Fictional Exemplars

In this paper I will argue that insofar as one holds that moral properties are defined by direct reference to exemplars (DRE), one should also hold that fictional characters can be moral exemplars. Since Linda Zagzebski's exemplarist virtue theory is the main advocate of DRE in the literature, I will focus on it. I will first pay particular attention to Zagzebski's exemplarist virtue theory to argue that there are practical benefits to holding that fictional characters can be moral exemplars and that these benefits outweigh the practical benefits of holding that exemplars have to be actual. Second, I will give a view of fictional characters and discovery of moral properties that allows fictional characters to function as moral exemplars. Last, I will raise what I take to be the most pressing problems for the view that DRE is compatible with fictional exemplars, and I'll propose solutions that a proponent of DRE can accept at little to no cost. I will conclude that insofar as a view endorses DRE, as Zagzebski's exemplarist virtue theory does, there are good reasons for proponents of that view to hold that fictional characters can be moral exemplars.

1   Practical benefits[1] of fictional exemplars

It seems natural to hold that fictional characters can be exemplars, since we often take fictional characters to be exemplars, and we use them that way in examples. Often it makes no practical difference whether we use fictional characters or biographies of real people. For example, in *Divine Motivation Theory*, Zagzebski uses the example of Ajax as a moral exemplar for the ancient greeks (Zagzebski 2004, 146-147). Whether or not Ajax actually existed seems to make no difference to whether the characteristics of Ajax helped ancient greeks recognize important moral properties. For another example, in her *Epistemic Authority*, Zagzebski uses a biographical account of Captain Oates, who, according to the accounts, displayed great courage

---

[1] There may be a metaphysical benefit, e.g. grounding moral obligations, that may privilege an actual exemplar theory over one that allows fictional exemplars. I don't discuss this benefit here since my primary targets are those views that hold that goodness is grounded in moral properties.

(Zagzebski mss, 23-24). She adds a footnote indicating that these accounts of Captain Oates are disputed. Suppose the disputers are correct and that there was no Captain Oates, or, if there was, he didn't perform the courageous act the accounts say he performed. The accounts would then be historical fiction, but this would not take away from our ability to read the account and recognize an act as courageous. Zagzebski says this best in her (2010): "We learn through narratives of both fictional and nonfictional persons that some people are admirable and worth imitating, and the identification of these persons is one of the pretheoretical aspects of our moral practices that theory must explain." (Zagzebski 2010, 51)

There are other practical benefits to holding fictional characters to be exemplars. If only actual people were moral exemplars, we wouldn't have access to nearly as many moral exemplars as we would if fictional characters were also exemplars. Through fiction, we will be able to see many characters with the same moral property we see in actual moral exemplars display that moral property many more times than we would if only actual people were moral exemplars. Seeing multiple instances of a property will help us to more clearly recognize that property, which is a practical benefit.

Further, fictional exemplars can realize their virtue or moral goodness in environments and settings in which we don't get to see actual exemplars display their goodness. Real people tend to act in certain characteristic ways, for example within the traditions they've received from their culture or time or location. Fictional characters act in settings we haven't been acquainted with before. We can see how fictional characters would act in a variety of unfamiliar contexts, which will enrich our understanding of what it is to have certain virtues or certain moral properties.[2]

Also, fictional characters can represent a specific virtue well without the distractions from that virtue that come from the complexity of being an actual agent. Fictional characters' virtues can, in a way, be isolated from other features that may obscure what is essential to a particular virtue so that we might more clearly recognize a virtue in a fictional exemplar than we would the same virtue in an actual exemplar.[3]

Last, we have access to a fictional character's mental life: thoughts, intentions, motives, etc., merely by virtue of the author or producer telling us about the character's mental life. Actual people may be wrong about their intentions or not have access to them. On the other hand, the narrative report of a fictional character's

---

[2] I give and reply to the problem (problem 4) that some of these contexts are too unrealistic in section 3 below.

[3] A reviewer suggested that Charles Dickens has Esther Summerson display generosity this way in *Bleak House*.

mental life is accurate by default. Here it is precisely the fictionality of the fictional character that helps the observer to recognize certain moral properties in an exemplar.[4]

In fact, I don't see *any* practical benefit that actual exemplar theories have over fictional exemplar theories. Many exemplars we cite are no longer living, so they are no longer actual. But someone's death shouldn't prevent us from being able to recognize important moral traits that person had when she was alive. Further, we don't have communicative access to many living exemplars we cite, so accounts we hear or read of them are practically no different than a realistic novel, with this possible exception: believing that an actual exemplar has a certain characteristic may be more motivating to us, because we may think that having that characteristic is more accessible to us than if we were merely to believe that a fictional exemplar has that characteristic. In addition to this exception, it seems that the only other cases where actual exemplars are more practically beneficial than fictional exemplars is when the exemplar is living and we communicate with her. That way, we can get answers to questions we have of her, whereas fictional characters don't answer our questions.[5] But the fictional exemplarist, since she believes there can be actual exemplars too, also has these benefits.

This brief sketch will suffice for my account of the practical benefits to holding that there can be fictional exemplars. What remains is to show that the disutility of the problems for fictional exemplarism doesn't outweigh the benefits to holding to fictional exemplarism. I'll address these problems in section three, but in order to do that, I will develop an account in the next section of what fictional characters are and how we define moral properties these fictional characters possess.

2 Fictional characters and discovery of moral properties

In this section, I'll first say which views of the metaphysics of fictional characters are compatible with the thesis of this paper, then I'll give an account of how we define moral properties these fictional characters possess that is compatible with a direct reference view.

Fictional exemplarism as I'm defending it is compatible with any non-eliminativist

---

[4]Some fictional characters are in movies but not in books. For simplicity, however, in the sequel I'll be focusing only on those fictional characters that are in books.

[5]Technically they could, should the author choose to have the fictional character enter into dialogue with real people, e.g. in an online chat session.

view of the metaphysics of fictional characters. We can hold that fictional characters are 1) just bundles of universals selected by an author (Wolterstorff 1980), 2) possible entities described by an author (Kripke 1963/1971, 65), or 3) artifacts created by an author (Thomasson 2004, 2009).[6] There are arguments for and against each of these views that I won't go into here. As long as we grant non-eliminativism about fictional characters, there are three more features of fictional characters that are important for my view.

First, fictional characters must in some sense have properties that can be discovered by those who learn about them.[7] If fictional characters did not have properties, then there would be no properties of fictional exemplars to be discovered, and if there were no properties of fictional exemplars to be discovered, DRE would be false. Fortunately, each of the three views of the metaphysics of fictional characters is a view in which fictional characters have discoverable properties.

Second, fictional characters cannot be identical to whatever properties she in fact has. If a fictional character is identical to whatever properties she in fact has, then the character cannot gain or lose any properties.[8] But in order for us to come to recognize moral properties in an exemplar, that exemplar needs to be able to gain or lose *some* properties or other. In fact, for many moral properties, a character cannot possess that moral property unless it is possible for her to gain or lose some property or other. Virtues in their fullest sense are displayed in actions, and someone's acting requires gaining or losing some property or other.

Fictional characters aren't identical to the properties they have if they have both essential and non-essential properties. I take fictional characters to have both essential and non-essential properties. We can take the distinction as to which properties are essential and which are non-essential to be set by the author, just as if the author created a new "natural" kind in her writings. This raises the question: how can we, as readers, discover which properties are essential and which are non-essential? I think we can do this in one of two ways. First, we can simply ask the author. This would be akin to asking an expert about the chemical structure of a natural kind. Second, we can try to discover the essential properties of fictional characters just like we try to discover the essential properties of natural kinds: we observe features of the natural kinds and investigate. Investigating may be difficult with a fictional entity,

---

[6]Artifactual views have also been defended by Salmon (1998) and Schiffer (1996). I am most persuaded by artifactual views of fictional characters.

[7]'Have' here can mean 'be constituted by' (as in the abstract object view) or 'instantiate' (as in possibilist and artifactual views).

[8]Thomasson (2009), 16, uses this as an objection to the abstract object and possibilist views (numbers 1 and 2 above, respectively). If her critique is successful, we can adopt the artifactual view (number 3 above), which is the majority view anyway.

since we only have access to limited information about that entity. But suppose we only had limited access to some natural kind due to, say, a defect in our abilities or the extinction of that natural kind or God's prevention. This would prevent us from being able to discover what the essential properties are, but it would not prevent there from being essential properties. Since my thesis doesn't require that we be able to tell whether *goodness* is an essential property of an exemplar or not, either of these two options will suffice for my purposes. I'll address further objections in the next section.

Third, fictional characters have to be able to have *moral* properties. This doesn't seem difficult to hold. If they have non-moral properties, I see no objection to holding that they have moral properties. After all, we do say, for example, that Samwise Gamges is brave, and Saruman the White is evil.

Now I'll discuss how we discover moral properties in exemplars. As is stipulated by DRE, moral properties are defined by direct reference, or ostensively.[9] For example, I point to an exemplar and say, 'good' is whatever is like that, and whatever other character is like that character is also good. Proponents of DRE hold that the direct reference account has features important for a moral theory.[10] Let's call these "DRE's benefits." First, DRE allows us to successfully refer without having any beliefs about a description that a term (such as 'good') might pick out. Specific to Zagzebski's exemplarist virtue theory, our emotion of admiration is directed at an exemplar construed in a certain way, even if we can't explicate the aspect of the exemplar to which our admiration is directed.[11] What is important to this first point is that we recognize that exemplars are good first without being able to provide an account of what 'good' means. Second, it allows us to successfully refer even when we associate a false description with the term. We do make mistakes, but we still ought to be able to successfully refer to and discover moral properties (such as goodness) by association with exemplars despite these mistakes. Third, it allows for narratives to help reveal to us the essential properties of a good person. By direct reference, we can identify characters in a narrative and by pointing to these characters and saying, for example, "She is generous," successfully refer to moral properties.

I need to say something specifically about defining *moral* properties by ostension, since the moral theories I am discussing endorse a direct reference view. According

---

[9]The direct reference view is defended, among other places, by Kripke (1972) for proper names and Putnam (1973) for natural kind terms. For a good summary of the view, see Reimer (2010), 13-15.

[10]The following features follow Zagzebski (2010), 51-52.

[11]Hare (2005): "[S]ince she also thinks that emotions introduce propositional contexts, which are therefore in the technical sense 'opaque', she would be better off saying that the intentional object is the combination, the thing only *qua* held under the 'thick' description as dangerous."

to the Kripke-Putnam model of direct reference to natural kinds, I fix the referent of a natural kind term to something that has an internal structure that all instances of that natural kind essentially have. For example, suppose I somehow point to a colorless, odorless liquid and call it 'water'. By doing so, I'm fixing the term 'water' to whatever essentially has the same internal structure as the stuff I'm pointing to essentially has—H_2O. If my thesis is correct, we need to be able to define terms like 'good' and 'good person' by ostension. The problem is that 'good' and 'good person' do not seem to refer to internal structures. There are two ways I think we can reconcile my thesis with the Kripke-Putnam model. The first is to take 'good' and 'good person' to refer to someone with a certain kind of internal structure, namely a motivational or psychological structure. Just as H_2O is the internal structure water essentially has, so a certain kind of psychological structure is what a good person essentially has. The second way avoids taking the essential feature of a good person to be an internal structure. To accommodate this view, we need to extend the Kripke-Putnam model so that we can fix a referent to something whose essential feature is not an internal structure. Here's one way to do this. First, take internal structures to be part of a description. For example, take 'being H_2O' to be a description. Second, hold that descriptions about essential features of a referent don't need to be about internal structures. Descriptions about structures work well for natural kinds, but descriptions are not about structures when the descriptions are about artificial kinds like fictional characters. This should allow us to define 'good' by reference to a fictional exemplar. Just as there can't be tigers without the internal structure of tigers, so there can't be an instance of 'good' without the descriptive features that the term 'good' refers to, even if the exemplar is fictional. Third, hold that moral properties do not need to be essential properties. (This is, in fact, what Zagzebski affirms in Zagzebski (2010)).[12] That is, unlike structures for natural kind terms, moral properties do not rigidly designate fictional characters. For example, if the chemical structure of water were to alter from being H_2O, the chemical would no longer be the same stuff it once was, but if the fictional character were to become vicious, the character would still be the same character she once was.

These modifications still leave us with the essential features of the Kripke-Putnam direct reference model. We're still able to define by ostension and maintain DRE's benefits given above. If the three conditions in the above paragraph hold, then we should be able to ostensively define 'good' and 'good person' by direct reference to fictional characters.

---

[12] "This [direct reference] theory is compatible with the possibility that paradigmatically good individuals are only contingently good. . . unlike water, which is probably water essentially." Zagzebski (2010), 52 & ftnt 7.

Now that I've given a sketch of what fictional characters can be and how we discover the moral properties these fictional characters possess, I'm in a position to address what I take to be the most pressing problems for the view that DRE is compatible with fictional exemplars. I'll try to show that the problems this view can be solved at little to no cost.

3   Problems and solutions

There are four problems I'll raise for the view that DRE is compatible with fictional exemplars. The first two problems concern defining moral properties like 'good' by direct reference to fictional characters. The first problem raises a metaphysical problem, and the second raises an epistemological problem, for trying to fix the reference of 'good' to properties of fictional characters. The next two problems concern what kind of investigative work we can do to find what deep moral properties fictional exemplars have. The third problem concerns whether we can do any investigation at all, and the last problem concerns whether our investigations would yield correct results.

The first problem is one Kripke raises for direct reference to fictional entities. The information we're given in writings of fiction underdetermines the reference of the term we want to apply. Kripke offers an example of a unicorn. "If we suppose, as I do, that the unicorns of the myth were supposed to be a particular species, but that the myth provides insufficient information about their internal structure to determine a unique species, then there is no actual or possible species of which we can say that it would have been the species of unicorns." (Kripke 1972, 157) The problem here is that when we take a term that is supposed to refer to a feature but that feature isn't determined, then there is no referring that can occur. 'Unicorn' is like that. So, 'unicorn' doesn't refer. Similarly, suppose an author writes about what she calls a non-actual substance named 'schmuranium'. To find out if 'schmuranium' refers, 'schmuranium' needs to have a chemical constitution. There may be no fact of the matter about what schmuranium's chemical constitution is. In that case, 'schmuranium' doesn't refer.

Does this problem apply to terms like 'good' or 'good person'? Perhaps it does. Perhaps 'good' or 'good person' is supposed to refer to a group of character traits that fictional characters don't have in extant writings about them. Perhaps some of these character traits are counterfactual: if a certain fictional character were in different circumstances, that character would do such-and-such. We're simply not given this information. The only information we're given about fictional characters is what the author has conveyed in her writings. What's on the surface is all there

is. So, there may be no fact of the matter whether a certain fictional character has deep moral properties. For example, there may be no fact of the matter whether a certain fictional character has certain character traits we're not told about by the author, or there may be no fact about what she would do in circumstances other than the ones the author has her in. The problem is not simply that we don't know what (counterfactual) character traits the character has; the problem is that there aren't any. If there aren't any, then 'good' when attributed to a fictional exemplar simply doesn't refer.

In reply, I think there *is* a fact of the matter about what a certain fictional character would do in counterfactual circumstances, and there is a fact of the matter about what non-revealed character traits that fictional character has. These facts are determined by what the author would write in counterfactual circumstances. Propositions of the form

> If fictional character F were in circumstances C, then F would perform A in C

are reducible to propositions of the form

> If author S of F were to have written about F in C, then S would have written that F performs A in C.[13]

This account doesn't require that the author has even thought about what a character may do in different circumstances. As long as someone thinks there are true counterfactuals about real people (and I do take authors to be real people), then there are true counterfactuals about fictional characters. So, just as we can look at real people and fix the reference of 'good' to their (counterfactual) properties, so we can fix the reference of 'good' to the (counterfactual) properties of fictional characters.

---

[13]Some comments: 1. I've omitted provisions for non-written literature for simplicity. 2. I'm assuming the worlds in which the author would have written about F in C are close to this one so that, for example, the author does not become crazy from this world to the next. 3. On the abstract object and possibilist views, different authors can write about the same character. In that case, we can modify the account so that for 'author S' we substitute 'the author(s)'. If we hold that two authors can write simultaneously about the same fictional character and that each author would have the character perform incompatible actions in C, then either a) there are two characters when we thought there was one or b) moral properties can only be had by fictional characters that do not have multiple authors such that each author would have the character perform incompatible actions in C. Adopting Thomasson's artifactual view will avoid the third complication in this comment. See Thomasson (2004).

All I need to overcome the objection is to hold that there is some fictional exemplar who is such that if she were in some close circumstances, she would do A, and there is no other fictional exemplar who would do not-A in those close circumstances.[14] It may be that for almost all fictional characters there is no fact of the matter about what that character's author would have said she does in certain remote circumstances. After all, exemplars only set a standard for behavior in the range of circumstances close to those she actually encounters. Some circumstances are so far from an exemplar's experience that the exemplar cannot be expected to set a standard in those circumstances. Similarly, there may be many circumstances in which an author would not have anything to say about what a fictional character would do. That is compatible with my account. We should not expect that every exemplar sets a standard in every circumstance she could encounter.

The second objection is epistemological. (Kripke 1972, 157) Suppose there's a story with a full description of gold but without a description of the internal structure of gold. That is, we're told there is a shiny yellow metal found in mountains, it has great financial value, pawn shops everywhere are trying to get you to trade it in, etc. Suppose also that something similar to the reduction in my reply in the last paragraph works and that if the author were to have written about the chemical structure of this element, she would have written that it has the atomic number 79, or whatever else the internal structure is. But she in fact doesn't write about this. How are we justified in believing that what she's referring to is gold? She could be referring to what is, in fact, fool's gold. Fool's gold, presumably for our purposes at least, has all the same descriptive features of gold except for its chemical structure.

Similarly, if we were to fix the referent of 'good' by pointing to a fictional character, even if that character does have all the properties to which 'good' refers, we wouldn't be able to justifiably believe that we were successfully referring. We could, after all, be referring to someone whose character appears to us to be good but who, because of her other hidden character traits, is not actually good. This has bad consequences. We want an exemplar theory that can serve as some sort of a guide for behavior (e.g. Zagzebski 2010, 49), and if we can't justifiably believe that something is good when what we see is in fact good, that's a high cost for the theory.

In reply, the same situation occurs with actual exemplars. With some actual exemplars, the only information we're given is what's on the surface. This happens, for example, when we're told about an actual exemplar but can't communicate with that exemplar. In fact, it may be practically impossible to communicate with some actual exemplars. Nevertheless, we still treat them as exemplars and can refer to them as 'good.' Further, it may be possible to investigate further into what deep

---

[14]Thank you to [BLIND] for making this point.

character traits a fictional character has by asking the author what her character would do in other circumstances. This only works if the author is living, but it is an option in some cases.

The third problem is this: since we can just ask the author about the character as long as the author is alive, there is no need to investigate fictional characters at all.[15] There is no investigating to do as long as one person, viz. the author, can simply create by virtue of her imagination what the investigators would find if there were anything an investigation would reveal. This would be bad. We want to be able to investigate to discover moral properties in theories that hold to DRE, especially in Zagzebski's exemplarist virtue theory.[16]

Someone may reply by saying that there may be some traits a fictional character has that surprise that character's author. So, the author doesn't know everything she creates about her character. For example, an author writes about a character who performs certain kinds of actions and, upon reading what she wrote, realizes the character has a certain trait she never before intended or noticed in that character.[17] However, this kind of discovery by the author does not occur for every fictional character. As long as the author has already intended or noticed the morally relevant features of the character, then the initial objection still stands. It seems there is no need to investigate characters whose morally relevant features have been noticed by a living author. Although this objection is forceful, this objection only applies to cases where the fictional character has a living author who has already noticed the morally relevant features of her fictional character.

In reply, in the cases where the objection applies we should treat asking the author about her character as parallel to the case of asking God about the chemical structure of a natural kind. Just as God creates natural kinds with a certain structure and does not fail to notice this structure, so a living author creates characters with morally relevant features and notices these morally relevant features. Would asking God about the chemical structure of a natural kind be worthy of the description 'investigation'? Perhaps it is, but if it's not, that's not so bad. Arguably, asking God about the chemical structure of a natural kind is preferable to investigating natural kinds themselves. It is, after all, less work, and we're more likely to get it right. At any rate, at least there is a parallel creative expert in the case of natural

---

[15][BLIND] offered this objection and the next one to my thesis in personal correspondence, Nov. 22-26, 2011. [BLIND] also offered a version of Kripke's epistemological objection.

[16]Correspondence with [BLIND], Nov. 23, 2011.

[17]This doesn't contradict my reply to the first objection. The author's surprise only shows the author that she did not realize that: If she, the author S of F, were to have written about F in a circumstance in which F's character needed to be revealed explicitly, then S would have explicitly revealed those character traits.

kinds, so the fact that we can ask the author to find all the details we need to find about fictional characters isn't unusual.

A similar reply should work even for those who think there is no God. If there is no God, we can still affirm that if there were a God with such-and-such qualities (communicative, correct, non-deceptive, etc.) who created natural kinds, then we could ask God about the chemical structure of a natural kind. More generally, if there were a creative expert of such-and-such qualities that created some object O, then we could ask that creative expert about O. The general case is applicable to both God (if God exists) and authors.

The fourth problem is that investigation into fictional characters would often yield wrong results, because authors are often wrong about how moral properties or virtues appear when someone possesses them. For example, an author might make a character courageous, but because of the unrealistic fictional context, the character's courage isn't admirable. For another example, an author might make temperance in a character always be accompanied by narrow-mindedness and stodginess so that the temperance isn't admirable in that character.

In reply, I propose we take authors to be experts about what non-moral characteristics their characters display but not always experts about what moral characteristics their characters display. At least, there seems not to be anything initially implausible about letting authors create fictional characters while still holding that they often make false moral assessments about those characters. Authors make moral assessments like everyone else, but they, also like everyone else, are able to create fictional characters and so be experts about what non-explicit non-moral features that character possesses.[18]

To summarize my replies to these objections, what fictional characters would do is determined by what their authors would have them do, and we can discover what these fictional characters would do by asking their author (or, in the case where the author is no longer living, we can't discover what these characters would do. But that's no different than the case of actual exemplars we can't communicate with.) Further, investigation of fictional characters isn't precluded by my account, or if it is, that's not so bad, since we can hold a similar view about natural kinds. And if we can investigate a fictional character by asking the author, we can avoid often

---

[18]This is a general statement. If an author gives contradictory non-explicit non-moral features of a character, she cannot be an expert about those features. For example, an author cannot, in a sequel that begins at the point the prequel ended, say she's writing about the same character but also give that "same" character vices where that "same" character had virtues before. Likewise, she can't make that "same" character a different species. Or, if she did, we would take her to have created a fictional environment with different moral or natural laws so not take the changes resulting from those laws to be bearing on our actual situation.

getting wrong results by taking the author to be an expert about the non-moral and non-explicit features of the character but not necessarily the moral features of her fictional character. I think these are solutions that someone can accept at little or no cost.

Conclusion

I've argued that there are practical benefits to holding that fictional characters can be moral exemplars, and, after offering an account of which views of the metaphysics of fictional characters are compatible with the thesis of this paper and how we can, by direct reference, discover their moral properties, I've argued that the problems for the view that DRE is compatible with fictional exemplars don't succeed. I conclude, then, that insofar as one holds that moral properties are defined by direct reference to exemplars, one should also hold that fictional characters can be moral exemplars.[19]

References

Hare, John. (2005) "Review: Divine Motivation Theory." *Notre Dame Philosophical Reviews.* http://ndpr.nd.edu/news/24314-divine-motivation-theory/

Kripke, Saul. (1972) *Naming and Necessity.* Cambridge: Harvard University Press.

Putnam, Hilary. (1973) "Meaning and Reference," *The Journal of Philosophy* 70:19, 699-711.

Reimer, Marga. (2010) "Reference," *Stanford Encyclopedia of Philosophy.* http://plato.stanford.edu/archives/spr2010/entries/reference/

Salmon, Nathan. (1998) "Nonexistence," *Nous* 32:3, 277-319.

Schiffer, Stephen. (1996) "Language-Created Language-Independent Entities," *Philosophical Topics* 24:1, 149-167.

Thomasson, Amie. (1999) *Fiction and Metaphysics.* Cambridge: Cambridge University Press.

Thomasson, Amie. (2004) "Fictional Characters as Abstract Artifacts," in *Philosophy of Literature*, ed. Eileen John and Dominic McIver Lopes. Blackwell Publishing, 144-153.

Thomasson, Amie. (2009) "Fictional Entities," in *A Companion to Metaphysics*, ed. Jaegwon Kim, Ernest Sosa, and Gary Rosenkrantz. Blackwell Publishing, 10-18.

Wolterstorff, Nicholas. (1980) *Works and Worlds of Art.* Oxford: Clarendon Press.

Zagzebski, Linda. (2004) *Divine Motivation Theory.* Cambridge: Cambridge University Press.

Zagzebski, Linda. (2010) "Exemplarist Virtue Theory," *Metaphilosophy* 41:1-2, 41-57.

Zagzebski, Linda. (mss) "Trust in Emotions," Wilde Lecture and chapter 4 of *Epistemic Authority.*