

An Evaluation of How Hartmann and Boven's Bayesian Probabilistic  
Criterion for Belief Revision Suggests an Epistemological Problem  
for a Divine Ethical Standard

Word Count: 4239

## Contents

Introduction .....	3
Background .....	4
Implications.....	6
Argument.....	7
Comparisons .....	12
Conclusion.....	14
Bibliography.....	16

## An Evaluation of How Hartmann and Boven's Bayesian Probabilistic Criterion for Belief Revision Suggests an Epistemological Problem for a Divine Ethical Standard

NOTE: You don't need to be a mathematician to understand this piece. (I am not a mathematician.) You can skip over the Bay's Theorem diagrams. The point of the piece is to point out the dangers of attempting to ascertain truth. My own conclusions are a bit more cryptic than those revealed in this text.

### Introduction

Luc Bovens and Stephen Hartmann, both professors at the London School of Economics, co-presented a lecture at the 'Third International and Interdisciplinary Conference on Modeling and Using Context' in Dundee, Scotland (2001). Their subsequent paper, entitled 'Belief Expansion, Contextual Fit and the Reliability of Information Sources,' attempts to develop a 'probabilistic criterion for belief expansion that is sensitive *to the degree of contextual fit* of the new information as well as the *reliability of our information source*' (1).

Bovens and Hartmann argue that 'when a scientific community is presented with new data, this does not occur within a vacuum... these new data are being assessed on the background of a context of beliefs.' Most importantly, they claim that 'whether these new data will be accepted or not is a question of the reliability of the sources... as well as of their contextual fit' (2).

While the focus of their paper is on the philosophy of science, its implications for the philosophy of ethics, particularly divine command ethics, are significant. In recent years, there has been a resurgence of interest in divine command theory. Robert Merrihew Adams's essay, 'A Modified Theory of Divine Command Wrongness,'<sup>1</sup> has sparked considerable debate. Others,<sup>2</sup> including Philip Quinn<sup>3</sup> and Richard Mouw,<sup>4</sup> have contributed to the growing dialog.

Adams has addressed that particular and well-known problem for the divine command ethicist posed by Plato's Euthyphro Dilemma.<sup>5</sup> Other proponents have argued along similar lines, correcting, embracing or extending Adams's position. But few of the recent papers have addressed the seminal issue underscoring all divine command theory – that is, even if the notion of a divine command ethics is viable, how does one know that a distinct command is of divine origin? More particularly, how does one judge the truth of two contrary claims, when both

---

<sup>1</sup> Cf. 'Finite and Infinite Goods'

<sup>2</sup> Cf. Stephen Clark, Robert Burch, Edward Wierenga and Paul Rooney

<sup>3</sup> For example, 'Religious Obedience and Moral Autonomy'

<sup>4</sup> For example, 'The God Who Commands'

<sup>5</sup> (1) An act is right because God commanded (or wanted or willed or approved) it. Or alternatively, is it the case that (2) God commanded (or wanted or willed or approved) this act because it is right? (Euthyphro 10a)

purport to be of divine origin? The issue is, in essence, a question of epistemology.<sup>1</sup> And until it is answered, these other arguments have limited utility.

The aim of this essay is to argue in favor of modesty in relation to alleged claims of divine revelation, without denying that the revelation *could* have supreme authority, and more particularly to demonstrate how Hartmann's and Boven's observations regarding belief revision in science suggests a very specific problem for the divine command ethicist. I will argue that this problem embodies something of an 'epistemological trap,' in that one's prior acceptance of a divine ethical standard limits one's capacity to evaluate a competing divine ethical standard. If one is wrong in his earlier conclusion, he is almost unable to evaluate and accept the new standard.

### **Background**

Boven and Hartmann begin their paper with what they call the 'parable about belief expansion.' The parable attempts to illustrate a scenario wherein a new item of information is being presented by a new source of information. The question arises, 'Are we justified to add this new information to what we already know?'

I go to a lecture about wildlife in Greenland which was supposed to be delivered by an expert in the field. When I arrive, I notice that the expert has excused himself and that the biology department has sent a newcomer to fill in for him. I have no beliefs about wildlife in Greenland, but I do have some about Greenland's climate and about the kinds of climate conditions that various types of wildlife favor. (2)

Having established the background, they posit a scenario where the newcomer proclaims three different propositions:

1. That a particular valley in Greenland has large colonies of wild elk
2. That a particular valley in Greenland has large colonies of wild armadillo
3. That a particular valley in Greenland has large colonies of wild boar

Boven and Hartmann suggest that they would be willing to accept Proposition 1 based on their existing belief set. They would not be willing to accept Proposition 2 from either the 'newcomer' or the absent 'expert.' They might be willing to accept Proposition 3, but only if it were espoused by the absent expert. The authors explain, 'when it comes to wild boars the difference between

---

<sup>1</sup> Kant emphasized the acuteness of this problem in his 'Religion Within the Boundaries of Mere Reason' (6:185-6:190).

the newcomer's and the expert's credentials simply make for the difference.' They conclude,<sup>1</sup> 'the more reliable the information source is, the less contextual fit is required for me to be justified to add the belief...' (2).

Next, they construct a mathematical model of the same concept. And while one does not need to be a mathematician to appreciate the conclusions of this model, to fully grasp their approach, one must refer to their paper directly. For the limited purposes of this investigation, however, their equations will only be summarized.

They begin by seeking to establish the following:

1. That a new set of propositions are independent of each other

$$\underline{I(\{REPR_i\}, \mathcal{M}) | \{R_i\}}$$

2. That (for the sake of simplification) all of their sources are equally unreliable

$$\underline{p > q > 0}$$

Then, they attempt to quantify the following:

3. The degree of confidence in the background information<sup>2</sup> (the existing belief set) before the new data

$$\underline{P^*(R_1, \dots, R_n) = P(R_1, \dots, R_n | REPR_1, \dots, REPR_n)}.$$

4. A threshold value for justifying their belief

$$\underline{P^*(R_1, \dots, R_n) \geq t}.$$

5. The degree of confidence in the background information conjoint with the new information<sup>3</sup>

$$\underline{P^{**}(R_1, \dots, R_{n+1}) = P(R_1, \dots, R_{n+1} | REPR_1, \dots, REPR_{n+1})}.$$

Lastly, they calculate the following:

6. That they are justified to expand their belief set just in case<sup>4</sup>

<sup>1</sup> For more on this see Duhem's 'The Aim and Structure of Physical Theory' and Quine's 'Two Dogmas of Empiricism.'

<sup>2</sup> Here they are allowing belief to correspond to a sufficiently high degree of confidence. This follows a tradition in epistemology that goes back to John Locke.

<sup>3</sup> This is the posterior joint probability of this information after the new report has come in.

<sup>4</sup> It can be shown that, for any set of propositions  $R_1, \dots, R_m$ , given the constraints on  $P$  in (1) and (2) [see illustration at end of note] in which the likelihood ratio  $x := q/p$  and  $a_i$  is the sum of the joint probabilities of all combinations of values

$$\underline{P^{**}(R_1, \dots, R_{n+1}) \geq t.}$$

Now, it is possible to postulate two scenarios wherein one has a set of ethical beliefs, and wherein one is challenged with a new ethical idea. In the first scenario, this new item has a high degree of contextual fit. In the second, this new item has a low degree of contextual fit.

**Scenario 1:**

$$\begin{aligned} P^{**}(R_1, R_2, R_3) &= P'(R_1, R_2, R_3 | REPR_1, REPR_2, REPR_3) \\ &= \frac{.15}{.15 \cdot .05^0 + 0 \cdot .05^1 + .5 \cdot .05^2 + .35 \cdot .05^3} = .99 \end{aligned}$$

**Scenario 2:**

$$\begin{aligned} P^{**}(R_1, R_2, R_3) &= P''(R_1, R_2, R_3 | REPR_1, REPR_2, REPR_3) \\ &= \frac{.05}{.05 \cdot .05^0 + .3 \cdot .05^1 + .45 \cdot .05^2 + .2 \cdot .05^3} = .76 \end{aligned}$$

If  $t = .80$  (see point 6), one may justify expanding their belief set based on the reliability of the sources in Scenario 1, but not in Scenario 2 (note the number .99 in Scenario 1, and the number .76 in Scenario 2). The point is relatively simple. To create a condition where the new ethical idea is accepted, one must either lower the standard (below .80) or increase the reliability of the idea's source (above .76).

**Implications**

Boven and Hartmann's primary contribution in their paper is the recognition of this ratio between contextual fit and source reliability.<sup>1</sup> They rightly argue that for a new idea to be accepted, it must fit well with an existing set of beliefs or it must have extremely reliable source.

This has significant implications for the ethicist. One might suppose that Abraham's compliance with God's command to murder Isaac was a very specific type of ethical situation wherein there

of the variables  $R_1, \dots, R_n$  that have  $i$  negative values and  $n - i$  positive values...The proof is straightforward: Apply Bayes' Theorem to the right-hand side of (3); simplify on grounds of the conditional independences in (1) and substitute in the parameters  $p$  and  $q$  as defined in (2).

$$P^*(R_1, \dots, R_m) = \frac{a_0}{\sum_0^m a_i x^i}$$

<sup>1</sup> They are, in fact, contrasting their viewpoint with traditional AGM-Style belief revision originated by Alchourrón, Gärdenfors, and Makinson (1985) and developed further by Gärdenfors (1988, 1992) and Alchourrón and Makinson (1982, 1985).

was an extremely low contextual fit. That is, (we might assume)<sup>1</sup> Abraham believed murdering innocent children was wrong, but that Abraham had extreme confidence in his information source, that is (he believed) God himself commanded the deed.

By its very nature, a divine command represents the claim of ultimate authority. It is either divine/ultimate or not. There is no room for middle ground. One may be semi-sure that a given command is divine, but this does not mitigate the fact that the command itself cannot be semi-divine. This is an important distinction. The purpose of this paper does not entail establishing a criterion for determining if a given command is divine, but rather it attempts to show certain implications that arise from the fact that a given command is *potentially absolutely* divine.

To move forward, one must 'stretch' Boven's and Hartmann's equation to the extreme. Of particular interest is the special case *where the contextual fit is almost nil but the source of the new idea is potentially infinitely reliable*. Consider a scenario wherein a particular divine command ethicist is committed to a particular set of ethical beliefs based on a particular set of divine commands. Then, for the purpose of this investigation, assume that this ethicist is in error, and that the truth lies with a competing set of divine commands.

How does this ethicist realize her error? Can she use the tool of logic to correct her thinking?

### **Argument**

This paper contends that one's prior acceptance of a divine ethical standard limits one's capacity to evaluate a competing divine ethical standard. The form of its argument is as follows:

Wherein:

- P. An ultimate authority claim (UAC) is self-attesting.
- Q. One cannot reasonably evaluate a UAC by a competing UAC.
- R. One's prior acceptance of a UAC limits one's capacity to evaluate a competing UAC.

1.  $P \rightarrow Q$
2.  $Q \rightarrow R$
3. P
4.  $P \rightarrow R$  (Hypothetical Syllogism 1, 2)
5. R (Modus Tollens 3, 4)

So that:

---

<sup>1</sup> This is presumed on the basis of pre-Mosaic moral code among the followers of YHWH as evidenced by Abraham's tithe to Melchizedek, and other religious observations dating from Seth forward. CF: Keil & Delitzsch, also Kaiser.

1. If a UAC is self-attesting, then one cannot reasonably evaluate a UAC by a competing UAC.
2. If one cannot reasonably evaluate a UAC by a competing UAC, then one's prior acceptance of a UAC limits one's capacity to evaluate a competing UAC.
3. A UAC is self-attesting.
4. If a UAC is self-attesting, then one's prior acceptance of a UAC limits one's capacity to evaluate a competing UAC.
5. Therefore, one's prior acceptance of a UAC limits one's capacity to evaluate a competing UAC.

From this form, we adopt this final syllogism:

6. One's prior acceptance of a UAC limits one's capacity to evaluate a competing UAC.
7. A divine ethical standard is a UAC.
8. Therefore, one's prior acceptance of a divine ethical standard limits one's capacity to evaluate a competing divine ethical standard.

Point 7 is reasonably evident. Clearly, this argument (especially points 1-5, 6, and 8) rests on the propositions P and Q. Moreover, since P is more a matter of explanation than attestation, it will briefly be examined, while Q will receive the most attention.

#### *The Case for Proposition P*

Boven and Hartman recognized that a problem with contextual fit may be remedied by an increase in the reliability of the source, though the two authors did not envision a scenario wherein the reliability of the source is absolute. But this challenging possibility is synonymous with the essence of a UAC.

A UAC is self-attesting. It occupies a distinct philosophical category precisely because it could be issued by an ultimate authority. If and only if a UAC was actually issued by an ultimate authority, then it would be supported by ultimate attributes (for instance, power and knowledge). If and only if a UAC were supported by ultimate attributes, then *its highest endorsement can only come from itself*.

So then, the position of an ultimate authority is not endangered by the objection of a limited authority. And one cannot subject an ultimate authority to a limited authority without implying that the limited authority is actually higher than the ultimate authority (there can be only one



ultimate<sup>1</sup>), so that to persist with an evaluation of a UAC by a limited authority is, in effect, to elevate the limited authority to the status of ultimate authority.

This has significant implications for the divine command ethicist. A divine command is essentially a UAC, and as such, it is self-attesting, so that one cannot *reasonably* evaluate a divine command by a competing divine command. This leads us to a discussion of Q.

### *The Case for Proposition Q*

One cannot reasonably evaluate a UAC by a competing UAC, because to employ an ultimate authority to validate/invalidate another ultimate authority is to start with a proposition that assumes the conclusion. But one cannot argue with a proposition that assumes the conclusion. Such an argument might be stated as follows:

Wherein:

x = an ultimate authority claim

y = a competing ultimate authority claim

(P) If x agrees with y, then y is true.

(Q) x does not agree with y.

(R) Therefore, y is not true.

In keeping with the special nature of a UAC, x and y are mutually exclusive (there can be only one ultimate), so to begin an argument with x is to summarily invalidate y. This approach results in the fallacy of *petito principii*. Therefore, one cannot rationally evaluate a UAC by a competing UAC.

Still, it might be objected that while the divine command ethicist may not substitute x with a competing *agent*, she may substitute x with a decisive *standard*, such as reason. So that, in favor of that *standard*, she might reject the previous UAC in favor of the new UAC. Indeed, the divine command ethicist might simply claim that, based on new facts, the competing UAC is the most reasonable choice.

---

<sup>1</sup> I take 'ultimate' here to be synonymous with the notion 'supreme.'

But the standard by which one evaluates a UAC must come from the UAC itself. As demonstrated earlier, a UAC is potentially supported by ultimate attributes, so that its highest endorsement can only come from itself. One cannot subject a UAC to the standard derived from a lesser authority claim without implying (with circular reasoning) that the lesser authority is actually higher than the ultimate authority. Moreover, for one to assert that a UAC is subject to a particular external standard, one must make an ultimate authority claim oneself. The whole attempt is ludicrous.

Nevertheless, it may be further objected (with considerable force) that one may derive some standard of reason from the UAC itself, as some form of reason is necessarily involved in any claim. Mind is an expression of reason. Indeed, even to communicate a particular claim, one must employ a proposition.<sup>1</sup> Without reason in its most basic form, one could hardly communicate at all. Accordingly, the divine command ethicist might draw a standard of reason from the original UAC with which to evaluate the new, competing UAC.

But to evaluate a UAC with reason implies both an *agent* and a *standard*. One cannot separate the two. The divine command ethicist, by virtue of her acceptance of the possibility of the divine, finds herself in the very precarious position of a severely limited agent evaluating a potentially unlimited agent (ultimate authority). Three observations are in order:

First, she must attempt the (tenuous) project of self-evaluation. Did she have an authentic experience? Or was there a mitigating factor in her mind or body that precipitated her experience?

Second, having determined her own reliability, she must somehow evaluate the reliability of the ultimate authority. And in most (if not all cases) the only way she has to evaluate the authenticity of the ultimate authority is to evaluate the new claim itself. Somehow, she must establish that she has an epistemic warrant for the authenticity of the claim, which in turn implies that she has an epistemic warrant for the authenticity of the authority.

Third, for this process to be plausible, her evaluation tool (reason) must be sufficient to discern truth. She must be able to subject the new claim to the powers of her reason and either reject or accept it.

Herein lies her challenge: reason *may* be necessary to corroborate truth, but it is not (always) sufficient. Reason is a process that is dependent upon information (principles, meanings,

---

<sup>1</sup> Even a single sentence requires predication.

premises). If one starts with incomplete or incorrect information, one may correctly employ reason and yet arrive at a wrong (though valid) conclusion.

Reason, then, is necessarily limited. This limitation comes not because of its deficiency, but because of its insufficiency. To some extent, the force of reason is limited by its agent's (a) capacity and (b) knowledge. The problem may be illustrated with a minor formula:

Wherein:

x = capacity

y = knowledge

$$x \rightarrow \infty + y \rightarrow \infty > 5x + 10y$$

The left side of the equation represents an ultimate authority with infinite capacity and knowledge, while the right represents a limited authority with an arbitrary quantity of capacity and knowledge. The differential<sup>1</sup> between the limited agent and the unlimited (ultimate) authority leads to a principle difficulty for the divine command ethicist seeking to evaluate a UAC.

There are potential cases wherein a claim may only appear irrational to an agent with lesser knowledge or capacity. For example:

An ill little girl may reason that it is cruel for a physician to force a needle into her arm. She may conclude that the physician is only increasing her already severe pain, but her conclusion is likely wrong for at least two reasons. One, compared to the physician, she has a limited knowledge of medicine. Two, compared to the physician, she has limited capacity to reason.

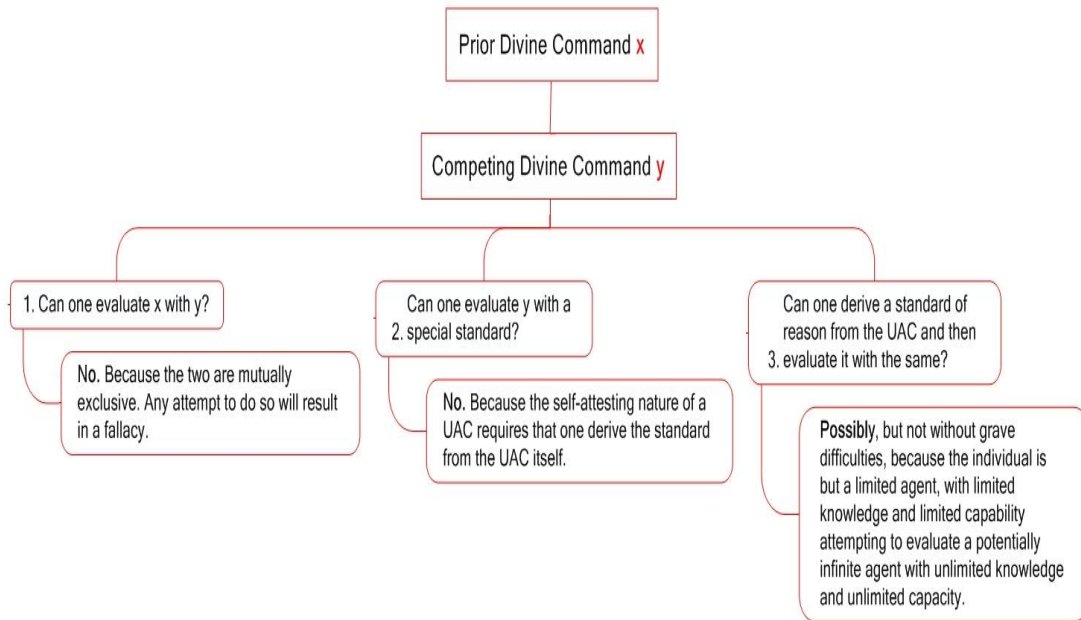
Theoretically, the little girl may structure a valid argument (things that hurt me are bad, the needle is hurting me, therefore, the needle is bad), but its conclusion is/could be wrong. This is analogous to the divine command ethicist who is but a *limited agent, with limited knowledge and limited capability, attempting to evaluate a potentially infinite agent, with unlimited knowledge and unlimited capacity.*

The attempt hardly seems promising.

---

<sup>1</sup> This differential could account for those claims of an ultimate authority that may appear to be a contradiction, when they may be, in actuality, a paradox.

In sum, the divine command ethicist is faced with a peculiar problem not envisioned by Boven's and Hartman's idealizations – that special situation wherein the reliability factor of the source is potentially infinite. By virtue of her own acceptance of a divine command, she must accept the special category of ultimate authority claims – with all of its implications. Accordingly, she will find it particularly difficult to evaluate the claim of a new (and contrary) divine command, as the following chart illustrates:



### Comparisons

To demonstrate the peculiar situation of one who accepts the possibility of a UAC, it may be helpful to contrast the position of the divine command ethicist with the 'reason-centric ethicist.'<sup>1</sup> This may best be accomplished by setting the two against each other in a brief argument:

Now, the reason-centric ethicist may begin his criticism of the divine command ethicist by arguing that unencumbered by an allegiance to the divine command of an ultimate authority, he is free to evaluate a new ethical claim with the tool of reason. Whereas the divine command ethicist is trapped by her obligation to the ultimate authority of the prior divine command and that for reasons already outlined, she cannot readily evaluate the new claim.

But the divine command ethicist is not without a defense. She may respond that for a reason-centric ethicist, (his limited level of) reason itself serves as the *ipso de facto* ultimate authority, so that in actuality, his position mirrors her own. After all, if reason is the ultimate authority for

<sup>1</sup> This may broadly refer to either the a/non-theist or the 'otherly' (perhaps Kantian) theist – that is the theist who leaves little or no room for God to interact in this world and who, therefore, relies mostly on his 'God-given' ability to reason.

the reason-centric ethicist, how can he use reason to evaluate a new, competing (perhaps divine-command-style) UAC that claims to transcend reason?

And the divine command ethicist may press her case further, by pointing out that she may reject outright a new competing divine command ethic on the basis of her allegiance to the prior divine command ethic (while candidly admitting she cannot do so with pure reason). But the reason-centric ethicist cannot make this move without engaging in an absurd, double-layered contradiction.

For if reason is his ultimate authority, and all he has is reason, then he must use reason to reject the UAC, but this itself engages a logical fallacy. In effect, he can neither use reason as the standard to evaluate the UAC, nor use his commitment to the ultimate authority of reason to reject the UAC. He is left with nothing, while she at least has her (perhaps blind) faith in the prior divine command.

But the reason-centric ethicist may counter with a telling blow: On what basis did the divine command ethicist accept the prior divine command ethic? Was reason involved in the decision (and how could it not be)? If so, then her decision to accept the ultimate authority of the divine command ethic was based on an external standard (reason) which contradicts the decision itself.

For if reason were the original 'decider,' then reason is really her ultimate authority, at which point her entire ethical system is based upon contradiction. If she responds that she derived her standard of reason from the UAC itself, he may point out that she is now in no better position than she left him in her last counter. Moreover, she is moving to a position that is indefensible, denying the supremacy of reason while employing it in an ultimate argument for her case.

At this point, the divine command ethicist is running out of rejoinders. Her best tactic may be to argue that her acceptance of the prior divine command ethic signifies her belief in the probability/actuality of ultimate authority, whereas the reason-centric ethicist cannot rule out the possibility of ultimate authority, and so there is a very real sense in which he is facing a dilemma that is similar to hers.

Yet in the end, the matter cannot be resolved with the arguments. For it is certainly possible to be right (in a realist sense) without being able to prove it, and in the special case where one is supporting a trans-rational claim, it is quite likely that one cannot prove it.<sup>1</sup> Indeed, if one can successfully reason for a trans-rational claim, then the claim is hardly trans-rational.

---

<sup>1</sup> Any arguments to the contrary are self-defeating, as the two concepts are mutually exclusive.

Nevertheless, the divine command ethicist must concede that she is in a very dangerous position. For by embracing the special category of the UAC, she expresses a willingness to commit to a standard of which (a) she cannot independently justify, and of which (b) she cannot be certain. This might well evoke a degree of modesty. After all, even if the reason-centric ethicist is wrong, it is unlikely that every divine command ethicist, from every differing religious position, can be right at the same time.

### **Conclusion**

Boven and Hartman's emphasis on the reliability of the source in belief expansion/revision points to a specific problem for the divine command ethicist. The divine command ethicist has necessarily accepted the possibility of an ultimate authority claim (UAC). The nature of a UAC suggests a condition wherein an extremely low contextual fit may be compensated by an infinitely reliable source.

*A UAC is self-attesting.* It occupies a distinct philosophical category precisely because it could be issued by an ultimate authority. If and only if a UAC was actually issued by an ultimate authority, then it would be supported by ultimate attributes (for instance, power and knowledge). If and only if a UAC were supported by ultimate attributes, then its highest endorsement can only come from itself.

*One cannot reasonably evaluate an ultimate authority claim by a competing UAC,* because the two (ultimate) claims are mutually exclusive, so that to validate/invalidate one by the other is to start with a proposition that assumes the conclusion.

One cannot reasonably evaluate the new claim with a special (external) standard because the new competing claim may actually be issued by an ultimate authority, and so it may be inclusive of ultimate attributes. And if it is inclusive of ultimate attributes, then its highest endorsement can only come from itself.

Moreover, if one tries to extract a standard of reason from the new claim itself, one must enforce the standard as a limited agent, with limited knowledge and limited capability, attempting to evaluate a potentially infinite agent. with unlimited knowledge and unlimited capability.

*Therefore, one's prior acceptance of a UAC limits one's capacity to evaluate a competing UAC,*  
and therefore, one's prior acceptance of a divine ethical standard limits one's capacity to  
evaluate a competing divine ethical standard.

## **Bibliography**

1. Adams, R., 'A Modified Divine Command Theory of Ethical Wrongness,' in Helm, P. (ed.) Divine Commands and Morality, Oxford: Oxford University Press, 1981, pp. 83-108
2. Adams, R., 'A Divine Command Metaethics as Necessary A Posteriori,' in Helm, P. (ed.) Divine Commands and Morality, Oxford: Oxford University Press, 1981, pp.109-119
3. Adams, R., Finite and Infinite Goods, Oxford University Press, New York, 1999.
4. Alchourrón, C. E., P. Gärdenfors, and D. Makinson: 1985, 'On the logic of theory change: Partialmeet functions for contraction and revision,' Journal of Symbolic Logic 50 (1985), 510-530.
5. Alchourrón, C. E., and D. Makinson, 'The logic of theory change: Contraction functions and their associated revision functions,' Theoria 48 (1982), 14-37.
6. Bovens, L. and S. Hartmann, 'Belief Expansion, Contextual Fit and the Reliability of Information Sources,' Lecture Notes In Computer Science 2116 (2001), 421-424. (Page number referenced are to the PDF download, available at this address: <http://philsci-archive.pitt.edu/archive/00002434/>)
7. Duhem, P., The Aim and Structure of Physical Theory, trans. Philip P. Wiener, Princeton, NJ, Princeton University Press, 1954.
8. Gärdenfors, P., 'Belief revision: a vade-mecum,' in Pettorossi, A. (ed.) Meta-Programming in Logic, Berlin: Springer, 1992, pp. 1-10.
9. Gärdenfors, P., Knowledge in Flux, MIT Press, Cambridge, MA, 1988.
10. Kant, I., Religion Within the Boundaries of Mere Reason (6:185-6:190), Cambridge University Press, Cambridge, 2004.
11. Keil, C. F., and F. Delitzsch, Commentary on the Old Testament (Vol. 1, p. 158), Hendrickson, Peabody, MA, 2002.
12. Kierkegaard, S., Concluding Unscientific Postscript to Philosophic Fragments, Princeton University Press, Princeton, NJ, 1992.
13. Kierkegaard, S., Philosophic Fragments, Princeton University Press, Princeton, NJ, 1985.
14. Mouw, R., The God Who Commands, University of Notre Dame Press, Notre Dame, IN, 1990.
15. Plato, Plato: Complete Works (Euthyphro 10a) ed. J.M. Cooper, Indianapolis, IN: Hackett, 1997.
16. Quine, W.V., 'Two Dogmas of Empiricism,' Philosophical Review 60 (1951), 20-34.
17. Quine, R., 'Religious Obedience and Moral Autonomy,' in Helm, P. (ed.) Divine Commands and Morality, Oxford: Oxford University Press, 1981, pp. 49-66