Vagueness and Truth

Mark Colyvan*

Abstract

In philosophy of logic and elsewhere, it is generally thought that similar problems should be solved by similar means. This advice is sometimes elevated to the status of a principle: the principle of uniform solution. In this paper I will explore the question of what counts as a similar problem and consider reasons for subscribing to the principle of uniform solution.

1 Introducing the Principle of Uniform Solution

It would be very odd to give different responses to two paradoxes depending on minor, seemingly-irrelevant details of their presentation. For example, it would be unacceptable to deal with the paradox of the heap by invoking a multi-valued logic, $Ł_\infty$, say, and yet, when faced with the paradox of the bald man, invoke a supervaluational logic. Clearly these two paradoxes are of a kind—they are both instances of the sorites paradox. And whether the sorites paradox is couched in terms of heaps and grains of sand, or in terms of baldness and the number of hairs on the head, *it is essentially the same problem* and therefore must be solved by the same means. More generally, we might suggest that similar paradoxes should be resolved by similar means. This advice is sometimes elevated to the status of a principle, which usually goes by the name of *the principle of uniform solution*. This principle and its motivation will occupy us for much of the discussion in this paper. In particular, I will defend a rather general form of this principle. I will argue that two paradoxes can be thought to be of the same kind because (at a suitable level of abstraction) they share a similar internal structure, or because of external considerations such as the relationships of the paradoxes in question to other paradoxes in the vicinity, or the way they respond to proposed solutions. I will then use this reading of the principle of uniform solution to make a case for the sorites and the liar paradox being of a kind.

*Department of Philosophy, University of Sydney, Sydney, NSW, 2006, Australia. *Email:* mcolyvan@usyd.edu.au

First, as I've already indicated, it is uncontroversial that we all subscribe to some version of the principle of uniform solution.[1] There may well be disagreement about the application, scope and details of formulation of the principle, but there is no doubt that some version is agreed to by everyone. Such agreement is rare in philosophy, so it is worth pausing over why there should be such agreement here. I take it that it is not merely a matter of simplicity. That is, I take it that the reason philosophers accept some version of the principle of uniform solution is not motivated by a concern to deal with as many paradoxes as possible with a minimum of resources. There are two reasons why this cannot be the motivation. The first is that this motivation makes no reference to, nor does it require that, the two paradoxes in question be of a kind. Simplicity consideration, if anything, would motivate a principle akin to treating all paradoxes with a minimum number of tools—use a hammer for every task, irrespective of whether all the tasks involve nails. The second reason that this cannot be the motivation for the principle of uniform solution is that simplicity considerations notoriously divide philosophers and it would be extraordinary if the agreement we have over the principle of uniform solution were due to something so controversial as simplicity.

I take the motivation for the principle of uniform solution to be a concern for treating the root of the problem—the disease—not just the symptoms. At least, this is what motivates my subscription to the principle, and I suspect that something like this is the motivation for others as well. The idea is that if you do not bring similar solutions to bear on similar paradoxes, you will be merely hiding the problem—masking the symptoms—rather than getting to the heart of the matter and treating the underlying pathology. An example will help. Let's suppose that you are tempted to treat the paradox of the bald man by stipulating that some specific number, n, hairs on the head is the sharp cut-off between bald and not bald. And further suppose that you are tempted to treat the paradox of the heap by denying that there are any heaps. Both paradoxes are solved. But the concern is that you have not fully appreciated the fact that in both cases the vagueness of the predicates in question gives rise to borderline cases and that these are tolerant to incremental creep. Without addressing the underlying pathology—classical logic's apparent inability to handle vague predicates—the problem will re-emerge under a new guise. The principle of uniform solution, thus motivated, seems like a sound methodological principle to ensure against further outbreaks of trouble.

There is still the question of what makes the different instances of the sorites paradox of a kind. Why is the paradox of the heap and the paradox

[1]It is also clear that this principle has broad applications, across all areas of philosophy where paradoxes arise: philosophy of logic, decision theory, ethics, philosophy of space and time and so on. See [5] for a recent implicit appeal to the principle to a family of paradoxes in decision theory.

of the bald man the same paradox? And what is the underlying pathology. Clearly both these paradoxes share some underlying form and the fact that one is about sand and the other about hair is irrelevant. And it's not too difficult to spell out what this shared underlying form is. But the problem gets more serious when we consider other, more difficult cases. Is the liar paradox of a kind with Russell's paradox? Is Curry's paradox of a kind with the liar? Is the totally-ordered (numerical) sorites of a kind with the non-totally ordered (multi-dimensional) sorites? And the case I will turn to in the last section of this paper: is the sorites of a kind with the liar? In order to make a start on these questions, let's revisit an old debate between Russell and Ramsey about the liar and Russell's paradoxes.

2 The Liar and Russell's Paradoxes

Russell and Ramsey disagreed about whether the liar and Russell's paradox were of a kind. Ramsey [20] stressed that the liar, along with Grelling's paradox, Berry's paradox and other definability paradoxes, were *linguistic*, in that they relied upon thought or language. He argued that the "fault" was ambiguity in key terms such as 'defines'. These paradoxes were to be contrasted with set theory paradoxes such as Russell's paradox and the Burali-Forti paradox. Letting his logicism come to the fore, Ramsey saw the paradoxes of set theory as only relying on logical notions (such as set membership and number).[2] He thus saw two quite different kinds of paradox here: linguistic and logical/mathematical. Russell [21], on the other hand, focussed more on the underlying structure of the paradoxes and saw them all as paradoxes of impredicativity. And we might add as a further consideration in favour of Russell, both the set theory paradoxes and the liar paradoxes admit extended forms (cyclical liar and cyclical set theory paradoxes) that appear to be generated by the same formal feature: essentially the violation of the vicious circle principle [9].[3]

So on one way of looking at the liar and Russell's paradox, we see a difference: one is about truth the other about set membership. But on another way of looking at them, they are of a kind: they both depend on self reference. What is the correct way to classify these paradoxes? Who won the day, Ramsey or Russell? It's fair to say that although Ramsey's classification was very influential for a considerable time, more recently it has fallen out of favour. For what it's worth, I'm inclined to agree with Russell and not with Ramsey here, but the matter is neither trivial nor settled. Agreeing with Russell is one thing, but putting your finger on why the liar and the set theory paradoxes are of a kind is not possible

[2] Moreover, Ramsey's logicism prevented him from recognising another possible distinction: between mathematics and logic.

[3] I will make more of this point in a later section.

without a substantial theory of similarity of paradoxes. Neither Russell nor Ramsey were attempting anything so ambitious as this, yet each, in their own way, was providing a partial theory about the similarity or dissimilarity of the two paradoxes under discussion. What was particularly interesting about Russell's approach, at least, is that he can be seen to be looking for the underlying mechanism and was considering structural similarities rather than the more superficial features of the paradoxes.[4] And Russell did not need to invoke logicism or any other substantial philosophical theory to defend his view.[5]

In any case, let's run with this (no doubt not entirely accurate) caricature of Russell and Ramsey on the liar and set theory paradoxes. My suggestion is that Russell was more interested in what made the paradoxes tick, whereas Ramsey was more concerned with the details of their construction (whether they were made from linguistic and belief terms or from terms from set theory). What is appealing about determining what makes the two paradoxes tick is that once this is known, a more systematic, unified treatment can be applied. We won't be merely treating the symptoms, we will know what the real problem is. A brief look at another unifier, Graham Priest, will help to see the benefits of getting at the root of the problem.

3 The Inclosure Schema

Graham Priest [18, p. 134], following a suggestion from Russell, has proposed a formal schema that can be used to help classify paradoxes: *the Inclosure Schema*.

Inclosure Schema: There are two properties φ and ψ, and a function δ such that

(1) $\Omega = \{y|\varphi(y)\}$ exists, and $\psi(\Omega)$

(2) if x is a subset of Ω such that $\psi(x)$, then

  (a) $\delta(x) \notin x$, and

  (b) $\delta(x) \in \Omega$.

Priest sees this schema as a paradox litmus test: if the paradox in question can be made to conform to this schema, then it is of a kind with others

[4] But in fairness to Ramsey, he probably did not see the fact that the liar involved truth and Russell's paradox involved set membership as superficial at all. Indeed, he may well have thought that this difference resulted in two quite different mechanisms.

[5] Again, in fairness, to Ramsey, it is easy to see reliance on logicism as a defect of Ramsey's account, but at the time, logicism was a real contender for *the* philosophical account of mathematics. Invoking such an account to do some work in the classification of paradoxes is perfectly understandable.

that satisfy the schema. It should not be surprising that according to this way of carving up the territory, many of the usual suspects turn out to be of a kind: Burali-Forti Paradox, Russell's Paradox, the Liar Paradox, Liar Cycle Paradoxes, Berry's Paradox, and Grelling's Paradox, to name a few. But interestingly, one is left out in the cold (Curry's Paradox) and one other has its status in dispute (Yablo's Paradox). Let's look at Yablo's paradox in a little more detail, because the debate over it is instructive.

Yablo's paradox also comes in both set- and truth-theoretic versions. Here is the truth-theoretic version:

($s1$) For all $k > 1$, sk is not true.

($s2$) For all $k > 2$, sk is not true.

($s3$) For all $k > 3$, sk is not true.

. . .

. . .

. . .

(sn) For all $k > n$, sk is not true.

. . .

. . .

. . .

The above infinite list is paradoxical and, on the face of it at least, there is no self-reference or cycles of reference. *But* the paradox is definitely liar like. Indeed, it is a kind of limiting case of longer and longer liar cycles. Why this paradox is so important is that if, as Stephen Yablo [25] originally suggested,[6] it is non self-referential yet liar like, then self reference does not seem to be the culprit after all. The source of the trouble in the so-called paradoxes of self-reference must be elsewhere. Moreover, if Yablo's paradox is not circular, it fails to fit the inclosure schema and thus fails to be classified as a kind with the liar. But everyone agrees that this paradox *is* liar like. Some such as Priest [16] and Beall [1] suggest that, despite initial appearances, the paradox is self-referential because a fixed-point construction is required in the derivation of the paradoxical conclusion. They argue that the Yablo list is self-referential in the sense that the derivation of the paradoxical conclusion invokes a fixed-point construction.

There are many interesting points arising from this debate. First, those who argue that Yablo's paradox is in fact circular, do so by invoking a quite

[6] And more recently supported by Roy Sorensen [23] and Bueno and Colyvan [3].

different notion of circularity than the more intuitive notion, appealed to in liar cycles. After all, "requiring a fixed-point construction in the derivation of the paradoxical conclusion" is quite different from "refers to another sentence that referees to ... that refers to the first sentence". I'm not claiming that thinking of circularity in terms of fixed-point constructions is illegitimate. Far from it. This seems a perfectly reasonable sense of circularity (or self reference). I am just noting that the inclosure schema is not doing all the work here; we can do a little massaging of terms like 'circularity'. But perhaps most interesting of all are the reasons for thinking that even when the structural features of the paradox (whether it is circular and whether it satisfies the inclosure schema) are in dispute, there can be agreement about whether Yablo's paradox is liar like. It is worth spelling out some of the reasons for this agreement.

1. Yablo's paradox is very naturally seen as the limiting case of longer and longer liar cycles. That is, it bears a certain close relationship with paradoxes that are liar like.

2. Yablo's paradox, like the liar paradox, has a set-theoretic analogue.[7]

3. Yablo's paradox, like the liar, comes in both unstrengthened and strengthened forms. That is, Yablo's paradox can be constructed in terms of falsity or in terms of untruth (respectively).

4. And last, but not least, both the liar and Yablo's paradox use the same basic machinery (i.e. truth, negation, T-schema, etc.).

These points of contact between the liar and Yablo's paradox suggest that features other than internal structural features (as displayed by the inclosure schema) are playing a role. We also saw this in relation to Russell's paradox and the liar, although there it was less clear because the two paradoxes in question also shared a common structure (as revealed by the inclosure schema). In the case of Russell's paradox and the liar, the extra-structural feature was that they both admitted extended forms. That is, the two paradoxes had similar, nearby neighbours. There is another point of similarity between all four of these paradoxes: they admit the same kind of solutions. I discuss this issue further in the next section. The important point to note for now is that although the inclosure schema (or other ways of revealing internal structural similarities) is clearly relevant, it cannot be the whole story. There is more to paradox classification than recognising similarity of structure.[8]

[7]Whatever you think about the Russell–Ramsey debate over the liar and Russell's paradox, it is clear that Russell's paradox is at the very least a set-theoretic analogue of the liar paradox.

[8]There is also the substantial issue of the level of abstraction at which the structural similarities are revealed. See Smith [22] and Priest [15, 17] on this issue.

4 What Comes First, the Diagnosis or the Treatment?

Let me return to the medical analogy that I used to motivate the principle of uniform solution: we want to treat the disease and not merely the symptoms. The idea is that we don't want to employ different solutions for what are essentially the same problem, for otherwise we are missing the "mechanism".[9] But sometimes we cannot find the mechanism before we consider treatments. Some psychiatric disorders, I'm told, are like this. The mechanisms are complex and often inaccessible, but classifying the diseases in terms of the treatments they respond to is still helpful. Indeed, it might be argued that insofar as a classification by treatments works, it works because there is a common mechanism and the treatment is latching onto that mechanism.

To give one rather dramatic example, before the days of quick, reliable, and readily-available brain imaging, there were two quite different types of acute cerebral disorders that presented identically (or near enough) clinically. Consider a patient with an acute onset of weakness down one side of their body, slurred speech, and relaxed facial muscles on the same side as the weakness. These are classic stroke symptoms, but, as it turns out, there are two quite different mechanism for this. One is a thrombotic event (a piece of clot that lodges in an inter-cerebral artery and blocks blood supply to all points distal to the block. The second mechanism is a burst blood vessel in the brain that results in an inter-cerebral bleed that puts pressure on the nearby areas of the brain (and cuts off blood supply to those regions). There is a very reliable test to tell these two apart: administer an anticoagulant such as heparin. If the stroke were due to the first mechanism (the thrombotic event), the patient would often make a dramatic recovery (if caught early enough) after the heparin. On the other hand, if the stroke were due to the second mechanism (the bleed), then the patient would see a rapid decline after heparin, and possibly die. There are, of course, less extreme cases of this type of diagnosis via cure (or failure to cure), but this one is useful to make the point because the underlying mechanisms are quite straight forward.

That's all well and good for medicine, but what's that got to do with classifying paradoxes? Well, we might take a leaf out of medicine's book and consider classifying paradoxes by the methods they respond to (and fail to respond to). For example, here is a reason why we might consider the liar and the set-theory paradoxes as alike: they both respond equally well to hierarchical treatments. They also both respond to three-valued

[9] In the cases of diseases "mechanism" talk is legitimate, but of course in the logical and set-theory paradoxes, such causal language is out of place. But still there is the thought that there is something that is at the root of the paradox—that which makes it tick (to invoke another analogy).

approaches (partial set membership and 3-valued logics) equally as well. Notice that I'm not claiming that either the hierarchical approach or the three-valued approach are genuine solutions; it's just that in each case these approaches do about as well, and face much the same problems. We could then suggest that the reason for this is that the treatments are latching onto, partially latching onto, or completely missing (as the case may be) the underlying mechanisms. So here we have the beginnings of a classification of the underlying problem in terms of the cure, partial cure, or failed cure (as the case may be).

But recall that we are trying to justify the principle of uniform solution. In the context of this justification, isn't classifying paradoxes by their solutions question begging or vacuous? After all, the suggestion seems to amount to something like: paradox A and paradox B are the same paradox because they admit the same solution, therefore we should use the same solution on them. This vacuousness objection would be a fair criticism of simply using *successful* solutions as a means of classifying paradoxes, but that is not what I'm advocating here. Rather, I am advocating that we should use treatments—*successful, partially successful, and unsuccessful*—as a tool in the classification of paradoxes. Nor am I suggesting that this should be the only tool. I've already suggested that Priest's inclosure schema is a very useful way to reveal internal structural similarities and I also mentioned that the relationship to nearby paradoxes can be revealing. Indeed, the latter is closely related to the present suggestion of classifying by treatment. In the next section I will examine the connections between these two proposals.

5 The Location of Paradoxes in "Dialectic Space"

So far I have argued that there are a number of methods we might use to show that two paradoxes are essentially the same. Two of these, classification via response to treatments and classification via near neighbours, are related. Let's return to the (simple) liar and Yabo's paradoxes, each couched in terms of falsity. I have already noted that one reason for thinking that they are essentially the same (or at least of a kind) is that they both have nearby neighbours—strengthened liar and strengthened Yablo paradoxes, couched in terms of untruth. But now consider the usual line of reasoning that relates the simple versions and the strengthened versions.

Consider a truth-value gap approach to both the simple liar and the simple Yablo paradox. According to this approach, the paradoxical sentences in question (the liar sentence and every sentence in the Yablo list) fail to take a truth value, and we invoke a suitable logic to support such claims (such as K_3). But, as is well known, this gappy approach suffers serious revenge problems at the hands of the strengthened versions of the paradoxes in question. I want to make it clear that I am not claiming that gappy solutions do not work (at least I'm not claiming that *here*); I am only pointing out that in both the liar and Yablo's paradoxes, there are nearby strengthened versions that hit the gappy solution very hard, and in much the same way. Now consider glutty approaches to the two paradoxes in question. According to this approach, the sentences in question are both true and false, and we invoke a suitable logic to support such claims (such as LP). As is also well known, these glutty solutions perform better when confronted with strengthened versions of the paradoxes.

What this suggests is that although the relationship of the paradoxes to their nearby neighbours is significant, the dialectic connecting the various paradoxes is also important. The way the nearby neighbours respond or fail to respond to proposed treatments tells us something about what makes the whole family tick. Indeed, the very reasons for introducing the neighbours is important. In the cases just discussed, the neighbouring strengthened paradoxes were introduced to raise problems for gappy approaches. So I am suggesting that we do not just consider the internal structure of the paradoxes—although that is undoubtedly important—we also consider the external relationships—the relationships to other nearby paradoxes. Moreover, I'm suggesting that we consider these latter relationships in terms of the dialectic that connects them all—the successes, partial successes, and failures of various treatments and the moves and counter-moves in the debate that ensues.

This proposal is not so neat and tidy as Priest's elegant inclosure schema. Be that as it may, the added complexity is needed. For a start, the inclosure schema is at best limited in its application; it is only designed to classify paradoxes into two categories: those that fit the schema (the liar-like ones) and those that don't. What about further classification among those that do not fit the schema? Priest is certainly not suggesting that they should all yield to the same treatment; he is only suggesting that those that fit the schema should be treated similarly. But we can reasonably ask about the similarity of paradoxes that do not fit the schema. Is the Prisoners Dilemma a Newcomb problem [13]? Is the two envelope paradox a St. Petersburg paradox [7]? Are the St. Petersburg and the Pasadena Paradoxes of a kind [4, 14]? And, as we will consider in the next section, are the sorites and the liar of a kind [2, 8, 11, 12, 24]? To answer these questions we need something more discriminating than the inclosure schema. This is not a failing of the inclosure schema, of course. It was never intended to answer such question. To answer these questions we need a more general framework for tackling the questions about similarity and identity of paradox. The proposal I've sketched here is a first step towards such a framework.

6 The Liar and the Sorites

It is not obvious that the liar and the sorites share structural similarity,[10] there are other considerations that push for their similarity. The first point of similarity is that both paradoxes have strengthened forms that raise problems for gappy solutions. So, as is well known in the case of the liar, if you are tempted to advance a gappy solution, whereby the liar sentence is neither true nor false, you face revenge problems in the form of the strengthened liar. The trick here is to take the tripartite division the gap theorist advances—true, false and neither—and collapse two of the categories and put pressure on the new bipartite division. The usual way this is done is by defining a new predicate, *untrue*, whereby a sentence is untrue iff it is either false or gappy. The revenge problem arises in the form of the strengthened liar sentence: *This sentence is untrue*. Gaps may solve the original liar but they fall foul of the corresponding revenge problem.[11] Moreover, it is agreed by everyone that the strengthened liar is of the same kind as the original liar and so wheeling in a different solution for the strengthened form of the paradox is not a viable option (for reasons of uniform solution).

Now notice how so-called higher-order vagueness and the related higher-order sorites can be seen as a kind of revenge problem for gap approaches to the sorites. Recall that the gap approach to the sorites posits a third category—neither true nor false—to accommodate all the borderline cases. The higher-order problem involves the same trick as sketched above: the tripartite division is collapsed to a bipartite one via a new predicate. In this case, the usual predicate is *determinately a heap*, where a heap is determinately a heap iff it is not determinately a non-heap and not a borderline case. Now the higher-order sorites is rolled out and we find that the boundary between the categories 'determinately a heap' and the (first-order) 'borderline case' also allows incremental creep. The dialectic structure is the same as with the strengthened liar, where we think of the higher-order sorites as a revenge problem, analogous with the liar revenge problems.

Next it is worth noting that, in both the case of the liar and the sorites, glutty solutions fare better with regard to these revenge problems. But in both cases there is an intuitive revenge problem in the vicinity that would do to gluts what the above revenge problems does to gaps. Take the liar first. Intuitively the predicate in question needs to collapse the three categories of true (only), both true and false, and false (only), to two. The predicate in question would appear to be "true-and-true-only", with the corresponding strengthened liar sentence: *This sentence is not true-and-true-only*. But here's the problem for this line of attack on the glutty approach: this sen-

[10]Although, see [6] for some reason to suspect that a continuous version of the sorites is structurally similar to the liar.

[11]Although I take such revenge problems to be very serious for gap accounts, I am not suggesting that it's game over for gappy accounts.

tence is also paradoxical so is both true (only), and not true-and-true-only. The revenge problem one is grasping for can be formulated, but the sentence in question doesn't convey the meaning you would like. Some have argued that this shows that the glutty solution is somewhat unsatisfying, in that it blocks revenge by what appears to be a trick akin to not allowing the opponent the expressive resources to formulate the relevant revenge problem.[12] I won't enter into that issue here. It is sufficient to note that the intuitive revenge problem for the glut solution does not deliver the serious blow that the analogous revenge problem does for the gap solution and that this lack of revenge is for a very specific reason.

Now return to the sorites. The situation here is analogous. We've already seen how higher-order sorites can be seen as a revenge problem for gappy approaches. But what about a revenge problem for glutty approached? Here the new predicate in question would need to be "true-and-true-only a heap" and the strengthened sorites series would be constructed in terms of this predicate. But as with the liar, there is no problem allowing that borderline sentences are both true-and-true-only and also not true-and-true-only. These hypercontradictions block the revenge problem in the case of the liar, and they do the same for analogous strategies in the case of the sorites. This is not to say that the glut approach solves the sorites. There is a lot more work to be done on spelling out the details of the glutty solution to higher-order vagueness. All I am claiming for present purposes is that, just as with the liar, revenge problems hit the gappy approach but the corresponding glutty revenge problems lack any bite. Moreover, they lack bite for the same reason—because of hypercontradictions.

I have argued that the liar and the sorites both have nearby neighbours in the form of strengthened forms—the strengthened liar and the higher-order sorites—that play similar roles in the dialectic against gaps (and for gluts). That is, there is a case to be made for the liar and the sorites occupying similar positions in their respective "dialectical spaces". Also, it seems that both the liar and the sorites respond to gappy approaches about as well as each other and that gap solutions have serious difficulties with strengthened forms. And, more tentatively, a case can be made that glutty approaches do somewhat better in the face of the relevant revenge problems in both cases.

This suggests that there may be a link between the sorites and the liar via extra-structural considerations. Is this enough to make the case that the liar and the sorites are of a kind and therefore require a uniform solution? So far we have similarity established by two of the three means I outlined in earlier sections. The final one is structural similarity, and that has not yet been established. Is two out of three good enough, or is it structural similarity the one that really matters? Perhaps the other two (position in dialectic space and diagnosis by cure) are mere signs of underlying structural

[12]For instance, Daniel Nolan has made this point in conversation.

similarity? These are interesting questions that are difficult to get much traction on at the present time. But whichever way the answers to these questions fall, I take it that there are some points of contact between the liar and the sorites and that these points of contact are sufficient to warrant further investigation into possible structural similarities[13] and to at least keep an open mind about a uniform treatment of both the paradoxes of vagueness and truth.[14]

References

[1] Beall, JC 'Is Yablo's Paradox Non-Circular?' *Analysis* 61 (2001), 176–187.

[2] Beall, JC and Colyvan, M. 'Heaps of Gluts and Hyde-ing the Sorites', *Mind*, 110 (2001), 401–408.

[3] Bueno, O. and Colyvan, M. 'Paradox Without Satisfaction', *Analysis*, 63 (2003), 152–156.

[4] Colyvan, M. 'No Expectations', *Mind*, 115 (2006), 695–702.

[5] Colyvan, M. 'Relative Expectation Theory', *The Journal of Philosophy*, forthcoming.

[6] Chase, J. 'A Continuous Sorites', to appear.

[7] Chalmers, D. 'The St. Petersburg Two Envelope Paradox', *Analysis*, 62 (2002), 155–157.

[8] Field, H.H. 'Semantic Paradoxes and the Paradoxes of Vagueness', in JC Beall and M. Glanzberg (eds.), *Liars and Heaps*, Oxford University Press, Oxford, 2003, pp. 262–311.

[9] Giaquinto, M. *The Search for Certainty: A Philosophical Account of Foundations of Mathematics*. Clarendon Press, Oxford, 2002.

[10] Hyde, D. 'From Heaps and Gaps to Heaps of Gluts', *Mind*, 106 (1997), 641–660.

[11] Hyde, D. 'Reply to Beall and Colyvan', *Mind*, 110 (2001), 409–411.

[12] Hyde, D. 'Are the Sorites and the Liar Paradox of a Kind?', to appear.

[13] Lewis, D.K. 'Prisoners' Dilemma is a Newcomb Problem', in R. Campbell and L. Sowden (eds.), *Paradoxes of Rationality and Cooperation*. University of British Columbia Press, Vancouver, 1985.

[13] See [19] for some preliminary work in this direction.

[14] I am grateful to James Chase, Dominic Hyde, Carrie Jenkins, Kristie Miller, Hugh Mellor, Daniel Nolan, Graham Priest, Nick Smith, Roy Sorensen and Zach Weber for correspondence and conversations on the issues dealt with in this paper. I am also indebted to audiences at the Truth and Reality Conference at the University of Otago in January 2007, as well as audiences at Macquarie University, the University of Sydney, the University of Miami, the Society for Exact Philosophy Meeting at the University of British Columbia in May 2007, and the Australasian Association of Philosophy Conference at the University of New England in July 2007. This research was funded by an Australian Research Council Discovery Grant (grant number DP0666020).

[14] Nover, H. and Hájek, A. 'Vexing Expectations', *Mind* 113 (2004), 237–249.

[15] Priest, G. 'The Structure of the Paradoxes of Self Reference', *Mind*, 103 (1994), 25–34.

[16] Priest, G. 'Yablo's Paradox', *Analysis* 57 (1997), 236–242.

[17] Priest, G. 'On the Principle of Uniform Solution: A Reply to Smith', *Mind*, 109 (2000), 123–126.

[18] Priest, G. *Beyond the Limits of Thought*, second edition. Oxford University Press, Oxford, 2002.

[19] Priest, G., Colyvan, M., Smith, N.J.J. and Weber, Z. 'Sorites and Inclosures', to appear.

[20] Ramsey, F.P. 'The Foundations of Mathematics' in D.H. Mellor (ed.) *Ramsey: Philosophical Papers*, Cambridge University Press, Cambridge, 1990 (first published in 1925).

[21] Russell, B. *The Principles of Mathematics*, George Allen and Unwin, London, 1903.

[22] Smith, N.J.J. 'The Principle of Uniform Solution (of the Paradoxes of Self-Reference)', *Mind*, 109 (2000), 117–122.

[23] Sorensen, R.A. 'Yablo's Paradox and Kindred Infinite Liars', *Mind*, 107 (1998), 137–155.

[24] Tappenden, J. 'The Liar and Sorites Paradoxes: Toward a Unified Treatment', *Journal of Philosophy*, 90 (1993), 551–577.

[25] Yablo, S. 'Paradox Without Self-Reference', *Analysis* 53 (1993), 251–252.