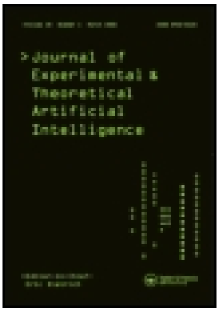


Journal of Experimental & Theoretical Artificial Intelligence
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/teta20

# Epistemic paradox and the logic of acceptance

Michael J. Shaffer [a]

[a] Department of Philosophy , St. Cloud State University , Centennial Hall 365, 720 4th Ave. S., St. Cloud , MN , 56301 , USA

To cite this article: Michael J. Shaffer (2013) Epistemic paradox and the logic of acceptance , Journal of Experimental & Theoretical Artificial Intelligence, 25:3, 337-353, DOI: 10.1080/0952813X.2013.783018

To link to this article: http://dx.doi.org/10.1080/0952813X.2013.783018

*Journal of Experimental & Theoretical Artificial Intelligence*, 2013
Vol. 25, No. 3, 337–353, http://dx.doi.org/10.1080/0952813X.2013.783018





# Epistemic paradox and the logic of acceptance

Michael J. Shaffer*

*Department of Philosophy, St. Cloud State University, Centennial Hall 365, 720 4th Ave. S., St. Cloud, MN, 56301, USA*

(*Received 29 June 2012; final version received 21 January 2013*)

Paradoxes have played an important role both in philosophy and in mathematics and paradox resolution is an important topic in both fields. Paradox *resolution* is deeply important because if such resolution cannot be achieved, we are threatened with the charge of debilitating irrationality. This is supposed to be the case for the following reason. Paradoxes consist of jointly contradictory sets of statements that are individually plausible or believable. These facts about paradoxes then give rise to a deeply troubling epistemic problem. Specifically, if one believes all of the constitutive propositions that make up a paradox, then one is apparently committed to belief in every proposition. This is the result of the principle of classical logical known as ex contradictione (sequitur) quodlibet – that anything and everything follows from a contradiction, and the plausible idea that belief is closed under logical or material implication (i.e. the epistemic closure principle). But, it is manifestly and profoundly irrational to believe every proposition and so the presence of even one contradiction in one's doxa appears to result in what seems to be total irrationality. This problem is the *problem of paradox-induced explosion*. In this paper it will be argued that in many cases this problem can plausibly be avoided in a purely epistemic manner, without having either to resort to non-classical logics for belief (e.g. paraconsistent logics) or to the denial of the standard closure principle for beliefs. The manner in which this result can be achieved depends on drawing an important distinction between the propositional attitude of belief and the weaker attitude of acceptance such that paradox constituting propositions are accepted but not believed. Paradox-induced explosion is then avoided by noting that while belief may well be closed under material implication or even under logical implication, these sorts of weaker commitments are not subject to closure principles of those sorts. So, this possibility provides us with a less radical way to deal with the existence of paradoxes and it preserves the idea that intelligent agents can actually entertain paradoxes.

**Keywords:** paradox; belief; acceptance and logic

## Introduction

Paradoxes have often played an important role both in philosophy and in mathematics. This can be seen from even the most casual inspection of any number of specific monographs on paradoxes and from inspection of numerous general works on the histories of philosophy and mathematics (see Kline, 1980; Sorensen, 2003). One needs to only consider the effect of Russell's famous paradox on mathematics to make this point. Of course, the Russell paradox gave rise to various radical revisions of set theory, and this fundamentally changed our

*Email: mjshaffer@stcloudstate.edu

understanding of the foundations of mathematics while scuttling Frege's logicist programme. The sorites paradox gave rise to the burgeoning study of vagueness and to the developments in the theory of supervaluations. Similarly, the liar paradox has given rise to a number of radical innovations in the theory of truth. These are, of course, merely examples of the important role that paradoxes have played in philosophy and mathematics. However, it should be amply clear that awareness of these logical and epistemic problems has served to guide the development of both philosophy and mathematics in noticeable and often painful ways.

Perhaps, the most important effect that paradoxes have had on mathematics and philosophy derives from paradox resolution. A number of important logical and mathematical developments can be explicitly traced back to specific attempts to resolve paradoxes, and the importance of these attempts to resolve paradoxes is a direct consequence of the problems that are supposed to follow from the presence of contradictions in our belief systems. This further indicates a slightly deeper and more troubling issue. Specifically, paradox resolution is taken to be a deeply important task because if such resolution cannot be achieved, we are threatened with the charge of debilitating irrationality. This is supposed to be the case because of the well-known reason discussed below.

Two of the striking features of true paradoxes are that such paradoxes consist of jointly contradictory sets of statements and that the individual propositions that make up such paradoxes are individually plausible or believable. These facts about paradoxes then give rise to a deeply troubling epistemic problem. Specifically, if one believes in all of the constitutive propositions that make up a paradox, then one is apparently committed to belief in every proposition. This is, of course, both the result of the principle of classical logic known as ex contradictione (sequitur) quodlibet (ECQ) – that anything and everything follows from a contradiction – and the plausible idea that belief is closed under logical or material implication (i.e. the epistemic closure principle). But, it is manifestly and profoundly irrational to believe every proposition and so the presence of even one contradiction in one's doxa appears to result in what seems to be total irrationality.

Let us refer to this general problem here as the *epistemic explosion problem* and let us refer to the more specific case of epistemic explosion in the context of paradoxes as the *problem of paradox-induced explosion*. In this paper, it will be argued that in many cases the *problem of paradox-induced explosion* can plausibly be avoided in a purely epistemic manner, without having to resort either to non-classical logics for belief (e.g. paraconsistent logics) or to the denial of the standard closure principle for beliefs.

The manner in which this result can be achieved depends on drawing an important distinction between the propositional attitude of belief and that of acceptance, although in a manner rather different from one in which Cohen (1992) has, for example, distinguished the two concepts. In effect, what will be argued for here is that the epistemic sting of many paradoxes can be ameliorated by recognising that while the propositions that constitute a paradox may be individually believable or plausible, they are not typically the objects of belief on the part of the epistemic agent who is considering a paradox. Rather, when simply considering paradoxes, the epistemic agent can be understood to be merely committing to the paradox constituting propositions in a weaker sense. We can then avoid paradox-induced explosion by noting that while belief may well be closed under material implication or even under logical implication, these sorts of weaker commitments are not subject to closure principles of those sorts. So, this possibility provides us with a less radical way of dealing with the existence of paradoxes and it preserves the idea that intelligent agents can actually entertain paradoxes.

In effect, it will be suggested here that it is possible that one can epistemically entertain paradox constituting sets of propositions in a perfectly reasonable sense without having to accept any sort of non-classical logic that denies ECQ for belief, without denying some form of the standard epistemic closure principle for belief and without ending up believing everything. This is a desirable sort of solution to the problem of paradox-induced explosion as it satisfies the principle of minimum mutilation in preserving both the classical logic for belief and a standard closure principle for belief. So at the very least, it is to be preferred on that basis. Moreover, as will be demonstrated in what follows, there is a good reason to believe that many propositions we entertain are not strictly believed and that the introduction of other weaker propositional attitudes is not merely an *ad hoc* way to avoid the problem of paradox-induced explosion. This solution also suggests that modelling truly intelligent agents should respect the desiderata that such agents can entertain paradoxes and that designed agents should not be exclusively modelled by belief structures. Two additionally interesting side effects of this solution are that our commitments to alternative logics can be better explained and that logical pluralism can be given an epistemic basis.



## The nature of paradox

Real paradoxes such as those noted above are sorts of epistemic and logical oddities. The basic nature of a paradox in this sense is that it involves a set of propositions $\Lambda$ each of which is *prima facie* reasonable to endorse, but where (in the context of background knowledge $\Sigma$) the set $\Lambda$ appears to imply a contradiction. So paradoxes are essentially sets of propositions that appear to be individually, rationally endorsable but which cannot collectively be endorsed. This can be because the set $\Lambda$ is itself internally inconsistent or because $\Lambda$ appears to imply some proposition $p$ and $\Sigma$ implies $\neg p$ (i.e. the denial of $p$). Let us refer to a given set $\Lambda_i$ as the *paradox constituting propositions* of paradox $i$. We can then also present paradoxes as deductive arguments where the members $\lambda_1, \lambda_2, \ldots, \lambda_n$ of a given set $\Lambda$ are the premises and where they either appear to directly imply ($p$ and $\neg p$) or where $\Lambda$ appears to imply $p$ and $\Sigma$ implies $\neg p$.[1] So let us now turn our attention to a specific example of a paradox in the sense that we have just explicated.

## Paradoxes

In order to illustrate the main points to be made here, let us consider a particular paradox: the lottery paradox. This paradox is a particularly troubling paradox that has apparently important implications for epistemology, especially for fallibilistic conceptions of knowledge. The lottery paradox, first formulated by Kyburg (1961), arises as a result of sceptical worries that afflict the traditional justified, true, belief account of knowledge, according to which justification is taken to require apodictic certainty. In so far as this account flies in the face of sceptical concerns such as Descartes' evil demon, it seems reasonable to assume that knowledge might not require satisfying such a strict justification condition.[2] So, the idea is that we ought to weaken that condition and accept that knowledge is fallible.

The lottery paradox then arises from the following two general propositions:

[$\lambda_1$] There exists a fair lottery with $i$ tickets;
[$\lambda_2$] There is a number $n$, $0.5 < n < 1$, such that if $P(p) = n$ for $S$, then $S$ knows that $p$.[3]

Given the suppositions $\lambda_1$ and $\lambda_2$, the set $\Lambda_L$ for this paradox, we can then specifically generate the lottery paradox by considering a particular lottery with say 1000 tickets and where we fix $n$ at a sufficiently high level, say 0.999.

From the instance of $\lambda_1$ such that $i = 1000$, it follows that we know that one ticket must necessarily win – because the lottery is fair. So we know that either ticket 1 or ticket 2, or, …, or ticket 1000 will win. However, it also follows from instance $\lambda_1$ that for all *i*, the probability that the *i*th ticket is a losing ticket is 0.999. Thus, by our instance of $\lambda_2$ we know that for all *i*, *i* will lose. In other words, we know that ticket 1 will lose, we know that ticket 2 will lose and, . . . , and that ticket 1000 will lose. However, as already noted when considering our instance of $\lambda_1$, we know that one ticket must win. So we have an explicit contradiction that derives from these two instances. We know that no ticket will win and we know that one ticket must win for this lottery and for this standard for probabilistic knowledge such that $n = 0.999$.

## Epistemic explosion: ECQ and the epistemic closure principle for belief

So, what is so worrying about paradoxes such as the lottery? The answer can be found in a deceptively simple form of argumentation. ECQ is a rather surprising but valid argument form in classical deductive logic. So it is a theorem of classical logic. Its proof is quite simple:

| | |
|---|---|
| P1: $(p \,\&\, \neg p)$ | assumption |
| P2: $p$ | conjunction elimination [P1] |
| P3: $\neg p$ | conjunction elimination [P1] |
| P4: $p \vee q$ | disjunction introduction [P2] |
| P5: $q$ | disjunctive syllogism [P3, P4] |
| $\therefore (p \,\&\, \neg p) \supset q$ | conditional proof [P1–P5] |

Of course, this generalises for any *q* whatsoever and so every contradiction validly implies every proposition in the context of classical logic. By our account of paradox introduced in Section 2, it is then apparent that any bona fide paradox, such as the lottery or the sorites – by implying a contradiction – will logically imply every proposition. However, this is not yet as brutally damaging from an epistemic perspective as it actually turns out to be, because we have not yet introduced the epistemic principle that ultimately gives rise to the problem of paradox-induced explosion, viz. the principle of closure for beliefs. The standard axiomatisations of the logic of belief are KD4–KD45. The relevant axioms that constitute these systems are as follows:

($K$) $B_S(p \supset q) \supset (B_S p \supset B_S q)$
(D) $B_S p \supset \neg B_S \neg p$
(4) $\mathrm{B_S p} \supset \mathrm{B_S B_S p}$
(5) $\neg B_S p \supset B_S \neg B_S p$
(0.2) $\neg B_S \neg B_S p \supset B_S \neg B_S \neg p$
(0.3) $B_S(B_S p \supset B_S q) \vee B_S(B_S q \supset B_S p)$
(0.4) $p \supset (\neg B_S \neg B_S p \supset B_S p)$

What is relevant for the purposes of this paper is that *K*, in particular, generalises to the principle of closure for belief.

The closure principle for belief is a principle that has been the subject of considerable controversy in a number of debates in epistemology including those about scepticism and fallibilism. This principle of epistemic rationality is often rendered simply as follows:

$$\text{(SCBM) If } \mathrm{B}_{Sp} \text{ and } \mathrm{B_S(p \supset q)}, \text{ then } \mathrm{B}_{Sq}.$$

Now this version of the principle, the subjective closure of belief under material implication, says nothing about justification and so it is quite generic. It simply states that if a believer *S* believes *p* and also believes that *p* materially implies *q*, then *S* believes *q*. There is, however, a stronger version of the principle, objective closure of belief under material implication

(OCBM):

$$\text{(OCBM) If } \mathrm{B}_{Sp} \text{ and } (\mathrm{p} \supset \mathrm{q}), \text{ then } \mathrm{B}_{Sq}.$$

This version of the closure principle states that if *S* believes a proposition, then he/she is committed to believe in all of the material implications of that belief. This principle seems rather less plausible, but there is also a weaker version of SCBM, the subjective closure of justified belief under material implication (SCJBM), and it is still sufficient to generate the problem of paradox-induced explosion and so it is all that we need to subscribe to here:

$$\text{(SCJBM) If } JB_{Sp} \text{ and } JB_S(\mathrm{p} \supset \mathrm{q}), \text{ then } JB_{Sq}.$$

This version simply states that if *S*'s belief in *p* is justified and he/she is also justified in the belief that *p* materially implies *q*, then *S* is justified in believing *q*.

Closure can also be understood to involve closure under *logical* implication, and so we get analogues of the closure principles framed in terms of material implication as follows (respectively, the subjective closure of belief under logical implication, the objective closure of belief under logical implication and the subjective closure of justified belief under logical implication):

$$\text{(SCBL) If } B_{Sp} \text{ and } B_S(\vdash p \supset \mathrm{q}), \text{ then } \mathrm{B}_{Sq}.$$

$$\text{(OCBL) If } \mathrm{B}_{Sp} \text{ and } \vdash p \supset \mathrm{q}, \text{ then } \mathrm{B}_{Sq}.$$

$$\text{(SCJBL) If } \mathrm{JB}_{Sp} \text{ and } \mathrm{JB}_S(\vdash p \supset \mathrm{q}), \text{ then } \mathrm{JB}_{Sq}.$$

So, what is crucial at this stage is that we understand that it is commonly believed to be the case that one of these principles governs belief.

Since, ex vi definitionis, paradoxes are constituted by propositions each of which is endorsable, they at least appear to be the propositions that we are individually justified in believing. Moreover, once we have been made aware of the paradoxical nature of that set of propositions, we are also apparently justified in believing that they imply the conclusion of the paradox, which is either an explicit contradiction or a proposition that contradicts a proposition that is an element of our background knowledge. So it appears that for any true paradox, we are justified in believing a contradiction and then by another appeal to SCJBM (or *modulo* by some other closure principle), we are justified in believing every proposition because according to ECQ, that contradiction implies every other proposition and that we know this and are also justified in believing it.

This is why paradox resolution is supposed to be so crucial. The presence of even one paradox in our belief system appears to commit us to the bluntly irrational implications of ECQ. This is a commitment to belief in every proposition. Why is this irrational? Belief in every proposition is irrational in the following sense. Following Popper and Wittgenstein, the information that a proposition carries is just what it rules out in the logical space of possibilities. So, for example, in terms of possible world semantics, the information carried by the proposition 'the sky is blue' is just the set of possible worlds in which it is true. By the same token, this idea can be expressed in the negative as well. The information carried by an expression is just the set of worlds it rules out – the partition it imposes on the set of worlds. An informative proposition is then informative in this sense because it is not true in many worlds. Now if ECQ commits us to believe in every proposition because we believe the premises of that argument and on that basis, we justifiably believe a contradiction, then the information content of our belief system would effectively be zero. The problem of paradox-induced explosion is utterly nihilistic in this sense. It is effectively the denial of rationality lock, stock and barrel. However, this problem can be readily defused without any baroque appeals to paraconsistent logics for belief, etc. by simply paying closer attention to the actual doxastic attitudes that we have with respect to many paradox constituting propositions.

## Beliefs and weaker commitments

The standard propositional attitudes that are typically dealt within epistemology and artificial intelligence are belief and knowledge. However, it is widely accepted that these are not the only propositional attitudes that one can have towards propositional contents, even if this appears to be an often forgotten or ignored point. We only need to consider the attitudes of considering *p*, grasping *p*, supposing that *p* or of wishing that *p* be the case and so on, in order to see that the taxonomy of doxastic/sub-doxastic states is really quite diverse and complex. However, the fact that there has been relatively little discussion of these other propositional attitudes is a rather serious lacuna in epistemology and artificial intelligence, and it is likely that it has given rise to the tendency to over ascribe belief and knowledge to epistemic agents where other propositional attitudes are really at work or should be at work.

It is an important task when we are considering situations or models that involve propositional commitments to distinguish cases involving belief from those that do not involve belief. One effective way of doing this is to distinguish commitments that involve the norm of truth from those that do not involve the norm of truth. This is, of course, because it is widely agreed that the norm of belief is truth. By distinguishing such cases, we can avoid attributing inappropriate features to such situations. This can be effectively accomplished in the case of belief by looking at instances of the following argument scheme (scheme 1):

P1: The operant and appropriate norm in situation $x$ involving $S$'s attitude $\beta$ towards proposition $p$ is $y$.
P2: $y$ is not truth.
P3: Truth is the norm of belief.

Therefore, $S$'s attitude $\beta$ towards proposition $p$ in situation $x$ is not belief.

So if we take seriously the claim that there are true (or even merely possible) substitution instances of this argument scheme, then it is reasonable to believe that we can make sense of the idea that there are propositions that are believable (i.e. it is possible to believe them) and even plausible (i.e. they are not known to be false or do not seem to be false), but that are not actually believed. This is because there can be non-truth-normed rational commitments and we shall see several examples of this in what follows. Some of these commitments are pragmatic in nature and so the sense of believability in those cases is broadly pragmatic, others involve commitments based on plausibility. We know, of course, that $-Bp$ does not entail $Bp$ as a matter of elementary modal logic, but it is also reasonable to suppose that we need not believe a proposition merely because it is plausible and rational to hold for some pragmatic reasons, or because it is merely plausible.[4] So the upshot of this is that it is reasonable to believe that propositions can be rationally entertained but not be believed. In the course of this paper, it will be argued that there are many cases of commitments that cannot reasonably be understood to be beliefs, but which allow us to achieve certain important epistemic and practical goals. These kinds of weaker propositional commitments are best understood as forms of acceptance – as opposed to belief – and once we see acceptance for what it is, it is quite a commonplace attitude to have towards propositions. However, determination of the right kind of weaker propositional attitude appropriate to ascribe in a given situation is crucial if these notions are to do the work required of them here with respect to paradox consideration.

In accordance with the recognition that many commonplace propositional commitments are not beliefs, Cohen (1992), in particular, has usefully distinguished belief from a particular form of acceptance. He treats the latter is voluntary, whereas the former as non-voluntary and he shows how belief and acceptance of this sort have often been conflated with serious negative implications for a number of philosophical issues. Here, it will be argued that two somewhat different but related varieties of acceptance play important roles in the resolution of the problem of paradox-induced explosion. So, let us then begin by looking at the various concepts of acceptance in contrast to the concept of belief. The first important distinction to make with respect to the various attitudes of acceptance concerns the extent of such commitments. So, *S*'s acceptance of *p* is *full*, if and only if *S*'s commitment to *p* is closed under logical or material implication. Where *S*'s commitment is not full in this sense, we will call such acceptance *limited*. The second important distinction to make among the various forms of acceptance concerns the norm that governs such cases of acceptance. So, *S*'s acceptance of *p* is *strong*, if and only if, *S*'s commitment to *p* is such that *S* is plausible for *p*. Where *S*'s commitment is not strong in this sense, we will call the commitment *weak* and (at least for the purposes of this paper) the norm that governs such weak forms of acceptance will be understood as pragmatic utility. So, these two important distinctions yield four important categories of acceptance: strong full acceptance, weak full acceptance, strong limited acceptance and weak limited acceptance. Further, more refined versions of each of these forms of acceptance can then be determined by specifying additional features definitive of each of these types of propositional attitude. But, for the purposes at hand, we can ignore these more fine-grained characterisations and focus on the relevant concepts of weak limited acceptance and strong limited acceptance that play a role in resolving the problem of paradox-induced explosion.

So, as it is to be understood here, *weak limited acceptance* is a propositional attitude like belief and knowledge. Its main features are as follows:

(1) Accepting *p* is purely voluntary.
(2) Accepting *p* is non-evidential.
(3) Accepting *p* is a form of supposition.
(4) Accepting *p* is a pragmatic matter.
(5) Accepting *p* is contextual.
(6) Accepting *p* is not a commitment to the literal truth of *p*.
(7) Accepting *p* is neither subjectively nor objectively closed under implication.[5]

Scheme 1 arguments can then help us discriminate truth-normed commitments, such as belief, from non-truth-normed commitments, such as this form of acceptance, on the basis of Principle 6. In any case, the view endorsed here is that accepting a proposition in this weak and limited way is a sort of voluntary, non-evidential, suppositional, pragmatic and contextual commitment that is something like epistemically 'trying out' or 'using' a proposition and some of its implications in a context, and while the account of weak limited acceptance offered here shares some features in common with Cohen's account, it is appreciably different because in Cohen's account acceptance is characterised by subjective closure under (material) implication. So, Cohen's form of acceptance is a form of weak *full* acceptance. Limited forms of acceptance can be distinguished from the sort of full acceptance that Cohen has described and from belief as well in virtue of the following argument scheme (scheme 2):

P1: In any situation $x$ involving $S$'s attitude $\beta$ towards proposition $p$, $S$'s commitment is governed by closure principle $k$.
P2: $S$'s commitment to $p$ in situation $C$ is not governed by $k$.

Therefore, $S$'s attitude towards proposition $p$ in situation $x$ is not $\beta$.

So, in the case where we plug Cohen acceptance in for $\beta$ and subjective closure of that attitude under material implication for $k$, we can demonstrate that a given commitment is not a case of Cohen acceptance. Of course, we can do the same for belief as well, by plugging in appropriate combinations of attitudes and closure principles, at least where there is a consensus on those matters.

In addition, we can distinguish cases of weak acceptance from cases of strong acceptance by determining whether they involve the requirement that $p$ is plausible for $S$ or whether $S$'s commitment to $p$ is pragmatically motivated. Given this difference, *strong limited acceptance* can be understood as being characterised in terms of the following principles:

(1) Accepting $p$ is purely voluntary.
(2) Accepting $p$ is non-evidential.
(3) Accepting $p$ is a form of supposition.
(4) Accepting $p$ requires that $S$ takes $p$ to be plausible.
(5) Accepting $p$ is contextual.
(6) Accepting $p$ is not a commitment to the literal truth of $p$.
(7) Accepting $p$ is neither subjectively nor objectively closed under implication.

Here, plausibility will be understood in the following sense. $S$ is plausible for $p$, if and only if, $S$ does not know that $\neg p$ and $p$ does not *prima facie* seem to be false to $S$.[6] In the case of weak commitments such as weak limited acceptance, agents need not have this attitude towards the proposition in question. Adopting the attitude of weak acceptance towards a proposition *may* involve propositions that are taken to be plausible or doing so may involve propositions that are, in fact, plausible despite the agent's not taking them to be so, but neither of these conditions are required to weakly accept a proposition. An agent might be pragmatically entertaining a proposition that happens to be plausible, but the plausibility of that proposition will not be the basis on which it is being entertained. Moreover, many such pragmatic commitments will involve propositions that are not plausible. So, let us then turn to the examination of some cases involving weak commitments in order to illustrate the utility and explanatory power that distinguishes these concepts from belief.

### *A scientific example of acceptance*

As it is understood here in its most basic sense, to accept $p$ is to suppose $p$ to be the case for some pragmatic purpose or because $p$ is plausible. Let us first consider the acceptance of Newton's laws of motion by some epistemic agent in order to describe the motion of a projectile fired from a cannon. Now the first thing we should specify here is that this case is intended to be one in which the agent is our contemporary and is fully aware that, strictly speaking, Newton's laws of motion are false. Nevertheless, there are numerous applications where those false laws are good enough to do the job. Suppose, for example, that our agent is doing this in the following context. Let us suppose that he is an instructor in a college teaching a basic physics course and that he is solving the problem on a chalkboard for the class.

Now, it should be clear that our agent does not literally believe Newton's laws to be true in this case. In fact, he is *ex hypothesi* fully aware that they are false and in fact also why they are false. Nevertheless, it is perfectly reasonable to accept them as a matter of supposition in this context while at the same time he believes and knows them to be false. He need not actually add those propositions to his doxastic system and make the appropriate revisions simply to employ those laws to solve the problem. Moreover, he does not need any evidence in order to voluntarily accept Newton's laws for the purpose of solving a problem in Newtonian mechanics. In this sense, his acceptance of those laws is merely a pragmatic matter that has nothing to do with the evidence to the effect that Newton's laws are true. However, he surely would not accept Newton's laws in various other contexts and so acceptance is context bound in the following sense. The instructor probably, for example, would not and should not accept Newton's laws in the context of solving a problem in a relativistic mechanics course for a case where the velocity of the projectile in the problem is supposed to be near the speed of light. So it appears to be the case that we often only accept certain propositions as a matter of their pragmatic usefulness in some specified context where it is, in fact, useful to do so.

With respect to the theory of acceptance introduced here, the real critical departure from Cohen's account concerns principle 7. For Cohen acceptance is closed subjectively under material implication and in this sense, it must be distinguished from the types of limited acceptance with which we are concerned at this point. For Cohen, the subjective closure of acceptance under material implication (SCAM) means that:

$$\text{(SCAM) If } A_{Sp} \text{ and } A_S(p \supset q), \text{ then } A_{Sq}.$$

As we saw earlier, this is to be distinguished from the objective closure of acceptance under material implication (OCAM):

$$\text{(OCAM) If } A_{Sp} \text{ and } (p \supset q), \text{ then } A_{Sq}.$$

Cohen believes that the latter principle is too strong when it comes to acceptance but he is clear that the acceptance of *p* is subjectively closed under material implication. However, this is problematic when it comes to many cases of acceptance, especially when we focus on the voluntariness and the contextuality of acceptance as it relates to closure principles.

In our example, it seems entirely reasonable to assert that the instructor might well accept Newton's laws and to accept that they have certain (classical) implications without actually *accepting all of those implications*. This might initially seem to be a strange thing to assert, but consider the following real possibility. Our instructor accepts Newton's laws for the purpose of solving the projectile motion problem in class, he accepts that Newton's laws (classically) imply that objects travelling near the speed of light will have Newtonian motions, but (1) he does not necessarily accept that objects travelling near the speed of light have Newtonian motions and (2) he certainly need not accept that objects travelling near the speed of light will have Newtonian motions or that the Earth has one moon. This is possible and wholly reasonable, because the instructor's acceptance of Newton's laws is contextually limited given that he is fully aware that the motions of objects near the speed of light are non-Newtonian and are correctly described by the laws of relativistic mechanics. By using Newton's laws, he is *not* thereby necessarily committed to the full set of classical implications of Newton's laws simply in virtue of his accepting them for the purpose of solving this one problem in Newtonian mechanics in his class – even given his explicit awareness of the implications of those mechanical laws.

The contextual 'boundary' imposed by his restricted supposition makes it such that he need not be thereby automatically committed to accept *all* of the accepted classical implications of those propositions that he explicitly accepts or even all of the relevant implications of his supposition. This also accords with the more voluntary nature of this form of acceptance. He could, of course, *voluntarily* accept what he accepts to be relevantly implied by Newton's laws, but he surely need not do so and he would be rationally committed only to accept those implications required by the context in which he is operating. So acceptance appears to be a matter of *extent*, where extent has something to do with the range of implications voluntarily accepted and perhaps the salience of those implications that is determined by context. Our instructor might, for example, want to accept all of those propositions relevantly implied by Newtonian mechanics if he were trying to 'think his way into the Newtonian frame of mind', say as an historian. This would be a plausible case of full acceptance. In that context, he might accept the full body of Newton's theory for the sake of, for example, a careful and non-anachronistic reading of the Leibniz(Clark correspondence. However, in most cases, we only need to accept a subset of the accepted implications of what we accept, and this is clearly not meant to be a subset all of all of the material or logical implications of what we accept. Limited acceptance then only involves accepting some of the implications of what we accept and what we accept to be implied by what we accept.

The propositions involved in this plausible scenario – while believable in the sense described earlier – are not literally believed. We can then see that in this kind of case, we are dealing neither with belief nor with Cohen acceptance. These two points can be made clear especially via the following two arguments: where *C* refers to the classroom situation described above, *I* refers to our instructor and *N* is the set of propositions that constitute Newtonian mechanics:

P1: The operant and appropriate norm in situation *C* involving *I*'s attitude $\beta$ towards proposition *N* is *y*.
P2: *y* is not truth.
P3: Truth is the norm of belief.

Therefore, *I*'s attitude $\beta$ towards proposition *N* in situation *C* is not belief.

In addition, for the case of subjective closure under material implication and Cohen acceptance, we can see that:

P1: In any situation *x* involving *S*'s Cohen acceptance of a proposition *p*, *S*'s commitment is governed by subjective closure under material implication.
P2: *I*'s commitment to *N* in situation *C* is not governed by subjective closure under material implication.

Therefore, *I*'s attitude towards proposition *N* in situation *C* is not Cohen acceptance.

So, we have a plausible case of a commitment to a proposition that is neither belief nor Cohen acceptance, and this case of acceptance is characterised by the main features of the sort of pragmatic acceptance described in Section 5. So, in this case, the instructor's acceptance of Newton's laws is both limited and weak. It is weak because the instructor knows that they are false, they do not even *prima facie* seem true to him/her and he/she is entertaining them purely for pragmatic reasons. But this is not the only sort of situation that can plausibly be thought of to involve these sorts of commitment that are weaker than belief.

### *A logical example of acceptance*

Let us consider another example of acceptance that is in itself quite interesting from the perspective of logic. Consider a logician who is fully familiar with the wide variety of extant logics and who is interested in seeing whether some of the paradoxes of deontic logic can be resolved. Let us further suppose that this logician is interested in determining what would happen if one were to combine the standard system of deontic logic with a specific temporal logic (for whatever reason). Now, in order to do this, it is clear that our logician is not required to literally believe these logics to be true and that he/she is not required to literally believe all of what follows classically from his/her commitment to those sets of propositions simply in virtue of entertaining them for the purpose of this heuristic exercise. Moreover, his/her endorsement of the propositions that constitute these systems of logic is not driven by evidence at all.[7] Of course, he/she must entertain at least *some* of what follows from this combined set of propositions if he/she is to determine whether such an approach would allow for some traction in dealing with the paradoxes of deontic logic, but this need not at all be because he/she thinks that they are true, justified or even plausible. In fact, what he/she should accept in this context in order to achieve his/her goal are all of the accepted consequences of what he/she accepts so that he/she can determine the fruitfulness of this idea. So, again, as with our historically minded scientist, our logician only needs to accept the accepted implications of what he/she accepts. But, his/her commitment to those propositions in the first place is purely voluntary, does not involve a commitment to their truth, does not commit his/her to all of what classically follows from the acceptance of those propositions in the sense of the classical closure of those propositions and need not even involve the idea that those propositions are plausible.[8]

As in the scientific example, the propositions involved in this rather mundane scenario – while believable in the sense described earlier – are not literally believed. We can again see that, in this kind of case, we are dealing neither with belief nor with Cohen acceptance by using schemes 1 and 2: where $M$ refers to the research situation described above, $R$ refers to our logician and ($D$ and $T$) is the set of propositions that constitute a specific system of deontic and temporal logic, respectively:

P1: The operant and appropriate norm in situation $M$ involving $R$'s attitude $\beta$ towards proposition ($D$ and $T$) is $y$.
P2: $y$ is not truth.
P3: Truth is the norm of belief.

Therefore, $R$'s attitude $\beta$ towards proposition ($D$ and $T$) in situation $M$ is not belief.

As in our last case with respect to subjective closure under material implication and Cohen acceptance (i.e. weak full acceptance), we can see that:

P1: In any situation $x$ involving $S$'s Cohen acceptance of a proposition $p$, $S$'s commitment is governed by subjective closure under material implication.
P2: $R$'s commitment to ($D$ and $T$) in situation $M$ is not governed by subjective closure under material implication.

Therefore, $R$'s attitude towards proposition ($D$ and $T$) in situation $M$ is not Cohen acceptance.

So, we have another plausible case of propositional commitment that is neither belief nor Cohen acceptance (i.e. weak full acceptance).[9]

### *A theatrical example of acceptance*

Let us consider a final example of the sort of weak limited acceptance that has been explicated here. John Gielgud famously portrayed Hamlet in a New York production in 1936. In doing so, Gielgud adopted the role of a Danish prince who is the son of his recently deceased father King Hamlet. At the opening of Shakespeare's play, Horatio, the prince's best friend, tells prince Hamlet that a ghost who appears to be the ghost of King Hamlet has been seen wandering about nearby. Prince Hamlet then seeks out the ghost and speaks with it. He identifies it as the ghost of his father. In the course of their conversation, Hamlet's father's ghost exhorts the prince to avenge his death at the hands of Claudius, the present King of Denmark and brother of the former King. Now when Gielgud had adopted the role of Hamlet in his performances during the 1936 production of the play, it does not seem to be reasonable to say that he literally believed that he was Hamlet, that he spoke to the ghost of his (i.e. Hamlet's) father, that his best friend is Horatio, that his uncle Claudius poured poison in King Hamlet's ear and so on. It would be sheer lunacy to suppose this to be the case and it would be sheer lunacy to claim that when he was in other contexts, he continued to maintain those same contextual commitments. Moreover, it does not seem to be reasonable to suppose that in adopting this role, Gielgud was committed to the full set of consequences of his commitment as Hamlet. He surely is not committed to the implications that there really are ghosts and witches and so on simply because he voluntarily adopted the persona of Hamlet in the context of a Shakespeare play.

Once again, the propositions involved in this familiar kind of scenario – while believable in the sense described earlier – are not literally believed, and they are not even plausible. On the basis of this, we can easily see that in this kind of case, we are dealing neither with belief nor with Cohen acceptance, and we can use schemes 1 and 2 to show this: where *N* refers to the theatrical situation described above, *A* refers to our actor and *H* is the set of propositions that constitute the Hamlet story:

P1: The operant and appropriate norm in situation *N* involving *A*'s attitude $\beta$ towards the propositions that are elements of *H* is *y*.
P2: *y* is not truth.
P3: Truth is the norm of belief.

Therefore, *A*'s attitude $\beta$ towards proposition *H* in situation *N* is not belief.

As in the previous two cases, with respect to subjective closure under material implication and Cohen acceptance, we can see that:

P1: In any situation *x* involving *S*'s Cohen acceptance of a proposition *p*, *S*'s commitment is governed by subjective closure under material implication.
P2: *A*'s commitment to *H* in situation *N* is not governed by subjective closure under material implication.

Therefore, *A*'s attitude towards proposition *H* in situation *N* is not Cohen acceptance.

So, we have yet another plausible case of a commitment to a proposition that is neither belief nor Cohen acceptance, and it is characterised by the main features of the sort of pragmatic acceptance described in Section 5. What these examples then show is that it is at least reasonable to believe that there are propositional attitudes that are (1) distinct from belief and Cohen acceptance and (2) that best explain a variety of cognitive commitments that cannot be explained easily as cases of belief or of Cohen acceptance. So let us then return to the issue of paradox consideration and how we can avoid paradox-induced explosion by appealing to the concepts of limited acceptance.

## Paradox-induced explosion defused

The manner in which paradox-induced explosion can be avoided – again without recourse to alternative logics for belief, etc. – should then be clear. As in the three cases discussed earlier in this paper, in the case of paradoxes such as the lottery (on careful analysis), we can see that the natural interpretation to adopt is that one or more of the instances of the general paradox constituting propositions that give rise to such paradoxes are also merely accepted in some limited way. Given the claims then that such limited acceptance is not governed by the norm of truth and that it is neither objectively nor subjectively closed under material or logical implication, paradox-induced explosion is blocked. This is because the explosion argument is not valid for cases of limited acceptance (see Beall & Restall, 2006, chap. 5). Let us see how this approach works in the specific case of the lottery.

### *The lottery paradox*

Remember that the propositions that constitute $\Lambda_L$ in this case and that give rise to the lottery paradox are:

[$\lambda_1$] There exists a fair lottery with $i$ tickets, and
[$\lambda_2$] There is a number $n$, $0.5 < n < 1$, such that if $P(p) = n$ for $S$, then $S$ knows that $p$.

However, while it *might* be plausible to suppose that we actually believe $\lambda_2$ as a general principle and even that we have good evidence in its favour, this is neither true of $\lambda_1$ nor of any particular instance of $\lambda_2$ (such as the claim that if $P(p) = 0.93$ for $S$, then $S$ knows that $p$) necessary to generate the paradox. $\lambda_1$ and many of its instances are purely hypothetical suppositions introduced only for the philosophical purpose of challenging the traditional theory of knowledge, or for challenging closure or agglomerativity. It is totally implausible to believe that anyone, in merely considering the lottery paradox, actually believes in the existence of some particular lottery that would raise the paradox. In this respect, our attitude towards this proposition is similar to our professor's attitude towards the propositions of Newtonian mechanics, to our logicians' attitudes towards the two logics she is studying and to our actor's attitudes towards the Hamlet story in the sense that they are clearly limited. Similar concerns can be raised about instances of $\lambda_2$ and it is also not plausible to claim that those instances are always believed, particularly when we are forced to select an arbitrary value for $n$ in order to generate the paradox. In many philosophical contexts – say in teaching a course on paradoxes – such instances of $\lambda_2$ are not strictly believed because they are introduced in a wholly arbitrary manner for pedagogical purposes and there is certainly no value of $n$ such that it is *known* to be necessary for $S$ knowing that $p$. This is most obviously the case because we do not even have a particular belief that there is an $n$ such that it is necessary for $S$ knowing that $p$. In such a context, instances of $\lambda_2$ are merely accepted for the pragmatic purpose of demonstrating that paradox for the students and the presenter may have no evidence to support the contention that $n$ should be set at say 0.999 or 0.988, or some other value. Such propositions then are plausibly viewed as being voluntarily entertained without recourse to evidence (whether they are in fact plausible) and we are not thereby committed to the acceptance of even those propositions that we accept or believe to follow classically from the paradox constituting set. One or more of the members of $\{\lambda_1, \lambda_2\}$ are merely accepted and so there is no real epistemic problem here. There is no epistemic explosion in the case of propositions that are accepted in the limited sense for pragmatic reasons.

As in our previous cases, we can see that here we are dealing neither with belief nor with Cohen acceptance. This can be seen by applying schemes 1 and 2: where $P$ refers to the philosophical classroom situation described above, $O$ refers to our teacher and $\Lambda_L$ is the set of propositions that constitute the lottery paradox (i.e. appropriate instances of $\lambda_1$ and $\lambda_2$):

P1: The operant and appropriate norm in situation $P$ involving $O$'s attitude $\beta$ towards proposition $\Lambda_L$ is $y$.
P2: $y$ is not truth.
P3: Truth is the norm of belief.

Therefore, $O$'s attitude $\beta$ towards proposition $\Lambda_L$ in situation $P$ is not belief.

As in the previous cases, with respect to subjective closure under material implication and Cohen acceptance, we can see that:

P1: In any situation $x$ involving $S$'s Cohen acceptance of a proposition $p$, $S$'s commitment is governed by subjective closure under material implication.
P2: $O$'s commitment to $\Lambda_L$ in situation $P$ is not governed by subjective closure under material implication.

Therefore, $O$'s attitude towards proposition $\Lambda_L$ in situation $P$ is not Cohen acceptance.

So, the problem of a paradox-induced explosion can be avoided – at least in the context of situations involving our merely entertaining sets of paradox constituting propositions for pragmatic purposes. The attribution of this commitment in these cases best explains them and allows us to make sense of how paradoxes can be understood and epistemically entertained without succumbing to the debilitating irrationality that explosion would bring.

So, we have seen that paradoxes such as the lottery can be entertained for pragmatic purposes by adopting the attitude of weak limited acceptance towards the propositions that constitute those paradoxes. However, this is not the only way in which such paradoxes can be rationally entertained without succumbing to the debilitating irrationality that comes with paradox-induced explosion. In fact, the suggestion that the problem of paradox-induced explosion can be avoided by appealing to *weak* limited acceptance alone overlooks the issue of entertaining paradoxes for non-pragmatic reasons, and of entertaining paradoxes as purely intellectual problems.[10]

Let us then consider the following case. Suppose that our teacher is simply considering the lottery paradox late one night in his office as an epistemic problem with no pragmatic goals in mind. In doing so, he grasps the paradox in a way that it is truly paradoxical in the sense that he is not sure of how to resolve it, as none of the elements of $\Lambda_L$ are such that he knows them to be false and none of them seems obviously to be in error (i.e. they are plausible for him). Here, the attitude that our teacher has towards the paradox is rather different than in the classroom situation, but it is similar in some respects as well. We can see this as follows: where $Q$ refers to our teacher working as a researcher thinking alone in his office, $O$ refers to our teacher qua researcher and $\Lambda_L$ is the set of propositions that constitute the lottery paradox (i.e. appropriate instances of $\lambda_1$ and $\lambda_2$):

P1: The operant and appropriate norm in situation $Q$ involving $O$'s attitude $\beta$ towards proposition $\Lambda_L$ is $y$.
P2: $y$ is not truth.
P3: Truth is the norm of belief.

Therefore, $O$'s attitude $\beta$ towards proposition $\Lambda_L$ in situation $Q$ is not belief.

As in the previous teacher's case, with respect to subjective closure under material implication and Cohen acceptance, we can see that:

P1: In any situation $x$ involving $S$'s Cohen acceptance of a proposition $p$, $S$'s commitment is governed by subjective closure under material implication.
P2: $O$'s commitment to $\Lambda_L$ in situation $P$ is not governed by subjective closure under material implication.

Therefore, $O$'s attitude towards proposition $\Lambda_L$ in situation $Q$ is not Cohen acceptance.

So, our teacher (qua researcher in this context) does not believe the propositions that make up $\Lambda_L$ and he does not fully accept them either. However, in this case, neither does he accept them in the weak sense that we saw in the pedagogical case nor does he entertain them for merely pragmatic reasons, but rather as an intellectual problem. Nevertheless, his acceptance is still limited. In simply considering the paradox, it is not reasonable to suppose that he is committed to the classical implications of his acceptance of those propositions. So, his commitment to them is then limited, but it is not weak because the elements of $\Lambda_L$ are entertained on the basis of their plausibility.[11] As a result, the two kinds of limited acceptance provide a solution to the problem of paradox-induced explosion that allows not only the pragmatic consideration of paradoxes, but also the purely intellectual consideration of paradoxes as well.

## The logic of acceptance

At this point, we are faced with a deeply interesting matter that has tangible implications for logic itself. While considerable work has been done concerning the logics of belief and knowledge under the rubric of epistemic logic, there has been virtually no work done on the logic of acceptance because these forms of commitment have been almost wholly overlooked by logicians and epistemologists alike.[12] The question then is: what is the logic of acceptance? A number of things can be said at this point in response to this question. But first, I want to make it clear that I am neither advocating nor presenting a fully worked out system of logic for any of the forms of acceptance introduced here that would correspond to the extant systems for belief or knowledge as understood in epistemic logic. Rather, I am simply indicating some possible approaches that look promising and/or are likely to be correct.

The first and most obvious suggestion to be made here – given what has already been said above – is that the underlying logic of limited acceptance might be a form of relevant logic, and this is significantly different from the standard assumption of classical logic as the underlying logic for both knowledge and belief in epistemic logic (see Mares & Meyer, 2001; Routley & Routley, 1972). But, since acceptance is very much like the epistemic operators *B* and *K* in epistemic logic, what we need is to develop a corresponding weaker system of logic for acceptance that involves the main features of acceptance noted above and one that was illustrated in various examples entertained here. This is a worthy project to pursue because the logic of acceptance allows us to defuse the problem of paradox-induced explosion, to explain a wide variety of commitments that are weaker than belief and, more specifically, to potentially explain the type of commitments we have to logical systems (particularly to non-classical logics). The introduction of these concepts and their accompanying logic(s) may also have serious and interesting implications for the construction of more realistic artificially intelligent systems.

## Conclusion

All of this then shows us that there is a possible way to address paradoxes in cases where there need not be any worries about paradox-induced explosion, but that this is due entirely to the fact that paradoxes are entertained as a sort of voluntary intellectual exercise that do not involve evidential matters, closure under classical implication or commitments to truth. So in many cases, it looks like the *problem of paradox-induced explosion* might be avoided in this purely epistemic manner, without having to resort either to non-classical logics for beliefs (e.g. relevance logics or paraconsistent logics) or to the denial of the standard closure principle for beliefs. This is a desirable sort of solution to the problem of paradox-induced explosion as it satisfies the principle of minimum mutilation in preserving both classical logic and the closure principle for belief. This solution is based on appealing to the concepts of limited (weak and strong) acceptance. Appealing to these concepts in order to defuse the problem of paradox-induced explosion is not *ad hoc*, however, and these attitudes seem to be present in many contexts where intelligent agents employ propositions as matters of convention or as intellectual exercises. This appears as if it might include cases where we use refuted but approximately true theories, where we are engaged in fictional performances and possibly cases where we adopt alternative logics.

Two important extensions of this suggestion are then, first, to examine whether folk psychology already includes acknowledgement of the existence of these sorts of weak propositional attitudes or if they are concepts that we *should* introduce as refinements of folk psychology based on their pragmatic utility and/or explanatory power. Second, the results here suggest that a truly intelligent artificial agent should be able to cope with paradoxes. This, in turn, suggests that there may be additional constraints on our conceptions of human intelligence related to the human ability to cope with a paradox. Implementing this suggestion might, however, require designed intelligent systems to be modelled in terms of representational structures that include both beliefs and other weaker propositional commitments such as acceptance. This may further require engineering systems that have voluntary control over at least some of their commitments, as this is also a defining feature of acceptance.



## Notes

1. Compare this account of paradox with those of Rescher (2001) and Olin (2003). The account offered here is broadly compatible with both of these accounts.
2. Kyburg took this paradox to show that we should reject the rule adjunction (see Kyburg, 1997).
3. '*S*' is used throughout this paper as a variable ranging over agents.
4. Compare Rescher (2011) with respect to the idea of the plausibility of beliefs/propositions.
5. Also, accepting *p* does not necessarily mean that *p* is empirically adequate in the sense understood by van Fraassen (1980).
6. See Rescher (2001) for a related concept of plausibility.
7. Given the burgeoning logical pluralism we find in contemporary logic, it is not even clear that the propositions of the various extant systems of logic are necessarily, if at all, true. I am inclined to believe that commitments to logics are matters of convention.
8. Notice that this view may help logical pluralists to provide a foundation for their view. On this reading of the grounding of logical systems, one needs to commit to such a system only up to a degree of acceptance in order to understand (or partially understand) or to use it for some practical purpose. Those logicians who argue that there is only one logic, classical logic, and that all other logics are incomprehensible are then faced with having to deny that we can entertain alternative systems of logic or that we can entertain logical systems with multiple ('conflicting') operators (some of which are non-classical) in the sense of acceptance introduced here. Obviously, I do not think that this manoeuvre is defensible and commitments in the sense of acceptance appear to be widespread. In any case, see Beall and Restall (2006) for a thorough discussion of logical pluralism and Haack (1996) for a discussion of the possibilities of comprehending and justifying alternative logics.
9. This example can, unlike the Newton case, easily be modified in order to illustrate a case of strong limited acceptance. All that we need to add is that the logics in question are plausible and that our logician is committing to them on the basis of their plausibility.
10. Recognition of this point is attributed to a very helpful referee.
11. So this case may then be more like a slightly modified version of our logician case mentioned in an earlier footnote. One might suppose that in the case of the logician, the logician accepts the two logics in the strong sense that they are plausible and that this is the basis of their acceptance.
12. For a good introduction to epistemic logic, see Meyer (2001) and for the classic presentation of work on that topic, see Hintikka (1962).

## References

Beall, J. C., & Restall, G. (2006). *Logical pluralism*. Oxford: Oxford University Press.
Cohen, L. J. (1992). *An essay on belief and acceptance*. Oxford: Clarendon Press.

Haack, S. (1996). *Deviant logic, fuzzy logic*. Chicago, IL: University of Chicago Press.
Hintikka, J. (1962). *Knowledge and belief*. Ithaca, NY: Cornell University Press.
Kline, M. (1980). *Mathematics: The loss of certainty*. Oxford: Oxford University Press.
Kyburg, H. (1961). *Probability and the logic of rational belief*. Middeltown, CT: Wesleyan University Press.
Kyburg, H. (1997). The rule of adjunction and reasonable inference. *The Journal of Philosophy*, *94*, 109–125.
Mares, E., & Meyer, R. (2001). Relevant logics. In L. Goble (Ed.), *The Blackwell guide to philosophical logic* (pp. 280–308). Oxford: Blackwell.
Meyer, J. (2001). Epistemic logic. In L. Goble (Ed.), *The Blackwell guide to philosophical logic* (pp. 183–202). Oxford: Blackwell.
Olin, D. (2003). *Paradox*. Montreal: McGill-Queen's University Press.
Rescher, N. (2001). *Paradoxes*. Chicago, IL: Open Court.
Routley, R., & Routley, V. (1972). Semantics of first degree entailment. *Nous*, *6*, 335–359.
Sorensen, R. (2003). *A brief history of the paradox*. Oxford: Oxford University Press.
van Fraassen, B. (1980). *The scientific image*. Oxford: Clarendon Press.