

Can We Learn about Real Social Worlds from Fictional Ones?

Todd Jones

University of Nevada Las Vegas

ABSTRACT: It is very common for social scientists to be asked whether their findings about human nature could also be learned from reading great works of literature. Literature teachers frequently assign readings partly to teach people important truths about the world. But it is unclear how looking at a work of *fiction* can tell us about the real world at all. In this paper I carefully examine questions about the conditions under which the fictional world can teach us about the real world.

KEYWORDS: social science, fiction, justification, knowledge, education

1. INTRODUCTION

When talking to colleagues from the English department, it is not unusual for social scientists to get some joshing along the lines of “but didn’t Dickens teach us this about human nature already?” More than a few professors of literature believe that, whatever other functions literature can serve, fiction can be an excellent teacher of the things that the social sciences are supposed to teach us. Hardly a day goes by when an essay is not published in a newspaper or blog talking about (in defense of a humanities-oriented education) how fiction is an excellent guide to understanding human events. The literary critic George Steiner was quite clear about this sentiment, once saying that “Molière and Stendhal would always have more to teach us about human thought and character than all the cognitive scientists there ever have been or ever will be.”¹

But is anything like this really true? It may be a routine assumption among people who have taken high school or college English classes that *The Red Badge of Courage* can teach us about the psychology of soldiers. But that shouldn’t stop us from noticing that, on the face of it, the claim seems immensely improbable. Stephen Crane wrote his novel about the inner life of a Civil War soldier, despite his having been born six years after the war ended. Crane himself wrote, “Of course, I have never been in a battle, but I believe that I got my sense of the rage of conflict on the football field.”² The thoughts and actions depicted in Crane’s Civil War novel are the product of his creative imagination (as were those of his

earlier novel, *Maggie: Girl of the Streets*). Why should we have any presumption that we can learn about real world thoughts and actions from looking at these?

The idea that studying fictional human beings is a good way to learn about real human beings, then, is an idea that, oddly, seems both trivially true and preposterous at the same time. In this paper I want to *carefully and systematically examine* questions of whether and how we can learn about the real world from fiction. The degree to which such an important question fails to be systematically examined is striking. Business minded “Philistines,” radical Oscar Wildean aesthetes, and skeptical post-modernists, whatever their differences, are united in their assumption that the function of fiction is *not* to produce knowledge. By contrast, the humanities-defending blogosphere and most literary critics seem to be united in taking the idea that we can learn from fiction as an axiomatic first principle. Detailed discussions of what this or that novel teaches us often proceed completely independently of any considerations of *how* a fictional world can say anything at all about a factual world. As Gregory Currie observes, “There is a puzzling mismatch between the strength of opinion on this topic and the state of the evidence. In fact, I suspect it is worse than that; advocates of the view that literature educates and civilizes don’t overrate the evidence—they don’t even think that evidence comes into it.”³ The Philosophical literature on learning from fiction does look more systematically at arguments.⁴ But on my reading of this literature, most philosophic defenses of learning from fiction do not sufficiently address the Achilles heel of such claims: It is very easy to pick up *beliefs* about the real world from reading about fictional worlds. Fiction could well lead us to believe that opposites attract. But this is a far cry from enabling us to have *justified true beliefs* about the world, which is required for acquiring *knowledge*. This is an issue that Plato worried about quite a lot, and I do not believe his concerns have been answered. Scholars like Nussbaum,⁵ Palmer,⁶ Gibson,⁷ and numerous others are very good at giving various defenses of fiction’s ability to give us new ideas about how the world might be. But it’s much less clear how fictional literature can give us *justified* beliefs about how the world truly is.

In this paper, rather than focus on these issues in the way they have (or haven’t) been debated in the literature, I want to proceed by directly looking at the conditions under which it could be possible for a fictional world to give us knowledge about the real world. I’ll then look at when and how well fiction might satisfy those conditions. Is it really likely that we can learn more about human happiness from Tolstoy than from Jonathan Haidt? It’s high time we seriously examine that question, rather than just assuming that it is obvious that we can, or obvious that we can’t.

2. LESSER (BUT STILL SERIOUS) PROBLEMS

Let’s begin by noting that, for us to gain knowledge about the real world from fiction, at least three things must happen:

- 1) a certain belief about our world must arise (because of seeing a similar fictional world situation represented)⁸

- 2) such beliefs must be new and non-trivial (so one is truly learning and gaining new knowledge, not merely reinforcing what's already well-known),
- 3) such beliefs must be justified (if the person doesn't have adequate evidence for thinking the belief is true, then she hasn't acquired *knowledge*).⁹

In this paper, I will concentrate on the difficulties surrounding this third step. The most important obstacle to be overcome by someone arguing that we can learn from fiction is to show how the beliefs we acquire from reading fiction are justified. We can form various and sundry beliefs about the world as we read fiction. But unless we have some *evidence* that these beliefs are true, the fiction has only increased the number of *possibly true/possibly false* beliefs we have about the world—not our *knowledge* of it. It's worth noting briefly, however, that meeting the other two requirements just mentioned is much more difficult than generally realized.

Which sorts of beliefs are likely to be formed, when someone reads a given work of fiction? We don't really know. When we watch our colleagues look at, say, a fire in an impoverished neighborhood, we have no idea what sorts of particular or general beliefs they will walk away having. We are at a similar loss as to what they will come to believe when they encounter a burning building in a novel like *The Museum of Extraordinary Things*. Will they come to have beliefs about the vulnerability to fire of similar looking buildings? Will they have no thoughts about social inequality? It's completely unclear. This is one of the reasons why most classic works of literature (e.g., *Macbeth*) leave people making dozens of different sorts of claims about what the work is all about. People can come to believe all kinds of different things upon reading such works. The situation is quite different for a typical non-fiction text where we have much more of an idea of which beliefs readers are likely to walk away with. The authors of essays are usually explicitly and deliberately trying to get people to believe particular propositions. They may or may not succeed. But we can usually tell what people will believe if they do succeed. But the events depicted in given work of fiction, like any observed sequence of events, can easily create all kinds of different general beliefs. The consequent probability of it tending to impart to readers *any given belief* is comparatively low. One of the obstacles to getting people to learn certain things about the world by reading fiction is that we just don't know what they will walk away believing.

There is also the problem that we can't really learn new things from works of fiction if we *already know* them. To increase our *knowledge*, we must acquire beliefs that we didn't have before.¹⁰ This worry can be thought of as a version of a problem that philosopher Jerome Stolnitz terms "the cognitive triviality of art."¹¹ Stolnitz argues that much of what we are said to "learn" from great literature turn out to be platitudes that we already knew. He argues, for example, that, in the end, *Pride and Prejudice* mainly conveys the idea that two attractive people who fit together can be kept apart for a long time by . . . pride and prejudice. One way to describe Stolnitz's worry is that whatever we may be gaining from looking at fiction, we can't describe it as acquiring *new knowledge*.

Now, it is, of course, also true that the short summarizing ideas that Stolnitz describes aren't the only ones we pick up from literature. Literature also often provides us with rich interconnected exemplars comprising *sets of* ideas. "Farm work is hard" could hardly count as new knowledge for most of us. But complex subtle pictures of the various difficulties of farming life that one acquires after reading Willa Cather's *O Pioneers!* might well be a candidate for something new one has learned from fiction. But, as we'll be discussing, such belief sets, however new and detailed and impressive cannot count as new *knowledge* unless and until such beliefs are *justified*. Meanwhile, Stolnitz's worry is certainly worth considering. In sections below, we will be looking mainly at how we can be justified in believing the ideas acquired from fiction. We've just seen however, that even if problems surrounding justification can be solved, we will still have to deal with the problem of whether readers are likely to form a given belief *to begin with*, and whether a belief picked up from fiction is likely to be one that they didn't already have. These problems will have to be focused on elsewhere.

We can start thinking about justification issues, by noting that there are many ways in which we could come to have justified beliefs about the world. Consider perception. When we look at an intersection in normal light, and form the belief that "a blue car is stopped at that traffic signal," our belief is justified if we have no reason to think our perceptual abilities aren't up to snuff. Or consider certain types of testimony. If a friend you know to be trustworthy tells you that Thomas Jefferson owned slaves, you can justifiably believe that Jefferson most likely owned slaves, since there is a presumption of truth-telling between friends. If you know your friend is an expert in history, then you have even more reason to believe this is a fact. Reading fiction is not like this. If we read a story in which the police burn a suspect they are interrogating with cigarettes and form the belief that police sometime burn suspects, such a belief is not immediately justified. Testimony and perception are processes *designed* to give people true beliefs. But that is not the primary function of fiction. Without more evidence, we don't have good reasons to believe that the events and circumstances in the fictional world tell us about things that are true in the actual world.¹²

I think there are two main families of ways to get the additional things needed for worlds created by a writer's imaginations to give us justified beliefs about the actual world: 1) looking at fiction can enable us to pick up hypotheses about the world we didn't have before, that get confirmed/justified elsewhere, and 2) we could come to have reasons for thinking that the fictional world closely resembles the actual one.

3. ACQUIRING JUSTIFIED TRUE BELIEFS FROM FICTION: PART 1—EXTRA-TEXTUAL CONFIRMATION

In what I'll call the "extra-textual confirmation" family, we may pick up *beliefs* about the real world from reading fiction, but none of the *justification* of those beliefs comes from the fiction. The justification comes from elsewhere. For example, *Caleb's Crossing* is a novel by Geraldine Brooks about the life of Caleb Cheesah-teaumuck, the first Native American to graduate from Harvard. Having read the

novel, I now believe that in colonial New England many towns were destroyed in fighting between settlers and Native Americans, that Harvard originally had an “Indian School” attached to it, and that the Puritans often drank small beer for breakfast. But if I now know (rather than just believe) these things, I didn’t come to know them upon completing the novel. I came to know them when, after the novel whetted my curiosity, I looked up various things that experts had written about these details. I came to know these things via 1) the novel creating a number of suppositions, and 2) my then coming to find many were true through evidence from outside of the text. Very different sorts of things about the world can be learned this way. Novels can lead people to form suppositions about everything from the nature of free will, to the causes of divorce, to the best way to cure deer hides. This sort of process could tell us about what has happened in the past or what could possibly happen in the future. Reading *Huckleberry Finn* probably *won’t* lead us to form beliefs what actual people named Jim and Huck did on the Mississippi before the Civil War, given we are well aware that they were fictional (though doing so certainly isn’t impossible). But many other suppositions we may happen to form—about particular events *resembling* these (particular runaway slaves in 1850) or about the properties of more abstract *types* of people and events (American slaves, moral conversions)—will be candidates for things we come to know through this supposition and verification process.

3.1 Different Types of Extra-Textual Confirmation

There are different kinds of extra-textual confirmation processes. One big subfamily could be called *active* extra-textual confirmation. These are processes where one does specific activities *with the aim of* finding out if certain claims are true. Let’s start with a simple example (while noting that the same point will also apply to complex ideas about psychology or morality). Suppose one reads in Oscar Wilde’s story, “The Fisherman and his Soul” of Tartars drinking mare’s milk, and so forms a supposition that mare’s milk can be harvested and be a good form of nutrition for a group. One could actively independently confirm this by milking and drinking the milk of the horses at a local farm. Or one could inductively reason that it’s likely that mare’s milk could serve as a source of nutrition, given how often cows’ or goats’ milk is a staple where cows and goats are raised, and how many of the world peoples keep herds of horses. Or one could consult a trustworthy expert on dairy practices throughout the world to see if mare’s milk is used anywhere. This sort of fiction-reading, supposition forming, and active extra-textual confirmation process is one way that one could come to have knowledge about the world on the basis of reading fiction.

Another subfamily could be called *passive* extra-textual confirmation. Here, the same sorts of justifying evidence comes to the agent as in active verification. But with passive verification, the agent doesn’t go out of her way to *seek out* the evidence. An example of passive verification would be reading about mare’s milk in fiction, then later happening to see someone milking a horse on a farm, or on a TV special. Someone might hear an anthropology professor mention it in a lecture. She could even have a stream of consciousness meandering about

shepherds, cows, and goats that leads to the thought that herded horses should be able to be used as a source of milk. She might then remember the supposition she had after reading “The Fisherman and His Soul” and how the plausible conclusion of her meanderings now confirms that supposition. Even though the evidence has passively come to her, it’s still evidence that the claim is true. And if it is indeed true, the person now has a new justified true belief, new knowledge about mare’s milk—an idea she first picked up from fiction.

There is also a third subfamily of extra-textual confirmation in between active and passive which one might call *background knowledge* extra-textual confirmation. In background knowledge extra-textual confirmation, memories of general or particular things are elicited by a situation in the fiction, and those memories can be used to help indicate that one can justifiably believe in the existence of the situation described. Reading about the Tartar’s drinking mare’s milk might remind one of a farm family one once knew, who, come to think about it, did some morning milking, but had no cows or goats. The memories of this family—whose activities weren’t really fully noticed or comprehended until the mare’s milk passage called attention to it—helps justify the fiction-induced supposition that some people might well nourish themselves with mare’s milk. Or reading that passage about mare’s milk might remind someone that horses are large mammals, that every kind of large mammal can produce large amounts milk for its young, and therefore, by statistical syllogism, horses are capable of supplying milk for people. This isn’t quite active verification, since one need not go out of one’s way to look for evidence that a claim is true. Nor is it passive verification since one doesn’t just happen upon the evidence. Here the supposition induced by the fiction itself automatically elicits stored information from associative memory, and such information can start producing inferences. Such memories and inferences can be used to justify suppositions that fiction can induce.

How likely is it that readers will learn from fiction via processes of extra-textual confirmation? When people go out of their way to do *active* verification, it’s very likely. It’s especially likely when they use especially reliable processes like consulting well-regarded experts. This kind of verification of ideas picked up from fiction certainly sometimes happens with especially intellectually curious readers. Unfortunately, there’s little evidence that most ordinary consumers of fiction spend more than the tiniest fraction of the myriad of hours they spend reading or watching and thinking about fiction attempting to actively verify suppositions.

When we engage in extra-textual confirmation via background knowledge, we can also learn from fiction. But that isn’t as likely to happen as one might think. We won’t, for example, genuinely be learning something new if the fiction elicits a memory that simply *confirms* a supposition. If we already knew there were X’s that have Y’s, and pull that information from memory—then the fiction isn’t teaching us something we didn’t already know. We do learn when we elicit something like *pieces* of a heretofore-unknown claim or the premises leading to it, and then bring them together. But for learning to happen this way, a number of circumstances must be in place. We need to be reading about things in fiction about which we happen to have lots of information in memory stored in memory about. (E.g., when a person knows lots of information about nineteenth-century

whaling and reads *Moby Dick*.) The fiction has to be structured in such a way, and one's memory has to be structured in such a way that the fiction actually elicits the information-pieces. And one's inferential structure has to be such that it puts the information together in a way that produces a new conclusion. Only if all those conditions are met can background information verify that something you read about in fiction is actually the case. One should note that some types of suppositions require much more background knowledge than others to confirm. The idea that some state of affairs is *logically* possible, for example, can be arguably demonstrated just by showing it's conceivable. (So *this* is something that fiction could be excellent at doing.) It takes more background to show that it's physically possible in our world, given its laws or given certain circumstances. It takes still more to show that such a state has actually existed somewhere.

What about passive extra-textual confirmation? It looks like this is unlikely to be a major way of gaining true knowledge. Recall that we should only talk about something as "learned" if we didn't already know it. That means that among the things most likely to provide *knowledge* is information about things that are relatively exotic to readers—things beyond their everyday experiences. One could potentially *learn* something about drinking mare's milk from fiction because it's not the sort of thing one likely already knows about. But one hasn't yet gained *knowledge* until one is *justified* in believing one's new supposition about the world. But that very newness and exoticness is just what makes it unlikely that one will come across a justification through passive extra-textual confirmation. An American reader is very unlikely to happen upon any Tartars or anyone else drinking mare's milk in their everyday lives.¹³ People are also unlikely to just happen to start *reasoning* about such matters. We could get a justified belief about such exotic matters by hearing *experts* in the realm talk about them. But how likely is it that we will coincidentally *happen to hear* experts we haven't sought out talk about a realm we heard about in fiction and remember that we did?

Suppose however that we are lucky enough to happen to receive just this sort of passive confirming evidence. Would we then have really gained new knowledge, new justified true beliefs from fiction? We would—but, with passive extra-textual confirmation, the evidence that happens to confirm the suppositions we've picked up from fiction, gives us information that would be provided, *whether or not we had ever gotten the supposition from the fiction*. Active confirmation has us going out and collecting information we would never have gathered, had the fiction not made us curious about it. Such justifying information gives knowledge we never would have had were it not for the fiction. But passive confirmation has justifying information coming our way with or without the fiction. While seeing the idea beforehand in fiction may prime us to pay more attention to it or remember it better if it's dramatic, we can't clearly say that such new information is really a matter of fiction teaching us something we wouldn't have otherwise known about the world.

It appears then, that one of the ways we can learn about the real world from fiction is when fiction gives us some suppositions about the world for which we later, independently of the fiction itself, get evidence of their truth. But if we do this without actively seeking out reliable evidence, we might or might not happen

upon the requisite justification of these suppositions. And if we do happen upon it, we gain no more information about the world than the evidence itself gives us, so we don't really learn about the world. We spend a lot of time acquiring ideas which may or may not be true, all the while running a considerable risk that we will have numerous false beliefs about the actual world. If one were to spend that time reading experts, making direct observations, or even just talking to friends, it seems we would likely learn far more (without nearly the risk of coming to have false beliefs).

Note that this point about ideas acquired from fiction needing justification to count as knowledge holds *no matter what form of representation* the ideas take (e.g., it doesn't matter if the representation is propositional or not). Indeed, the same point can be made even if one thinks that what fiction does is create *dispositions* to successfully interact with the world. Fiction might well tell us "how to" interact with the world to get a certain result. But these dispositions will have been formed in response to simulated interactions with a *different* world from ours (a fictional one). What we need are reasons to think that ways of acting would likely work to produce their intended effects in the possibly very different *real* world. The fiction alone will not give us any such reasons. We would still need some version of extra-textual confirmation or the other justifying reasons discussed below.¹⁴

4. ACQUIRING JUSTIFIED TRUE BELIEFS FROM FICTION: PART 2—RESEMBLANCE

What other ways are there of getting knowledge of the world from fiction besides various forms of extra-textual confirmation? Well, we could know that certain things about our world are likely to be the case if we "saw" that certain things were true in the fictional world, and we had good reasons to think that such things in our world were the same as, or very much like those of the fictional world. In Victor Hugo's *Les Misérables*, Inspector Javert achieves his lifelong goal of recapturing Jean Valjean, but he finds it gives him no peace. If things in the real world work like they do in the fictional world, perhaps you should warn your obsessed-with-an-escaped-suspect policeman uncle that capturing that suspect is not the thing most likely to make his life better. Maybe your wrestling-championship-focused son would be better off not putting all his eggs in one basket.

We could be justified in forming beliefs this way, if we have good reasons for thinking that the relevant things in our world are really like those in the fictional world. If we could be confident there was such a likeness, we would not need extra-textual confirmations of the suppositions we get from fiction. Do we ever have good reasons for thinking that our world is like the fictional one?

4.1 Necessary Truths

One of the reasons we might think that *our* world has certain properties that we've seen in a fictional world, is that we note that certain properties must be necessarily connected in *any* possible world. We might see a set of circumstances *abc* in the fictional world and note that it is a *necessary truth* that property *X* must be present there, too. If properties *X* and *abc* are necessarily connected, that means

that if we have *abc* in our world, then, we must also have *X*, just as they were necessarily connected in the fictional world. Indeed, for *necessary* truths about *abc*, all possible worlds with *abc* will have these properties. If *all* possible worlds must be like the worlds in the fiction in this respect, then we automatically know that *our* world must be like it in this respect. So, for example, if one notices that, for the characters in a work of fiction, defecting is necessarily always the rational strategy in prisoner's dilemma type situations, then it must necessarily be the rational strategy in our world as well.

I will not discuss learning necessary truths from fiction extensively here. The kind of knowledge about human thinking and behavior we are discussing here, is mostly about *contingent* things that happen to be true of people, not analytic truths about all possible beings. Nevertheless, a brief discussion of some of the problems of uncovering analytic truths from fiction is in order.

The underlying idea of how one might learn necessary truths from fiction is familiar to anyone who has ever had an introductory philosophy class. A standard way that philosophers try to demonstrate conceptual necessary truths to people is by giving them *fictional vignettes* such as Gettier cases or Trolley problems, and inviting them to see that in these cases and in all relevantly similar cases, properties *X* and *abc* (e.g., proper action and refraining from treating someone as a means to an end) must be linked. If short vignettes could teach us about necessary truths, why couldn't more elaborate works of fiction?

But there are many problems with assuming that fiction can teach us about the world this way. In one of the most influential papers ever published in philosophy, Willard Quine argued that *no claims were entirely analytic*. Our justifications for believing any claim always require some empirical information gleaned from examining the world. Claims that were once thought to be necessary truths like "a whale is a big fish" have turned out not only to be not analytically true, but not true at all. Quine argued that even our most firmly held beliefs might one day be undermined by empirical findings in other realms. If Quine is right, and there are no analytic truths to be found *by any methods*, then there are no analytic truths to be found by contemplating fictional examples. We can't unproblematically assume that fiction can tell us about analytic truths, without first making some headway against decades of arguments by Quineans concluding that there aren't any analytic truths to be had.

A further reason for skepticism about the idea that we could uncover necessary truths by looking at fictional situations and consulting our intuition about what seems necessary comes from experimental philosophy. Philosophical theorizing often consists of describing a (fictional) scenario, say one in which a person pushes a large man off a footbridge to stop a runaway train from killing five innocent people.¹⁵ We are invited to examine our intuitions about the situation to see what heretofore unknown properties of the situation would be entailed by the existence of that scenario. Philosophers typically claim that fully envisioning the fictional situation enables most people to intuit necessary truths like the wrongness of killing to save a larger number of lives. Perhaps what we do when we read more elaborate fictional scenarios like the one in *Huck Finn* where Jim is captured, is to consult our intuitions in just this manner and see the moral wrongness of slavery

for any people. Perhaps reading Conrad tells us necessary truths about the nature of self-knowledge.

But for the last two decades, however, many experimental philosophers have argued that intuitions are not stable reliable indicators of analytic truths at all. These philosophers have claimed that intuitions about what must be true about scenarios described in philosophical vignettes are influenced by factors such as gender,¹⁶ native language,¹⁷ the order in which cases were presented,¹⁸ the smells and cleanliness of the environment,¹⁹ and even font size.²⁰ If intuitions about the cases examined are reliably tracking necessary characteristics that exist in all possible worlds they shouldn't vary at all, and they certainly shouldn't be influenced by things like the type of font the scenario is described in. If intuitions are as malleable as many X-phi practitioners claim they are, then our thoughts about what must or must not be the case in a situation described in a fictional vignette are not reliable guides to what can and can't be counted on to be necessarily present. Philosopher Noel Carroll has defended the use of looking at literature for understanding necessary truths on the basis of literature transmitting ideas using essentially the same techniques as philosophical thought experiments: "If philosophy conducted by means of thought experiments is an adequate source of knowledge and education, then so should literature be."²¹ But that means, conversely, that if the vignette-examining and intuition pumping methods of philosophy are an *unreliable* source of knowledge, then they are unreliable source via literary study as well.

And even if we could pick up necessary truths from fiction, we need to ask whether the kinds of fiction we can best learn them from are the sorts of literary works often lauded as sources of wisdom (long tales whose meanings even experts argue over). Are we really more likely to clearly see what is and isn't necessary in all possible worlds from these, rather than from short vignettes carefully written for that purpose, where the intuitions elicited are explicitly discussed and debated?

Knowing that the world and fictional examples are similar through sharing common analytically necessary structures does seem to be one way that fiction could potentially tell us about the world. But we need to keep in mind that there have long been many challenges to ideas about uncovering necessary truths through the use of intuitions and non-empirical investigation. And it seems likely that even if ideas about necessary truths can be gained by looking at fictional examples, we would be much better justified in our ideas, if there were further explicit examination of the sources of the intuitions or of possible counterexamples. Such discussions are available in essays; they rarely are in the fictional stories themselves.

4.2 Analogy

We needn't think that our world must have certain features because we've come to see that *all* worlds of a certain sort do. We might well come to think that our world has that feature because it is very much like another *particular* world that has that feature—and that other world can be a fictional one. Perhaps we can sometimes learn about the world from fiction through straightforward analogical reasoning.

A standard way of reasoning by analogy about anything is to start with the observation that some familiar sample item and a lesser-known target each have a set of properties (call them *wxy*) in common. We know that the sample also has some other property (call it *Z*) that has some (unspecified) lawful connection with properties *wxy*. We don't think there are significant differences that would make the target item different in the relevant respects. So it stands to reason that the target will have property *Z* too.

Regarding fiction, we might reason by analogy, to come to know things about the world in ways like the following: Jumpa Lahiri's book *The Lowlands* contains a detailed fictional account of radicalism in India in the seventies. When one compares the activities happening in Lahiri's account with the historical record, they seem to match remarkably well. So we have good reason to think that the sample, Lahiri's fictional Indian political radicals, and the target, actual Indian political dissidents, are very similar in many ways. Straightforward analogical reasoning gives us reasons to believe they must be similar in other ways. Now Lahiri's book implies it was not unusual for Indian police sometimes to summarily execute political dissidents. Given that the reality and the fiction were similar in other ways, it stands to reason that they are similar in this way too. Perhaps we should thereby conclude that police did sometimes perform summary execution of political dissidents in the seventies.

Reasoning by analogy this way has many nice features. Unlike with extra-textual confirmation, the *fiction itself* gives one a fair amount of the evidence that the claim about the real world is true. The more the work of fiction gives us evidence that it is like the real world we want to know about, the more we will be justified in thinking that other features that are there in the fictional world will also be there in the real world. Unlike with necessary truths, we are not limited to learning about features that are true in all possible worlds. Analogies are able to tell about things that are true in *highly particular* situations—here, in highly particular fictional situations and similar real-world situations. The features might be very concrete, like what people tended to eat for breakfast in particular regions and times, or they can be highly abstract, like the type of courage certain sorts of people display.

I suspect that people pick up ideas from fiction by reasoning by analogy this way all the time. They see things happen in a fictional sample. They see that a realm of the real world closely resembles the fictional sample in many features, and so they assume that the real-world situation will resemble it in not-yet-seen features as well. I, myself, seem to have been reasoning by analogy this way when campus police came to my office to investigate a bicycle that had been stolen from there. As they went to leave, I asked if they were going to dust for fingerprints. I likely asked this because this is how I'd seen burglaries investigated on fictional TV crime shows.

But a problem with reasoning this way is that the conditions required for a good inductive conclusion by analogy are very hard to meet (especially for the kind of analogies we are interested in here). In my burglary situation, the police were somewhat amused at my inquiry about fingerprints, because, unlike the investigations in analogous fictional burglaries I had in mind, these real police rarely did

fingerprint dusting.²² One of the problems is that between almost any two mildly complex situations there are numerous differences as well as similarities. Differences make any analogy weaker. The number of differences existing between even any two *actual* situations is so great that Alex Rosenberg,²³ in an interesting book chapter titled “History Debunked,” argues that almost no historical situation can really give us much information about what to expect from analogous situations. And the differences between any given real situation and an analogous *fictional* situation that one is reasoning from are likely to be at least equally great. Indeed, there are many reasons they are likely to be greater. The people and situations described in fiction have to be more striking, interesting, and dramatic than real ones for fiction to hold our interest. And even if authors are *trying hard* to create realistic worlds, having an accurate detailed comprehensive knowledge of the real world is no easy matter. And fiction writers are not required by anyone to have the level of knowledge and expertise about any particular realm that social and natural scientists do. So Conan Doyle could write a Sherlock Holmes story in which the villain was a snake trained to follow whistles—without having to check if this is possible. (It isn’t.) George Bernard Shaw could write *Pygmalion* assuming a psychological fact about opposites attracting. (They don’t—ask the researchers at dating site companies who make matches for a living.) A novelist who creates a gripping story about a paranoid schizophrenic might or might not know things that a typical psychologist treating schizophrenics does. A reader of this novel who notices that many features of the main character match those of the brother of a friend could easily come to all kinds of faulty conclusions about how the brother is thinking, by assuming that he’s likely to think like the analogous character in the novel.

There are many *prima facie* reasons, then, to worry about whether fictional situations that seem similar on the surface to real ones really have enough similarities and few enough differences to warrant believing certain things about the real world because of what is there in the fiction.²⁴ These worries could be overcome, however, if we could look carefully at the two cases to see that, unlike what is often the case, they are actually *very* similar. The similarities between the situation in the fiction and the complementary situation in real life might be so numerous that we have good reasons to expect that other features that we see in the fiction will be there in real life as well.

But this is something that’s actually very difficult to know. First, it’s often hard to know what exactly is and isn’t there in the fictional world (including possible differences from real-world situations) on the basis of the verbal descriptions provided by authors. More problematic, however, is that we seldom know enough about the *real-world counterpart* to the fictional situation to know how similar it is, or whether it contains numerous dissimilarities. Are Lahiri’s fictional Indian clandestine radicals really like the ones existing then? To know, I would have to know a lot about India in the seventies. Are Brooks’s Wampanoag really like the historical ones of the seventeenth century? If the fictional ones took hallucinogens, does that mean the real ones did? I discussed Brook’s case earlier in the context of extra-textual confirmation, because nothing in my background knowledge could tell me whether the fictionalized Wampanoag

were similar enough to the real Wampanoag for me to be able to use this similarity as a basis for a warranted *analogical* inference. I did not know enough about the actual Wampanoag. This is the typical situation we find ourselves in when reading fiction. We might sometimes read about situations we already do know a lot about. But the more we know, the less new there is for the fiction to tell us about. We typically learn when we get information about *exotic* realms that we know relatively little about. But the more exotic the world we are reading about is, the less we know about its real-world counterpart, and the harder it is to tell if the two are really highly analogous.

Seeing that various scenarios in fictional worlds seem very like those in the actual world invites us to reason analogically. Schoolyard bullying in a work of fiction might seem so similar to incidents of playground bullying going on in a nearby playground that we naturally come to think that the two situations must be alike in additional ways. If the bully in the fiction bullies out of a deep sense of insecurity, then it is easy to think this is true of today's typical schoolyard bullies (even though lots of evidence indicates that it is not). Analogical reasoning comes naturally to us, and when the conditions for a good analogy are met, we can use such reasoning to make good justified inferences. It is possible for fiction to teach us things through analogical reasoning. But if we want to come to good conclusions this way, we need to be cognizant of how hard it is for comparing real and fictional worlds to meet these conditions.

4.3 Author Trustworthiness

A different way of knowing that some things in a fictional world are really like certain things in our world might help us to get around problems with analogical reasoning. If something could *directly tell us* that certain characteristics of fictional scenarios are very much like their counterparts in the real world, then learning about those characteristics in the fiction would automatically give us justified beliefs about similar characteristics in the actual world. One thing that could do this is having some knowledge of *authorial trustworthiness*. If and when writers of fiction could be trusted to create fictional worlds that were quite like the actual world in certain respects, then we could reliably infer things about the actual world when we saw them in the fiction world.

There are different ways that authors, intentionally or by happenstance, could make sure that their fictional worlds resemble the real world. It is common practice in historical fiction, for the notable historic events taking place in the background world of the fiction to be close to those of the real world. Most authors in this genre usually respect this convention. And certain authors can come to have particularly good reputations for accuracy. Many movie and television producers are also well known to try to be scrupulously accurate regarding their depictions of the habits, customs, and fashion styles of a particular era. *Mad Men* director Matt Weiner was so obsessed with accuracy that it became a competition among viewers to see if they could find any historical inaccuracies in a *Mad Men* episode's choice of stove tops or suit coat buttons.

Even where none of the particular events are the same, outside of science fiction, it's a general convention among most writers that the *physical laws* of the fictional world are the same as those of the real world. So if we come to be aware of a law-like pattern in the world of the fiction, odds are good that this is a law in the real world as well. There also tends to be pressure to have the social and psychological laws governing the behavior of the characters in the fiction to be very similar to those governing people's behavior in the real world—or the fiction will not be believable.²⁵ Where we can be relatively sure that the authors are creating activities governed by the psychological regularities of the real world, then, as we become aware of why the fictional characters do what they do, we could become aware of the rules that make real people do what they do.²⁶

Our ability to really know about the world through authorial trustworthiness will, of course, be a matter of degree in a number of ways. First, the amount of similarity between the fictional world's particular events, general customs, and physical and psychological laws and those of the real world can range from identical to only a little alike. The more alike they are, the more knowing about the fictional world tells us about actual world. And a writer's disposition to depict a world that is like the actual one can run the full gamut. A writer can be utterly obsessed with historical or psychological accuracy or totally indifferent. Even an obsessed writer might make him or herself into a leading expert in an area or prove to be a poor student of it. A writer who is strikingly accurate in one realm might be utterly indifferent to accuracy in another.²⁷ How accurate a writer tends to be in making sure her depictions of a fictional world match aspects of the actual world will partially determine how well we can learn about the world from the fiction. Our degree of knowledge also depends on the degree of *our own knowledge about* how accurate various writers tend to be. We might have a rough idea of the genre conventions of historical fiction, which tells us something about the likelihood of a civil war battle having gone the way it did in *The Killer Angels*. Or we might have a detailed knowledge of what liberties a particular writer like Richard Price does and doesn't feel he can take regarding the typical home lives of police officers. An author's degree of trustworthiness grows as we ourselves come to have greater reasons to trust his created fictional worlds are like our real ones. The greater the degree of authorial trustworthiness, the greater our knowledge of real events based on fictional ones can be.

Having a high degree of authorial trustworthiness in a certain era then, seems to be an important way of learning about the world from fiction. But here, too, it's important to realize the limitations and obstacles we face in trying to learn about the world this way. We need to start by remembering that gauging authorial trustworthiness requires that we acquire a *specific type of knowledge about writers*. Most fiction writers assume that readers will fill in details that they don't specify by assuming that the fictional world is like our world, unless the writer specifies otherwise (Kendall Walton²⁸ calls this the *reality principle*). But the myriad of realistic information that readers bring to the table are things they *already know or believe* about the world, so the fiction isn't teaching them this information. And the things that the author talks about specifically are often the things that are *not* just assumed to be realistic and might or might not be fanciful inventions of the

author. To know whether what the author *specifically* describes is like our world or not, we need to specifically find out whether or not the author tends to be accurate about this or that particular realm. This isn't information most people carry around in their heads, however. Most casual readers or watchers of fiction, then, will possess very little information about authorial trustworthiness that they can use to make sure they can learn about the world from fiction.

But such information, with some effort, might be acquired. Perhaps the quickest and easiest way to learn about whether an author can be trusted in a certain realm is to get information about *general genre conventions*. If we learn that most of the authors writing historical fiction tend to be completely accurate about, say, geographic information, then odds are good that any given writer of historical fiction in a given book will be accurate in their geography.²⁹ If the *physical* laws of fictional worlds (outside of science fiction) are mostly the same as ours, then what we find out about physical laws of a fictional world probably applies to our physical laws, too. We need to keep in mind, however, that this knowledge will be limited in the way that all statistical aggregate knowledge is. Perhaps *most* historical fiction authors will not alter geographic details. But, of course, any given writer or director can choose to disregard genre conventions in a given area for the sake of a better story (though it certainly upset *me* to see Kevin Costner in *No Way Out* run down the C and O Canal in central Washington D.C. where it doesn't run). Knowing that most authors in a particular realm tend to be accurate makes it only somewhat reasonable to infer that a particular one is.

Much more useful would be to acquire *specific* information about the dispositions of specific authors to be accurate in specific realms. Getting this information, however, is much more difficult than learning about general genre conventions. One good way of getting this information, for example, would be to look at the track record of an author's descriptions of certain sorts of specific and general features of a fictional world and compare them with their closest counterparts in the real world to see how similar they are. Meanwhile, figuring out how true-to-life various authors make their fictional worlds by comparing them to the actual world, also requires independently learning a lot about areas of the *real* world. And even if we acquire this information, we must also remember that knowing an author has, in the past, been very accurate about a realm doesn't mean we can count on her to continue to be. An author can always change his habits and needn't feel bound by past practice. William Kennedy, for example, is the Pulitzer Prize-winning author of a series of books about fictional events in Albany, New York over the past several centuries. As a frequent visitor to Albany, I've found these books to be a marvelous guide to the history and geography of the region. I've often lugged the books around to help figure out which historical marvel once stood where a strip mall now stands. I've worked hard to corroborate what's in the book with historical sources. I once had the pleasure of meeting Mr. Kennedy at a book signing and at the question and answer session I asked him what he thought it was and wasn't permissible to alter in fiction. "In fiction," he told me, "you can change anything."

If we do have good information that a given author can be trusted to create a fictional world that is much like the real world, this enables us to infer things about

the real world from acquaintances with that fictional one. We see, however, that it is difficult to get good information about this. Still, we could get some degree of rough statistical knowledge of authorial trustworthiness, if we work at getting it. This could enable us to learn some measure of the facts about the real social world from fictional ones.

5. CONCLUDING REMARKS

The purpose of this paper was to carefully examine whether and how we can learn about the social world from fiction, rather than just assume we do (or don't). I've tried to make clear exactly how works of fiction could give us justified true beliefs about the actual world. We can learn about the real world from a fictional one if we have good reasons for thinking that the fictional world strongly resembles the real world. This is possible in a number of ways. We can have good reasons for thinking there is a resemblance if the fiction has shown us necessary truths. But it is difficult to know that it has. Knowing that *authors* have created a fictional world that is similar to ours gives us another reason to believe our world is similar. But it is hard to know if they have. If we know, through direct comparison, that our world and the fictional world are similar in some ways, we can analogically infer that they are likely similar in other ways. But it is difficult to know if the two worlds are sufficiently similar. If readers want to feel like they have learned something from the fiction via the resemblance shortcut, they need to make sure they do the extra work needed to see if such similarities are there.

We can also learn about the real world through fiction if the fiction leads us to form certain suppositions about what the world is like, and then we do the extra work to gather evidence that such suppositions are indeed true of the actual world. There is no reason to think of such work as difficult. But it is usually work we neglect to undertake.

Since learning about the world from fiction will usually require our doing much more than just reading about the world, we are faced with a tough choice. On the one hand, we can "double down" and invest more resources, trying to make sure that when we study fiction, we also take the extra time to gather the needed additional information that enables us to transform the new beliefs we pick up from fiction into justified knowledge. On the other hand, the time it takes to collect the additional information required to turn the fiction into knowledge could be used to get much more information about the real world by reading about it directly. And all of the time spent reading the fiction in the first place is also time *taken away from* what could have been used for learning new additional things, or for re-looking at the same things. The simple fact that any time spent in any reading or other information gathering activity is time taken away from other sorts of reading or information gathering also means that we *can't* be content to simply say "we can do both" when we are thinking about whether more can be learned through fiction or non-fiction. If, in the time it takes to learn something through fiction and supplementary research, one could learn 20 comparable things through non-fictional information gathering, then trying to "do both" will likely entail an enormous loss of knowledge.

If studying fiction to learn about the real social world turns out to be a very inefficient way to learn about it, it may be more useful to learn about it through the social sciences, and use the study of fiction for different purposes. Few think, after all, that the best way to teach students science would be to focus on science fiction. Perhaps the study of fiction should concentrate more on teaching people the craft of writing, giving us ideas about how the world *could be*, or inspiring people to great deeds. Perhaps we could study fiction to learn more about how certain types of writing causally effects individuals or historical events. Whether the difficulty of really learning facts from fiction should cause literature scholars and teachers to work hard to gather the extra information necessary for enabling literature to teach us about the social world, or whether it means that the focus of the study of literature should be elsewhere should be a *central debate* scholars should be having. Whatever we decide, what we clearly should *not* do is just assume that Moliere and Stendhal, by themselves, will really teach us more about human thought and character than social science ever will, however common this assumption may be among English teachers or other segments of the educated public.

ENDNOTES

1. See quote in Philip Kitcher, *Science in a Democratic Society* (Amherst, NY: Prometheus Books, 2011), 250.
2. George Monteiro, *Stephen Crane's Blue Badge of Courage* (Baton Rouge: Louisiana State University Press, 2000), 86.
3. Gregorie Currie, "Does Great Literature Make Us Better?" *New York Times*, June 1, 2013.
4. See, for example, Berys Gaut, *Art, Emotion and Ethics* (Oxford University Press, 2007). See also Gilbert Plumer, "The Transcendental Argument of the Novel," *Journal of the American Philosophical Association* 3.2 (2017): 148–67.
5. Martha Nussbaum, *Love's Knowledge* (Oxford University Press, 1992).
6. Frank Palmer, *Literature and Moral Understanding* (Oxford: Clarendon Press, 1992).
7. John Gibson, *Fiction and the Weave of Life* (Oxford University Press, 2012).
8. The issue I am looking at here is how people can learn about the world because they come to have true beliefs about our world *based on seeing similar things depicted in a fictional world*. There are, of course, other ways fiction could lead people to have true beliefs about the world. Perhaps only certain kind of historical circumstances and/or certain kinds of authorial mindsets can produce certain sorts of fiction. And perhaps works of fiction tend to have certain systematic effects on readers' minds and subsequent real-world events. If this is the case, and we can learn how these relationships work, then studying works of fiction will be able to tell us about the likely existence of certain sorts of antecedent or consequent circumstances. I do think these sorts of "causal stream" examination are important ways we can learn about the world from fiction. But this is not the sort of learning I will be focusing on in this paper.
9. Traditionally, knowledge has been considered to consist of justified true belief (and some additional criteria have been proposed in recent years). But exactly what "be-

lief," "justification" and "truth" amount to are contentious notions that epistemologists and metaphysicians have debated for centuries. For our purposes here, we can make do with rough everyday notions. One should note that I am not committed to there being anything like sentences in the head (indeed, I doubt there are such things). Traditionally, however, justification has been thought of a property of belief-sentences. Given that I don't think that we actually have sentences in the head, the way I will be thinking of the issue of "justified beliefs," here, is whether our belief sentences—best thought of as statements describing *states of affairs we think are there*—are warranted to be asserted by the evidence agents have. If the agent isn't warranted in saying "the cat is on the mat," he doesn't have a justified belief that it is. Note that some philosophers have explored the relation between fiction and fine-grained depictions of other mental states that aren't necessarily beliefs. See, for example, Tamar Szabó Gendler, "Alief and Belief," *Journal of Philosophy* 105.10 (2009): 634–63. See also Derek Matravers, *Fiction and Narrative* (Oxford: Oxford University Press, 2015). For our purposes, the notion of belief that I have described should suffice.

10. This is not to say that seeing an example of a known idea depicted in fiction couldn't improve our ability to access it from memory or give us other cognitive improvements. But this is different from the fiction providing us with new *knowledge*.

11. Jerome Stolnitz, "On the Cognitive Triviality of Art," *The British Journal of Aesthetics* 32.3 (1992): 191–200.

12. In this paper I am assuming a traditional "internalist" notion of justification for which it is required that an agent be aware of evidence for the truth of her belief in order for it to count as justified. There are, of course, also "externalist" theories of justification that use a more lenient requirement of the belief being formed by a generally reliable process. I believe that the arguments against externalist theories of justification are persuasive. See, for example, Earl Conee and Richard Feldman, "Internalism Defended," *American Philosophical Quarterly* 38.1 (2001): 1–18. For those not persuaded, I think that many of the arguments I discuss here can be adapted to apply to externalist theories. But to discuss this in detail would require another paper. Whatever one's views about internalism and externalism, the main points I want to make here concern whether beliefs acquired from fiction are justified in the sense of people having good reasons for believing them.

13. When we do see confirming instances of some X doing Y, note that it doesn't tell us anything about *how often* an X does Y. It just tells us that it *sometimes* does. Indeed "confirming instances" is a notorious logical *fallacy*, where one isn't in a position to know about the characteristics of a set of X, because one isn't looking at an unbiased sample (one is only looking at samples of X that do have Y). Reading that Tartars drink mare's milk, and then seeing an example of someone (a Tartar even), and concluding "The Tartars *often* drink mare's milk" would be just this sort of unjustified inductive generalization. And, of course, coming to the conclusion that they often drink it on the basis of a *single* example—and a fictional example at that—would be an even more unjustified inductive generalization. Yet many enthusiasts about learning from fiction seem to repeatedly reason in just this manner.

14. Of course, reading fiction might still lead to your being successful in the world in other ways—say making you love or hate certain kinds of people or situations. But this is something different than giving you justified *knowledge* about what the world is like.

15. Judith Jarvis Thomson, "Killing, Letting Die, and the Trolley Problem," *The Monist* 59.2 (1976): 204–17.

16. Stephen Stich and Wesley Buckwalter, "Gender and the Philosophy Club," *The Philosophy Magazine* 52 (2011): 60–5.

17. Krist Vaesen, Martin Peterson, and Bart Van Bezooijen, "The Reliability of Armchair Intuitions," *Metaphilosophy* 44.5 (2013): 559–78.

18. See Walter Sinnott-Armstrong, "Framing Moral Intuitions," in *Moral Psychology*, Vol. 2, *The Cognitive Science of Morality: Intuition and Diversity*, ed. Walter Sinnott-Armstrong (Cambridge, MA: MIT Press). See also Joshua Knobe and Richard Samuels, "Thinking Like a Scientist: Innateness as a Case Study," *Cognition* 126.1 (2013): 72–86.

19. Simone Schnall, Jonathan Haidt, Gerald L. Clore, and Alexander H. Jordan, "Disgust as Embodied Moral Judgment," *Personality and Social Psychology Bulletin* 34.8 (2008): 1069–109.

20. Jonathan M. Weinberg, Joshua Alexander, Chad Gonnerman, and Shane Reuter, "Restrictionism and Reflection," *The Monist* 95 (2012): 200–22.

21. Noël Carroll, "The Wheel of Virtue: Art, Literature, and Moral Knowledge," *The Journal of Aesthetics and Art Criticism* 60.1 (2002): 3–26, see 7.

22. Apparently, this happens quite frequently and has been alleged to cause serious problems for the criminal justice system. On the basis of their consumption of crime fiction, juries apparently often expect that truly guilty people leave behind a great deal of evidence that police can find. In the absence of such overwhelming evidence, juries are sometimes reluctant to return guilty verdicts, even in very strong cases. This phenomenon has come to be known as the "CSI Effect," after the popular TV show. See Evan W. Durnal, "Crime Scene Investigation (As Seen on TV)," *Forensic Science International* 199.1–3 (2010): 1–5.

23. Alex Rosenberg, "History Debunked," in *The Atheist's Guide to Reality* (New York, Norton, 2011), 242–74.

24. See Onora O'Neil, "The Power of Example," *Philosophy* 61, no. 235 (1986): 5–29.

25. See Gilbert Plumer, "The Transcendental Argument of the Novel," *Journal of the American Philosophical Association* 3.2 (2017): 148–67.

26. For externalist epistemologists, we don't have to have evidence of an author's trustworthiness to gain knowledge of the things they write about. They just need to be reliable purveyors of such information whether or not *we* know that they are. See, for example, Stacy Friend, "Believing in Stories," in *Aesthetics and the Sciences of Mind*, edited by Greg Currie, Matthew Kieran, Aaron Meskin, and Jon Robson (Oxford University Press, 2014): 227–48. Friend develops and defends an interesting version of such an externalist view.

27. See Friend, "Believing in Stories."

28. Kendall L. Walton, *Mimesis as Make-believe* (Cambridge, MA: Harvard University Press, 1990).

29. See Friend, "Believing in Stories."