A Paper About Musical Works and Geometric Objects
(Note: actual title removed.)

Michaela McSweeney

ABSTRACT: I argue for two theses about abstract objects like *the equilateral triangle* and Allegri's choral work *Miserere Mei Deus*. The negative thesis says that (e.g.) *the equilateral triangle* is not triangular, nor does it in any sense have an abstract "analogue" of triangularity. The positive thesis says that (e.g.) *the equilateral triangle* does, however, have a nature, and that nature is purely structural. I argue for both theses by examining the relationship between these abstract objects and the concrete objects that resemble them.

Introduction

Many philosophers give negative accounts of abstract objects, that is, they try to define abstract objects in terms of the properties they uniquely lack.[1] While some philosophers give positive accounts of abstract objects, they usually don't say much about what abstract objects are like.[2]

I argue for both a negative and a positive thesis about two kinds of abstract objects, geometric shapes and musical works, using *the equilateral triangle* and Allegri's choral work *Miserere Mei Deus* as examples. To motivate both theses, I show that we can learn about abstract objects from examining the concrete objects that resemble them.

The negative thesis is that these abstract objects lack certain kinds of properties which I'll call **what-it's-like** properties. What-it's-like properties are those properties that (a) aren't fundamentally relational in nature and (b) can be (but needn't actually be) instantiated in spacetime. Properties like *being greater than, being to the left of,* etc. are not what-it's-like properties (even though *being to the left of* can be instantiated in spacetime, it is relational). Neither is the property *being a natural number,* because it is not instantiable in spacetime (it might or might not be fundamentally relational). Notice that *being identical to Mars* is a what-its-like property (because it is instantiated in spacetime), but *being identical to the number 7* is not (because it can't be instantiated in spacetime).

My negative thesis hence explains why it is so hard to give positive accounts of what abstract objects are like. What-it's-like properties are properties that can serve to give us a sense of the character of the object that instantiates them—of what that object is like. As the term suggests, if these objects don't instantiate any what-it's-like properties (or, as I'll also argue, quasi-what-it's-like properties), then there is a sense in which they aren't like anything at all. (Note: 'what it's like' here is about what the object is like, and has nothing to do with what it is like for some creature to perceive the object.)

---

[1] E.g. Goodman and Quine (1947) seem to assume a negative account, where abstract objects lack spatiotemporal location, (though they don't think such things exist), as does Quine (1948). Armstrong (1978) discusses them as lacking causal power. Lewis (1986) calls this the "way of negation".
[2] I am thinking here of "abstractionist" and related accounts, e.g. Hale (1987), Wright (1983); Dummett (1973). Plato is an exception. Zalta (1983) and Cowling (2017) may also be.

If you're concerned that the negative thesis is trivial given this definition of what-it's-like properties, note that it rules out *the equilateral triangle* having three angles (which is a property that is instantiable in spacetime and is not fundamentally relational). It also rules out that musical works, if they are abstract objects, could be loud, beautiful, sonic, etc.

Take a **structural property** to be any non-what-it's-like property. The positive thesis says that the abstract objects in questions do have natures—but that they are purely structural natures. This means that they have some intrinsic properties, but they are only structural properties. Shapiro (1997) and Resnik (1997) advocate forms of mathematical structuralism on which numbers are nodes in mathematical structures. My positive thesis is similar: structuralism is true of both geometry (point and line are nodes in the structural object *the equilateral triangle*) and musical works (some note might be a node in *Miserere*); but *the equilateral triangle* itself cannot be a node in a structure, since it itself has internal structure. Instead, *the equilateral triangle* is akin to Shapiro's *natural number system;* it is itself a Platonic structure.

So, my thesis is the conjunction of two claims:

> (a) objects like *Miserere* and *the equilateral triangle* have no what-it's-like properties, but (b) they do have intrinsic, structural properties and hence have natures.

This leaves open whether there are other abstract objects that either have what-it's-like properties or lack internal structure. While this means that I don't (yet) have a unified account of abstracta, this is simply a result of what I take to be the correct methodology for investigating the nature of entities we know little about. Just as it would be a mistake to assume that all unobservable entities that we posit in physics have something important and illuminating in common, it is a mistake to assume that abstracta have something important and illuminating in common. But even if abstracta *do* have something important and illuminating in common, I think that the right way to arrive at that commonality is by examining individual abstract objects first, not by attempting to say what it is that abstracta lack.

In §1, I explain what it is for an object to be a best mime of a given target, and I show that one consequence of this definition is a principle, MIMESIS. In §2, I introduce 'mimetic translation' and use it to motivate the claim that there are prima facie ties for best mimes of abstract objects. In §3, I argue that there are no candidates for breaking these ties. In §4, I apply MIMESIS, and show that these objects have no what-it's-like (or quasi-what-it's-like) properties. I also argue that these objects have internal structure. I conclude that *Miserere* and *the equilateral triangle* (and objects like them) have purely structural internal natures.

1. Mimes

Crucial to my argument is the notion of a mime. A **mime** is a model of a target (the object it is mimicking); it resembles that target either by (a) sharing intrinsic properties with that target or (b) by "mimicking" those properties. Mimes are better or worse depending on how well they do (a) or (b).

I take (a) to be fairly obvious. For example, consider two identical twins, Tegan and Sara. It seems clear that Tegan does a good job mimicking Sara precisely because Tegan and Sara share many intrinsic properties. To see that that we need (b), suppose that are wondering which of the following is a better mime of a beech tree: the linguistic description 'a beech tree', a child's drawing of the beech tree, a rubber ball, and an accurate three-dimensional clay model of the beech tree. The model is the best mime among these, because it directly models many of the properties of the beech tree, e.g. having green leaves of a specific shape, a trunk, and a certain pattern of bark, being three-dimensional, etc.

Some of these properties (e.g. being three-dimensional) the model literally shares with the beech tree; others (e.g. having green leaves) it does not, since its "leaves" are not really leaves. So why is the model a better mime than the ball, which is also three-dimensional?

To answer this, I first must introduce **predicate-sharing**: sharing a predicate in virtue of sharing similar, but distinct, properties. The beech tree has green leaves*,* and the model has green leaves*,* but 'has green leaves' picks out distinct properties in each case. What matters is that it is *appropriate* for us to use the same predicate to pick out both properties, at least in the context of treating the model tree as a model of the beech tree. Predicate-sharing is appropriate here because the property of the model mimics the property of the tree. I don't think I'm stipulating predicate-sharing into existence—our actual linguistic practices already include it. It is also similar to how we use the same names to pick out both parts of models and parts of their targets, e.g., we might say 'Ireland is pink' when we are looking at a world map, but it does not follow that we think that the pink blob on the map is identical to the country Ireland.

When it comes to concrete objects, we can often be good judges of when predicate-sharing is appropriate; we often just look and see whether the mime's property looks like (or smells like, sounds like…) the target's property.[3] But this is much harder with abstract objects.[4]

---

[3] One way of understanding Plato's views about self-predication is along these lines: a concrete object and a Form predicate-share when the concrete object is a good mime of the Form; but the Form and the concrete object actually share a property (because the form of the Large is not, itself large—it is perhaps *large**). The idea is that while the Form and the concrete object don't share a property, it is appropriate to use the same predicate, 'large', to refer to *large* and *large*,* because the concrete object is modeling the Form.

[4] Mimeticity is similar to Dretske's (1981, ch. 1) 'analog representation'. I don't use this term here for two reasons: first, it is unclear that being a mime is a kind of representation; second, it suggests that an analogy is being drawn between the mime and the target; but in some instances, mimeticity is just identity. Mimeticity *also* bears some similarity to Quine's notion of canonical representation, or Shapiro's (e.g. 1985); but again, I use the term differently. Quine and Shapiro are targeting linguistic (or mathematical-linguistic) representations, whereas I'm concerned with cases where the best mimes will typically not be linguistic at all.

Predicate-sharing is not a feature of the mind-and-language-independent world; we have to judge when two objects predicate-share their properties, and predicate-sharing necessarily involves language. Still, the appropriateness of predicate-sharing is determined by objective, mind-and-language-independent facts about objects' properties. That the model mimics the tree is not a fact about the way we use language; it is a fact about reality: about the model and the tree themselves. Predicate-sharing is a way of capturing objective property similarity, even when there is no property sharing. Inappropriate predicate-sharing occurs when we err in choosing to pick out two properties with the same predicate (because there is no genuine resemblance between the two properties).[5]

Predicate-sharing might not be enough to get the modeling relation we want between mimes and their targets. Consider our clay model of the tree, and a second model, whose roots are growing out of its branches and whose leaves are buried under a clay ground. This second model might predicate-share with the tree just as much as the first does; it has green leaves, it has roots, and so on. But the way that these properties are distributed—in this case, their location—does not accurately model the tree. Best mimes' properties are distributed in the same way that their target's properties are; all this means is that if the objects in question have (either physical or non-physical) parts, the properties of the objects are distributed over those parts in similar ways.

Given these claims, here are two definitions:

> The **best mimes** of a given target are the objects (a) predicate-share the most with the intrinsic properties of the target, and (b) are such that the properties of the mime are distributed in as similar as possible a way to the properties of the target.[6]
>
> **Mimeticity** is a measure of how good of a mime x is of a target y. (I will often use 'mimeticity' without specifying the target, when context makes this obvious.)

The best mimes of a target are the ones that perfectly mimics it. How should we think about best mimes of abstract objects? Suppose that the abstract object *the equilateral triangle* is spatial or quasi-spatial--that it is extended either in real space or some abstract analog of real space. (I will argue against this claim later in the paper; here I use it as an illustrative example.) It follows that pictorial mimes of *the equilateral triangle* are more mimetic than algebraic ones, despite the fact that they don't share

[5] Dodd (2007), following Wollsterstroff (1980), claims that these uses of (e.g.) 'loud' are analogical predicates: when we say 'musical work x is loud' what we mean is 'x is such that y cannot be a properly formed token of x unless it is loud'. The crucial question here is what it is to be a properly formed token of x. I prefer predicate-sharing, where we don't try to give an analysis or definition of the property that we are attributing to the abstract object, in part because it is unclear how we determine what it is to be a properly formed token of x. Kleinschmidt and Ross (2012) argue against Dodd's analog predication, which would affect predicate-sharing, but which does not affect my final conclusion.

[6] Here I invoke the general notion of an intrinsic property: it is, as Lewis puts it, a property "which things have in virtue of the way they themselves are" (1986 p. 61). I'm sidestepping the vexed issue of sharpening that definition. (See e.g. Sider 1996, Lewis and Langton 1998, Cameron 2009, Francescotti 2014.)

properties with *the equilateral triangle*. They have a property—being concretely two-dimensional—that they predicate-share with the property—being abstractly two-dimensional—and that *the equilateral triangle* also has. Pictorial mimes have physical correlates of properties of their abstract targets—in this case, for example, they have three (roughly) equal angles and three (roughly) equal sides. So, if we know that *the equilateral triangle* has three equal angles and three equal sides, then we should think a pictorial mime is better than an algebraic one, as the latter lacks angles and sides.

There is no obvious relation between the type-token relationship and mimeticity. If some or all the tokens of a given type are the best mimes of that type, then that type-token relationship might be subsumed by mimeticity. Yet mimeticity is not reducible to the type-token relationship, since mimeticity is a general notion that applies to any target, not just to types. Further, there is no reason to think that all of the tokens of a given type are equally best mimes of that type—there can be paradigmatic tokens of a type. One standard view in the ontology of music is that musical works are types, whose tokens are their performances (and perhaps other related "sound events"). Wolterstorff (1980) and Dodd (2007), come closest to addressing the relationship between mimeticity and type-token relationships. Both think that there are appropriate and inappropriate tokens of musical types. But mimeticity does not track the "appropriate token" relation either. Being the best mime of a target is distinct from, and neither necessarily includes or excludes, being a token of a type.

1.2 Epistemology of Mimes: Motivating MIMESIS

How can we know that one mime of a target is better than another? In one sense it is easy to identify the best mimes of any object: every object is a best mime of itself. (Note that there might be many mimes that are tied for best--qualitative duplicates of the object.) But it is harder to identify the second-best, third-best, etc. mimes. Some mimes will seem in some ways but not others. Still, identifying best mimes is only as hard as identifying the intrinsic properties of the target. We typically don't know what the intrinsic properties of abstracta are. So what do we do?

To begin to answer this, I will motivate the following principle:

> **MIMESIS:** Let *a* be any object, and *m* be any potential candidate mime of *a*.
> **If**: we have some *m, m', m''*… that are candidates for being best accessible-to-us mimes of *a,* and no good reason to think any of the *m*s is better than the others, **then:** we should not believe that *a* has any properties that predicate-share with some but not all of the *m*s.

MIMESIS might initially appear trivial, given that every object (along with its qualitative duplicates) is its own best mime. This is what "accessible-to-us" is doing in its definition. I introduce MIMESIS only as an epistemic aid in determining what an object's intrinsic properties are, when we can't directly examine the object. It is epistemically unhelpful to say "we can learn a lot about the earth's core from directly examining the thing that best models the earth's core, which is the earth's core itself", since we cannot directly examine the earth's core. Likewise, it is unhelpful to say this kind of thing about any abstract object. I am interested in cases in which we can learn things about the properties of epistemically inaccessible objects from those mimes that we have access to. In the abstract case, for example, what we are looking for (at least initially) to satisfy the antecedent of MIMESIS are concrete objects that "mimic" the abstract object *better than all other concrete objects do.*

There are only two ways for the consequent of MIMESIS to be false. Given one assumption (that I will outline momentarily), in both cases, the antecedent must be false. First, suppose that *m, m',* and *m''* are our candidate best mimes of *a.* Now suppose that *m* has a property, p, that it predicate-shares with a property, p*, that *a* has. (So '*m* is *P'* and '*a* is *P'* are both true*,* where *P* is the predicate we use to refer to both p and p*.) Further, suppose that *m'* and *m''* are not *P.*

There are now two possibilities. First: excluding P, *m, m',* and *m''* predicate-share all and only the same predicates P', P"… with each other and with *a.* Then the antecedent is immediately falsified: Since *m* predicate-shares an additional property with *a* that *m'* and *m''* do not, and all else is equal, *m* is by definition a better mime of *a* than *m'* and *m''* are. (Example: suppose *a* is a (real) plum tree that has plums growing on it, *m, m',* and *m''* are models that are qualitatively identical except that *m'* and *m''* can't appropriately said to have plums, whereas *m* can (e.g. it has small purple clay balls hanging from its clay branches). Then *m* is more mimetic of *a.*)

The second possibility is that *each* of some subset of the *m*s has a distinct property that *a* predicate-shares with, and the rest of *a*'s properties are either predicate-shared with all of the *m*s, or with none of them. (So, in addition to P, there is some predicate P' which *m'* predicate shares with *a* but with neither of *m* or *m'',* and similarly for P" and *m''.*) Here we need to invoke the following assumption:

> **Construction:** when we have two (or more) equally best mimes of a target, each of which shares a property (that the other lacks) with the target, we can always generate (or find, or imagine) a new, better mime.

Construction just entails the following sort of thing: if we are trying to find the best mime among the models of the plum tree, and none of our models have plums, we can add some little purple clay plums to get a better model.

Given construction, in the case in which each of *m, m',* and *m''* predicate-share in distinct ways with *a,* we should simply be able to construct (or find) a new *m'''* that is more mimetic than any of *m, m',* or *m''*: we construct an *m'''* that has (i) predicate-shares everything that *a, m, m',* and *m''* already predicate-share, and (ii) also has all of the properties that some but not all of the original *m*s predicate-share with the target. *m'''* will now be a better mime of *a* than any of the original *m*s, thus falsifying the antecedent. (Example: suppose *a* is a grafted tree and has plums, apricots, and peaches growing on it. Suppose *m* has plums (but not apricots or peaches), *m'* has apricots (but not peaches or plums), and *m''* has peaches (but not plums or apricots). Solution: construct a new model, *m''',* that has all three. Now *m'''* supplants the other *m*s as being the best mime of *a*.)[7]

These two cases (and relevantly similar variations) are the only two ways that the consequent can be false, and in either case, the antecedent is also false. So MIMESIS is true.

Why believe in construction? I think the better question is to ask under what circumstances it could be false. One obvious case in which it seems to be falsified is if it were somehow simply too difficult for us to construct a better mime (and there wasn't some already existing object that could do the job). But remember that MIMESIS is supposed to be an epistemic principle, that allows us to learn about what target objects are like. And so we don't actually need actual, concrete objects to serve as better mimes, but rather simply an understanding of what those objects would be like if they were to exist; hence the "or imagine" clause. (For example, suppose we lived in a world in which it was impossible to get clay to be any color but orange, green, or brown. We wouldn't be able to construct a clay model of our grafted fruit tree that had (ripe) plums. But we could easily imagine such a model, and know that the imagined model was a better mime than our actual models that only had apricots and peaches.)

Two notes: first, one might wonder if, rather than employing a construction claim that involves our ability to imagine mimes, I should instead state MIMESIS in modal terms. I do not do so because I don't want to confuse my reader into thinking that MIMESIS is a principle about metaphysics. Possible objects that we can't conceive of (or aren't in a position to conceive of qua mimes of our targets) are irrelevant to what MIMESIS is for, which is to enable us to learn about the nature of objects. And this is true even if those possible objects are actually the best mimes of a target.

Second, in constructing and finding mimes, we will never need to represent something as having logically or metaphysically incompatible properties. For recall that the target is going to have all the properties that we want our best mimes to have (or at least, will predicate-share them), and nothing can have logically or metaphysically incompatible properties. Hence, constructing a mime never requires that we do the (logically or metaphysically) impossible.

As I said above, in the case of concrete objects that we have easy access to (e.g. our grafted fruit tree), MIMESIS is not particularly useful to us. In part, this is because MIMESIS will lead us back to the object itself: the effort to find or construct a better mime in the kind of case described above will eventually lead to the tree itself. In the next section, though, I will show that MIMESIS is quite useful when we focus on objects that we can't directly observe or come into causal contact with.

---

[7] One might worry here about cases in which there infinitely better mimes approaching the target. If such cases are possible, they don't complicate MIMESIS, since in order for its antecedent to be true, we need an identifiable set of candidate mimes; and once we have those, the consequent will hold regardless of whether there is an infinitely-approaching-closeness problem.

2. Mimetic Translation

For any *a,* MIMESIS doesn't tell us whether any *m* is a best mime of *a*, except by comparing what we already know about how good *a's* other mimes are. This is as it should be, given that MIMESIS follows from the definition of a best mime (plus construction). MIMESIS applies only once we already believe that some *m*s are candidates for being best mimes of *a,* and have no way to choose which is most mimetic of *a.* But why think that we will encounter such cases with abstract objects? In this section, I answer this question by introducing mimetic translation.

A **mimetic translation** is a translation—a mapping between two things--that preserves the structure of a third thing: their target. Though I don't claim that mimes are representations, it is useful to temporarily think in representational terms: a mimetic translation is like a mapping between representations that preserve their target's structure. (I'll sharpen this in a moment.)

There are two kinds of work that mimetic translations can do. First, mimetic translations can explain how we know what we know. Consider a good candidate for being a best mime of *the equilateral triangle:* a thinly-outlined drawing of an equilateral triangle, produced with a compass and a straightedge. This drawing is a good candidate because it preserves *the equilateral triangle's* structure. A small green sphere, on the other hand, does not. Likewise, a near-perfect performance of *Miserere* preserves its structure, and a child's rendition of *Row Your Boat* does not. One way to explain why we know this (if, as I claim, we do know it) is that there is no mimetic translation between the drawing and the small green sphere. The lack of a mimetic translation explains why we know that certain objects are not among the best mimes of a given target: if x is a best mime of z, a necessary condition on y's being a best mime of z is that there is a mimetic translation between x and y.

Second, mimetic translations can provide us with new evidence. If x is a best mime of z, then the presence of a mimetic translation from x to y gives us a strong, but defeasible, reason to believe that y is also a best mime of z. There are mimetic translations between the drawing and the formula, and between the score and the performance. These mimetic translations demonstrate a structural equivalence between two different mimes, which share a target by sharing a structure.

Why is this defeasible? Structure isn't all that matters. In many cases, we have independent reasons to believe that z has some property, P, that it predicate-shares with x but not y. For example, consider DNA molecules from our two identical twins, Tegan and Sara. Sara's DNA is a best mime of Tegan's DNA. And there is a mimetic translation between Sara's DNA and a perfect plastic model of that DNA. This gives us reason to believe that the plastic model is also a best mime of Tegan's DNA. But we have a defeater: we know that Tegan's DNA has certain intrinsic properties that the plastic model lacks, but which Sara's DNA has.

In light of the work we need mimetic translations to do, I propose three constraints on mimetic translations. The first is that we (or at least, experts) must be able to produce mimes of one kind from mimes of the other kind. In the mathematical case, anyone who knows enough algebra and geometry can produce a token of an algebraic formula representing *the equilateral triangle* after looking at a picture of one, and vice-versa.

For *Miserere,* this is less straightforward. It is easy enough (for some of us) to go in one direction: to look at a copy of the score, and then perform it. It is harder to go in the other direction: to listen to a particular performance, and to produce a written score. However, I picked this example because, allegedly, the fourteen-year-old Mozart did just this: he listened to a single performance of *Miserere,* and then produced a near-perfect transcription.

The second constraint is that mimetic translations cannot involve ad-hoc stipulation: the translations can't be a result of stipulation or memorization of ad-hoc "matches" between mimes. For example, we can't just use a translation process which requires us to memorize dictionaries like the below, which tell us that when we see the label 'Budweiser', we should produce the square, when we see the label 'Miller Lite', we should produce the triangle, etc.:

 = Budweiser

 = Miller Lite

 = Michelob Ultra

(In the case of the small green sphere, without the second constraint, we could make this exact move there to show that there was a translation manual between near-perfect drawings of triangles and green spheres.)

Good structure-preserving translations do require both memorization (e.g. memorizing which placement of a mark on the staff corresponds to which tone produced on an instrument) and stipulation. Without knowing of at least a few nodes on each structure that they correspond to one another, we can't reproduce one from the other. Compare this to representing mass: without stipulating that an object in the world has mass of one gram, we can't apply the gram scale to mass-in-the-world. But the memorization and stipulation involved (in both the mass case, and the cases under discussion here) are not ad-hoc, in part because they are general, rather than specific: once we know how to read and perform music, they enable us to move back and forth between any score and performance, without a special translation manual for that particular score and performance.

In good cases, we might stipulate some minimal things about the relationship between nodes in a structural mime of kind K and nodes in a structural mime of kind K'; but doing so allows us to translate between any mime of kind K and the corresponding mime of kind K', and vice-versa. So whatever stipulation and memorization the translation involves, it must be stipulation and memorization of general frameworks for translation rather than specific cases. The translation process itself must generate the right results when it is applied to *any* mimes of kind K. It must give us a method for translating a random mime of kind K into one of kind K', and vice-versa, where experts recognize the resulting translation as correct.[8]

The third constraint is that there is minimal deletion and re-introduction of structure in the translation procedure. To see why we need this, imagine a translation procedure that says that if you have a mime that has the property of being red, you should ignore all of its properties, use a complex algorithm that determines how to add nineteen distinct properties, and translate the resulting twenty properties via some elaborate code into musical notation. One can imagine that we could develop systematic and perfectly general procedures like this, perhaps satisfying the second constraint; and such a procedure wouldn't be ruled out by the first constraint. But this is not a mimetic translation.

The third constraint is an attempt at ensuring that the translation procedure preserves everything that we take to matter about the target, once the relevant background is in place. Think of a translation procedure as a kind of transformation from one object to another. The constraint says that this transformation can't be too complicated, or take too long, or involve too many wild changes.

There is a straightforward objection to these constraints: if we don't know anything about the nature of the target of our mimes, then how do we know what structure in a given mime is mere noise and should be ignored in the translation procedure? Here are two responses. The first is to claim that the third constraint is a kind of parsimony constraint: translation procedures that are simple, not complex or gerrymandered, are better. Insofar as parsimony can guide us in theory choice, it should also guide us here.

The second, better response is to justify the constraints by clarifying my methodology. I don't assume that we don't know anything about music or mathematics. Independently of the questions of whether mathematical or musical objects exist, and if so, what they are like, we know a lot about mathematics and music. Given this, the first two constraints can be thought of as naturalizing constraints: mimetic translations must be dependent on the bodies of knowledge we have and the methods of inquiry we already use. We know that algebraic formulas and near-perfect drawings both do a good job mimicking certain mathematical objects (or doing whatever the nominalistic version of this is). Algebra and geometry are both well-developed bodies of human inquiry and knowledge, and moreover, the relationship between them is also well understood. Contrast this with the relationship I have stipulated into existence between beers and pictures of shapes. The stipulation in the latter case, but not the former, is unacceptable: in the latter case, it is ad-hoc. More specifically, it is not in the service of highlighting a known pre-existing relationship between two domains of inquiry, but rather attempts to create such a relationship.

This is not to say that mimetic translations cannot sometimes be surprising. We must work enough leniency into our constraints so that discovery is possible. I don't pretend to have made these

---

[8] There is now an issue of what counts as a 'kind', but I will not address that here. I hope that what I have just said is somewhat intuitive; since these translations are serving an epistemic and not a metaphysical purpose, I am not worried that the borders between kinds of mimes are not metaphysically robust.

constraints precise enough to have struck the perfect balance between allowing for the possibility of discovery, and respecting the bodies of knowledge (and knowledge about relations between them) that we already have. But the constraints are motivated if we help ourselves to the claim that we know a fair amount about (e.g.) music, despite not knowing what musical works really are. And this claim seems sensible to me.

In light of this, I claim that there are facts of the matter—which, conditional on the existence of abstract objects, in many cases, we know--about which mimes are best. Consider the claim that a careful drawing is a better mime of *the equilateral triangle* than a small green sphere is. The proposed constraints on mimetic translation make sense of this. While it is easy to construct some translation procedure that maps the small green sphere to the formula, or to the drawing, and easy enough to imagine that we could become experts at translating between the two, it is hard to imagine how such a procedure could (a) allow for systematic translations between solid colored three-dimensional shapes and two-dimensional drawings, without requiring us to stipulate an ad-hoc translation manual into existence, and without a background context of a developed body of knowledge, and (b) not involve a very complex translation procedure.

How do we come to know that a mime of an abstract object is anywhere near most mimetic of that object? This is a version of a general objection to platonism that lies outside of the scope of this paper. I assume that a near-perfect drawing is a good candidate for being most mimetic of *the equilateral triangle,* and ask what follows. It seems that we can rank mimes once we fix the facts about how good a single mime is. More generally, neither MIMESIS nor mimetic translation will help us at all unless we presuppose that we have at least some idea of which concrete mimes are candidate best mimes. However, I think that we already have this kind of information. We can't learn from a mimetic translation that we were entirely wrong about our initial choice of a best mime. Instead, we learn conditionals: if the drawing is a good candidate for being a best resemblance of the triangle, then so is the formula.

Mimetic translations help establish what belongs in the antecedent of MIMESIS. And there are strong mimetic translations—translations that respect the three constraints I've proposed—between, on the one hand, the perfect drawing of *the equilateral triangle* and the algebraic formula, and, on the other hand, the performance of *Miserere* and the copy of the written score.

Suppose that I define a mass unit, the *shram,* which is equivalent to 1.4 grams. Gram representations of mass facts and shram representations of mass facts are equally good ways of representing these facts; part of how we determine that they are equally good is by seeing the structure-preserving translations between them. The structure-preserving translation between gram and shram measurements does not entail that mass structure is neither gram-like nor shram-like. What it does is massively burden-shift: we need to provide a reason to think that (e.g.) gram talk is more metaphysically perspicuous than shram talk is, or, put differently, we need a defeater for the claim that gram talk and shram talk are not equally good. Otherwise, these structure-preserving translations show that it is arbitrary to claim that gram talk is (metaphysically) better.

Likewise, the fact that we can produce mimetic translations between the formula and the drawing, and the score and the performance, puts the argumentative burden on those who claim that (e.g.) the performance is more mimetic than the score. Suppose that we start out thinking that a perfect performance of *Miserere* is a best mime. We know that there is a structure-preserving translation (that respects our two constraints) from that performance to a copy of the written score. It follows that we now need a reason for thinking that one is more mimetic.

The central claim in this section is this: in the case of both *Miserere* and *the equilateral triangle,* there are strong mimetic translations that do burden-shifting work. Focus on *Miserere.* We have a very strong mimetic translation between the token score and the token performance. So, absent a defeater, we should think they are equally most mimetic. If a written copy of its score and a particular performance of it are equally most mimetic, it will follow that these two mimes of *Miserere* satisfy the antecedent of MIMESIS; and that means that we should not believe that *Miserere* predicate-shares with any of the properties that the performance has but the written score lacks. (So, it will turn out that *Miserere* is not "sonic" in even an abstract sense, and that *the equilateral triangle* doesn't have three sides.)[9]

Let me consider a final objection before moving on, in part to help clarify things. Even if one is convinced by what I've said about the musical case, one might not be as convinced in the mathematical case. One might think that formulas, or equations, are not the sorts of things that wear their structure on their sleeve. Instead, one might think, they are more like sentences of a language.

A copy of a musical score and a corresponding performance seem (close to) structurally isomorphic; the written staff's left-to-right order corresponds to the performance's temporal order; the height of a note on the staff corresponds to the pitch of a note in the performance; the shape of a note on the page corresponds to how long it is held in the performance; and so on. This preservation of structure is what enables us to reproduce mimes of one type from another: each preserves the same structural core.[10] So the objection goes: musical notation is a language that preserves musical structure. But algebraic notation does not preserve geometric structure, so how could a token instance of an algebraic formula be a best mime of *the equilateral triangle?*

Certainly it seems at first glance that while the drawing of the triangle displays the triangle's

---

[9] A fruitful way to examine the argument here would be to compare it to literature on scientific representation. French (2003), French and Ladyman (1999), Van Frassen (2008), among others, have, for example, argued for various structural notions of scientific representation, and some have used these to motivate structural realism, which is directly related to the view I argue for here. But mimeticity itself bears much more in common with Giere's (e.g. 2004) notion of scientific representation as similarity (or "fit").

[10] There won't be a perfect structural correspondence; the score contains certain marks that correspond in no way to any feature of the performance, and many features of the performance (particularities that the conductor and performers add) won't either. But this supports my view: the target itself doesn't have these extra features. We can imagine an idealized version of the score, where such things are edited out; perhaps we can imagine an idealized performance, though that is trickier. I return to this issue in section four, and to a related objection about the triangle case.

structure, the formula represents the structure in some other way (e.g. not by literally looking just like the triangle). Here, though, I remind my reader that we are supposed to be starting with what we know. And we can do algebra without doing pictorial geometry; that we can understand the structure that the formula encodes that don't involve envisioning or depicting a figure with three equal sides and three equal angles that lies on a two-dimensional plane. Mathematicians (and mathematics students) aren't simply manipulating meaningless marks on a page (or in their heads) when they think about algebra non-geometrically. Moreover, we (or at least, some of us) can think algebraically when we look at pictorial representations of geometric figures. So it seems that there is some structure that our formula captures that is not simply reducible to our ordinary pictorial concept of a two-dimensional equilateral triangle. And also remember that the conclusion we are inching towards is that *the equilateral triangle* is not triangular; so its pictorial representations do no better here than its algebraic ones.

To see that the geometric and musical cases are not as far away from each other as you might think, imagine what it would be like to learn how to read musical notation if you could not either hear sound or imagine it. I claim: there is still a very clear sense in which you could learn to read music. And you would develop a concept of what it was that the marks on the page encoded or represented. This, I claim, is exactly what it would be like to learn algebra without knowing that it could also be represented spatially (or without the ability to see, feel, or imagine shapes). In both cases, you would be missing something important, but would also know quite a bit about the structure of the abstract objects. So the geometric and musical cases are not so far apart after all.

However, this objection is important to consider in order to clarify things. To be more careful, it is not, in fact, a token instance of a formula that itself is a candidate for being a best mime of *the equilateral triangle*. Rather, it is the content that that token instance represents. The formula represents some content which (a) is much more general than e.g. a two-dimensional pictorial representation of the triangle and (b) itself is a best mime of *the equilateral triangle*.

Some mimes wear their structure on their sleeve more than others do. Some work more like e.g. sentences in a language—where, most of the time, it is the content that that sentence represents that we really mean to be focusing on when we think of the sentence as a mime of a given target. Some work more like drawings or musical notation, where the very method of representation displays the structure of its intended target (though, of course, where we still have to do some interpretation). Varying levels of interpretation are required for different mimes. None of this seems problematic, given that all of the physical mimes in question involve some conventional bits that need to be interpreted before they can be understood to bear the right relationship to one another. For example, it is an arbitrary convention that if A and B are notational marks on a piece of sheet music, and A is to the left of B, then the note that A represents is to be played before the note that B represents.

Notice also that there is not just one drawing of an equilateral triangle and one token of an algebraic formula (nor one specific type of formula) that are candidates for being best mimes of *the equilateral triangle*. So we have to consider all of them. In addition to allowing us to conclude something from the drawing and the formula's relationship, my argument shows something about the relationships among (e.g.) the various algebraic mimes themselves: there is some more minimal structure that they all have in common, and this is their target. What I am trying to bring home is that the minimal structure that is the shared content of all of the token algebraic mimes is just as good a candidate for mimicking the structure of *the equilateral triangle* as the minimal structure that is the shared content of all of the token pictorial mimes.

I have so far argued that there is a strong but defeasible reason for thinking that the drawing and the formula (and the score and the performance) are equally mimetic. I now want to turn to potential defeaters for these claims.

3. Defeaters?

I will consider two potential defeaters for the claim that the drawing and the formula (and the score and the performance) are equally mimetic of their targets, and argue that neither of them are successful.

First, we could appeal to intuitions. But part of the point of the burden-shifting argument in §2 is to show that intuitions aren't sufficient to fix on the correct concept of (e.g.) an equilateral triangle. Further, note that there is a longstanding challenge for the platonist about why our beliefs about what abstract objects there are would align with what abstract objects there actually are. Almost no one takes intuitions to answer that challenge (at least not without further defense of the reliability of those intuitions). The challenge I am trying to answer in this paper is related: why would our concepts of the *natures* of abstract objects align with their actual natures? I take it that brute appeals to intuitions here are bad for all of the same reasons that they are bad with respect to the more standard challenge.

Moreover, there are debunking arguments about our intuitions about both music and geometry. Many of us have an intuition that the drawing is more mimetic than the token formula. The debunking argument says that this is because of the way most of us were taught about triangles. My elementary school teachers never told me about abstract Euclidean geometric objects. First, they pointed to lots of two-dimensional triangular objects and drawings and told me they were triangles; then, they taught me about angles, sides, and how to differentiate types of triangles. This whole time, I was being taught that these two dimensional concrete mimes of triangles were the real triangles. Only much later did I learn that there is a way to represent an equilateral triangle algebraically. The way we acquire concepts certainly helps determine what those concepts are like, so it's reasonable to think that given the way I was taught what a triangle is, my concept itself came to be of two-dimensional object with three sides and angles that sum to 180 degrees. And if I consult my intuitions about what triangles are like, I am very likely to report on this concept.

So intuitions can't be defeaters, at least absent arguments for their reliability and against this kind of debunking challenge. What about intentions? Perhaps we intend to represent something spatially extended, or quasi-spatially extended, with both the drawing and the formula. Can our intentions guarantee that our mimes are best mimes, by using them to fix their target?

No, for familiar reasons: it would be a bizarre coincidence if the very abstract objects that existed were all and only those which we intended our concrete objects to serve as best mimes to. The only way I know of to overcome this bizarre coincidence is to adopt **Plenitudinous Platonism**: the claim that there is an abstract object for every concept we could possibly have (or: any abstract object that could exist, does exist), and which concept we have determines which object our terms refer to (or, in this case, what our mime targets).[11]

The Plenitudinous Platonist (PP) has an easy answer to which of our mimes are best (whichever we intend to be best are best). But the PP faces a related problem. Mathematics is full of non-trivial identities, and the PP wants to respect mathematical practice. However, for her view to be truly non-arbitrary, the PP must be committed to there being a distinct object for each mime. So, if Plenitudinous Platonism were true, there would be no non-trivial identities (any two concrete items, like marks on a page, would be best mimes of distinct abstract objects, because each *could* correspond to a different target). The PP could bite the bullet and claim that there really is a distinct abstract object for every concrete mime; but at the cost of falsifying mathematical claims about non-trivial identities.

If the PP does bite the bullet, then for her, a mimetic translation shows that, for any two mimes, in addition to their distinct intended targets, there is a third abstract object that they are equally best mimes of: the object that shares their structure, but lacks their what-its-like properties. The PP must be committed to the existence of this object, as it is possible. So, if the PP were considering the case of the drawing vs. the formula, she should think: the drawing is the best mime of object *a*; the formula is the best mime of *a'*; and they are equally best mimes of a third object, *a''*.

Which of these three objects is *the equilateral triangle?* Perhaps the answer is indeterminate. It turns out that each of our concrete objects is a best mime of many abstract objects, despite our intentions. In each case, one of those abstract objects—for example, the one that the formula and the drawing equally best mime—is the same object that I claim that *the equilateral triangle is*. So: the PP needn't accept the earlier arguments in this paper; but she does need to posit many more abstract objects than she might have initially thought, including exactly the objects that I claim *Miserere* and *the equilateral triangle* are. For the PP, the arguments I have given thus far might then be taken to suggest that it is those objects that 'Miserere' and 'the equilateral triangle' typically refer to.

I conclude that our intentions can't matter to what the best mime of a target is. Only the PP can take our intentions to matter; but the PP, I claim, must also believe in the objects I've argued *Miserere* and *the equilateral triangle* are (and perhaps I have supplied her with a reason to think that they are the best candidates for being the referents of 'Miserere' and 'the equilateral triangle').

[11] This view was first advocated by Balaguer (1998), who argues that, if Platonism is true, then Plenitudinous Platonism must be true.

One might think that *Miserere* is special here, because it is a creative artifact, and perhaps its nature is determined by the intentions of its creator. While it might be that certain things about the natures of creative artifacts are determined by the intentions of their creators, I am assuming that, for the platonist, it is a mistake to think that creators can determine the properties of the abstract object (especially the kinds of properties under discussion here) itself as opposed to its intended tokenizations. For example, I might create a piece of music that I think is essentially sonic—but depending on what abstracta are like, and what the property of sonicity is like, it might be impossible for it to be sonic at all (if, for example, sonicity is a property that necessarily involves physical soundwaves). I can intend all I like to make an abstract object sonic, but if doing so is metaphysically impossible then my intentions won't make a difference (even if they somehow make a difference to other features of the object). I don't have the space to fully defend this assumption, and readers can certainly choose to get off the boat here (though note that the argument would still apply to *the equilateral triangle*).

Perhaps there are other potential defeaters for our translations. But I do not know what they would be. It seems to me that the only kind of evidence we could have for thinking that one mime was more mimetic would have to come from the abstract object itself, independent of its mimes. We lack access to the abstract object independent of its mimes. So we lack a way to decide which mime is best.

4. What Should We Conclude?

I have argued that, if the drawing is a best mime of *the equilateral triangle,* then so is the formula; and likewise for the performance, the score, and *Miserere*. To conclude, I'll briefly motivate two further claims. First, given that the score and the performance (and the drawing and the token formula are) don't seem to share any what-it's-like properties, we should think that the abstract objects in question have no what-it's-like or quasi-what-it's-like properties. Second, *Miserere* and *the equilateral triangle* do have structural (non-what-it's-like) properties, and hence have internal natures. From these two claims, it follows that these two objects (and objects relevantly like them) have internal, non-what-it's-like, structural natures. Call this claim **internal structuralism**.

Internal structuralism contrasts with three kinds of views one might have about abstracta: (i) **what-it's-like-ism***:* there is some sense in which abstract objects in question instantiate what-it's-like properties, or abstract "analogues" of what-it's-like properties. (According to the what-it's-like-ist, for example, there is some sense in which Sherlock Holmes has the property of enjoying cocaine.) (ii) **pointillism***:* abstract objects are featureless/structureless points in abstract "space". And (iii)

relationalism: abstract objects are mere nodes in an abstract structure that is, in some sense or other, prior to them.[12]

Recall MIMESIS:

> **MIMESIS:** Let *a* be any object, and *m* be any potential candidate mime of *a*. **If**: we have some *m, m', m''…* that are candidates for being best accessible-to-us mimes of *a,* and no good reason to think any of the *ms* is better than the others, **then:** we should not believe that *a* has any properties that predicate-share with *some but not all of* the *ms*.

The perfect drawing and the algebraic formula are two *m*s that satisfy, for the object *the equilateral triangle,* the antecedent of MIMESIS, and likewise for a written copy of the score of *Miserere* and a performance of it.

I'll focus on the triangle case: according to MIMESIS, we should not believe that *the equilateral triangle* predicate-shares with any properties that the drawing has but the formula lacks; and vice-versa. So we should not believe that *the equilateral triangle* is triangular, or that it has three sides. Nor should we believe that it has abstract analogues of these properties, if we think that they are restricted to the concrete realm. Hence, there is no sense in which we should think that either *the equilateral triangle* is triangular, or that it is trilateral. There is not even some abstract sense in which the equilateral triangle is triangular.

Further, it is unclear what properties the formula and the drawing could share that would justify us attributing any what-it's-like or quasi-what-it's-like properties to *the equilateral triangle.* (Perhaps being spatially extended is a property that the written formula and the drawing share? No; an utterance of the formula, or a mental picture, is also a best mime…) If so, then we might think that it follows that *the equilateral triangle* has no what-it's-like or quasi-what-it's-like nature.

But this is too quick. The consequent of MIMESIS is about what we shouldn't believe; MIMESIS is a negative epistemic principle. It follows that there is another option besides rejecting what-it's-like-ism: we might instead claim that the objects do have what-it's-like or quasi-what-it's-like properties, but that we simply can't know which of those properties they have.

There are two versions of this claim. First, **epistemicism:** while we can't know which what-it's-like properties the abstract object has, we can know that they are among its mimes' properties, and hence we do know that its properties are among the disjunction of the what-it's-like properties of its mimes. So, e.g., we can know that *the equilateral triangle* is either like the pictorial mime or the algebraic

[12] This is not the normal way to carve up views in the ontology of abstracta. But most platonist views fall into one of these categories. Dodd (2007, ch. 2), for example, argues that musical works are unstructured types, and so is a pointillist. Zalta (1983) gives a 'two modes of predication' view on which 'x is F' is ambiguous between x encoding and exemplifying F. So he is a what-it's-like-ist. Thomasson (1999) seems to be a what-it's-like-ist about e.g. fictional characters, which she takes to be contingent abstracta that depend on their creators. Resnik's (1997) structuralism seems to involve a commitment to both relationalism (about "patterns" and pointillism (about objects).

one (or some other mime that is tied with the two). Second, **alienism**: the abstract object does have what-it's-like properties, but that they are all alien properties—what-it's-like properties that aren't instantiated in our concrete world (and that aren't easily imagined by us, e.g. are not things like having mass of exactly 2.7690543… grams).

If our options are epistemicism, alienism, or my negative thesis (that *the equilateral triangle* has no what-it's-like properties), then we should believe my negative thesis.

One vice that a theory can have is to commit us to worldly arbitrariness: to commit us to there being a fact of the matter that p, without there being any reason why p rather than q. Epistemicism is committed to arbitrariness, and for that reason alienism seems preferable. If we have a theory that commits us to arbitrariness, and another that is just as theoretically virtuous except that it does not, we should prefer the latter. I think that alienism does this, and so all else equal we should prefer alienism to epistemicism.

However, epistemicism and alienism have another vice: each theory commits us to the properties of abstracta being unknowable. Epistemicism allows us to know large disjunctive truths about abstracta, but not which disjunct is true. Alienism tells us that the nature of abstract objects is wholly mysterious to us (since their natures involve properties we could not know about). Call the general vice of unknowability (that captures both kinds of unknowability) 'ineffability'.

Sometimes arbitrariness or ineffability is taken as a knock-down reason to reject a theory. I don't think this is right. Both of these vices should be traded off and weighed against other virtues and vices. If we have a theory that commits us to ineffability or arbitrariness, but is significantly more theoretically virtuous than other theories, then we should adopt it. If, however, we have a theory that does all the same work, that is equally (or more) parsimonious, elegant, informative, explanatorily powerful, unifying, etc., and rids us of this ineffability, then we should adopt that one.

And there is a better theory: the one I am putting forward here. I have argued that *the equilateral triangle* has no what-it's-like properties. Momentarily, I will argue that we know that *the equilateral triangle* does have structural properties, which are distinct from those of, e.g. *the square*; and similarly for *Miserere* (which has distinct structural properties from *Row Your Boat*). Together, these entail that *Miserere* and *the equilateral triangle* are **internally structural objects:**

> **internally structural objects** are objects have purely structural, internal natures: they have intrinsic properties that distinguish them from other abstract objects, but none of those properties are what-it's-like properties.

Internal structuralism (the view that some objects, in this case *Miserere* and *the equilateral triangle* are internally structural) is just as explanatorily powerful as alienism and epistemicism, but is more virtuous. First, it is more parsimonious than alienism (doesn't posit the properties that alienism does); second, it doesn't commit us to ineffability as both alienism and epistemicism do; and third, it does not commit us to arbitrariness as epistemicism does. So, I claim, we should adopt it.

I want to note that readers could get off the boat here (by denying that arbitrariness or ineffability are theoretical vices, two assumptions that I have not defended here) if they wanted to embrace either alienism or epistemicism. I argue for these assumptions in (reference redacted). Further, they are commonly taken to be theoretical vices, and are often taken to be reasons to reject theories altogether even if an alternative theory is not available.[13]

The defense of internal structuralism is also incomplete, so I will now defend the "internal" part of internal structuralism: the claim that *Miserere* and *the equilateral triangle* have internal structure, that is, they have at least some structural (non-what-it's-like) properties intrinsically.

The ingredients for this defense are all contained in §2, where I claimed that the rendition of *Row Your Boat* is not as mimetic of *Miserere* as a near-perfect performance is, nor as a copy of the written score is. If this is right, it follows that *Miserere* has a nature, and, given the rest of what I have argued, a structural one.

Even if we embed the child's rendition of *Row Your Boat* in a complex structural system (trying to treat it, essentially, as a "mere node" in a larger structure), it still does not approach being a best mime of *Miserere*. My account of mimetic translation explains this: embedding *Row Your Boat* in a complex structural system in a way that attempts to make it relate to other things in the same way that *Miserere* does will violate at least one of the three constraints I propose on mimetic translation. The best explanation of this is that *Miserere* has an internal structure that its best mimes succeed in preserving.

So pointillism and relationalism are false of *Miserere* and of *the equilateral triangle,* because both objects have intrinsic internal structure. What-it's-like-ism is also false of these objects, because they don't have any what-it's-like (or quasi-what-it's-like) properties—and specifically, there is no sense in which *the equilateral triangle* is triangular, nor any sense in which *Miserere* is sonic.

Are all abstract objects internally structural? Maybe not. I suspect that words are; they have structure, but we could run a similar argument on them. Other cases look harder: if non-eliminative ante rem mathematical structuralism is true, then the number nine lacks an internal nature—relationalism is true of it, so it doesn't have an internal nature. Or consider *the tiger*. Perhaps all of the best mimes of *the tiger* have stripes (they are the most 'tigery' tigers).[14] If this is true, then I haven't provided an argument that *the tiger* has no stripes, and what-it's-like-ism might be true of it. However these questions get settled, though, many abstract objects will turn out to have purely structural

---

[13] e.g. with respect to arbitrariness, Parfit (1984) (re: fission); Benacerraf (1965) (re: reducing numbers to sets); (Jubien 2001, Cowling 2017, Keller 2019) (re: Benacerraf-style arguments to propositions); Kripke (1980) (re: possibilism about fictional entities). Sider (2011, ch. 10) also discusses arbitrariness and arrives at a more moderate conclusion, though I note that he takes the burden to be on him to argue *against* arbitrariness being a reason to reject a metaphysical view.

[14] The thought here is that there might not be anything playing the role, for *the tiger,* that the token score plays for *Miserere.*

internal natures.

I have argued for two theses: first, that *Miserere* and *the equilateral triangle* lack what-it's-like or quasi-what-it's-like properties, and so (e.g.) *the equilateral triangle* is in no sense triangular. And second, these objects *do* have natures: we can know enough about them, by knowing which of their mimes are better and worse, to know this. Since they don't have what-it's-like or quasi-what-it's-like natures, but they do have internal natures (they can't be mere nodes in a larger structure, and they have intrinsic properties), they must have purely structural internal natures. As I've pointed out, though, my argument does not obviously generalize to other abstract objects. In order to develop a full theory, I would need to both consider whether MIMESIS will always deliver informative results when applied to abstracta, and also to look at the results of similar arguments (applications of MIMESIS to other kinds of abstract objects). Consider this, then, a first step in constructing a theory of abstract objects from the ground up.

References

Armstrong, D.M. Nominalism and Realism (Universals and Scientific Realism, Volume I). Cambridge. 1978.

Benacerraf, P. 'What Numbers Could Not Be', *The Philosophical Review* 74 (1) (47-73). 1965.

Balaguer, M. Platonism and Anti-Platonism in Mathematics. Oxford University Press. 1998.

Cameron, R. 'Intrinsic and Extrinsic Properties', in Le Poidevin (ed.), *The Routledge Companion to Metaphysics*. Routledge. 2009.

Cowling, S. Abstract Entities. Routledge. 2017.

Dodd, J. Works of Music: An Essay in Ontology. Oxford University Press. 2007.

Dretske, F. Knowledge and the Flow of Information. MIT Press. 1981.

Dummett, M., Frege: Philosophy of Language. Duckworth. 1973.

Francescotti, R. (ed.), Companion to Intrinsic Properties. De Gruyter. 2014.

French, S. 'A Model-Theoretic Account of Representation (Or: I Don't Know Much About Art… But I Know It Involves Isomorphism)'. *Philosophy of Science* 70 (5) (1472-83). 2003.

French, S. and Ladyman, J. 'Reinflating the Semantic Approach'. *International Studies in the Philosophy of Science* 13 (2) (103-21). 1999.

Giere, R.N. 'How Models are Used to Represent Reality'. *Philosophy of Science* 71 (5) (742-52). 2004.

Goodman, N. and Quine, W.V.O. 'Steps Towards a Constructive Nominalism'. *The Journal of Symbolic Logic* 12 (4) (105-122). 1947.

Hale, B., Abstract Objects. Basil Blackwell. 1987.

Hale, B. Necessary Beings: An Essay on Ontology, Modality, and the Relationship Between Them. Oxford University Press. 2013.

Jubien, Michael. 'Propositions and the Objects of Thought', *Philosophical Studies* 103, pp. 47-62. 2001.

Keller, L. J., 'What Propositional Structure Could Not Be', *Synthese* 196 (4) (1529-1533). 2019.

Kleinschmidt, S. and Ross, J. 'Repeatable Artwork Sentences and Generics' in C. Mag Uidhier, ed., *Art and Abstract Objects* (125-156). Oxford University Press. 2012.

Kripke, S. *Naming and Necessity.* Oxford: Blackwell. 1980.

Langton, R. and Lewis, D.K., 'Defining 'Intrinsic''. *Philosophy and Phenomenological Research* 58 (2) (333-345). 1998.

Lewis, D.K., On The Plurality of Worlds. Basil Blackwell. 1986.

McDaniel, K. The Fragmentation of Being. Oxford University Press. 2017.

Meinwald, C. Plato's Parmenides. Oxford University Press. 1991.

Parfit, Derek. Reasons and Persons. Oxford University Press. 1984.

Pasnau, R. and Shields, C. The Philosophy of Aquinas. Westview Press. 2004.

Pelletier, F.J. and Zalta, E.N. 'How to Say Goodbye to the Third Man'. *Nous* 34 (2) (165-2002). 2000.

Resnik, M. Mathematics as a Science of Patterns. Oxford University Press, 1997.

Shapiro, S. Philosophy of Mathematics: Structure and Ontology. Oxford University Press. 1997.

Sider, T. 'Intrinsic Properties'. *Philosophical Studies* 83 (1-27). 1996.

Sider, T. *Writing the Book of the World.* Oxford University Press. 2011.

Thomasson, A. *Fiction and Metaphysics.* Cambridge University Press. 1999.

Van Frassen, B. Scientific Representation: Paradoxes of Perspective. Oxford University Press. 2008.

Wolsterstorff, N. Works and Worlds of Art. Clarendon Press. 1980.

Wright, C. Frege's Conception of Numbers as Objects. Aberdeen University Press. 1983.

Zalta, E.N. Abstract Objects: An Introduction to Axiomatic Metaphysics. Dordrecht: D. Reidel. 1983.