fact, unacceptable? In this case, there is good reason to hold that a mental “voicing” can be a well-formed instance of a non-visual novel. Note, however, that the fact that a mental “voicing” can be a well-formed instance of a non-visual novel does not pose any real threat to the essence of the view defended in this chapter. If non-visual novels can be properly instanced by mental “voicings,” this view can easily be rectified—by (a) rejecting the foregoing argument against treating mental “voicings” as well-formed instances of non-visual novels and (b) replacing the current account of “a reading” with an account according to which readings are not only sequences of particular actual sounds generated as a result of reading aloud but also mental “voicings,” or sequences of particular mental “sounds” (or “sounds” produced in one’s mind with the help of imagination) generated as a result of reading silently.

Before proceeding further, it is worth pointing out that treating readings as well-formed instances of non-visual novels does not give us reason to think that “the lookings at a picture and the listenings to a performance” (Goodman, 1968, mind cannot surpass a reading of a text/actual music in some non-sonic (say, emotional) vividness.

114) are well-formed instances of paintings and musical works, respectively.¹⁴ Before explaining why this is so, it is necessary to clarify what is meant here by “a looking” and “a listening.” Each of these expressions can be used in either of two senses. In one sense, “a looking” and “a listening” refer to particular acts: the act of looking and the act of listening, respectively. In the second sense, they refer to particular kinds of experience: the experience of looking at something and the experience of listening to something, respectively.

So why think that treating readings as well-formed instances of non-visual novels does not give us reason to think that “the lookings at a picture and the listenings to a performance” (Goodman, 1968, 114) are well-formed instances of paintings and musical works, respectively? Suppose that “a looking” and “a listening” are used in the first sense. Then the lookings at a painting and the listenings to a performance are particular acts. However, such acts, obviously, do not manifest the properties bearing on the appreciation of the corresponding works and, hence, the primary properties that must be experienced to fully appreciate these works. So, understood in this sense, neither the lookings nor the hearings can be well-formed instances of paintings and musical works, respectively. Suppose next that “a looking” and “a listening” are used in the second sense. In this case, the lookings at a painting and the listenings to a performance are particular perceptual or quasi-perceptual (imaginary) experiences, namely the experience of looking at a painting and the experience of listening to a performance, respectively. However, such experiences do not manifest the properties of the objects of which they are experiences. Thus, when you

¹⁴See Goodman (1968).

are looking at a black painting, the experience that you get—the experience *of* this painting—is not itself black and, hence, does not manifest the property of blackness (what manifests this property is the painting). Similarly, when you hear a particular sound, the experience that you get—the experience *of* this sound—does not emit any sound and so does not manifest the property of sounding a particular way (it is the sound that manifests this property). Meanwhile, if the lookings at a painting and the listenings to a performance do not manifest the properties of the objects of which they are experiences, then neither the lookings nor the listenings manifest the properties that must be experienced to fully appreciate the corresponding works, and so neither the lookings nor the listenings can be well-formed instances of paintings and musical works, respectively.

4.2.2 Mereological Sums of Readings and Graphic Elements as Well-Formed Instances of Visual Novels

Thus, the entities that *can* serve as well-formed instances of *non-visual* novels are readings. What are the entities that *can* serve as well-formed instances of *visual* novels?

The first thing to note is that these entities cannot be readings. To be an instance of an artwork, an entity must manifest all the primary properties that must be experienced to fully appreciate this work. As shown in Section 4.1, the primary properties that must be experienced to fully appreciate visual novels include certain visual properties. Thus, the entities that serve as well-formed instances of visual

novels must be capable of manifesting visual properties. However, readings cannot manifest such properties. For (a) an entity can manifest a property only if this property is apprehensible by directly perceiving this entity, but (b) no visual property is apprehensible by directly perceiving a reading (an essentially sonic entity).

What are then the entities that serve as well-formed instances of visual novels? As is clear from what has been said earlier, the primary properties that must be experienced to fully appreciate a visual novel are certain sonic properties, experienceable properties that provide experiential access to the semantic content, and visual properties. As noted in the previous subsection, the sonic properties and experienceable properties that provide experiential access to the semantic content can be manifested by readings. Now, what about the visual properties? What can they be manifested by? The answer, I think, is as follows: The visual properties can be manifested by various graphic elements—for example, pictures of certain unreadable singular symbols and sequences of symbols, in the case of *The Da Vinci Code*; words printed in blue and purple, in the case of *House of Leaves*; a drawing of the map of Treasure Island, in the case of *Treasure Island*; pictures of things like the yin-yang symbol, female underpants, and Eliot Rosewater’s handwriting, in the case of *Breakfast of Champions*; and photographs, maps, and drawings, in the case of *Schwindel. Gefühl*.

In light of what has been said, the entity that can manifest all the primary properties that must be experienced to fully appreciate a visual novel is whatever involves a reading of this novel and particular graphic elements. Now, of course, a natural question arises: Ontologically speaking, what is this entity?

One possible answer to the foregoing question is that it is *a set* of a reading and particular graphic elements. This answer, however, is unacceptable. The consensus is that sets are abstracta. Meanwhile, as is generally agreed, abstracta are directly imperceptible and, hence, cannot manifest any properties. So if the entity that involves a reading and graphic elements is a set, then it cannot manifest any properties—and, hence, cannot be the entity that can manifest all the primary properties that must be experienced to fully appreciate a visual novel.

One could respond to the foregoing objection by adopting a view according to which sets are concreta—say, the view advocated by Maddy (1990). However, this response is unappealing. As shown by Balaguer (1998) and Carson (1996), among others, Maddy (1990)’s view is highly problematic. And, to my knowledge, there is no other potentially acceptable view according to which sets are concreta.

Another possible answer to the question posed above is that the entity in question is *a mereological sum* (or fusion) of a reading and graphic elements. This answer has at least two advantages. First, it is categorially transparent: It entails that the entities containing readings and graphic elements fall under the category of “a mereological sum,” and this category is reasonably well understood in metaphysics. Second, the answer being discussed does not face the analog of the problem of the previous answer. It would face this analog only if a mereological sum of a reading and graphic elements were abstract. However, such a sum is concrete—for, (a) it is a mereological sum of certain concreta—namely, a reading and graphic elements—and (b) any mereological sum of concreta is concrete.

Here, one could object as follows. The foregoing answer implies that mereolog-

ical sums of readings and graphic elements are metaphysically unproblematic. But is this implication, in fact, true? Mereological sums of readings and graphic elements are concreta constituted by readings and graphic elements. As has been demonstrated above, readings are events. And graphic elements are objects. So mereological sums of readings and graphic elements are concreta constituted by events and objects. However, no standard concretum—concretum included in what might be called “the ontology of ordinary language”—is constituted this way. Meanwhile, if this is so, then mereological sums of readings and graphic elements are rather strange entities and, hence, are metaphysically problematic.

This objection assumes that the fact that a mereological sum of a reading and graphic elements, being an entity constituted by an event and an object, is a strange concretum entails that such a sum is metaphysically problematic. Is this assumption true? Surely, a mereological sum of a reading and graphic elements is, in some sense, a strange concretum—a concretum that, unlike “ordinary” concreta, is constituted by both an event and objects. But why think that this entails that such a sum is problematic, from a metaphysical viewpoint? The fact that an entity is strange does not, *by itself*, show that there is something wrong with this entity. For there are a number of entities that are considered strange but at the same time universally acknowledged as part of our ontology (consider, for instance, quarks). Thus, a mereological sum of a reading and graphic elements cannot be metaphysically problematic merely because it is strange. Can it be metaphysically problematic because it is *particularly* strange—that is, because, unlike “ordinary” concreta, it is constituted by an event and objects? *Prima facie*, the answer is “No.” The mere

fact that a concretum is constituted by an event and objects does not seem to be a reason to hold that there is something wrong with this entity, from a metaphysical viewpoint.

In light of what has been said, the question of whether the assumption being discussed is true should be answered in the negative. If this answer is correct—and there seems no reason to think otherwise—the objection being discussed involves a false assumption and, hence, fails.

One could also object to the answer that the entity involving a reading and graphic elements is a mereological sum as follows. This answer requires us to add a new kind of entity—the mereological sum of a reading and graphic elements—to our ontology and thereby make this ontology less parsimonious. However, making an ontology less parsimonious is highly undesirable.

Like the previous objection, this objection fails. It assumes that making an ontology less parsimonious is undesirable. But this assumption is false. Making an ontology less parsimonious is undesirable only if doing so is unreasonable. However, in our case, there is a good reason to make the ontology less parsimonious. The reason is that we need to explain what sort of entity serves as a well-formed instance of a visual novel.

Thus, neither objection to the answer that the entity that involves a reading and graphic elements is a mereological sum is satisfactory. And there seem to be no other potentially satisfactory objections. Meanwhile, as mentioned earlier, this answer has at least two advantages. First, it is categorially transparent. Second, it avoids the objection to the answer according to which the entity that contains a

reading and graphic elements is a set. Given this, let us agree that the entity that contains a reading and graphic elements is a mereological sum.

Here, one might ask: What sort of thing is a mereological sum of a reading and graphic elements, ontologically speaking? To answer the foregoing question, it is necessary to answer the following questions:

- (1) Are mereological sums of readings and graphic elements concrete or abstract?
- (2) Are mereological sums of readings and graphic elements objects, events, or something else?
- (3) What sort of things are the constituents of these sums—readings and graphic elements, ontologically speaking?

The first of these questions has already been answered: Mereological sums of readings and graphic elements are concrete. Let us, therefore, turn to question (2). It is reasonable to assume that a mereological sum is (a) an event if and only if it is constituted *solely* by events, and (b) an object if and only if it is constituted *solely* by objects. However, as shown above, mereological sums of readings and graphic elements are constituted by both events and objects. As a result, such sums cannot be classified as objects or events. How should they be classified then? The best answer to this question, I think, is to classify them as some third kind of entity—an entity that is neither (purely) an event nor (purely) an object but that is *partly* an event and *partly* an object.

Let us now address question (3) (“What sort of things are readings and graphic elements, ontologically speaking?”). The nature of readings has already been clar-

ified in the previous subsection. The nature of graphic elements is relatively clear: Graphic elements are particular concrete objects, namely those concrete objects that manifest visual (primarily, color) properties.

Given what has been said, a mereological sum of a reading and graphic elements can be characterized as:

- a concretum;
- an entity that is neither (purely) an event nor (purely) an object but that is partly an event and partly an object;
- an entity that is constituted by a reading and graphic elements, where a reading is a sequence (a mereological sum) of certain sounds (= particular disturbance events, namely events in which the air is disturbed by certain movements performed by a reader), and graphic elements are certain concreta that manifest the relevant visual properties.

Having clarified the nature of mereological sums of readings and graphic elements, let us return to our main question—the question “What are the entities that can serve as well-formed instances of visual novels?” As has been established, the entities that can manifest all the primary properties that must be experienced to fully appreciate visual novels are mereological sums (hereafter: “sums”) of readings and graphic elements. Recall now that if an entity manifests all the primary properties that must be experienced to fully appreciate an artwork, then this entity is a well-formed instance_e of this work. So sums of readings and graphic elements can be well-formed instances_e of visual novels.

Can sums of readings and graphic elements be well-formed instances_p of visual novels? An entity is a well-formed instance_p of an artwork if this entity (a) manifests all the primary properties that must be experienced to fully appreciate this work and (b) stands in an appropriate historical-intentional relation to it. As already mentioned, sums of readings and graphic elements can manifest all the primary properties that must be experienced to fully appreciate visual novels. Furthermore, such sums doubtless can stand in any appropriate historical-intentional relation to the corresponding visual novels. So the question posed above can be answered in the affirmative.

4.2.3 Conclusion

Thus, readings *can* be well-formed instances of non-visual novels, while sums of readings and graphic elements *can* be well-formed instances of visual novels. In light of this, the questions posed at the beginning of Section 4.2—“What entities serve as well-formed instances of visual novels?” and “What entities serve as well-formed instances of non-visual novels?”—can be answered as follows: The entities that serve as well-formed instances of visual novels are sums of readings and graphic elements; the entities that serve as well-formed instances of non-visual novels are readings. Now, since, as has already been mentioned, answering these questions is sufficient to answer the question “What entities serve as well-formed instances of novels?,” the foregoing answer is also the answer to this latter question.

4.3 The Ontological Status of Non-Well-Formed Instances of Novels

Having established what entities serve as *well-formed* instances of novels, let us now establish what entities serve as *non-well-formed* instances of novels. The first thing to note is that establishing what entities serve as non-well-formed instances of novels makes sense only if we accept that there can be non-well-formed instances of novels. And we do not have to accept that. Thus, following Goodman (1968), we could maintain that novels as well as works of other notational arts do not have non-well-formed instances—for example, instances that involve minor errors (such as a single wrong letter or note)—since any entity that serves as an instance of a work of a notational art must comply perfectly with the canonical notation of this work. However, denying the possibility of non-well-formed instances of novels goes against the widely endorsed view according to which novels can have non-well-formed instances. In light of this, let us assume that non-well-formed instances of novels are possible.

What entities can serve as non-well-formed instances of novels? To be a non-well-formed instance of an artwork, it is sufficient (a) to be an instance of this work; (b) to manifest sufficiently many, but not all, of the primary properties that must be experienced to fully appreciate this work; and (c) to be, in principle, capable of manifesting all sensory kinds of these properties. Consider now readings. As already established, readings can be *well-formed* instances of non-visual novels. Consequently, readings can be instances (*simpliciter*) of non-visual novels. Furthermore, as shown above, readings can manifest all the primary properties that

must be experienced to fully appreciate a non-visual novel. So, clearly, readings can manifest sufficiently many, but not all, of these properties and are capable of manifesting all sensory kinds of these properties. Thus, the entities that can serve as non-well-formed instances of *non-visual* novels are readings.

Now, what about *visual* novels? The entities that can serve as non-well-formed instances of such novels are sums of readings and graphic elements. The argument in favor of this thesis is analogous to the foregoing argument in favor of the thesis that readings can serve as non-well-formed instances of non-visual novels.

Given what has been said, the answer to the question of what entities serve as non-well-formed instances of novels is the same as the answer to the analogous question about the entities that serve as well-formed instances of novels: The entities that serve as non-well-formed instances of novels are readings and sums of readings and graphic elements.

4.4 Objections

Let us now consider potential objections to the view that instances of novels are readings and sums of readings and graphic elements.

Objection 1. Suppose there is a manuscript that manifests the text of some novel *L*. Suppose also that this manuscript has never been read aloud. Then the view that instances of novels are readings and sums of readings and graphic elements entails that no instance of *L* has ever existed. But this entailment is mistaken. For a novel comes into existence only if its first instance comes into existence. So if there

have never been any instances of L , then L does not exist. But L does exist—qua an entity that is embodied (or contained) in its manuscript.

Response. This objection is based on the thesis that a novel comes into existence only if its first instance comes into existence. Is this thesis true? It is natural to assume that a novel comes into existence only if an entity that encodes the essential artistic properties of this novel (e.g., an inscription, a reading, or an audio recording (such as a CD or a computer file)) comes into existence. But not all encodings are instances of a novel. Consider, for example, an inscription. It is an encoding of a novel but, as shown in the previous chapter, cannot be an instance of a novel. Thus, assuming the foregoing assumption is correct—and there seems no good reason to think otherwise—the answer to the question posed above is “No”: A novel can come into existence even if no instance of this novel comes into existence.

It is also worth mentioning that the thesis that a novel comes into existence only if its first instance comes into existence is rejected by some ontologies of literature. Thus, this thesis is rejected by the ontology advanced by Urmson (2004). According to Urmson (2004), “for a literary work to exist it is a necessary and sufficient condition that a set of instructions should exist such that any oral performance which complies with that set of instructions is a performance of the work in question” (Urmson, 2004, 92). Now, the entity that serves as such a set of instructions is the inscription. So according to Urmson (2004), if an inscription of a novel exists, then the corresponding work exists. Meanwhile, as shown in the previous chapter, inscriptions cannot be instances of novels. Thus, Urmson (2004)’s ontology of literature entails that a novel can come into existence even if no instance of this novel

comes into existence.¹⁵

Furthermore, theses parallel to the thesis that a novel comes into existence only if its first instance comes into existence are rejected by a number of non-literary ontologies. Thus, Rohrbaugh (2003)'s ontology of analog photographs (hereafter: "photographs") rejects the thesis that photographs come into existence only if the corresponding instances come into existence. According to Rohrbaugh (2003), photographs

come into existence when they are taken. At the moment the button is pressed, the shutter opens and closes, exposing the film, and we say that we have 'taken a photograph.' The phrase has what is called 'success grammar.' If I forget to load the film and blithely snap away at your birthday party, then I should correct my claim to have taken photographs of it; without exposed negatives, no such photographs exist. What is so important about the moment at which the film is exposed? This event determines certain crucial facts about what the photograph is like, in particular, structural facts. What is in the frame and what is the composition? How much depth of field is there? Is it in color or black and white? All the qualities of a particular photograph have their start here. Before this moment, there are only shifting possibilities of photographs that might be taken; after it, the actual initial qualities are fixed for this one. (2003, 190)

¹⁵Another example of a literary ontology that rejects the thesis that a novel comes into existence only if its first instance comes into existence is the ontology advanced by Ingarden (1973).

So, in Rohrbaugh (2003)'s view, a photograph comes into existence when the film is exposed to light. Is this film an instance of the photograph? Surely, as Rohrbaugh (2003) points out, the film *encodes* the primary properties that must be experienced to fully appreciate the photograph. However, it does not *manifest* these properties. For it cannot manifest the visual properties that bear on the appreciation of the photograph,¹⁶ and there is no doubt that these properties constitute the primary properties that must be experienced to fully appreciate this photograph. Meanwhile, to be an instance of an artwork, an entity must manifest the primary properties that must be experienced to fully appreciate this work. So the film is not an instance of the photograph. As a result, on Rohrbaugh (2003)'s ontology, it is not necessary for an instance of a photograph to come into existence in order for the photograph itself to come into existence.

Similarly, a number of ontologies of musical works—in particular, the ontologies defended by Dodd (2004, 2007), Kivy (1983, 1987), and Levinson (1980, 2011b, 2012)—are incompatible with the thesis that a musical work comes into existence only if an instance of this work comes into existence. According to Dodd (2004, 2007)'s and Kivy (1983, 1987)'s ontologies, musical works are eternal entities (in particular, eternal abstract sound structures) and, hence, do not come into existence at all. Hence, if either ontology is correct, then no thesis that presupposes that musical works come into existence—including the thesis that they come into

¹⁶If the film could manifest the visual properties that bear on the appreciation of the photograph, then these properties would be apprehensible by directly perceiving this film. But, in fact, they cannot be apprehended that way.

existence only if the corresponding instances come into existence—can be true.

According to the ontology propounded by Levinson (1980, 2011b, 2012), a musical work comes into existence when, and only when, the composer creates the score of this work. So this ontology is compatible with the thesis that a musical work comes into existence only if an instance of this work comes into existence *only if* this score is an instance of the work. However, no score can be an instance of a musical work. Scores cannot manifest any sonic properties.¹⁷ Meanwhile, such properties are doubtless contained in the set of the primary properties that must be experienced to fully appreciate a musical work. So scores cannot manifest the primary properties that must be experienced to fully appreciate a musical work. But manifesting these properties is necessary to be an instance of this work.

Objection 2. Most of those who read novels do not listen to readings of these novels. Meanwhile, if instances of novels are readings and sums of readings and graphic elements, then one can fully appreciate novels only by listening to readings of these novels. So most readers do not fully appreciate the novels they read. However, this consequence is problematic.

Response. This objection attacks the view that instances of novels are readings and sums of readings and graphic elements by claiming that this view has a problematic consequence—that most of those who read novels do not fully appreciate them. Why is this consequence problematic? Presumably, it is problematic because it is false. But why think that it is false? That is, why think that most

¹⁷If a score could manifest sonic properties, then such properties would be apprehensible by directly perceiving this score. But, in fact, they cannot be apprehended this way.

of those who read novels *fully* appreciate them? One could answer this question as follows: If readers did not fully appreciate the novels they read, they would be aware of that, but that is not how things are: Most readers do not think that they fail to fully appreciate the novels they read. This answer, however, can hardly be called satisfactory, as there seems no real reason to hold that if readers do not fully appreciate the novels they read, then they must be aware of that.

There seems no other potentially plausible explanation of why the consequence being discussed could be false. At the same time, there is good reason to consider this consequence true. As shown in the previous chapter, a full appreciation of a novel presupposes an experiential engagement with the sonic properties of this novel. But the only way to experience these properties is to listen to a reading of this novel. Thus, contrary to what Objection 2 suggests, readers cannot fully appreciate a novel without listening to a reading of this novel.¹⁸

Objection 3. Imagine a novel M that is composed entirely of meaningful but unpronounceable strings of characters and, as a result, cannot be read aloud. According to the view that instances of novels are readings and sums of readings and graphic elements, no instance of M is possible. But, surely, M can have instances.

Response. This objection assumes that novels that cannot be read aloud are covered by the view that instances of novels are readings and sums of readings and graphic elements. But this assumption is false. As noted in the previous chapter,¹⁹

¹⁸A similar response can be used to defuse the objection based on the idea that novels have clearly been intended, and for centuries, to be read, not listened to via readings of them aloud.

¹⁹See Chapter 3, Footnote 1.

the term “a novel” is restricted to novels that can, in principle, be read aloud. And so the view that instances of novels are readings and sums of readings and graphic elements applies only to such novels.²⁰

²⁰Here, one might ask: What are then the entities that serve as instances of novels that cannot be read aloud? Like novels that can be read aloud, novels that cannot be read aloud can be divided into non-visual novels, or novels that do not contain any aesthetically relevant visual elements, and visual novels, or novels that do contain such elements. Let us first establish what entities serve as well-formed instances of *non-visual* novels that cannot be read aloud.

To be an instance_e of an artwork, it is sufficient to manifest all the primary properties that must be experienced to fully appreciate this work. What entities are capable of manifesting all such properties, in the case of non-visual novels that cannot be read aloud? To answer this question, we first need to clarify what primary properties must be experienced to fully appreciate such novels. Clearly, since novels that cannot be read aloud cannot, in principle, be heard, no sonic property needs to be experienced to fully appreciate them. Thus, sonic properties are not contained in the set of the primary properties that must be experienced to fully appreciate non-visual novels that cannot be read aloud. Presumably, the only properties that must be contained in this set are certain experienceable properties that provide experiential access to the semantic content. Now, inscriptions are doubtless capable of manifesting such properties. So inscriptions can manifest all the primary experienceable properties that must be experienced to fully appreciate non-visual novels that cannot be read aloud—and, hence, can be well-formed instances_e of such novels.

Furthermore, inscriptions can be well-formed instances_p of non-visual novels that cannot be read aloud. As shown in the previous paragraph, inscriptions can manifest all the primary properties that must be experienced to fully appreciate such novels. And there is no doubt that inscriptions can stand in an appropriate historical-intentional relation to them. Meanwhile, to be an instance_p of an artwork, it is sufficient to (a) manifest all the primary properties that must be experienced to fully appreciate this work and (b) stand in an appropriate historical-intentional relation to it.

What entities serve as well-formed instances of *visual* novels that cannot be read aloud? Clearly,

Objection 4. Suppose there is some novel N that is written in an archaic language. Suppose also that no one who reads N aloud pronounces all of its words correctly. Then, according to the view that well-formed instances of novels are read-

the set of primary properties that must be experienced to fully appreciate such novels does not contain any sonic properties. What this set contains are (a) certain experienceable properties that provide experiential access to the semantic content and (b) certain visual properties. Now, it is clear that mereological sums (hereafter: “sums”) of inscriptions and graphic elements are capable of manifesting the foregoing properties. Hence, such sums can manifest all the primary properties that must be experienced to fully appreciate visual novels that cannot be read aloud. Meanwhile, if an entity manifests all the primary properties that must be experienced to fully appreciate an artwork, then this entity is an instance_e of this work. So sums of inscriptions and graphic elements can be well-formed instances_e of visual novels that cannot be read aloud.

Furthermore, sums of inscriptions and graphic elements can be well-formed instances_p of visual novels that cannot be read aloud. As mentioned above, these sums can manifest all the primary properties that must be experienced to fully appreciate such novels. Furthermore, these sums doubtless can stand in an appropriate historical-intentional relation to them. Meanwhile, if an entity (a) manifests all the primary properties that must be experienced to fully appreciate an artwork and (b) stands in an appropriate historical-intentional relation to this work, then this entity is a well-formed instance_p of this work.

Thus, given what has been said, it is reasonable to hold that the entities that serve as well-formed instances of non-visual novels that cannot be read aloud are inscriptions, whereas the entities that serve as well-formed instances of visual novels that cannot be read aloud are sums of inscriptions and graphic elements. In light of this, and taking into account the fact that the distinction between non-visual and visual novels is exhaustive, the question “What entities serve as well-formed instances of novels that cannot be read aloud?” can be answered as follows: The entities that serve as such instances are inscriptions and sums of inscriptions and graphic elements.

Now, what entities serve as non-well-formed instances of novels that cannot be read aloud? To

ings and sums of readings and graphic elements, there are no well-formed instances of N . Meanwhile, if there are no such instances, then N cannot be appreciated. But, surely, this consequence is false: It is possible for us to appreciate N .

Response. According to this objection, the fact that there are no well-formed instances of a novel implies that this novel is completely inaccessible to our appreciation. Is this implication true? Suppose there are no ideal performances of *The Rite of Spring* (each performance contains at least one wrong note)²¹ or any other be a non-well-formed instance of an artwork, it is sufficient (a) to be an instance of this work; (b) to manifest sufficiently many, but not all, of the primary properties that must be experienced to fully appreciate this work; and (c) to be, in principle, capable of manifesting all sensory kinds of these properties. As has been shown, inscriptions are well-formed instances of non-visual novels that cannot be read aloud and, hence, are instances of such novels. Furthermore, given what has been said above, it is clear that inscriptions can manifest sufficiently many, but not all, of the primary properties that must be experienced to fully appreciate novels. Finally, inscriptions doubtless can, in principle, manifest all sensory kinds of these properties. Thus, inscriptions can be non-well-formed instances of non-visual novels that cannot be read aloud.

Now, what are the entities that can be non-well-formed instances of visual novels that cannot be read aloud? These entities can be sums of inscriptions and graphic elements. The argument in favor of this thesis is analogous to the foregoing argument in favor of the thesis that inscriptions can be non-well-formed instances of non-visual novels that cannot be read aloud.

Given what has been said, the answer to the question of what entities serve as non-well-formed instances of novels that cannot be read aloud is the same as the answer to the analogous question about the entities that serve as well-formed instances of novels that cannot be read aloud: The entities that serve as non-well-formed instances of novels that cannot be read aloud are inscriptions and sums of inscriptions and graphic elements.

²¹Given the extreme complexity of *The Rite of Spring*, this supposition may well be true.

entities manifesting all the primary properties that must be experienced to fully appreciate *The Rite of Spring*. Then there are no well-formed instances of *The Rite of Spring*. Despite this, however, we may still be able to appreciate it. Of course, we will not be able to appreciate it *in full*.²² But we will be able to appreciate it *partially*—say, by listening to one of its imperfect performances.

Thus, we can appreciate a musical work even if there are no well-formed instances of this work. And the same can be said not only about musical works but about all other artworks. The absence of well-formed instances of an artwork does not necessarily preclude us from appreciating this artwork. As a result, Objection 4 is based on a false assumption—that a novel cannot be appreciated if there are no well-formed instances of this novel.

Here, one could object as follows. The foregoing response assumes that a novel cannot be fully appreciated without appreciating a well-formed instance of this novel. But this assumption is false. In fact, we can fully appreciate a novel even if it does not have any well-formed instances.

This objection, however, fails. A full appreciation of a novel presupposes an appreciation of all the primary properties that must be experienced to fully appreciate this novel. Meanwhile, if there are no well-formed instances of a novel, then *all* of these properties cannot be appreciated. Thus, contrary to what the objection states, the possibility of a full appreciation of a novel presupposes the

²²To be able to do this, we must listen to an entity manifesting all the primary properties that must be experienced to fully appreciate *The Rite of Spring*. But, by assumption, no such entity exists.

existence of at least one well-formed instance of this novel.

It is worth noting that besides the thesis that an artwork can be appreciated even if there are no well-formed instances of this work, we can adopt a stronger thesis—that an artwork can be appreciated even if there are no instances of this work whatsoever. Suppose the original canvas of the *Mona Lisa* has been completely destroyed and, as a result, there are no instances of this painting. Suppose next that there are a number of reproductions of the *Mona Lisa*.²³ Then it is still possible for us to appreciate the *Mona Lisa*. We “might reasonably claim to be knowledgeable about [this] painting and even have an appreciation of its beauty and power through studying reproductions alone” (Lamarque, 2010, 59). Now, what has been said about the *Mona Lisa* applies to other artworks as well. Given this, it is reasonable to hold that the absence of instances of an artwork does not necessarily preclude us from appreciating this work.²⁴

Objection 5. When we apprehend epistolary novels (such as Samuel Richardson’s *Clarissa*, Fanny Burney’s *Evelina*, and Thornton Wilder’s *The Ides of March*), it is reasonable to assume that we are supposed to imagine ourselves apprehending letters that were not intended to be read aloud. But if this is so, then there is

²³It is assumed here that the only instance of a painting is its original canvas. This assumption reflects the consensus on what entities serve as instances of paintings.

²⁴Perhaps the same can also be said about at least some non-artistic objects. Suppose there are no Benz Patent-Motorwagens and, hence, no instances of the Benz Patent-Motorwagen. In this case, it seems natural to suppose that we can still appreciate—at least, to a certain extent—the Benz Patent-Motorwagen—say, by appreciating a die-cast Benz Patent-Motorwagen model or photographs of a Benz Patent-Motorwagen.

good reason to believe that the sonic dimension of such novels is irrelevant to their aesthetic appreciation, and, as a result, there seems to be no motivation to hold that their instances are identical to or incorporate readings.

Response. The reason why letters are not supposed to be read aloud is that their content is not intended for anyone other than the reader, and not because their sonic properties are aesthetically irrelevant. So perhaps, when we apprehend epistolary novels, we should assume that we are supposed to imagine ourselves apprehending letters which were not intended to be read aloud *to anyone other than ourselves*. However—contrary to what the objection states—we do not have any reason to believe that when apprehending such novels, we should assume that we are supposed to imagine ourselves apprehending letters which were not intended to be read aloud *at all*.

Objection 6. Suppose we listen to a reading of *War and Peace* in English. In this case, we cannot fully appreciate *War and Peace*, since we have access only to an instance of a translation of *War and Peace* (a particular English reading of it), and not to an instance of *War and Peace* (a particular Russian reading of it). However, this result seems wrong. Intuitively, when we listen to a reading of *War and Peace* in English, we can fully appreciate *War and Peace* itself, and not just a translation of this novel.

Response. This objection assumes that when we listen to a reading of *War and Peace* in translation, we can fully appreciate *War and Peace*. Is this assumption true? Here is a possible argument in favor of answering this question in the affirmative. *War and Peace* and its translation are identical. Therefore, whenever

we listen to a reading of *War and Peace* in translation, we listen to a reading of *War and Peace*—and, hence, can fully appreciate *War and Peace* (assuming that *War and Peace* can be fully appreciated by listening to its reading).

Is this argument satisfactory? If the artistic value of an artwork depends on some intrinsic property of this work, then the identity of this work depends on this property. Meanwhile, as shown in the previous chapter, the artistic value of a novel depends on the sonic and semantic properties of this novel, which are doubtless intrinsic.²⁵ Thus, given that both *War and Peace* and its translation are novels, their identities depend on these properties. Now, it is clear that the language of *War and Peace* and the language of the translation are different. So at least some sonic and semantic properties of *War and Peace* differ from the sonic and semantic properties of its translation. But then, contrary to what the foregoing argument states, *War and Peace* cannot be identical to its translation.

There seem no other potentially satisfactory arguments in favor of the assumption that we can fully appreciate *War and Peace* by listening to a reading of *War and Peace* in translation. At the same time, there is a strong consideration against this assumption. To fully appreciate an artwork, it is necessary to appreciate what the aesthetic value of this work is grounded in. The aesthetic value of a novel is

²⁵Following Lewis (1983), “an intrinsic property” can be defined as follows: For all x , x is an intrinsic property of y just in case x is a property possessed by y “in virtue of the way [y] itself, and nothing else, is” (Lewis, 1983, 112). For a detailed analysis of the expression “an intrinsic property” (as well as the related expression “an extrinsic property”), see Weatherson and Marshall (2014).

grounded in certain sonic and semantic properties. So a proper appreciation of a novel is impossible without an appreciation of these properties. Can we appreciate the sonic and semantic properties of *War and Peace* by listening to a reading of a translation of this novel? Since the sonic and semantic properties of the translation of *War and Peace* are different from the sonic and semantic properties of *War and Peace*, a reading of this translation cannot have the sonic and semantic properties of *War and Peace*. But if this is so, then listening to such a reading cannot enable us to appreciate the sonic and semantic properties of *War and Peace*. Consequently, since *War and Peace* cannot be fully appreciated without appreciating its sonic and semantic properties, *War and Peace* cannot be fully appreciated by listening to its translation.

Note that the impossibility of a *proper* appreciation of *War and Peace* by listening to a reading of *War and Peace* in translation does not entail that this novel cannot be appreciated *at all* by listening to such a reading. Listening to a reading of *War and Peace* in translation can enable us to apprehend some core semantic properties (such as the properties that constitute the characters and plot), properties related to the structure, and perhaps even certain sonic properties (for example, some properties concerned with rhythm). Meanwhile, apprehending these properties makes it possible for us to appreciate *War and Peace*.

Note also that Objection 6 does not pose a *special* threat to the view that instances of novels are readings and sums of readings and graphic elements. This objection would pose such a threat only if it were applicable to this view, and not the Orthodox View—that is, the view that instances of novels include inscriptions.

However, this objection can be applied to the Orthodox View as well. For one can say: “Suppose we read an inscription of *War and Peace* in English. In this case, we cannot fully appreciate *War and Peace*, since we have access only to an instance of a translation of *War and Peace* (a particular English inscription), and not to an instance of *War and Peace* (a particular Russian inscription). However, this result seems wrong. Intuitively, when we read an inscription of *War and Peace* in English, we can fully appreciate *War and Peace* itself, and not just a translation of this novel.”

4.5 Some Consequences

Let us now examine some potential consequences of the view that instances of novels are readings and graphic elements.

One of these consequences is that to fully appreciate a novel, one must listen to a reading of this novel—a particular sequence of sounds, which can be generated by the reader herself, or someone else, or some electronic device (such as a computer or an e-reader). For suppose the foregoing view is true. Then a well-formed instance of a novel either is identical to or incorporates a reading. Meanwhile, to fully appreciate an artwork, we must experientially engage with its well-formed instance.²⁶ Thus, a full appreciation of a novel requires an experiential engagement with a reading of

²⁶Suppose we do not experientially engage with a well-formed instance of an artwork. Then, since only such an instance manifests all the primary properties that must be experienced to fully appreciate this work, we do not experience all of these properties. But if this is so, then we do not fully appreciate the work.

this novel. How can one experientially engage with a reading? The only way to do this is to listen to it. So if the view that instances of novels are readings and graphic elements is true, then a full appreciation of a novel requires listening to a reading of this novel.

Another consequence of the view that instances of novels are readings and sums of readings and graphic elements is that a novel cannot be fully appreciated by means of silent reading. Suppose this view is true. Then, as mentioned above, to fully appreciate a novel, we must listen to its reading. However, silent reading does not involve the production of any physical sounds and, hence, does not enable us to listen to a reading. Thus, by reading a novel silently, we cannot fully appreciate it.

Yet another consequence of the view that instances of novels are readings and sums of readings and graphic elements is that we may want to stop reading novels silently and start listening to them. Suppose this view is true. Then, as has been shown above, we cannot read a novel silently if we want to fully appreciate it; if we want that, we should listen to its reading. But if this is so, then, assuming that we would like to fully appreciate novels, we may want to stop reading novels silently and start listening to their readings.

Here, a natural question arises: Is realizing this change in our way of appreciating novels practicable? There are two potential reasons to answer “No.” First, it is hard to get access to readings of novels; second, listening to a reading of a novel is less convenient than reading it silently. Is either reason persuasive?

Consider the first reason. Why is it hard to get access to readings of novels? Is it hard because the means to generate such readings are unavailable? That cannot

be true. For there are numerous electronic devices—computers, audio players, cell phones, etc.—that can be used to produce and/or re-produce readings. Another potential explanation of why it might be hard to get access to readings of novels is that it is hard to get access to the sources of such readings. But, like the previous explanation, this explanation is unsustainable. For there are numerous audio versions of various novels (“audiobooks”), which can easily be accessed over the Internet. True, audio versions are not available for *all* novels. But the “audiobook” market is expanding: More and more audio versions of novels are becoming available as time goes by. So it is reasonable to suppose that in the not too distant future, there will be audio versions of all novels.

Thus, neither explanation of why it is hard to get access to readings of novels is satisfactory. Meanwhile, there seems to be no other potentially plausible explanation of that. So it is unclear why it is hard to get access to readings of novels. Hence, the first reason in favor of the claim that switching from reading a novel silently to listening to its reading is problematic appears unpersuasive.

What about the second reason? Is listening to a reading of a novel less convenient than reading this novel silently? Perhaps that is so under certain circumstances. However, in some cases, it is more convenient to listen to a reading of a novel rather than read this novel silently. Clearly, when you drive a car, you do not want to read a novel silently; but, at the same time, you can easily listen to a reading of this novel. Also, when you are tired, you may prefer listening to a reading of a novel to reading this novel silently. Furthermore, listening to a reading of a novel is preferable to reading this novel silently if you want to reduce your eye strain.

Given what has been said, the second reason in favor of the claim that switching from reading novels silently to listening to their readings is problematic also seems unpersuasive.

Thus, neither reason is persuasive. Meanwhile, there seems to be no other potentially satisfactory reason to hold that it will be problematic for us to change our way of appreciating novels by switching from reading them silently to listening to their readings. But if this is so, then nothing stops us from thinking that changing this way is, in fact, quite practicable.

Another consequence of the view that instances of novels are readings and sums of readings and graphic elements is that, given certain plausible assumptions, novels do not depend for their existence on their instances. Suppose this view is true. Then, since most novels have never been sounded out, most novels do not have instances. Meanwhile, if this is so, then the supposition that a novel exists only if at least one instance of this novel exists entails that most novels do not exist. But most novels do exist—or so it seems. So we arrive at a contradiction. How can it be eliminated if we are to assume that instances of novels are readings and sums of readings and graphic elements? One way to do this is to reject the thesis that most novels exist. Another way is to reject the idea that a novel cannot exist if at least one instance of this novel does not exist. Which way is preferable? If we choose the first way, we must violate one of our deeply entrenched artistic intuitions—that novels that we consider existent actually exist. Meanwhile, as is generally agreed, violating deeply entrenched artistic intuitions is highly undesirable. So the first way of eliminating the contradiction is unsatisfactory. Consider now the second way.

Given what has been said in the previous section (in particular, the response to Objection 1), there is no real reason against choosing this way. But if this is so, then the second way of eliminating the contradiction being discussed is preferable to the first way. Meanwhile, this way presupposes a rejection of the thesis that novels depend for their existence on the existence of their instances.

Here, one might ask: If the existence of novels does not depend on the existence of their instances, then what, if anything, does it depend on? There can be different answers to this question. One possible answer, advocated by pure platonists, is that novels, being eternal and indestructible abstracta of some kind, do not depend for their existence on anything (other than perhaps God). Another possible answer, endorsed by Levinson (2011b, 2012), is that novels, being initiated types, depend for their *coming into* existence on their authors' acts of indicating particular abstract structures (by means of creating certain inscriptions), but do not depend on anything for their *subsequent* existence.²⁷ According to yet another possible answer, upheld by

²⁷Levinson (2011b, 2012) does not explicitly say that he endorses the idea that novels do not depend on anything for their subsequent existence. However, that he actually endorses this idea can be shown as follows. In "What a Musical Work Is, Again," he points out that he is not sure "what would be the rationale" (Levinson, 2011b, 262) for holding that "the permanent elimination of all records and memories of [the musical work] would suffice to destroy it" (Levinson, 2011b, 262). He then says:

... The residual pull of [the idea that "nothing can destroy [the musical work], once created" (Levinson, 2011b, 262)] is hard for me to deny. Once a PSS [performed-sound structure], in a coherent musical context, it *might* just inhabit the abstract realms of the universe, it seems, forever. Why should it lapse into nonexistence, one might

Rohrbaugh (2003), novels, being “historical individuals,” depend for their existence on their “embodiments”—particular physical objects that ground the facts about the novels’ essential properties.

Finally, the view that novels are properly instanced solely by readings and sums of readings and graphic elements entails that the art of the novel has a lot more in common with performing arts (especially, with music) than usually thought. We generally believe that the art of the novel is considerably different from any of the performing arts. Thus, in our view, the most typical, or paradigmatic, instances of novels (inscriptions) differ, in their fundamental nature, from instances of works of performing arts (performances). In particular, we believe that paradigmatic instances of novels are objects, whereas instances of performing arts are events. Furthermore, we think that the way paradigmatic instances of novels are generated is different from the way instances of performing arts are generated: Paradigmatic instances of novels are generated by means of writing or copying, whereas instances of works of most (though not all²⁸) performing arts are generated by means of

ask, just because we do? It is perhaps a comforting thought that the nonmaterial products of culture, once given start, may be logically destined to outlast us—at least in the rarefied sense here in question. (Levinson, 2011b, 263)

Thus, according to Levinson (2011b), musical works, once created, are indestructible and, hence, do not depend for their existence on anything. Meanwhile, as is clear from his “Indication, Abstraction, and Individuation,” in his view, the fundamental existence conditions of novels (as well as all other literary works) are the same as the fundamental existence conditions of musical works.

²⁸In jazz, which is doubtless a performing art, instances are normally produced without recourse to instructions.

performing in accordance with particular written and/or oral instructions (or quasi-instructions). Finally, we usually hold that paradigmatic instances of novels differ from instances of works of performing arts with regard to the art status: Paradigmatic instances of novels are not artworks themselves, since these instances do not have any artistically significant properties besides those that belong to the novels; whereas instances of works of performing arts do possess such properties and, therefore, can be treated as artworks in their own right.²⁹

Now, if the view that novels can only be properly instanced by readings and sums of readings and graphic elements is adopted, then the mentioned dissimilarities largely (though not completely) disappear. For suppose that instances of novels are readings and sums of readings and graphic elements. Then, since readings are events, instances of novels are identical to or contain events and, hence, are close (albeit not the same), in their fundamental nature, to instances of performing arts. Furthermore, since readings are generated with the help of inscriptions, which can be characterized as instructions (or quasi-instructions), the way instances of novels are generated is similar to the way instances of works of performing arts are generated.

²⁹Are instances of works of performing arts, in fact, artworks? There are powerful reasons to answer this question in the affirmative (see, e.g., Alperson (1984) and Kivy (1995)). But what if the correct answer is “No” (for reasons in favor of this answer, see Thom (1993) and Kania (2011))? In this case, the theses “Paradigmatic instances of novels are not artworks themselves” and “Instances of works of performing arts possess artistic properties and, therefore, can be treated as artworks in their own right” should be replaced with the theses “Paradigmatic instances of novels do not deserve artistic appreciation” and “Instances of works of performing arts possess artistic properties and, therefore, can be treated as objects that deserve artistic appreciation,” respectively.

Finally, because readings have particular artistically relevant properties that belong to them, rather than to the corresponding novels, readings can be treated as artworks in their own right—and so, with regard to the art status, instances of novels are similar to instances of works of performing arts.³⁰

That the art of the novel has so much in common with performing arts does not, by itself, mean that it is one of them. For it also differs from them in at least two important respects. First, not all instances of novels are pure events; some are sums of events and objects. However, instances of works of performing arts are generally considered to be pure events. Second, instances of novels can be easily created by most of those who speak English (or any other sound-based language) and know how to generate graphic elements using a printer or some other copying device. But most instances of works of performing arts can be created only by a relatively small number of specialists (actors, stage directors, musicians, dancers, etc.).

³⁰There seems no good reason to think that readings cannot be artworks. However, if there is, in fact, such a reason, then the thesis “Because readings have particular artistically relevant properties that belong to them, rather than to the novels, readings can be treated as artworks in their own right” should be replaced with the thesis “Because readings have particular artistically relevant properties that belong to them, rather than to the novels, readings can be treated as objects that deserve artistic appreciation.”

4.6 Final Remarks

In closing, it is worth making two additional remarks concerning the view that instances of novels are readings and sums of readings and graphic elements. First: This view naturally invites the question “What makes a reading a reading of a given novel?” This question is doubtless important. Despite this, however, I will not address it—for two reasons. First, a satisfactory answer to it would require an investigation that, given space limitations, is beyond the scope of this dissertation. Second, answering this question here is unnecessary, since the truth of the thesis that instances of novels are readings and sums of readings and graphic elements does not depend on whether we know what makes a reading a reading of a given novel. If the truth of this thesis depended on whether we know what makes a reading a reading of a given novel, then, by analogy, the truth of the thesis that instances of musical works are performances would depend on whether we know what makes a performance a performance of a given musical work. But the latter kind of dependence does not hold. For if it did, we would have to question the thesis that instances of musical works are performances—given the fact that there is no generally accepted theoretic answer to the question of what makes a performance a performance of a given musical work. But, of course, we do not want to question this thesis. (Surely, musical works are properly instanced by performances.)

Before proceeding further, however, I would like to mention three important constraints on any satisfactory answer to the question of what makes a reading a reading of a given novel. First, any such answer must be compatible with the idea

that to be a reading of a given novel, a reading must manifest the text of this novel without omitting important parts of this text or adding extraneous elements. Next, any satisfactory answer must entail that a reading is a reading of a particular novel only if it follows the norms of pronunciation prescribed by the linguistic context of the novel. Finally, any satisfactory answer must allow for some phonetic variability between readings of the same novel. For suppose an answer does not allow for this. Then according to this answer, only sonically identical readings can be readings of the same novel. But this is too restrictive. Surely, we would like to regard certain appropriate readings with minor phonetic differences (such as slight differences in intonation or in pronouncing [s] and [z]) as readings of the same novel.

The second remark concerns a potential extension of the view that instances of novels are readings and graphic elements. Clearly, this view has a rather limited scope: It applies only to novels. But it can be extended to include other textual entities with similar primary appreciation-relevant experienceable properties. The most obvious candidates for such inclusion are poems and short stories. Other possible candidates are speeches (for example, Cicero's *Catilinarian Orations* and Martin Luther King's "I Have a Dream"), newspaper and magazine articles, pamphlets, letters, and any other textual entities that are subject to aesthetic appreciation.

Chapter 5: The Novel as a Performing Art

Introduction

Our analysis has shown that instances of novels are readings and sums of readings and graphic elements. If this result is correct, then, as noted in the previous chapter,¹ the novel is rather similar to performing arts (primarily, music, dance, and theater). In light of this, a natural question arises: Is the novel one of such arts?

My goal in what follows is to show that there is good reason to answer this question in the affirmative. I begin with a critique of the existing arguments in favor of the view that the novel is a performing art—namely, the arguments advanced by Kivy (2006) and Urmson (2004) (Section 5.1). Next, I put forward a new argument in favor of this view. I argue, in particular, that the novel is a performing art, since (a) it is an art such that to fully appreciate its work, one must experientially engage with either a performance, a playing of a recording of a performance, or a reproduction of a performance of this work,² and (b) being such an art is sufficient

¹See Chapter 4, Section 4.5.

²By “a reproduction of a performance” here is meant a sequence of sounds or movements that is generated, with the help of a set of instructions (such as a score), by a computer or some other mechanical device.

for being a performing art (Section 5.2). Finally, I make a remark concerning the question of whether the view that the novel is a performing art accords with the historical development of literature (Section 5.3).

5.1 Urmson's and Kivy's Arguments: A Critical Analysis

The view that the novel is a *non*-performing art is accepted by an overwhelming majority of theorists. Yet it is not accepted by all of them. Some—in particular, Urmson (2004) and Kivy (2006, 2010)—endorse the opposite view—that the novel, along with music, theater, and dance, is a performing art. To support this view, both Urmson (2004) and Kivy (2006, 2010) offer particular arguments. Let us examine these arguments.

5.1.1 Urmson's Argument

Urmson (2004)'s argument can be formulated as follows:

1. The novel is relevantly analogous to a particular kind of music, namely music intended primarily for score reading rather than performance.
2. Music intended primarily for score reading rather than performance is a performing art.
3. If (1) and (2) are true, then the novel (broader: all literature intended for silent reading) is a performing art.
4. So the novel (broader: all literature intended for silent reading) is a performing

art.³

The crux of this argument is clearly premise (1). Is this premise true?

Before answering this question, two remarks are worth making. First, (1) is not equivalent to the thesis that the novel is relevantly analogous to music. This premise is rather narrow: It implies that the novel is relevantly analogous only to *a particular kind* of music—music intended primarily for score reading rather than performance. The foregoing thesis, on the other hand, is much broader: It implies that the novel is relevantly analogous not just to the mentioned kind of music but to music *in general*.

Second, according to Urmson (2004), the thesis that the novel is relevantly analogous to music in general is false. The reason for this, in his view, is that there is an apparent disanalogy between the novel, on the one hand, and “ordinary” music (i.e., music that is intended for performance rather than score reading), on the other. In the case of the latter, there are performers, or, to use Urmson (2004)’s term, “executant artists”—those who perform a musical work for an audience. In the case of the former, however, there seem to be no such artists. For,

...who could such artists be? When one... reads a novel to oneself, there seems to be only oneself and the novelist involved. Is the reader in fact the executant artist with himself as audience as the pianist who can play to himself as audience and the dancer who can dance for his own satisfaction? But I do not seem to myself to be exhibiting any technical

³Premises (2) and (3) are not explicitly stated by Urmson (2004). However, they (or their equivalents) must be included in the argument in order for it to be valid.

or interpretative skills when I read to myself, and there are other grave objections to this suggestion.... (Urmson, 2004, 90)

Let us now return to the question of whether premise (1) (“The novel is relevantly analogous to a particular kind of music, namely music intended primarily for score reading rather than performance”) is true. To answer this question, we first need to clarify the expression “score reading.” By “score reading,” Urmson (2004) means the process of (a) apprehending what a musical score expresses and (b) generating an experience that enables us to realize what we would hear were this score actually performed. Thus, when one reads a musical score—say, the score of Bach’s *Musical Offering*—she apprehends the content of this score and generates an experience that enables her to realize what she would hear if she heard a performance of the *Musical Offering*.⁴

It is important to stress that according to Urmson (2004), the experience generated in the process of score reading—the experience that enables one to realize what one would hear were the score actually performed—is not a musical *performance*. As Urmson (2004) points out, to be a musical performance, an entity must satisfy a number of conditions related to pitch, dynamics, tone color, duration, and some other aspects of the corresponding musical work. However, at least some of these conditions are invariably violated by the foregoing experience.

⁴According to Urmson (2004), score reading can be viewed as analogous to a particular way of reading a recipe by an experienced cook—the way according to which the cook apprehends what the recipe expresses—a particular set of instructions—and generates an experience that enables him to “recognize what the confection would taste like” (Urmson, 2004, 91).

Having clarified what is meant by “score reading,” let us now consider Urmson (2004)’s argument in favor of premise (1). As is clear from what has been said above, in the case of music intended primarily for score reading, we apprehend the content of a musical score and generate an experience that enables us to realize what we would hear were this score actually performed. In Urmson (2004)’s view, that is very similar to what we do in the case of literature and, in particular, in the case of the novel. In this case, we also apprehend the content of the “score” of a novel—a particular inscription—and generate an experience that enables us to realize what we would hear were this novel actually performed by reading it aloud. Thus, when we read *War and Peace* silently, we grasp the content of an inscription of this novel and generate an experience that enables us to realize what we would hear if this novel were actually read aloud. The foregoing similarity between the novel and music intended primarily for score reading is sufficient, according to Urmson (2004), to substantiate the thesis that the novel is relevantly analogous to music intended primarily for score reading—and, hence, premise (1).

Is Urmson (2004)’s argument acceptable? The key thesis of this argument is that when we read novels silently, we generate an experience that enables us to realize what we would hear were this novel actually read aloud. However, there is reason to doubt that this thesis is true. When I silently read a novel, I do not generate an experience that enables me to realize what I would hear if this novel were actually performed (by reading it aloud). And most other readers, I suppose, do not generate any such experience when reading novels silently either.

Thus, Urmson (2004)’s argument in favor of premise (1) involves a dubious

claim and, hence, cannot be accepted. At the same time, there is a powerful argument against this premise. It seems uncontroversial that in order for the novel to be relevantly analogous to music intended primarily for score reading, the way we read novels must be analogous to the way we read scores of musical works intended for score reading. However, in fact, the way we read novels is considerably different from the way we read scores of such musical works. As Urmson (2004) rightly notes, when we silently read the score of a musical work intended for score reading, our goal is to acquire a non-performing experience that enables us to realize what we would hear if the work were actually performed. But that is not what our goal is when we silently read novels. In this case, we do not try to acquire a non-performing experience that enables us to realize what we would hear if the work were performed. Our goal, when reading a novel silently, is to apprehend the story of this novel.⁵ (One might object that achieving this goal presupposes a generation of an experience that enables us to realize what we would hear if the novel were actually performed. But this objection fails, since, in fact, a generation of such an experience neither precedes, not accompanies, nor follows the process of apprehending the story of a novel.)

Thus, it is reasonable to consider premise (1) false. Meanwhile, if this premise is false, then Urmson (2004)'s argument in favor of the view that the novel is a performing art fails.

But what if premise (1) is true? Should Urmson (2004)'s argument then be

⁵Note that this claim is purely descriptive. It does not imply that when we read a novel silently, our goal *should be* to merely apprehend the story of this novel.

considered acceptable? No. For there is good reason against another premise of this argument—premise (2) (“Music intended primarily for score reading is a performing art”). Presumably, the reason why an art is called “performing” is that works of this art are to be appreciated through a performance or an entity that manifests all the relevant properties manifested by a performance.⁶ Consider now music intended primarily for score reading. Surely, it can be appreciated through a performance or an entity that manifests all the relevant properties manifested by a performance. But that is not the proper way to appreciate it. Since it is intended *primarily for score reading*, the proper way to appreciate it is by appreciating what results from score reading, which, as has been mentioned above, is a particular *non-performing* experience. But if this is so—if music intended primarily for score reading is to be appreciated through such an experience, and not through a performance or a performance-like entity—then, given the above characterization of “a performing art,” this kind of music is not, in fact, a performing art—and, hence, premise (2) is false.

5.1.2 Kivy’s Argument

Let us now turn to Kivy (2006, 2010)’s argument in favor of treating the novel as a performing art. This argument can be formulated as follows:

5. We appreciate novels through readings—particular sequences of sounds generated either in the external (spatiotemporal) world (when reading aloud) or

⁶This is a preliminary account of “a performing art.” A complete account of this expression is presented in Section 5.2.

in our mind (when reading silently)⁷.

6. These readings are artistic performances.
7. So we appreciate novels through artistic performances.
8. If we appreciate novels through artistic performances, then the novel is a performing art.
9. So the novel is a performing art.⁸

One could object to Kivy (2006, 2010)'s argument as follows. There is no doubt that sometimes—for example, when we listen to audiobooks—we appreciate novels through their readings. But what about the case when we read novels silently? In this case, there is no reason to hold that we generate any readings. But if this is so, then when we read silently, we do not appreciate novels through readings—and, hence, premise (5) is false.

This objection may seem strong, but, in fact, it is not. It assumes that there is no reason to think that we generate any silent readings in the case of reading novels silently. However, as shown in Chapter 3, this assumption is questionable.

⁷As might be recalled from Chapter 3, according to Kivy (2006), a reading of a novel is a particular kind of act—namely, a sounding of the text of the novel—that usually results in an artistic experience of this novel. Although Kivy (2006)'s concept of “a reading” is not the same as the concept that I use when formulating the ontology of instances of novels in the previous chapter, there is no principled difference between these concepts.

⁸Premise (8) is not explicitly stated by Kivy (2006, 2010). However, it (or perhaps some other premise equivalent to it) must be contained in the argument, since otherwise, this argument is invalid.

According to Kivy (2006, 2010), a reading is a sounding of a novel that usually results in an artistic experience of this novel. So we generate silent readings in the case of reading novels silently if in this case, (a) we generate silent soundings of these novels and (b) these soundings normally result in artistic experiences of these novels. In light of the psychological evidence that supports the idea that our silent reading involves the mental “voicing” of what we read,⁹ it is reasonable to maintain that in the process of reading novels silently, we do, in fact, generate silent soundings of these novels. Furthermore, there is no doubt that a silent sounding generated in the process of reading a novel silently normally results in an artistic experience of this novel. So, contrary to what the assumption being discussed states, there is a reason to hold that when we read novels silently, we do, in fact, generate readings.

A more serious objection to Kivy (2006, 2010)’s argument concerns premise (6) (“Readings through which we appreciate novels are artistic performances”). Perhaps readings through which we appreciate novels *in the case of reading aloud* are artistic performances. But what about readings through which we appreciate novels *in the case of reading silently*? Are these readings artistic performances?

Kivy (2006, 2010)’s argument in favor of answering this question in the affirmative is as follows. When we read novels silently, we read *while impersonating the storyteller*¹⁰ and, as a result, *con espressione*. Meanwhile, whenever one reads *con espressione*, the reading one generates possesses certain artistic properties (includ-

⁹See, e.g., Petkov and Berlin (2013).

¹⁰Kivy (2006, 2010)’s definition of “impersonation” is as follows: For all x and for all y , x impersonates y just in case x plays the part of y .

ing the property of being “expressive”) and, hence, is artistic. Thus, when we read novels silently, the readings we generate are artistic. Furthermore, these readings are performances. For to be a performance, it is sufficient to be an action that is intended to be presented to an audience. And our silent readings of novels doubtless satisfy this condition: They are actions—activities that are generated by agents (readers) and directed at a goal (to generate a silent reading)—and each of them is intended to be presented to an audience (though this audience consists of just one person—the reader herself). Thus, given what has been said, our silent readings of novels are artistic performances.

However, as shown in Chapter 3, the foregoing argument fails. It assumes that when we read a novel silently, we impersonate the storyteller—similar to how Ion the Rhapsode impersonates Homer when telling the stories of the *Iliad* and the *Odyssey*. But this assumption is false. When reading a novel silently, we apprehend the story told by this novel and sometimes imagine the characters and certain events of this story, but we do not impersonate the storyteller.

At the same time, as pointed out in Chapter 3, there are at least two strong reasons against treating silent readings of novels as artistic performances. First, an artistic performance comes into existence only if its creator has a particular intention—the intention to perform. But when we read a novel silently and thereby generate a reading, we do not have this intention. Second, according to a widely accepted view, to be artistic, an entity must be created with some art-relevant intention. However, our silent readings are not created with any such intention.

Thus, there is good reason against the thesis that readings we generate in

the process of reading novels silently are artistic performances.¹¹ Meanwhile, if this thesis is false, then premise (6) of Kivy (2006, 2010)'s argument is false—and so this argument fails.

5.2 The Novel as a Performing Art

Thus, neither Urmson (2004)'s nor Kivy (2006, 2010)'s argument stands up to criticism. This does not mean, however, that there is no satisfactory argument in favor of the view that the novel is a performing art. In fact, there is such an argument. And to this argument I now turn.

5.2.1 What Is a Performing Art?

Let us begin by clarifying the expression “a performing art.” According to Davies (2011), what distinguishes a performing art from other arts is that in the case of a performing art, “our access to, and appreciation of, *works* (as receivers) is at least in part mediated by performances of those works, and thus by the activities of those in the *performing arts* such as conductors, directors, musicians, dancers, and actors” (Davies, 2011, 18–19). Now, our access to, and appreciation of, *y* is mediated by *x* just in case to fully appreciate *y*, it is necessary to experientially engage with *x*. So Davies (2011)'s account of “a performing art” can be formulated as follows: What makes an art a *performing* art is the fact that to fully appreciate a

¹¹Note that the thesis that readings we generate in the process of reading novels silently *are not* artistic performances does not entail that readings generated in the process of reading novels silently *cannot* be such performances.

work of this art, it is necessary to experientially engage with a performance of this work. Or, in other words: For all x , x is a performing art just in case (a) x is an art and (b) a full appreciation of a work of x requires an experiential engagement with a performance of this work.

Is Davies (2011)'s account of "a performing art" satisfactory? One consequence of this account is that to fully appreciate a performable work, it is necessary to experientially engage with a performance of this work. This consequence, however, is problematic. Consider, for instance, the case of a musical work. Surely, by listening to a performance of such a work, one can (under appropriate conditions) fully appreciate this work. But is listening to a performance of a musical work *necessary* to fully appreciate this work? The answer to this question is "Yes" *only if* (a) to fully appreciate an artwork, one must experientially engage with an entity that manifests all the primary properties that must be experienced to fully appreciate this work and (b) in the case of musical works, the only entity that can manifest all such properties is a performance of a musical work. There is no doubt that (a) is true. But what about (b)? Is a performance of a musical work, in fact, the only entity that can manifest all the primary properties that must be experienced to fully appreciate this work?

Consider a playing of a recording of a musical work. Surely, not any such playing manifests all the primary properties that must be experienced to fully appreciate this work. Nevertheless, it is plausible to suppose that there are playings of recordings that manifest all such properties. Moreover, even if there are no such playings of recordings, they doubtless *can* (and, most likely, *will*) be created. Meanwhile, a

playing of a recording of a musical work is not a performance of this work; rather, it is a *representation* of such a performance.

The foregoing argument assumes that playings of recordings are not performances. According to Mag Uidhir (2007), however, this assumption is mistaken: In fact, under certain circumstances, playings of recordings *are* performances. To support his view, he advances the following argument. Suppose that a number of members of the audience of a live concert have been stricken by a strange condition, as a result of which “they are incapable of hearing the sounds produced by the orchestra, fellow audience members, or even themselves” (Mag Uidhir, 2007, 308) and “can only hear sounds produced within their own ear canals” (Mag Uidhir, 2007, 308). Suppose next that there is a hearing device such that “upon placement into the ear canals, [it] first records all incoming sounds and then plays the recording” (Mag Uidhir, 2007, 309).¹² Can the audience members, with the help of this device, hear the performance? Given the above assumptions, it is clear that they can hear only a playing of a recording of this performance. So if playings of recordings are representations of performances, and not themselves performances, then the answer to the foregoing question is “No.” But this answer seems wrong. It is natural to say that in the case being discussed, the audience members do hear the performance. So it seems that playings of recordings can be performances.

Mag Uidhir (2007)’s argument, however, is unpersuasive. According to this argument, the fact that it is natural to characterize the audience members as “hearing

¹²It is also assumed here that “the hearing device flawlessly preserves the wearer’s normal... coordination of aural input with visual input (and any other sort)” (Mag Uidhir, 2007, 309).

the performance” supports the claim that they, in fact, hear this performance. But this fact does not really support this claim. Suppose Mary puts on a virtual reality headset generating the visual experience indistinguishable from an experience one could get by actually looking at the Eiffel Tower, and says: “Look! It’s the Eiffel Tower.” Clearly, her phrase sounds quite natural. But, of course, this fact does not support the claim that what she sees is the actual Eiffel Tower (what she sees is a particular *representation* of this tower). Meanwhile, this case is relevantly analogous to the case described by Mag Uidhir (2007). Given this, in the latter case, the fact that it is natural to say that the audience members hear the performance does not support the claim that they actually hear this performance.

Thus, as shown above, besides a performance, there is at least one entity that can manifest all the primary properties that must be experienced to fully appreciate a musical work—a playing of a recording of this work. Are a performance and a playing of a recording the only entities that are capable of manifesting all such properties? No. There is, in fact, another such entity—a reproduction of a performance, or, in other words, a sequence of sounds that is generated, with the help of a score, by a computer or some other mechanical device. Such a reproduction can, under appropriate conditions, manifest all the primary properties that must be experienced to fully appreciate the corresponding work. At the same time, it is neither a performance nor a playing of a recording of a performance. (Here, one might ask: What is the difference between a reproduction of a performance and a playing of a recording of a performance? The key difference between them concerns the status of the sound. In the case of a playing of a recording, the sound is, in some

sense, “contained” in some medium for playback (say, an electronic file or a DVD) and, hence, exists before the playing is actualized. In the case of a reproduction, however, the sound does not exist before this reproduction is actualized; it comes into existence in the process of actualizing the reproduction.)

Thus, contrary to what Davies (2011)’s account of “a performing art” implies, to fully appreciate a musical work, it is not necessary to listen to a performance of this work. As a result, this account cannot be accepted.

But what is then an acceptable account of “a performing art”? To answer this question, we first need to answer the following question: What is it necessary to experientially engage with in order to fully appreciate a work of a performing art?

Consider the case of a musical work. A full appreciation of a musical work requires engaging experientially with (in particular, listening to) whatever manifests all the primary properties that must be experienced to fully appreciate this work. As mentioned above, the entities that are capable of manifesting all such properties include performances, playings of recordings of performances, and reproductions of performances. And there seem to be no other entities capable of that. So to fully appreciate a musical work, it is necessary to experientially engage with (in particular, to listen to) a performance, a playing of a recording of a performance, or a reproduction of a performance of this work.

Now, what has been said about musical works applies to works of other performing arts as well. As a result, the question posed above can be answered as follows: To fully appreciate a work of a performing art, it is necessary to experientially engage with a performance, a playing of a recording of a performance, or a

reproduction of a performance of this work.

Taking this into account, a satisfactory account of “a performing art” can now be formulated:

Performing art (PA): For all x , x is a performing art just in case (a) x is an art and (b) to fully appreciate a work of x , it is necessary to experientially engage with a performance, a playing of a recording of a performance, or a reproduction of a performance of this work.

Regarding the foregoing account (hereafter: “**PA**”), two remarks are worth making. First, this account covers the paradigmatic performing arts—namely, music, dance, and theater. To fully appreciate a musical work, it is necessary to listen to a musical performance, a reproduction of a musical performance, or a playing of a recording of a musical performance. A work of dance cannot be fully appreciated without seeing a dance performance, a reproduction of a dance performance, or a playing of a recording of a dance performance. Finally, a full appreciation of a work of theater is impossible without seeing and listening to a theatrical performance, a reproduction of a theatrical performance, or a playing of a recording of a theatrical performance. At the same time, there is no doubt that music, dance, and theater are arts. Meanwhile, according to **PA**, if x is an art and to fully appreciate a work of x , it is necessary to experientially engage with a performance, a playing of a recording of a performance, or a reproduction of a performance of this work, then x is a performing art.

Second, **PA** entails that those arts that are doubtless non-performing—painting,

drawing, sculpture, etching, and photography—are not performing arts. It is clear that a full appreciation of works of these arts—namely, a painting, a drawing, a sculpture, an etching, and a photograph—is possible without an experiential engagement with a performance, a playing of a recording of a performance, or a reproduction of a performance. Meanwhile, according to **PA**, an art is not performing if to fully appreciate a work of this art, it is not necessary to experientially engage with a performance, a playing of a recording of a performance, or a reproduction of a performance of this work.

One could object to **PA** as follows. Consider a film. The only way to properly appreciate it is by engaging experientially with—in particular, by watching and listening to—a screening of this film. Meanwhile, a screening of a film is a performance of this film. So if **PA** is true, then cinema is a performing art—which is, of course, false.

However, this objection fails. It assumes that a screening is a performance. But this assumption is false. Consider paradigmatic performances—a dance performance, a performance of a musical work, and a theatrical performance. Each of these performances is a result of the performer’s interpretation of the instructions provided by the author. And the same, I think, can be said about other kinds of performances. Thus, a performance is essentially interpretative. But a screening of a film is not interpretative. So such a screening cannot be a performance.

An opponent of **PA** could respond as follows. Let us agree that a screening of a film is not a performance. However, such a screening is a playing of a recording of a performance. Meanwhile, if that is the case, then **PA** is false. For suppose that

a screening of a film is a playing of a recording of a performance. Then **PA** entails that cinema is a performing art. However, as already mentioned, this consequence is doubtless false.

This response is based on the thesis that a screening of a film is a playing of a recording of a performance. But this thesis is false. If it were true, then a film would be a performable entity. However, a film is not such an entity; it is not something that is performed, even once. Surely, in the case of a film, we can speak of *the actors'* performance. But this kind of performance is not a performance *of the film*; rather, it is a performance *that is used in the creation of the film*.

Here, a natural question arises: If the relation of a screening of a film to this film is neither that of a performance of a work to this work nor that of a playing of a recording of a performance of a work to this work, then what exactly is this relation? Clearly, a screening of a film is a playing of some recording. Meanwhile, it seems right to consider this recording to be the film itself. Given this, the foregoing question can, I think, be answered as follows: The relation of a screening of a film to this film is that of a playing of a work to this work.¹³

¹³It is also worth noting that a screening of a film is an instance of this film. That a screening of a film is capable of being such an instance can be shown as follows. There is no doubt that a screening of a film can manifest all the primary properties that must be experienced to fully appreciate this film. So such a screening can be a well-formed instance_e of the film. Furthermore, a screening of a film can stand in an appropriate historical-intentional relation to this film. As a result, such a screening can be a well-formed instance_p of the film. Meanwhile, anything that is either well-formed instance_e or a well-formed instance_p of an artwork is an instance of this work. So a screening of a film can be an instance of this film.

5.2.2 What Are Performances of an Artwork?

Having clarified the expression “a performing art,” let us now show that the novel falls under this expression.

According to **PA**, the novel is a performing art if (a) it is an art and (b) to fully appreciate a novel, it is necessary to experientially engage with a performance, a playing of a recording of a performance, or a reproduction of a performance of this novel. There is no doubt that the novel is an art. So the novel is a performing art if to fully appreciate a novel, it is necessary to experientially engage with a performance, a playing of a recording of a performance, or a reproduction of a performance of this novel.

Now, as shown in the previous chapter, a novel cannot be fully appreciated without engaging experientially with its reading—a sequence of particular sounds generated as a result of reading aloud. So, given what has been said, if readings of novels are performances, reproductions of performances, or playings of recordings of performances, then the novel is a performing art. Are such readings, in fact, performances, reproductions of performances, or playings of recordings of performances?

To answer this question, we first need to clarify the expression “performances of an artwork.” What characteristic features does a performance of an artwork have? The first thing to note is that a performance of an artwork—as well as any other performance—is an event. As mentioned in the previous chapter,¹⁴ something is an event if it (a) is said to *occur*, or *happen*, or *take place*, (b) has relatively

¹⁴See Chapter 4, Section 4.2.

vague spatial boundaries and relatively crisp temporal boundaries, (c) tolerates co-location, (d) cannot move, and (e) takes up time and persists by perduring, that is, by having distinct temporal parts (or stages) at different times.¹⁵ Meanwhile, a performance of an artwork has all of these characteristics. It is said to *take place* somewhere (“A performance of Beethoven’s Symphony No. 5 is taking place at the John F. Kennedy Center for the Performing Arts”). Its temporal location (its beginning and end) can be identified quite accurately, while its spatial location can be identified only approximately. It appears to tolerate co-location—in particular, co-location with other events. It cannot move. Finally, it takes up time and persists by perduring.

Thus, a performance of an artwork is an event. Furthermore, it is a particular kind of event—namely, *an action*. Here is a possible definition of “an action”: For all x , x is an action just in case x is an event that is generated by an agent or agents and has a goal. This definition is rather intuitive. It entails that (a) events we would normally call “actions”—biking, talking, scratching, and the like—are, in fact, actions and (b) events we would not normally call “actions”—for example, hurricanes, floods, rains, and the movement of the sun—are not actions.¹⁶ Suppose, therefore, that the foregoing definition is correct. Then a performance of an artwork is an action if it is an event that is generated by an agent or agents and has a goal. Meanwhile, a performance of an artwork satisfies this condition. As has been shown earlier, it is an event. Furthermore, it is generated by an agent (the performer) and

¹⁵See Casati and Varzi (2015).

¹⁶For a detailed account of “an action,” see Wilson and Shpall (2012).

has a particular goal (to make it possible to appreciate the artwork).

Thus, a performance of an artwork is an action. But, of course, not every action is a performance of an artwork. So to be such a performance, an action must possess some additional properties. What might these properties be? To answer this question, let us consider a *typical* performance of an artwork. What characteristic features does such a performance have?

First, a typical performance of an artwork is usually generated with the help of a set of (written or spoken) instructions—such as a musical score, a theatrical script, and a score that describes a dance. This is not to say, of course, that the use of instructions is *necessary* to generate a performance. Surely, a performance can be generated without using any instructions whatsoever. However, common performing practices normally involve the use of certain instructions.

Second, a typical performance of an artwork is interpretation-driven, that is, it is, in part, a result of the performer's *interpreting* this work. By “interpreting an artwork” here is meant the process that consists of two consecutive activities. The first activity is that of comprehension: A performer understands the work—by apprehending instructions on how to perform it, its semantic meaning, potential intentions of the author, and any other relevant factors. The second activity is that of creative modification: Based on her understanding of the work, a performer makes modifications to the performance that are not specified in the work's instructions.¹⁷

¹⁷The process of a performer's interpreting a work can be illustrated with the help of the following example. Suppose John (an actor) interprets the following excerpt from Tennessee Williams's *A Streetcar Named Desire*:

Third, creating a typical performance of an artwork requires the application of special skills. Thus, in most cases, to create a musical performance, it is necessary to be able to play a musical instrument and/or read scores; to create a dance performance, it is necessary to be able to perform certain elaborate bodily movements; and to create a theatrical performance, it is necessary to be able to play the roles of various fictional characters. (This is not to say, of course, that a performance cannot be created by someone who lacks any special skills.¹⁸ Rather, the point is that creating a *typical* performance requires an exercise of certain special skills.)

Fourth, a typical performance of an artwork is intended to be presented before an audience. As Godlovitch (1998) puts it,

Performances are not reflective activities savoured by their agents in solitude. Performances... are other-directed, or, in the idiom, “given.” Unlike rehearsals, exploratory sight-seeing, recreational practice, and other

Blanche: What kind of bed’s this—one of those collapsible things?

In this case, John performs two actions. First, he comprehends Blanche’s phrase—by grasping its semantic meaning as well as various contextual factors that might be relevant to Blanche’s voice, gesticulation, and facial expression. Second, in his performance of the phrase, John adds certain features that are not specified in Williams’s play—such as making a particular facial expression when pronouncing this phrase (for example, smiling) or pronouncing this phrase in a particular way (for example, loudly or with a certain accent).

¹⁸Arguably, an example of a performance that does not require any special skills is Yvonne Rainer’s dance *Room Service*, in which “the dancers perform a series of ordinary movements that involve... the moving, arranging, and rearranging of objects such as mattresses and ladders” (Davies, 2011, 13).

player-centered activities, performances are specifically and directly intended, designed, or meant for audiences. As purposive activities, their *telos* is to be experienced by those for whom the performer prepares them. (Godlovitch, 1998, 28)¹⁹

Fifth, a typical performance of an artwork is presented before an audience. Here, one might ask: Must the audience before which a performance is presented be actually present? Or, in other words: Can there be a performance without an actual audience? According to Thom (1993), the answer is “No.” In his view, if x is presented before a non-actual audience, then x is not a performance. To support this view, he offers the following argument:

In performing I believe myself to be referring to present persons, to whom I am in effect saying, “You, attend to me.” If no one is present at the performance, there is a failure of reference... [Thus,] the audience is not a mere dispensable accessory to performance. (Thom, 1993, 192)

Alternatively, Thom (1993)’s argument can be presented as follows. x is a performance only if its author makes a particular “address” to an audience—namely, an “address” that has the form of the imperative “You, attend to me.” Meanwhile, such an “address” can be made only if it successfully refers to an audience. And this reference is possible only if the audience is actually present. Thus, without an actual audience, x is not a performance.

¹⁹Godlovitch speaks here of musical performances. But what he says, I think, can be said about all kinds of artistic performance.

Thom (1993)'s argument assumes that a successful reference of a performer's "address" to an audience is a necessary condition of being a performance: If x does not involve such a reference by a performer, then x cannot be a performance. But why think this assumption is true? Suppose a performer of a musical work *imagines* that he is making an "address" ("You, attend to me") to an actual audience. Suppose next that he is, in fact, alone. In this case, his "address" fails to refer to the audience. However—assuming that his performance corresponds to the score of the work and, hence, is correct—there seems to be no real reason to think that what he is doing is not a genuine performance.

Thus, Thom (1993)'s argument rests on a highly questionable assumption and, hence, cannot be accepted. At the same time, as far as I am aware, there is no other potentially acceptable argument in favor of the thesis that there cannot be a performance without an actual audience. Given this, as well as what has been said in the previous paragraph, the question "Can there be a performance without an actual audience?" should be answered in the affirmative.

Sixth, a typical performance of an artwork is usually not a mere means to appreciate the corresponding work; in many cases, it is something that is worth appreciating in its own right. Consider, for instance, a typical musical or theatrical performance. Such a performance usually is worthy of appreciation qua independent art object,²⁰ and is not just a means to appreciate the corresponding work.

Finally, by perceiving a typical performance of an artwork, we can appreciate this work. For example, by listening to a performance of Jean Sibelius's *Finlandia*,

²⁰Note that this does not imply that a performance is an artwork.

we can appreciate *Finlandia*. By watching and listening to a performance of *Hamlet*, we can appreciate *Hamlet*. And by watching and listening to a performance of *Swan Lake*, we can appreciate *Swan Lake*.²¹

Suppose now that an action has all of the foregoing features of a typical performance of an artwork. Is this action a performance? There can be no doubt about that. Thus, the properties that an action must possess to be a performance of an artwork can (though, perhaps, do not have to) be the properties of being:

- usually generated with the help of some set of (written or spoken) instructions
- interpretation-driven
- created by those who have special skills
- intended to be presented before an audience
- presented before an audience
- often worth appreciating in its own right
- such that by perceiving it, we can appreciate the corresponding artwork

In light of what has been said, the following account of “performances of an artwork” can be provided:

Performances of an Artwork (POA): For all x , x 's are performances of an artwork

A if x 's are actions that (a) are normally generated with the help of some sets

²¹It is worth noting that the claim that by perceiving a performance of an artwork, we can appreciate this artwork neither amounts to nor entails the claim that by perceiving a typical performance of an artwork, we can *fully* appreciate this work.

of (written or spoken) instructions, (b) are interpretation-driven, (c) can be created only by those who have special skills, (d) are intended to be presented before an audience, (e) are presented before an audience, (f) are, in many cases, worth appreciating in their own right, and (g) are such that by perceiving them, we can appreciate *A*.²²

One could object to the foregoing account (hereafter: “**POA**”) as follows. Consider live readings of a philosophy text (say, the *Critique of Pure Reason* or the *Nicomachean Ethics*). Such readings are actions that (a) are usually generated with the help of sets of (written or spoken) instructions (namely, particular inscriptions), (b) are interpretation-driven, (c) can be created only by those who can read and, hence, have special skills, (d) are intended to be presented before an audience, (e) are presented before an audience, (f) are often worth appreciating in their own right, and (g) enable us to properly appreciate the corresponding works. Meanwhile, according to **POA**, any actions that have the foregoing properties are performances. So if **POA** is true, then readings of a philosophy text are performances. But if that is so, then philosophy is a performing art. However, this consequence is doubtless false. So **POA** is false.

However, this objection fails. According to it, the fact that live readings of a philosophy text are performances of this text entails that philosophy is a performing art. But this entailment does not hold. According to **PA**, philosophy is a performing art only if (a) philosophy is an art and (b) to fully appreciate a philosophical text,

²²Note that this is not a *definition* of “performances of an artwork,” as it provides only a sufficient condition of being such performances.

it is necessary to engage experientially with a performance, a playing of a recording of a performance, or a reproduction of a performance of this text. However, even if live readings of a philosophy text are performances of this text, both (a) and (b) are false: Philosophy is not an art, and to fully appreciate a philosophical text, there is no need to engage experientially with a performance, a playing of a recording of a performance, or a reproduction of a performance of this text.

One could also argue against **POA** in the following way. Suppose **POA** is correct. Then a mother reading *Winnie the Pooh* to her four year old son is giving a performance. But this consequence is false.

This objection is also unsuccessful. It implies that it is wrong to treat the mother's reading as a performance. But this implication is ungrounded. There is no real reason against treating the mother's reading that way. One could say that such treatment is wrong, since the reading may not possess much aesthetic value—at least, when compared to readings of professional readers, such as Jim Dale, Garrison Keillor, and Penny Marshall. But this explanation can hardly be considered acceptable. It assumes that whether x is a performance depends on whether x has significant aesthetic value: If x has such value, then x can be a performance; otherwise, x is not a performance. But this assumption is false. An entity lacking any significant aesthetic value can still be a performance (a mediocre performance is still a performance!).

One could also try to explain the inappropriateness of regarding the mother's reading as "a performance" by saying that this reading has no aesthetic value. But this explanation is also unsatisfactory. It assumes that the mother's reading of

Winnie the Pooh has no aesthetic value. But this assumption is false. In fact, any intelligible reading of a literary work is aesthetically valuable at least to some degree. If a reading has no aesthetic value at all, then it just cannot be a reading of a literary work.

5.2.3 Readings as Performances

Taking into account **POA**, let us now address the question of whether readings of novels are performances of artworks.

Readings—that is, sequences of particular sounds generated as a result of reading aloud—can be divided into two categories: (1) readings that are directly generated by agents or agents and non-agents (for example, computers and stereo systems) and (2) readings that are directly generated only by non-agents. Call the first kind of readings “live readings” and the second kind “mechanical readings.”²³ Are readings of either of these kinds performances?

Let us first consider whether live readings of a novel are performances. According to **POA**, to be performances of a novel, it is sufficient for live readings to satisfy two conditions. First, live readings must be actions. Second, they must be (a) in most cases, generated with the help of some set of (written or spoken) instructions, (b) interpretation-driven, (c) capable of being created only by those who have special skills, (d) intended to be presented before an audience, (e) presented before an intended audience, (f) often worth appreciating in their own right, and

²³Note that the fact that a mechanical reading is not *directly* generated by an agent does not entail that such a reading is not an action.

(g) such that by perceiving them, we can appreciate the corresponding novels. Do live readings of a novel satisfy these conditions?

As already mentioned, an entity is an action if it is an event generated by an agent and is directed at a goal. According to the definition of “live readings,” such readings are generated by agents. Furthermore, live readings are events. As noted earlier, something is an event if it (a) is said to *occur*, or *happen*, or *take place*, (b) has relatively vague spatial boundaries and relatively crisp temporal boundaries, (c) can be spatially co-located with other events, (d) cannot move, and (e) takes up time and persists by perduring, that is, by having distinct temporal parts (or stages) at different times. And there is no doubt that live readings possess all of the foregoing properties.

Thus, live readings of a novel satisfy the first condition. Do they satisfy the second condition? There is good reason to answer this question in the affirmative. First, in an overwhelming majority of cases, a live reading of a novel is generated with the help of a (physical or mental) inscription of this novel. Meanwhile, an inscription of a novel is essentially a set of written instructions—the instructions on how to correctly read this novel. Thus, *live readings of a novel are usually generated with the help of instructions.*

Second, *live readings of a novel are interpretation-driven.* As is clear from what has been said in the previous section, to be interpretation-driven, x must satisfy two conditions. First, x 's creator must have an understanding of the work associated with x . Second, x 's creator must make creative additions to x —that is, additions that are not specified in the corresponding set of instructions and that

accord with the creator's conception of the work associated with x . In the case of live readings of a novel, both conditions are at least minimally satisfied. A reader of a novel doubtless understands what she reads. Furthermore, it is uncontroversial that a reader of a novel makes various creative additions to her reading—for example, by varying the speed and intonation of her reading (usually neither the speed nor intonation are specified by inscriptions) or by using different voices when reading the words of different characters (again, inscriptions do not normally provide any explicit description of what the characters' voices are).

Third, *live readings of a novel require an exercise of a special skill*—the skill of reading aloud. One could object that this skill is not really special, since virtually everyone has it. But this objection fails. First, although there are considerably more people who can read aloud than there are people who can, say, play guitar or dance salsa, not everyone can read aloud; in fact, there are a considerable number of people who cannot do it. Second, the fact that all or virtually all people can do x does not entail that x is not special. A skill is special just in case it is not acquired in a natural way. Thus, the skill of walking on two legs is not special, since acquiring it is a natural process. On the other hand, the skill of playing guitar *is* special—for, it has not been acquired as a result of a natural process. Now, what about the skill of reading? Is it special? To have a minimal capacity to read, one must have a number of skills. First of all, one must be capable of recognizing the letters of the alphabet. Also, one must know the sounds associated with these letters. Furthermore, one must know and be able to follow the rules on how to pronounce sounds of two or more concatenated letters. Finally, one must know punctuation and capitalization

conventions. Clearly, the mentioned skills cannot be acquired in a natural way—merely by following the instincts. Acquiring these skills is an artificial process—a process that requires one to go beyond nature. As a result, the question posed above should be answered in the affirmative: The skill of reading is special.

Fourth, in the process of reading a novel, a reader has the intention to present what he is reading to himself and/or to those who listen to his reading. Thus, *live readings of a novel are intended to be presented to an audience.*

Fifth, when reading a novel, a reader actually presents what he is reading to someone—either just to himself or, more often, to several people. So *live readings of a novel are presented before an audience.*

Sixth, *live readings of a novel are usually worth appreciating in their own right*, and not just as entities that make it possible to appreciate another artistic object (the novel). Of course, not every live reading possesses properties that are *high* in artistic value. It may even be right to say that most live readings do not possess such properties. But despite this, virtually every live reading has a sufficient number of artistic properties that are worthy of artistic appreciation. Hence, virtually every such reading is worth appreciating qua an independent artistic object.

Finally, by listening to live readings of a novel, we can appreciate certain sonic and semantic properties. Meanwhile, as has been established in the previous chapter, these properties provide direct experiential access to the sonic and semantic components of the novel and, hence, to the novel itself. Thus, *by listening to live readings of a novel, we can appreciate this novel.*

Thus, live readings of novels are performances. Now, what about mechanical

readings of novels—that is, readings that are directly generated solely by non-agents? Are such readings performances? It is reasonable to assume that to be a performance of an artwork, an entity must be, at least minimally, interpretation-driven. Are mechanical readings interpretation-driven? Consider playings of recordings. It is clear that generating them does not require an understanding of the associated works. Meanwhile, as mentioned above, to be interpretation-driven, an entity must be generated with an understanding of the work associated with this entity. Thus, playings of recordings are not interpretation-driven. What about reproductions of performances? Are they interpretation-driven? No. The argument here is analogous to the foregoing argument concerning playings of recordings. Thus, neither playings of recordings nor reproductions of performances are interpretation-driven. Meanwhile, there are no other kinds of mechanical readings. So no mechanical reading is interpretation-driven. But if this is so, then mechanical readings of novels are not performances.

5.2.4 Conclusion

As has been established above, the novel is a performing art if readings of a novel are performances, playings of recordings of performances, or reproductions of performances of this novel. Thus—given the fact that live readings are performances, whereas mechanical readings are not, as well as the fact that readings are either live or mechanical—if mechanical readings of novels are playings of recordings of performances or reproductions of performances, then the novel is a performing

art. As already mentioned, a mechanical reading is either a playing of a recording of a live reading or a reproduction of a live reading. So, given that live readings are performances, mechanical readings of novels are playings of recordings of performances or reproductions of performances. Consequently, the novel is a performing art.

Before proceeding, it is worth noting that the fact that the foregoing conclusion contradicts the consensus view on whether the novel is a performing art cannot be used to undermine this conclusion and, hence, the argument presented above. There are two ways x can undermine y —by (logically) entailing that y is problematic and by counting against y . Clearly, the fact that the conclusion of the argument presented in this section contradicts the consensus view on whether the novel is a performing art does not entail that this conclusion is problematic. Does this fact count against this conclusion? The answer to this question is “Yes” *only if* there is no plausible explanation of why most people are mistaken in believing that the novel is a non-performing art. However, there is such an explanation. What accounts for the fact that most people mistakenly believe that the novel is a non-performing art is their commitment to the following unsound argument:

10. For all x , x is a performing art if (a) x is an art and (b) to fully appreciate works of x , it is necessary to experientially engage with performances or some surrogates thereof of these works.
11. It is false that to fully appreciate novels, it is necessary to experientially engage with any performances or surrogates thereof.

12. So the novel is not a performing art.²⁴

5.3 The Novel as a Performing Art and the History of Literature

One might object that the view that the novel is a performing art does not accord with the history of literature, as this history is the history of a non-performing art. This objection, however, does not pose a real threat to this view. First of all, in the case of the arts, theory trumps practice. Thus, even if the history of literature contradicts the view that the novel is a performing art, this view can still be true—if it is supported by our best theoretical considerations regarding the novel. Is it, in fact, supported by these considerations? Given what has been said in this and the previous two chapters, the answer to this question is doubtless “Yes.”

Furthermore, there is good reason to think that the history of literature is, in fact, largely the history of a performing art. To see this, let us first have a look at the history of the reading practice.

In the ancient world, the practice of reading a literary text silently was uncommon;²⁵ the normal way to read such a text was to read it aloud, whether before an audience or alone.²⁶ This is not to say, of course, that at that time, *everyone* read aloud. Thus, St. Ambrose—one of the doctors of the Church, who lived in the fourth century—could read silently, as is evidenced by the following excerpt from St. Augustine’s *Confessions*:

²⁴Given the above analysis, it is clear that the reason why this argument is unsound is that premise (11) is false.

²⁵See Balogh (1927), Manguel (2014), Nietzsche (2002), Saenger (1997), and Thomas (1992).

²⁶See Manguel (2014).

When [Ambrose] read, his eyes scanned the page and his heart sought out the meaning, but his voice was silent and his tongue was still. Anyone could approach him freely and guests were not commonly announced, so that often, when we came to visit him, we found him reading like this in silence, for he never read aloud. (Finkelstein and McCleery, 2013, 38)

However, Ambrose was the exception, rather than the rule.^{27,28,29}

Thus, in the ancient world, it was uncommon to read silently. Then the

²⁷See Manguel (2014) and Saenger (1997).

²⁸Interestingly, one piece of evidence supporting this thesis is contained in the foregoing quote. As is clear from this quote, Augustine finds Ambrose's manner of reading rather odd. (For, otherwise, why would Augustine stress the fact that Ambrose was silent when reading and never read aloud?) Meanwhile, the fact that Augustine finds Ambrose's manner of reading odd is a clear indication of the fact that silent reading was considered an oddity, rather than the norm, at that time.

²⁹The fact that the practice of silent reading was so uncommon in the ancient world is doubtless puzzling. Why didn't the ancients (at least, the Greeks) read silently? In the literature on the history of reading, there are two main answers to this question. According to one answer, endorsed by Knox (1968), among others, the ancients did not do that because at that time, it was generally agreed that a proper appreciation of a literary text required appreciating the sonic properties of this work through listening to the work's sound. The second answer, defended, most notably, by Saenger (1997), is that the ancients did not read silently because by reading silently, it was impossible for them to understand what they read. (Regarding the latter answer, a natural question arises: Why couldn't they understand what they read by reading silently? A possible answer to this question is as follows. The ancients read texts "written in what is known to linguists and paleographers as *scriptura continua*, which is to say, uninterrupted writing" (Kivy, 2006, 16), or, other words, texts, there was no separation between words. Meanwhile, as Saenger (1997) points out, comprehending such texts is extremely hard without hearing them and, hence, through silent

situation began to change:

During the Middle Ages, one group of readers after another mastered the technique of silent reading. The first were the copyists working in the monastic scriptoria. Then, around the middle of the twelfth century, scholars in the universities acquired the ability. Two centuries later the lay aristocracy learned to read silently. By the fifteenth century silent reading was the norm, at least for readers who also knew how to write and who belonged to segments of society that had long been literate. For others, who belonged to groups that slowly learned to read and for whom books remained strange, rare objects, the old way of reading no doubt remained a necessity. As late as the nineteenth century, neophytes and maladroit readers could be identified by their inability to read silently. In Labiche's play *La Cagnotte* (1864), the farmer Colladan replies to a person who loses patience when he reads a very private letter out loud: "If I read out loud, it's not for you, it's for me... Whenever I don't read out loud... I don't understand what I am reading." (Chartier, 2002, 125–126)

As is clear from the foregoing description, the proliferation of silent reading was far from swift. In fact, it was not until the early twentieth century that silent reading became the usual way to read for nearly everyone. Even in the nineteenth century, there was a considerable number of people who read aloud, rather than

reading.)

silently, to themselves on a regular basis.

So silent reading was virtually nonexistent for a long time and began to slowly gain popularity only in the Middle Ages; reading aloud, on the other hand, was the dominant reading practice before the Middle Ages and then existed, as a popular alternative to silent reading, for a continuous period—before eventually going out of fashion in the twentieth century.

Thus, the history of literature is largely (though, of course, not solely) the history of literature that is read aloud. But literature that is read aloud is literature that is appreciated through live readings. And, as has been shown in Section 5.2, the latter are performances. So the history of literature is, to a large extent, the history of a performing art.

PART III

Introduction to PART III

As mentioned in the introduction to PART II, to answer the main question of this dissertation, it is sufficient to answer the following questions: “What is the ontological status of novels?” and “What is the ontological status of instances of novels?” The answer to the latter question has been provided in PART II. The purpose of PART III is to answer the question “What is the ontological status of novels?”

Chapter 6: What a Novel Is Not

Introduction

Possible views on the fundamental nature of a novel include:

1. the view that a novel is a set of embodiments of this novel¹ (e.g., a set of inscriptions/readings/inscriptions *and* readings/certain mental states appropriately related to this novel)
2. the view that a novel is a property (e.g., the property of having a particular syntactic/ semantic/syntactic-semantic/sonic structure)
3. the view that a novel is a pure type² (e.g., a pure abstract syntactic/semantic/syntactic-semantic/sonic structure)
4. the view that a novel is an initiated type³ (e.g., a pure abstract syntactic/semantic/syntactic-semantic/sonic structure indicated⁴ by the author in a particular cultural con-

¹By “an embodiment of a novel” is meant a concrete singular entity that possesses all, or at least sufficiently many, of the relevant artistic properties of this novel. Entities that can serve as embodiments of novels include inscriptions, readings, electronic files, and/or mental states appropriately related to this novel.

²The expression “a pure type” is clarified in Section 6.1.3.

³The expression “an initiated type” is clarified in Section 6.1.4.

⁴The expression “to indicate” is clarified in Section 6.1.4.

text)

5. the view that a novel is a “historical individual”⁵ (e.g., a syntactic/semantic/syntactic-semantic/sonic abstractum that depends for its essence and existence on certain concreta)
6. the view that a novel is an “abstract artifact”⁶ (e.g., a syntactic/semantic/syntactic-semantic/sonic abstractum that depends for its essence and existence on certain concreta)
7. the view that a novel is an embodiment of this novel (e.g., the original manuscript)
8. the view that a novel is a mereological sum of embodiments of this novel (e.g., a mereological sum of inscriptions/readings/inscriptions *and* readings/certain mental states appropriately related to this novel)
9. eliminativism, or the view that a novel does not exist (and so does not have any nature)

Views (1)–(6) are *abstractionist*, or ones according to which a novel is *an abstractum* (where an abstractum can roughly be characterized as an entity that has at least some of the following properties: (a) being nonspatial, (b) being atemporal, (c) being causally impassive (= being incapable of being causally affected by anything), (d) being causally inactive (= being incapable of causally affecting anything), and

⁵The expression “a historical individual” is clarified in Section 6.1.5.

⁶The expression “an abstract artifact” is clarified in Section 6.1.5.

(e) being modally inflexible (= having all intrinsic properties⁷ essentially⁸). Views (7) and (8) are *concretist*, or ones according to which a novel is *a concretum* (where a concretum can roughly be characterized as whatever has some or all of the following properties: (a) being in space and/or time, (b) being causally efficacious (= being capable of causally affecting something and of being causally affected by something), and (c) being modally flexible (= having some intrinsic properties nonessentially)).⁹ The last view, eliminativism, is *neither concretist nor abstractionist* (for, (a) it implies that novels do not have any nature at all, whereas (b) according to both concretist and abstractionist views, novels have some—abstract or concrete—nature).

My goal in this chapter is to show that none of the views (1)–(9) stands up to criticism. I begin with a critique of the abstractionist views (Section 6.1). Next, I argue against the concretist views (Section 6.2). Then I show that eliminativism is untenable (Section 6.3). Finally, I make a remark concerning the view that, I believe, is successful (this view is discussed in the next chapter) (Section 6.4).

⁷Following Lewis (1983), “an intrinsic property” can be defined as follows: For all x and for all y , x is an intrinsic property of y just in case x is a property possessed by y “in virtue of the way [y] itself, and nothing else, is” (Lewis, 1983, 112). For a detailed analysis of the expression “an intrinsic property” (as well as the related expression “an extrinsic property”), see Weatherson and Marshall (2014).

⁸“An essential property” can be defined as follows: For all x and for all y , x is an essential property of y just in case if x is a property of y and if y is deprived of x , then y goes out of existence.

⁹For a detailed analysis of the expressions “abstract” and “concrete,” see Burgess (1997), Cowling (2017), Hale (1988), Rosen (2012), and Yablo (2002).

6.1 Abstractionist Views

Let us begin with an examination of the abstractionist views.

6.1.1 A Novel as a Set of Embodiments

Consider first the view that a novel is *a set of embodiments of this novel*.¹⁰ One objection to this view is as follows. Sets have their members *essentially*. In other words, (a) “no set could *gain or lose* any members (or exchange some members for others) and still be the set it was” (Van Cleve, 1985, 585) and (b) “no set *could have had* members other than the ones it does” (Van Cleve, 1985, 585).¹¹ So if a novel is a set of embodiments of this novel, then the identity of a novel is determined by the embodiments contained in this set. But is the identity of a novel, in fact, determined by that? If that is the case, then whenever one of the embodiments of a novel is modified or destroyed, the novel is also modified or destroyed, respectively. However, this result is unacceptable. If I modify/destroy one

¹⁰One might object that since the existence of a set of embodiments of a novel implies that this novel must exist *independently* of this set, the view that a novel is a set of embodiments of this novel is incoherent. However, this objection is ungrounded. It assumes that the existence of a set of embodiments of a novel implies that this novel must exist *independently* of this set. But this assumption is false. Clearly, a set of concreta that manifest or encode particular textual properties (and perhaps are appropriately related to the author) can exist even if the corresponding novel exists only qua an entity identical to this set. Meanwhile, a set of embodiments of a novel is just a set of such concreta.

¹¹For an explanation of why this is so, see Van Cleve (1985).

of the existing embodiments of (say) *Moby-Dick*, I will not thereby modify/destroy *Moby-Dick*. Thus, the answer to the foregoing question is “No.” But then the view that a novel is a set of embodiments of this novel cannot be true.

Here is another objection to this view. Suppose a novel is a set of embodiments of this novel. Then a novel is either a set that contains more than one embodiment of this novel or a set that contains just one embodiment of this novel. Suppose a novel is a set that contains more than one embodiment (perhaps *all* embodiments) of this novel. Then to fully apprehend a novel, it is necessary to apprehend *all* of the embodiments contained in the set with which this novel is identified. But this consequence is false. A complete apprehension of a novel requires apprehending *only one* of its embodiments (assuming, of course, that this embodiment is of satisfactory quality).

Suppose, on the other hand, that a novel is a set that contains just one embodiment of this novel. Then to apprehend a novel, it is necessary to apprehend the embodiment that is contained in the set with which this novel is identified. But this consequence is doubtless false. A novel can be apprehended not only by apprehending the embodiment that is contained in the set with which this novel is identified but also by apprehending any other embodiment of this novel (assuming, of course, that this latter embodiment is of satisfactory quality).

Thus, a novel is neither a set that contains just one embodiment of this novel nor a set that contains more than one embodiment of this novel. But if this is so, then the view that a novel is a set of embodiments of this novel must be false.

Here, a proponent of this view could respond as follows. The foregoing ob-

jection is based on the assumption that if a novel is a set that contains more than one embodiment of this novel, then a complete apprehension of a novel requires apprehending *all* of the embodiments contained in the set with which the novel is identified. But is this assumption true? Consider the following principle: For all x and for all y , if y is appropriately related to x , then it is possible to fully apprehend x by apprehending y . This principle looks uncontroversial. Suppose, therefore, that it is true. Then if an embodiment contained in the set with which a novel is identified is appropriately related to this novel, it is possible to fully apprehend a novel by apprehending *just one* of the embodiments contained in the set with which this novel is identified. Meanwhile, any embodiment of a novel (and, hence, any embodiment contained in the set with which a novel is identified) is appropriately related to this novel. As a result, the answer to the question posed above is “No.” But if this is so, then the objection being discussed involves a false assumption and, hence, fails.¹²

The foregoing response states that an embodiment of a novel is appropriately related to this novel. But what exactly does it mean for an embodiment of a novel to be *appropriately related* to this novel?

Prima facie, this question can be answered as follows: An embodiment of a novel is *appropriately related* to this novel just in case it provides all information relevant to the artistic appreciation of this novel. Suppose this answer is correct. Then the response being discussed implies that an embodiment of a novel provides all information relevant to the artistic appreciation of this novel. However, this

¹²This response draws upon Dodd (2000, 2004)’s response to a similar objection. For an analysis of that objection and Dodd (2000, 2004)’s response, see Caplan and Matheson (2006).

implication is incompatible with the view that a novel is a set of embodiments of this novel. For suppose this view is true. Then, since all information relevant to the artistic appreciation of a novel includes information about the ontological composition of this novel, all information relevant to the artistic appreciation of a novel includes the information that the novel is a set composed of embodiments of this novel. However, an embodiment of a novel does not normally contain the information that the novel is such a set. So such an embodiment does not normally provide all information relevant to the artistic appreciation of the corresponding novel. As a result, if the view that a novel is a set of embodiments of this novel is true, then the implication under consideration—that an embodiment of a novel provides all information relevant to the artistic appreciation of this novel—is false.

Thus, if the foregoing answer to the question “What exactly does it mean for an embodiment of a novel to be appropriately related to this novel?” is true, then the response being discussed is not available to a proponent of the view that a novel is a set of embodiments of this novel. Meanwhile, there is no satisfactory alternative to this answer. Given what has been said, as well as the fact that the response being discussed cannot be successful without a satisfactory answer to the foregoing question, this response fails.

Can any of the foregoing objections to the view that a novel is a set of embodiments of this novel be plausibly countered by a proponent of this view? There seems no reason to think so. Given this, the view that a novel is a set of embodiments of this novel should be rejected.

6.1.2 A Novel as a Property

Let us now turn to the view that a novel is *a property*. This view faces the following objection. If x is a property, then x cannot have any of the following properties:

- (a) having a particular number of words
- (b) telling a particular story
- (c) having particular sonic properties
- (d) being in English

So if the view being discussed is true, then a novel does not have any of the properties (a)–(d). But this consequence is doubtless false. Take, for instance, *Pride and Prejudice*. It has all of (a)–(d): It has a particular number of words, tells a particular story, has certain sonic properties, and is in English.

Can a proponent of identifying a novel with a property plausibly respond to this objection? No—or so it seems. Given this, the view that a novel is a property cannot be accepted.

6.1.3 A Novel as a Pure Type

Consider now the view that a novel is a *pure type*, or, in other words, an abstract type^{13,14} that is eternal (= has always existed and will always exist) and acausal (= cannot stand in causal relations).¹⁵ This view faces the following objection. If a novel is a pure type, then, since such a type is, by definition, eternal, a novel is eternal. However, this consequence is false, for (a) if a novel is eternal, then it cannot be created, it can only be discovered; however, (b) novels are not discovered, they are created. Thus, the view that a novel is a pure type is false. Call this “the Creation Objection_{pt}.”¹⁶

¹³The use of the word “abstract” here may seem redundant, but, in fact, it is not. Perhaps it would be redundant if *all* types were, by definition, abstract. However, according to my terminology, besides *abstract* types, there are also *concrete* types. (For a clarification of the expression “a concrete type” as well as a defense of the view that some entities are best viewed as concrete types, see Chapter 7.)

¹⁴For an account of what a type is, see Rohrbaugh (2003, 179–181) and Wetzel (2009, 2014).

¹⁵Proponents of this view include Strawson (2003) and Currie (1989). Also, a parallel view on the fundamental nature of musical works is endorsed by Dodd (2000, 2002, 2007) and Kivy (1983, 1987).

¹⁶Another common objection to the view that a novel is a pure type stems from the idea that novels are *modally flexible* (where “modally flexible” can be defined as follows: For all x , x is modally flexible just in case x does not have all of its intrinsic properties essentially).^{*} Consider some novel—say, *Don Quixote*. Had Cervantes’s decisions been different, *Don Quixote* could have been shorter or contained slightly different expressions. Meanwhile, the properties of being shorter and containing slightly different expressions are intrinsic.^{**} So at least some novels could have intrinsic properties other than the ones they actually have. But if this is so, then at least some novels are modally flexible.^{***} Meanwhile, pure types have *all* of their intrinsic properties

Response to the Creation Objection_{pt} (1). A proponent of the view that a novel is a pure type could respond to the Creation Objection_{pt} as follows. This objection implies that if a novel is eternal, then it cannot be created. But this essentially and, hence, are not modally flexible. So the view that a novel is a pure type is false.

Yet, unlike the Creation Objection_{pt}, this objection—which might be called “the Objection from Modal Flexibility”—can, I think, be defused. Here is a plausible response available to a proponent of the view that a novel is a pure type. The Objection from Modal Flexibility assumes that novels are modally flexible. If this assumption is true, then some of the intrinsic properties of a novel can be changed without thereby causing it to go out of existence. Meanwhile, *prima facie*, the intrinsic properties of a novel amount to certain semantic, sonic, syntactic, and, in some cases, visual properties. Thus, the truth of the foregoing assumption entails that at least some of these properties can be changed without thereby causing the novel to go out of existence. Is this entailment true? No. The semantic, sonic, syntactic, and, in some cases, visual properties are essential to a novel. Meanwhile, if this is so, then any change in these properties—regardless of how minimal this change is—causes the novel to go out of existence. Thus, novels are not modally flexible—and so the Objection from Modal Flexibility fails.

Here, one could object as follows. Consider those syntactic properties that do not have any effect on the meaning, sound, artistically relevant graphic elements, or artistic context of a novel (for instance, the properties of having certain commas, periods, and quotation marks). These properties can be possessed by this novel. Furthermore, they are intrinsic. But changing them does not cause the novel to go out of existence. So, contrary to what the foregoing response states, a novel *can* be intrinsically changed without being destroyed.

In response to this, a proponent of the view that a novel is a pure type could say the following. The foregoing objection assumes that those syntactic properties that have no effect on the meaning, sound, artistically relevant graphic elements, and artistic context of a novel can be possessed by this novel. But is this assumption true? The artistic value of a novel does not depend on any property that has no effect on the meaning, sound, artistically relevant graphic elements, or artistic

implication is false. Suppose a novel is eternal. Then it can be a result of the process of discovering a particular linguistic structure. Meanwhile, this process is creative, since (a) it requires a considerable amount of non-standard, highly original thinking, and (b) any such process is creative. So a novel can be a result of a creative process. But if x is a result of such a process, then x is created. So, contrary to what the Creation Objection_{pt} implies, a novel can be created even if it is eternal.¹⁷

This response is based on a particular account of creation—namely, the account context of this novel. So those syntactic properties that have no effect on the meaning, sound, artistically relevant graphic elements, and artistic context of a novel are not directly relevant to the artistic value of this novel. Meanwhile, if x is an intrinsic property, then x is possessed by an artwork only if x is directly relevant to the artistic value of this work. Thus, those syntactic properties that have no effect on the meaning, sound, artistically relevant graphic elements, or artistic context of a novel cannot be possessed by this novel. As a result, the objection being discussed involves a false assumption and, hence, fails.

* The view that novels are modally flexible has been defended by Hazlett (2012) and Rohrbaugh (2003).

** The reason for this is as follows. The properties of being shorter and containing slightly different expressions are possessed by novels by virtue of the way these novels themselves are, not by virtue of something else. Meanwhile, as already mentioned (see Footnote 7), a property is intrinsic if it is possessed by a thing by “virtue of the way that thing itself, and nothing else, is” (Lewis, 1983, 112).

*** This is not to say, of course, that the modal flexibility of a novel is unrestricted. (Take, once again, *Don Quixote*. It could not have the property of describing 2017 England or portraying Alonso Quixano as one who does not know anything about chivalry.) The point is that novels have *at least some* modal flexibility with regard to their intrinsic properties.

¹⁷Cf. Fisher (1991, 130).

according to which being a result of a creative process is sufficient for being created. Is this account true? Consider the mass–energy equivalence formula $E = mc^2$. It was not created (rather, it was discovered). But it is a result of a creative process—the process that required a considerable amount of non-standard, highly original thinking. So the foregoing account is false. But then the response being discussed is unsatisfactory.

Response to the Creation Objection_{pt} (2). Here is another possible response available to a proponent of the view that a novel is a pure type. According to the account of creation advanced by Deutsch (1991), “to create is to freely stipulate” (Deutsch, 1991, 220), or, in other words, to provide a description that is invariably correct, whatever its content. Clearly, if this account is true, then even an eternally existing novel can be created—by means of freely stipulating the textual content of this novel. Is this account, in fact, true? There seems no reason to think otherwise. Thus, a novel can be created even if it is eternal. But if this is so, then one of the assumptions of the Creation Objection_{pt}—that if a novel is eternal, then this novel cannot be created—is false, and, hence, this objection fails.

This response is also unsatisfactory. It is based on Deutsch (1991)’s account of creation. But this account is misguided. First of all, it entails that a novel can be created (by stipulating particular textual content) *even if this novel already exists*. But this entailment is problematic. Creation implies coming into existence. Meanwhile, nothing that already exists can come into existence. So a novel that already exists cannot be created.

Furthermore, if Deutsch (1991)’s account is true, then *all* instances of creation

result from stipulation. However, this consequence is false. Consider, for instance, painted canvases or copies of musical scores. They are doubtless instances of creation. But they are not results of stipulation.¹⁸

Response to the Creation Objection_{pt} (3). Following Currie (1989),¹⁹ a proponent of the view that a novel is a pure type could also respond to the Creation Objection_{pt} as follows. Imagine there are two planets—Earth and Twin Earth. Imagine next that the cultural history of Twin Earth is exactly the same as the cultural history of Earth, except that the culturally relevant events that take place on Twin Earth occur later than the corresponding culturally relevant events that take place on Earth. Consider now Melville living on Earth (hereafter: “Melville”) and Melville living on Twin Earth (hereafter: “Twin Melville”). Did Twin Melville create *Moby-Dick*? According to the foregoing assumptions, his compositional activity followed Melville’s. But “a work cannot be created and then created again at a later time” (Currie, 1989, 61). Thus, Twin Melville did not create *Moby-Dick*. Now, what about Melville? Did *he* create it? By assumption, the cultural histories of Earth and Twin Earth are identical. So Melville’s compositional activity is the same as Twin Melville’s compositional activity. But if this is the case, then, given that Twin Melville did not create *Moby-Dick*, Melville did not create it either. Thus, writing a novel does not presuppose creating this novel. But if this is so, then there is good reason to hold that novels are not created.²⁰ Meanwhile, if novels are not

¹⁸See Predelli (2001, 287).

¹⁹See Currie (1989, 61–64).

²⁰The original version of this argument—the one advanced by Currie (1989)—is concerned with music. In a nutshell, this version is as follows. Imagine there are two planets—Earth and Twin

created, then one of the premises of the Creation Objection_{pt} is false—and, hence, this objection fails.

According to this response, the thesis that (a) Twin Melville did not create *Moby-Dick* and (b) Melville’s compositional activity is the same as Twin Melville’s compositional activity entails that *Moby-Dick* was not created by Melville. Is this entailment valid? Suppose Twin Melville did not create *Moby-Dick*. Suppose next that Melville’s compositional activity is the same as Twin Melville’s compositional activity. Why can’t it then be true that, since Melville’s compositional activity occurred before Twin Melville’s compositional activity, Melville created *Moby-Dick*, whereas Twin Melville merely discovered it? Meanwhile, if this can be true, then the foregoing entailment is invalid—and, hence, the response being discussed is Earth. Imagine next that the cultural history of Twin Earth is exactly the same as the cultural history of Earth, except that the culturally relevant events that take place on Twin Earth occur later than the corresponding culturally relevant events that take place on Earth. Consider now Beethoven living on Earth (hereafter: “Beethoven”) and Beethoven living on Twin Earth (hereafter: “Twin Beethoven”). Did Twin Beethoven create the *Hammerklavier*? According to the foregoing assumptions, his compositional activity followed Beethoven’s. But “a work cannot be created and then created again at a later time” (Currie, 1989, 61). Thus, the answer to the foregoing question is “No”: Twin Beethoven did not create the *Hammerklavier*. Now, what about Beethoven? Did *he* create the *Hammerklavier*? By assumption, the cultural histories of Earth and Twin Earth are identical. So Beethoven’s compositional activity is the same as Twin Beethoven’s compositional activity. Meanwhile, if this is so, then, given that Twin Beethoven did not create the *Hammerklavier*, Beethoven did not create it either. Thus, composing a musical work does not presuppose creating this work. But if this is so, then there is good reason to hold that musical works are not created. (See Currie (1989, 61–64).)

unsatisfactory.

Here, a proponent of this response could object as follows. The foregoing argument assumes that Twin Melville discovered *Moby-Dick*. But this assumption is false. If it were true, then Twin Melville's writing *Moby-Dick* would presuppose discovering this novel. However, in fact, discovering a novel is not involved in writing this novel.²¹

This objection commits a proponent of the foregoing response to the Creation Objection_{pt} to the view that at least some artworks are neither created nor discovered. Is this view acceptable? The most intuitive view is that all artworks are created. If, for some reason, this view is false, then the only satisfactory option left is that some artworks (for example, paintings and non-cast sculptures) are created, whereas other artworks (such as musical works and works of literature) are discovered. Given this, there is a strong prima facie reason against the view that at least some artworks are neither created nor discovered. At the same time, there seems no good reason in favor of this view.²² Taking this into account, the objection being discussed seems untenable.

Here is another problem facing the foregoing response to the Creation Objection_{pt}. According to this response, Melville's *Moby-Dick* and Twin Melville's *Moby-Dick* are numerically identical (hereafter: "identical"). But why think that they are, in fact,

²¹Cf. Currie (1989, 63).

²²A potential reason in favor of the view that artworks are neither created nor discovered has been offered by Currie (1989). Yet this reason is unsustainable. It is based on the view that artworks are action types, but, as shown by Levinson (1992), this view is seriously problematic.

identical?

One possible answer is as follows: Melville's *Moby-Dick* and Twin Melville's *Moby-Dick* are particular pure types, and these types are identical. However, this answer is question-begging, since it presupposes the truth of the view being defended—that novels are pure types.

Another possible answer is this: All identity-relevant factors (including contextual ones) concerning Melville's *Moby-Dick* and Twin Melville's *Moby-Dick* are the same; meanwhile, if all identity-relevant factors with regard to some artworks are the same, then these artworks are identical. However, this answer is based on a false thesis—that all identity-relevant factors concerning Melville's *Moby-Dick* and Twin Melville's *Moby-Dick* are the same. For suppose this thesis is true. Then, since Melville and Twin Melville belong to different (though qualitatively identical) worlds and, hence, different cultures—those of Earth and Twin Earth, respectively—the factor of belonging to a particular culture must be identity-irrelevant. However, as demonstrated by Levinson (1992, 2007, 2012), there is good reason to hold that this factor is actually identity-relevant.

Thus, neither answer is satisfactory. Meanwhile, there is no other potentially satisfactory answer. So there is no reason to treat Melville's *Moby-Dick* and Twin Melville's *Moby-Dick* as one and the same novel. Moreover, there are a number of reasons to consider them nonidentical. Here is one such reason. Given the fact that Melville wrote *Moby-Dick* on Earth, it seems odd to say that his *Moby-Dick* has no location at all or that it is located not only on Earth but at any other place, including Twin Earth. From an intuitive viewpoint, Melville's *Moby-Dick* has a

rather specific location—it is located on Earth. Likewise, since Twin Melville wrote *Moby-Dick* on Twin Earth, it does not seem right to say that his *Moby-Dick* is located everywhere or that it does not have any location at all. Intuitively, like the location of Melville’s *Moby-Dick*, the location of Twin Melville’s *Moby-Dick* is rather specific: Twin Melville’s *Moby-Dick* is located on Twin Earth. Thus, our intuitions support the thesis that (a) Melville’s *Moby-Dick* is located on Earth, whereas (b) Twin Melville’s *Moby-Dick* is located on Twin Earth.²³ Meanwhile, if this thesis is true, then Melville’s and Twin Melville’s *Moby-Dick*’s differ in their essential properties, and, as a result, cannot be identical.

Here is another reason against identifying Melville’s *Moby-Dick* and Twin Melville’s *Moby-Dick*. Melville doubtless intended the names of various places—“the Pacific Ocean,” “Cape Horn,” “New Bedford,” etc.—involved in his *Moby-Dick* to refer to the corresponding places on Earth, not on Twin Earth. But, surely, Twin Melville did not have this intention. He did not intend “the Pacific Ocean” to refer to an *earthly* place (namely, the Pacific Ocean located on Earth). Nor did he intend “Cape Horn,” “New Bedford,” or any other geographical name contained in his *Moby-Dick* to refer to any such place. His intention was to refer to certain places *on Twin Earth*.²⁴ Thus, given what has been said, the geographical references in Melville’s *Moby-Dick* are essentially different from the geographical references in

²³Cf. Burgess (1997, 21–22).

²⁴Clearly, what has been said here about Melville’s and Twin Melville’s intentions with regard to the names of places contained in *Moby-Dick* can be said about their intentions with regard to the names of certain other things (e.g., people, animals, and practices) contained in this novel.

Twin Melville's *Moby-Dick*: The former are concerned with Earth, whereas the latter are concerned with Twin Earth. But if this is so, then, given that the references contained in a novel determine its semantic content and, hence, are essential to it, Melville's *Moby-Dick* cannot be identified with Twin Melville's *Moby-Dick*.

Yet another reason against identifying Melville's *Moby-Dick* and Twin Melville's *Moby-Dick* is concerned with *potential* artistically relevant differences in the properties of Melville's *Moby-Dick* and Twin Melville's *Moby-Dick*.²⁵ Clearly, even if *Moby-Dick* is assumed to have been regarded in the same way by the peoples of Earth and Twin Earth, it might still be regarded differently by these peoples *in the future*. In particular, it is possible that at some future time *t*, *Moby-Dick* is regarded as (a) a highly *original* literary work by those who live on Earth and (b) an *unoriginal* literary work by those who live on Twin Earth. Suppose now that at *t*, *Moby-Dick* is regarded that way. Then at *t*, if Melville's *Moby-Dick* = Twin Melville's *Moby-Dick*, *Moby-Dick* is not highly original *simpliciter* (rather, it is highly original *according to the culture of Earth* and is not highly original *according to the culture of Twin Earth*). Thus, if Melville's *Moby-Dick* = Twin Melville's *Moby-Dick*, then it is possible for *Moby-Dick* not to be highly original *simpliciter* (and perhaps to be unoriginal). However, this consequence seems wrong. *Prima facie*, being highly original *simpliciter* is an *essential* property of *Moby-Dick*, and so it is impossible for *Moby-Dick* not to be highly original *simpliciter*.²⁶ Given this, there is good reason to hold that Melville's *Moby-Dick* and Twin Melville's *Moby-Dick*

²⁵Cf. Levinson (1992, 221–222).

²⁶Similar arguments are given by Levinson (1980, 10–13).

are different novels.

Thus, all things considered, Melville's *Moby-Dick* and Twin Melville's *Moby-Dick* are different novels. Meanwhile, as already mentioned, the foregoing response to the Creation Objection_{pt} implies that they are identical. So this response is unsatisfactory.

Response to the Creation Objection_{pt} (4). A proponent of the view that a novel is a pure type could also respond to the Creation Objection_{pt} in the following way. This objection assumes that novels are created, and not discovered. But why think that this assumption is true? Until a proponent of the Creation Objection_{pt} provides a satisfactory answer to this question, this objection does not pose any real threat to the view that a novel is a pure type.

The foregoing response is successful only if the thesis that novels are created, and not discovered, cannot be substantiated. But this thesis can, in fact, be substantiated. As Levinson (1980) points out,

There is probably no idea more central to thought about art than that it is an activity in which participants create things—these things being artworks. The whole tradition of art assumes art is creative in the strict sense, that it is a godlike activity in which the artist brings into being what did not exist beforehand much as a demiurge forms a world out of inchoate matter. The notion that artists truly *add* to the world, in company with cake-bakers, house-builders, law-makers, and theory-constructors, is surely a deep-rooted idea that merits preservation if at all

possible. The suggestion that some artists... instead merely *discover* or *select* for attention entities they have no hand in creating is so contrary to this basic intuition regarding artists and their works that we have a strong prima facie reason to reject it if we can. (Levinson, 1980, 8)²⁷

Thus, one of our core intuitions about art is that artworks are not something that is already there and to which artists merely draw our attention; rather, artworks are added to the world by artists. Meanwhile, if this is so, then there is good reason to think that novels are created, and not discovered.

Following Dodd (2000), one could object as follows. The foregoing argument assumes that one of our core intuitions about art—that artists add certain entities, artworks, to the world, and not merely draw our attention to something that already exists—supports the thesis that novels (as well as other artworks) are created, rather than discovered. But this assumption is dubious. It is true only if when we claim that artists add artworks to the world, we mean that they create these works. But do we, in fact, mean this in the given case? No: When we claim that artists add artworks to the world, we do not mean that they create these works; rather, we mean that they bring artworks *to the notice of the artworld*, or, in other words, introduce these works to our cultural life and open them up for appreciation.²⁸

This objection is based on the idea that when we claim that artists add artworks to the world, we mean *solely* that they bring these works to the notice of

²⁷Besides Levinson (1980), the view expressed in this quote is accepted by many other theorists, including Currie (1989, 63–64), Fisher (1991), and Predelli (2001, 282).

²⁸See Dodd (2000, 430).

the artworld. However, this idea is highly questionable. Surely, when claiming that artists add artworks to the world, we imply that they bring these works to the notice of the artworld. But, *prima facie*, we do not imply *just* that. We also imply that artists actually bring the works into existence and, hence, create them. In light of this, it is reasonable to hold that unless a proponent of the objection being discussed offers a plausible explanation of why our claim that artists add artworks to the world should not be interpreted as implying that artists create artworks, this objection cannot be considered successful.

Another argument in favor of the thesis that novels are created is as follows. Consider some great novel, say, *Emma*. If we conceive of it as existing eternally, before Austen's creative act, then a small part of the glory that surrounds this act seems to be removed.²⁹ But if this is so, then if *Emma* were discovered, and not created, it would not have the artistic value that it actually has. So *Emma* was created, and not discovered. Now, clearly, what has been said about *Emma* applies to many (though perhaps not all) other novels. So at least some novels were created, and not discovered.

Finally, the thesis that novels are created, and not discovered, is supported by the following consideration. "Works in the fine arts and in a number of other arts as well are, as either physical objects or events..., literally created" (Levinson, 2011b, 220). But if this is so, then it just seems wrong to accept the idea that novels are not created but rather discovered. As Levinson (2011b) puts it:

Shall paintings, drawings, etchings, sculptures, palaces, dances, films,

²⁹See Levinson (1980, 9).

and so on all be truly creatable, in the full sense of the word, and only... novels denied this possibility? There would be little profit, and false economy, in that. (Levinson, 2011b, 221).

Thus, a satisfactory substantiation of the thesis that novels are created, and not discovered, *can* be provided. Meanwhile, if this is so, then the foregoing response to the Creation Objection_{pt} fails.

Conclusion. Given what has been said, none of the responses to the Creation Objection_{pt} examined above stands up to criticism. At the same time, there seem no other potentially successful responses to this objection. As a result, there is good reason to consider the view that a novel is a pure type unacceptable.

6.1.4 A Novel as an Initiated Type

Let us now turn to the view that a novel is *an initiated type*,³⁰ or, in other words, an abstract type that is created as a result of a particular intentional action—namely, that of indicating, determining, or fixing a pure type (such as a pure sound structure or a pure syntactic structure)³¹ by means of creating an appropriate con-

³⁰This view has been put forward and defended by Levinson (1980, 2011b, 2012).

³¹According to Levinson (2012), the act of indicating a pure type normally amounts to the following acts: “a deliberate *choice*, an act of *appropriation*, an attitude of *approval*, and the *establishment* of a rule or norm” (Levinson, 2012, 54). Thus, when Chopin indicates the pure type of his Mazurka in A minor, Op. 17, no. 4, “he *chooses* or *selects* [certain] notes—here including pitches, rhythms, timbres, and dynamics...” (Levinson, 2012, 53). Furthermore, he

has a certain attitude—in part approval, in part appropriation—toward those particular notes. He doesn’t in effect merely say: ‘here are some sounds’ but rather, ‘here

cretum (say, a score or an inscription).^{32,33} Is this view acceptable? Here is a possible reason to answer “No.” Suppose a novel is an initiated type. Then, since initiated types are abstract, a novel is an abstractum. Meanwhile, the consensus is that abstracta are incapable of standing in causal relations.^{34,35} Thus, a novel cannot stand in such relations. Meanwhile, on the standard account of creation, x can be created only if x can stand in causal relations. So, assuming that this account is true, a novel cannot be created. But, surely, novels are created. So the view that a novel

are some sounds, they are now specifically mine, I embrace them, and in this exact sequence.’ (Levinson, 2012, 53)

Finally, “he establishes a *rule*, a *norm*, a miniature *practice*, whereby pianists play a piece *by Chopin* and not just any piece of music when they play *that* sequence of notes chosen by Chopin, and do so precisely *because* that sequence was chosen by Chopin” (Levinson, 2012, 54).

³²An example of an initiated type is the Ford Thunderbird—“a metal/ glass/plastic structure-as-indicated (or determined) by the Ford Motor Company on such and such a date” (Levinson, 1980, 22). Other examples of such a type include (a) Beethoven’s Quintet Opus 16—a sound structure-as-indicated by Beethoven in 1797; (b) the Lincoln penny—a visual structure-as-indicated-by the U.S. Government in 1909; and (c) the hedgehog—“a biological structure-as-determined-or-fixed by natural terrestrial evolution at a particular point in history” (Levinson, 1980, 22).

³³Although an initiated type is created, a significant component of this type—a particular pure type—is not. In this respect, an initiated type differs from many “ordinary” created objects—desks, cars, watches, etc.

³⁴See Cameron (2008, 296), Dodd (2000, 431), Dummett (1973, 491), Rosen (2012), and Mag Uidhir (2013, 10).

³⁵As Mag Uidhir (2013) points out, “should any general characterization of abstracta have a plausible claim to being standardly held, it clearly must be that abstracta are non-causal (especially given the standard, broad characterization of concreta as causally efficacious material inhabitants of space-time)” (Mag Uidhir, 2013, 10).

is an initiated type is false. Call this the Creation Objection_{it}.

Response to the Creation Objection_{it} (1). A proponent of the view that a novel is an initiated type could respond to the Creation Objection_{it} with the help of the arguments against the creatability of a novel that are discussed in Section 6.1.3. However, as has been shown, none of these arguments stands up to criticism, and, in addition, there are a number of strong considerations in favor of the view that novels are created. So the foregoing response to the Creation Objection_{it} fails.

Response to the Creation Objection_{it} (2). Another response available to a proponent of the view that a novel is an initiated type is as follows. The Creation Objection_{it} assumes that if x is created, then x can stand in causal relations. However, this assumption is false: Creating x does not require x 's capacity to stand in such relations.³⁶

Is this response satisfactory? Suppose x is not caused to exist. Then x either (a) does not exist, or (b) came into existence without any cause, or (c) is eternal. Suppose first that x does not exist. In this case, obviously, x is not created.³⁷ Suppose, on the other hand, that x came into existence without any cause. In this case, x is not created either. For suppose x comes into existence without any cause.

³⁶Note that although a proponent of this response *could* use one of the “acausal” accounts of creation discussed in Section 6.1.3 to support the thesis that creating x does not require x 's capacity to stand in causal relations, she is not committed to either account.

³⁷The thesis that if x does not exist, then x is not created is not universally accepted. Thus, following Meinong (1960), Parsons (1975) argues that being created is compatible with being nonexistent. However, Parsons (1975)'s argument is highly controversial. It is based on Meinong (1960)'s theory of objects, and the latter faces a number of serious problems (see Reicher (2016)).

Then nothing can be held responsible for x 's coming into existence and, hence, x cannot have an author. However, creation—at least, the kind of creation we are interested in—presupposes authorship: If x is created, then x has an author.

Finally, suppose that x is eternal. Then, pace Deutsch (1991), x is not created. For suppose x is eternal. Then there was no time at which x did not exist. However, creating x requires that there be such a time, since nothing that already exists can be created.

Thus, if x is not caused to exist, then x is not created. By contraposition, if x is created, then x is caused to exist. Meanwhile, if x is *caused* to exist, then x is capable of standing in causal relations. So, contrary to what the response being discussed suggests, the thesis that x is created only if x can stand in causal relations cannot be rejected. As a result, this response is unsatisfactory.

Response to the Creation Objection_{it} (3). Here is a more promising response available to a proponent of the view that a novel is an initiated type. The Creation Objection_{it} is successful only if abstracta cannot stand in causal relations. However, at least some abstracta can, in fact, stand in such relations. So the Creation Objection_{it} fails.

This response is based on the thesis that at least some abstracta can stand in causal relations. Clearly, this thesis requires substantiation. Why think that at least some abstracta can stand in causal relations?

(*Argument (i).*) Here is one possible reason to think so. Consider the singleton whose member is the Eiffel Tower. If this singleton cannot stand in causal relations, then it was not caused to exist and, hence, either (a) has always existed or (b) came

into existence without any cause. However, neither option is acceptable. Suppose the singleton containing the Eiffel Tower has always existed. Then, since the Eiffel Tower has not always existed, there was a time when the singleton existed, but the tower did not exist and had never existed. But this consequence is implausible. If an impure set exists at t , then its members exist at t or existed before t .³⁸ Consequently, there could be no time when the singleton containing the Eiffel Tower existed and the tower did not exist and had never existed. Thus, this singleton has not always existed.³⁹

Suppose, alternatively, that the singleton containing the Eiffel Tower came into existence without any cause. Then this singleton is a purely accidental object.

³⁸Note that this thesis is compatible with the idea that when no members of an impure set exist at t , this set can nevertheless exist at t if all of its members existed at some time before t .

³⁹According to Howell (2002), not only impure sets can have temporary existence; some *properties* can have such existence as well. Thus, consider the property of having a particular sound structure and being produced in a way that is properly connected to Beethoven's 1804–1808 acts of indication. In Howell (2002)'s view:

It surely... cannot be eternal. How can it already have existed in, say, 1600—or at the moment of the Big Bang—when the specific concrete entities to which it essentially relates had themselves not yet come into existence? To suppose that it can would be like supposing that your signature—not just inkmarks geometrically congruent to it, but actual marks that attest to you, to your own personal identity—could exist a million years before you do or that the set consisting of last night's thunderstorm and today's gusting of the wind could pre-exist both these events. (Howell, 2002, 112–113)

However, this consequence seems implausible. It is odd to claim that the singleton being discussed is a purely accidental object—an object that could have come into existence at any moment, even at the moment of the Big Bang, without a sufficient reason. The natural view is that this singleton came into existence when its constitutive element—the Eiffel Tower—came into existence.

Thus, neither the supposition that the Eiffel Tower has always existed nor the supposition that this tower came into existence without any cause is true. But then, given what has been said above, the singleton containing the Eiffel Tower was caused to exist. Meanwhile, if this is so, then, since singletons are sets and, as is generally agreed, the latter are abstract, at least some abstracta can stand in causal relations.

The foregoing argument, which might be called “Argument (i),” assumes that if an impure set exists at t , then its members exist either at t or at some time before t . Is this assumption true? Presumably, the main reason to say so comes from the natural inclination to treat impure sets as constituted by, or made of, their members. However, this inclination does not have any real basis. Although it is uncontroversial that a set *depends on* its members for its essence, there is nothing that would force us to accept the idea that a set is *constituted by*, or *made of*, its members.

Thus, the mentioned reason is unsatisfactory. And there seems no other potentially satisfactory reason to think that if an impure set exists at t , then its members exist either at t or at some time prior to t . In light of this, the assumption being discussed can be rejected in favor of the assumption that an impure set can exist at

t if its members exist not only at *t* or some time before *t* but at *any other time* as well.⁴⁰

The above reasoning may not seem very compelling. It states that an impure set can exist at *t* even when its members do not exist at *t* and have never existed before *t*. But why think that such a set can, in fact, exist at *t* in that case?

Here is a possible reason to think so. Suppose determinism—that is, the view that every event is caused to exist by certain previous events and the laws of nature—is true. Then for each future event, there is a causal chain that leads to this event. But if this is so, then at any time, there is, ontologically speaking, all relevant information about every future event. Meanwhile, if this consequence is true, then it seems natural to hold that sets—as well as perhaps other abstracta—that depend for their essence on certain contingent concreta can exist before these concreta come into existence.⁴¹

Following Howell (2002), one could object that this reason is unsatisfactory, since it is based on determinism and the latter is false.⁴² But this objection fails. It assumes that determinism is false. However, this assumption is unsubstantiated: There is no compelling argument that demonstrates the falsity of determinism.⁴³

But what if this objection is successful—and, hence, the foregoing reason in favor of the thesis that an impure set can exist at *t* even if its members do not exist

⁴⁰On this view, the singleton whose member is the Eiffel Tower exists not only after the creation of the Eiffel Tower but also before its creation—since the Eiffel Tower exists at *some* time.

⁴¹See Howell (2002).

⁴²See Howell (2002, 113).

⁴³See Hoefer (2016).

at t and have never existed before t is unsatisfactory? Should this thesis then be considered ungrounded? No. For there is an alternative reason that *is* satisfactory.

Before formulating this reason, let us first distinguish between two senses of the expression “to exist”: the presentist and the ontological.^{44,45} According to the *presentist* sense, x exists at t just in case x is present at t . In this sense, Socrates does not exist in 2017 but exists in 413 B.C. According to the *ontological* sense, x exists (hereafter: “exists_o”⁴⁶) at t just in case x is in the domain of our most unrestricted quantifier. In this sense, Socrates exists both in 2017 and 413 B.C.

So why think that an impure set can exist⁴⁷ at t even if its members do not exist at t and have never existed before t ? Suppose eternalism, or the view that not only entities present now but also past and future entities (for instance, Aristotle, the Great Library of Alexandria, World War II, lunar orbital stations, and interstellar flights) exist_o, is true. Then at every moment, all future concrete entities exist_o. But if this is so, then at every moment, *any* impure set containing one or more of the future concrete entities exists_o. Meanwhile, if at every moment, any such set exists_o, then there is good reason to think that impure sets exist even if their members have not come into existence yet. For suppose that at every moment, any

⁴⁴This distinction is made by Markosian et al. (2016). Cf. also Tillman (2011, 22–23).

⁴⁵Perhaps, given Quine (1948)’s view on the univocality of “exist,” it would be more accurate to speak of different uses, rather than senses, of “exist.” However, for the sake of convenience, I speak of different senses of “exist.” If the reader finds my talk of different senses of “exist” inappropriate, she is free to paraphrase it in terms of uses of “exist.”

⁴⁶The subscript “o” means that the expression is used in the ontological sense.

⁴⁷Here and in what follows, when “exist” is used without a subscript, it is used in the presentist sense.

impure set containing one or more of the future concrete entities exists_o. Then, given the non-physical character of sets, it seems natural to maintain that at every moment, impure sets containing one or more of the future concrete entities are present and, hence, exist. At the same time, the opposite idea—that such sets come into existence when, and only when, their members come into existence—seems unmotivated. Presumably, the main motivation to endorse this idea is that before the members of impure sets come into existence, these members are absent from our ontology (and hence, cannot serve as the constitutive “material” of the corresponding sets). But this motivation is absent if eternalism is true. For, as has already been mentioned, in this case, all future entities—including the future members of impure sets—exist_o and, hence, are contained in our ontology.

Here, one could object that the foregoing reasoning is based on a false assumption—that eternalism is true. Is this objection successful? Surely, eternalism (like virtually all philosophical views) has not been conclusively proven, and there are powerful rival theories—in particular, presentism (the view that only present objects exist) and past-and-presentism (the view that only past and present objects exist). However, there are strong arguments in favor of eternalism.⁴⁸ Furthermore, both presentism and past-and-presentism face serious problems.⁴⁹ Given this, the objection being discussed can hardly be called successful.

(*Argument (ii).*) Thus, in light of what has been said, it should be clear that Argument (i) fails. Yet this is not the only possible argument in favor of the thesis

⁴⁸See, e.g., Sider (2001).

⁴⁹See, e.g., Markosian et al. (2016) and Sider (2001).

that at least some abstracta can stand in causal relations. Another such argument is as follows.⁵⁰ Consider the following claims:

- (1) John's reading an Easton Ellis novel caused him to buy copies of all of Easton Ellis's other works.
- (2) John's reading that copy/a copy of *American Psycho* caused him to buy copies of all of Easton Ellis's other works. (Walters, 2013, 470)

Is either John's reading an Easton Ellis novel or his reading that copy/a copy of *American Psycho* the cause of his buying copies of all of Easton Ellis's other works? x can be the cause of y just in case x is *proportional* to y , where "proportional" is defined according to the following principle advanced by Yablo (1996):

Principle of Proportionality: C [a potential cause] is proportional to E [a potential effect] iff C is both required and enough for E , where

- C is required for E iff none of C 's determinables⁵¹ screens it off from E ,⁵² and
- C is enough for E iff C screens off all of its determinates⁵³ from E . (Yablo, 1996, 266–267; as cited in Walters, 2013, 470)

⁵⁰This argument has been advanced by Walters (2013).

⁵¹For all x and y , x is a determinable of y iff x is necessitated by y because x is immanent or included in y .

⁵²" C_1 screens C_2 off from E iff, had C_1 occurred without C_2 , E would still have occurred" (Yablo, 1996, 266-267; as cited in Walters, 2013, 470).

⁵³For all x and for all y , y is a determinate of x iff y necessitates x because x is immanent or included in y (see Walters (2013, 470)).

Is either John's reading an Easton Ellis novel or his reading that copy/a copy of *American Psycho* proportional to his buying copies of all of Easton Ellis's other works? Consider first John's reading an Easton Ellis novel. Is this reading both required and enough for his buying copies of all of Easton Ellis's other works? No. Since John only liked *American Psycho*, his reading an Easton Ellis novel cannot be *enough* for buying copies of all of Easton Ellis's other works. What about John's reading that copy/a copy of *American Psycho*? Is this reading both required and enough for his buying copies of all of Easton Ellis's other works? Again, the answer is "No":

Reading a particular copy, or indeed some copy or other, is not required to secure the effect, since [John] could have read the first three chapters of one copy and the remainder of another and the effect would still have occurred. Moreover, it seems that if John reads some copy of *American Psycho*, there is some copy of *American Psycho* that John reads. But John would have gone book shopping if he'd read *American Psycho* without having read one of the actual copies, and so [(2)] brings in too much in another way too. (Walters, 2013, 471).

Thus, neither John's reading an Easton Ellis novel nor his reading that copy/a copy of *American Psycho* is proportional to his buying copies of all of Easton Ellis's other works. But if this is so, then, given that x can be the cause of y only if x is proportional to y , neither John's reading an Easton Ellis novel nor his reading that copy/a copy of *American Psycho* can be the cause of his buying copies of all

of Easton Ellis's other works.

What is then the cause of John's buying these copies? Consider John's reading *American Psycho*. It is both required and enough for his buying copies of all of Easton Ellis's other works. So it is proportional to this buying. Given this, as well as the fact that x can be the cause of y if x is proportional to y , John's reading *American Psycho* can be the cause of his buying copies of all of Easton Ellis's other works. At the same time, there are no other events that are proportional to his buying copies of all of Easton Ellis's other works and, hence, that can be the cause of this buying. Thus, "it was [John's] reading *American Psycho* that caused him to buy the other novels" (Walters, 2013, 470). But *American Psycho* is an abstractum. So at least some abstracta can stand in causal relations.

The foregoing argument—call it "Argument (ii)"—assumes that the only plausible candidate for the cause of John's buying copies of all of Easton Ellis's other works is John's reading *American Psycho*. Is this assumption true? Consider the following claim:

- (3) John's reading an inscription of *American Psycho* caused him to buy copies of all of Easton Ellis's other works.

Is John's reading an inscription of *American Psycho* enough for his buying copies of all of Easton Ellis's other works? The effects from reading an inscription of *American Psycho* are equivalent to the effects from reading *American Psycho*. Meanwhile, by assumption, reading *American Psycho* is enough for buying copies of all of Easton Ellis's other works. Thus, John's reading an inscription of *American Psycho* is

enough for his buying these copies.

Is John's reading *required* for that? Following the logic of the explanation given in the case of (2), two reasons in favor of answering this question in the negative could be given. First, John would have bought copies of all of Easton Ellis's other works even if he had not read an inscription of *American Psycho*—say, if he had read the first three chapters of one inscription and the remainder of another. Second, John would have bought copies of all of Easton Ellis's other works even if he had not read any of the actual inscriptions of *American Psycho*—provided that he had read *American Psycho*. Is either reason successful?

The first reason is successful only if John did not read an inscription of *American Psycho* when he read the first three chapters of one inscription and the remainder of another. But is it true that he did not read such an inscription in this case? The sum of the concrete text of the first three chapters of one inscription of *American Psycho* and the concrete text of the remainder of another inscription of *American Psycho* is the concrete text that instances the text of *American Psycho*. Meanwhile, an inscription of a literary work is, by definition, any concrete text that instances the text of this work. So the sum of the concrete text of the first three chapters of one inscription of *American Psycho* and the concrete text of the remainder of another inscription of *American Psycho* is an inscription of *American Psycho*. As a result, the answer to the foregoing question is “No”: When John read the first three chapters of one inscription of *American Psycho* and the remainder of another, he did read an inscription of *American Psycho*. So the first reason fails.

Let us now turn to the second reason. It supports the claim that John's

reading an inscription of *American Psycho* is not required for his buying copies of all of Easton Ellis's other works only if (3) entails that John read one of the *actual* inscriptions of *American Psycho*. But (3) does not entail that. According to (3), John read an inscription *simpliciter*, and not necessarily an *actual* inscription, of *American Psycho*. Thus, the second reason also fails.

One could replace the second reason with the following alternative: John would have bought copies of all of Easton Ellis's other works even if he had not read an inscription of *American Psycho*—provided that he had read *American Psycho*. However, this alternative is doubtless unacceptable. It implies that a novel can be read without reading an inscription of this novel (regardless of whether this inscription is one of the actual inscriptions). But this implication is false. The only way to read a novel is to read one of its inscriptions.⁵⁴

Thus, John's reading an inscription of *American Psycho* is proportional to his buying copies of all of Easton Ellis's other works—and, hence, can be the cause of this buying. So the assumption involved in Argument (ii)—that the only plausible candidate for the cause of John's buying copies of all of Easton Ellis's other works is his reading *American Psycho*—is false. As a result, this argument fails.

(*Argument (iii).*) Here is yet another possible argument in favor of the thesis that at least some abstracta can stand in causal relations.⁵⁵ Consider the claim:

⁵⁴This is not to say, of course, the only way to *apprehend the content* of a novel is to read an inscription of this novel. The content of a novel can also be apprehended by listening to a reading of this novel.

⁵⁵Like the previous argument, this argument has been put forward by Walters (2013).

(4) Conan Doyle wrote *The Hound of the Baskervilles*.

(4) cannot be paraphrased as:

(4*) “Conan Doyle wrote MANUSCRIPT” (Walters, 2013, 471), where “MANUSCRIPT” refers to “a single manuscript from which all subsequent copies of *The Hound of the Baskervilles* were generated” (Walters, 2013, 471).

If (4) could be paraphrased as (4*), then Conan Doyle could not have written *The Hound of the Baskervilles* if he had not written MANUSCRIPT. But, of course, he could have written *The Hound of the Baskervilles* in the given case—say, by producing a different manuscript containing the text of *The Hound of the Baskervilles*.

Furthermore, the following potential paraphrase of (4) is also unsatisfactory:

(4**) Conan Doyle wrote a manuscript indiscernible from MANUSCRIPT.

If (4**) were an acceptable paraphrase of (4), then Conan Doyle could not have written *The Hound of the Baskervilles* if he had not written a manuscript indiscernible from MANUSCRIPT. However, since such “indiscernibility is neither necessary nor sufficient for the identity of novels” (Walters, 2013, 471), Conan Doyle could have written *The Hound of the Baskervilles* in this case—say, by producing a manuscript relevantly similar to, but not indiscernible from, MANUSCRIPT.

Thus, neither (4*) nor (4**) is an acceptable paraphrase of (4), and there are no other potentially acceptable paraphrases of (4) that do not require us to assume that *The Hound of the Baskervilles* can stand in causal relations. So (4) cannot be paraphrased as a claim that does not require us to assume that. Meanwhile, *The*

Hound of the Baskervilles is an abstractum. So some causal claims that presuppose the causality of abstracta cannot be paraphrased as claims that do not presuppose such causality. But if this is so, then there is good reason to believe that least some abstracta are capable of standing in causal relations.

The foregoing argument—call it “Argument (iii)” —assumes that (4) cannot be paraphrased as a claim that does not presuppose the causality of abstracta. Is this assumption true? Consider the following:

(4^{***}) Conan Doyle was the first to write an inscription of *The Hound of the Baskervilles*.⁵⁶

Is (4^{***}) an acceptable paraphrase of (4)? Here is one possible reason to say “No”: The equivalence of (4) and (4^{***}) has a false implication—that Conan Doyle could not have written *The Hound of the Baskervilles* if he had not written an inscription of *The Hound of the Baskervilles* from which all subsequent inscriptions of *The Hound of the Baskervilles* were generated. But this reason is unsatisfactory. Contrary to what it states, the equivalence of (4) and (4^{***}) implies that Conan Doyle could not have written *The Hound of the Baskervilles* if he had not written an inscription of *The Hound of the Baskervilles simpliciter*, not an inscription of *The Hound of the Baskervilles from which all subsequent inscriptions of The Hound of the Baskervilles were generated*.

Another possible reason in favor of the thesis that (4^{***}) is not an acceptable paraphrase of (4) is as follows: The equivalence of (4) and (4^{***}) has a

⁵⁶Note that (4^{***}) does not imply that Conan Doyle wrote a *particular* inscription of *The Hound of the Baskervilles* (such as the manuscript he actually wrote).

false implication—that Conan Doyle could not have written *The Hound of the Baskervilles* if he had not written an inscription *indiscernible* from the inscription from which all subsequent copies of *The Hound of the Baskervilles* were generated. This reason is also unsatisfactory. According to it, the equivalence of (4) and (4^{***}) implies that Conan Doyle could not have written *The Hound of the Baskervilles* if he had not written an inscription *indiscernible* from some other inscription. But, in fact, the equivalence of (4) and (4^{***}) does not imply that.

Here is yet another reason in favor of the thesis that (4^{***}) cannot be accepted as a paraphrase of (4): The equivalence of (4) and (4^{***}) has a false implication—that Conan Doyle could not have written *The Hound of the Baskervilles* if he had not written an inscription of *The Hound of the Baskervilles*. Like the previous reasons, this reason fails. To write a literary work, it is necessary to write an inscription of it.⁵⁷ So, contrary to what the reason being discussed assumes, Conan Doyle could not have written *The Hound of the Baskervilles* if he had not written an inscription of *The Hound of the Baskervilles*.

As a result, none of the foregoing reasons against the thesis that (4^{***}) is an acceptable paraphrase of (4) stands up to criticism. Meanwhile, there are no other potentially satisfactory reasons against this thesis. So, all things considered, (4^{***}) can be considered an acceptable paraphrase of (4).

Thus, given that (4^{***}) does not presuppose the causality of abstracta,⁵⁸ the

⁵⁷This is not to say, of course, that to *create* a literary work, it is necessary to write an inscription of it. A literary work can be created without writing any inscriptions (for instance, it can be created in one's imagination).

⁵⁸One might object as follows. (4^{***}), if true, commits us to the existence of an inscription

assumption involved in Argument (iii)—that (4) cannot be paraphrased as a claim that does not presuppose the causality of abstracta—is false. Meanwhile, if this is so, then this argument fails.

(*Argument (iv).*) Another possible argument in favor of the thesis that at least some abstracta can stand in causal relations is as follows.⁵⁹ It can be assumed that events are Kim-style ordered triples of objects, properties, and times. Meanwhile, if this assumption is true, then the following criterion of an object’s participation in an event in a causally relevant way seems appropriate: “An object participates in an event in a causally relevant way if and only if it is a member of the ordered triple that is the event” (Caplan and Matheson, 2004, 120). Suppose, therefore, that this criterion is true. Then, since abstracta can be members of ordered triples that are events, abstracta can participate in events in a causally relevant way and, hence, can stand in causal relations.

According to this argument, which might be called “Argument (iv),” a par-

The Hound of the Baskervilles. Meanwhile, there can be no inscription of *The Hound of the Baskervilles* unless there is the abstract type *The Hound of the Baskervilles* and, hence, a particular abstractum. Thus, (4^{***}), if true, commits us to the existence of a particular abstractum. But if this is so, then (4^{***}) does presuppose the causality of abstracta.

However, this objection fails. It assumes that there can be no inscription of *The Hound of the Baskervilles* unless there is the abstract type *The Hound of the Baskervilles* and, hence, a particular abstractum. But this assumption is false. An inscription of *The Hound of the Baskervilles* is just a particular concretum—namely, a particular concrete sequence of symbols written on paper (or some other material) or displayed on the screen of an electronic device. And there seems no real reason to hold that the existence of such a concretum entails the existence of any abstracta.

⁵⁹This argument has been advanced by Caplan and Matheson (2004).

ticular Kim-style view on the nature of events—that events are ordered triples of objects, properties, and times—is acceptable. But is this view, in fact, acceptable?

Before addressing this question, it is worth noting that characterizing the view that events are ordered triples of objects, properties, and times as “Kim-style” is somewhat inaccurate. There is reason to characterize this view as “Kim-style” only if according to Kim (1973, 1976), events are either ordered triples of objects, properties, and times or perhaps some similar ordered triples. But, strictly speaking, Kim (1973, 1976) does not treat events as ordered triples. In his view, an event is “a concrete object (or n -tuple of objects) exemplifying a property (or n -adic relation) at a time” (Kim, 1973).^{60,61}

So can the view that events are ordered triples of objects, properties, and

⁶⁰This is not to say, of course, that ordered triples do not have a role in Kim (1973, 1976)’s account of events. On this account, ordered triples of objects, properties, and times *individuate events*.

⁶¹Here is another reason to reject the thesis that Kim (1973, 1976) treats events as ordered triples. An ordered triple is an abstractum. So if Kim (1973, 1976) treated events as ordered triples, he would treat them as abstracta. However, in his view, events are concreta, as is clear from the following quote:

My events are “particulars” and “dated.” That they are dated is obvious. I am not clear what “particulars” are; but events in my sense have locations in space, namely the locations of their constitutive substances.... And my events are not “eternal” objects; they do not exist in all possible worlds; they exist only if the existence condition is met, which is a contingent matter of fact. *If this doesn’t show that something is “concrete” or “particular,” what does?* (Kim, 1976, 315; italics added)

times be accepted? According to an overwhelming majority of theorists, including Davidson (1969), Kim (1976), Lewis (1986), and Dretske (1977), events are spatiotemporal (or—to use Schaffer (2016)’s term—immanent). But ordered triples of objects, properties, and times are sets, and—at least, on the standard account⁶²—sets are not spatiotemporal. Thus, there is a *prima facie* reason against the view that events are ordered triples of objects, properties, and times. At the same time, there seems no good reason in favor of this view. As a result, all things considered, the question posed above should be answered in the negative. Meanwhile, if this is so, then Argument (iv) fails.

The foregoing objection can be avoided if the account of events qua ordered triples of objects, properties, and times is replaced with Kim (1973, 1976)’s account of events—the account according to which events are concrete objects exemplifying properties at particular times. But can Argument (iv) be accepted if it presupposes the truth of the latter account?

Consider the event described by the sentence “John apologized” and the event described by the sentence “John said ‘I apologize.’” Clearly, under at least some circumstances, these events are the same.⁶³ However, as Davidson (1969) points out, Kim (1976)’s account entails that they *must* be different. For suppose this account is true. Then “two sentences are about the same event [only] if they assert truly of

⁶²On the account put forward by Maddy (1990), sets exist in spacetime. However, this account is highly controversial. For powerful objections to it, see Balaguer (1998), Milne (1994), Riskin (1994), Carson (1996), and Maddy (1997).

⁶³Note that it is not claimed here that the event described by the sentence “John apologized” and the event described by the sentence “John said ‘I apologize.’” are *always* the same.

the same particulars (i.e. substances) that the same properties (or relations) hold of them” (Davidson, 1969, 300). The sentences “John apologized” and “John said ‘I apologize’” do not assert that the same properties hold of John. According to the first sentence, John has the property of having apologized, whereas the second sentence states that he has a different property—that of having said “I apologize.” Thus, the events described by these sentences cannot be the same.

Thus, Kim (1973, 1976)’s account of events has an unpalatable consequence—that the event described by the sentence “John apologized” and the event described by the sentence “John said ‘I apologize’” are necessarily different. In a similar way, it can be shown that this account has other unpalatable consequences—for instance, that a murder cannot be a killing, that the signing of a check cannot be the paying of a bill, and that the death of Walter Scott cannot be the death of the author of *Waverley*.^{64,65} Given this, there is good reason against Kim (1973, 1976)’s account of events. Meanwhile, if this is so, then, clearly, Argument (iv) cannot be accepted if this argument presupposes the truth of this account.

Despite what has been said, however, let us assume, for the sake of argument, that Kim (1976)’s account of events is true. Can the modified version of Argument (iv) (the version that presupposes the truth of Kim (1973, 1976)’s account of events) be accepted in this case? The original version of this argument assumes that an object participates in an event in a causally relevant way if and only if it is a member of the ordered triple that is the event. In the context of this version, this assumption

⁶⁴See Davidson (1969, 300).

⁶⁵For an exposition of other problems with Kim (1976)’s account, see Davidson (1969).

is reasonable. But why assume that an object participates in an event in a causally relevant way if and only if it is a member of the ordered triple that is the event, in the context of the modified version of Argument (iv)? Prima facie, the construal of events according to this version warrants the following criterion of an object's participation in an event in a causally relevant way: An object participates in an event in a causally relevant way if and only if it is either (a) the object constituting this event or (b) the exemplification of the property⁶⁶ constituting this event. However, if the modified version of Argument (iv) employs the mentioned criterion, then this version is inconsistent. For suppose this version actually employs this criterion. Then this version entails that abstracta can stand in causal relations if and only if an object constituting an event or the exemplification of a property constituting an event can be abstract. However, neither the object nor the exemplification of the property can, in fact, be abstract. The reason for this is that (a) they constitute an event and (b) according to Kim (1976)'s account, anything that constitutes an event is concrete and, hence, non-abstract. So according to the modified version of Argument (iv), abstracta cannot stand in causal relations. But, at the same time, this version claims that abstracta *can* stand in such relations. As a result, it is inconsistent.

Thus, the modified version of Argument (iv) is either (a) based on a completely unwarranted assumption—that an object participates in an event in a causally relevant way if and only if it is a member of the ordered triple that is the event—or

⁶⁶Here and in what follows, by “a property,” I mean a property-universal, not a property-instance.

(b) inconsistent. Given this, this version cannot be accepted.

Here, one could respond as follows. The foregoing objection is successful only if in the case of the modified version of Argument (iv), the criterion of an object's participation in an event in a causally relevant way is equivalent to (a) the thesis that an object participates in an event in a causally relevant way if and only if this object is an object constituting this event or the exemplification of the property constituting this event, or (b) some other thesis that precludes abstracta from being capable of standing in causal relations. But, in the mentioned case, this criterion does not have to be equivalent to either (a) or (b). It can be equivalent to a thesis that does entail (together with the other assumptions of Argument (iv)) that abstracta can stand in causal relations—namely, (c) the thesis that an object participates in an event in a causally relevant way if and only if this object is the object constituting this event, the exemplification of the property constituting this event, *or the exemplified property itself*.

This response assumes that in the case of the modified version of Argument (iv), the criterion of an object's participation in an event in a causally relevant way can be equivalent to (c). Is this assumption acceptable? The relata of causal relations are events, and on Kim (1976)'s account, the latter are concrete. Given this, it is natural to hold that participating in events in a causally relevant way requires being concrete. Meanwhile, if participating in events in a causally relevant way requires that, then no property (or any other abstractum) can participate in events *in a causally relevant way*—and, hence, (c) is false. Thus, there is a strong *prima facie* reason against (c). At the same time, there is no real reason in favor

of (c). Given this, the mentioned assumption is false. But then the response being discussed fails.

(Argument (v).) Here is yet another argument in favor of the thesis that at least some abstracta can stand in causal relations.⁶⁷ It can be assumed that events are Lewis-style sets of spacetime points. Meanwhile, if this assumption is true, then the following criterion of an object's participation in an event in a causally relevant way seems right: "An object participates in an event in a causally relevant way if and only if the set of points in the spacetime region that it occupies is a subset of the set of spacetime points that is the event" (Caplan and Matheson, 2004, 121). Suppose, therefore, that this criterion is true. Then, given that an abstractum can occupy a set of points in the spacetime region that is a subset of the set of spacetime points that is the event, abstracta can participate in events in a causally relevant way and so can stand in causal relations. Call this "Argument (v)."

Is Argument (v) successful? The answer is "Yes" only if abstracta can occupy a set of spacetime points and, hence, have a spatiotemporal location. But can abstracta, in fact, have such a location? There is a good reason to hold that they cannot. The reason is that if abstract objects can be located in spacetime, then the very concept of abstractness as well as the abstract/concrete distinction become extremely vague, if not meaningless.

One can respond that the meaningfulness of the concept of abstractness as well as the abstract/concrete distinction can be preserved by assuming that those abstracta that are located in spacetime (hereafter: "spatiotemporal abstracta") are

⁶⁷Like Argument (iv), this argument has been put forward by Caplan and Matheson (2004).

located in spacetime in a way that is *essentially different* from the way in which concreta are located in spacetime. Regarding this response, a natural question arises: How exactly are spatiotemporal abstracta located in spacetime?

Following Goodman (2007), one could answer this question by saying that spatiotemporal abstracta have precise temporal locations but vague spatial ones (where x has a vague spatial location at some spatial region y just in case (a) it is indeterminate that x is located at any particular spatial region within y , but (b) it is determinately true that x is located within y , and (c) it is determinately false that x is located outside y). However, this answer is incompatible with the response being discussed. Suppose that spatial abstracta have vague spatial locations. Then, given that the abstract/concrete distinction is exhaustive, the way spatiotemporal abstracta are located in space can be essentially different from the way concreta are located in space *only if* there are no concreta that have vague spatial locations. However, there are, in fact, such concreta. Consider elementary particles (such as quarks). They are doubtless concrete. At the same time, according to standard views of quantum mechanics, the spatial locations of these particles are vague.⁶⁸ Or consider Mount Everest. As Tye (1990) points out,

It seems obvious that there is no line which sharply divides the matter composing Everest from the matter outside it. Everest's boundaries are fuzzy. Some molecules are inside Everest and some outside. But some have an indefinite status: there is no objective, determinate fact of the matter about whether they are inside or outside. (Tye, 1990, 535)

⁶⁸See Kane (2005, 8–9).

If Tye (1990) is right—and there seems no real reason to think otherwise—then Everest has a vague spatial location. Meanwhile, there is no doubt that Everest is a concretum.

Thus, the answer being discussed entails that the way spatiotemporal abstracta are located in space is not essentially different from the way concreta are located in space. However, if this entailment is true, then the foregoing response does not work, as this response is based on the assumption that the way spatiotemporal abstracta are located in space *is* essentially different from the way concreta are located in space.

So the foregoing answer to the question “How exactly are spatiotemporal abstracta located in space?” is unsatisfactory. At the same time, there seems no other potentially satisfactory answer to this question. Meanwhile, if there is no such answer, then the foregoing response is too metaphysically obscure and, hence, cannot be accepted.

(*Argument (vi).*) Besides the arguments discussed above, there is yet another argument in favor of the thesis that at least some abstracta can stand in causal relations.⁶⁹ This argument is as follows. It can be assumed that events are Lewis-style sets of spacetime points. Meanwhile, this assumption is compatible with the following criterion of an object’s participation in an event in a causally relevant way: “An object participates in an event in a causally relevant way if and only if either [(a)] the set of points in the spacetime region that it occupies is a subset of the set

⁶⁹Like the previous two arguments, this argument has been put forward by Caplan and Matheson (2004).

of spacetime points that is the event, or [(b)] it is appropriately related to another object—for example, a token, an instance, or a performance—such that the set of spacetime points in the region that that object occupies is a subset of the set of spacetime points that is the event” (Caplan and Matheson, 2004, 121). Suppose, therefore, that this criterion is true. Then, since an abstractum can be appropriately related to an object such that the set of spacetime points in the region that this object occupies is a subset of the set of spacetime points that is the event, abstracta can participate in events in a causally relevant way and, hence, can stand in causal relations.

This argument—call it “Argument (vi)” —is based on a particular criterion of an object’s participation in an event in a causally relevant way—namely, the criterion according to which an object participates in an event in a causally relevant way if and only if either (a) the set of points in the spacetime region that it occupies is a subset of the set of spacetime points that is the event, or (b) it is appropriately related to another object such that the set of spacetime points in the region that that object occupies is a subset of the set of spacetime points that is the event. Is this criterion acceptable? It entails that an object can participate in an event in a causally relevant way if this object is appropriately related to another object such that the set of spacetime points in the region that that object occupies is a subset of the set of spacetime points that is the event. Here, a natural question arises: Under what conditions is an object *appropriately* related to that other object (hereafter: “*O*”)? Here is a possible answer: An object is appropriately related to *O* if (and only if) *O* or one of *O*’s properties is an instance of this object. But why think that

being related to O in this way is sufficient for participating in an event in a causally relevant way? There seems no satisfactory explanation of that. Meanwhile, if there is no such explanation, then the foregoing answer cannot be accepted. Are there any other potentially acceptable answers to the question “Under what conditions is an object *appropriately* related to O ?” No—or so it seems. But if this is so, then the criterion employed in Argument (vi)—and, hence, the argument itself—cannot be accepted.

Here, one might ask: What criterion of an object’s participation in an event in a causally relevant way should we then adopt, if events are Lewis-style sets of spacetime points? It is natural to say that an object participates in an event in a causally relevant way if this object is, in some sense, a “part” of this event. So if events are Lewis-style sets of spacetime points, then an object participates in an event in a causally relevant way if this object is, in some sense, a “part” of the set of spacetime points that is the event. Meanwhile, an object is a “part” of a set of spacetime points just in case this object occupies a subset of this set. Given this, if events are Lewis-style sets of spacetime points, there is good reason to endorse the criterion of Argument (v): An object participates in an event in a causally relevant way if and only if this object occupies a set of spacetime points that is a subset of the set of spacetime points that is the event.

(Conclusion.) Thus, none of the Arguments (i)–(vi) stands up to criticism. Meanwhile, so far as I am aware, there are no other potentially acceptable arguments in favor of the thesis that at least some abstracta can stand in causal relations. As a result, there seems no real reason to uphold this thesis. But if there is no such

reason, then the foregoing response to the Creation Objection_{it} is unsatisfactory.

Response to the Creation Objection_{it} (4). A proponent of the view that a novel is an initiated type could also respond to the Creation Objection_{it} as follows. There is no good reason in favor of the thesis that abstracta cannot stand in causal relations. As a result, this thesis is ungrounded. But if this is so, then the Creation Objection_{it} is itself ungrounded and, hence, does not pose any threat to the view that a novel is an initiated type.

The foregoing response assumes that there is no good reason in favor of the thesis that abstracta cannot stand in causal relations. Is this assumption true?

According to Dodd (2000), the mentioned thesis can be substantiated using the following argument.

Statements seemingly reporting causal relations between abstracta can always be paraphrased in such a way as to reveal that the relata of the causal relation are really concrete.... We may well say, for example, that the bitter taste of a certain substance is caused by the shape of its molecules, but in saying this we do not commit ourselves to the idea that an abstract object—a certain shape—causes the bitter taste; what causes the bitter taste is the presence of a molecule of that shape. Likewise, it is not an abstract entity—bitterness—which is causally produced but a substance's bitter taste. (Dodd, 2000, 431–432)

Thus, for any statement that seems to imply the existence of causal relations between abstracta, there is a satisfactory paraphrase that does not involve this implication.

Meanwhile, if this is so, then it is unreasonable to treat abstracta as capable of standing in causal relations.

However, Dodd (2000)'s argument is unsatisfactory. For, as Caplan and Matheson (2004) point out, it is invalid. Its validity holds only if the paraphrasability claim—"that sentences that seemingly report causal relations between abstract objects can always be paraphrased as sentences that actually report causal relations between concrete objects" (Caplan and Matheson, 2004)—entails that abstracta cannot stand in causal relations. But this claim does not entail that. For suppose it is true. Then abstracta might be capable of standing in causal relations even if all sentences that seemingly report causal relations between abstract objects could be paraphrased as sentences that actually report causal relations between concrete objects.

Yet there is an argument similar to Dodd (2000)'s that does not face the foregoing problem. This alternative argument can be formulated in the following way. One reason to believe that abstracta can stand in causal relations is that our ordinary talk seems to require abstracta to have the capacity to stand in such relations. But, in fact, our ordinary talk does not require that—because it can be paraphrased so that abstracta do not stand in causal relations. Thus, the foregoing reason to treat abstracta as capable of standing in causal relations is false.⁷⁰ Are

⁷⁰Note that it is not assumed here that the possibility of paraphrasing our ordinary talk as talk that does not require abstracta to be capable of standing in causal relations is incompatible with the idea that there could be abstracta capable of that. The point is that because of this possibility, the mentioned reason to hold that abstracta can stand in causal relations is false.

there are any other plausible reasons to treat abstracta as capable of that? No—or so it seems. As a result, it is not *necessary* to treat abstracta as capable of standing in causal relations. But then, in light of the principle of ontological simplicity,⁷¹ abstracta should be treated as incapable of standing in such relations.

The foregoing argument faces two potential objections. One of these objections is that this argument involves a false claim—that all sentences that require abstracta to be capable of standing in causal relations can be paraphrased away (that is, paraphrased as sentences that do not require abstracta to be capable of that). The second objection is that one of the claims of this argument—that it is possible to paraphrase away our ordinary talk that requires abstracta to be capable of standing in causal relations—cannot be substantiated. Is either objection successful?

According to the first objection, there are sentences that (a) require abstracta to be capable of standing in causal relations but (b) cannot be paraphrased away. But why think that there are such sentences? As already shown, Walters (2013)'s answer to this question is unsatisfactory. Following Caplan and Matheson (2004) and Dodd (2007), one could give the following answer. Consider sentences reporting causal interactions between events. Since events are abstract, these sentences require abstracta to be capable of standing in causal relations. At the same time, it is unclear how these sentences can be paraphrased away.

The foregoing answer assumes that events are abstract. But is this assumption true? As noted in the introduction, the distinguishing features of an abstractum are

⁷¹Roughly, the principle of ontological simplicity is that View A should be preferred to View B if, other things being equal, View A is less ontologically complex than View B.

as follows:

- being nonspatial
- being atemporal
- being causally impassive (= being incapable of being causally affected by anything)
- being causally inactive (= being incapable of causally affecting anything)
- being modally inflexible (= having all intrinsic properties essentially)

—while the distinguishing features of a concretum are as follows:

- being in space and/or time
- being causally efficacious (= being capable of causally affecting something and of being causally affected by something)
- being modally flexible (= having some intrinsic properties nonessentially).

Consider now events. As pointed out in Chapter 4,⁷² they are to be construed as entities that (a) are said to *occur*, or *happen*, or *take place*, (b) have relatively vague spatial boundaries and relatively crisp temporal boundaries, (c) tolerate co-location, (d) cannot move, and (e) take up time and persist by perduring, that is, by having distinct temporal parts (or stages) at different times.⁷³ Given this, it is

⁷²See Chapter 4, Section 4.2.

⁷³See Casati and Varzi (2015).

clear that events have the first of the abovementioned features of a concretum—that of being in space and time. Furthermore, on the standard view—endorsed by Davidson (1963, 1967), Kim (1973), and Lewis (1986), among others—events have the second of these features—that of being causally efficacious.⁷⁴ Do events have the third feature—that of being modally flexible? To be modally flexibly, it is sufficient to have at least some intrinsic properties nonessentially. Meanwhile, some events possess such properties nonessentially, since some events could have occurred in a way slightly different, with regard to their intrinsic properties, from the way they actually occurred (for example, World War II could have started a few minutes earlier). So the answer to the foregoing question is “Yes.”

Thus, events have all the distinguishing features of concreta. At the same time, given what has been said, it is clear that events do not have any of the distinguishing features of abstracta. Given this, events should be treated as concreta, not abstracta. Meanwhile, if this is so, then the assumption of the foregoing answer to the question “Why think that there are sentences that (a) require abstracta to be capable of standing in causal relations but (b) cannot be paraphrased away?” is false—and, as a result, this answer is unsatisfactory.⁷⁵

⁷⁴Moreover, on this view, events are the paradigmatic entities that possess causal efficacy.

⁷⁵Here, one could object as follows. The foregoing account of events is unsatisfactory. It assumes that there is only one type of events—events qua particulars (John’s flicking his finger at noon on September 15, 2016; Mary’s riding her bicycle tomorrow morning). But this assumption is false. Besides events qua particulars, there are also what might be called “generic events”—events such as finger-flicking and bicycle-riding.

Is this objection successful? It is based on the idea that “generic events” are events. Yet this

Thus, neither answer to the question “Why think that there are sentences that (a) require abstracta to be capable of standing in causal relations but (b) cannot be paraphrased away?” is satisfactory. At the same time, there are no other potentially satisfactory answers to this question. Given this, the first of the objections under consideration is unsubstantiated and, hence, fails.

The second objection also fails. It assumes that the claim that it is possible to paraphrase away our ordinary talk that requires abstracta to be capable of standing in causal relations cannot be substantiated. However, this assumption is false. Here is how the mentioned claim can be substantiated. There is a good inductive reason to hold that it is possible to paraphrase away our ordinary talk that seems to require idea is not obvious. For it seems more natural to treat “generic events” not as events but as entities that are *instanced* by events—*types of events*.^{*} In light of this, the question posed above should, I think, be answered in the negative.

But what if “generic events” *are* events? Should we then accept the foregoing answer to the question “Why think that there are sentences that (a) require abstracta to be capable of standing in causal relations but (b) cannot be paraphrased away?”? No. If by “events” is meant *only* events qua particulars or *both* generic events and events qua particulars, then this answer falsely assumes that events are abstract (for, clearly, events qua particulars are not abstract). If, on the other hand, by “events” is meant *only* generic events, then another assumption of this answer—that sentences reporting causal interactions between events cannot be paraphrased away—is false (for such sentences can always be paraphrased away in terms of *instances* of these events—namely, certain events qua particulars).

* On this construal, finger-flicking is the type of event that is instanced by events such as John’s flicking his finger at noon on September 15, and bicycle-riding is the type of event that is instanced by events such as Mary’s riding her bicycle tomorrow morning.

abstracta to be capable of standing in causal relations. The reason is that it is possible to provide numerous examples of paraphrasing away this talk (a number of such examples have been given above). At the same time, there are no reasons to believe the opposite—that paraphrasing away our ordinary talk that seems to require abstracta to be capable of standing in causal relations is impossible.⁷⁶ In light of this, it is reasonable to believe that it is possible to paraphrase away our ordinary talk that requires abstracta to be capable of standing in causal relations.

Thus, given what has been said, it is clear that there is at least one good reason in favor of the thesis that abstracta cannot stand in causal relations. But if this is so, then one of the assumptions of the above response to the Creation Objection_{it} is false, and, hence, this response fails.

Conclusion. Our analysis has shown that none of the foregoing responses to the Creation Objection_{it} is successful. Meanwhile, there are no other potentially successful responses to this objection. Given this, the view that a novel is an initiated type cannot be accepted.

6.1.5 A Novel as a Historical Individual/an Abstract Artifact

There are two abstractionist views left. One of these views, advanced and defended by Rohrbaugh (2003), is that a novel is “*a historical individual*”—something

⁷⁶Presumably, there would be reason to believe this only if there were a satisfactory counterexample—a sentence that (a) requires abstracta to be capable of standing in causal relations but (b) cannot be paraphrased away. But, so far as I am concerned, there is no such counterexample.

that is (a) *abstract* and, hence, external (nonmental) and nonspatial; (b) *modally flexible*, or such that at least some of its intrinsic properties could have been different; (c) *temporally flexible*, or such that it is “subject, in principle, to change in [its intrinsic] properties over time” (Rohrbaugh, 2003, 186); (d) *temporal*, or such that it comes into and goes out of existence; and (e) *ontologically dependent on certain concreta*—namely, “embodiments,” or, in other words, *concreta* that ground its essential properties. According to the second view, advanced and defended by Thomasson (1999, 2004), a novel is “*an abstract artifact*,” where the latter is very similar to a historical individual: Like a historical individual, it is abstract (in the sense of being nonmental and nonspatial), modally and temporally flexible, temporal, and ontologically dependent on certain *concreta*;⁷⁷ its only difference from a historical individual is that unlike the latter, it ontologically depends not only on its *embodiments* but also on certain other *concreta*—for example, a particular language, linguistic capacities required to be able to understand this language, and knowledge of relevant background information.⁷⁸

⁷⁷According to Thomasson (1999), an abstract artifact can be characterized as anything that (a) is “abstract in the sense of lacking a spatiotemporal location” (Thomasson, 1999, 38), (b) depends for its existence and essence on contingent entities (Thomasson, 1999, 38), and (c) is “not timeless but instead [is] created at a particular time in particular circumstances, can change, and can once again cease to exist even after [it has] been created” (Thomasson, 1999, 38). Meanwhile, this characterization entails that an abstract artifact is abstract (in the sense of being nonmental and nonspatial), modally and temporally flexible, temporal, and ontologically depends on certain *concreta*.

⁷⁸See Thomasson (1999, 11).

Is either of the foregoing views acceptable? Clearly, both of them face objections analogous to the Creation Objection_{it}. Can these objections be successfully countered by a proponent of the view that a novel is a historical individual or the view that a novel is an abstract artifact? From a theoretical perspective, both proponents are in no better position than a proponent of the view that a novel is an initiated type. They would be in a better position only if they were not committed to the abstractness of novels. But they are committed to that: Both the view that a novel is a historical individual and the view that a novel is an abstract artifact entail that a novel is an abstractum. Given this as well as the critical analysis provided in the previous subsection, the answer to the foregoing question is “No.” As a result, both the view that a novel is a historical individual and the view that a novel is an abstract artifact are *at least* as problematic as the view that a novel is an initiated type. Therefore, neither view can be accepted.⁷⁹

6.2 Concretist Views

Thus, none of the abstractionist views examined above is acceptable. Let us now turn to the concretist alternatives mentioned at the beginning of this chapter:

⁷⁹A view similar to the views discussed here has been provided by Ingarden (1973). According to it, a literary work “is a ‘purely intentional formation,’ derived from the sentence-forming activities of its author(s), and founded on some public copy of these sentences, and also depending for its existence and essence on a relation to certain ideal meanings attached to the words of the text” (Thomasson, 2008). However, Ingarden (1973)’s view is unacceptable. The reason why this is so is analogous to the reason why the views discussed in this section are unacceptable.

the view that a novel is an embodiment of this novel and the view that a novel is a mereological sum of embodiments of this novel.

6.2.1 A Novel as a Particular Embodiment

Consider first the view that a novel is *an embodiment of this novel*.^{80,81,82}

One reason against this view is as follows. Suppose some novel N is one of its embodiments—namely, embodiment E . Now, a question arises: Why is N (identical to) E , and not E_{alt} , where the latter is some embodiment of N other than E ?

One could say that N is E , and not E_{alt} , since E has all the relevant artistic properties of N . But this explanation is unsatisfactory, as nothing stops us from supposing that E_{alt} also has all of these properties.

Alternatively, one could try to explain why N is E , and not E_{alt} , by saying that E is the first embodiment of N to come into existence (the original manuscript created by N 's author). But this explanation is also unsatisfactory. First of all,

⁸⁰Recall that by “an embodiment of a novel” is meant a concrete singular entity that possesses all, or at least sufficiently many, of the relevant artistic properties of this novel.

⁸¹Proponents of this view are Collingwood (1958), Mag Uidhir (2013), and (arguably) Sartre (2004), among others.

⁸²One might object that since the existence of an embodiment of a novel implies that this novel exists *independently* of this embodiment, the view that a novel is an embodiment of this novel is incoherent. However, this objection fails. It assumes that the existence of an embodiment of a novel implies that this novel exists independently of this embodiment. But this assumption is false. An embodiment of a novel is just a concretum that has particular textual properties (and perhaps is appropriately related to the author). And, surely, such a concretum can exist even if the corresponding novel exists *only* qua an entity identical to this concretum.

it depends on the assumption that E comes into existence before E_{alt} . However, this assumption can be rejected in favor of the assumption that E and E_{alt} come into existence simultaneously. Furthermore, according to the explanation being discussed, what explains the fact that N is E , and not E_{alt} , is a particular purely temporal and/or logical difference in the order of coming into existence between E and E_{alt} . But why is such a difference sufficient to explain this fact? Whether x is N (or some other novel) depends *solely* on what *artistic* properties x has. However, neither the temporal nor the logical property of coming into existence after or before some embodiment is, *in itself*, an artistic property. (To see this, imagine that E and E_{alt} are correctly produced in the same way and in the same relevant cultural-historical context but at different times. In this case, E and E_{alt} differ in the temporal and logical properties related to the order of E 's and E_{alt} 's coming into existence. However, there is no reason to ascribe different artistic properties to E and E_{alt} .) So the fact that E and E_{alt} differ in the order of coming into existence is not sufficient to explain why N is identical to E , and not E_{alt} .

Thus, neither explanation can be accepted. Meanwhile, there is no other potentially acceptable explanation of why N is E , and not E_{alt} . As a result, the thesis that N is E is essentially ungrounded. But if this is so, then, since this thesis is a direct consequence of the view that a novel is an embodiment of this novel, this view is itself essentially ungrounded.

Here is another consideration against the view that a novel is an embodiment of this novel. Suppose *Moby-Dick* is one of its embodiments—say, some embodiment E . Suppose next that E is completely destroyed. In this case, *Moby-Dick* is also

completely destroyed. However, this result is problematic. Surely, if x is not the only embodiment of a novel, then destroying x cannot be sufficient to destroy this novel. Meanwhile, E is not the only embodiment of *Moby-Dick* (in fact, there are many other embodiments of *Moby-Dick*). So destroying E cannot be sufficient for the destruction of *Moby-Dick*. Thus, *Moby-Dick* cannot be E . But if this is so, then the view that a novel is an embodiment of this novel is false.

Yet another objection to this view is as follows. Suppose *Moby-Dick* is one of its embodiments—say, some embodiment E . Suppose next that E is modified in some way—say, by tearing a number of pages out of it (if it is an inscription), or by forgetting some of its text (if it is a mental entity), or by mispronouncing some of its text (if it is a reading). Then *Moby-Dick* is itself modified. However, this consequence is unacceptable. The foregoing modification of E cannot affect the artistic properties of *Moby-Dick* and, hence, cannot cause *Moby-Dick* to be modified.⁸³ Thus, *Moby-Dick* cannot be E . But then the view that a novel is an embodiment of this novel is false.

Finally, this view faces the following objection. If a novel is an embodiment of this novel, then *Moby-Dick* can be touched/smelled/tasted—by touching/smelling/tasting the embodiment *Moby-Dick* is identical to. But this result is wrong. A novel (as opposed to, say, a text of this novel printed on paper) cannot be touched/tasted/smelled—by touching/tasting/smelling its embodiments or by any

⁸³Of course, under certain circumstances, a novel can be modified by modifying its embodiment. For example, a novel can be modified in this way if the modification has the form of editing and is carried out by the author(s) of the novel.

other means.⁸⁴ So the view that a novel is an embodiment of this novel is false.

There seems no way for a proponent of the view that a novel is an embodiment of this novel to plausibly respond to the foregoing objections. In light of that, this view cannot be accepted.

6.2.2 A Novel as a Mereological Sum of Embodiments

Let us now turn to the view that a novel is *a mereological sum of embodiments of this novel*.^{85,86,87} One objection to this view is as follows. Suppose a novel is a

⁸⁴Here, it is assumed that mental entities amount to certain physical states of the brain. If this assumption is rejected and mental entities are taken to have a special, nonphysical nature (of the sort that precludes such entities from being touched, tasted, or smelled), then the objection being discussed does not apply to one version of the view that a novel is an embodiment of this novel—the version according to which a novel is a particular mental entity.

⁸⁵Although, to my knowledge, this view has not been defended by anyone, a similar view on the nature of a musical work has recently been defended by Alward (2004), Caplan and Matheson (2006, 2008), Tillman (2011), and Tillman and Spencer (2012).

⁸⁶Note that the view that a novel is a mereological sum of embodiments of this novel is essentially different from the view that a novel is a set of embodiments of this novel.* According to the latter view, a novel is an abstractum. For, on this view, a novel is a set, and, as is generally agreed, sets are abstract. According to the former view, however, a novel is a concretum. The reason for this is that on this view, a novel is a mereological sum of concreta (namely, of certain embodiments of this novel), and any mereological sum of concreta is a concretum.

* The view that a novel is a set of embodiments of this novel has been discussed in Section 6.1.1.

⁸⁷One might object that since the existence of a mereological sum of embodiments of a novel implies that this novel exists *independently* of this sum, the view that a novel is a mereological sum of embodiments of this novel is incoherent. However, this objection fails. The explanation of why this is so is analogous to the explanation of why the objections discussed in Footnotes 10 and

mereological sum of embodiments of this novel. Then a complete apprehension of a novel requires an apprehension of all the embodiments that make up the mereological sum with which this novel is identified. But this consequence is false. To fully apprehend a novel, it is sufficient to properly apprehend *just one* of its embodiments (provided that this embodiment is correct). As a result, the view that a novel is a mereological sum of embodiments of this novel is false.

Here is another objection to this view. Consider the following principle: For all x , if x 's part is modified at t , then x is modified at t . This principle is uncontroversial. Thus, suppose that a part of a table is modified at t —for example, by making a crack in its top. Then it seems right to say that the table itself is modified at t . Likewise, if John's part is modified at t —for instance, by means of a surgery—then it is natural to say that John himself is modified. Suppose now that a novel is a mereological sum of embodiments of this novel. Then, given the foregoing principle, whenever one modifies an embodiment of *Moby-Dick* at t , one modifies *Moby-Dick* at t . But this consequence is false. Modifying an embodiment of *Moby-Dick* does not entail modifying *Moby-Dick* itself. So the view that a novel is a mereological sum of embodiments of this novel is false.

Here is yet another consideration against this view. Suppose a novel is a mereological sum of embodiments of this novel. Then whenever one touches/smells/tastes an embodiment of *Moby-Dick*, one touches/smells/tastes a part of *Moby-Dick*. But this consequence does not seem right. Intuitively, no part of a novel can be touched/smelled/tasted—by touching/smelling/tasting an embodiment of this novel or in some other way.

82 fail.

Thus, the view that a novel is a mereological sum of embodiments of this novel is false.

Additionally, the “touch” version of the foregoing objection can be strengthened as follows. It seems uncontroversial that when we touch an object, we often do not (and cannot) touch it in its entirety but, at the same time, can truthfully say that we touch *it*, and not just one or more of its parts. For instance, when we touch an apple, we touch its part, and not all of it. Nevertheless, we can truthfully say that we touch *it*, and not just one of its parts. This suggests that the following principle is true: For all x , to touch x , it is sufficient to touch a part of x . Suppose now that a novel is a mereological sum of its embodiments. Then, given the foregoing principle, *Moby-Dick* can be touched by touching its embodiments. But, as mentioned above, this consequence seems unacceptable.⁸⁸ Given this, the truth of the view that a novel is a mereological sum of embodiments of this novel is far from obvious.

There is little doubt that the foregoing objections cannot be plausibly replied to a proponent of the view that a novel is a mereological sum of embodiments of this novel. Given that, this view cannot be accepted.

⁸⁸Here, it is assumed that mental entities are certain physical states of the brain. If this assumption is rejected and mental entities are considered to be certain nonphysical entities, then the last two objections apply to all versions of the view being discussed except the version according to which a novel is a mereological sum of certain mental embodiments.

6.3 Eliminativism

Thus, neither the abstractionist nor the concretist views stand up to criticism. Given this as well as the fact that the abstract/concrete distinction is exhaustive, one might want to adopt eliminativism, or the view that there are no novels. But can this view be adopted?

Here is a powerful reason to answer this question in the negative. Since eliminativism contradicts one of our strongest artistic intuitions—that novels exist—there is a strong *prima facie* reason against it. At the same time, there are no good reasons in its favor (after all, why think that novels do not exist?). Given this, eliminativism should be rejected. Call this “the Existence Objection.”

Response to the Existence Objection (1). Following Hazlett (2012), a proponent of eliminativism could respond to the Existence Objection as follows. This objection is successful only if there is no real reason in favor of eliminativism. However, there is, in fact, such a reason. Consider repeatable artworks, or artworks that have more than one instance. If such artworks exist, they must be abstracta. Meanwhile, no abstractum has accidental intrinsic properties. So if repeatable artworks exist, then none of them has such properties. However, repeatable artworks do have at least one accidental intrinsic property. So there are no repeatable artworks.⁸⁹ Meanwhile, if this is so, then, since novels are repeatable artworks, eliminativism is true.

This response is based on the thesis that repeatable artworks must be abstracta. But why think that this thesis is true?

⁸⁹See Hazlett (2012, 162).

Here is the answer given by Hazlett (2012): Repeatable artworks must be abstracta, since (a) such artworks have instances, but (b) “no concrete object has instances” (Hazlett, 2012, 163). Is this answer satisfactory? Although Hazlett (2012) is right that repeatable artworks have instances, his claim that no concrete object has instances is puzzling. Given the account of instances of artworks defended in Chapter 2,⁹⁰ there is no conceptual barrier to treating some concreta—say, paintings—as capable of having instances. Perhaps Hazlett (2012) uses some other account of instances of artworks. But he does not explicate this account, and it is unclear what

⁹⁰In a nutshell, this account is as follows:

Instance_e: For all x , x an instance_e of some artwork A if and only if x is either:

- a well-formed instance_e of A —an entity that manifests all the primary properties that must be experienced to fully appreciate A ; or
- a non-well-formed instance_e of A —an entity that (a) manifests sufficiently many, but not all, of the primary properties that must be experienced to fully appreciate A and (b) could, in principle, manifest all sensory kinds of these properties.

Instance_p: For all x , x an instance_e of some artwork A if and only if x is either

- a well-formed instance_e of A —an entity that (a) manifests all the primary properties that must be experienced to fully appreciate A and (b) stands in an appropriate historical-intentional relation to A ; or
- a non-well-formed instance_e of A —an entity that (a) manifests sufficiently many, but not all, of the primary properties that must be experienced to fully appreciate A ; (b) could, in principle, manifest all sensory kinds of these properties; and (c) stands in an appropriate historical-intentional relation to A .

this account could possibly be.

Given what has been said, Hazlett (2012)'s answer cannot be considered satisfactory. At the same time, there are no other alternative answers that could be potentially satisfactory. Thus, the thesis that repeatable artworks must be abstracta is unsubstantiated. Meanwhile, if this is so, then the foregoing response to the Existence Objection is itself unsubstantiated and, hence, fails.

Response to the Existence Objection (2). A proponent of eliminativism could also respond to the Existence Objection with the help of Cameron (2008)'s account of existence-entailing claims. Before examining how exactly she could do that, let us first clarify what this account is.

The key idea behind Cameron (2008)'s account of existence-entailing claims is that at least some of our common sense claims entailing that there are particular entities are not ontologically committing with regard to these entities. This idea can be illustrated with the following example. Consider the claim that there are statues. It is a common sense claim that entails that there are statues. However, this entailment does not necessarily imply that there are, *in fact*, statues. Put otherwise, it is possible that the truth of the claim that there are statues is compatible with there being no statues in the actual world.

In light of what has been said, a natural question arises: By virtue of what are common sense claims that (a) entail that there are certain entities but (b) are not ontologically committing with regard to these entities true? Cameron (2008) answers this question as follows: Such claims are true by virtue of the corresponding facts about the world. Thus, the claim that there are statues is true by virtue of the

fact that the world contains “entities [namely, certain enduring simples] that, for at least some portion of their life, are arranged into a statue shape (and are arranged thus because of the actions of an intentional agent)” (Cameron, 2008, 299).⁹¹

One might wonder whether Cameron (2008)’s account is consistent: For how can (a) a common sense claim entailing that there are certain entities and (b) the claim that there are, in fact, none of these entities be true at the same time? Cameron (2008) replies to this in the following way. A common sense claim entailing that there are certain entities and the claim that there are, in fact, none of these entities belong to different ways of describing the world, or different “languages”: The former claim is part of the language of “common sense” (or “ordinary English,” as Cameron (2008) calls it)—the language that is intended to reflect our intuitive view of the world and does not necessarily describe how things really are; the latter claim is a claim of “*Ontologese*—the language we use to describe how the world is at its fundamental level”⁹² (Cameron, 2008, 300–301). Given that our intuitive view of the world and our view of how the world actually is differ, claims of the language of “common sense” and claims of *Ontologese* do not always have the same truth conditions. A claim of the language of “common sense”—call it “*p*”—is true just in case *p*. However, “*p*” is not necessarily true just in case *p* (where “*p*” is a claim of *Ontologese*); “*p*” can be true even if it is not the case that

⁹¹Although Cameron (2008) endorses the view that the claim that there are statues is true by virtue of the mentioned fact, he leaves open the possibility that this claim could be true by virtue of some other fact.

⁹²According to Cameron (2008), *x* exists at the fundamental level just in case *x* is part of our ontology (or, in other words, has being) (see Cameron (2008, 303)).

p. For example, the claim of the language of “common sense” “There are statues” can be true even if **there are no statues** (where “**there are statues**” is a claim of Ontologese). Thus, the simultaneous truth of (a) a common sense claim entailing that there are certain entities and (b) the claim that there are, in fact, none of these entities is unproblematic (provided that the truth conditions of the former claim do not presuppose the actual existence of these entities, or, in other words, the truth of the claim that **there are these entities**).

Using Cameron (2008)’s account of existence-entailing claims, a proponent of eliminativism could respond to the Existence Objection as follows. The intuition that there are novels amounts to the common sense claim that there are novels. Meanwhile, according to Cameron (2008)’s account of existence-entailing claims, this claim can be true even if there are, in fact, no novels. Thus, the intuition that novels exist does not *necessarily* contradict eliminativism. Does this intuition *actually* contradict eliminativism? The answer is “Yes” only if the common sense claim that there are novels is true, at least in part, by virtue of the fact that novels really exist. However, this claim is not true by virtue of this fact; it is true by virtue of the fact that there really are certain abstract eternally existing linguistic structures that are not novels but that are appropriately related to them.⁹³ Thus, there is no contradiction between the intuition that novels exist and eliminativism. But if this is so, then, contrary to what the Existence Objection states, eliminativism

⁹³This thesis is analogous to Cameron (2008)’s thesis that the common sense claim that there are musical works is true not by virtue of the fact that musical works exist but by virtue of the fact that there are certain abstract eternally existing sound structures.

accords with this intuition.

This response involves the following assumption: (a) There really are certain abstract linguistic structures, but, although these structures are appropriately related to novels, they are not novels; (b) in fact, novels do not exist. Consider now the following alternative to this assumption: (b*) Novels actually exist; (a*) they are identical to certain abstract eternally existing linguistic structures. Which of the mentioned assumptions is preferable? The alternative assumption is supported by the intuition that novels exist. The assumption involved in the response under consideration is supported by the intuition that novels do not exist eternally. Now, it seems a lot easier to accept the claim that novels exist eternally rather than the claim that they do not exist at all. Given this, it is reasonable to hold that the intuition that novels exist overrides the intuition that they do not exist eternally. But if this is so, then the support of the assumption of the response being discussed is weaker than the support of the alternative assumption—and, hence, the question posed above should be answered as follows: The latter assumption should be preferred to the former one. Meanwhile, if this answer is correct, then the response being discussed is unsatisfactory.

Conclusion. Thus, neither response to the Existence Objection is successful. Meanwhile, there are no other potentially satisfactory responses to this objection. Therefore, eliminativism is unacceptable.

6.4 Toward a Satisfactory Ontology of Novels

Our analysis has shown that a novel cannot be identified with:

- a set of embodiments of this novel
- a property
- a pure type
- an initiated type
- a historical individual
- an abstract artifact
- an embodiment of this novel
- a mereological sum of embodiments of this novel

At the same time, according to our analysis, denying the existence of novels is untenable. But what is then a novel, ontologically speaking? What basic sort of entity can it be identified with?

Prima facie, if a novel is a concretum, then it is either a particular embodiment or a mereological sum of embodiments. However, given what has been said above, a novel is neither. Should we then conclude that a novel is not a concretum? No. For there is, I think, a concretum with which a novel can be identified. Here, of course, the following questions arise: (a) What exactly is this concretum? and (b)

Why think that a novel can be identified with it? Answers to these questions are provided in the next chapter.

Chapter 7: An Ontology of Novels

Introduction

In the previous chapter, I have suggested that a novel can be understood as a concretum of a particular kind. My goal in this chapter is to expound this suggestion. I begin by arguing that a novel can be regarded as what I call *a concrete type* (Section 7.1). Next, I examine potential objections to identifying a novel with such a type, arguing that none of these objections stands up to criticism (Section 7.2). Then I turn to an examination of the essential elements of the concrete type to which a novel is identical (Section 7.3). Finally, I summarize the view defended in this chapter (Section 7.4).

7.1 A Novel as a Concrete Type

As has already been said, a novel can be understood as a concretum, or, in other words, an entity that has at least some of the following properties: (a) being in space and/or time, (b) being causally efficacious (= being capable of causally affecting something and of being causally affected by something), and (c) being

modally flexible (= having some intrinsic properties¹ nonessentially²). But what *kind* of concretum is a novel? To answer this question, it is necessary to answer the following questions:

1. What spatial and/or temporal regions does a novel occupy?
2. How does a novel occupy them?
3. What is a novel constituted by?
4. What exactly is the constitution relation between a novel and whatever constitutes this novel?

Let us answer the first of these questions. It is natural to suppose that if an artwork occupies any spatial or temporal regions, then these regions are those that are occupied by instances of this work. Now, given what has been said in Chapter 4, instances of non-visual novels are readings, while instances of visual novels are mereological sums (hereafter: “sums”) of readings and graphic elements.³ In light of what has been said, question (1) can be answered as follows: If a novel is non-visual,

¹As mentioned previously, “an intrinsic property” can be defined as follows: For all x and for all y , x is an intrinsic property of y just in case x is a property possessed by y “in virtue of the way [y] itself, and nothing else, is” (Lewis, 1983, 112). For a detailed analysis of the expression “an intrinsic property” (as well as the related expression “an extrinsic property”), see Weatherson and Marshall (2014).

²As mentioned previously, “an essential property” can be defined as follows: For all x and for all y , x is an essential property of y just in case if x is a property of y and if y is deprived of x , then y goes out of existence.

³Recall that a novel is *non-visual* just in case its primary properties that must be experienced to fully appreciate it are certain (a) sonic properties and (b) experienceable properties that provide

then it occupies those spatiotemporal regions that are occupied by its readings; if a novel is visual, then it occupies those spatiotemporal regions that are occupied by sums of its readings and certain graphic elements.

Let us now turn to question (2) (“How does a novel occupy the spatial and temporal regions it occupies?”) Before answering this question, we first need to clarify how an object *simpliciter* (not necessarily a novel) can occupy a spatial/temporal region. Following Tillman (2011), we can identify the following ways of occupying a spatial/temporal region by an object:

- An object can *pertend* a spatial/temporal region it occupies, where x *pertends* a spatial/temporal region y iff x occupies y and x has a proper part at every proper spatial/temporal subregion of any spatial/temporal region that x occupies (roughly: an extended object with parts at every spatial/temporal region).
 - An object can *span* a spatial/temporal region it occupies, where x *spans* a spatial/temporal region y iff x occupies y and does not have a proper part at any proper spatial/temporal subregion of x 's path, or, in other words, of the fusion of spatial/temporal regions occupied by x (roughly: an extended object without proper parts).
 - An object can be *multiply located* (hereafter: “*multilocated*”) at a spatial/temporal
-
- experiential access to the semantic content; a novel is *visual* just in case its primary properties that must be experienced to fully appreciate it are certain (a) sonic properties, (b) experienceable properties that provide experiential access to the semantic content, and (c) visual properties.

region it occupies, where x is *multilocated* at a spatial/temporal region iff x occupies two disjoint spatial/temporal regions and does not occupy their fusion (roughly: a part-but-not-path occupier of an extended spatial/temporal region).⁴

Does a novel span the spatial/temporal regions it occupies? Suppose the answer is “Yes.” Then, according to the foregoing account of “spanning,” a novel does not have proper parts. However, this consequence seems unacceptable. For suppose it is true. Then it is unclear how novels, being spatiotemporal, can have certain properties that, from an intuitive viewpoint, they doubtless have—say, the properties of being short/long and having one or several chapters (sections, paragraphs). Given what has been said, there is reason to answer the above question in the negative: A novel does not span the spatial/temporal regions it occupies.

In light of the foregoing result, there are four potentially acceptable options:

1. A novel pertends the spatial regions it occupies and is multilocated at the temporal regions it occupies.
2. A novel pertends the temporal regions it occupies and is multilocated at the spatial regions it occupies.
3. A novel pertends both the spatial and temporal regions it occupies.
4. A novel is multilocated at both the spatial and temporal regions it occupies.⁵

⁴See Tillman (2011, 17).

⁵Cf. Zemach (1970)’s “four ontologies.”

Which of these options reflects the actual state of affairs? According to the account of “pertending” provided above, if x pertends a spatial region, then x has a proper spatial part at every proper spatial subregion of any spatial region that x occupies. So if the first option reflects the actual state of affairs, then, given that novels occupy the spatiotemporal regions occupied by their instances (readings or sums of readings and graphic elements), a novel has a proper spatial part at every spatial region occupied by instances of this novel. However, this consequence is problematic. For suppose a novel has a proper spatial part at every spatial region occupied by instances of this novel. Then to completely apprehend a novel, all of its instances must be apprehended. However, as pointed out in Chapter 6,⁶ a complete apprehension of a novel requires apprehending *just one* of its instances (assuming that this instance is well-formed, of course), not all of them.⁷

Consider now the second option. According to the foregoing account of “pertending,” if x pertends a temporal region, then x has a proper temporal part at every proper temporal subregion of any temporal region that x occupies. So if the second option reflects the actual state of affairs, then, given that novels occupy the spatiotemporal regions occupied by their instances, a novel has a proper temporal part at every proper temporal subregion of the temporal regions occupied by instances of this novel. However, this consequence is problematic. Suppose a novel

⁶See Chapter 6, Section 6.2.2.

⁷Furthermore, if a novel has a proper spatial part at every spatial region occupied by instances of this novel, then a complete apprehension of an overwhelming majority of novels is impossible (for, in most cases, some or all of the instances have already ceased to exist). However, this result seems wrong. Intuitively, most novels can be fully apprehended.

has a proper temporal part at every proper temporal subregion of the temporal region occupied by this novel. Then a novel has the property of having the duration equivalent to the sum of the durations of all of its readings. But, *prima facie*, this result is wrong.

Consider now the third option. It entails that a novel pertends both spatial and temporal regions occupied by instances of this novel. However, as is clear from what has been said about the previous two options, this entailment has at least two problematic consequences—(a) that a novel must have a proper spatial part at every spatial region occupied by instances of this novel and (b) that a novel has the property of having the duration equivalent to the sum of the durations of all of its readings.

Finally, let us consider the fourth option. According to it, and given the fact that novels occupy the spatiotemporal regions occupied by their instances, a novel is multilocated at each of the spatial and temporal regions occupied by instances of this novel: In the temporal case, a novel occupies the times at which its instances exist without occupying their fusion; in the spatial case, a novel occupies the spatial regions occupied by its instances, without occupying their fusion. This option does not face the problems of the options discussed above. Furthermore, it does not encounter any other problems. In light of this, the best answer to the question “How does a novel occupy the spatial and temporal regions it occupies?,” I think, is this: A novel occupies the spatial/temporal regions it occupies by being multilocated at these regions.

Let us now address question (3) (“What is a novel constituted by?”). As has

been established above, a novel is multilocalized where each of its instances is wholly located. In light of this, there is good reason to answer question (3) as follows: A novel is constituted by its instances; in particular, a non-visual novel is constituted by its readings, whereas a visual novel is constituted by sums of its readings and graphic elements.

This brings us to question (4) (“What exactly is the constitution relation between a novel and whatever constitutes this novel?”). Let us first say what the constitution relation between a novel and whatever constitutes this novel is not. There is no doubt that this relation can be neither *the relation of identity simpliciter* nor *a combination of the relation of identity and some other relation or relations*. Suppose the constitution relation between a novel and whatever constitutes this novel is the relation of identity. Then a novel is identical either (a) to *each of its instances* or (b) to *a mereological sum of its instances*. Suppose now that the first possibility is true. Then, by the transitivity of identity, instances of a novel must be identical to each other. But, obviously, this result is false. Suppose, on the other hand, that a novel is identical to a mereological sum of its instances. Then: (a) a novel cannot be fully apprehended unless all the instances that compose the mereological sum to which this novel is identical are apprehended, (b) modifying an embodiment contained in the mereological sum to which the corresponding novel is identical entails modifying this novel, and (c) whenever one touches/smells/tastes one of these instances, one touches/smells/tastes a part of this novel. However, as shown in Chapter 6,⁸ these consequences are problematic. The first consequence is

⁸See Chapter 6, Section 6.2.2.

problematic, since a complete apprehension of a novel does not require apprehending *all* of its instances. The reason why the second consequence is problematic is that in an overwhelming majority of cases, modifying an embodiment of a novel does not entail modifying this novel. Finally, the third consequence is problematic because it seems impossible to touch/smell/taste a part of a novel (as opposed to a physical instance of this novel)—by touching/smelling/tasting its instances or by any other means.

Now, suppose that the constitution relation between a novel and whatever constitutes this novel is a combination of the relation of identity and some other relation or relations. Then the only plausible option seems to be this: A novel is identical to one particular instance of this novel and is related to the other instances in a way that does not presuppose being identical to them. Is this option tenable? As is clear from the previous chapter,⁹ there is good reason against identifying a novel with a particular instance of this novel. First of all, there is no satisfactory explanation as to why a novel is identical to the instance to which it is considered to be identical, and not some other instance of this novel. Furthermore, identifying a novel with a particular instance of this novel contradicts our intuitions regarding destroying, modifying, and experiencing novels. Given this, the question posed above should be answered in the negative.

Thus, the constitution relation between novels and their instances cannot be the relation of identity *simpliciter*; nor can it be a combination of the relation of identity and some other relation or relations. Now, what about treating the con-

⁹See Chapter 6, Section 6.2.1.

stitution relation between novels and their instances as *the relation that a whole bears to its proper parts*? Is this suggestion tenable? No. If the constitution relation between novels and their instances were the relation that a whole bears to its proper parts, then to fully appreciate a novel, we would have to appreciate all of its instances. However, this consequence is problematic, since, as already mentioned, a complete apprehension of a novel requires apprehending *only one* of its instances (assuming, of course, that this instance is well-formed), not all of them.

So what is then the constitution relation between a novel and its instances? As is clear from what has been said, a novel (a) is wholly located where each of its instances is wholly located but (b) is neither identical to any one of these instances or their mereological sum nor related to them in the way in which a whole is related to its proper parts. In light of this, the foregoing question can, I think, be answered as follows: The constitution relation that holds between a novel and its instances is *the relation of coincidence*, where the latter is defined in the following way:

Coincidence For all x and for all y , x coincides with y just in case x (a) is wholly located where each y is wholly located but (b) is neither identical to y or a mereological sum that contains y nor related to y in the way in which a whole is related to its proper part.

It is worth stressing that the relation of coincidence, as defined above, is neither strange nor unnatural. In fact, this relation accords perfectly with common sense. Consider a clay statue. If someone smashes it into multiple pieces, it will be destroyed, but the clay that constitutes it will survive. So the statue is not identical

to the clay. Furthermore, the statue does not seem to be identical to a mereological sum that contains the clay (after all, what could this sum be?). Finally, the statue is not related to the clay as a whole is related to its proper part. For suppose the opposite is true. Then if the clay is eliminated, there must remain something else—something that could serve as another proper part of the statue (for, otherwise, the clay cannot be a *proper* part of the statue). It is clear, however, that nothing will be left if the clay is eliminated.

Thus, the statue is neither identical to the clay or a mereological sum containing the clay nor related to the clay as a whole is related to its proper part. At the same time, the statue is wholly located where the clay is wholly located. Thus, in this case, the conditions of **Coincidence** are satisfied. As a result, it is quite natural to characterize the statue as *coincident* with the clay.

Let us now return to the question posed at the beginning of this section: What sort of concretum is a novel? As is clear from what has been said above, novels fall under the category of concreta that (a) are capable of spatiotemporal multilocation (that is, can occupy several disjoint spatiotemporal regions without occupying their fusion) and (b) coincide with their instances (that is, are wholly located where each of their instances is wholly located but are neither identical to any of them or a mereological sum of them nor related to them in the way in which a whole is related to its proper parts). Let us call this category “concrete types.” In light of this, the foregoing question can be answered as follows: A novel is a concrete type.¹⁰

Before proceeding further, three remarks are worth making. First, the term

¹⁰A similar view on the nature of properties is advocated by Armstrong (1978).

“type” is used here in a way that does not commit its referents to being abstract.¹¹ Second, the category of concrete types does not imply that all existent types are concrete; it is compatible with the idea that certain entities—such as various mathematical entities like numbers, functions, and pure sets—could be abstract types. Finally, the category of concrete types should be distinguished from the category of types introduced by Zemach (1970, 1992). This is not to say, of course, that these categories do not have anything in common. According to Zemach (1970, 1992), types—that is, entities such as the Taxpayer, the Lion, the American Woman, and the Letter A—are not abstract; rather, they “are *particulars* recurring both at many different times and in many different places” (Zemach, 1992, 7):

Mr. Jones is a material thing, a particular, although he can be, *all* of him, in two distinct spatiotemporal locations (e.g., in his office at 9 AM and at home at 8 PM). Similarly, the type-entity The Cat is, all of it, in many distinct spatiotemporal locations. . . It is the same type-entity The Cat which is seen first on the mat, then on the couch, and at the same time climbing a tree in the yard. (Zemach, 1992, 7)

Thus, Zemach’s category of types implies that whatever falls under this category is concrete and can be multilocalized at both spatial and temporal regions. Meanwhile, that is exactly what is implied by the category of concrete types. So in this regard, these categories are similar. Yet there is a crucial difference between them. The difference concerns the interpretation of the relation that holds between an entity

¹¹Although such ontologically neutral usage is not orthodox, it is not uncommon. Thus, it is adopted by Peirce (1931), Davies (2012), and Wollheim (1971), among others.

falling under each of the corresponding categories and whatever constitutes this entity. According to Zemach (1992)'s category, the relation that holds between types and what constitutes them is that of *nesting*: Types are nested in what constitutes them (where y is *nested* in x iff every essential property of x ¹² is also an essential property of y (Zemach, 1992, 148)). On the other hand, the category of concrete types does not require an entity falling under this category to be *nested* in its constituents; an entity can fall under this category even if it is not nested in its constituents—provided that it is *coincident* with them.

7.2 Some Objections

Let us now examine possible objections to the view that novels are concrete types.

Objection 1. Consider the following principle:

Localization Principle: For all x , if x is concrete, then x cannot be spatially multi-located.

This principle is supported by our intuition: Intuitively, concreta—both objects (tables, people, trees, animals, etc.) and events (dances, musical performances, fights, etc.)—cannot be wholly located at two distinct places at the same time.¹³

¹²Here, by “an essential property of x ” is meant a property that x has “at any time, place, or possible world where $[x]$ exists” (Zemach, 1992, 148).

¹³Thus, suppose we see John wholly present at two different places at the same time. In this case, most likely, we will not believe what we see; that is, we will not believe that John is, *in fact*, at two different places at the same time.

So there is a good reason in favor of the **Localization Principle**. Meanwhile, this principle entails that novels cannot be concrete types. For suppose it is true. Then there can be no object that is concrete and capable of spatial multilocation. But concrete types are, by definition, concrete and capable of spatial multilocation. So there can be no such types. But if this is the case, then, of course, no novel can be a concrete type.

Response. Let us distinguish between two kinds of possibility: logical and physical. Logical possibility can be defined as follows:

Logical Possibility : For all x , x is logically possible iff x does not violate the law of non-contradiction ($\neg(A \wedge \neg A)$).

Here is the definition of physical possibility:

Physical Possibility : For all x , x is physically possible iff x does not violate any physical laws.

Which of these kinds of possibility is used in the **Localization Principle**? Suppose it is logical possibility. Then the **Localization Principle** entails that if x is concrete and spatially multilocalized, then x violates the law of non-contradiction. However, this entailment is false. Suppose that, with the help of a time machine, John, currently located in 2017, travels to the past—say, to 2011. Then, in 2011, there are John’s younger body and John’s older body. Now, what about John himself? What is his spatial location in 2011?¹⁴

¹⁴It can be objected that this example is meaningless, as backwards time travel is logically impossible. Is this objection successful? The answer is “Yes” only if there is a *compelling* argument

There is little doubt that in 2011, John cannot be spatially located just where his younger body is located. For, in that case, he would not be spatially located where his older body is spatially located in 2017, which contradicts one of our assumptions.

One might say that in 2011, John is spatially located solely where his older body is spatially located. But this answer is problematic. If it is true, then the person who is wholly spatially located where John's younger body is spatially located is not John—and, hence, in 2017, it is false that John was younger in 2011. However, this consequence is doubtless false.

Another possible answer is that in 2011, John is (a), in part, spatially located where his younger body is spatially located and (b), in part, spatially located where his older body is spatially located. But this answer is also problematic. For suppose it is correct. Then in 2011, John amounts to the mereological sum of his instances—his younger body and his older body—(or to whatever is constituted by these instances) and, as a result, has two completely independent minds. However, no person (as opposed, perhaps, to a human being) can have two completely independent minds.

Presumably, the best answer to the question “What is John's spatial location against the logical possibility of backwards time travel. However, there seems to be no such argument (for an exposition of the most promising arguments against the logical possibility of backwards time travel—including the arguments put forward by Lewis (1976)—as well as a critique of these arguments, see Smith (2016)). Given this, the question posed above can, I think, be answered in the negative.

in 2011?” is that in 2011, John is spatially multilocated: He occupies the spatial regions occupied by each of his instances—namely, his younger and older bodies—and does not occupy their fusion. (Note that this answer does not face the problem of the previous answer. According to that problem, the previous answer has an unpalatable consequence—that in 2011, John has two completely independent minds. However, the answer being discussed does not have this consequence. According to this answer, it is not true that in 2011, John is, (a) *in part*, spatially located where his younger body is spatially located, and (b), *in part*, spatially located where his older body is spatially located. Rather, this answer states that in 2011, John is *multilocated*, or, in other words, is *wholly* present both where his younger body is spatially located and where his older body is spatially located. Meanwhile, if this is so, then, according to this answer, in 2011, John has a single mind (which is wholly present where each of John’s bodies is wholly present).)

Taking into account what has been said, let us agree that in 2011, John is multilocated. Is the law of non-contradiction violated in this case? There seems no real reason to think so. John’s multilocation is easy to imagine. So, assuming that conceivability entails logical possibility,¹⁵ John’s multilocation is possible and, hence, non-contradictory. Furthermore, presumably, to get a contradiction in the given case, we should accept that in 2011, John does not occupy either the spatial region occupied by his younger body, or the spatial region occupied by his older body, or both of these regions. But we do not have to accept that. The fact that

¹⁵This assumption is, of course, controversial. For reasons to endorse it as well as responses to potential objections, see, e.g., Chalmers (1996, 2002, 2009).

in 2011, John is multilocalized at the spatial regions occupied by his younger and older bodies does not entail that he is not spatially located at each of these regions. Meanwhile, this is the only fact that is relevant in the given case.

Thus, in 2011, John does not violate the law of non-contradiction. Meanwhile, as already mentioned, John is multilocalized in 2011. Furthermore, there is no doubt that he is concrete. So the implication of the **Localization Principle**—that if x is concrete and spatially multilocalized, then x violates the law of non-contradiction—is false. As a result, the **Localization Principle** is itself false.

Suppose now that the **Localization Principle** employs the concept of physical possibility. Is this principle true in this case? If the answer is “Yes,” then in 2011, John—as well as any other multilocalized concretum—violates at least some physical laws. But there seems no real reason to think that in 2011, he violates any such laws. Meanwhile, the **Localization Principle** entails that if a concretum is multilocalized, then this concretum violates at least some physical laws. In light of what has been said, there is good reason to consider this principle false.

One could defend the **Localization Principle** employing the physical concept of possibility (hereafter: “the **Localization Principle_p**”) by saying that this principle is *intuitively* correct and, hence, should be regarded as true by default unless a satisfactory reason against this principle is provided. This defense is successful only if there is no satisfactory reason against the **Localization Principle_p**. However, there is, in fact, such a reason. Suppose the **Localization Principle_p** is true. Then any spatially multilocalized concretum must violate at least some physical laws. Put otherwise, the physical world *cannot* contain spatially multilocalized

concreta. However, this consequence can be rejected. It is reasonable to hold that the physical world can—and, in fact, does—contain spatially multilocated concreta. Consider, for instance, the plague. It is a concretum present in the physical world. But it may well be treated as capable of spatial multilocation. Thus, imagine that at the same time, some disjoint spatial regions *A* and *B* are affected by the plague. Where is the plague in this case? It seems odd to say that a *part* of the plague is spatially located at *A* and a *part* of the plague is spatially located at *B*. Intuitively, the plague is *wholly* spatially located at both *A* and *B*—and, hence, is multilocated.

The same kind of reasoning that has been applied to the plague is applicable to many other concreta, for example, water, gold, beer, steel, wood, air, etc. Thus, it is reasonable to hold that the physical world does contain spatially multilocated concreta. But if this is so, then, given that the **Localization Principle_p** entails that no spatially multilocated concretum can be present in the physical world, this principle can be rejected.

Objection 2. The view that novels are concrete types forces us to add a radically new ontological category to our ontology—that of concrete types. However, adding this category to our ontology is unjustified.

Response. Objection 2 assumes that the category of concrete types is *radically* new. Is this assumption true? Recall that a concrete type is an entity that is capable of spatiotemporal multilocation and is coincident with its instances (where *x* is coincident with *y* just in case *x* (a) is wholly located where what constitutes *x* is wholly located but (b) is neither identical to what constitutes *x* nor related to *x* as a whole is related to its proper parts). Given this, a concrete type has the following

features. First, it can be spatiotemporally multilocated. Second, it is wholly located where what constitutes it is wholly located. Third, it is neither identical to any of its instances or their sum nor related to them in the way in which a whole is related to its proper parts.

Consider now ordinary concreta. They share most of the mentioned features. First, ordinary concreta are capable of *temporal* multilocation. Thus, my desk has been temporally multilocated since it came into existence: Since that time, it has occupied several disjoint temporal regions without occupying their fusion (for example, it was wholly located at a particular moment on March 12, 2012, when it was created, and then at all subsequent moments until the present moment). Second, ordinary concreta are wholly located where what constitutes them is wholly located. A statue, which is constituted by a lump of clay, is wholly located where this lump is wholly located. Likewise, my desk, which is constituted by some wooden matter, is wholly located where this matter is wholly located. Third, ordinary concreta are neither identical to what constitutes them nor related to it in the way in which a whole is related to its proper parts. Thus, a cat is not identical to the matter that constitutes it, since this matter existed before the cat came into existence and will exist after the cat dies. Nor is a cat related to its constitutive matter as a whole is related to its proper part. If a cat were related to this matter that way, then after eliminating the matter, some other part of the cat would remain. But, surely, no part of a cat will be left if the cat's matter is eliminated.

Given the mentioned similarities between concrete types and ordinary concreta, it seems unreasonable to characterize the category of concrete types as *radically*

neasatiw. Meanwhile, if this is so, then there is reason against Objection 2.

Here, however, one might object as follows. Although ordinary concreta and concrete types are similar, there are at least two essential differences between them. First, ordinary concreta can have just one instance, whereas concrete types can have more than one instance. Second, ordinary concreta are incapable of spatial multilocation, whereas concrete types can be spatially multilocated. In light of this, contrary to what the above response states, characterizing the category of concrete types as *radically* new is justified.

Is the foregoing objection successful? Assuming that backwards time travel is possible,¹⁶ every ordinary concretum—a statue, a desk, a cat, a person, etc.—can, in principle, be transported into the past. Meanwhile, if this is so, then every such concretum (a) can be spatially multilocated, or, in other words, can be wholly spatially located where its younger instance is spatially located and wholly spatially located where its older instance is spatially located, and (b) can have more than one instance. So if backwards time travel is possible, then the mentioned differences between ordinary concreta and concrete types are nonessential—and, hence, the objection being discussed fails.

The above response depends on the possibility of backwards time travel. But what if such time travel is impossible? Can Objection 2 be defused in this case?

¹⁶This assumption is, of course, controversial. However, it is not untenable. True, there are a number of seemingly strong arguments against the possibility of backwards time travel (one of them is the well-known Grandfather Paradox (Lewis, 1976)). But there is good reason to hold that each of these arguments can be plausibly defused (see Smith (2016)).

The answer to this question, I think, is “Yes.” Objection 2 assumes that adding the category of concrete types to our ontology is unjustified. Is this assumption true? Perhaps it would be true if by adding the category of concrete types to our ontology, we could account *solely* for the ontological status of novels. But, in fact, doing this can help us elucidate the ontological status of many other concreta—such as the plague, gold, water, air, beer, and sand. In light of this, as well as the fact that the category of concrete types is relatively transparent (especially, compared to the rather obscure categories of initiated types, historical individuals, and abstract artifacts¹⁷), adding this category to our ontology is justified. Meanwhile, if this is so, then Objection 2 contains a false assumption and, hence, fails.

Objection 3. Suppose there is a reading of some novel, say, *Moby-Dick*, that came into existence at 11 am and went out of existence at 9 pm on June 13, 2017. If novels are concrete types, then *Moby-Dick* is constituted by this reading. Meanwhile, any concretum inherits all the properties of what constitutes this concretum. So if novels are concrete types, then *Moby-Dick* inherits the properties of the foregoing reading, including the properties of having come into existence at 11 am on June 13, 2017 and of having gone out of existence at 9 pm on June 13, 2017. But, of course, *Moby-Dick* did not come into existence at 11 am on June 13, 2017; nor did it go out of existence at 9 pm on June 13, 2017. So novels are not concrete types.

Response. Objection 3 assumes that any concretum inherits all the properties of what constitutes this concretum. Put otherwise, according to this objection, the following principle is true:

¹⁷For an analysis of these categories, see Chapter 6.

Inheritance: For all x and for all y , if (a) x and y are concreta and (b) x is constituted by y , then x inherits y 's properties.

But is **Inheritance** true? Consider a clay statue. It is constituted by a lump of clay. So if **Inheritance** is true, then, since both the statue and the clay are concrete, the statue must inherit all of the properties of the clay. However, in fact, the statue does not inherit all of these properties. Thus, the clay came into existence before the statue came into existence. Furthermore, the clay can survive the destruction of the statue. Thus, **Inheritance** is false. But if this is so, then the objection being discussed fails.

Here, one might object as follows. Suppose a novel does not inherit its properties from what constitutes it—its instances. But why, then, does a novel have the properties that it has? Until a proponent of the view that novels are concrete types provides a satisfactory answer to this question, this view cannot be considered acceptable.

The foregoing objection assumes that the view that novels are concrete types cannot be accepted unless there is an explanation of why novels have the properties that they have. However, this assumption seems too strong. Clearly, an explanation of why novels have the properties that they have is necessary in order for the view that novels are concrete types to be *complete*. However, the absence of such an explanation is not a reason to consider this view *unacceptable*. After all, we may know what nature an entity has without knowing *why* this entity has this nature.

But let us agree, for the sake of argument, that the foregoing assumption is

correct: The view that novels are concrete types cannot be accepted unless there is an explanation of why novels have the properties they have. Can the objection being discussed be defused in this case? Yes. This objection is successful only if there is no way to plausibly explain why novels have the properties they have. However, there is, in fact, such a way. According to it, novels inherit their properties from the corresponding canonic embodiments (say, inscriptions), or embodiments that manifest or encode the correct textual properties of these novels.¹⁸ Thus, *War and Peace* inherits its properties—having the title “War and Peace,” being in Russian, being authored by Leo Tolstoy, etc.—from a canonic embodiment of *War and Peace*.¹⁹

Objection 4. Suppose novels are concrete types. Then the existence of a novel depends on the existence of a reading of this novel: A novel exists when, and only when, its reading exists. Meanwhile, if this is so, then (a) there can be no novel that has never been read aloud and (b) a novel has discontinuous existence: It comes into existence when one starts reading it aloud, goes out of existence when one stops reading it aloud, and comes into existence again when one resumes reading it aloud.²⁰ But are (a) and (b) true?

¹⁸For a defense of a similar explanation concerning musical works, see Tillman and Spencer (2012, 257–258).

¹⁹The thesis that novels inherit their properties from the corresponding canonic embodiments does not imply that novels inherit *all* of the properties of these embodiments. So although *War and Peace* does inherit certain properties from its canonic embodiment, it does not inherit all of these properties from this embodiment (for instance, it does not inherit the property of being made of something).

²⁰That (a) follows from the thesis that a novel exists when, and only when, its reading exists is

Suppose someone creates a manuscript of some novel *N*. Suppose next that *N* has never been read aloud. Does *N* exist? Yes—or so it seems. But if *N* exists, then there can be a novel that has never been read aloud—and, as a result, (a) is false.

Now what about (b)? Suppose there is a manuscript of some novel *M*. Suppose next that *M* has been read aloud but is not being read aloud now. Does *M* exist *now*? Again, the answer seems to be “Yes.” Meanwhile, if this answer is correct, then a novel does not necessarily go out of existence when one stops reading it aloud—and, hence, (b) is false.

Thus, both (a) and (b) are false. But if this is so, then novels cannot be concrete types.

Response. Objection 4 assumes that the view that novels are concrete types has two problematic consequences—that (a) there can be no novel that has never been read aloud and that (b) a novel comes into existence when one starts reading it aloud, goes out of existence when one stops reading it aloud, and comes into existence again when one resumes reading it aloud. Is this assumption true?

Before addressing this question, let us first clarify the expression “to exist.”

uncontroversial. Suppose this thesis is true. Then a novel exists only when its reading exists. So if a novel has never been read aloud, it has never existed. Similarly, there is no doubt that (b) follows from the thesis that a novel exists when, and only when, its reading exists. Suppose this thesis is true. Then each time a reading of a novel occurs, this novel exists, and each time a reading of a novel does not occur, this novel does not exist. As a result, a novel comes into existence when one starts reading it aloud, goes out of existence when one stops reading it aloud, and comes into existence again when one resumes reading it aloud.

As mentioned in the previous chapter, there are two senses of this expression: the presentist and the ontological. According to the presentist sense, x exists (hereafter: “exists_{pr}”) at t just in case x is present at t ; according to the ontological sense, x exists (hereafter: “exists_o”) at t just in case x is in the domain of our most unrestricted quantifier.²¹ But besides these senses, there is another sense of the expression “to exist”—the possibilist. According to this sense, x exists (hereafter: “exists_{po}”) at t just in case there is a (correct) embodiment (that is, an entity that manifests or encodes the relevant properties of the object it is an embodiment of)—say, a set of instructions or a recording—that can be used to bring x into existence_{pr}. Thus, Beethoven’s Symphony No. 7 exists_{po} now *even if it is not being performed*—since there is an embodiment of it (say, a correct copy of its score) that can be used to bring it into existence_{pr}. At the same time, if all embodiments of Beethoven’s Symphony No. 7 were destroyed (say, as a result of some global catastrophe), then it would not exist_{po}.²²

²¹As mentioned in the previous chapter, in light of Quine (1948)’s view on the univocality of “exist,” it might be more accurate to speak of different uses, rather than senses, of “exist.” However, for the sake of convenience, I speak of different senses of “exist.” If the reader finds my talk of different senses of “exist” inappropriate, she is free to paraphrase it in terms of uses of “exist.”

²²Here, one could object as follows. Take a (healthy) human sperm and a (healthy) unfertilized human egg. If the sperm fertilizes the egg, then, under appropriate circumstances, a particular human being—call this being “John”—will come into existence_{pr}. Now, can it be said that John exists—in the possibilist, or any other, sense—*before* the fertilization occurs? Of course not! However, the foregoing account of “exist” forces us to give an affirmative answer to this question. For, on this account, an entity exists_{po} if there is an embodiment of this entity that can be

So does the view that novels are concrete types entail that (a) there can be no novel that has never been read aloud and that (b) a novel comes into existence when one starts reading it aloud, goes out of existence when one stops reading it aloud, and comes into existence again when one resumes reading it aloud? Let us first consider whether this view entails that there can be no novel that has never been read aloud. Clearly, it entails this only if a novel that (a) is a concrete type and (b) has never been read aloud cannot exist. However, in fact, such a novel can exist—*if* the expression “to exist” is used in the possibilist or the ontological sense. Take some novel N . Suppose N is a concrete type. Suppose next that N has never been read aloud. Finally, suppose that there is an inscription that can be used to bring N into existence_{pr}. Then N exists_{po}. For, according to the definition of “to exist_{po},” if there is an embodiment that can be used to bring x into existence_{pr}, then x exists_{po}. And, by assumption, there is, in fact, at least one embodiment that can used to bring this entity into existence_{pr}. Meanwhile, in the case under consideration, there is an embodiment of John—the mereological sum of the sperm and the unfertilized egg—and this embodiment can be used to bring John into existence_{pr}.

This objection assumes that the mereological sum of the sperm and the unfertilized egg is an embodiment of John. Is this assumption true? If the answer is “Yes,” then the mereological sum of the sperm and the unfertilized egg must encode or manifest the relevant properties of John (qua a particular biological organism). However, this sum neither encodes nor manifests these properties. It would do that only if it contained John’s DNA. Yet, unlike the fertilized egg, it does not contain his DNA (the DNA’s contained in the sperm and the unfertilized egg are essentially different from John’s DNA). Thus, the foregoing question should be answered in the negative. Meanwhile, if this is so, the the objection being discussed contains a false assumption and, hence, fails.

be used to read N aloud and, hence, to bring N into existence_{pr}.

Furthermore, N exists_o. Since it exists_{po}, it is in the domain of our most unrestricted quantifier. And according to the definition of “to exist_o,” if x is in this domain, then x exists_o.

Thus, the claim that unread novels cannot exist is not a consequence of the view that novels are concrete types. Now, what about the claim that a novel comes into existence when one starts reading it aloud, goes out of existence when one stops reading it aloud, and comes into existence again when one resumes reading it aloud? Is this claim entailed by the view that novels are concrete types? The answer is “Yes” only if a novel that (a) is a concrete type and (b) is not being read aloud cannot exist. But is it true that such a novel cannot exist? Take some novel M . Suppose M is a concrete type. Suppose next that M is not being read aloud. Finally, suppose that there is an inscription of M . Then M exists_{po}. For, by assumption, there is at least one embodiment that can be used to read M aloud and, hence, bring M into existence_{pr}. And, as mentioned already, if there is an embodiment that can be used to bring x into existence_{pr}, then x exists_{po}.

Furthermore, M exists_o. The explanation of why this is so is analogous to the foregoing explanation of why N exists_o.

Given what has been said, the question posed above should be answered in the negative: That a novel comes into existence when one starts reading it aloud, goes out of existence when one stops reading it aloud, and comes into existence again when one resumes reading it aloud is not entailed by the view that novels are concrete types.

Thus, contrary to what Objection 4 assumes, the view that novels are concrete types entails neither that (a) there can be no novel that has never been read aloud nor that (b) a novel comes into existence when one starts reading it aloud, goes out of existence when one stops reading it aloud, and comes into existence again when one resumes reading it aloud. As a result, this objection fails.

A proponent of Objection 4 could respond as follows. Perhaps the foregoing response to Objection 4 is successful. But let us assume that the sense of “to exist” employed in Objection 4 is presentist. In this case, the foregoing response does not apply, and, at the same time, there seems no other potentially plausible way to respond to Objection 4.

Is this defense of Objection 4 successful? Suppose Objection 4 employs the presentist sense of the expression “to exist.” Then according to this objection, it is false (a) that a novel that has never been read aloud cannot exist_{pr} and (b) that a novel comes into existence_{pr} when one starts reading it aloud, goes out of existence_{pr} when one stops reading it aloud, and comes into existence_{pr} again when one resumes reading it aloud. However, the falsity of (a) and (b) is far from obvious. If a novel has never been read aloud, then has it ever been present, assuming that an inscription or some other embodiment other than a reading of this novel has been present? Answering this question in the negative seems right, especially given the fact that being present does not amount to being existent *simpliciter*. Meanwhile, if novels that have never been read aloud have never been present, then, since, by definition, being present is equivalent to being existent_{pr}, they cannot exist_{pr}—and, as a result, (a) is true.

Consider now (b). It does not seem unnatural to hold that a novel becomes present when one starts reading it aloud, disappears (though does not go out of existence *simpliciter*) when one stops reading it, and becomes present again when one starts reading it aloud again. Meanwhile, if a novel becomes present when one starts reading it aloud, disappears (though does not go out of existence *simpliciter*) when one stops reading it, and becomes present again when one starts reading it aloud, then, since, by definition, being present is equivalent to being existent_{pr} , (b) is true.

Thus, if the presentist sense of the expression “to exist” is used in Objection 4, then this objection involves a questionable claim—that (a) a novel that has never been read aloud can exist_{pr} and (b) a novel does not come into existence_{pr} when one starts reading it aloud, go out of existence_{pr} when one stops reading it aloud, and come into existence_{pr} again when one resumes reading it aloud. Meanwhile, according to the response being discussed, this sense is, in fact, used in Objection 4. So this response cannot be considered successful.

7.3 The Essential Elements of the Concrete Type to Which a Novel Is Identical

Thus, a novel can be understood as a concrete type—a concretum that is coincident with each of its instances, or, in other words, that (a) is wholly located where each of its instances is wholly located but (b) is neither identical to any of its instances or a sum of them nor related to them in the way in which a whole

is related to its proper parts. Now, a natural question arises: What *exactly* is the concrete type to which a novel is identical? In particular, what are the essential elements that compose this type?

Let us first clarify the expression “an essential element of x .” This expression can be defined as follows: For all x and for all y , x is an essential element of y just in case if y is deprived of x , then y goes out of existence. Thus, suppose that a particular shape is an essential element of some figure F_1 . Then, according to the foregoing definition, if F_1 is deprived of this shape, then F_1 goes out of existence. On the other hand, suppose that some figure F_2 is deprived of a particular shape but does not go out of existence. Then the foregoing definition entails that this shape is not an essential element of F_2 .

So what essential elements does the concrete type to which a novel is identical have? As demonstrated in Chapter 3,²³ the artistic value of a novel depends on the sound of this novel. So the sonic properties of a novel are essential to this novel: If a novel loses some of its sonic properties, then (even if these properties are minor ones) it goes out of existence. Meanwhile, if the sonic properties of a novel are essential to it, then the sonic element is an essential element of a novel and, hence, of the concrete type to which this novel is identical.

Thus, one of the essential elements of the concrete type to which a novel is identical is *the sonic element*. Another such element is *the semantic one*. There is no doubt that the artistic value of a novel depends primarily on what the novel tells, or, in other words, on its content. Meanwhile, the content of a novel is constituted

²³See Chapter 3, Section 3.1.2.

by certain semantic properties. So the artistic value of a novel depends on the semantic properties of this novel—and, as a result, such properties are essential to novels. Meanwhile, if this is so, then the semantic element is an essential element of a novel and, hence, of the concrete type to which this novel is identical.

Are there any other elements that are essential to the concrete type to which a novel is identical? Consider the following quote from *Moby-Dick*:

S: “I thought so. All right; take a seat. Supper?—you want supper?
Supper’ll be ready directly.” (Melville, 1922, 18)

Suppose now that the first instance of the question mark is deleted, while the second instance is replaced with a period:

S_{modified}: “I thought so. All right; take a seat. Supper—you want
supper. Supper’ll be ready directly.”

Clearly, the foregoing change causes S and, hence, *Moby-Dick* to acquire new semantic and sonic properties (S_{modified} does not mean the same as S and has a different sound). Meanwhile, this change is purely *syntactic*. So the semantic and sonic properties of *Moby-Dick* depend on certain syntactic properties.

Now, arguments analogous to the argument given above can be made with regard to any novel. Therefore, the syntactic properties of a novel determine some of its semantic and sonic properties. Meanwhile, as already mentioned, the semantic and sonic properties of a novel determine, in part, the artistic value of this novel. So the artistic value of a novel depends on the syntactic properties of this novel. But then the syntactic properties of a novel are essential to this novel. As a result,

the syntactic element is an essential element of a novel and, hence, of the concrete type to which this novel is identical.

Now, what about *the visual element*? Is it an essential element of the concrete type to which a novel is identical? An answer to this question depends on whether the novel in question is visual or non-visual. Consider a visual novel, or a novel whose primary properties²⁴ that must be experienced to fully appreciate it are certain (a) sonic properties, (b) experienceable properties that provide experiential access to the semantic content, and (c) visual properties. As pointed out in Chapter 4,²⁵ the artistic value of such a novel depends on certain visual properties. As a result, at least some visual properties of a visual novel are essential to it. Meanwhile, if this is so, then the visual element is an essential element of a visual novel and, hence, of the concrete type to which this novel is identical.

At the same time, the visual element is not an essential element of the concrete type to which a non-visual novel is identical. Consider a non-visual novel, or a novel whose primary properties that must be experienced to fully appreciate this novel amount to certain (a) sonic properties and (b) experienceable properties that provide experiential access to the semantic content. As noted in Chapter 4,²⁶ the artistic value of such a novel does not depend on any visual properties. Meanwhile, *x* is an

²⁴As might be recalled from Chapter 2, primary properties are those through which the primary content—or, in other words, the set of “those contentful properties that may be the ground of other contentful properties but which are not themselves grounded in contentful properties” (Davies, 2010, 411)—of an artwork is articulated.

²⁵See Chapter 4, Sections 4.1 and 4.2.

²⁶See Chapter 4, Sections 4.1 and 4.2.

essential element of an artwork only if the artistic value of this work depends on the properties possessed by this element. So, given that visual properties are the only relevant properties possessed by the visual element, this element is not an essential element of a non-visual novel and, therefore, of the concrete type to which such a novel is identical.

Besides the essential elements mentioned above, the concrete type to which a novel is identical has another essential element—the *contextualist one*, or the one concerned with certain essential contextual properties, primarily the property of being created by a particular individual or individuals in a particular cultural context. Why think that this is the case? This question can be answered with the help of the following variation of the thought experiment offered by Currie (1991).²⁷ Consider Jane Austen’s *Northanger Abbey*. It was written “in 1803 as a burlesque on the Gothic novel” (Currie, 1991, 328). Imagine now that “a hitherto unknown manuscript by Anne Radcliffe, entitled *Northanger Abbey* (circa 1793), and word for word the same as Austen’s, turns up in the attic, that we conclude (never mind on what evidence) that this is in fact a coincidence, that Austen had no knowledge of Radcliffe’s work, and that, far from being a satire, Radcliffe’s *Abbey* was meant as a serious contribution to the genre” (Currie, 1991, 328). Given what has been said, there is little doubt that Austen’s and Radcliffe’s novels have different artistic values. But how can this difference be explained? Presumably, a difference between the artistic values of two novels can be explained by appealing to the difference between the syntactic, sonic, semantic, visual, and/or contextualist elements of one

²⁷See Currie (1991, 328).

of these novels and the respective elements of the other novel. However, in the case under consideration, the syntactic, sonic, semantic, and visual elements of Austen's novel do not differ from the respective elements of Radcliffe's novel—for, (a) by assumption, these novels have the same text and, hence, the same syntactic, sonic, semantic, and visual properties, and (b) if two novels have the same syntactic, sonic, semantic, and visual properties, then the syntactic, sonic, semantic, and visual elements of one of these novels and the respective elements of the other novel are the same. Thus, the fact that Austen's and Radcliffe's novels have different artistic values can only be explained by appealing to the difference between the contextualist elements of these novels. In light of this, the answer to the foregoing question is as follows: Austen's and Radcliffe's novels have different artistic values because the contextualist element of Austen's novel and the contextualist element of Radcliffe's novel are different.

Here, one might object that the foregoing answer is unsatisfactory, since it involves a false claim—that the contextualist elements of Austen's and Radcliffe's novels are different. Is this objection successful? Austen's novel has the property of being a satire on the genre of Gothic novel and does not have the property of being a serious contribution to this genre, whereas Radcliffe's novel has the latter property and does not have the former one. Meanwhile, both of these properties are contextual, as they are grounded in the context in which Austen's and Radcliffe's novels came into existence, rather than in the non-contextual, structural features of these novels. Thus, contrary to what the objection being discussed states, the contextualist element of Austen's novel and the contextualist element of Radcliffe's

novel are different: The former has the property of being a satire on the genre of Gothic novel and does not have the property of being a serious contribution to this genre, whereas the latter has the latter property and does not have the former one.²⁸ As a result, this objection fails.

Thus, Austen's and Radcliffe's novels have different artistic values because the contextualist element of Austen's novel and the contextualist element of Radcliffe's novel are different. But if this is so, then the artistic values of these novels depend on certain contextualist properties—and so these properties are essential to these novels. Meanwhile, if the contextualist properties of a novel are essential to it, then the contextualist element is an essential element of this novel. As a result, the contextualist element of Austen's novel and the contextualist element of Radcliffe's novel are essential elements of these novels and, hence, of the concrete types to which these novels are identical.

Clearly, argumentation analogous to the one provided above can be applied to any other novel. So the following general thesis must be true: The contextualist element is essential to a novel and, hence, the concrete type to which this novel is identical.

Are there any essential elements of the concrete type to which a novel is identical, besides the ones examined above? The answer to this question, I think, is "No." Thus, given what has been said, the question posed at the beginning of this section can be answered as follows: The essential elements of the concrete type

²⁸For a powerful defense of the idea that contextual properties are relevant to the artistic value of an artwork, see Levinson (1980, 2007).

to which a non-visual novel is identical include the sonic, semantic, syntactic, and contextualist elements; the essential elements of the concrete type to which a visual novel is identical include all of the mentioned elements *and* the visual element.

7.4 The Ontological Status of a Novel

Our analysis has shown that:

1. A novel can be treated as a concrete type—a concretum that (a) is capable of spatiotemporal multilocation (that is, can occupy several disjoint spatiotemporal regions without occupying their fusion) and (b) coincides with its instances (that is, is wholly located where each of its instances is wholly located but is neither identical to any of them or a mereological sum of them nor related to them in the way in which a whole is related to its proper parts).
2. The essential elements of the concrete type to which a novel is identical include the sonic, semantic, syntactic, and contextualist elements if this novel is non-visual, and the sonic, semantic, syntactic, contextualist, and visual elements if this novel is visual.

Given this, from an ontological point of view, a novel can be regarded as a concrete type that, depending on whether this novel is non-visual or visual, is composed of (a) the sonic, semantic, syntactic, and contextualist elements or (b) the sonic, semantic, syntactic, contextualist, and visual elements.

Bibliography

- Alpers, P. (1984). On Musical Improvisation. *Journal of Aesthetics and Art Criticism*, 43(1):17–29.
- Alward, P. (2004). The Spoken Work. *Journal of Aesthetics and Art Criticism*, 62(4):331–8211.
- Aristotle (2010). *Rhetoric*. Cosimo, Inc.
- Armstrong, D. (1978). *A Theory of Universals*. Cambridge University Press.
- Austin, J. L. (1962). *How to Do Things with Words*. Clarendon Press.
- Balaguer, M. (1998). *Platonism and Anti-Platonism in Mathematics*. Oxford University Press.
- Balogh, J. (1927). Voces paginarum. *Philologus*, 82:84–109, 202–240.
- Beardsley, M. (1981). *Aesthetics: Problems in the Philosophy of Criticism*. Hackett Publishing Company.
- Beardsley, M. (1983). An Aesthetic Definition of Art. In Curtler, H., editor, *What Is Art?* Haven Publications.
- Beardsley, M. (2004). The Concept of Literature. In Eileen, J. and Lopes, D., editors, *The Philosophy of Literature: Contemporary and Classic Readings*. Wiley-Blackwell.
- Bergeron, V. and Lopes, D. M. (2009). Hearing and Seeing Musical Expression. *Philosophy and Phenomenological Research*, 78(1):1–16.
- Burgess, J. P. (1997). *A Subject with No Object: Strategies for Nominalistic Interpretation of Mathematics*. Oxford University Press.
- Cameron, R. (2008). There Are No Things That Are Musical Works. *British Journal of Aesthetics*, 48(3):295–314.

- Caplan, B. and Matheson, C. (2004). Can a Musical Work Be Created? *British Journal of Aesthetics*, 44(2):113–134.
- Caplan, B. and Matheson, C. (2006). Defending musical perdurantism. *British Journal of Aesthetics*, 46(1):59–69.
- Caplan, B. and Matheson, C. (2008). Defending ‘Defending Musical Perdurantism’. *British Journal of Aesthetics*, 48(1):80–85.
- Carroll, N. (1998). *A Philosophy of Mass Art*. Clarendon Press.
- Carson, E. (1996). On Realism in Set Theory. *Philosophia Mathematica*, 4(1):3–17.
- Casati, R. and Dokic, J. (2014). Sounds. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Fall 2014 edition.
- Casati, R. and Varzi, A. (2015). Events. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Winter 2015 edition.
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- Chalmers, D. J. (2002). Does Conceivability Entail Possibility? In Gendler, T. S. and Hawthorne, J., editors, *Conceivability and Possibility*, pages 145–200. Oxford University Press.
- Chalmers, D. J. (2009). The Two-Dimensional Argument Against Materialism. In McLaughlin, B. P. and Walter, S., editors, *Oxford Handbook to the Philosophy of Mind*. Oxford University Press.
- Chartier, R. (2002). *The Book History Reader*, chapter The Practical Impact of Writing, pages 118–141. Routledge.
- Cliff-Hodges, G. (2015). *Researching and Teaching Reading: Developing Pedagogy Through Critical Enquiry*. Taylor & Francis.
- Collingwood, R. G. (1958). *The Principles of Art*. Oxford University Press.
- Cowling, S. (2017). *Abstract Entities*. Routledge.
- Currie, G. (1985). What is Fiction? *Journal of Aesthetics and Art Criticism*, 43(4):385–392.
- Currie, G. (1989). *An Ontology of Art*. St. Martin’s Press.
- Currie, G. (1991). Work and Text. *Mind*, 100(3):325–340.
- Danto, A. C. (1981). *The Transfiguration of the Commonplace: A Philosophy of Art*. Harvard University Press.

- Davidson, D. (1963). Actions, Reasons, and Causes. *Journal of Philosophy*, 60(23):685–700.
- Davidson, D. (1967). Causal Relations. *Journal of Philosophy*, 64(21):691–703.
- Davidson, D. (1969). The Individuation of Events. In Rescher, N., editor, *Essays in Honor of Carl G. Hempel*, pages 216–34. Reidel.
- Davies, D. (2008). *The Performance of Reading: An Essay in the Philosophy of Literature* by Kivy, Peter. *Journal of Aesthetics and Art Criticism*, 66(1):89–91.
- Davies, D. (2010). Multiple Instances and Multiple ‘Instances’. *British Journal of Aesthetics*, 50(4):411–426.
- Davies, D. (2011). *Philosophy of the Performing Arts*. Wiley-Blackwell.
- Davies, D. (2012). What Type of ‘Type’ Is a Film? In Mag Uidhir, C., editor, *Art and Abstract Objects*. Oxford University Press.
- Davies, S. (1991). *Definitions of Art*, volume 42. Cornell University Press.
- Davies, S. (2001). *Musical Works and Performances: A Philosophical Exploration*. Oxford University Press.
- Davies, S. (2003a). Ontology of Art. In Levinson, J., editor, *The Oxford Handbook of Aesthetics*. Oxford University Press.
- Davies, S. (2003b). *Themes in the Philosophy of Music*. Oxford University Press.
- Dean, J. T. (2003). The Nature of Concepts and the Definition of Art. *Journal of Aesthetics and Art Criticism*, 61(1):29–35.
- Deutsch, H. (1991). The creation problem. *Topoi*, 10(2):209–225.
- Dickens, C. (1902). *A Tale of Two Cities*. James Nisbet & Co., Limited.
- Dilworth, J. (2008). The Abstractness of Artworks and Its Implications for Aesthetics. *Journal of Aesthetics and Art Criticism*, 66(4):341–353.
- Dodd, J. (2000). Musical Works as Eternal Types. *British Journal of Aesthetics*, 40(4):424–440.
- Dodd, J. (2002). Defending Musical Platonism. *British Journal of Aesthetics*, 42(4):380–402.
- Dodd, J. (2004). Types, Continuants, and the Ontology of Music. *British Journal of Aesthetics*, 44(4):342–360.
- Dodd, J. (2007). *Works of Music: An Essay in Ontology*. Oxford University Press.

- Dretske, F. I. (1977). Referring to Events. *Midwest Studies in Philosophy*, 2(1):90–99.
- Dummett, M. (1973). *Frege: Philosophy of Language*. Duckworth.
- Eby, C. P. (2013). Literary Movements. In Moddelmog, D. A. and del Gizzo, S., editors, *Ernest Hemingway in Context*. Cambridge University Press.
- Feagin, S. L. (2008). Critical Study: Reading and Performing. *British Journal of Aesthetics*, 48(1):89–97.
- Finkelstein, D. and McCleery, A. (2013). *An Introduction to Book History*. Routledge.
- Fisher, J. A. (1991). Discovery, Creation, and Musical Works. *Journal of Aesthetics and Art Criticism*, 49(2):129–136.
- Fisher, J. A. (1995). Is There a Problem of Indiscernible Counterparts? *Journal of Philosophy*, 92(9):467–484.
- Godlovitch, S. (1998). *Musical Performance: A Philosophical Study*. Routledge.
- Goodman, J. (2007). A Novel Category of Vague Abstracta. *Metaphysica*, 8(1):79–96.
- Goodman, N. (1968). *Languages of Art*. Bobbs-Merrill.
- Goodman, N. and Elgin, C. Z. (1987). *Reconceptions in Philosophy and Other Arts and Sciences*. Hackett Publishing Company.
- Grice, H. P. (1975). Logic and Conversation. In Cole, P. and Morgan, J. L., editors, *Syntax and Semantics: Speech Acts*, volume 3. Academic Press.
- Grierson, H. J. C. (1906). *Periods of European Literature: The First Half of the Seventeenth Century*. William Blackwood and Sons.
- Grisham, J. (2011). *Bleachers*. Dell.
- Hale, B. (1988). *Abstract Objects*. Blackwell.
- Hazlett, A. (2012). Against Repeatable Artworks. In Uidhir, C. M., editor, *Art and Abstract Objects*. Oxford University Press.
- Hicks, J. (2009). ‘Fire, Fire, Fire Flowing Like a River, River, River’: A History and Postmodernism in Truman Capote’s Handcarved Coffins. In Bloom, H., editor, *Truman Capote*. Infobase Publishing.
- Hoefer, C. (2016). Causal Determinism. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2016 edition.

- Howell, R. (2002). Types, Indicated and Initiated. *British Journal of Aesthetics*, 42(2):105–127.
- Ingarden, R. (1973). *The Literary Work of Art*. Northwestern University Press.
- Kane, R. (2005). *A Contemporary Introduction to Free Will*. Oxford University Press.
- Kania, A. (2011). All Play and No Work: An Ontology of Jazz. *Journal of Aesthetics and Art Criticism*, 69(4):391–403.
- Kim, J. (1973). Causation, Nomic Subsumption, and the Concept of Event. *Journal of Philosophy*, 70(8):217–236.
- Kim, J. (1976). Events as Property Exemplifications. In *Action Theory*, pages 310–326. D. Reidel.
- King, S. (2000). *On Writing*. Simon & Schuster.
- Kivy, P. (1983). Platonism in Music. *Grazer Philosophische Studien*, 19:109–129.
- Kivy, P. (1987). Platonism in Music: Another Kind of Defense. *American Philosophical Quarterly*, 24(3):245–252.
- Kivy, P. (1995). *Authenticities: Philosophical Reflections on Musical Performance*. Cornell University Press.
- Kivy, P. (2002). *Introduction to a Philosophy of Music*. Clarendon Press.
- Kivy, P. (2006). *The Performance of Reading: An Essay in the Philosophy of Literature*. Blackwell Publishing.
- Kivy, P. (2010). The Experience of Reading. In Hagberg, G. and Jost, W., editors, *A Companion to the Philosophy of Literature*. Wiley-Blackwell.
- Knox, B. M. W. (1968). Silent Reading in Antiquity. *Greek, Roman and Byzantine Studies*, 9:421–435.
- Kowalewski, M. (1993). *Deadly Musings: Violence and Verbal Form in American Fiction*. Princeton University Press.
- Kripke, S. A. (1980). *Naming and Necessity*. Harvard University Press.
- Lamarque, P. (2009). *The Philosophy of Literature*. Blackwell Pub.
- Lamarque, P. (2010). *Work and Object: Explorations in the Metaphysics of Art*. Oxford University Press.
- Levinson, J. (1980). What a Musical Work Is. *Journal of Philosophy*, 77(1):5–28.

- Levinson, J. (1989). Refining Art Historically. *Journal of Aesthetics and Art Criticism*, 47(1):21–33.
- Levinson, J. (1992). *An Ontology of Art*, by Gregory Currie. *Philosophy and Phenomenological Research*, 52(1):215–222.
- Levinson, J. (2006). *Contemplating Art: Essays in Aesthetics*. Oxford University Press.
- Levinson, J. (2007). Aesthetic Contextualism. *Postgraduate Journal of Aesthetics*, 4(3).
- Levinson, J. (2011a). Autographic and Allographic Art Revisited. In *Music, Art, and Metaphysics*. Oxford University Press.
- Levinson, J. (2011b). What a Musical Work Is, Again. In *Music, Art, and Metaphysics*. Oxford University Press.
- Levinson, J. (2012). Indication, Abstraction, and Individuation. In Uidhir, C. M., editor, *Art and Abstract Objects*. Oxford University Press.
- Lewis, D. (1976). The Paradoxes of Time Travel. *American Philosophical Quarterly*, 13(2):145–152.
- Lewis, D. (1983). Extrinsic Properties. *Philosophical Studies*, 44(2):197–200.
- Lewis, D. (1986). Events. In Lewis, D., editor, *Philosophical Papers Vol. II*, pages 241–269. Oxford University Press.
- Lewis, R. (1985). Money, Love, and Aspiration in *The Great Gatsby*. In Brucoli, M. J., editor, *New Essays on The Great Gatsby*. Cambridge University Press.
- Llosa, M. V. (1987). *The Perpetual Orgy: Flaubert and Madame Bovary*. Farrar, Straus and Giroux.
- Locke, J. (1975). *An Essay Concerning Human Understanding*. Clarendon Press.
- Lyas, C. A. (1969). The Semantic Definition of Literature. *Journal of Philosophy*, 66(3):81–95.
- Maclachlan, D. L. C. (1989). *Philosophy of Perception*. Cliffs Prentice-Hall.
- Maddy, P. (1990). *Realism in Mathematics*. Oxford University Press.
- Maddy, P. (1997). *Naturalism in Mathematics*. Oxford University Press.
- Mag Uidhir, C. (2007). Recordings as Performances. *British Journal of Aesthetics*, 47(3):298–314.
- Mag Uidhir, C. (2013). *Art & Art-Attempts*. Oxford University Press.

- Manguel, A. (2014). *A History of Reading*. Penguin Books.
- Markosian, N., Sullivan, M., and Emery, N. (2016). Time. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2016 edition.
- Meinong, A. (1960). On the Theory of Objects (Translation of ‘Über Gegenstandstheorie’, 1904). In Chisholm, R., editor, *Realism and the Background of Phenomenology*, pages 76–117. Free Press.
- Melville, H. (1922). *Moby Dick*. The St. Botolph Society.
- Meskin, A. and Robson, J. (2011). Videogames and the Moving Image. *Revue Internationale de Philosophie*, 4:547–564.
- Milne, P. (1994). The Physicalization of Mathematics. *British Journal for the Philosophy of Science*, 45(1):305–340.
- Nanay, B. (2012). The Multimodal Experience of Art. *British Journal of Aesthetics*, 52(4):353–363.
- Nannicelli, T. (2013). *A Philosophy of the Screenplay*. Routledge.
- Nietzsche, F. (2002). *Beyond Good and Evil: Prelude to a Philosophy of the Future*. Cambridge University Press.
- O’Callaghan, C. (2009). Sounds and Events. In Nudds, M. and O’Callaghan, C., editors, *Sounds and Perception: New Philosophical Essays*, pages 26–49. Oxford University Press.
- O’Callaghan, C. (2010). Constructing a Theory of Sounds. *Oxford Studies in Metaphysics*, 5:247–270.
- Ohmann, R. (1971). Speech Acts and the Definition of Literature. *Philosophy and Rhetoric*, 4(1):1–19.
- Parsons, T. (1975). A Meinongian Analysis of Fictional Objects. *Grazer Philosophische Studien*, 1:73–86.
- Pasnau, R. (1999). What is Sound? *Philosophical Quarterly*, 50(196):309–24.
- Peirce, C. S. (1931). *Collected Papers of Charles Sanders Peirce*. Cambridge: Harvard University Press.
- Petkov, C. I. and Berlin, P. (2013). Silent Reading: Does the Brain ‘Hear’ Both Speech and Voices? *Current Biology*, 23(4):R155–R156.
- Plimpton, G. (1966). The Story Behind a Nonfiction Novel. *New York Times*.
- Predelli, S. (2001). Musical Ontology and the Argument From Creation. *British Journal of Aesthetics*, 41(3):279–292.

- Putnam, H. (1970). Is Semantics Possible? *Metaphilosophy*, 1(3):187–201.
- Putnam, H. (1975). The Meaning of ‘Meaning’. *Minnesota Studies in the Philosophy of Science*, 7:131–193.
- Quine, W. V. (1948). On What There Is. *Review of Metaphysics*, 2(1):21–38.
- Reicher, M. (2016). Nonexistent Objects. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition.
- Rey, G. (1999). Concepts and Stereotypes. In Margolis, E. and Laurence, S., editors, *Concepts: Core Readings*, page 261. MIT Press.
- Ribeiro, A. C. (2007). Intending to Repeat: A Definition of Poetry. *Journal of Aesthetics and Art Criticism*, 65(2):189–201.
- Riskin, A. (1994). On the Most Open Question in the History of Mathematics: A Discussion of Maddy. *Philosophia Mathematica*, 2(2):109–121.
- Rohrbaugh, G. (2003). Artworks as Historical Individuals. *European Journal of Philosophy*, 11(2):177–8211.
- Rosch, E. (1978). Principles of Categorization. In Rosch, E. and Lloyd, B., editors, *Cognition and Categorization*. Lawrence Elbaum Associates.
- Rosen, G. (2012). Abstract Objects. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Spring 2012 edition.
- Rowling, J. K. (2004). *Harry Potter and the Philosopher’s Stone*. Bloomsbury Publishing Plc.
- Saenger, P. (1997). *Space Between Words: The Origins of Silent Reading*. Stanford University Press.
- Sartre, J.-P. (2004). *The Imaginary: A Phenomenological Psychology of the Imagination*. Routledge.
- Schaffer, J. (2016). The Metaphysics of Causation. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2016 edition.
- Searle, J. R. (1975). The Logical Status of Fictional Discourse. *New Literary History*, 6(2):319–32.
- Searle, J. R. (2012). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press.
- Sider, T. (2001). *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press.

- Smith, E. E. and Medin, D. L. (2002). The Exemplar View. In Levitin, D., editor, *Foundations of Cognitive Psychology: Core Readings*, pages 277–292. Mit Press.
- Smith, N. J. (2016). Time Travel. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2016 edition.
- Stecker, R. (2004). What Is Literature? In John, E. and Lopes, D., editors, *Philosophy of Literature: Contemporary and Classic Readings*. Blackwell Pub.
- Stevenson, R. L. (2011). *The Art of Writing*. Indo-European Publishing.
- Strawson, P. (2003). Aesthetic Appraisal and Works of Art. In Lamarque, P. and Olsen, S. H., editors, *Aesthetics and the Philosophy of Art*. Blackwell.
- Swinnerton, F. (1915). *R. L. Stevenson*. Mitchell Kennerley.
- Thom, P. (1993). *For an Audience: A Philosophy of the Performing Arts*. Temple University Press.
- Thomas, R. (1992). *Literacy and Orality in Ancient Greece*. Cambridge University Press.
- Thomasson, A. (2008). Roman Ingarden. In *Stanford Encyclopedia of Philosophy*.
- Thomasson, A. L. (1999). *Fiction and Metaphysics*. Cambridge University Press.
- Thomasson, A. L. (2004). The Ontology of Art. In *The Blackwell Guide to Aesthetics*. Blackwell Pub.
- Tillman, C. (2011). Musical Materialism. *British Journal of Aesthetics*, 51(1):13–29.
- Tillman, C. and Spencer, J. (2012). Musical Materialism and the Inheritance Problem. *Analysis*, 72(2):252–259.
- Tolstoy, L. (2007). *War and Peace*. Random House, Inc.
- Tye, M. (1990). Vague Objects. *Mind*, 99(396):535–557.
- Urmson, J. O. (2004). Literature. In John, E. and Lopes, D., editors, *Philosophy of Literature: Contemporary and Classic Readings*, pages 88–92. Blackwell Pub.
- Van Cleve, J. (1985). Why a Set Contains its Members Essentially. *Noûs*, 19(4):585–602.
- Walters, L. (2013). Repeatable Artworks as Created Types. *British Journal of Aesthetics*.
- Walton, K. L. (1990). *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Harvard University Press.

- Watt, I. (1967). *The Rise of the Novel*. University of California Press.
- Weatherson, B. and Marshall, D. (2014). Intrinsic vs. Extrinsic Properties. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2014 edition.
- Weitz, M. (1956). The Role of Theory in Aesthetics. *Journal of Aesthetics and Art Criticism*, 15(1):27–35.
- Wetzel, L. (2009). *Types and Tokens: On Abstract Objects*. MIT Press.
- Wetzel, L. (2014). Types and Tokens. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2014 edition.
- Wilson, G. and Shpall, S. (2012). Action. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Summer 2012 edition.
- Wollheim, R. (1971). *Art and its Objects*. Harper & Row.
- Wollheim, R. (1980). *Art and its Objects: With Six Supplementary Essays*. Cambridge University Press.
- Yablo, S. (1996). Wide Causation. *Philosophical Perspectives*, 11(11):251–281.
- Yablo, S. (2002). Abstract Objects: A Case Study. *Noûs*, 36(s1):220–240.
- Yúrkievich, S. (2003). Saying and Singing: An Interview with Julio Cortazár and Saul Yúrkievich. In *Background Noise*. Catbird Press.
- Zangwill, N. (1995). The Creative Theory of Art. *American Philosophical Quarterly*, 32(4):307–323.
- Zemach, E. M. (1970). Four Ontologies. *Journal of Philosophy*, 67(8):231–247.
- Zemach, E. M. (1992). *Types: Essays in Metaphysics*. E.J. Brill.