MICHAEL GLANZBERG

THE LIAR IN CONTEXT

(Received in revised form 6 July 1999)

About twenty-five years ago, Charles Parsons published a paper that began by asking why we still discuss the Liar Paradox. Today, the question seems all the more apt. In the ensuing years we have seen not only Parsons' work (1974), but seminal work of Saul Kripke (1975), and a huge number of other important papers. Too many to list. Surely, one of them must have solved it! In a way, most of them have. Most papers on the Liar Paradox offer some explanation of the behavior of paradoxical sentences, and most also offer some extension for the predicate 'true' that they think is adequate, at least in some restricted setting. But if this is a solution, then the problem we face is far from a lack of solutions; rather, we have an overabundance of conflicting ones. Kripke's work alone provides us with uncountably many different extensions for the truth predicate. Even if it so happens that one of these is a conclusive solution, we do not seem to know which one, or why. We should also ask, given that we are faced with many technically elegant but contradictory views, if they are really all addressing the same problem. What we lack is not solutions, but a way to compare and evaluate the many ones we have.

This is a rather strange situation to be in. One would think that a solution to a philosophical puzzle would involve some explanation of what the puzzle is, and how the proposed solution is successful. This, it seems should be enough to evaluate and choose among proposed solutions. Though some work on the Liar Paradox has been like this, a great deal of it has not. The reason, I think, is that many paradoxes, and the Liar Paradox in particular, have a sort of dual nature. On the one hand, paradoxes can be merely logical puzzles. Solving a paradox, so understood, is a matter of finding some logical device that avoids whatever contradiction the paradox produces. On the other hand, paradoxes can be much more. Beyond posing some

Philosophical Studies **103:** 217–251, 2001.

logical puzzle, they can indicate deep problems of some kind. A solution to a paradox of this sort involves more than just finding an appropriate logical trick. It requires identifying the source of some apparent inconsistency, and explaining why it is merely apparent. Though we may well rely on some logical devices to do so, there is more to the solution. Many discussions of the Liar Paradox, though they grant that the problem is of the deeper kind, pay too little attention to the extra issues that such a problem raises. The result is the situation we have: to many logical devices for avoiding contradiction, and too few explanations of the nature of the problem they are supposed to solve.

My goal in this paper is to try to explain what I think the problem posed by the Liar Paradox is, what an adequate solution to it requires, and what the philosophical interest of the problem is. I shall begin by defending the thesis that the Liar Paradox is a problem about the relation between a proposition and the context in which it is expressed – it is problem of context-dependence. There are already several proposals for solutions along these lines in the literature; but in the interest of explaining what the problem is, and what is needed to properly solve it, I shall concentrate on the question of what we need to know about context to adequately formulate and defend such a solution.[1] I shall argue that the Paradox demands that the domain of truth conditions, out of which propositions may be formed, must itself by context-dependent. For any context, there must be another context relative to which the domain of truth conditions available for forming propositions is strictly wider. This is what I shall call the *resolution problem*: it requires that our representations of content and context be able to provide for indefinitely fine resolution. The resolution problem, I shall maintain, is among the basic theoretical problems posed by the Liar Paradox. Finally, I shall address the philosophical importance of this problem. I shall briefly examine what its metaphysical consequences might be, and I shall argue that it leads us to call into doubt a basic assumption in the philosophy of language, which underlies the semantics/pragmatics distinction. My interest here is in understanding the nature of the Paradox. Though I hope to reveal something about what a solution to the Paradox must be like, I shall not here develop the solution in any detail.

My strategy in this paper is to work by elimination. In section (I), I shall argue that the Liar Paradox should not be understood as primarily a problem about truth, as it often is; but must be understood at least as a problem about truth and truth bearers (propositions). Building on this, in section (II), I shall argue that the problem about truth bearers that generates the Paradox is one of the relation between propositions and the contexts in which they are expressed. To uncover what sort of problem about this relation, I shall consider in section (III) how the Liar Paradox fares in the setting of one of the more far-reaching theories of context and context-dependence. I shall argue that in fact this theory cannot provide a solution to the Paradox, but its failure to do so will help to uncover the nature of the underlying problem. I shall argue that the theory fails to resolve the Paradox precisely because it lacks a way to solve the resolution problem. This, I maintain, is the heart of the Paradox. Finally, in section (IV), I shall consider the wider philosophical significance of the resolution problem. I shall suggest that it need not be understood as suggesting an anti-realist conclusion, as it might be taken to; but that it does force us to reconsider a basic version of the semantics/pragmatics distinction.

I. TRUTH AND THE LIAR PARADOX

To figure out what problem the Liar Paradox presents, we ought to start at the beginning, by asking if there really is any deep or serious problem posed by the Paradox at all. A common reply is that the paradox indicates some fundamental inconsistency in our notion of truth. Alfred Tarski, for example, writes that in colloquial language "it seems to be impossible to define the notion of truth *or even to use this notion in a consistent manner and in agreement with the laws of logic*" (Tarski, 1983, p. 153, emphasis mine). As it stands, it is not clear that this present a terribly important problem. If what Tarski meant is that our use of 'true' is inconsistent, it is not news. We encounter inconsistent uses of the term 'true' all the time. I can produce one for you right now: Everything said by Republicans is not true, but Oliver North truly said that Bush and Reagan were responsible for the Iran-Contra scandal. Saying this may reflect on my political acumen, but it hardly seems to pose any deep philo-

sophical problem about truth. Perhaps what Tarski was trying to say is that we are somehow *forced* into inconsistent uses of 'true'. If the matter really is just use, then this is not an interesting problem either. If I were held hostage alternatively by the Democratic Leadership Council and Lawrence Walsh, I might be forced to say just what I said above, and again this does not seem to pose any significant problem.

It appears that if there is problem posed by the Liar Paradox, it is not one about our use of the term 'true'. Nor would the situation change if we shifted our attention to the commitments or beliefs our uses reflect. We might make more progress by asking generally where inconsistency really can pose an important problem. The answer is relatively simple. Inconsistency may be unfortunate or undesirable when it appears in our utterances or beliefs or commitments, but when it appears in our *theories* it is devastating. Formally speaking, a theory is inconsistent if it contains some sentence and its negation; but more importantly, under some very weak logical assumptions, an inconsistent theory contains every well-formed formula of the language in which it is framed. Stepping back from technical details, what this means is that inconsistent theories are pretty much useless. One plausible suggestion is that a paradox produces an inconsistency in a theory, and so makes it worthless. On this suggestion, if the Liar Paradox is to pose an interesting problem, there must be some theory, which itself has some interesting status, that the Paradox shows us we must abandon.

The question, of course, is *what* theory. The usual way Tarski's remark is understood is this: there is some theory that is supposed to describe our intuitive notion of truth, and the Paradox shows this theory to be inconsistent. I am going to argue that this is not a very satisfying answer. To do this, we need a little more detail about the kind of theory the Paradox is supposed to undermine. What should a theory that explicates our intuitive notion of truth tell us? The tradition tends to observe that 'true' is a predicate, let us say for now a predicate of sentences. Accordingly, the philosophically modest goal is set of explaining the behavior of our ordinary predicate. Indeed, the goal is often taken to be only that of developing a theory that suffices to fix the extension of the predicate 'true'.

Tarski observed that a theory will meet this goal if it suffices to prove all instances of the schema:

‘*s*’ is true if and only if *s*,

where ‘*s*’ is replaced by a sentence. (This is Tarski’s famous Convention T.) Tarski’s observation suggests what I shall call the *naive theory of truth*. The interesting component of the theory is the axiom schema:

(N) T‘*s*’ ↔ *s*.

(Technically, we must think of this as being added to some theory of arithmetic or syntax strong enough to prove the diagonal lemma. I shall suppress all these details, though I shall make use of the diagonal lemma.) The Liar Paradox shows the naive theory of truth to be inconsistent. The Liar Sentence is a sentence that says of itself that it is not true, e.g. ‘this sentence is not true’. Formally, this is a fixed point of the predicate ‘¬T’, i.e. a sentence *l* such that we have *l* ↔ ¬T‘*l*’. We normally can produce such sentences (using the diagonal lemma), and as I said, I shall skip the details. Assuming we have one, we just observe that from (N) we then get ¬T‘*l*’ ↔ T‘*l*’. Contradiction![2]

So, one reasonable way to understand the Liar Paradox might be to note that because of it the naive theory of truth is inconsistent. Hence, it might be said, what is needed for a solution to the Paradox is a consistent theory of truth. There is no question that this is a problem, and very difficult one. However, I shall argue that it is a misunderstanding of the Liar Paradox to think that this is really the basic problem it poses. The reason is brought out by the fact that the Liar Paradox still can be a difficulty even for consistent theories of truth. Constructing such theories is a matter of some technical delicacy, so I shall have to be sketchy.

On idea for constructing a consistent theory of truth, derived from Kripke, is to allow for some of the effect of truth-value gaps in a certain way. Truth is construed as a *partial* predicate. The basic components of such a theory, call it P, can be:

(P1) The four inference rules: T‘*s*’/*s*, *s*/T‘*s*’, ¬T‘*s*’/¬*s*, ¬*s*/ ¬T‘*s*’.

We can then go on to define falsehood as the truth of a negation by adding:

(P2) F‘s’ $\leftrightarrow$ T‘$\neg s$’.[3]

It is worth noting that what is important here for achieving consistency is restricting the kinds of inferences that can be made with truth. Adding the axiom schema T‘s’ $\vee$ F‘s’ retains consistency, but neither T‘s’ $\rightarrow$ s nor $s \rightarrow$ T‘s’ can be consistently added. In fact, we cannot even add T‘s’ $\rightarrow$ T‘T‘s’’ or T‘T‘s’’ $\rightarrow$ T‘s’. But the partiality idea can still be seen in the background. The idea behind the restricted inference rules we have is that one can only draw the usual inferences governing truth if it is determined by the theory that the sentence in question is true or determined that it is false. Take our Liar Sentence l such that we have $l \leftrightarrow \neg$T‘l’. If we have P $\vdash$ T‘l’ or P $\vdash$ F‘l’, then we can recreate the Liar Paradox to show the theory inconsistent. If P $\vdash$ T‘l’, then by (P1) P $\vdash$ l. But then P $\vdash$ T‘l’, so P is inconsistent. We can reason similarly for P $\vdash$ F‘l’. But this just shows P $\nvdash$ T‘l’ and P $\nvdash$ F‘l’ (assuming P is consistent). If we add T‘s’ $\vee$ F‘s’, we get a theory for which the disjunction property obviously fails, but this is not a problem. To see intuitively why the theory as it stands is consistent, observe that unless P proves something about the truth-status of l, we cannot really make any interesting inference from $l \leftrightarrow \neg$T‘l’. In fact, with the theory as it stands we cannot really make any interesting inferences with truth at all. (Once we make clear to what theory the axioms are added, we find we get a conservative extension of it.) A real proof is a decidedly technical matter, but not very difficult.[4]

We may grant that P is a consistent theory of truth. Though the Liar Paradox cannot be used to show it inconsistent, it can be used to create another sort of difficulty. We can reason as follows. According to P, l is neither true nor false, i.e. P $\nvdash$ T‘l’ and P $\nvdash$ F‘l’. Insofar as P is supposed to capture our notion of truth, this tells us that l is not true. But this is just what l says (we have P $\vdash$ ($l \leftrightarrow \neg$T‘l’)), so it seems that according to P, l is true after all. We are back in contradiction. This sort of reasoning, and the problem it poses, are called the *Strengthened Liar*. It does not show that P is inconsistent, as the argument takes place outside the theory. It starts by noting the P $\nvdash$ T‘l’. From there, it goes on to claim $\neg$T‘l’, which formally speaking is not acceptable. But formal notions aside, it appears perfectly reasonable. When we observe that P $\nvdash$ T‘l’, we do not merely note that P is incomplete (as it surely is) and happens

not to prove something in which we are interested. The entire point of constructing the partial theory in the first place was to arrange it so that P $\nvdash$ T‘*l*’ and P $\nvdash$ F‘*l*’. Thus, we may conclude, though from outside the theory, that according to P, *or any consistent extension of it*, we do not have *l* true. As P is supposed to be the theory of our notion of truth, this appears to be the conclusion that *l* is not true, i.e. ¬T‘*l*’. To avoid this, we would need to appeal to two different kinds of not being true, one corresponding to P $\nvdash$ T‘*l*’ (and the same for any consistent extension) and one corresponding to ¬T‘*l*’ simpliciter. Yet the theory that is supposed to capture our notion of truth does not give us two different kinds – as far as we can learn from it, not being true is just not being true.

This reasoning cannot be carried out in P, and so the consistency of P is preserved. But it becomes a problem when we attempt to apply the theory. Recall, we are considering a suggestion that a theory of truth like P is supposed to explain our notion of truth, at least insofar as it is supposed to determine the extensional behavior of the predicate ‘true’. It is supposed to do so by combining with other things we know, presumably about the way things are, to tell us what is true. For example, we know the fact that grass is green. Our theory then tells us that we may conclude that ‘grass is green’ is true. Something similar happens when we come to the Liar Sentence. We do not, in this case, appeal to some empirical facts, but all the same we use the theory, plus some other information, to come to a conclusion about what sentences are true. We know P $\nvdash$ T‘*l*’, and that the same holds for any consistent extension of it. It appears we have concluded, by using our theory of truth, that the Liar Sentence does not come out true. But now we have our problem: it seems we must conclude that the Liar Sentence is not true. Though this is to ignore the P $\nvdash$ *s* versus P $\vdash$ ¬*s* distinction, it seems we have no choice but to do so in the way we understand what the theory is telling us. To respect the distinction, we would need to have some understanding of two different ways of not being true, and the theory that is by supposition to characterize these aspects of our notion of truth, P, gives us none we can use. (N.B. in the theory we are considering, not being true is not the same as being false in the sense of F‘*s*’.) Yet without it, even if the theory is consistent, our applications of it will not be. We use P, combined with other information, to come

to conclusions about what is true, and it seems that among the conclusions we can draw are both that the Liar Sentence is true and not true.

We are faced with a dilemma here. We can conclude either that our theory cannot be *used* consistently (even if it is consistent), or that the theory leaves out some important aspect of different ways sentences can not be true, or we can give up on the idea that the theory really captures our notion of truth at all (and so not take seriously the problem of how we understand or apply it). Any of these is a problem. P is supposed to be a theory which explains our notion of truth, and as such, any of these three options shows it to be inadequate. We may still insist that it captures something important about truth, but we cannot hold that it fully captures this notion.

What is important is that we have come to this conclusion on the basis of the Liar Paradox. Thus, we have used the Liar Paradox to show a theory to be inadequate, even though the theory is consistent. At least, this shows that our initial characterization of the problem of the Liar as that of constructing consistent theories of truth must be wrong. We need to try to find a better way of explaining the problem the Liar poses, and we also need to find a way to talk about truth that allows us to respect the formal distinction we have found ourselves unable to observe here.

We have already seen, in part, what is lacking in our application of the theory of truth. We must be able to distinguish *ways of not being true*. With such a distinction, we might hope to say that the fact that any theory of truth must fail to say T‘*l*’, and so that in one sense *l* is not true, does not tell us that *l* is not true in the sense of ¬T‘*l*’. This distinction, however, resists easy implementation. We might suggest, for instance, that to implement it we think of our theory as embedded in a three-valued logic, with the value **True**, and two other values **False** and **N**. **False** and **N** formally represent different ways of not being true. But this is of no help at all. Indeed, the theory already marked a similar distinction. The theory associates three values to sentences: P $\vdash$ T‘*s*’ corresponds to assigning the value **True**, P $\vdash$ F‘*s*’ corresponds to assigning the value **False**, and P proving neither corresponds to the value **N**. We already observed that keeping P consistent requires it to have non-trivial instances of the behavior corresponding to **N**. so P already has the some impor-

tant features of the three-valued logic, yet we still encounter the Strengthened Liar. It is of no more help to have these distinctions among truth values elsewhere in the theory. We can still observe that any theory like P must not assign the Liar Sentence the value **True**. Insofar as this is the value corresponding to T‘*s*’, we once again have to conclude, by the same reasoning as before, ¬T‘*s*’. As our theory already mimicked important features of three-valued logic, adding a three-valued logic explicitly does not help matters at all.

It is clear that we need some distinction among ways of not being true. But it is now also clear that not any such distinction will do, and indeed we already have a version of it that will not. What is needed is not just some version of the distinction, but one which is suitably robust to be respected in our applications of the theory of truth. The obvious suggestion is that our theories of truth have lacked a substantial notion of truth bearers: objects which are candidate bearers of truth or falsehood which can be distinguished from other objects, including sentences. (Indeed, by design, the theories we have looked at say nothing about objects of any kind – truths or truth bearers.) If our theory gave us the resources to draw such a distinction, it would provide precisely the kind of distinction between ways of not being true we lacked above, which we could rely on in the applications of the theory we found so troubling. If the Liar Sentence were not true because it is a truth bearer, but one which no reasonable theory of truth could declare true, then we could conclude reasonably that it is not true. We could then pass from P $\nvdash$ T‘*l*’ to ¬T‘*l*’, getting the Strengthened Paradox. On the other hand, if our theory tells us that *l* is simply not a truth bearer, we can easily respect this in our application of the theory and so *not* conclude ¬T‘*l*’, When faced with the Liar Sentence, and the observation that according to our theory it is not true, we could *decline* to infer that it is true after all, on the grounds that truth only applies to truth bearers, and it is not one. This would provide a way to observe the P $\nvdash$ T‘*s*’ versus P $\vdash$ ¬T‘*s*’ distinction in our application of the theory.

The moral of all this is that to avoid the Strengthened Liar Paradox, we need to look at theories that say something substantial about the domain of truth bearers. Where does this leave the problem of the Liar? It is obvious by now that it cannot be characterized as the

problem of constructing consistent theories of truth. We now have a hint as to where to look for a better characterization. It is the failure of a theory of truth like P to give a substantial characterization of the domain of truth bearers that leaves it vulnerable to the Strengthened Liar. This is not a statement of the problem, and certainly not anything like a solution to it. It just points out something that is lacking in our initial characterization, and where we should look to correct it. We still must face the task of figuring out what problem of truth and truth bearers is at the root of the phenomena we have so-far observed.

II. PROPOSITIONS AND CONTEXT

So far we have a negative conclusion. I have argued that the Liar Paradox cannot be understood simply as a problem about truth. At the very least, it must be understood as relating to truth and truth bearers. We still lack a clear statement of what problem about truth bearers the Liar presents. It is to this question that we must now turn.

It will be helpful to reconsider the kind of inference we make in producing the Paradox, now with some more attention to truth bearers. But first, it is necessary to say a word about how truth bearers will be treated. This is especially important, as in a moment, I am going to talk about context. It is generally recognized that sentences cannot play the role of truth bearers if we admit any context-dependence, as we would find that the same sentence is both true and false. Our choices are to apply truth to something like pairs of sentences and contexts, or to utterances, or to propositions expressed by utterances. It is, at least, immensely convenient to take the last option: treating truth as applying to propositions. This allows us easily to distinguish between a sentence, the context in which it is uttered, and the truth-evaluable claim that might be made by uttering it. The other options tie context and truth bearer too closely together, and hence make distinguishing among them all the more awkward (not to mention the standard difficulties encountered by utterance theories, like how to handle predications of truth to claims that have not actually been uttered). Many have expressed worries about assuming the existence of such intensional entities as propositions. Though I myself am not worried, let me stress that

my point here is not primarily ontological. To say that a proposition is expressed by an utterance is to say no more than that a truth-evaluable claim is made, which leaves open whether the truth bearer has *separate* ontological status or can be *reduced* to something else. I shall assume that it makes sense to quantify over propositions, so indeed I may take on some ontological commitments. But nothing in the argument to follow depends on precisely what they are, or would be disturbed by reductionist attempts to explain them away.

To proceed, I shall take up, in somewhat abbreviated form, the challenge of the last section. I shall sketch a very minimal theory of truth and truth bearers – a minimal theory of propositions – and consider how the Liar Paradox fares in light of it. It will turn out that to address the Paradox, we need only rely on a few fairly straightforward principles governing propositions, so in the end we will not need to complete the task of developing a really good theory of propositions. I shall argue that the problem posed by the Liar is one about the relation of propositions to contexts. But to reach this conclusion, I shall start with the assumption that context plays no role. I shall then argue that under this assumption, we can make no progress on the Liar.

So, for the moment, let us assume that, somehow, *context plays no role of any kind*. In this case, we may concentrate on the relation between a sentence and the proposition it expresses, if it expresses one. Write 'Exp(s,p)' for 's expresses the proposition p'. As propositions are truth bearers, we may write 'Tp' for 'the proposition p is true'.

We may now state some principles concerning the relation Exp and the predicate T. The first two capture the ideas that there is only one proposition expressed by a sentence (remember, we are assuming no effects of context), and that truth predication is non-opaque:

(U-Exp) $(\mathrm{Exp}(s,p)\wedge\mathrm{Exp}(s,q)) \rightarrow p=q$

(T-Id) $p=q \rightarrow (\mathrm{T}p \leftrightarrow \mathrm{T}q)$.

We also need a principle that looks strikingly familiar:

(T-Prop) $\mathrm{Exp}(s,p) \rightarrow (\mathrm{T}p \leftrightarrow \Sigma)$ for 'Σ' a sentence and 's' a name for 'Σ'.

This is the version for propositions of the T-schema we considered above. Unlike the version we considered in section (I), it has the significant antecedent 'Exp(s,p)'. With this, the principle becomes much weaker, and much harder to challenge. It expresses merely the idea that propositions may be evaluated for truth at the circumstance of utterance. I doubt anything that failed to validate it could count as a reasonable theory of propositions.

Finally, we will require one principle that tells us that in certain cases, we can be sure that an utterance of a sentence does express something. The idea is that if we prove some sentence, on the basis of true premises, via truth-preserving steps, it is itself true. But to be true, it must express a proposition, particularly a true one. So we have:

(Exp-Prov) If s is provable from true premises, then $\exists p(\text{Exp}(s,p))$.

Each of these principles is eminently plausible, and we should expect to find something like them as part of any reasonable theory of truth and propositions. However, they are chosen not because they present a particularly complete or elegant picture of propositions, but because they are important for examining the Liar. We will also need to reformulate our Liar Sentence, in light of our refined approach to truth bearers. It should now be a sentence that says of itself that it does not express a true proposition, or more formally, let l be the sentence:

$$\neg\exists p(\text{Exp}(l,p)\wedge \text{T}p).$$

(I am reusing the name of the Liar Sentence, and so shifting my notation slightly.)

With our theory as it stands, we can reproduce the Paradox. We first argue that if l expresses a proposition, then it is true just in case it is not true. Suppose l does express a proposition, say q. Then if q is true, there is a true proposition expressed by l, contrary to what l says. Hence, q, the proposition expressed by l, is not true. Conversely, if q is not true, then the proposition expressed by l is not true. Hence, appealing to uniqueness, what l says is true, so q is true.

To be a little more careful, we can reason as follows. Suppose there is proposition q such that we have Exp(l,q). Then, we find $Tq \leftrightarrow \neg Tq$, a contradiction.

1. Suppose Tq. Then by (T-Prop), $\neg\exists p(\text{Exp}(l,p)\wedge Tp)$.
2. From (1), and the assumption that Exp(l,q), $\neg Tq$ follows by logic.
3. Suppose $\neg Tq$. By (T-Prop), $\neg l$, i.e. $\exists p(\text{Exp}(l,p)\wedge Tp)$.
4. From (3), (U-Exp), and Exp(l,q), p=q.
5. From (3), (4), and (T-Id), Tq.

One of the benefits of putting the issue as I have here is that we can see plainly that this is not yet a paradox at all. It is a proof that an utterance of l cannot express a proposition (be truth-evaluable, if you like). So our conclusion from this is only:

6. $\neg\exists p(\text{Exp}(l,p))$.

But we can push forward. If an utterance of l cannot express a proposition at all, it cannot express a true proposition. From (6), by logic, we have:

7. $\neg\exists p(\text{Exp}(l,p)\wedge Tp)$.

The problem is that the sentence on line (7) just is l. We have:

8. l.

This conclusion, that an utterance of l does not express a true proposition, is the result of deduction, from no (undischarged) premises. So in drawing this conclusion, we say something true. But observe, this means that in uttering l, I say something true. Thus, an utterance of l expresses a proposition, in fact a true one. This is the principle codified in (Exp-Prov), which gives us:

9. $\exists p(\text{Exp}(l,p))$.

Now we have a problem. From (9), we can return to line (1) and genuinely derive a contradiction. Indeed, we have several contradictions from which to choose. We have proved that an utterance of l both does and does not express a proposition. As we have proved that it does, we know that it expresses something that is both true and not true.

My appeal to propositions (or truth-evaluablility, or some such thing) is by no means a solution to the Paradox. It only helps to

make clear what the real issue is. One way to see this is to notice how the parts of the inference we have just seen correspond to those that I discussed with respect to theories of truth. The inference that shows a theory of truth to be inconsistent appears in (1–5), where it tells us that a proposition expressed by the Liar Sentence is true just in case it is not true. We saw before that we could construct consistent theories, and still be vulnerable to Paradox. By building in the appeal to truth bearers explicitly, we do not even need to detour through theory construction. We can just conclude that nothing is expressed, in (6). But neither the exercise in theory building, nor this conclusion, avoids the paradoxical problem. Above, the thing that makes the theory consistent, not proving T‘*l*’ or F‘*l*’, gave us a way to argue that *l* is true after all, and hence get a paradox. We do exactly the same thing here in (7–9).

Even so, I think that putting things in the terms I have does help us to see what is going on. If there is a contradiction, it is because we have found a sentence such that an utterance of it at one moment appears unable to say anything at all, but then appears to say something true. Indeed, we first *prove* that it cannot express a proposition, and then subsequently prove that it can.

At a sufficiently high level of abstraction, it is not yet clear that we really have a problem. The behavior the Liar Sentence exhibits it just the kind of behavior we expect from certain sorts of context-dependent sentences. On a commonly held view, sentences involving demonstratives behave this way. The sentence ‘that snowplow is loud’ can (normally) only be used to express something in the presence of a snowplow. We can show that an utterance of it fails to say something in a given context by showing that there is no snowplow to be demonstrated, and yet use it, and perhaps argue that we have used it successfully, in another context. We have seen that the Liar Sentence behaves in just this way. Thus, we are forced to conclude that this sentence exhibits some aspects of context-dependence.

I say ‘forced’, and I mean it. We have proved that the Liar Sentence cannot say something, and we have also actually used it successfully. The only way this can be is if it behaves in a context-dependent way. Our proof is based on our little theory of propositions; yet re-examining the theory leads to the same con-

clusion. In the absence of any context-dependence, none of the principles used in the inference from (1) to (9) seem questionable. (T-Id) and (U-Exp) are mere technicalities.[5] (Exp-Prov) just reflects the idea that we can prove things, and so see them to be true. Though principles looking like (T-Prop) are routinely challenged, it must be stressed that (T-Prop) itself, expressed in terms of propositions, is much harder to deny. If we think of propositions as determining truth conditions, for instance, it merely tells us that the truth of an utterance of a sentence is determined by whether the actual circumstance in which the sentence is uttered is among the truth conditions determined. So, it would appear that we have some unquestionable principles that lead to a contradiction. This is unacceptable.

However, this conclusion is reached only under the assumption that context-dependence plays no role. Bearing this in mind, we should re-consider the building blocks of the theory: 'T' and 'Exp'. The logical form of the truth predicate is hardly up for grabs, as we introduced propositions just to be truth bearers. But that of 'Exp' is. If we drop the assumption of no context-dependence, we must see 'Exp' as having an additional argument place, giving a ternary relation Exp(s,c,p) that holds if sentence s, in context c, expresses proposition p. If we do this, we can accommodate the conclusions of the two parts of the inference, first that l does not express a proposition, and then that it expresses a true one. So long as the contexts are different for the two conclusions, this is to have $\neg\exists p$(Exp(l,c,p) and $\exists q$Exp(l,c',q). When c and c' are distinct, this is not a contradiction, and only an equivocation makes it appear as one.

Unless we recognize some role for context-dependence, we cannot avoid the contradictory conclusions our theory of propositions seems to imply. In section (I), I argued that the Liar is a problem about truth and truth bearers, and it could only be avoided by a theory that says something substantial about truth bearers. We can now refine this point. We can at least conclude that the features of truth bearers – propositions – that our theory must address to avoid the Paradox must include something about the relation of propositions to contexts. The Liar Paradox is a problem of context-dependence.

But just to say that we must take context into account, and so 'Exp' must have an extra parameter, is hardly to explain what is

happening with the Liar. Notwithstanding my last argument, the comparison between context-dependent behavior and the behavior of the Liar Sentence appears to be somewhat strained, if not outright absurd. First, the Liar Sentence contains no indexical terms or other devices that we usually think of as generating context-dependence, so it is not clear how it can exhibit any context-dependence at all. To explain away the contradiction by appeal to a context shift, we would have to say the shift takes place between the step of concluding that an utterance of *l* does not express a proposition, and the next step of noting that then it does not express a true one. It hardly seems there could be a context shift there, as all that happens is the drawing of an inference. Perhaps more importantly, the kinds of things we do to show that the Liar Sentence cannot say something, and then that it does, seem to make no reference to the context involved, unlike the kinds of things we might do in determining whether a demonstrative utterance says something.

Where does this leave us? We began by looking at theories of truth, and found that without taking adequate account of truth bearers, we were not able to address the phenomenon presented by the Liar. We then re-cast the Liar within a theory including some principles about truth bearers (propositions). There, we found that the Liar Sentence must behave in a strongly context-dependent way: it must in some contexts fail to express a proposition, and in some contexts not. Yet it appears that there is neither context shift nor context-dependence to be found. This is the problem we face. Two things are required to resolve it: (1) a notion of context such that we can reasonably say a change of context occurs in situations like the inference I discussed a moment ago, and (2) an account of expression in context, so understood, that allows us to explain why the Liar Sentence behaves as it does. It is worth mentioning that the problem of constructing consistent theories of truth is still with us, in a way. It is certainly crucial to the Liar Sentence that it contains a truth predicate, so an account of expression that will do the job will have to say something about the semantics of the term ‘true’. But the significance of the Liar is broader than this.

This is, I believe, an accurate statement of the problem posed by the Liar Paradox, and what is required for a solution to it. Accurate, but not sufficiently informative. It does not really tell us what we are

up against. The notion of context itself is very plastic, and it is well-known that context-dependence can have many sources, sometimes deeply hidden. At least in ordinary parlance, the range of things we are inclined to call matters of context is huge and various; including anything from the referent of an indexical like 'I', to understanding why a joke is funny, to understanding the political significance of an off-hand remark. To note that we must say something about context and context-dependence to resolve the Liar Paradox tells us very little about what sorts of issues are at stake, and what might be needed to resolve them. To get a more revealing characterization of the problem posed by the Liar, we must ask what it is about the general phenomenon of context-dependence that is at work in it.

III. CONTEXT CHANGE

To answer this question, we must see how the problem of the Liar Paradox relates to the theory of context. It will turn out that the problem of the Liar is in fact quite far-reaching for such theories. It will force us to reconsider a basic assumption in philosophy of language, about the kinds of effects context can have on content. In the end, this will amount to questioning at least one form of the familiar semantics/pragmatics distinction.

Let us consider once more the problematic inference. The crucial steps are:

A. The conclusion that an utterance of *l* does not express a proposition.
B. From (A), the conclusion that an utterance of *l* does not express a true proposition.

At (B), we assert *l* and express a proposition, and so run into paradox. I have argued that to avoid the Paradox, we must recognize a context shift between (A) and (B). The issue at hand is to explain how context can shift between (A) and (B), and how *l* can be sensitive to this shift. Above, I raised this as a puzzle. It is not clear how *l* could be sensitive to context at all, nor how the inference that gets us from (A) to (B) could have anything to do with a shift in context. It is, after all, merely a sequence of assertions.

However, in the past several years, a view of context and context change has been developed that might bear on this puzzle.

Robert Stalnaker (1978) and others have argued that quite generally an assertion does have the effect of changing the context for the remainder of a conversation in which it takes place. We might reasonably hope that a theory of this sort could resolve the problem posed by the Liar. It certainly has the right form. If assertion can change context for the remainder of a conversation, it could perhaps change it in such a way as to make the Liar Sentence able to express something. In the end, I am afraid, Stalnaker's ideas do not provide a solution, but the reason they do not will help to reveal the basic issue about context posed by the Liar, so I shall now turn to a brief presentation of Stalnaker's theory, and how the Liar fares in it.

The first component of the theory is a notion of context. Contexts are collections of propositions. For the issues under discussion here, it will be safe to follow Stalnaker in thinking of propositions as sets of possible worlds.[6] Hence, we may represent a context as the set of possible worlds compatible with all the proposition in it, i.e. their intersection. Call this set the *context set*. Stalnaker argues that the propositions that form a context are those that are presupposed, taken to be common ground, by speakers in a conversation at the point of utterance.[7]

The second component is an account of what proposition is expressed by an assertion in a context. The idea, roughly, is that asserting a proposition divides the class of possible worlds into two parts: those in which what is said holds, and those in which it does not. But what class of world? Those that are relevant possibilities at the point in the conversation where the assertion is made. These are just the worlds in the context set. What determines how the relevant possibilities are divided by making an assertion? Again roughly, the meaning of the sentence uttered. In this setting, we may think of the meaning as given by a *character*: a functions from worlds to propositions (sets of worlds). By uttering a sentence (assertively), we express the proposition determined by the intersection of the value of its character on the world of utterance and the context set of the utterance.[8]

This theory has some far reaching consequences. Perhaps the most important for us is that *each assertion made in a conversation changes the context*. What has already been said in a conversation

is normally taken for granted by speakers in just the way that makes it part of the context on Stalnaker's view. Thus, once an assertion is made and accepted, the proposition asserted is added to the context of the conversation from that point on. Suppose we have some context set C, and relative to it an assertion is made, whose content is $p \subset C$. Then the new context set C' is the set of worlds in C that are also in p (consistent with what p says). As $p \subset C$, we have $C' = C \cap p = p$. The resulting context set is exactly the content of the assertion.[9]

This theory appears to offer just the kinds of things I said we need to explain the Liar Paradox. As practically any assertion changes the context according to this theory, we might hope that it could explain the change in context between (A) and (B) in the Liar inference. Unfortunately, the promise of the outline is taken back by the details. The theory does not provide a solution to the Liar. (The suggestion that Stalnaker's theory even might be helpful for the Liar Paradox is mine. Stalnaker never said anything of the kind.)

There are generally two sorts of reasons it does not succeed. The first is that the possibility of the utterance of the Liar Sentence made at (B) expressing something is left a mystery. At (A), we come to the conclusion that an utterance of *l* cannot express a proposition. The context set gets updated to the set of worlds in which this holds. Call it C. But then how can the utterance to *l* in the actual world, at (B), express anything relative to C? If the actual world is in the context set, it appears it cannot. Generally, the actual world need not be in the context set, as speakers can be mistaken about the world, or presuppose things they know to be false. But in this case it does not seem that anything like this occurs. After all, the conclusion at (A) is the conclusion of a *proof*. Unless we can find some flaw in the argument, it seems we must grant that the conclusion is correct, and so the actual world is as we say at (A). But then we must grant that an utterance of *l* cannot express a proposition in this world, at (B), and we are certainly back in contradition.[10]

The second sort of problem arises if we ignore the apparent impossibility of the Liar Sentence's expressing something, and ask what would happen were we to find something for it to express. Suppose there is some proposition expressed by the utterance of *l* in the actual world *a* at (B). Call it *p*. We immediately get into trouble. If *p* relates to the Liar Sentence *l* as we intuitively expect it to, then

we have $a \in p$ just in case $a \notin p$. Suppose $a \in p$. Then the proposition expressed by an utterance of l in a at (B) is true in a. But given what l says, this means that p must be false in a, i.e. $a \notin p$. Conversely, if $a \notin p$, then the proposition expressed by l in a at (B) is not true in a. Thus it is not the case that l expresses a true proposition in a, so what l says is true in a after all, i.e. $a \in p$.

Let us consider this reasoning a little more carefully. We are in a setting where a sentence expresses a proposition only relative to a context, so we need to relativize the relation 'Exp' to context. In the theory we are considering, this is carried out in two ways. First, sentences have characters which determine propositions relative to worlds. Then, the proposition expressed is determined by this proposition and the context set. To introduce some notation, let us write 'ChExp(s,w,p)' for 'the character of s, evaluated at world w, determines proposition p'. We may then think of the expression relation as the relativization of this relation to a context set. So we may write '$\text{Exp}_\text{D}(s,w,p)$' for '$\exists q(\text{ChExp}(s,w,q) \wedge q \cap \text{D} = p)$'. I will assume a version of the unique expression principle for ChExp:

(U-ChExp) $(\text{ChExp}(s,w,p) \wedge \text{ChExp}(s,w,q)) \rightarrow p = q$.[11]

Except for defective contexts, where ChExp is not constant on worlds in the context set, this implies a unique expression principle for utterances. It will also be useful now to work with a relativized notion of truth, written 'T(w,p)' for 'p is true in w', i.e. $w \in p$.

Our Liar Sentence needs once again to be refined. We want a Liar Sentence that still says 'an utterance of this sentence does not express a true proposition', where we now expect this sentence to be context-dependent, both in having a world parameter a and a context set parameter D. So we must consider the sentence l:

$\neg \exists q(\text{Exp}_\text{D}(l,a,q) \wedge \text{T}(a,q))$.

We will need to make one assumption about the character of l. Suppose ChExp(l,a,p). Intuitively, this means p is the set of words in which l (as used in a) comes out true. This idea corresponds to (T-Prop), but with context now playing a significant role, it is not an entirely trivial matter to state it properly. For l in particular, we may note that the only dependence on context comes from 'a' being an argument of 'Exp' and 'T', and Exp being relativized to D. Hence, for this case, we may safely assume that so long as D

is held fixed throughout our reasoning, if ChExp(l,a,p) then $a \in p$ if and only if $\neg\exists q(\text{Exp}_D(l,a,q) \wedge T(a,q))$. This an awkward principle to work with, as it starts with an assumption about ChExp, but winds up using 'Exp_D' in the biconditional. However, observe that if $a \in D$, $\exists q(\text{Exp}_D(l,a,q) \wedge T(a,q))$ if and only if $\exists q(\text{ChExp}(l,a,q) \wedge T(a,q))$. Hence, we have:

(T-l) If $a \in D$, and ChExp(l,a,q), then $a \in p$ if and only if $\neg\exists q(\text{ChExp}(l,a,q) \wedge T(a,q))$.

(This is enough for our purposes here, so I shall not work out the general principle corresponding to (T-Prop).[12])

From this, it follows that an utterance of l in context set D and world a cannot express a proposition if $a \in D$. Suppose it does. Then it must express $p = D \cap r$ where ChExp(l,a,r). We now get a contradiction, as we find that $a \in p \leftrightarrow a \notin p$. As $a \in D$, $a \in p$ just in case $a \in r$. Suppose $a \in r$. Then by (T-l) $\neg\exists q(\text{ChExp}(l,a,q) \wedge T(a,q))$. As ChExp($l$,$a$,$r$), we have $\neg T(a,r)$. So $a \notin r$. Conversely, Suppose $a \notin r$. Then again by (T-l), $\exists q(\text{ChExp}(l,a,q) \wedge T(a,q))$. As ChExp($l$,$a$,$r$), by (U-ChExp), we may conclude $T(a,r)$. So $a \in r$.

An utterance of l in world a and context set D cannot express a proposition if $a \in D$. We can only avoid the contradictory conclusion of the argument by holding that any world in which an utterance of l expresses a proposition is not in the context set of the speakers in that world. This is a generalization of a point we made a moment ago for the actual world, and it is equally problematic. It forces us to conclude that relative to a given context, the Liar Sentence cannot express a proposition in any world in the context set (nor in the actual world whether or not it is in the context set); but yet that it does express a proposition. This is a violation of the very conception of context with which we have been working. It is thus no wonder that when we look at what happens with the actual world, we get absurd answers.[13]

We encounter with the Liar Paradox a general breakdown of the theory. Whatever it is that makes the difference in the transition from context (A), where the Liar Sentence does not express a proposition, to (B), where it does, is something the theory cannot represent. Likewise, it cannot provide a proposition for l to express at (B).

It must be stressed, the problem here is not generated by having a restricted context set D. Rather, the problem is that the theory cannot

provide a proposition, or a difference in context, *at all*, making use of the usually restricted domain D or *any other* worlds. To dramatize this, consider the case in which at context (A), the context set is the entire set of worlds, say W. Then *l* cannot express a proposition, relative to any world, period. But we know, on pain of contradiction, that at context (B) an utterance of *l* does in fact express a proposition. Yet there is no world at all relative to which it can, and so no way for the theory to explain the difference between (A) and (B). Perhaps more importantly, there is simply no proposition for it to express, for we have seen that starting at any world, and attempting to form propositions from all of W, the set of all worlds, there is no proposition that *l* can express.[14]

The theory under consideration provides for lots of context shifts; yet in spite of this, all the problems of the Liar are still with us. Why is this so? There is a fixed limit to the differences in context the theory can represent. Any difference in context *must* be a difference in context set – a difference in what is presupposed – so there must be a proposition expressing the difference. When we come to the Liar, it cannot find a proposition that expresses the difference between contexts (A) and (B) coherently. Likewise, it cannot find a proposition at all for the Liar Sentence to express at (B). The theory simply runs out the propositions. As propositions are sets of worlds, what it really runs out of are worlds themselves.

However, it would not help matters to simply give the theory more worlds. Given we did not say how many worlds we started with, it is hard to see even what this would be. Regardless, the problems we have encountered do not depend on the size of the domain of worlds W. This is made most clear by the case where the context set at (A) is W, the background domain of all worlds. In this case, there is no proposition available to be expressed by *l* at (B), and no way to refine the context that makes any sense of *l*'s expressive power. Had we started with some larger W′, it would have made no difference. We would come to the same conclusion about W′. We do not need to ever have a context in which the context set is all of W to encounter this problem, though. *Whatever* the context set at (A) may be, we have no worlds in W that can make sense of the transition to context (B), or provide a proposition to be expressed at (B). Again, having started with a larger set W′ would have made no difference.

What is needed is for there to be an *expansion* of W, the background domain of all worlds, in the shift from (A) to (B). We need, from (B), to have available a proposition for *l* to express, that could not be any subset of W. This requires not a bigger collection of worlds at (A), but to have an expanded W′ available at (B). We must see the *background domain* of worlds as *expanding*. This is not a shift in the domain of relevant worlds at a given point in a conversation. Stalnaker's theory allows for easy shifts in that set. Rather, it must be a shift in the background domain of all worlds W to a distinct expanded background domain W′.

Stalnaker's theory does help to show what it is to have such a shift. In shifting to a larger background domain of worlds, we essentially increase our ability to distinguish contexts and express propositions. We have more propositions, both to be expressed, and to be presupposed. To put it figuratively, we increase the resolution of theory, allowing it to display differences in finer detail. Stalnaker's theory works at a fixed resolution. It is highly flexible within its limits, but the outer limits of its resolution, both for what propositions it can represent and what contexts it can distinguish, are fixed by the background domain of worlds. The Liar Paradox requires that at least some shifts in context also be shifts in resolution: particularly, *increases* in resolution. It requires that for any context, there be another context relative to which resolution has been increased, so there are more worlds in the background domain and hence more propositions available. Once we have an expanded domain of worlds, we can make better sense of the Liar in the previous context. We can note that given the background domain of worlds we had to work with in the previous context, it could not, from that context, express a proposition. Relative to the new context, with more worlds in the background domain, it can.[15]

It is important to stress that the Paradox requires that this be the case *for any context*. We can begin the Liar reasoning in any context, which shows that for any context, there must a distinct context relative to which we have finer resolution, giving an expanded background domain of worlds. We must thus have *open-ended* refinability of resolution. Let us call the problem of explaining this phenomenon the *resolution problem*.

The Liar Paradox confronts us with the resolution problem, hence solving the resolution problem is basic to any solution to the Liar Paradox. Now, to say this is not by any means to provide a solution to the Liar Paradox. Nor is it the case that solving the resolution problem alone would constitute a solution to the Liar. But the resolution problem is one of the basic conceptual problems that underlies the Paradox. What we learned from considering Stalnaker's theory is that without a solution to the resolution problem we cannot begin to make sense of the behavior the Paradox exhibits.

I began this section with a question. Having argued that the Liar is a problem of context-dependence, I asked what sorts of issues about context are involved. Identifying the resolution problem answers the question. The issue is not merely one of finding context shifts, but of explaining a special kind of context shift: resolution shift.[16]

IV. THE RESOLUTION PROBLEM

I have now done most of what I promised I would. I have tried to show what the problem of the Liar Paradox is, and what is required for a solution to it. I have argued that it is a problem of context-dependence, and that solving it requires us to address the resolution problem. It is the resolution problem that gives us some sense of what we are up against. It tells us, at least in part, what the underlying difficulty about context is. But the resolution problem itself raises questions. What does it tell us about context or truth conditions? Does it have metaphysical implications? To conclude, I shall examine the resolution problem itself, to try to shed some light on these questions.

The resolution problem confronts us with a problem about truth conditions. The domain of truth conditions – possible worlds in many semantic theories – is apt to expand as context changes. One might well ask if this has some significant metaphysical implications. Indeed, one might wonder if it implies some form of anti-realism. If we ask why it should be that the domain of truth conditions itself must always be apt to shift as context shifts, one natural response is because the truth conditions are in part constituted by something specific to agents, which shifts as context shifts. If they are in part constituted by context-specific capacities

of agents to verify or recognize claims as assertible, we might find something that looks much like a global form of by-now traditional semantic anti-realism.[17] That would indeed be a remarkable result, for it would tell us that the Liar Paradox leads to anti-realism.

This may be one way to approach the resolution problem. Certainly nothing I have argued rules it out. But it is striking that my arguments have had nothing to do with issues, such as speakers' capacities to grasp evidence-transcendent truth conditions, which mark the usual ground for the discussion of semantic anti-realism. Though I have focused on truth conditions, the issue has been the relation of truth conditions to context, not the nature of truth conditions *per se*. This suggests we might see the resolution problem as one that is much more about language and language use than about the metaphysics of realism and anti-realism. I shall attempt now to sketch out such a way of seeing it. The understanding of the problem I shall suggest will not leave it entirely free of metaphysical import, but it will shift the primary emphasis from anti-realism to language use.

Suppose we have speakers in two distinct contexts S and T. Relative to these two contexts, the resolution problem forces us to admit, there may be different domains of truth conditions available. Hence, there may be some proposition expressible in S that is not expressible in T. There is thus *a* way in which speakers at T cannot express or understand what is being said at S. Nonetheless, it may be that speakers at T need only find a different context T′, into which one can move from T, from which they can understand what is said at S. If the passage from T to T′ is sufficiently routine, it would be an exaggeration to say that the thought expressed at S was truly unthinkable to speakers at T. They merely have to do something routine to have the thought.

Call this the *modest scenario* for solving the resolution problem. It is modest, in that there are not thoughts had by anyone that are ultimately unthinkable by someone else. The modest scenario does not appear to suggest an anti-realist view. It requires that some truth conditions be inaccessible from some context, but it does not require that the truth conditions themselves be anti-realist. It does not suggest that once expressed, a proposition cannot be evidence-

transcendent, satisfy the principle of bivalence, or have any of the other marks of realist truth conditions.

The modest scenario is also modes in that it indicates full commensurability between domains of truth conditions. As far as the resolution problem itself demands, domains of truth conditions need only be able to shift by expansion. Hence, the modest scenario may have it that T′ differs from T only in that it expands the domain of truth conditions available at T to include all those available at S. There cannot be a top, a domain of all the truth conditions, but if the differences in domains result only from expansion, it is open for the domains and contexts from which they are available to form a linearly ordered hierarchy. Moving to another context, from which a previously inaccessible proposition becomes available, may be simply a matter of moving higher up in a hierarchy.[18]

Though it does not appear to be anti-realism, that there cannot be a top of the hierarchy is not trivial matter. This is to require an *open-ended* hierarchy of contexts, and truth conditions available in them. Speakers *cannot* get outside of this hierarchy. This places genuine limits on what speakers can express within their own contexts. Fixing a context, there will be something expressible in another context that a speaker simply cannot, in any way, express without *changing* the context.

To see the importance of this, it is worth pausing for a moment to consider a common assumption about the way truth-conditional content and context relate. It is a widespread supposition in thinking about language that there are important ways in which content – at least truth-conditional content – can be *separated* from context. A proposition, once expressed, is something that can be grasped by others, in other contexts. Though context is involved in the production and apprehension of the message, the message itself, the proposition, is independent of the context in which it was produced, and it makes sense to ask about its being understood by someone in any other context. *What* is said can be separated from *how* it was said. This is a very basic version of the distinction between semantics and pragmatics.

There are, of course, some things about the significance of an utterance which are not separable from the context of utterance. An utterance may carry non-truth-conditional content which cannot

be grasped independently of the context. Jokes provide examples of this sort of phenomenon, as can reports of tragic or horrific events. Hence, for argument's sake, we may put the distinction only to require that truth-conditional content be separable from context (though I am not aware of any theory of non-truth-conditional content that entirely settles the matter).

So understood, this version of the semantics/pragmatics distinction is deeply entrenched in much thinking about meaning and intentionality. It is certainly not the only distinction that might bear this name, and particular theoretical frameworks might provide different ways of describing it. But it is, I believe, a fundamental philosophical starting point for drawing such distinctions. Even those theorists of pragmatics who believe the relation between an utterance and the truth conditions it expresses to be primarily to pragmatic matter do not challenge it. Though they take proposition *expression* to be part of pragmatics, they grant that propositions themselves are independent of pragmatics.[19]

The resolution problem calls into question this basic semantics/pragmatics distinction. Even on the modest scenario, the domain of truth conditions, out of which contents are constructed, must be context-dependent. Hence, we must recognize that certain propositions are simply not available for expression or comprehension, in any way, in some contexts. The result is that there are cases where what is said, even truth-conditionally construed, cannot be wholly separated from the context in which it is expressed.

However, the modest scenario requires only limited disruption of the semantics/pragmatics distinction. If we fix a range of truth conditions, and look only at contexts relative to which that domain is available, then restricted to those contexts things are as the semantics/pragmatics distinction would have them. In terms of the hierarchy of contexts, the violation of the semantics/pragmatics distinction amounts to the point that speakers cannot get out of the hierarchy, or reach a top. Even truth-conditionally, what they can express in constrained by where they happen to be in it. But restricted to fixed levels of the hierarchy, the distinction holds.

If I am right, the resolution problem is one of the relation of content to context, and not a direct case for global anti-realism. But even the modest scenario is not completely neutral as to the nature

of truth conditions. It makes puzzling views according to which the truth conditions are simply there, as objects. If that were the case, it seems, we would have a maximal domain of truth conditions. The modest scenario rules this out, as there cannot be a top to the hierarchy of contexts. This is a challenge for modal realists, who would agree that in some sense the possible worlds are simply there. At least, it makes trouble for any modal realist who also insists that truth conditions are possible worlds.[20]

There are views of possible worlds that are more easily reconcilable with the resolution problem, but do not embody semantic anti-realism. A prominent example is the one defended by Stalnaker himself.[21] Rather than delve too deeply into the dense matter of modal ontology, I shall just say something in very general terms about how we may think of truth conditions in a way compatible with the resolution problem. Borrowing a metaphor from Stalnaker (1984), it is useful to think in term of construction. Speakers in contexts construct contents, in that they formulate and express thoughts. The resolution problem leads us to add that the truth conditions speakers use to construct contents are themselves the sorts of things that are constructed within contexts.

According to this picture, insofar as we think of truth conditions as objects, they are constructed objects. A good model is that of artifacts. What artifacts there are is a matter of what artifacts have been made; yet this is not a reason to see artifacts in any kind of anti-realist light. That buildings only exist if we construct them should not lead us to anti-realism about buildings. It might be possible to think of truth conditions similarly. What truth conditions are available, in a given context, are those that can be constructed in the context, but once constructed, they work as the realist would have it. This is hardly an account of truth conditions. I merely want to make the point that there is room to face the metaphysical consequences of the resolution problem without resorting to global semantic anti-realism.

What metaphysics we ultimately opt for is not to be decided just by thinking about the resolution problem. But the problem does raise some important philosophical issues. It raises some genuine questions about the metaphysics of truth conditions, and more about

the nature of context, and what speakers do in contexts to express thoughts.

What is remarkable is that we get to these conclusions by examining the Liar Paradox. This tells us something about what is needed to solve the Paradox. But it also tells us something about what makes the Liar Paradox philosophically interesting. At its root, if I am right, are some genuinely difficult philosophical issues.[22]

NOTES

[1] Parsons (1974) and Burge (1979) are the classic sources for context-based approaches to the Liar. More recent work includes Barwise and Etchemendy (1987), Koons (1992) (building on Gaifman 1988), and Simmons (1993). (See notes (16) and (21) for some more discussion.)

In attempting to understand the problem posed by the Liar, and what is needed for a solution to it, this paper shares some of the spirit of Chihara (1979), though I come to very different conclusions than he does.

[2] Where clear and convenient, I am intentionally confusing use and mention. There are enough quotation marks in this paper already.

[3] This kind of theory is discussed by Vann McGee (1989, 1991), though he uses it as an intermediate step on the way to a view that is supposed to address some of the problems I am about to raise. I should mention that I choose it primarily for its convenience. Much stronger theories have been developed, particularly by Solomon Feferman (1991). A cleaned-up version of the theory presented here, extending a theory of arithmetic, is equivalent to Peano arithmetic (see Friedman and Sheard 1987), whereas Feferman's have the strength of various fragments of ramified analysis.

[4] The technical results in this paragraph are drawn from Friedman and Sheard (1987).

[5] In an interesting paper, Ian Rumfitt (1993) has challenged a principle like (U-Exp). However, his concerns are with Kripke-like belief puzzles, and I do not think they bear on the sorts of uses of the principle I make here.

[6] The Liar Paradox is strictly about what, if any, truth conditions may be assigned to a sentence in a context, so it is insensitive to issues of how finely structured propositions must be. Thus, we may take the possible worlds view of propositions as at least a simplifying assumption, regardless of whether the familiar arguments, such as those of Soames (1987), ultimately show propositions to be structured entities.

I have subtly slipped into the language of possible worlds. As with propositions, I wish to do so without taking on any great philosophical commitments. I need to discuss circumstances in which sentences are uttered and what is said is true or false. The language of possible worlds is useful for this, and is used by most of the theories I discuss, so I shall follow suit.

[7] This is a pragmatic notion of presupposition that has become quite common. Stalnaker (1974) discusses presupposition in more detail. A survey of some more recent work may be found in Beaver (1997). Christopher Gauker (1998) argues against the idea that speakers' attitudes determine contexts, though I do not think his arguments affect what I argue here.

[8] This leads to two (of Stalnaker's three) requirements on context sets and assertions: (1) when an utterance of a sentence expresses a proposition, it should express one relative to each world in the context set, and (2) it should express the same proposition in each. These make compatible the ideas that the context set is a representation of the context of utterance and that characters are functions on individual worlds.

For purposes of this discussion, I am equating David Kaplan's (1989) characters with Stalnaker's propositional concepts. It should be noted that there are a few trivial formal differences between the two, and more important differences between the phenomena they are used to explain, but none of these differences are relevant here.

[9] There has been a great deal of further development of this approach to context and context change, much of it motivated by Irene Heim's File Card Theory (1988) or Hans Kamp's Discourse Representation Theory (1984) or Jeroen Groenendijk and Martin Stokhof's Dynamic Logic (1991). (More recent work is surveyed in the handbook articles van Eijck and Kamp (1997) and Muskens, van Bentham, and Visser (1997).) However, these developments do not affect the arguments I offer here. As I have already mentioned, the issue before us is entirely one of truth conditions. These theories depart from Stalnaker's original idea primarily by finding ways to assign to sentences distinct semantic values (or discourse representation structures) but identical truth conditions. In Heim's way of presenting the theory, for instance, the semantic values of sentences are sets of pairs of a world and a sequence of variables. The presence of the variables is important for explaining certain phenomena of anaphora. Because of them, sentences can have distinct semantic values that determine the same truth conditions. For problems that are specific to truth conditions, such as the Liar, this will not matter.

[10] In cases like this, where the context does not seem to be working right, there is a general strategy that sometimes helps. The idea is to find some systematic way to reinterpret problematic sentences so as to avoid the defect in the context. The usual method is to apply some context-shifting operator to a sentence's character, which shifts the input world. In particular, Stalnaker (1978) argues that in some cases, we should apply the diagonal operator, which essentially causes a character to behave as if it had as input the world of evaluation. This is useful in cases where Stalnaker's three principles of assertion are violated, as they are, for instance, with identity statments.

If we, somehow, applied this strategy to the Liar, the result would have to be to evaluate the utterance of *l* at (B) as if it were made in some other world. But no world presents itself. It will not help to shift to the world as speakers see it, as the speakers presuppose that the utterance cannot express a proposition. Yet what

other worlds should we consider than the actual one and those consistent with what speakers persuppose? It appears that the theory does not point us to any world, or context-shifting operation, that could explain the ability of the utterance at (B) to be expressive, and so the basic problem of the Liar Paradox remains.

[11] Actually, all that is needed for the argument that follows is $(\mathrm{ChExp}(s,w,p)\wedge\mathrm{ChExp}(s,w,q)) \rightarrow (w\in p \leftrightarrow w\in q)$, but (U-ChExp) is an extremely natural principle in the framework under consideration.

[12] I have assumed here that the only way $\mathrm{Exp}_D(s,w,p)$ is context-dependent is in the parameters w and D. Given the way the theory under consideration works, this is the right assumption to make, but note that the conclusions I come to in this paper may ultimately call it into question. For a discussion of some of the standard issues for formulating a general principle like (T-Prop) in the framework under consideration, see Chierchia and McConnell- Ginet (1990).

[13] If we simply insist that there is some proposition for l to express, we quickly encounter a break-down of our technical apparatus. Suppose that evaluated at a, with context set D, an utterance of l expresses a proposition, say g. If we assume $a\in$D, we get a contradiction. If we assume $a\notin$D, our technical apparatus malfunctions catastrophically. Suppose $a\notin$D. We seem to have a quick argument that g is true at a, as I have argued it should be. But the reasoning is suspicious. All we need is to note that for any proposition p, $a\notin p\cap$D, since $a\notin$D. Tracing through our definition, this implies that $\neg\exists q(\mathrm{Exp}_D(l,a,q)\wedge\mathrm{T}(a,q))$, which in turn looks as if it shows $a\in g$. But this is unimpressive. Our set-up required us to ask about the truth of the intersection with D, but assuming $a\notin$D, this is clearly the wrong question to ask. Once we lose the assumption that $a\in$D, we can no longer take the proposition expressed by l at a simply to be the intersection of the value of its character with D. We know that the expression relation Exp_D itself must be context-dependent, but without this assumption, we hardly know at all what happens to the expression relation as context shifts. As we don't know what happens to the expression relation, which is mentioned in l, we have no idea what proposition l might express.

[14] This is to suppose that (A) is somehow a null context. I doubt it is possible for speakers to ever really be in such a context. But it is, on the theory we are discussing, a coherent example to consider.

Stalnaker's theory, as I have described it, provides only for monotonic context update. One moves from context to context only by restricting the context set. The Liar leads us to expect non-monotonic context shifts. This example shows, however, that merely providing for non-monotonic context update cannot be enough to solve the Liar. Without other changes, doing so would not help with this case.

[15] Hence, a modification of Stalnaker's theory to allow for such shifts in domain would give it something to say about the shift in context involved in the Liar Paradox. I should note, however, that I am not sanguine about the prospects of a purely representational theory of context, like Stalnaker's, providing a full explanation. The shift in the case of the Liar is marked by the drawing of an inference, which representational approaches to context, especially Stalnaker's, have difficulty explaining.

[16] One of the few discussions of the Liar Paradox to address the question of what phenomenon of context-dependence is involved is that of Tyler Burge (1979). There is a great deal that is extremely valuable in this paper, but on the issues that concern us here, I see some serious difficulties. It attempts to explain the context-dependence of the Liar by way of Gricean pragmatic mechanisms. As traditionally understood, these are based on the combination of propositions expressed and conversational principles of cooperative rationality governing what is appropriate to express. But the Liar poses a problem about how a proposition can be expressed at all, and until we answer it, we cannot appeal to traditional Gricean devices. Stalnaker's theory embodies many Gricean ideas in a more flexible and general setting, so its failure makes me doubt that even a non-traditional variant on Grice's ideas could do the job Burge asks them to do.

The same points hold for work in the Burgean tradition, including Koons (1992) and Simmons (1993). Neither work seems to me to be particularly concerned with the pragmatic foundations of the context-based approach to the Liar; but to the extent that they mention it, they acquiesce to Burge's Gricean account.

[17] As found, for instance, in Dummett (1976), among many places.

[18] According to the modest scenario, where we only have domains of truth conditions shifting by expanding, we will get a genuine linear hierarchy. The arguments I have presented do not rule out something more like a partial ordering of domains. They do require that whatever the ordering, it cannot have *any* top elements.

A solution to the resolution problem by any sort of hierarchy or ordering of domains, such as we are considering, indicates something like a hierarchical solution to the Liar Paradox itself; though, obviously, one based on contexts and domains of propositions rather than a hierarchy of truth predicates. This would bear more similarities to the proposal defended by Parsons (1974) than to Tarski's original one, but would differ in important ways. However, I promised I would not present a solution to the Paradox, so I shall not pursue this further here.

[19] See, for example, Carston (1988) or Bach (1994).

[20] Other puzzles of this sort, such as that from Davies (1981), which purport to show that modal realism is incompatible with the demands of adequately representing content, have been known for some time. Able defenders of modal realism, prominently Lewis (1986), have attempted to answer them. One might defend modal realism against my challenge by insisting that all the worlds there are cannot form a single domain out of which speakers can form propositions. This is stronger than the idea that we simply do not use some of them. It requires that in some way we *cannot* do so. It would be fair to ask, in response, what work the modal realist component of the view is doing for the explanation of content.

[21] In some important ways, Jon Barwise and John Etchemendy's (1987) situation-theory-based work on the Liar Paradox comes to similar conclusions to my own (as does Willem Groeneveld's (1994) combination of their ideas with ideas from dynamic semantics). They came to the conclusion that there must be a sequence of (actual) situations $s_1, s_2, s_3, \ldots$ which is open-ended in much the way that I conclude contexts must be. (The open-endedness is captured especially nicely by Groeneveld's idea of the failure of his *success* condition.) In Barwise

and Etchemendy's framework, which does not handle expression and expression failure quite as I do, the sequence of situations is generated by adding *semantic facts*. In their notation, $s_{i+1} = s_i \cup \{\langle Tr, f_{s_i}; 0\rangle\}$. Non-well-founded set theory makes it particularly easy to find these sorts of semantic facts.

I have some serious worries about the account of context that is implicit in situation theory, thought it would take us too far afield to discuss them here. Of more immediate concern is what situation theory might say about the resolution problem. To give a solution to it along these line, we would have to say something about these semantic facts. Though non-well-founded set theory provides formal models of them, this seems to present us with a version of the resolution problem, rather than a solution. When we look at more foundational work on situation theory, we do not get a clear solution. Versions of situation theory which see situations as partial possible worlds or regions of space-time (as in Kratzer (1989) and perhaps Barwise (1989)) have their resolution fixed by the domain of possible worlds and the structure of space-time, which certainly appears to imply a maximal resolution, and so re-poses the resolution problem in its old form. The abstract situations and states of affairs of Barwise and Perry (1983) seem to have their resolution fixed by the domains of individuals, relations, and locations. So the result is to move the resolution problem to these domains.

In commenting on semantic paradoxes, Barwise (1989, p. 84) mentions the "impossibility of, in general, making claims about the whole world" and that the paradoxes "always generate facts outside the scope of the part of the world the claim was about." Perhaps the sensitivity to set/class distinctions that is always present in Barwise's work could help to fill in this idea. Hence, I do not claim that situation theory cannot solve the resolution problem; only that it does not come ready-made to do so. Regardless, it strikes me that the line hinted at in Barwise's remarks is not one that is available only to situation-theorists. Replace "the whole world" with "domain of possible worlds" and we have an idea that modal realists might try to use as well. They might try to argue that there is some reason that all the possible worlds cannot form a single domain of quantification.

[22] I am grateful to Lenny Clapp, Warren Goldfarb, Richard Heck, Darryl Jung, Charles Parsons, Sanford Shieh, Jason Stanley, Zoltán Gendler Szabó, Carol Voeller, and an anonymous referee for *Philosophical Studies* for many helpful comments, suggestions, and discussions. I am especially grateful for many years of advice and encouragement from the late George Boolos.

REFERENCES

Bach, K. (1994): 'Conversational Impliciture', *Mind and Language* 9, 124–162.

Barwise, J. (1989): 'Situations and Small Worlds', in *The Situation in Logic*, pp. 79–92, Stanford: CSLI Publications.

Barwise, J. and Etchemendy, J. (1987): *The Liar*, Oxford: Oxford University Press.

Barwise, J. and Perry, J. (1983): *Situations and Attitudes*, Cambridge: The MIT Press.
Beaver, D. I. (1997): 'Presupposition', in van Bentham and ter Meulen (eds.), *Handbook of Logic and Language*, pp. 939–1008, Cambridge: The MIT Press.
Burge, T. (1979): 'Semantical Paradox', *Journal of Philosophy* 76, 169–198.
Carston, R. (1988): 'Implicature, Explicature, and Truth-Theoretic Semantics', in Kempson (ed.), *Mental Representation: The Interface Between Language and Reality*, pp. 155–181, Cambridge: Cambridge University Press.
Chierchia, G. and McConnell-Ginet, S. (1990): *Meaning and Grammar*, Cambridge: The MIT Press.
Chihara, C. (1979): 'The Semantic Paradoxes: A Diagnostic Investigation', *Philosophical Review* 88: 590–618.
Davies, M. (1981): *Meaning, Quantification, Necessity*, London: Routledge and Kegan Paul.
Dummett, M. (1976): 'What is a Theory of Meaning? (II)', in Evans and McDowell (eds.), *Truth and Meaning*, pp. 67–137, Oxford: Oxford University Press.
van Eijck, J. and Kamp, H. (1997): 'Representing Discourse in Context', in van Bentham and ter Meulen (eds.), *Handbook of Logic and Language*, pp. 179–237, Cambridge: The MIT Press.
Feferman, S. (1991): 'Reflecting on Incompleteness', *Journal of Symbolic Logic* 56: 1–49.
Friedman, H. and Sheard, M. (1987): 'An Axiomatic Approach to Self-Referential Truth', *Annals of Pure and Applied Logic* 33: 1–21.
Gaifman, H. (1988): 'Operational Pointer Semantics: Solution to the Self-Referential Puzzles I', in Vardi (ed.), *Theoretical Aspects of Reasoning about Knowledge*, pp. 43–60, Los Angeles: Morgan Kauffman.
Gauker, C. (1998): 'What Is a Context of Utterance?' *Philosophical Studies* 91, 149–172.
Groenendijk, J. and Stokhof, M. (1991): 'Dynamic Predicate Logic', *Linguistics and Philosophy* 14, 39–100.
Groeneveld, W. (1994): 'Dynamic Semantics and Circular Propositions', *Journal of Philosophical Logic* 23, 269–306.
Heim, I. (1988): *The Semantics of Definite and Indefinite Noun Phrases*, New York: Garland Publishing, Inc.
Kamp, H. (1984): 'A Theory of Truth and Semantic Representation', in Groenendijk, Janssen and Stokhof (eds.), *Truth, Interpretation, and Information*, pp. 1–41, Dordrecht: Foris Publications.
Kaplan, D. (1989): 'Demonstratives', in Almog, Perry and Wettstein (eds.), *Themes from Kaplan*, pp. 481–563, Oxford: Oxford University Press. This is the first publication of a widely-circulated manuscript dated 1977.
Koons, R.C. (1992): *Paradoxes of Belief and Strategic Rationality*, Cambridge: Cambridge University Press.
Kratzer, A. (1989): 'An Investigation of the Lumps of Thought', *Linguistics and Philosophy* 12, 607–653.

Kripke, S. (1975): 'Outline of a Theory of Truth', *Journal of Philosophy* 72, 690–716.
Lewis, D. (1986): *On the Plurality of Worlds*, Oxford: Basil Blackwell.
McGee, V. (1989): 'Applying Kripke's Theory of Truth', *Journal of Philosophy* 86, 530–539.
McGee, V. (1991): *Truth, Vagueness, and Paradox*, Indianapolis: Hackett Publishing Company.
Muskens, R., van Bentham, J. and Visser, A. (1997): 'Dynamics', in van Bentham and ter Meulen (eds.), *Handbook of Logic and Language*, pp. 587–648, Cambridge: The MIT Press.
Parsons, C. (1974): 'Content and Context: The Liar Paradox', *Journal of Philosophical Logic* 3, 381–412.
Rumfitt, I. (1993): 'The Paratactic Theory Revisited and Revised', *Mind* 102, 429–454.
Simmons, K. (1993): *Universality and the Liar*, Cambridge: Cambridge University Press.
Soames, S. (1987): 'Direct Reference, Propositional Attitudes, and Semantic Content', *Philosophical Topics* 15, 47–87.
Stalnaker, R. (1974): 'Pragmatic Presuppositions', in Munitz and Unger (eds.), *Semantics and Philosophy*, pp. 197–230, New York: New York University Press.
Stalnaker, R. (1978): 'Assertion', in Cole (ed.), *Syntax and Semantics, vol. 9: Pragmatics*, pp. 315–332, New York: Academic Press.
Stalnaker, R. (1984): *Inquiry*, Cambridge: The MIT Press.
Tarski, A. (1983): 'The Concept of Truth in Formalized Languages', in Corcoran (ed.), *Logic, Semantics, Metamathematics*, 2nd ed., pp. 152–278, Indianapolis: Hackett Publishing Company. English translation by J.H. Woodger of 'Der Wahrheitsbegriff in den formalizierten Sprachen', *Studia Philosophica* 1 (1935).

Department of Linguistics and Philosophy
MIT
Cambridge, MA 02139
USA
E-mail: glanzber@mit.edu