

ABSTRACT

In general, properties realize certain roles in the workings of nature. For example, *mass* makes objects resist acceleration. But what is the relationship between these roles and the properties that realize them? According to ‘quidditism’, the roles are *contingently* realized by the properties that in fact realize them. Opponents charge that quidditism implies the existence of epiphenomenal and unknowable “quiddities” or “inner natures”. The purpose of this dissertation is to argue in favor of quidditism and explore its epistemic and pragmatic consequences.

I begin by showing that certain considerations from physics show that properties cannot be individuated via their nomological roles, as anti-quidditism maintains. I then argue that quidditism can be had without epiphenomenal and unknowable quiddities. The arguments to the contrary proceed either from a misconception of what quidditism is, or they rely on an implicit assumption to the effect that quidditism implies that there is *something* in virtue of which properties are individuated from one another.

I turn next to an examination of quidditism’s epistemic implications. Here I develop and defend a posthumously published paper by David Lewis. Following Lewis, I argue that quidditism implies that we will never know which properties realize at least some of the nomological roles. At the core of my case is an argument against the common claim that quidditist skepticism is merely a species of traditional external-world skepticism.

Finally, having argued that quidditistic skepticism importantly differs from traditional external-world skepticism, I explore which, if either of these two types of skepticism have *practical* implications—that is, implications for how we ought to live our lives. I begin by getting clearer on the distinction between traditional external-world

skepticism and quidditistic skepticism, showing that the former is of a kind I call ‘Cartesian skepticism’, while the latter is of a kind I call ‘Kantian skepticism’. A skepticism is Cartesian, on my terminology, if it posits ignorance of *graspable* facts; while it is Kantian if it posits ignorance of *ungraspable* facts. Finally, I argue that, in general, Cartesian skepticisms have drastic practical implications while Kantian skepticisms have none at all.

**INTRODUCTION:
WHAT IS QUIDDITISM?**

What is the relationship between a property and its role in the workings of nature? For example, what is the relationship between *mass* and its role of making objects resist acceleration? Is *mass* inextricably linked with its role, or is there a sense in which the two are distinct? According to one view, properties and their roles can in some sense ‘come apart’. The unfortunate name for this thesis is ‘quidditism’, and its recently come in for some heavy criticism.

Now I can’t do anything about quidditism’s awkward name, but this dissertation is largely an attempt to defend its reputation. In the chapters that follow I will be primarily interested in three questions: (1) Is quidditism *true*? (2) Does quidditism have any *epistemic* implications? And (3) do the epistemic implications of quidditism, if there are any, have any *practical* significance? These, respectively, are the questions of the three chapters that follow, and my answers will be ‘yes’, ‘yes’, and ‘no’.

But before we get to these topics, we need to know more precisely what quidditism is. Although the term has been in the literature since David Armstrong’s (1989), it has been less than clear what quidditism is supposed to be. In particular, there seems to be a tension between the ways that ‘quidditism’ has been *defined* and the way that ‘quidditism’ has been *used*. Getting clear on these matters is the primary task of the present chapter. At the end, there will be a brief outline of the chapters to come.

I said that the rough idea behind quidditism is that properties in some sense ‘come apart from’ their roles in nature (hereafter, their ‘nomological roles’). To make this

precise, we need to answer two questions: what do we mean by ‘a property’s nomological role’, and what do we mean when we say that properties ‘come apart from’ their nomological roles?

Section One: What do we Mean by ‘a Property’s Nomological Role’?

This question is relatively easy. To take a simple example, suppose we live in a Newtonian world. Then the property *mass* figures in the law that between any two objects with mass there is force proportional to the product of their masses divided by the square of the distance between them (Newton’s Law of Gravitation). It also figures in the law that the net force on a given object is equal to the product of its mass and acceleration (Newton’s Second Law of Motion). The nomological role of *mass* is simply the role of being the property x such that between any two objects with x there is a force proportional to the products of their respective amounts of x , divided by the square of the distance between them, and the net force on any object is equal to the product of its amount of x and its acceleration.

More generally, when we say that a property plays a certain nomological role, we mean that it is the property that does such and such, where the ‘such and such’ is defined in the following way. Start by considering each fundamental natural law in which the given property figures. Now conjoin the statements of each of these laws into one long sentence. Now replace each occurrence of the word for the property in question with a free variable. That gives us an open sentence ‘... x ...’. The role of the given property is thus the role of being the property x such that... x

Every definition of ‘quidditism’ that I know of agrees that roles are to be defined in this way. The more difficult question is the second question mentioned above: what do we mean when we say that a property ‘comes apart from’ its role? I will be answering this question in the course of the next two sections. My strategy will be to begin with the most

widely cited definition of ‘quidditism’, which I will then tweak, step by step, into an acceptable definition.

Section Two: Paying Our Respects to Black’s Definition

One common way for philosophers to precisify rough ideas about distinctness, or about one thing’s ‘coming apart from’ another, is in terms of what *could have been*. Moreover, philosophers have tried to understand talk about what *could have been* in terms of *possible worlds*. (For now, simply think of a possible world as a maximally specific way that things could have been.) Putting these two ideas together, philosophers often formulate theses about distinctness directly in terms of possible worlds.

The most widely cited,¹ and the first *explicitly* given definition of ‘quidditism’,² does exactly this. In Robert Black’s ‘Against Quidditism’ (2000) he writes

[According to one view,] nothing *constitutes* the fact that a certain quality playing a certain nomological role in that world is identical with a certain quality playing a different role in ours; they just are the same quality, and that's all that can be said. Since the fashion is to mine the Scottish tradition for technical terms in this area, let us use the word 'quidditism' for the acceptance of primitive identity between fundamental qualities across possible worlds. (p.92)

According to this definition, quidditism is the thesis that when a property that plays a certain nomological role in one possible world is identical to a property that plays a

¹ For example, Stephen Mumford (2004: 151) and David Lewis (2008: 7) both cite Black’s definition. And although she doesn’t cite Black, Ann Whittle (2006) defines quidditism as ‘the acceptance of primitive property identity across possible worlds’ (p. 463).

² Although Armstrong (1989) uses the term ‘quidditism’, he never explicitly defines it. It is clear from the context, however, what he means by it. Armstrong’s (implicit) definition of ‘quidditism’ is discussed at length in section four of the present chapter.

different nomological role in another possible world, there is nothing that *constitutes* the fact that the one property is *the same property* as the ‘other’.

Notice that, taken quite literally, quidditism on Black’s definition is compatible with the view that it is *never* the case that a property playing one role in one world is identical with a property playing a different role in a different world—it merely says that *when* that happens, there is nothing that constitutes the fact that the one property is identical with the other. Indeed, if that were never the case, then quidditism, on Black’s definition, would be vacuously true.

Almost certainly Black did not intend his definition to be taken quite so literally. Rather, he intended quidditism to be the thesis that there *are* cases in which a property playing one role in one world is identical with a property playing a different role in another world, and when this happens there is nothing that constitutes the fact that the one property is identical with the other. Let’s call this

Black’s Definition. Sometimes, a property that plays one role in one world is identical with a property that plays a different role in a different world, and when this happens there is nothing that constitutes the fact that the one property is identical with the other.

Formulating Black’s definition this way, we can see that the second conjunct—the conjunct which we originally took to be *all* of quidditism on Black’s definition—is actually quite trivial. Suppose that a property playing a certain nomological role in one world really is *identical* with a property playing a certain nomological role in another world. When we ask whether something constitutes this fact, what are we asking? Taken literally, we are asking this: given a property P that plays one role in one world and a different role in another world,

1. What constitutes the fact that $P = P$?

That the answer to this question should be ‘nothing’ seems quite obvious. As Lewis (1986) has put it,

There is never any problem of what makes something identical to itself; nothing can ever fail to be. And there is never any problem about what makes two things identical; two things never can be identical. (pp.192 – 93)³

We should say exactly this in response to Black’s suggestion that the debate over quidditism should be understood, in part, as a debate over whether there is anything that *makes* (i.e. constitutes the fact that) a property in one world is identical with a property in another. As it is with respect to individuals, the answer to *that* question is easy: nothing *makes* something identical with itself, nothing *constitutes* the fact that a property is identical with itself.

The first, and implicit, part of Black’s definition of ‘quidditism’ is the important part: there is a non-trivial debate over whether a property in one world is *ever* identical with a property playing a different role in different world. According, let us drop the second conjunct from Black’s definition and call the resulting definition

The Transworld-identity Definition. Sometimes, a property that plays a certain nomological role in one world is identical with a property that plays a different nomological role in a different world.

³ Lewis is here responding to the suggestion that the debate over haecceitism (which is to individuals as quidditism is to properties—see below) is a debate over whether there is something that *makes* an individual in one world identical with an individual in another, supposing that there are such individuals. He is responding by saying that the answer to *that* question is easy: nothing ever *makes* something identical with itself.

Now we're getting closer. But not close enough.

First, this definition makes quidditism *seem* like a thesis about *identity*—in particular, about transworld identity. But this puts the emphasis in the wrong place. Here we can take another page from Lewis's (1986) book.

We *do* state plenty of problems in terms of identity. But we *needn't* state them so. Therefore, they are not problems about identity. Is it ever so that an F is identical to a G? That is, is it ever so that the same thing is an F, and also a G? More simply, is it ever so that an F is a G? The identity drops out. (p. 193)

It's misleading to say that we are interested in the question of whether one property playing one role in one world is ever *identical* with a property playing a different role in another world. To make the proper emphasis clear, we should follow Lewis's advice and let 'the identity drop out'. Let's call the resulting definition of 'quidditism'

The Same-property/Different-role Definition. Sometimes, a property that plays one nomological role in one world plays a different nomological role in a different world.⁴

I think that the same-property/different-role definition successfully captures something one might mean when he says that properties 'come apart from' their nomological roles.

Unfortunately, this definition is still unsatisfactory, and for two reasons.

Here's the first reason. Above I suggested that you think of possible worlds, for the time being, as maximally specific *ways* that things might have been. But not all theorists agree that possible worlds should be understood in this way. So-called *modal realists*

⁴ Note that this definition of 'quidditism' is equivalent to the former. Their different names merely reflect the difference in emphasis. (The identity doesn't really drop out; the *mention* of identity drops out.)

claim that possible worlds are things just like the actual world: they are hunks of concrete space and time (or spacetime) that may or may not contain concrete material objects or physical fields. Unfortunately, such a view makes trouble for the same-property/different-role definition of quidditism. And although modal realism is by no means the most popular view of possible worlds, we need to make room for it.

To see how modal realism causes trouble for our above definition of ‘quidditism’, suppose you’re a modal realist and you think that properties are ‘world-bound’ entities—that is, suppose you think that one and the same property is never present at more than one possible world. In that case, ‘quidditism’ on the same-property/different-role definition will be trivially false: since *mass* exists only in this world, there is no world where *mass* plays a different role. Still, you might think that the proper analysis of sentences like ‘*Mass* could have played a different role’ is in terms of counterparts: ‘*Mass* could have played a different role’ is true if and only if there is a possible world in which a *counterpart* of *mass* plays a role distinct from the role played by *mass* in this world. But to say that *mass* could have played a different role is certainly to say that *mass* is in some sense distinct from its role. Thus, although the same-property/different-role definition seems to force this kind of a modal realist into rejecting quidditism, there is a very important sense in which such a modal realist can still claim that properties come apart from their roles.

No later than we suggest this problem do we see its solution. It seems that ‘quidditism’ ought to be defined as follows:

The Anti-supervenience Definition. For any property P and its nomological role R, P could have played some role other than R.

According to this definition, quidditism is the view that nomological roles do not *supervene* on the properties that realize them. To say that one thing supervenes on another is just to say that things could not have been different with respect to the former unless they were different with respect to the latter. On the above definition, quidditism is the view that

things *could have been* different with respect to which role is realized *without* having been different with respect to which property realized the role.

Here we choose to define quidditism not in terms of possible worlds, but simply in terms of what could have been. This leaves room for modal realists and others to disagree about how best to understand talk of what could have been, without immediately forcing them to accept/reject quidditism. The modal realist who thinks properties are world-bound and uses counterparts to evaluate possibility claims, for example, can either accept quidditism, saying that there is a world in which a counterpart of *mass* plays a role distinct from *mass*'s actual role, or reject quidditism, saying there is no such counterpart of *mass*.

Now we're getting pretty close to a workable definition of 'quidditism'. Still, I said that there is a second problem for the same-property/different-role definition, and that problem is not avoided by moving to the above anti-supervenience definition. I discuss the problem and its solution in the next section.

Section Three: Haecceitism, Quidditistic Skepticism, and a New Definition

To see why there is a problem for the above definition of 'quidditism', I want to briefly discuss the analogy between quidditism and haecceitism. Haecceitism, quite roughly, is the view that individuals 'come apart from' their qualitative characters. There have been several attempts to make haecceitism more precise, but most theorists now follow Lewis (1986) in defining 'haecceitism' as an anti-supervenience thesis.

The Anti-supervenience Definition of 'Haecceitism'. For any individual A and its qualitative character C, C might have been had by some individual other than A.⁵

⁵ This is almost Lewis's (1986) definition of 'haecceitism'. A bit more precisely, Lewis defines 'haecceitism' as the view that there are at least two possible worlds w_1 and w_2 such that (1) w_1 and w_2 are qualitatively exactly alike, and yet (2) there is some actual individual A such that the counterpart of A in w_1 has a different

Quidditism, it is often said, is meant to be to properties as haecceitism is to individuals. And the above anti-supervenience definition of ‘quidditism’ preserves this analogy quite nicely. But not perfectly: while haecceitism denies that individuals *supervene on* their qualitative characters, quidditism, on the above definition, denies that a property’s *role* supervenes on *it*. In other words, the *direction* of supervenience denied by haecceitism is not analogous to the *direction* of supervenience denied by quidditism.

Of course, this disanalogy is little reason to reject one definition of ‘quidditism’ in favor of another. Why should we take the alleged analogy between haecceitism and quidditism so seriously? After all, the two theses are only supposed to be analogous in so far as they both hold that one type of thing ‘comes apart from’ another type of thing, and the phrase ‘comes apart from’ is quite ambiguous: *one way* to make it precise is to deny supervenience in one direction; *another way* is to deny supervenience in the other direction. Why not simply rest content with having haecceitism deny supervenience in one direction and quidditism deny it in the other? Fortunately, we have good reason to take the analogy seriously: the way quidditism has been *used*.

The primary use of quidditism within the philosophical literature is to support a particular epistemological position known as ‘quiddistic skepticism’. Roughly, quiddistic skepticism is the view that we will never come to know *which properties* realize the fundamental nomological roles realized at our world. David Lewis (2008) is the champion of quiddistic skepticism, and his argument can be summarized as follows.

If there are facts about which properties realize which roles—that is, facts over and above the mere facts about which roles are realized—then those fact are beyond the realm of empirical inquiry. Let us call these worrisome facts, if there are any,

qualitative character than the counterpart of A in w_2 . Are these definitions equivalent? I think they are, but the difference between them, if there is any, needn’t concern us here.

‘quidditistic facts’. According to quidditism, there *are* quidditistic facts, because quidditism is the view that properties ‘come apart from’ their roles. But quidditistic facts are contingent facts about our world and, hence, if they are beyond the realm of empirical inquiry, then they are simply unknowable.

Of course, this line of reasoning needs much explaining and defense, and this is *the* task of Chapter Two. For now, let us simply focus on the role of quidditism in the argument, which is to imply that there are facts *beyond* the mere facts about which roles are realized—namely, there are facts about *which* property realizes those roles. Our question is thus this: does quidditism, on the anti-supervenience definition given above, imply that there are such facts?

Unfortunately, it does not. Consider *mass* and its nomological role, call it ‘R’. Suppose that *mass* could have realized some role distinct from R. If so, then quidditism, on the anti-supervenience definition above, is true. Still, it is consistent with this that, necessarily, if R is realized, then R is realized by *mass*. (Just imagine a possible world in which R is not realized and *mass* or its counterpart realizes some role other than R.) Now consider (1) the proposition that R is realized and (2) the proposition that *mass* realizes R. Since at every possible world where R is realized, it is realized by *mass* (we are supposing), (1) necessarily implies (2). And since at every possible world where *mass* realizes R, R is realized (trivially), (2) necessarily implies (1). Hence, propositions (1) and (2) are necessarily equivalent. Thus, if facts are individuated no more fine-grainedly than necessarily equivalent propositions, then the fact that *mass* realizes R *just is* the fact that R is realized.⁶ In that case, there is no fact about which property realizes R that is *beyond* the mere fact that R is realized—that is, there is no quidditistic fact about the role R.

⁶ What if facts are individuated more finely than necessarily equivalent propositions? In that case, there might still be a fact about which property realizes R that is beyond the mere fact that R is realized. But if there is such a fact, its existence has nothing to do with quidditism on the above definition. Again, we’re in

What we have just seen is that denying that roles supervene on the properties that realize them does imply that there are quiddistic facts. But now the solution should be coming into view: to get quiddistic facts we need to deny that *properties* supervene on their *roles*.

The Revised Anti-supervenience Definition. For any property P and its nomological role R, R might have been realized by some property other than P.

On this definition, quidditism clearly implies that there are quiddistic facts. If there is a possible world where R is realized but not realized by *mass* or a counterpart of *mass*, then propositions (1) and (2) are non-equivalent—the proposition that *mass* realizes R is *more specific* than the proposition that R is realized. Thus, if facts are individuated at least as finely as propositions, there is a fact about *which* property realizes R that is beyond the mere fact that R is realized.

We are now several steps removed from Black’s definition of ‘quidditism’. Nonetheless, the above definition is to be preferred to all the others. Unlike Black’s original definition, it makes quidditism non-trivial; unlike the transworld-identity definition, it captures the intuitive idea of properties ‘coming apart from’ their roles; unlike the same-property/different-role definition, it’s as congenial to those who think that properties are world-bound as those who don’t; and unlike the original anti-supervenience definition, it makes quidditism imply that there are quiddistic facts. In addition to all this, the revised anti-supervenience definition makes the analogy between quidditism and haecceitism (on Lewis’s definition) just about as tight as it can be. For these reasons, I henceforth adopt the revised anti-supervenience definition as my official definition of ‘quidditism’.

search of a definition of ‘quidditism’ according to which *quidditism* plays a key role in implying the existence of such facts.

We arrived at our preferred definition of ‘quidditism’ by starting with Black’s definition and working out the kinks. But there is another place we might have started in search for a definition, and it wouldn’t do justice to our topic not to consider that starting point here.

Section Four: Quiddityism

I noted above that although Black (2000) was the first to explicitly define ‘quidditism’, he was not the first to use the term. That title would seem to go to David Armstrong (1989), who first used the term ‘quidditism’ as follows.

Haecceitism for individuals is parallel to Quidditism for universals. Quidditism for universals seems very plausible. Each universal must surely have its own nature.
(p. 59)

Although Armstrong does not explicitly *define* ‘quidditism’, it’s clear from the context how he understands it:

Armstrong’s Definition. Each property has its own nature.⁷

Armstrong had previously introduced the term ‘quiddity’, which he understands as a property’s ‘nature’ (p. 44, 55). So for Armstrong, quidditism is simply the acceptance of quiddities.

What might Armstrong mean here by ‘nature’ or ‘quiddity’? Clearly, he cannot take a property’s quiddity to be its nomological role. If we understood ‘quiddity’ or

⁷ If not all properties are ‘genuine universals’, I hereby stipulate that I use the term ‘property’ to refer only to the latter.

‘nature’ in that way, then quidditism on Armstrong’s definition clearly wouldn’t capture anything one might mean when she says that properties ‘come apart from’ their roles.

Much more plausibly, *quiddity* is to be understood on analogy with *haecceity*. Indeed, on p. 44 Armstrong writes, ‘[P]roperties and relations do have their own haecceity, or, better, their quiddity or nature.’ A haecceity is supposed to be an individual’s *thisness*, where an individual’s *thisness* is a non-qualitative property that serves to individuate that individual from all other individuals. If we like, we can say it is the property of being *this individual*. Now Armstrong denies that there are such properties as haecceities. But he seems to think that we have a clear enough grasp of what a haecceity *would be* to understand, by analogy, what a quiddity is supposed to be. Taking the analogy seriously, let’s say that a quiddity is a property’s *suchness*, where a property’s *suchness* is a non-role-involving (second-order) property that serves to individuate that property from all other properties.

Many authors have understood ‘quiddity’ in exactly this way, and many have followed Armstrong in taking quidditism to be the acceptance of quiddities. These authors accept a definition of ‘quidditism’ that I’ll call

Quiddityism. Each possible property has its own *quiddity*, where a property’s *quiddity* is a non-role-involving (second-order) property that serves to individuate that (first-order) property from all other properties.

What then is the relationship between quidditism, on our official definition, and quiddityism? Some authors seem to have treated them as though they were equivalent, or as though one entails the other. In the next chapter I argue that neither is the case: quiddities can be had without quidditism, and, even more importantly, quidditism can be had without quiddities. I think that much of the resistance to quidditism has stemmed from an understandable fear of ‘quiddities’. But if what I say in Chapter One is correct, fear of quiddities ought not become fear of quidditism.

Section Five: Chapter Outline

Having settled on an acceptable definition of ‘quidditism’, it’s time to get down to brass tacks. I begin in Chapter One by defending quidditism against its two most common objections. The first objection is an argument from how properties are individuated. According to this argument, properties are individuated by their nomological roles, and hence any property that realizes, e.g., the role of *mass* must therefore be *mass*. Against this I argue that, given the possibility of certain symmetries in nomological role, properties simply cannot be individuated by their roles. The second common argument against quidditism charges that quidditism implies a certain kind of epiphenomenalism. This argument, I contend, confuses quidditism with quiddityism.

In Chapter Two I turn to the question of whether quidditism has any epistemic implications—in particular, whether it has any *skeptical* implications. Here I develop and defend David Lewis’s posthumous ‘Ramseyan Humility’ (2008), where Lewis argues for quiddistic skepticism. My discussion here is aimed at getting clear on exactly what quiddistic skepticism is, why Lewis thinks quiddistic skepticism follows from quidditism, and how quiddistic skepticism differs from more familiar types of skepticism. The last of these is especially important, as several philosophers have charged that quiddistic skepticism is akin to traditional skepticism about the external world (and can be refuted as such). I argue that this is not the case.

In Chapter Three, the final chapter, I turn to the question of whether quiddistic skepticism has any practical significance. But I do so only within the context of a much larger discussion. Chapter Three is really a *general* discussion of the practical significance of all types of skepticism and thus ought to be of interest even to those epistemologists who couldn’t care less about quidditism. My strategy in this chapter is to first make two, cross-cutting, mutually exclusive and exhaustive distinctions between types of skepticism. I then

argue that skepticisms of only one of these types have practical implications. As it happens, quiddistic skepticism is of one of the types *without* practical significance.