UCLA School of Law

Law & Philosophy Program

Legal Theory Workshop
UCLA School of Law

Mark Schroeder

University of Southern California

"Might Laws Be Interpretive Works in Progress? Judicial Review and Respect for Precedent"

Thursday, April 3, 2025 1:10 – 2:40pm
Law School Room 1314

Draft, March 2025. For UCLA Workshop.
Please don't cite or quote without permission.

Mark Schroeder
University of Southern California
February 24, 2025

Might Laws Be Interpretive Works in Progress? Judicial Review and Respect for Precedent[1]

My goal in this paper is to consider a question in the ontology of law – what *kind of thing* are laws? According to a natural, familiar, and widespread view, laws are either *artifacts* or *social constructs*. They are things that are made. The idea that laws are things that are made looms large for some – in the hands of some positivists it is the indubitable observation that must shape our understanding of many questions about the shape and content of the law. For others it lurks in the background – a truth, but not one that teaches us anything important about questions of legal theory. In this paper I will be interested in exploring an alternative possibility: that laws are what I call *interpretive objects*. In particular, that they are what I call *interpretive works in progress*.

I am not quite going to argue, in this paper, that laws are interpretive works in progress. Doing so would require a careful consideration of the alternatives, but because the idea of interpretive objects is an unfamiliar one, I will be too preoccupied with explaining this idea in its own right to be able to adequately consider its alternatives. Moreover, a thorough argument that laws are interpretive works in progress would require conducting a complete survey of different sources of evidence – different ways in which the answer to this question could make a difference for different areas of legal theory, and a separate assessment of which of those answers are independently plausible. But I won't have the space here for this survey, either.

Instead, I will confine myself to calling attention to two topics where I find the applications of the idea that laws are interpretive works in progress particularly intriguing – to simple questions about statutory interpretation, judicial review, and horizontal precedential respect. Because I can't fully consider the alternative treatments of these topics made possible by alternative accounts of the ontology of laws, I can't claim to argue that the thesis that laws are interpretive works in progress offers the best explanation of them. And because I can't full consider the other upshots of the idea that laws are interpretive works in progress, I

[1] For the UCLA Legal Theory workshop, for pre-read only – not for wider distribution. Please excuse the breezy and drafty nature of this manuscript, which is an attempted summary of my thinking for purposes of this audience in particular, as well as the complete absence of scholarly engagement or references.

can't claim to argue that any costs of this thesis are worth it. But I hope that you will find the applications of this idea to be sufficiently intriguing, as I do, that it is worth it to consider them in their own right, and to see what we can learn from them, whether or not we end up accepting this idea.

1.1 Law and Interpretivity

The idea that law is in some way essentially interpretive is not in any way novel. Famously, Dworkin claimed that the concept of law is what he called an *interpretive concept*. Applying this concept requires working out an interpretation of *law*. Some philosophers of language and mind claim, in a similar vein, that the properties of *meaning* or *content* are interpretive – that identifying the meanings of words or the contents of thoughts requires interpreting a language or a person. This idea, too, could be applied directly to law – perhaps because content is interpretive more generally, identifying the content of law requires interpreting a legal practice.

The idea that I will be interested in here, in contrast, is the idea not that *law* is an interpretive concept, or that the legal *content* relation is an interpretive relation. Rather, it is the idea that laws themselves are interpretive *objects*. The theses that *law* or *content* are interpretive concepts or properties are theses in *higher-order* metaphysics. They concern the nature of properties or relations – properties of actions like *being legal*, properties of rules like *being in force*, relations between people and legal regimes like *falling under the jurisdiction of*, and relations between laws and actions like *prohibiting*. On each of these views, it is these properties or relations like these (among others) that are interpretive in nature. But the idea that I will be interested in is an idea in *first-order* metaphysics. It is an idea about what *laws* are.

These two ideas are not exclusive. Although proponents of the theses that concepts like *law* and relations like *prohibiting* are interpretive do not typically also say that laws are interpretive objects, as I understand and will shortly go on to explain that thesis, nothing prevents them from adopting either answer. And likewise, even those who deny that *law* is interpretive could either adopt or deny this thesis. I expect that as a matter of sociology, many who resist the idea that *law* is interpretive for positivist reasons will also be unfriendly to the idea that laws are interpretive objects, but we should wait and understand the consequences of this thesis before deciding which arguments for positivism work as arguments against it and which do not.

The idea that some objects are interpretive in nature – and not merely some properties or concepts – is an unfamiliar one. I have not yet discovered anyone who has clearly stated and developed it before me. And the difference between first-order and second-order metaphysics can make it seem unpalatable.

Compare Dworkin's own defense of his brand of metaethical realism from the charge that it is ontologically profligate. It is not, Dworkin points out, like he is postulating an extra realm of *entities* – as if there were an additional kind of particle, "morons", that make things moral or immoral. All that his realism says is that some things *are* moral and some are immoral. The contrast that he is drawing is between first-order and higher-order metaphysics. As a thesis in first-order metaphysics, i.e., the idea that there are extra things – the "morons" – metaethical realism would indeed be absurd. But it is not so absurd, Dworkin thinks, to think that some things are really moral and some immoral.

I admit to being somewhat pulled by the rhetoric of this contrast. Indeed, as someone whose first ever undergraduate reading in analytic philosophy was Quine's essay 'On What There Is', I have always been cautious about avoiding saying things about first-order metaphysics so long as I could. And in the philosophy of law, there are both enough plausible candidates for what laws might be, and also enough candidate properties to be interpretive, for it to be highly non-obvious, at best, that any questions about the law could ever push us to wade into these murky metaphysical waters.

But my own path to the idea that laws are interpretive objects began outside the philosophy of law. Some years ago I came to the conclusion that people – including you and I – are interpretive objects. This led me to take seriously the theory of interpretive objects as a tool that we need in our theoretical toolkit. And what I discovered – once I had this theoretical framework in my toolkit – was that in many different areas of philosophy interpretive objects are often confused for artifacts or social constructs. This happens, I believe, for words, concepts, and contextual parameters; it happens for literary works and other aesthetic objects, and it happens in the theory of persons. But the philosophy of law, of course, is home to one of the kinds of thing most famously and pervasively identified by philosophers as an artifact or social construct – laws. Could this also be a similar mistake? To evaluate that question, we will need first to understand more about what it means for something to be an interpretive object, and then to explore what kinds of consequences this idea might have, in legal theory.

1.2 Interpretive Objects

Objects come in different kinds. Some are biological objects, some functional in nature, some are artifacts, and others social constructs. There are also natural objects like rocks, molecules, protons, planets, and galaxies. Which kind an object belongs to tells us something important about where it is, how long it lasts, what its parts are, and what its modal profile is. Natural objects, for example, are unified by the fact that their parts adhere more with one another than with other things. This is true of all of rocks, molecules,

protons, planets, and galaxies. This principle tells us where to look for the parts of a rock, as well as for its boundaries – you look for which things adhere as much with each other as they do with other things. It also tells us why when you break a rock in half you no longer have one rock, but two.

Functional objects are different. The parts of a functional object do not need to be more stuck to each other than to other things. What make them parts of a single functional object is that they combine together to perform a single function. A pair of shoes, for example, are not more stuck together than to other things, and that is a good thing, because that would prevent them from together performing their function of protecting your feet while you walk. And biological objects are different yet again. Biological objects are unified by their role in systems of self-maintenance. A herd of wildebeests is unified not by together performing some function, for example, but by the fact that they interbreed and engage in group behaviors to ward off predators. And an ecosystem is unified by the cycles of codependence that allow its various denizens to persist and reproduce in similar patterns over time. Artifacts, by analogy, are unified by their acts of creation – and social constructs, in some way by the social practices in which they figure.

Once we appreciate the diversity of kinds of objects, and that each kind of object is unified in a different way, we have the key to make sense of how an object could be interpretive in nature. Interpretive objects are simply unified by their interpretability. In the northeast corner of Los Angeles there is a neighborhood called Eagle Rock. This neighborhood stretches below the San Rafael hills between the cities of Glendale and Pasadena, and it is named for a large prominent rock that juts out from these hills, rising above the 134 freeway. You will not be surprised to learn that this rock looks like an eagle. It is the "eagle rock" after which the neighborhood is named. Or to be more precise, it looks like there is an eagle within the rock – with wings stretched, it looks (so long as you squint hard enough) like it is about to fly out of the rock and soar out over the neighborhood of Eagle Rock.

I claim that there is an eagle in this rock – the eagle after which the neighborhood of Eagle Rock gets its name. Not a real eagle, of course – it's not alive. But an interpretive eagle. The interpretive eagle is not a natural object – it is within the rock, but it does not encompass the whole rock. The rock itself is not shaped like an eagle, but only a ledge that juts out of the rock. It's also not a biological object, for the obvious reason that it does involve any systems of self-maintenance. It's not a functional object. It's not an artifact – no one made it. And although there is a social practice of calling it the "eagle rock", it is clearly not a social construct, either – it existed even before this social practice, and for all I know the ancestors of the Tongva people may well have been delighted to discover it there in the rock thousands of years ago. It doesn't exist because of how it sticks together or maintains itself or what function it serves or how it was

created or any social practice that it is a part of. It exists because it looks enough like an eagle – because, in short, it is *sufficiently interpretable* as an eagle. It is an interpretive object.

Interpretive objects are unified by their interpretability. In order to identify the boundaries of the interpretive eagle, you have to interpret the rock – you have to identify what makes it look the most like an eagle. The bottom parts of the rock are not part of it, because eagles are not shaped like that. But the curved ledges sweeping out of the rock face are, because they are what look enough like wings to make it look like there is an eagle there. Similarly, in order to identify when the eagle came to be and when it ceases to be, you have to decide which answer to this question makes it the most like an eagle. Interpreting what makes it the most like a (real) eagle is also what tells us the shape of its modal profile – what possible changes it could survive and which would destroy it, for example.

There is nothing puzzling about there being objects whose nature is interpretive, I contend, because the space of interesting kinds of objects includes kinds of objects that are unified in different kinds of ways, and being unified by interpretability is just as natural and important a way of being unified as any of these other ways. So we should allow that the idea of an interpretive object makes sense. But more: the example of the eagle rock shows us that the concept of an interpretive object doesn't just make sense. It also has useful applications. It allows us to make sense of what the Tongva people (as I imagine) discovered, when they first noticed the eagle in the rock all of those years ago. They discovered something that was already there, and there was no natural object, biological object, functional object, artifact, or social construct that was already there. The thing that was already there was an interpretive object. And they discovered it by interpreting the rock.

This is not, I admit, a particularly interesting application for the theory of interpretive objects. I mean, it is interesting to me because I drive past it several times a day taking my children to swim and water polo practice. But in the space of things to be particularly interested in from a more objective point of view, I freely admit that it is marginal, at best. But what is special about the Eagle Rock is only that it is relatively easy to see that it *must* be an interpretive object, because it is relatively easy to see that it *can't* be any of our other familiar kinds of object. It takes a lot of coincidence for the rock to look so much like an eagle. Most things that look like eagles do so only through acts of artifice. But that doesn't mean that they are not interpretive objects – it just means that it is less obvious that they are interpretive objects, because they are (approximately) co-located with artifacts.

And if that is right, then we may sometimes fail to identify that something is an interpretive object because we confuse it with an artifact with which it is co-located or closely associated. This point is important; if laws are interpretive objects, they are interpretive objects that are often co-located or closely

associated with artifacts. In part two I will show how we can distinguish between an interpretive object and an artifact, even when they closely overlap. That will provide us with an important tool for thinking through the consequences of the idea that laws are interpretive objects. But first, there is one more important distinction that I need to make – and a closely related observation about what is special about interpretive objects.

I.3 Inherence and Incidentality

In the last section I implied that some interpretive objects might be difficult to distinguish from artifacts because they are located in the same place as those artifacts. One of the traditional examples used to illustrate the possibility that two objects might be located in the same place is the example of the statue and the clay. A lump of clay is formed into a statue. The statue is made wholly out of the clay, and so they share all of the same parts. The clay doesn't disappear when the status is formed – once the statue is squashed it is still there. But they are located in exactly the same place for as long as the statue persists.

In this case we can tell the difference because the clay lasts for longer than the statue – it comes into existence before the statue and persists after the statue is gone. But even if the statue and lump are created at the same time – for example by joining partial statues made out of different lumps – and destroyed at the same time – for example by breaking the statue – the statue and the clay can be distinguished by their different *modal profiles* – differences in what possible changes the statue *could have* survived, from which the lump of clay could have survived.

But there is yet another way in which we can distinguish the statue from the clay. The statue (following a famous example from Gibbard let us call it 'Goliath') is bold, original, baroque, grey, and heavy. The lump of clay (again following Gibbard, let us dub it 'Lumpl') is also bold, original, baroque, grey, and heavy. But clearly there is a difference in how Goliath is baroque and how Lumpl is. Goliath is baroque in its own right, because it is a statue. But Lumpl counts as baroque only because it is formed into a statue that is. Whereas Lumpl is grey and heavy in its own right, and Goliath gets to be these things only because it is made out of Lumpl.

We need some terminology for this difference. I will say that Goliath is *inherently* bold, original, and baroque, but grey and heavy only *incidentally*, and that Lumpl is grey and heavy inherently, but bold, original, and baroque only incidentally. Philosophers will interpret this talk of inherence and incidentality differently. Some will say that it is a distinction between two ways of having a property. (I will talk this way.) Others will say that strictly speaking, only Goliath is original and Lumpl is not, but that we can speak loosely, with

a derivative meaning for 'original' in order to say loosely that Lumpl is. (I don't rule this out, but it requires too many words to say.) And others may say instead that Goliath is bold-qua-statue but heavy-qua-lump. (This is not my preferred view, but may blessings go with you if you try it on for size.) I will not be concerned with the differences between these theoretical accounts of this distinction, but rather with what they have in common – the distinction that I have marked between Lumpl and Goliath, of which each of these more detailed accounts are theories.

Whatever the nature of the distinction between inherence and incidentality, it is clear that the difference between which properties Goliath has inherently and which Lumpl does must have something to do with the kinds of objects that they are. Goliath gets to be baroque and original inherently because it is a *statue*, and Lumpl gets to be grey and heavy inherently because it is a natural object – a lump of clay.[2]

So just as objects of different kinds are extended in space and time according to different principles, have their parts determined in different ways, and have different kinds of modal profiles, similarly which of its properties an object has inherently and which only incidentally must depend on the kind of object that it is. Every object is located – extended – in some way in space, and lasts for some period – its "extension" in time. Every object is related to some other objects as its parts – its "extension" in what we might think of as "mereological space". And every object has a modal profile – a range of possible circumstances in which it would have existed – which again we can think of, at least metaphorically, as its "extension" in "modal space". Similarly, every object has some of its properties inherently, and some incidentally – an "extension" in what we might call "quality space". Together, I call these the *five spaces* – space, time, mereology, modality, and quality. Every object is distinguished by its extension in the five spaces, and different kinds of objects are distinguished by the manner in which they are extended in the five spaces.

Interpretive objects, I claimed earlier, are extended in space, time, mereology, and modality in whatever way makes them most interpretable as what they are interpretable as. The answer to where the interpretive eagle is, is given by whatever makes it the most like an eagle, and similarly for the answers to when it first came to be or when it will cease to be, what its parts are, and under what conditions it would have been. Similarly, I suggest, the boundaries of an interpretive object in quality space come from whatever makes it the most interpretable as what it is interpretable as. The eagle has whichever properties incidentally make it the most like an eagle.

---

[2] My characterization of natural objects in section I.2 actually closely follows Gibbard's own explanation of the continuity conditions for lumps of clay like Lumpl in Gibbard [1975].

The rock of the Eagle Rock, for example, is a little bit blotchy in color. But although there are many species of eagle with different kinds of coloring, none of them are blotchy in quite the way that the rock is. So being blotchy does not make the interpretive eagle more like an eagle. It is not, therefore, inherently blotchy. It is not a blotchy eagle, but just an eagle in blotchy colored rock. If, on the other hand, the part of the rock that looks like the eagle's head had been markedly lighter than the part corresponding to its wings, then that would have made it look more like an eagle, by making it look like a bald eagle, in particular. In that case, the interpretive eagle would have been, I suggest, inherently bi-colored, and not just an eagle that happened to be in bi-colored rock. Similarly, if there is a striated rock that looks like a cat, then if the striations run in the right direction they will be part of what makes it look not just like any old cat, but like a tiger. In that case there would be an interpretive tiger that was inherently stripey, and not just an interpretive cat in stripey rock.

These examples illustrate the point that the boundaries of interpretive objects in quality space are set, as we expected, interpretively. They go in whichever way makes the interpretive object most interpretable as the kind of thing as which it is interpretable – just as its boundaries in space, time, mereology, and modality.

But they also illustrate something really important about what is special about interpretive objects and their boundaries in quality space. Interpretive objects are always, in a sense, imperfectly embodied. The interpretive eagle is not perfectly like an eagle – it is only as much like an eagle as its embodiment in rock allows it to be. Because rocks are not very much like eagles, this form of embodiment substantially limits how much like an eagle the interpretive eagle can be. Seeing the interpretive eagle for what it is therefore requires acknowledging the limits of its embodiment. We have to ignore its blotchiness, for example, and the fact that it doesn't move.

But an interpretive object is always *as much like* what it is interpretable as, as its imperfect embodiment allows. If the rock is bi-colored in just the right way, then the interpretive eagle automatically gets to be inherently colored – that is all that it takes for there to be an interpretive bald eagle in the rock. This, in a way, is the ultimate hallmark of interpretive objects. Because their boundaries in space, time, mereology, modality, and quality are set in *whatever* way makes them most interpretable as what they are interpretable as, we are bound to find them wherever things are imperfectly embodied. And their shape is always going to reflect this imperfect embodiment. But they are always going to be as much of what they are interpretable as, as their embodiment allows.

Many things, I think, are imperfectly embodied in this way. We are, for one. We are embodied in imperfect ways in bodies that are subject to stress, hormones, and mood swings. Relating successfully to one another requires understanding the limits that come from our embodiment, and working around them.

Works of literature – including works of philosophy – are also imperfectly embodied. They are created at particular times and places, by imperfect people, through imperfect processes. They contain typos and are infected by background ideas that come from the time and place of their authors. Relating successfully to these works requires understanding the limits that come from this embodiment, and knowing how to navigate around them. Laws are also created by imperfect people, in imperfect ways. Could they be imperfectly embodied, in a similar way? So says the idea that laws are interpretive objects.

2.1 Literature

So far I have explained what interpretive objects are. I have shown that there is nothing extravagant about postulating them as a new kind of object, by situating the way in which they are unified and extended in the five spaces in comparison to the ways in which other familiar kinds of objects are so unified and extended. I have given one clear example of an interpretive object, in the case of the interpretive eagle in the eagle rock, which we know must be an interpretive object because it cannot be any other more familiar kind of object. And I have introduced the distinction between inherence and incidentality, in order to highlight the fact that interpretive objects are distinguished by their imperfect embodiment – and by the fact that they are always as much of what they are interpretable as, as the limits of their embodiment allow. What remains to be seen, is whether we can fruitfully identify any examples of interpretive objects that are more interesting for people who do not share my regular drive back and forth between Glendale and Pasadena.

Identifying more examples of interpretive objects is less straightforward, because in more interesting cases, it is much harder than in the case of the eagle rock to distinguish them from artifacts and from social constructs. Nevertheless, I believe, careful examination reveals that we can use the fact that the boundaries of interpretive objects are set by interpretability, in order to get leverage on whether there are also interpretive objects in more interesting kinds of cases in which they are co-located with other kinds of objects, and also on whether the objects that are and should be most interesting to us are the interpretive objects or those other objects. In this section, I'm going to illustrate this leverage by thinking through the case of works of literature.

By works of literature I mean to include at the very least novels, poems, articles, series, and philosophy monographs; I leave open exactly what might count as a work of literature as a question that might fruitfully be informed by our reflection on core cases such as these. Works of literature exist because of acts of creation. Someone authors them. This process creates artifacts – particular copies of a book in print or electronically, on hard drives. It also exists as part of a social practice. Works of literature belong to particular genres and

are framed by expectations about how they compare to other works and what our interest is in producing and reading them. And they perform interesting functions – works of philosophy exist to advance arguments or ideas, for example, and pulp fantasy novels help us to ignore our problems. So wherever there is a work of literature, there are artifacts, social constructs, and functional objects.

Any of these might naturally be called the "work". By "Gulliver's Travels" you might mean the particular paperback sitting on your bedside table, or anything that serves the function of satisfying this week's reading assignment for English Lit, or the ordered sequence of characters that Albert Rivero sent to his copyeditor for the *Norton Critical Edition*. But among all of these eligible things to be called "Gulliver's Travels", I suggest, is an interpretive object – something whose boundaries in each of the five spaces is set only by whatever makes it most interpretable as a satirical novel. And it is the interpretive object in which we are interested, when we want to engage with a work of literature.

There are many things that you can do to *Gulliver's Travels*. You can carry it, transcribe it, or use it as a doorstop. You can also illustrate it, translate it, read it, or produce a critical edition of it to compete with Rivero's. These things are not all on a par. Some of these things that you can do with *Gulliver's Travels* count as forms of *engagement* with the novel, but some do not. Illustration, translation, reading, and producing a critical edition of are all forms of engagement. Carrying, transcribing, and using as a doorstop do not – these can all be done by people who are completely illiterate. Forms of engagement are all ways that we can relate *inherently* to works of literature. These other ways of relating to works are ways that we relate to the work only incidentally – because they are ways of relating to its copies, or to the letters that make it up.

The main reason to think that works of literature are interpretive objects is that all of these forms of engagement – the ways that we can relate to works inherently, and not just incidentally – require interpretation. In order to illustrate or translate a satirical novel like *Gulliver's Travels* you have to make interpretive choices – you have to decide what is important and what is not, and how those things fit together. And as I'll now show, these interpretive choices affect how you understand the boundaries of the work in space, time, mereology, modality, and quality.

Observe, for example, that not all copies of a work have exactly the same text. In the first sixty or so copies of *Harry Potter and the Philosopher's Stone*, for example, Harry's letter from Hogwarts listed "one wand" twice, at both the beginning and end of his shopping list. If your copy of the book contains this typo, then you are rich – but you also have a genuine copy of *Harry Potter and the Philosopher's Stone*. Indeed, you have a copy of this same book even if yours was printed in North America, where it is called *Harry Potter and the Sorcerer's Stone*, a completely different title. But arguably, if you have a copy of the *Adulterer's Bible*, whose ten commandments enjoin that "Thou shalt commit adultery," then you don't have a copy of the Bible at all.

How could some typos or variations prevent an edition from being a copy of a book, but others not? The question of which do and which do not, I contend, must be an interpretive one. Answering it requires identifying whether counting this as a copy of the book helps, or interferes with, it being the most like what it is interpretable as. And arguably, whether there is a typo on Harry's shopping list does little to interfere with the kind of book that *Harry Potter and the Philosopher's Stone* is – indeed, it might even be argued to be foreshadowing of the climactic events of *Harry Potter and the Deathly Hallows* six volumes later. But plausibly, enjoining adultery does interfere with the kind of book that the Bible is, at least according to many. Since we can't evaluate what counts as a copy of a work or not without interpreting, the work itself can't have any privileged text – the boundaries of what is or could be part of its text is interpretive.

The same thing goes for the modal profile of a work. Some works exhibit less variation between editions. Perhaps they had a lazy copyeditor who declined to make any changes and went out of print after one edition. But they *could* have varied. There could have been a typo at press, or the author could even have made different choices. Tolkien, for example, might have had Galadriel give Gimli just one strand of her hair, rather than three. Had he done so, he would have failed to evoke a comparison to the Silmarils and *The Lord of the Rings* would have been slightly less rich. But it would still, for all of that, been *The Lord of the Rings*. Or he could have had Frodo slip the Ring off of his finger before Gollum bit it off.

That would have been another small change in the text, but it would arguably have made it a completely different book – as different as if he had made it all about Tom Bombadil, an idea that he in fact flirted with for some time in early drafts before settling on the story about the rings of power. This other book would have satisfied his publishing contract for a sequel to *The Hobbit*, but it would not have been *The Lord of the Rings*. Which possible (but counterfactual) books could have been *The Lord of the Rings*? Only interpretation can say. Our answer to what variations *The Lord of the Rings* could have admitted of tells us what is essential to it. And what is essential to it is a matter of interpretation – it is whatever makes it most interpretable as a high fantasy trilogy.

Works of literature are also extended in quality space. Some of their features are inherent, and some incidental. Paradigmatically, it is natural to think that typos are incidental features of works – features that they really do have, but only because of the imperfect features of their creation and embodiment. To engage with a work, you have to be able to identfy that something is a typo, and overlook it. For example, when reading that last sentence, you had to correctly see that it was about identification, which makes sense in the context of the sentence, and not about some technical philosopher's concept of identfication that you don't recall me introducing.

One of the leads in Shakespeare's *Cymbeline*, for example, is a strong female character from whom we get the modern name 'Imogen'. But it turns out we think so only because when the First Folio was typeset, the two n's in 'Innogen' were set too closely together. When contemporary directors stage *Cymbeline*, they have to decide whether the character is named 'Imogen' or 'Innogen', and it turns out that they split – it's an interpretive choice, not one that is settled by learning facts about the history of the text. On either view, the text is imperfectly embodied. It involves one spelling of the character's name only incidentally – because of how it has been transcribed in the First Folio.

Mark Danielewski's novel *House of Leaves* is famously scattered with apparent typos, some of which look like author or copyediting errors, but others of which, like "he stayed up searching all night until exhaustion caught up with me," could be the *narrator's* slip of the tongue giving a clue to their identity. Which are which? I say that it's an interpretive matter. It might be relevant to study Danielewski's notes, to answer this question, but it's not dispositive. The book does contain the word "me", but this feature of the text could be inherent, or it could be incidental.

Other features of a work can be incidental as well. If you read a philosophy article by someone who received their PhD at USC, for example, you may see that they use the word 'reason' a lot. This could be important to their thesis and argument. Or it could be just how people who received their PhDs at USC talk, and can be safely ignored when re-expositing their argument. Whether the word 'reason' is an inherent or an incidental feature of their paper is an interpretive question – it depends whether this contributes to, or gets in the way of, seeing their paper as making an important contribution – or something that most closely approximates an important contribution. When people disagree about what we can learn from Aristotle that is independent of his racism and sexism, they are likewise disagreeing about whether his racism and sexism are inherent features of his works, or instead just symptoms of his historical and social situatedness. Which of these is the case is not a purely textual question, and it's not a purely historical one. It's an interpretive question, about which reading makes his work most closely resemble a great contribution to philosophy.

The reason that all of these questions are interpretive, I contend, is that it is very, very difficult to produce a perfect work of literature. We misspell words by accident, overlook typos, and import irrelevant ideas make it into our texts because we imbibed them from our PhD supervisors or from the unquestioned background sexism of our age and social circle. Given all of the imperfection of the world in which works of fiction and philosophy are created, it is no wonder that the works that we create will be imperfectly embodied. Like the eagle, they will be in some respects embodied in ways that get in the way of seeing them for what they are. A good interpretation of them reveals what those are, and that those features are merely incidental, and not inherent. Consequently we can ignore those features when engaging with the work.

2.2 Interpretable as What?

So the evidence that works of literature are interpretive objects comes from multiple directions. By distinguishing between ways that we relate to works inherently and ways that we relate to them only incidentally, we can see that the former require the exercise of interpretation, while the latter do not. That is evidence that even if we can use "work" to refer to different kinds of thing, the interesting objects with which we engage inherently and not merely incidentally are interpretive. We have also seen that when we interrogate the boundaries of the work in space, time, mereology, modality, and quality, there are interesting questions about what count as copies of a work, what changes it could survive, when it first came into existence, and which of its features are inherent and which incidental that can only be settled interpretively. And finally, because interpretive objects are always only imperfectly what they are, the fact that our processes of creating works of literature are always only imperfect is evidence that it should not be surprising to find interpretive objects there.

The eagle is merely interpretive because it is not perfectly an eagle. But it is inherently as much like an eagle as its imperfect embodiment in the rock will allow. So what is it that works of literature are interpretable *as*? The answer is different, I think, for different genres of work. In the case of philosophy articles, for example, I think that it is as *making a contribution*. The perfect contribution would be many things: original, true, well-reasoned, well-explained, insightful, relevant, creative, recognized and taken up by other people, and more. It is very hard to be all of these things. Nothing is perfect in all of these ways. But we can find contributions in philosophy books and articles by seeing them as imperfectly embodied contributions. We use *charity* in order to identify what to read into a text as important for the contribution that it makes, and what to read out of it, as a feature of its imperfect embodiment that gets in the way of seeing its true contribution.

In the case of satirical novels, in contrast, some of the key elements are humor and criticism. The perfect satirical novel must be truly funny, and it must truly bite. It has to reach beyond itself to identify plausible critiques of people or social systems in the actual world. Real satirical novels do this only imperfectly – their authors' senses of humor are limited, for example, or only imperfectly perceptive about the full dimensions of social ills or oblivious to other forms of injustice, or overly preoccupied with matters of the author's own personal experience. But the greatest satirical novels rise above this. They begin as critiques of some particular social system, for example, but then transcend that because they work just as well or even better as critiques of other social systems as well.

The perfect idea of a law also involves multiple parts. Law aspires to systematicity – to uniformity in application across different circumstances. If different cases are to be treated differently, then the perfect law itself tells us how, and when. This gives law by its nature a kind of necessity. But law also has by nature an important kind of contingency. It solves coordination problems and makes decisions about matters that could have gone in different ways – and do so at different times and in different places. It must be contingent, because it embodies *our* practice about how to live together.

Like the ideal of a perfect philosophical contribution and the ideal of the satirical novel, these features of law are difficult to fully satisfy together. The contingency of law means that the law changes over time. But changes in the law over time conflict with the aspiration of the law to systematicity – to treating like cases alike. We can regain systematicity only if we understand changes in the law over time as themselves law-governed. This is why Hart ultimately postulates a "rule of recognition" – an ultimate rule that tells us how changes in the law themselves count as rule-governed, and hence which reconciles the contingency and evolution of law with its aspirations to systematicity. But in order to do this, the rule of recognition must itself be constant over time. We have to be stuck with whatever rule of recognition our ancestors worked with, or else the shift in this rule will itself fail to be rule-governed, and hence will threaten the systematicity of law. But the idea that being ruled by laws requires working with the same rule of recognition our ancestors selected is in tension with the aspiration of law to be the embodiment of our practice about how to live with one another. It cannot embody *our* practice, if it gives all authority to our ancestors' plan.

So these aspirations of the law, I think, are in tension with one another. Nothing can perfectly satisfy them. But some things might get close. Those things are laws. And the question to what makes them laws is settled by what makes them better at satisfying the ideals of systematicity, adaptability, and shared practice. Or so we can say if we accept the idea that laws are interpretive objects, and supplement it with the thesis that what laws are interpretable as, is exemplifications of a systematic but adaptable shared practice.

2.3 Interpretive Works in Progress

So far I have been explaining what interpretive objects are. The idea that I want to introduce and explore in this paper is that laws are interpretive objects, and I have just hinted at enough to see what this thesis could amount to. It amounts to the idea that laws, though often imperfectly embodied, are as much like exemplifications of a systematic but adaptable shared practice as their imperfect embodiment will allow. We'll see in part three that this idea offers us a way of thinking about several interesting questions about

statutory interpretation and judicial review. But before getting to those applications, there is one more important piece of the theory of interpretive objects that we need to introduce.

At the beginning of this paper I didn't just say that I would explore the idea that laws are interpretive objects. I said that I wanted to explore the idea that laws are a particular *kind* of interpretive objects – interpretive *works in progress*. An interpretive work in progress is just an interpretive object whose interpretive "text" extends into the future. It turns out that in interesting cases, when the text of an interpretive object extends into the future, we ourselves play a role in shaping that future by virtue of our interpretations of the object.

Umbhali, for example, is writing a serial novel. She has published seven chapters so far, which have been well-received, and is hard at work on chapter eight in order to meet her publisher's deadline. Because the first seven chapters are already in circulation, she can no longer make changes in them. But the rest of the novel still lies in front of her, and she could take it in any direction. Luckily, she has a plan. She wrote the first seven chapters with an interpretive vision about where it was going. But while procrastinating she falls down a rabbit hole of fan TikToks developing various theories about where the novel is going. She might even discover that one of her fans has an even better plan for the book than she did – and nothing prevents her from adopting that plan going forward.

In order to do so, Umbhali has to change her mind about how to understand what she has written so far. Some events from chapter three that she had originally intended merely as misdirection, for example, or only as character development, might now come to be central to the plot, and others which she had conceived of as central may on the new plan only create dramatic tension. She may even decide that a copyediting error in chapter four can really be turned into an important clue about the plot resolution.

These changes engender changes in the modal profile of her novel – in what changes it could have endured without becoming another book entirely. And they change what is inherent and what is incidental – turning a typo into an inherent feature of the text. Because she is the author, Umbhali has a lot of power over which of these interpretations is ultimately correct. But she gets this power not by having the intrinsic authority to settle interpretive questions about her own novel, but rather simply by the fact that she is the one with the power of the pen – by choosing some words over others to go into chapters eight and following, she is steering the interpretive text of the novel to favor one interpretation over others.

In doing so, Umbhali doesn't have to respect her own prior interpretation. And she doesn't even have to respect the most plausible interpretation of chapters one through seven. She just has to respect the interpretation that she reasonably hopes to *make* most plausible, through her own current choice about what

to write in chapter eight, and her expected later choices about how to follow up in chapters nine and following.

If Umbhali were to stick to her original plan even though she can now see a better way forward, she would be abnegating her responsibility as author to craft the best novel that she can. To do so would be a conscious decision to create an inferior work. Authors do sometimes decide to create inferior works. There are tradeoffs of time and energy against other goals in life, and we all must triage our efforts. But telling herself that this was her original plan, once she sees a better way that it could go is just rationalizing.

But Umbhali can't just write anything going forward. It has to make sense in the context of what she has already written. If she leaves off her characters in a gritty, historically realistic fifteenth-century Istanbul and starts over in chapter eight with a goth-punk setting in an alternative postapocalyptic Tokyo in which some people are born with magical powers, then it is going to be really hard to avoid the interpretive conclusion that she has simply begun a new book entirely, even if her publisher lets her get away with publishing them under the same title. She is, as Heidegger might say, *thrown* into the present by her past. She inherits the first seven chapters of the book as material that she must work with. She just doesn't have to stick to any prior interpretation of those chapters. It's up to her whether to start chapter eight with "and then she woke up".

You might think that what Umbhali ought really to do, in writing chapter eight, is to imagine all of the possible ways that the rest of the book could go, and to choose the very best of these, and then to write the version of chapter eight that fits into this very best possible future version of the book as a whole. But that is not quite right, either. Just as present-Umbhali is not bound by past-Umbhali's decisions about how to interpret her novel, future-Umbhali cannot be bound by present-Umbhali, either. Because present-Umbhali doesn't have the power to choose now how all of the future chapters will go, she is subject to the fortune of how her future self follows through on carrying out those chapters.

But maybe Umbhali is afraid of changing her mind again. If her new vision for the book is bold enough, she may fear that when it comes time to write chapter nine, she will get cold feet and chicken out. If chapter nine starts with "and then she woke up", all of her efforts to set up a better version of the novel could get foiled, and she could end up with a worse novel overall, for her inconstant efforts alternating between different directions in which she was trying to take it. Alternatively, Umbhali's new vision for her book could be so compelling and easy to carry out that she has no doubts that she will end up following through. It's important, in deciding how to craft chapter eight, how she can expect the future to go – and how she can shape that future by what she writes now, in the same way that her decisions about what to write

now are shaped by her decisions last month about how to write the chapter seven with which she has now been thrown.

If laws are interpretive works, then whether they are interpretive works in progress is going to depend on what their interpretive text is. If the interpretive text of a law is just the word of a statute, then arguably, laws could turn out to be interpretive objects, but not interpretive works in progress in any interesting sense. For statutes, like most novels that are not published serially, are "in progress" only before their texts are fully formed and they become enacted. But if the interpretive "text" of a law might include a broader practice of applying and following that law, then this is a practice that itself has a shape over time – and whose shape could go in different ways, in the time that has yet to pass. In that case, if laws are interpretive objects at all, then they are interpretive works in progress.

3.1 Statutory Interpretation

Everything, as I've argued, has some properties inherently, but others only incidentally. The inherent properties of an object are the properties that it has in its own right, and so identifying and responding to these properties is a way of identifying and responding to *it*. But incidental properties are features that an object counts as having only because of the way in which it is associated with something else that has these properties. Identifying and responding to them is not a way of identifying and responding to something in its own right, but only of responding to the other thing with which it is associated.

The distinction between inherent and incidental properties, I believe, is particularly helpful in thinking through an important aspect of what judges do in completing the task of statutory interpretation. For it allows us to recognize that laws genuinely have some features, even though they do so only incidentally, and not inherently.

Of course, many of the commonly acknowledged principles of statutory interpretation presume by default that all of the language of the statute is inherent. When interpreting a philosophy paper, it is sometimes important to acknowledge the possibility that an author uses the words 'desire' and 'want' interchangeably, and so should be read as talking about a single thing, even if we the reader would use those words differently. This is a way in which philosophical works can be imperfectly embodied – they can be embodied through their author's sloppiness about how to use the words 'desire' and 'want'. But lawyers often apply the principle of *Against Surplusage* to statutory interpretation. This principle says that if a statute uses two different words, then they should be interpreted as having different meanings.

If the principle of *Against Surplusage* were universally applied, that would mean that proper statutory interpretation never allows for acknowledging that distinctions between words in a statutory text are merely incidental. But not only does this principle have recognized exceptions in practice, but it *cannot* be universally applied. For it easily runs into conflict with other canonical principles of statutory interpretation. For example, the principle of Ordinary Meaning says that unless defined within a statute, terms must be given their "ordinary meaning", the principle of Consistent Use requires that terms be interpreted as used consistently across different statutes regarding the same subject, and the principle of Noscitur a Sociis says that a word is known by its associates and so must be interpreted consistently with how its closely associated words are defined. And each of these principles, just to take a few examples, can conflict, all by itself, with the principle of Against Surplusage, in the right interpretive context. Ordinary Meaning can create the conflict if a statute uses two undefined terms that share an ordinary meaning. Consistent Use can create it if a single statute borrows terms that are each clearly defined in other laws on similar topics to mean the same thing. And Noscitur a Sociis can create it if each of the terms have associates whose clear definitions would force the same meaning on this pair. And conflicts of principles of statutory interpretation can also arise from sets of three or more principles, in the right kinds of contexts.

So there can be no avoiding the conclusion that each of these rules can only be a presumption. We can *presume* by *default* that all features of the statutory language are inherent. But once we end up in a situation in which our principles of statutory interpretation lead us in different directions, we have to use judgment in how to apply them. And sometimes we will find that the right judgment requires overruling or identifying exceptions to these generally helpful interpretive principles. If we overrule the principle of Against Surplusage, then we are deciding that the property of being formulated by statutory language that uses two distinct words is not, after all, an inherent feature of this law. Seeing this about the law does not help us to identify it as the law that it is. Rather, it is a true feature of the statute that enacts the law, but merely an incidental feature of the law.

Similar examples can be constructed for other principles of statutory interpretation to which we might identify that we have found an exception. For example, we might overrule the principle of Ordinary Meaning for a statute that does not define its terms but uses terms that are commonly defined in closely related statutes, by appeal to the principle of Consistent Use. In that case the plain meaning of the statute is being construed as a merely incidental feature of the law, rather than an inherent one. It may still have that plain meaning – and it may be a defect in the law that it was formulated in this way that required judgment about whether to prioritize Consistent Use over Ordinary Meaning. But the plain meaning is not an inherent

feature of the law itself. Identifying the plain meaning is not, we decide in this case, helping us to identify the law.

So the distinction between inherent and incidental properties is helpful in describing what we are doing, in applying principles of statutory interpretation. These principles are principles about what is inherent to law, and when they conflict, navigating that conflict helps us to identify features of a law that may instead be merely incidental. The fact that these features are merely incidental, in turn, explains why it is that they can henceforth be ignored when interpreting and applying that law. Ignoring these features of the statutory language and how it is connected to other things like common understanding, legislative intent, and consistency in other statutory language, is completely consistent with utter and total respect for the law itself, because what we have concluded is that these features are not features that the law has in its own right, but only features that the law counts as having because of how it is associated with the statutory language that enacts it. The distinction between inherent and incidental properties also offers a striking unifying principle about the canons of statutory interpretation. They are all, in a way, articulations of different particular applications of the principle that features of statutory language are to be interpreted as inherent to the greatest extent that this is possible.

None of these points depends on the idea that laws themselves are interpretive objects. They are simply straightforward applications of the distinction between inherent and incidental properties – a distinction that, as I have argued, can be applied to *every* kind of object. Every object, no matter its kind, has some properties inherently and some incidentally. But I did argue, earlier, that *which* properties of an object turn out to be inherent and which incidental depends a lot on the kind of thing that it is. The statue has different properties inherently than the clay, and the reason why this is so is that the statue and the clay are different *kinds* of thing. One is an artifact, and one is a natural object. Similarly, if we grant that the task of statutory interpretation consists, *inter alia*, in discriminating between the inherent and merely incidental features of a law, then whether laws are interpretive objects or artifacts will have important consequences for which kinds of properties we should expect to be inherent, and which merely incidental.

In the case of interpretive objects, as I said, the inherent/incidental distinction tracks the difference between which features contribute toward making something interpretable, and which get in the way. I have called the features that get in the way of interpreting something its "imperfect embodiment". On this account of the nature of laws, the principles of statutory interpretation are not arbitrary, nor is the matter of how to adjudicate their conflicts. Rather, they identify features that would make statutory laws more interpretable *as laws* – more law-like. We apply them all by default – treating everything by default as if it is significant or meaningful – because it is an imperfection in law *qua law* for these principles of statutory interpretation

not to apply. Failures in these principles, broadly speaking, make for failures in the kinds of systematicity, adaptability, and shared practice that are at the heart of law itself. That is why we always respect these principles by default.

But they are also *only* defaults. For laws are created by an imperfect process, and sometimes the choices made by the legislators – even the purposive ones – can get in the way of seeing the product of their efforts as law. The conflicts between the principles of statutory interpretation are contingent, rather than principled – a more perfect legislative process, carried out by a more omniscient legislator with infinite temporal resources – would not require adjudicating conflicts of these kinds. So the view that laws are interpretive objects explains this very well. It says that it is no coincidence that we get into trouble just because the legislative body is imperfect. For it says that identifying what is inherent and what is incidental is precisely a matter of what features of imperfect embodiment get in the way of seeing the product of statutory enactment as being the most law-like.

The view that laws are artifacts must give a different, and I believe less satisfactory, answer to the question of which features of the law are which, and why. This can at first be hard to see, because many of the principles of statutory interpretation look like they are really epistemic principles about how to best discern legislative intent. And it is natural to think that whereas for interpretive objects, the inherent/incidental distinction should line up with what contributes to interpretability, for artifacts the inherent/incidental distinction should line up with the intentions of the artificer – which in the case of statutory law, would be the legislator. So there is a kind of harmony between the idea that laws are artifacts and the idea that statutory interpretation is about discerning intent.

But it isn't, I think, exactly right that the principles of statutory interpretation are good principles for discerning legislative intent. The principles of statutory interpretation tend to be decent principles about how to discern the intentions of an ideal legislator who never makes mistakes, has all relevant knowledge including perfect knowledge of every other statute and how that statute has been interpreted, and unlimited time in which to assemble this information in order to concoct ideal legislative language. But of course we know that these assumptions are not true of the actual legislators who enacted any of the laws that we seek to interpret. If statutory interpretation were all about legislative intent, then it would not make sense to apply principles like Against Surplusage and Ordinary Meaning widely until they come into conflict with other principles – on the contrary, sometimes it would make sense to appeal to what we know about the imperfections of the legislative body in order to disregard some of these interpretive principles even when they could be cleanly and uncontroversially applied with no conflicts whatsoever.

So although I don't think by any means that this cleanly adjudicates the question of whether laws are artifacts or interpretive objects, I do think that it casts the role of imperfections of the legislature in an illuminating light, and offers the potential to think more systematically about the explanatory unity in common principles of statutory interpretation. It is potentially useful work for the idea that laws are interpretative objects.

3.2 Judicial Review

If the inherent/incidental distinction can be applied to features of statutory laws, we might also wonder whether it can be applied to general features of legal systems, such as whether one or another law is *in force*. Indeed I think that it can, and the application to this issue is particularly fruitful.

In a standard application of judicial review, a court is determining whether an existing law was properly enacted. This law, such as the District of Columbia's Firearms Control Regulation Act of 1975, may have been on the books and followed for some time. In the case in which this act was overturned by the US Supreme Court in Heller vs. D.C. in 2008, this law had been in effect for a third of a century. But not all laws that are on the books and commonly accepted and followed fall within the legislative authority of the original body that enacted them. According to the 5-4 majority in Heller, for example, the District of Columbia has and had no authority to enact a blanket ban on the possession of handguns – and indeed, no legislative body in the United States has the power to do so – because this is forbidden by the proper reading of the second amendment.

There is a kind of puzzle about how to understand what happens when a court strikes down a statute as in Heller. It can be generalized into a more general puzzle about what happens when a court strikes down *part* of a statute, or when it overrules and revises a governing and longstanding interpretation of that statute (as Heller itself does, with the language of the second amendment). But I want to focus for purposes of discussion on what seems to me to be the purest and simplest case – the case of laws being stricken down in toto.

Here is one naïve way of describing what has happened when the court strikes down the Firearms Control Regulation Act. Legislative bodies are empowered to make laws. But they are not empowered to make just any laws. Their legislative powers are restricted by the scope of their proper legislative authority – as laid out, perhaps, in something like a Hartian rule of recognition. It is possible for legislatures to abuse this power by attempting to enact legislation that goes beyond their proper legislative authority. And it is possible for legislators – or even for those whom they govern – to mistake the exact bounds of the legislature's

authority, and to think that things fall within its scope when really they do not. But just because a legislative body votes on and "passes" statutory language, that does not by itself make it law. To be law, the statute has to fall within the scope of the legislative body's (legal) *authority* to shape the law.

On this way of thinking, the judgment of the 5-4 court majority was that the District of Columbia had exceeded its constitutional legislative authority, in attempting to ban the possession of handguns in the district. So what they held themselves to be discerning, on this interpretation, was that the Firearms Control Regulation Act was never actually law. It was never *in force*. It was just widely believed to be in force, and the court was merely correcting this misimpression and taking a step toward mitigating its harmful or unjust consequences on residents of DC who might have liked to have kept a handgun. Call this picture the *Pure Corrective* conception of judicial review.

One of the most attractive features of the Pure Corrective picture is that it captures – indeed, it embodies – the idea that courts are not legislatures. Legislatures have the power to enact or change laws. So if we take the idea that courts are not legislatures to its simplest, purest, interpretation, then it requires that courts do not have the power to enact or change laws. Revoking a law that has been in place for a third of a century would clearly be a significant change in the law. And so it seems like to do full faith to the idea that courts are not legislatures, we need to deny that the court changes the law in *Heller*. That is what the Pure Corrective picture does. And it tells us what the court is doing instead. It is discerning what the law already is, and disseminating epistemically authoritative advice about that. It is correcting the widespread misimpression that the Firearms Control Regulation Act of 1975 was ever the law. The fact that it embodies so perfectly the idea that courts are not legislatures is a powerful piece of support for the Pure Corrective picture of judicial review.

But on the other hand, the Pure Corrective picture paints a strange picture of the period between 1975 and 2008 in the District of Columbia – indeed, of all of the U.S. in the entire pre-2008 period in which no one suspected that personal handguns were relevant to the necessity of a well-regulated militia for the security of a free state. The facts about what the law was during that period, according to this interpretation, were just vastly different than was exhibited in anyone's actual practices during this time. People were arrested, fined, and jailed for violating a nonexistent law, and it was used as leverage to elicit plea bargain deals and to tack on time to sentences for other underlying crimes. Not only were all of these practices unjust and inappropriate, on the Pure Corrective picture, but they were flagrant violations of the rule of law.

Worse, I think, the Pure Corrective picture tells us that if the 5-4 majority of the court were in fact wrong in their discernment – as believed by four of their distinguished colleagues who had not been subjected

to prior screening by special interest groups with financial interests in selling firearms before being nominated to the court and hence whom we might have *a priori* grounds to better trust to have impartial judgment about this issue – then the DC Firearms Control Regulation Act of 1975 is in fact *still in force*. For according to the Pure Corrective picture, the court does not have the power to change or remove existing legislation, but only to correct our misimpressions about which laws are in fact in force. So if our impression was not in fact mistaken, then they have not corrected it, but in fact created a misimpression about what is the law in the District of Columbia. But just as it is difficult to see what people were doing between 1975 and 2008 if not following the law, it is difficult to see in virtue of what it is still the law, once the court struck it down.

Those who are compelled by this reasoning may be led to detect a basic mistake in the reasoning motivating the Pure Corrective picture. Given the possibility of judicial review, they may conclude, it is simply impossible to do such simpleminded justice to the idea that courts are not legislatures. Legislative review, as revealed by reflection on what the law was in DC between 1975 and 2008 and what it is now reveals, requires the power to revoke laws. So it simply follows, on this view, that courts do have in this sense some legislative powers. The idea that courts are not legislatures may imperfectly describe the idea that the legislative power of courts is only negative, but it is not a perfect description, because revoking laws *is* itself a legislative power. Call this the *Legislative* picture of judicial review.

The attractions and costs of the Legislative picture of judicial review are exactly the opposite of the attractions and costs of the Pure Corrective Picture. It faithfully captures the intuitive force of the idea that the law in the District of Columbia in fact changed in 1975, and again in 2008 – whether or not it should have. But it leaves unexplained why it might be intelligible to aspire to a separation between legislative and judicial powers, by stretching thin the idea that it can be possible for judges to interpret and apply the law without making it. And it leaves puzzling why, if judges in fact have a legislative power to revoke legislation, we might think that it is proper for this power to be exercised so selectively, and only when evidence is provided about the proper enactment of that legislation in the first place.

The puzzle of judicial review concerns how to resolve the tension between the strengths of these two opposing interpretations of the practice. It is not so sharp a puzzle so as not to admit of many different resolutions. But I now want to suggest that the distinction between inherent and incidental properties is particularly helpful for resolving it. According to a third picture that draws on this distinction, judges *do* have a power to change the law, but it is strikingly limited. For it is important for judges to respect and leave intact the *inherent* features of law. They can change the law, but only in order to correct for ways in which its inherent features come apart from what it is, inherently. Call this, for reasons that I will articulate shortly, the *Michelangelo* picture of judicial review.

According to the Michelangelo picture of judicial review, the Legislative picture is right that judges with the power of judicial review have a kind of broadly legislative power – they can change the law. But what the Legislative picture misses is that this power is radically constrained by the fact that in a very important sense, it is the job of judges *not* to change the law. It reconciles these two ideas by restricting the former to incidental features of the law, and the latter to inherent features. Judges *can* change the law, but their job is to do so only to help the law be more fully what it truly is.

Likewise, according to the Michelangelo picture the Pure Corrective picture is right, in a very deep and general sense, that judges are not legislatures. This is true not just in the shallow sense that their legislative power is only negative and not positive, but in the deep sense that their role is tightly constrained by their job of respecting the law for what it is. But what the Pure Corrective view misses, is that changing the law is not incompatible with the deepest respect for what the law is – on the contrary, a full respect for what the law is should respect and acknowledge what the law is *inherently*, and not preoccupy itself with what the law is *incidentally* – just as an adequate appreciation of the statue Goliath should focus on the aesthetic features that it has inherently, and not the features that it has only incidentally, such as that the paint on its belt is fading. It may not be wrong – indeed, it may actually be called for out of responsibility and respect for Goliath as a great work of art – for the museum to restore the paint on Goliath's belt, if centuries of neglect or improper museum lighting have faded it.

The inherent/incidental distinction therefore offers us a powerful and elegant way of distinguishing between what is right about each of the Pure Corrective and Legislative pictures of judicial review – one that offers not just some sense in which each is right, but one that respects each of their motivations. As in our discussion of statutory interpretation, this distinction is helpful whether we think of laws or "the law" as interpretive objects or not. Every kind of object, I've argued, admits of an inherent/incidental distinction amongst its properties. Everything has some properties in its own right, but counts as having others only because of how it is associated with other things that have those properties. So in the strictest sense, it is the distinction between inherent and incidental properties that creates room for the Michelangelo view, rather than the hypothesis that laws are interpretive objects.

We could try to adjudicate whether the hypothesis that laws or artifacts or the hypothesis that they are interpretive objects better fits with how judges actually determine which statutes are inherently in force, and which merely incidentally. But I'm not optimistic about getting new leverage on this question that has not already been appealed to by one or the other side of existing debates about the proper scope and standards of judicial review. Fortunately, I think that there is a different place that we can get leverage on this question

– one that is internal to the distinctive way that the Michelangelo picture makes sense of the force of each of the Legislative and Pure Corrective pictures.

Recall that the Michelangelo picture grants to the Legislative picture that judges do exercise a legislative power in overturning legislation. They do change the law. But it also respects the impulse behind the Pure Corrective view by insisting that properly deployed, this power must only be exercised in pursuit of a kind of correction. It just allows that this correction may consist in better conforming what is incidental to the law to what it is inherently. That is why judges can carve away laws that are in force only incidentally and not inherently. Because these laws are only incidentally in force, removing them does not change what is inherently the law – it only makes the inherent law easier to discern.

But *why would it be* that it is the role of judges to make the inherent features of the law more closely match its incidental features? The hypothesis that laws are interpretive objects has a tight answer to this question. In the case of interpretive objects, inherent features and those that contribute toward the object being interpretable as what it is interpretable as, and incidental features are those that get in the way. So if laws are interpretive objects, then their inherent features are the things that make them more lawlike – more fully themselves – and their incidental features are the things that get in the way of their being lawlike – that make them less fully what they are, in the same way that the hardness of the rock of the eagle rock makes it less like an eagle. So in the case of interpretive objects – and only in the case of interpretive objects – carving away what is merely incidental is how to make an object more fully and wholly what it is.

This is the perspective that Michelangelo describes taking toward his sculptures. He claims to have been able to discern the figure lurking within each block of marble, and to have seen his task as peeling away the layers of stone in between until the figure already lying within was revealed to the rest of us. By his own conception, Michelangelo's job was not to decide what the figure would look like, but only to reveal it to us. But revealing it to us was not just a matter of telling us where to look – he actually had to carve away a lot of the stone. Good laws should not be as difficult to discern as *David* was within the block of stone called 'il gigante' that lay abandoned in the yard of the Duomo in Florence for decades before Michelangelo took up the commission. They should need less trimmed away. But the diagnosis offered by the Michelangelo picture of the force of the appeal of the Pure Corrective picture comes from the way that it builds on Michelangelo's self-conception as a sculptor. Its justification pushes us to think of laws as interpretive objects.

3.3 Horizontal Precedential Respect

Sometimes courts have to weigh in on questions that they have previously answered. Horizontal precedential respect is the consideration that is owed to the court's own prior reasoning. But the grounds for horizontal precedential respect are notoriously difficult to articulate. In Dobbs vs. Jackson, the U.S. Supreme Court struggled with these grounds, and with how to apply them. In that case, the court was considering whether to overturn its own prior decision in Roe vs. Wade, as well as, prominently, its own prior affirmation of the heart of Roe in a number of other cases. Fortunately, the court had in the past articulated some of its own principles about when to respect precedent – a kind of horizontal precedent about when to respect precedent, if you will. But unfortunately, these principles were themselves articulated most clearly in Planned Parenthood vs. Casey – one of the very cases at risk of being overturned by overturning Roe itself.

In order for horizontal precedential respect to have any force at all, courts must be willing to respect their own prior decisions even when their members believe that those prior decisions were incorrectly decided. If you are not willing to accept any past decision that you believe was incorrectly decided, then you are not giving any respect at all to horizontal precedent. Casey articulated some general principles about when precedents should stand. One of these was that no new relevant information has come to light. If there are no new facts or arguments to bear on a case, then revisiting that case does not promise to bring anything new to it – only, perhaps, a differently composed court. Alito ripped apart this principle in oral arguments with a hypothetical about what would have been the proper course if, the morning after publishing their decision in Plessy vs. Ferguson, the justices of the court unanimously woke up regretting their decision and appreciating that they had made a mistake the day before. The lack of new information should not prevent us, he argued, from correcting clear errors when we recognize them. Another principle articulated by Casey is that widespread reliance on a prior precedent is an important reason to refrain from overturning it. Again, Alito scoffed at this justification. The precedent set in Plessy vs. Ferguson establishing "separate but equal" as the law of the land had been in place for nearly six decades and reliance on it was deeply engrained in social structures throughout the American South.

The central puzzle of horizontal precedent is that it appears to ask courts to limit themselves in ways that conflict with their own best judgment. It presents itself, on its face, as a constraint on the power of courts to do the right thing. When the right thing is right there in front of you, staring you in the face, why would you limit yourself in pursuing it? The conservatives on the Roberts court were not able to come up with any answer that could satisfy them on this score. Rather than carving out restrictions on the principles

about precedential respect articulated in Casey, they rejected all such principles without offering anything in their place that might offer guidance about when they might respect horizontal precedent in the future.

Because the force of the problem of precedential respect comes from the fact that it presents itself, on its face, to the deliberating members of a court as a limitation of their powers, it feels like answers to the problem have to justify that limitation. They have to explain what value is lost, or what right violated, by failing to respect precedent. But the problem with all such explanations, is that no matter what value we say is lost, and no matter what right we say is violated, the practical problem of when to respect horizontal precedent is going to arise in cases in which the underlying values and rights at stake are just as important – if not more important, in the considered views of the justices trying to evaluate whether to take this constraint seriously. In the case of abortion, for example, if your personal view is that fetuses have a moral right to be carried to term, then asking you to respect a precedent that does not acknowledge this right for the sake of some kind of procedural justification or for the sake of consistency in how the law is applied over time does not look like the right kind of value to address your underlying concern about the injustice or incorrectness of the original decision. It has the wrong scale to answer what you think is at stake in the size and consequences of the original mistake.

That is why I think that it is wrong-headed to try to justify respect for horizontal precedent by appeal to its value. It may have all sorts of values, but when those values come into competition with the values at stake or perceived to be at stake in the underlying issue of law about which a precedent has been established, these are arguments that persuade only people who already think that the precedent was correctly established, or who at least like what it does. None of them have any teeth with a jurist who regrets the earlier decision.

A better answer to the problem of respect for horizontal precedent must reject the intuitive way that it presents itself to us – as a limitation on the power of the court to do the right thing. Respect for precedent *feels* like a limitation on our power to do the right thing, because it gives our past a kind of power over what we do now. And it is natural to think that the more power our past has over us now, the less power that we have – for that power is usurped by our past. But this is wrong. It is an elementary mistake about how agency works when extended over time. The power that our past has over our present is the same power that we now have over our future. If our *past* has no power to constrain what we do now, then we *also* have no power to enact our judgment onto the *future*. A court – indeed, any decision-making body – that holds no respect for its own precedents is *not* more empowered to do the right thing. On the contrary, it is *less* empowered to do the right thing.

The idea that laws are interpretive objects offers, I think, a particularly powerful way of working out this thought. Return to Umbhali's decisions about how to write chapter eight of her serial novel whose first seven chapters are already in circulation. Umbhali can't carry out writing the rest of her novel now – she doesn't have enough time before the publication deadline for chapter seven. So she doesn't have direct control over how the rest of the novel will go. What she does have control over right now, is the text of chapter eight. By crafting chapter eight in a way that gives her future self the tools to make the most of it, she can steer her future decisions – indirectly. But the power that she has now to shape her future decisions is built on the same reasons why her decisions now must make sense of what she wrote in chapters one through seven. Her power over the future comes from the same place that gives her past power over what she does now.

It's not that Umbhali can't throw out her original plan for the novel. She can. She can also throw out her former interpretation of the events of chapter seven. And she can even retract the events of chapter seven entirely. She can start chapter eight with "and then she woke up", or introduce new evidence of the unreliability of the narrator. But what she can't do, is pretend that chapter seven was never published. Like it or not, chapter seven is part of what she has inherited when it comes time for her to write chapter eight. Its text is already out there and part of the novel that she is crafting. If she has regrets about the choices that she made then, she has a lot of tools at her disposal for finessing the consequences of those choices. But there are consequences, in turn, of how she deploys those tools.

Suppose, for example, that Umbhali decides to throw the events of chapter seven out. It was such a bad direction for the novel to take, she now decides, that chapter eight must begin with "and then she woke up". So far, so good. But next month, when it comes time to write chapter nine, Umbhali might again re-think her plan. Influenced by disappointed fan tiktoks, she might try the same strategy again, revealing in chapter nine that the true dream sequence was the one in chapter eight which apparently revealed chapter seven to have been a dream. Nothing that she can do now in crafting chapter eight can prevent her future self from writing chapter nine in this way. But at some point, if she tries revealing in chapter ten that chapter nine was a dream sequence and so the revelation in chapter eight that chapter seven was a dream sequence is veridical after all, or perhaps if she switches tracks again in chapter eleven, her novel is going to devolve into incoherence. The tenability of the interpretability of her book as a single work extended across chapters that are published serially – and not just as a series of disconnected short stories that take inspiration from one another – depends on how those chapters fit together. So by starting chapter eight with "and then she woke up", Umbhali is placing a bet that she will not change her mind again.

The justices who overturned Roe were also placing a bet that future courts would not overturn their own decision. Because they couldn't articulate any principle for respecting past decisions that are now judged to have been improperly decided, they not only can't rule out that this could happen, but by their own reasoning, it will be the right way for any future court to proceed, if the members of that court disagree with their current decision on the merits.

This is precisely the problem that Justice Kagan articulated in oral arguments for Dobbs:

> General, Justice Breyer started with *stare decisis*, an important principle in any case, and, here, for the reasons that Casey mentioned, especially so, to prevent people from thinking that this Court is a political institution that will go back and forth depending on what part of the public yells loudest and -- and -- and preventing people from thinking that the Court will go back and forth depending on changes to the Court's membership.

Kagan seems to be presupposing that there is something wrong with the court being a political institution that can change its mind depending on changes in its membership, and I suggest that the picture of laws as interpretive objects can help us to say what it is. *Law* can't be the kind of thing that oscillates simply on the basis of who is applying it. Or at least, anything that oscillates in that way is thereby less law-like. Starting each successive chapter of her serial novel with "and then she woke up" threatens to undermine whether Umbhali is really writing a novel at all. Similarly, oscillating in its interpretation of constitutional principles and in which precedents it affirms threatens whether there are any laws at all.

It is important to distinguish two different points that I want to make, here. The first is that respecting precedent is not properly understood as a way in which the court is *limited* by its past, but rather, as a way in which the limits imposed by its past actually give it greater power over the future. The idea that laws are interpretive objects tells us something about where these limits come from. If the court's past decisions are like Umbhali's earlier chapters, then they are not just *commentary* on the interpretive text of the law – they are actually part of the text to be interpreted. So it follows from this picture that even when a law is originally grounded in statutory language, the ongoing existence of the law, as it is reviewed and applied, comes to be grounded not only in the statutory language, but in the ongoing body of practices of application and review. *That* is how the court's earlier verdict comes to be part of the interpretive text that the justices must now confront – even if they now reject it. Rejecting it requires adding something like "and then she woke up" to the interpretive text of the law.

The second point that I want to make goes further. Counting the court's past decisions as part of the interpretive text of the law does not require endorsing those decisions. Requiring otherwise – a kind of absolute principle of horizontal *stare decisis* – would be like requiring Umbhali not to change her mind about

what really happened in chapter seven. The reasons why the court doesn't have to always take their earlier decisions at face value are therefore exactly the same as the reasons why Umbhali doesn't have to. But just as Umbhali can't throw everything out and start over in goth-punk postapocalyptic Tokyo, there are also limits on how much can be thrown out. Every time the court throws something out they are threatening the respect in which the practice that they are carving out resembles law in the first place – and hence threatening whether they are really helping a law to emerge more clearly as itself, Michelangelo-style, or instead replacing it with some other law, Legislature-style – or even undermining the existence of law in the first place.

4 Laws as Interpretive Works in Progress

And this brings me, at last, to the question of whether laws are not just interpretive objects, but what I have called interpretive *works in progress*. The Michelangelo Picture suggests in a way that they are. For it allows that statutory enactment does not complete a law, but leaves room for the process of judicial review to *refine* that law. Not for the arbitrary purpose of turning it into something else, of course. But in order to help it to emerge more clearly as itself – in order to remove incidental features that get in the way of seeing its inherent features for what they are.

The problem of horizontal precedent is, in a way, the flip side of the way that the Michelangelo picture encourages us to see laws as unfolding over time. If a law unfolds over time, then even the most careful efforts to refine that law only in ways that allow it to emerge more clearly as itself can go astray – we can make mistakes in our judgments about these matters, and judges can be the ones who make those mistakes. But if these mistakes become part of the interpretive text of the law, then they can inadvertently *change* what is inherent to the law – or at least, they can make it the case that the answer to what is inherent to the law is different than any evidence would have revealed that was available at the time of that decision.

These decisions, we may think, are in an important sense improperly decided. They shape the law in ways that would not have been licensed by strict attention to the goal of only changing the law in ways that help its incidental features to more closely match its inherent features as they have already been revealed. But just because the decision was improperly decided at the time doesn't mean that it hasn't become part of the interpretive text of the law. Similarly Umbhali can make mistakes about how to best carry forward her story – and she can realize later that she made these mistakes. But the mistakes that she has made have still become part of the interpretive text that she inherits when she sits down to write chapter eight, and the fact that it was a "mistake" to make them at the time that she wrote chapter seven is neither here nor there when it comes to her current task of doing the best she can with what she has been given.

Similarly, once we acknowledge that judicial verdicts become part of the interpretive text of the law, we have to acknowledge that the task of shaping the law into the most authentic version of itself and the task of shaping the law into the most authentic version of what we thought it was going to be before an earlier judicial verdict was released must come apart. The latter task requires revoking verdicts that we believe were improperly decided. But as we've seen, it gives us power over the past only at the cost of robbing us of power over the future. And it is an adulteration of the judicial duty to only carve the law in ways that render it more authentically itself. It is the former task, in contrast, that is the most direct articulation of the contrast between the roles of judges and legislators. Similarly, when Michelangelo inherited the block of marble that turned into David, it had already been blocked out by earlier sculptors and contained imperfections that prevented him from carving a figure with the kind of musculature that is characteristic of most of his work. He wasn't in a position to try to glue the missing stone back on, in order to leave room for David to be more muscular, as he would have liked before earlier sculptors carved off too much. Instead, he had to work with what he inherited. That is the role of judicial precedents.

Still, even though I argued earlier that the job of judges is different from that of legislators – and that it is different because of how it requires respect for what the law *is* – once we recognize that judicial decisions *can* shape what is inherent to the law, and what is merely incidental, judges must confront the fact that what is inherent to the law is *not*, after all, independent of their decisions, but depends constitutively on it – as well as on what happens as future courts consider similar issues.

When Umbhali's fans are publishing tiktoks speculating about the direction of her novel, all that they have to go on are the clues laid down in chapters one through seven. But Umbhali is not so constrained. She is licensed to adopt an interpretation that would only be second- or third-best of chapters one through seven, but can be *made* best for chapters one through eight by writing chapter eight in the way that is required in order to sustain it. It would be bad faith for Umbhali to explain to her editor that she couldn't write chapter eight in such a way, because that would fail to be faithful to the novel as revealed in chapters one through seven. This disavows real ongoing freedom and responsibility that Umbhali has over what kind of novel she is writing.

Similarly, once we acknowledge the role of judicial decisions in setting precedents, we can see that it is also bad faith for judges to decide only on the basis of past decisions – as if their own present decision has no power to help to shape the law. Once you appreciate that what shape the law has can be affected by your present decision, the contrast between judges and legislators and corresponding duty to only shape the law in ways that allow its incidental features to more closely match its inherent ones can leave your present course undetermined. Because what is inherent to the law depends on what you decide now, it can turn out that

either of two verdicts can count, once you make it, as helping to reveal the law more closely for its inherent features.

And this observation leads me back to a different lesson than that which Dworkin takes from the case of *Riggs vs Palmer*, which he discusses so prominently as evidence that judicial reasoning can and should draw on moral principles. In this case, the justices of the New York Court of Appeals reversed a lower-court decision allowing Elmer Palmer to inherit his grandfather's estate as a consequence of murdering him, on the grounds of the principle that no one may profit by their own crime. Dworkin famously argued that this case reveals in a stark way something that all judges implicitly understand in their daily practice – that the content of the law is shaped not just by statute and practice, but also by moral principles.

I'm not sure that the law *isn't* partly shaped by moral principles – if laws are interpretive objects, as I have been suggesting here, then the question of whether they are is essentially the question of whether being more consonant with moral principles is the right kind of thing to make a practice *more lawlike*. And I'm very open to the idea that it is. But I don't think that Elmer's case is a very good one for establishing this result, for the very simple reason that the principle that no one should profit by their own crimes is such a compelling one that it seems reasonable for the justices of the New York Court of Appeals to have been quite confident that once they had articulated this principle, successive courts would go along with it. And so given this confidence that others would go along with it, they were in a position where either decision that they made could have counted, in retrospect, as carving out the law in a way that allowed it to more closely incidentally approximate its true inherent shape. And this put them in a position in which they had a true decision to make about which law would be better to have.

Is it bad for judges to have this power? Is this too much like the power of the legislature? My answer to these questions is that judges can't help but have this power. We all have limited versions of it, insofar as what is truly the law depends in complex ways on social practices of conformity, and each of us plays a small role in these social practices of conformity. And as I have argued, judges can't help but have this power. It comes from their ability to shape precedent, and their ability to shape precedent is what gives them power to decide for the future, and not just the present. And it sounds like a legislative power. But it is a power that judges take up only when they are truly in a position, through their present actions, to make the law inherently either of two ways. And so it is tightly delimited by the principle of respect for what the law truly – that is, inherently – is.

As I have argued, judges must sometimes be in such a position. But they will not always be in a position to know that they are in such a position – and whether they are depends on the position that they are in to bet on the future. Because respect for precedent is not absolute, any future court could revoke their

decision, and then it would turn out after all that their decision did *not* hew our legal practice more closely to the law's true inherent nature. So I am very far from suggesting that judges – even Supreme Court justices – should be thinking of their decisions often in practical terms. But I do think that if laws are interpretive works in progress, then denial of the judicial power to shape the law, and of judicial responsibility when given such opportunities to shape it well, is only an exercise in bad faith.