

**Appropriate Attitudes and the Value Problem**  
(forthcoming in *American Philosophical Quarterly*, January 2006)

*Michael S. Brady*

*University of Glasgow*

The 'value problem' in epistemology is to explain why knowledge is more valuable than true belief.<sup>1</sup> This problem, as a number of philosophers have recently pointed out, is thought to be especially troubling for species of reliabilism.<sup>2</sup> This is because reliabilist theories define knowledge as true belief that is the product of a reliable process or faculty or agent. However, the value of a product does not in general seem to be enhanced by the fact that it comes from a reliable source, and so these theories seem unable to explain why true belief produced by a reliable source is more valuable than merely or accidentally true belief.<sup>3</sup> In support, note that the value of a photograph is not enhanced by the fact that it comes from a reliable camera: if two photographs look equally good, it is irrelevant whether one was produced by an expensive and reliable SLR whilst the other was produced by a cheap and temperamental instamatic. By the same token, the value of a true belief would not seem to be enhanced by the reliability of the process or faculty which produced it: if two beliefs are true, it is irrelevant whether one was produced by a reliable cognitive faculty whilst the other was not. The reason for this is straightforward: the value of reliability *depends upon* the value of the relevant

product, in which case reliability seems incapable of conferring additional value on products.<sup>4</sup>

This paper seeks to defend reliabilism against this objection, by appealing to an axiological principle relating value with appropriate or fitting valuing. Given the truth of this principle, reliabilists can argue that reliably-formed true belief is more valuable than accidentally-formed true belief on the grounds that it is more appropriate or fitting to value the former over the latter. In §1 the principle is introduced, explained, and employed in a defence of reliabilism. In §2 a number of objections to this defence are considered and rejected.

## 1

The task facing the reliabilist is to explain why knowledge is more valuable than true belief. But it seems as though the reliabilist *can* explain this, given a particular way of understanding the relation between the motives a person can have when pursuing true belief, and the value of true belief itself.

According to one prominent and plausible tradition in axiology, to be valuable is to be a fitting or appropriate object of positive evaluative attitudes. On this view, the goodness of x consists in its being a fitting or appropriate object of positive attitudes like love and admiration.<sup>5</sup> So according to this tradition, we are to understand the fact that x is valuable in terms of the fact that it is fitting or appropriate to love (or otherwise value) x; and we are to

understand the fact that x is more valuable than y in terms of the fact that it is fitting or appropriate to love x more than y (or fitting or appropriate to prefer x to y). Let us assume that this way of understanding the relation between values and evaluative attitudes is correct. If so, the value problem can be understood as the problem of explaining why it is more appropriate or fitting to love or value knowledge over mere true belief. But if we understand the value problem in this way, the reliabilist is in a good position to explain why knowledge is more valuable than true belief, since the reliabilist is well placed to explain why it is more appropriate to prefer reliably-formed true belief to accidentally-formed true belief.

In order to see this, we should note first that there are different standards according to which it is appropriate or fitting to love or value something, and in addition a variety of ways in which one can love or value something. To take the first case, it might be appropriate from a purely pragmatic or instrumental standpoint to value knowledge over true belief (perhaps valuing knowledge has evolutionary advantages), but this does not mean that we can solve the problem by pointing to this kind of explanation. The value problem is a problem in epistemology, and as such an explanation of the appropriateness of valuing knowledge over true belief should be made from the epistemic standpoint. In order to avoid non-epistemic solutions, any such explanation should be grounded in an account of what knowledge *is*. So what we seek is an account of knowledge which enables us to explain why it is more appropriate to love or value knowledge over mere true belief.<sup>6</sup>

But before we can get to this point, we must also acknowledge that there are a variety of ways in which we can love or value something. In other words: to say that a person loves or values the truth is not necessarily to describe any *one* evaluative attitude or emotion, since there are a whole range of positive evaluative attitudes covered by the terms 'valuing' and 'loving.' Here, in support, is what Robert Nozick has to say about evaluative attitudes:

Let us specify a class of verbs, call them the V verbs (for valuing). These include: bringing about, maintaining, saving from destruction, prizing, contemplating, valuing. Depending on the magnitude and nature of the values, other attitudes and relationships are also appropriate...We are to care about, accept, support, affirm, encourage, protect, guard, praise, seek, embrace, serve, be drawn toward, be attracted by, aspire towards, strive to realize, foster, express, nurture, delight in, respect, be inspired by, take joy in, resonate with, be loyal to, be dedicated to, and celebrate values. With the very highest values, we are to be elevated by, enthralled by, love, adore, revere, be exalted by, be awed before, find ecstasy in these highest values.<sup>7</sup>

Once we recognise the variety of different ways in which we can love or value the truth, we open up the possibility that there are different ways in which the truth can *be* valuable.<sup>8</sup> This suggests how reliabilism can solve the value problem. In short, the argument is that reliably-formed true belief is more valuable than accidentally-formed true belief because certain evaluative attitudes are fitting or appropriate with regard to the former, but inappropriate with regard to the latter. Given the connection between value and the appropriateness of loving or valuing, we thereby explain why reliable true belief is more valuable than accidental true belief.

To explain, consider someone who loves or values the truth, and who as a result is motivated to adopt reliable belief-forming strategies which enable her to attain true beliefs. Here we can characterise this person as having certain evaluative attitudes, all with a *practical* bent: she seeks to bring about, strives to realize, dedicates herself to truth. These V-verbs capture the motivational force of her loving; they illustrate the sense in which her valuing involves *pursuit*. Other evaluative attitudes towards the truth seem much less practical. Thus we contemplate, accept, affirm, embrace, and respect truths. Here there appears to be no practical impulse involved, no motive or movement towards the truth. If we class seeking and striving and the like as *active* evaluative attitudes, these others are best characterised as *passive* evaluations.

The distinction between active and passive evaluative attitudes can be employed by the reliabilist in an attempt to solve the value problem. For let us assume that reliabilism is correct, and so let us characterise knowledge as true belief which results from a reliable process or faculty. The problem for the reliabilist is thus to explain why reliably-formed true belief (hereafter RTB) is better than accidentally-formed true belief (ATB). Given our understanding of the value problem from the beginning of this section, this comes down to the problem of showing how loving or valuing RTB is more appropriate or fitting than loving or valuing ATB. It seems that reliabilists can here appeal to the difference between active and passive evaluative attitudes, and maintain that whilst it is appropriate to love RTB in an active sense, such attitudes are

*inappropriate* with regard to ATB. Reliabilists can maintain, in other words, that whilst it makes sense to value RTB in an active way, it does not make any sense to value or pursue ATB in the same way. Given the connection between value and the appropriateness of valuing, we can thus explain a difference in value between RTB and ATB.

But why are active evaluative attitudes inappropriate with regard to ATB? It is inappropriate for us to value ATB in an active sense, since there is nothing we can do to *attain* such beliefs which is compatible with our valuing them. Insofar as we reach true beliefs as a result of adopting some reliable method, these beliefs will not be accidentally true, in which case there is no reliable method we could adopt in order to strive for or pursue accidentally true beliefs. Given this, we might attempt to adopt an *unreliable* method for acquiring true beliefs, such as guessing: we might in this way form the beliefs that there are more red cars than blue cars in Tucson, that Oliver Hardy had a keen interest in gardening, and so on. However, although we could acquire a number of accidentally true beliefs in this way, we will also acquire a much larger number of false beliefs, and this is incompatible with genuinely loving true beliefs.<sup>9</sup> There is, therefore, nothing that we can do in order to attain ATB which is compatible with genuine love, and so it makes no sense to have an active evaluative attitude towards ATB. Valuing ATB in an active way, or having ATB as an epistemic goal or target, is thus highly irrational or inappropriate. Moreover, since there is no incompatibility between valuing RTB and attaining this goal, RTB *can* be a legitimate epistemic goal. Valuing

RTB in an active way, or having reliable true belief as an epistemic goal or target, is thus fitting or appropriate.

This might appear to be too quick. If someone values having true beliefs, then she should adopt reliable strategies in order to attain them. But the fact that someone values having true beliefs and is motivated to adopt reliable belief-forming strategies as a result does not entail that the object of her pursuit, or her epistemic goal, is to attain *reliably-formed true beliefs*. Instead, we might think that her epistemic goal is simply to attain *true beliefs*. (To return to the analogy: if someone values good photographs, then she should get a reliable camera; but it does not follow from this that her goal in getting a reliable camera is reliably-produced good pictures.) As a result, we can deny that RTB is a legitimate epistemic goal, and hence deny that it is somehow more appropriate to actively value RTB instead of ATB.

But this objection fails. Suppose that Bob earns his living as a wedding photographer. It is certainly true that he values taking good photographs, and as a result uses a reliable camera. But it is also true that it *matters* to him that the photographs he takes are reliably-produced: when he hands over the wedding album to the newly-married couple, it would *disturb* him if the best photographs in the album were pure flukes, in the sense that on most other similar occasions he wouldn't have managed to take such good photographs. In such a situation, although he might take pleasure in the couple's delighted reaction to the album, this would be a pleasure tempered by the thought that the photographs, good as they are, might so easily have been poor. This

suggests that what Bob wants isn't just that his photographs turn out to be good, but that they would turn out to be good in most nearby possible worlds: what he wants, in other words, is *reliably-produced* good photographs.

Reliably-produced good photographs can therefore be a perfectly legitimate or understandable photographic goal. By the same token, if Jane values true beliefs, then she had better adopt reliable strategies for attaining them. But it can matter to her that the beliefs that she forms are, on certain important issues, reliably-produced. It would disturb her, as a teacher, if her true belief about the correct grade to award a student paper was a pure fluke, in the sense that on most other occasions she wouldn't have managed to get the grade right. This suggests that what Jane wants isn't just that her beliefs about the grades to be awarded to student papers turn out to be true, but that they would turn out to be true in most nearby possible worlds. What she wants, in other words, is *reliably-produced* true beliefs. If so, there is no problem in regarding RTB as a perfectly legitimate or understandable epistemic goal.<sup>10</sup>

Given, therefore, a plausible axiological principle linking the value of true belief with the appropriateness or fittingness of loving true belief, and given a distinction between different ways in which one can love or value true belief, we can maintain that there are different ways in which true beliefs can be valuable.<sup>11</sup> This allows reliabilists to explain why knowledge is more valuable than true belief. For RTB is an appropriate target of active evaluative attitudes, and so it is valuable as an epistemic goal or end. Since ATB is not

an appropriate target for active evaluative attitudes, it lacks value as an epistemic goal or end.

## 2

In this section a number of objections to the previous argument will be considered, and rejected.

(i) The fact that it makes sense to pursue x instead of y does not entail that x is more valuable than y. This is because we might have very good *pragmatic* reasons for pursuing x instead of y, but these are reasons which hook up in the wrong way with intrinsic value. For instance, if Bill is in an environment which is rich in emeralds but poor in diamonds, it makes sense for him to mine for emeralds rather than diamonds. But it doesn't follow from this that emeralds are more valuable than diamonds. So the fact that it makes sense to pursue RTB instead of ATB does not entail that RTB is more valuable than ATB. Here too the reasons might be purely pragmatic (e.g., one has a pragmatic reason against doing something that it is impossible to do), and therefore might hook up in the wrong way with intrinsic value.

This objection is not fatal, since the analogy is misleading. (Indeed, this objection relies upon ignoring the remarks made in §1 about different standards of appropriateness.) The reason for pursuing emeralds rather than

diamonds is genuinely pragmatic, in the sense that it depends upon something *other than* the nature of emeralds and diamonds themselves, namely the structure of the relevant *environment*. There is nothing in the nature of emeralds and diamonds to dictate that it makes more sense to pursue the former rather than the latter, and this is why we cannot infer, from the fact that it makes more sense to pursue emeralds rather than diamonds, that emeralds are more valuable than diamonds. Things are different in the epistemic case, however, since here the appropriateness or otherwise of pursuit *does* depend straightforwardly upon *intrinsic* differences between RTB and ATB. This is why we can infer, from the fact that it makes more sense to value RTB than ATB, that the former is more valuable than the latter. So the obvious point that a whole host of *extrinsic* factors (such as environment, cost, resources, and the like) can affect what it makes sense to value and pursue does not prove damaging to the previous argument. For this argument insists that the explanation of appropriateness depends upon the intrinsic properties of knowledge and accidentally true belief.

(ii) The value problem is the problem of explaining the difference in value between knowledge and true belief. In §1 it was argued that we can explain this difference by employing a distinction between active and passive evaluative attitudes: since active evaluative attitudes are only appropriate with regard to RTB, then only RTB can be valuable as an epistemic goal or target.

However, knowledge is not only valuable as an epistemic goal or target, but as something that we *possess*. And if we inquire as to the difference in value between RTB and ATB that we possess, the active-passive distinction no longer seems helpful. This is because passive evaluative attitudes are *equally* appropriate with respect to the reliable or accidental true beliefs that we have. We accept and embrace the truth, whether or not it was reliably-formed; we contemplate and affirm the truth, regardless of the process or method by which it was attained. (Indeed, this seems to be the point of the original argument against reliabilism.) Moreover, given that it makes no real sense to strive for or seek out truths we already have, it would seem to be irrational to actively value such true beliefs. Passive evaluative attitudes, and *only* passive evaluative attitudes, are appropriate with respect to our present true beliefs.<sup>12</sup>

This gives rise to the following problem. On the previous account, RTB is more valuable than ATB as an epistemic goal or end. However, it seems that the reliabilist is also committed to holding that RTB and ATB are equally valuable when attained, since they are equally appropriate objects of contemplative or passive evaluative attitudes. But this implies that whatever it is that makes actively valuing RTB over ATB correct or appropriate has nothing really to do with the *value* of RTB or ATB themselves. For it is hardly plausible to maintain that *there is something about RTB* that makes it more worthy of pursuit than ATB, but that the (axiological) influence of this thing disappears once we have such beliefs – as it must if RTB and ATB are

*equally* valuable when possessed. The reason why it makes more sense to pursue RTB over ATB again seems to hook up in the wrong way with the values of RTB and ATB; if so, the solution to the value problem fails.

This is a powerful objection, but it can be overcome. The first thing to note is that the conclusion does not follow from the premises: the facts that (on the previous account) RTB is more valuable than ATB as an epistemic goal, but RTB and ATB are equally valuable when attained, do not necessarily entail that whatever it is that makes actively valuing RTB over ATB appropriate has nothing to do with the value of RTB or ATB themselves (on the grounds that this difference in value is not apparent when it comes to contemplating beliefs that we have). The premises might, instead, entail that whatever it is that makes RTB and ATB equally appropriate objects of contemplation has nothing to do with the value of RTB and ATB themselves (on the grounds that this difference in value *is* apparent when it comes to pursuit). We cannot, *without argument*, assume that it is the appropriateness of passive evaluative attitudes which really counts in a determination of value. If the objection is to be convincing, we need an argument as to why passive attitudes have priority over active attitudes. It seems that this will be a difficult argument to make, given that it is the concept of knowledge as an epistemic *goal* – or of knowledge as the *end* of inquiry – which is of central importance to epistemology. This suggests that it is knowledge as a goal rather than knowledge as an object of contemplation which should be of particular interest

to contemporary epistemologists.<sup>13</sup> Indeed, this is supported by one of the critics of reliabilism – Linda Zagzebski – who herself maintains that “the contemporary model of cognitive activity is inquiry rather than contemplation.”<sup>14</sup> If anything, then, we might think that active evaluative attitudes have priority over passive when it comes to a determination of epistemic value.

The reliabilist has a further response to the objection, which is to deny that RTB and ATB are equally valuable when possessed. The reliabilist can maintain instead that there is something about RTB that makes it more worthy of pursuit than ATB, and that the influence of this thing does *not* disappear once we have such beliefs.<sup>15</sup> In this way reliabilism need not be charged with the task of accommodating the fact that RTB and ATB differ in value as goals, but have the same value when possessed. To explain, in §1 it was argued that it is appropriate to value RTB as an end or goal, whereas it makes no sense to value ATB in this way. If so, there is an obvious sense in which RTB is more valuable than ATB when the relevant beliefs are possessed. Since we value RTB as an end or goal, our possessing RTB constitutes the *fulfilment* of our goal, and such fulfilment has value for us. This is a particular instance of the more general point that *getting what we want* has value for us, at least provided that what we want is suitably free from irrational influences. Since it makes no sense to value ATB as an end or goal, however, our possessing ATB does not constitute the fulfilment or satisfaction of our goal, and so lacks

this additional kind of value. We thus have a straightforward explanation of why RTB is more valuable than ATB when possessed, which is grounded in the nature of these beliefs: RTB, and only RTB, is valuable both because it is true *and* because it fulfills or satisfies our epistemic goal.

(iii) The defence of reliabilism might seem to be suspect when we consider Gettier cases: that is, cases where a subject arrives at a true belief by a method which renders her belief justified, but where (due to the interference of some anomalous or lucky factor) she fails to know the relevant proposition. The problem is this: we value knowledge more than Gettierized true belief, but we cannot explain this difference in value by appeal to a difference in the appropriateness of active evaluative attitudes towards knowledge (as RTB) and towards Gettierized true belief (GTB). Whilst it is appropriate to value RTB over ATB in an active sense, on the grounds that there is nothing that we can do to attain ATB which is compatible with our valuing it, it does not seem appropriate to value RTB over GTB in an active sense, since there is nothing that we can do to attain RTB *as opposed to* GTB: we cannot, for instance, strategize to avoid the kind of anomalous or lucky factors involved in Gettier cases. The attempt to explain differences in value in terms of differences in the appropriateness of active valuing therefore fails to account for the difference in value between RTB and GTB.<sup>16</sup>

Now we can grant that, as a matter of strategy, there is nothing that we can do to attain RTB as opposed to GTB. However, it doesn't follow from this

that active evaluative *attitudes* are equally appropriate with respect to RTB and GTB. That would only follow if it were appropriate to actively value GTB in the first place – which it is not. This is because a subject cannot actively *seek out* the anomalous or lucky factors which lead to GTB, and hence cannot actively pursue GTB. But a subject *can* actively seek out or pursue RTB, and hence can value RTB in an active sense, since valuing RTB in an active sense does *not* require a subject to have a second (impossible) goal of avoiding anomalous or lucky factors. (At least, valuing RTB in an active sense does not require this unless *all* goal-directed behaviour additionally involves the attempt to do what is impossible, namely to avoid unavoidable luck.) As a result, we can explain the difference in value between RTB and GTB in terms of a difference in the appropriateness of active evaluative attitudes towards RTB and GTB: whereas it makes sense to actively value RTB, it does not make sense to actively value GTB. This remains true even though there is nothing that a subject can do to attain RTB as opposed to GTB. The objection generated by Gettier cases therefore fails to be troubling to the defence of reliabilism.

## 3

In conclusion: reliabilist accounts of knowledge have been criticised on the grounds that they cannot explain why knowledge is more valuable than mere true belief. But reliabilists can appeal to the different ways in which we value

things, and thereby the different ways in which things can be valuable, in order to respond to this criticism. In particular, reliabilists can employ a distinction between active and passive evaluative attitudes to argue that only RTB can be valuable as an epistemic goal or aim, and in this sense is more valuable than ATB. In addition, reliabilists can appeal to the fact that RTB constitutes the satisfaction or fulfilment of a subject's goals, in order to explain the difference in value between RTB and ATB that we possess. Reliabilists can therefore solve the value problem.<sup>17</sup>

## Notes

<sup>1</sup> For discussions of the value problem, see M. DePaul (2001), W. Jones (1997), J. Kvanvig (1998), (2003), W. Riggs (1998), (2002), E. Sosa (1997), (2003), and L. Zagzebski (1996), (2000), (2003a), (2003b).

<sup>2</sup> Kvanvig, Jones, and Zagzebski all take this line.

<sup>3</sup> The following argument is due to Zagzebski (2000), (2003a).

<sup>4</sup> A number of reliabilists have attempted to address this problem. See in particular Riggs (2002) and Sosa (1997), (2003). The cogency of their attempts will not be discussed here, since a different kind of defence of reliabilism is proposed in §1. But for an excellent critical discussion of these and other ways of responding to the value problem, see Kvanvig (2003), especially chapter 4.

<sup>5</sup> The traditional has its roots in the writings of Franz Brentano, who maintained that 'We call a thing good when the love relating to it is correct. In that broadest sense of the term, the good is that which is worthy of love, that which can be loved with a love

that is correct.” (Brentano (1969), p. 18.) See also Broad (1930), Ewing (1947), Chisholm (1986), Wiggins (1987), Gibbard (1990), Anderson (1993), Lemos (1994), and Scanlon (1998).

<sup>6</sup> This restriction can be seen as a way of avoiding a ‘conflation problem’ in epistemology. (See D’Arms & Jacobson, 2000a and 2000b.) ‘Fitting attitude’ analysis of value have been criticised recently, on the grounds that it might be right or fitting to have certain attitudes towards objects which are *not* valuable: suppose an Evil Demon threatens to punish us if we do not value a saucer of mud. Such valuing in this instance might be correct, but it does not follow that the saucer is thereby valuable. To avoid this, we can hold that the reason for having the relevant attitude must be grounded in the intrinsic properties of the object itself. For a good discussion of recent attempts to address the problem, see Rabinowicz and Ronnow-Rasmussen (2004).

<sup>7</sup> Nozick (1981), pp. 429-430

<sup>8</sup> See Ewing (1947), p. 166, for the view that different attitudes can fit different kinds of valuable objects, and that ‘good’ therefore has different senses.

<sup>9</sup> Since, as James pointed out, loving the truth and hating falsehoods are separable attitudes, this conclusion might not follow. It might be possible for one to love the truth and yet pursue the truth in an inefficient manner, in which case there is no necessary conflict between loving the truth and amassing a large number of false beliefs. Now although this is a possibility, we might wonder as to the depth of someone’s love of truth were they to decide to pursue it inefficiently. Moreover, since we are talking about intellectually virtuous motives, and we can assume that the virtuous person abhors falsehood in addition to loving the truth, there will be a conflict

between adopting an unreliable method and intellectually virtuous motives. The Jamesian point does not therefore affect the argument here.

<sup>10</sup> Indeed, this point seems to generalise to cover most of the beliefs we form on issues which are important to us. If so, then our *usual* epistemic goal is to attain reliably-formed true beliefs.

<sup>11</sup> The idea that there are different ways in which x can be valuable should not, of course, be surprising. After all, x can have instrumental value whilst lacking intrinsic value, and vice versa. Furthermore, x can have extrinsic value whilst lacking instrumental value, in the case where x's value is relational.

<sup>12</sup> This point is due to Phyllis McKay.

<sup>13</sup> For further support, see Crispin Sartwell (1992), esp. p. 174.

<sup>14</sup> Zagzebski (2003), p. 141.

<sup>15</sup> In addition, this will be fair to the original criticism raised by Zagzebski *et al*, which suggests that we understand the value problem as referring to beliefs we already possess. Reliabilism will not therefore stand accused of dodging the issue.

<sup>16</sup> This objection was suggested to me by an anonymous referee for this *Journal*.

<sup>17</sup> Earlier versions of this paper were presented at staff seminars at the Universities of Leeds, Rijeka, and Stirling, and at a conference on Values and Virtues held at the University of Dundee in 2004. I am grateful to the audiences and delegates for their feedback, and in particular would like to thank Elvio Baccharini, Tim Chappell, Antony Duff, Fiona Macpherson, Andy McGonigal, Phyllis McKay, Duncan Pritchard, Peter Sullivan, Linda Zagzebski, and two anonymous referees for this *Journal* for their helpful comments and criticisms.

## Bibliography

- Anderson, Elizabeth (1993). *Value in Ethics and Economics*, Harvard University Press.
- Brentano, Franz (1969). *The Origin of Our Knowledge of Right and Wrong*, Routledge & Kegan Paul, translated by Roderick Chisholm.
- Broad, C.D. (1930). *Five Types of Ethical Theory*, Routledge & Kegan Paul.
- Chisholm, R. (1986). *Brentano and Intrinsic Value*, Cambridge University Press.
- D'Arms, Justin and Daniel Jacobson (2000a). 'Sentiment and Value', *Ethics*, 100, 722-48.
- (2000b). 'The Moralistic Fallacy: On the "Appropriateness" of Emotions', *Philosophy and Phenomenological Research*, 61, 65-90.
- DePaul, Michael (2001). 'Value Monism in Epistemology', in Matthias Steup (ed), *Knowledge, Truth, and Duty*, Oxford University Press, 170-83.
- Ewing, A.C. (1947). *The Definition of Good*, MacMillan.
- Gibbard, Allan (1990). *Wise Choices, Apt Feelings*, Clarendon Press.

- Jones, Ward (1997). 'Why do we Value Knowledge?', *American Philosophical Quarterly* 34, pp. 423-39.
- Kvanvig, Jonathan (1998). 'Why Should Inquiring Minds Want to Know? *Meno* Problems and Epistemological Axiology', *The Monist*, 81, 426-51.
- (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge University Press.
- Lemos, Noah (1993). *Intrinsic Value: Concept and Warrant*, Cambridge University Press.
- Nozick, Robert (1981). *Philosophical Explanations*, Harvard University Press.
- Rabinowicz, Wlodek and Toni Ronnow-Rasmussen (2004). 'The Strike of the Demon: On Fitting Pro-Attitudes and Value', *Ethics*, forthcoming.
- Riggs, Wayne (1998). 'What are the "Chances" of Being Justified?', *The Monist*, 81, 452-72.
- (2002). 'Reliability and the Value of Knowledge', *Philosophy and Phenomenological Research*, 64, 79-96.
- Sartwell, Crispin (1991). 'Why Knowledge is Merely True Belief', *Journal of Philosophy* vol. 89, 167-80.
- Scanlon, T.M. (1998). *What We Owe to Each Other*, Harvard University Press.
- Sosa, Ernest (1997). 'Reflective Knowledge in the Best Circles', *Journal of Philosophy*, 94, 410-30.

- (2003). 'The Place of Truth in Epistemology', in M. DePaul and L. Zagzebski (eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford University Press, 155-179.
- Wiggins, David (1987). 'A Sensible Subjectivism?' in *Needs, Values, Truth*, Basil Blackwell.
- Zagzebski, Linda (1996). *Virtues of the Mind*, Cambridge University Press.
- (2000). 'From Reliabilism to Virtue Epistemology', in G. Axtell (ed), *Knowledge, Belief, and Character*, Rowman and Littlefield, 113-122.
- (2003a). 'The Search for the Source of Epistemic Good', *Metaphilosophy* 34, 12-28.
- (2003b). 'Intellectual Motivation and the Good of Truth', in M. DePaul and L. Zagzebski (eds), *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford University Press, 135-54.