

**Solve The Paradox? –
A challenge to presumptions of paradox solution
existence with particular reference to the case of the liar**

**Colin Peter Johnston
MPhil Philosophy
University College London
2002**


ProQuest Number: U643298

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest U643298

Published by ProQuest LLC(2015). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code.
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

This thesis is an argument, primarily by example, that we have no good reason to suppose that the semantic paradoxes (specifically, the liar and the Grelling) have solutions as such things are standardly conceived. Rather, the matter of solving the liar is, I suggest, a purely practical issue which will arise only in very unusual circumstances.

Central to my argument is the example of a club C whose entry criterion is as follows: a person is eligible to join C if, and only if, he or she is secretary to a club which they are ineligible to join. If we enquire about the eligibility of C's own secretary, a lady who happens not to be secretary to any other club, the criterion ties itself in knots and we are landed in contradiction. However, we are not disposed in this case to find an error in the paradoxical reasoning about C's secretary (a faulty premise or inference). And a solution to the paradox will, if the secretary applies for membership, consist in taking a practical decision either to accept or to refuse her application.

Having expanded my position into a more comprehensive diagnosis of semantic paradox through, *inter alia*, a consideration of Kripke's notion of groundedness, I defend it from attack from a number of quarters. I discuss the thought that a predicate determines a condition which an object in its domain of applicability must either meet or fail to meet; I also address concerns provoked by Curry's paradox and the principle of *ex contradictione quodlibet*.

There is a further section in which I compare my stance with that adopted by Quine in 'The Ways of Paradox'.

Table of Contents

Section	Page
Chapter 0	
0.0	4
0.1	6
0.2	10
0.3	13
Chapter 1	
1.0	17
1.1	17
1.2	22
1.3	28
1.4	36
1.5	40
1.6	42
Chapter 2	
2.0	44
2.1	44
2.2	50
2.3	56
Chapter 3	
3.0	57
3.1	57
3.2	60
3.3	63
3.4	64
References	66

Solve the Paradox?

'The civil status of a contradiction: there is the philosophical problem'

Wittgenstein, 'Philosophical Investigations' I §125

Chapter 0

0.0

Paradoxes, it is often assumed, come in two varieties: the veridical and the falsidical. In the former, the surprising conclusion arrived at by seemingly valid argument from plausible premises is in fact true; in the latter it is false. If the paradox is veridical we must come to terms with the conclusion, if it is falsidical we must either throw out as false one of the innocuous looking premises, or reject as invalid one of the inferences made in the course of the argument. Once it is seen how it is that the conclusion is acceptable after all, or why it is that the premise or inference was defective, the paradox is solved. I shall call this conception, a conception at once of the nature of paradoxes and of the type of response they demand, the standard model.¹

Of all the well known paradoxes, that of the liar has proved peculiarly intransigent. Several of the last century's most powerful minds armed with all manner of logical machinery, let alone the less equipped but equally powerful minds of the ancients and medievals, have spent thousands of pages doing battle with the problem, but to date none have returned with an unmixed victory. It remains, for many, one of logic's greatest unsolved mysteries. I present it as follows:

¹ The introduction to R. M. Sainsbury's 'Paradoxes' (Sainsbury 1988), contains an unhesitatingly endorsed statement of 'the standard model':

This is what I understand by a paradox: an apparently unacceptable conclusion derived by apparently acceptable reasoning from apparently acceptable premises. Appearances have to deceive, since the acceptable cannot lead by acceptable steps to the unacceptable. So, generally, we have a choice: Either the conclusion is not really unacceptable, or else the starting point, or the reasoning, has some nonobvious flaw. (Sainsbury 1988 p1)

Premise(α): If x is a sentence consisting of the referring expression y immediately followed by the words 'is not true', and if y denotes the sentence z , then x is true if, and only if, z is not true

Premise(β): There is a sentence L consisting of a referring expression r immediately followed by the words 'is not true', and r refers to the sentence L itself

From these two premises we conclude immediately that L is true if, and only if, L is not true. From which we argue that because L , if true, is not true, and vice versa, L is both true and not true. A contradiction.

Now contradictions, it is commonly supposed, are necessarily false, and so the liar is a falsidical paradox. Thus, because premise (β) is an empirical fact (consider, for example, the following sentence L : The sentence you are reading right now is not true), we are forced either to reject either premise (α) or the argument from (α) and (β) to the contradiction.

If we treat the paradox as a reductio of premise (α) then we must conclude that the logical structure of truth, that is the logical relations of sentences containing the truth predicate both to each other and to sentences not containing the truth predicate, a structure part of which (α) is an attempted statement, is not so simple as was supposed. In this case the solution to the paradox will uncover the correct structure of truth, making it clear why and how premise (α) is false. Alternatively, if it is the reasoning and not the premises that is considered faulty, we are committed to a rejection of classical logic at least in the vicinity of the truth predicate; a move which would, one might imagine, appeal to a motivation and/or intuition beyond the existence of the paradox itself. Either way, the paradox demonstrates an error in our naive theory of the logic of truth, an error which, it is hoped, a logician or philosopher will some day discover and correct, thereby solving the paradox.

The sentiment of these last two paragraphs, the standard model as applied to the liar, is explicit in several philosophers:

Keith Simmons: This points to the significance of the liar: It suggests that we do not have a proper understanding of our ordinary notions of truth and falsity. An

investigation of the liar promises to correct and deepen our understanding of our basic semantic concepts.²

Vann McGee: However attractive this account [our naïve theory] may be, it cannot be right... The naïve theory of truth has an observably false consequence, viz., ‘The starred sentence is not true’ ≠ the starred sentence. Therefore, the naïve theory of truth is incorrect.³

Scott Soames: We should regard it [the liar paradox] as the impetus for replacing initially plausible but ultimately inaccurate assumptions about the truth predicate with an acceptable theory of how the predicate really works in the [English] language.⁴

As a means of examining this dominant view, the standard model take on the liar paradox, I want to introduce some other concrete examples of paradox. They will, I hope, shed some light on the logic of the liar.

0.1

Firstly the famous paradox of the village barber. Russell presents this as the suggestion that there is a village in which the barber shaves all and only those men who don't shave themselves. This seems an unextraordinary fiction until we consider the barber's own beard, and notice that the supposition entails both that it is shaved and that it is not shaved by its owner. The solution to the paradox, then, is that contrary to first impressions such a village could exist only if the barber is a woman, and if it is added to the initial story that the barber is a man, then it becomes clear that there could no more be such a village than there could be a stick which was both longer than twenty foot and shorter than ten.

Now consider the following reformulation: A village council, concerned about public health, decrees that the village barber is to shave all and only those men who don't shave themselves. The decree has the consequence that the barber is to shave himself if and only if

² Simmons 1993 p1

³ McGee 1990 p2

⁴ Soames 1999 p54

he does not shave himself, and so whatever he does he will fall foul of the law. It might be possible to interpret the decree as applying only to those men who shave, in which case he could grow a beard or perhaps pluck his chin in order to maintain legality. But if there is a further decree that all men are to shave then there seems no way out. The solution to the paradox, of course, is to point out the problem to the council who will, being reasonable people, amend the initial decree, perhaps to: The only man allowed to shave another is the barber. And they will do this because the point of the legislation was to prevent amateurs hacking away at each other's necks, not to persecute the poor barber.

My next example of paradox is taken from the novel 'Don Quixote'. At a certain point in the tale, the Don instates his loyal manservant, Sancho Panza, as governor of the Island of Barataria. After passing several judgements of astounding Solomonesque wisdom, Sancho is finally stumped by the following problem: A large river divides in two a lordship of the island and upon that river there is a bridge. At one end of the bridge there lies a gallows, and next to the gallows a court of justice where four judges sit to enforce a certain law made by the lord of the land. And the law runs as follows:

Whoever intends to pass from one end of this bridge to the other, must first upon his oath declare whither he goes, and what his business is. If he swear truth, he may go on; but if he swear false, he shall be hanged, and die without remission upon the gibbet at the end of the bridge.⁵

Now for a while the law served well enough, all travellers swearing true and being allowed to pass, until one day a certain passenger, on being sworn, declared that 'by the oath he had taken, he was come to die upon that gallows, and that was all his business'. This 'put the judges to a nonplus', for they reasoned that if he is allowed to pass freely he will have sworn false and so should have been hung, but if they hang him he will have sworn true and so should have been allowed to pass.

Sancho's initial suggestion is to cut the man in two, hanging one half and letting the other pass, but this falls flat when he is reminded that half a man cannot cross a bridge, and so there could be no letting or preventing in this case. Soon, then, he realises there is no more reason in the law to hang than to allow passage or vice versa, and so he makes recourse to the precept of his master Don Quixote that when the scales of justice are even to prefer

⁵ Cervantes: 'Don Quixote de la Mancha' Part II Chapter LI. (Cervantes 1993 p330)

mercy before rigour. Thus the traveller is allowed to pass unhindered, and insofar as we can talk of solution in this case, the paradox is solved.

A final example of paradox, before we take a pause for thought, is the case of the Secretaries' Liberation Club:⁶ A society, called the Secretaries' Liberation Club, or SLC for short, is founded as a protest against social elitism. The stipulated entry criterion is as follows: A person is eligible to join the club if and only if he or she is a secretary of a club which they are ineligible to join. Now as it happens, the club appoints a new secretary, Mrs Finline, who is not a secretary for any other club. This all seems very merry, until we ask whether or not Mrs Finline is eligible to join the SLC, at which point we embark upon the familiar paradoxical reasoning: If she is eligible for SLC then there is no club which she is ineligible to join and for which she is a secretary, and so she is ineligible for the SLC. But if she is ineligible to join then there is such a club, namely the SLC itself, and so she is eligible to join after all. What then should be done? Two courses of action seem possible: either sack Mrs Finline and ensure that in the future the club employs only secretaries who happen also to be secretaries of other clubs which they are ineligible to join, or amend the entry criterion. The possibilities have their pros and cons – the latter seems more sensible and less likely to run into difficulty with employment laws, but it is not without undesirable consequences, for to bar Mrs Finline would clearly be against the spirit of the club's foundation, but if she is admitted the uniformity of the club will be disrupted. Note, though, that positive action on this matter will be necessary only if Mrs Finline applies for membership; if she does not the potential problem may pass completely unnoticed.

So now for the thoughtful pause promised earlier. I have presented three possible situations each with the characteristic paradoxical feel, and indeed each case can be seen as an argument from a certain premise, or set of premises to an undesirable conclusion: that the barber is both required to shave himself and required not to shave himself; that the traveller may neither be allowed passage nor denied passage; that Mrs Finline is both eligible and ineligible to join the SLC. But contrary to the standard model of paradox I outlined at the start of this chapter, I shall suggest that for none of these three arguments should we throw

⁶ This example is taken from Chihara 1979. He does not claim the paradox to be of his own invention, but nor does he indicate any alternative origin.

out a premise or inference as in any sense incorrect – not even in the case of Mrs Fineline where the argument's conclusion is a direct contradiction.

What we have in each case is a stipulated rule, a law which tells us who may shave who, in what circumstances a bridge may be crossed, or when a person may be admitted to a club. And the paradox, I suggest, lies in the fact that despite having a completely unproblematic application in the vast majority of circumstances, it is possible to imagine cases in which each rule ties itself in knots and yields contradictory pronouncements. This, of course, is slightly bewildering, but only, I think, for the very simple reason that we are not used to such a conspiracy of rules and circumstances. Once we recognise that this is all that is going on, we are able to view these three paradoxes simply as entertaining quirks and certainly not as presenting profound problems for the philosophy of logic.

In the problematic cases (the barber, the traveller, Mrs Fineline) the rules, as I just remarked, fail to yield a coherent pronouncement. Thus to ask which way the rules pronounce in these cases is pointless; the question has no informative answer. But I take it as obvious that there is nothing more to SLC-eligibility than is given in the entry criterion, and nothing more to being allowed to cross the bridge or shave another than is given in the law. To suggest otherwise is surely absurd. It is therefore pointless to ask whether or not Mrs Fineline is eligible to join the club, whether the traveller should be hung or whether the barber should shave himself. The temptation to ask such questions should disappear once it is seen that the rules fail us in these cases. All there is to say about the matter is that this is the rule and that these are its consequences.

This conceded, it still seems reasonable to maintain that because in each case an unacceptable conclusion is derivable from the combination of a certain rule with unfortunate empirical facts the rule cannot be true. In the case of the SLC we derive a flat contradiction and so it might seem a matter of simple logic that the rule cannot be true. But this is an uncomfortable result, for if it is not true that a person may join the SLC if, and only if, they are secretary of a club they are ineligible to join, then we are left asking after the circumstances in which a person *is* eligible to join the SLC. And to ask this is to assume that there is, after all, more to SLC-eligibility than is given in the entry criterion, a suggestion I have just rejected as patently absurd.

So what then for the truth or falsity of the rules themselves? This question, I suggest, is a red herring: the rules are not of themselves either true or false in the sense implied. If such rules can in any way be called true or false, that could only mean that they are, or are not, the rules we follow. And in the above three paradoxes the rules are stipulated as those we are to follow, and so even this weakened sense is inappropriate. The rule in each case is a marker from which the truth or falsity of other sentences flow, it is not itself liable to truth-evaluation. It can be seen, perhaps, as antecedent to truth.⁷

We can see, then, that the standard model of paradox I set out above fails to fit my three examples. As a consequence the issue of solving the paradoxes will not consist in finding a faulty premise or an error in the reasoning. Instead, as I suggested at the time, solutions will be purely practical matters – the rules that were supposed to direct our behaviour have failed us and we need to make a decision as to how to proceed. In this light, then, these three paradoxes present practical problems not theoretical ones. But they may not even do this: Mrs Finline may be admitted or rejected without anyone noticing the contradiction; the barber may carry on shaving himself without realising that he is required not to; the traveller may be hung without the judges seeing that the very act of hanging makes itself illegal. Which is all to say that the unnoticed problem is no problem at all (except, perhaps, for the man swinging from the gallows).

0.2

As a means of easing us back towards the liar, and also as a matter of considerable independent interest, I shall next examine Grelling's heterological paradox. The following is, I hope, a fairly standard presentation:

First a few terminological definitions to get us off the ground: I shall call anything that can be truly predicated of an object a property of that object. Anything that is a property of some object I shall call a property. And I shall say that an adjective *A* *names* a property *p* if, and

⁷ If these last two paragraphs seem to be directed rather pointlessly at a position which no one would seriously maintain, then it is perhaps worth my remarking that in the only discussion I have found of the SLC paradox, Chihara claims that a statement of the eligibility criterion is indeed false. (See Chihara 1979 p594). Austin seems to have encountered similar confusion: in 'How to do things with Words' he denounces the 'widespread obsession that the utterances of the law... *must* somehow be statements true or false' (Austin 1976 p19).

only if, p can be truly predicated of some object O by means of the sentence: 'O is A'. Thus, for example, redness, or the property of being red, is a property named by the adjective 'red'.

Now note that adjectives not only name properties, but may also have properties. The adjective 'short', for instance, names the property of shortness, and has the properties of being an English word, of containing an 'o', and, as it happens, of being short.

This observation paves the way for the following definitions:

defn.(a): An adjective A is autological if, and only if, A has a property which it names

defn.(b): An adjective A is heterological if, and only if, it is not the case that A is autological

So the adjective 'short' is autological, as is 'polysyllabic' and 'English'. Examples of heterological adjectives are 'long', 'rude' and 'stubborn'.

The paradox reasoning is then initiated by asking whether the adjective 'heterological' is autological or heterological. By the initial definitions we see that 'heterological' names the property of being heterological, or heterologicality for short. Next, definition (a) tells us that 'heterological' is autological if, and only if, 'heterological' has the property of heterologicality, i.e. that 'heterological' is autological if, and only if, 'heterological' is heterological. And applying (b) we conclude that 'heterological' is autological if, and only if, it is not the case that 'heterological' is autological: the standard paradox result.

Those who seek to solve this paradox typically take one of the following two lines: either deny that the adjective 'heterological' names a property – that there is no property of heterologicality, or claim that 'heterological' does name a property, but that an adjective does *not* have this property if, and only if, it names none of its own properties. This second approach is, to me, quite baffling, for it entails a direct denial of the veracity of definitions (a) and (b), and I fail even to see what such a denial might mean. The definitions are simply not up for being denied – they are a stipulation of when we are to call a word heterological. We might decide, of course, that they constitute a poor definition if they turn out to be difficult to use, or of little use, or even to lead to contradiction; but to suggest that they are

not true is to fail to see that the notion of truth is inappropriate here. Again, the definitions are antecedent to truth.

Which leaves the paradox solver in the following position: she must deny that there is a property of heterologicality; she is forced to hold that the failure of an adjective to have a property of which it is a name, does not itself constitute a property of this word. But as I have presented the paradox, this attempt must also miss the mark, for I am explicitly using the word 'property' to include anything that can be truly predicated of an object, and we can truly predicate heterologicality of the adjective 'long', for instance. If this is felt to be an abuse of the word 'property', then we can substitute another word here, perhaps 'attribute'.

Another way of seeing the hollowness of this no-such-property 'solution' is by reformulating the initial definitions in terms of objects falling under concepts as opposed to having properties. We could say that an adjective 'a' names the concept of a-ness, or whichever noun form is most appropriate, and then adjust definition (a) to:

defn.(a'): An adjective A is autological if, and only if, it falls under a concept which it names

On this formulation the paradox solver will be forced either to deny the definitions, a move she is simply not in a position to make, or to declare that there is no such concept as heterologicality. But this second response is also beyond the pale, for to suggest that there is no concept of heterologicality implies that the word 'heterological' is in some important sense meaningless and/or unintelligible. And if this is the case how is it that we can both mean and understand sentences containing the word? – true sentences even. We can imagine a community where the word is used, and in whose dictionary we find (a') and (b), but where the paradox is never noticed. Is it to be suggested that even in this community there is no concept of heterologicality?

There is, then, no such thing as a solution to the Grelling conundrum:⁸ the rules lead to contradiction if we ask a certain, slightly bloody-minded question, and that is all there is to the matter. It is senseless to continue to ask whether or not 'heterological' is heterological,

⁸ As a further illustration of the unsolvability of the heterologicality problem, consider the following introduction of the concept applying, this time, to predicates not adjectives:

(Schema H): If X denotes the predicate 'pr', then X is heterological if, and only if, it is not the case that 'pr' pr

From this it is immediately visible that 'is heterological' is heterological if, and only if, it is not the case that 'is heterological' is heterological, no counter argument being possible here.

as if there were some hidden fact out there waiting to be discovered. The tendency to assume that there must be a solution most likely lies in the tacit assumption that there must always be some fact regarding whether or not an adjective is heterological – that there must always be an answer to the question ‘Is A heterological?’. I hope I have said enough to dislodge this prejudice. The only type of solution appropriate here, if circumstances arise in which we *require* an answer to the question of whether or not ‘heterological’ is heterological, that is if the paradox were to become a practical problem, is to take a practical and fully ad hoc decision either to call ‘heterological’ heterological or alternatively to call it autological – a move which would amount to adding a third clause to the definitions (a) and (b) along the lines of:

defn.(c): Definitions (a) and (b) do not apply to ‘heterological’; ‘heterological’ is heterological (autological)

0.3

It should be fairly obvious by now what I intend to say about the liar: that, like the paradoxes above, it fails to fit the standard model, and that therefore it has no solution as such things are traditionally conceived. There is, though, an immediately visible difference between the previous puzzles and that of the liar: In the case of the liar there are no *stipulated* rules to which we can turn as definitive of what is to count as a true sentence or proposition. Over the centuries, however, several people have attempted to provide some sort of definition of truth. Most famous, perhaps, is Aristotle’s dictum:

To say of what is that it is not, or of what is not that it is, is false, while to say of what is that it is, or of what is not that it is not, is true⁹

As an alternative, Tarski suggests:

A true sentence is one which says that the state of affairs is so and so, and the state of affairs indeed is so and so¹⁰

Both of these formulations would, I am fairly confident, be more or less universally accepted; their main drawback being their vagueness. In an attempt to tighten things up we can, however, bring quotation marks into the game, and remark that, ignoring complications

⁹ Aristotle, ‘Metaphysics’ Γ, 7, 27 in Aristotle 1984 p1597

¹⁰ Tarski 1983 p155

of tense, indexicals and demonstratives, the assertoric sentence 'p' says that p. And from here we arrive very swiftly at Tarski's convention (T):

(Schema T): If X denotes the (assertoric) sentence 'p', then X is true if, and only if, p

Now I don't want to get embroiled in the minimalism debate by suggesting that convention (T) exhausts the meaning of the word 'true', but I do suggest that the convention is fully faithful to the 'definitions' offered by Aristotle and Tarski. I also suggest that the inferences it codifies – the inferences of quotation and disquotation – are inferences that we do actually and unthinkingly make when using the truth predicate. This much, I take it, is undeniable. The problem, of course, is that from this schema we can derive the paradox-generating premise (α) stated at the start of this essay. And so the convention, when allied with classical logic, yields a contradiction on being faced with a liar sentence. Thus, it is commonly assumed, we must reject at least some of its instances, and the fact that the convention, once understood, would be immediately and unreservedly assented to by any man or woman on the street who has not read the literature on the liar is overridden.

The object of our investigation at this point is the correct rules of inference which surround the truth predicate. And depending on one's view of the nature of meanings, it will be held either that these are determined by, follow from, the meaning of the truth predicate, or alternatively that those rules are in part constitutive of the meaning of the truth predicate. Either way, what reason do we have to suppose that these rules must be consistent with all worldly facts, in particular with the existence of the liar sentence? The rules for SLC eligibility are not consistent with the existence of Mrs Fineline, and the rules for heterologicality are not even consistent with the existence of the very adjective they define. What reason, then, do we have to suppose the T-schema to be incorrect?

Consider the possibility of a community which speaks a fragment of English in which the word 'true' does not exist. And suppose that one day a member of that community decides to introduce the word by means of the T-schema (in much the same way that we could introduce the word 'heterological' by means of the H-schema suggested in footnote eight). The T-schema for the community's language now becomes a stipulation, and to reject it as incorrect would not be possible. But how would the meaning of a sentence of theirs containing the word 'true' differ from the meaning of its homophonic translation into

English? And if it is acceptable to allow that there is no solution to their liar paradox, why suppose there must be one to ours?

It is my suggestion that the appropriate response to the liar is not to fault a premise or an inference, but rather that solving the liar will be a purely practical matter unavoidable only in very unusual circumstances.¹¹ And that if circumstances do arise in which a solution is required, then a fully ad hoc decision to call any particular semantically problematic sentence either true or false will not only be all that is necessary, but all that is possible.¹² This contended, however, it is not my intention in this essay to argue positively that there could be no standard model solution to the liar – my concern is simply to combat the prejudice which insists that there must be such a solution out there to be found. To this end I

¹¹ Compare Popper 1954:

Socrates. Is there no way whatever of eliminating these paradoxes [the liar paradoxes]?

Theaetetus. There is a very simple way, Socrates.

Socrates. What is it?

Theaetetus. Just avoid them, as nearly everybody does, and don't worry about them.

Socrates. But is this sufficient? Is it safe?

Theaetetus. For ordinary language and for ordinary purposes it seems sufficient and quite safe. At any rate, you can do nothing else in ordinary language, since paradoxes can be constructed in it.

(Popper 1954 p167)

¹² Circumstances can be dreamt up in which liar-type problems become practical issues requiring resolution, as opposed to matters we can merrily ignore. Consider the following tale adapted from Kripke 1975: A journalist named Jones writes about Watergate. His line on the scandal is considerably more sympathetic to Nixon than most other journalists, and as a result Nixon claims:

(1) Everything Jones writes about Watergate is true

But then in his final article on the matter Jones states:

(2) The majority of Nixon's assertions up until this point regarding Watergate are false

Nixon takes offence at this remark and decides to sue. In the courtroom the following facts are ascertained: Firstly all Jones' other assertions regarding Watergate are true, and secondly, apart from assertion (1), Nixon's claims regarding Watergate are evenly balanced between the true and the false. Nixon's lawyers then reason as follows: Suppose the contentious remark, claim (2), is true. Then all Jones' Watergate assertions are true which makes Nixon's assertion (1) true, and so Jones' assertion (2) false. Thus claim (2)'s truth implies its falsity and so the claim must be false and Jones is proved guilty of libel.

But Jones' lawyers ask the judge to suppose that claim (2) is false. It then follows that (1) is also false which in turn makes (2) true. Thus (2) is true, and Jones is proved innocent.

As can be seen, then, the judge is going to have something of a headache. Jones can incontestably be proved both guilty and innocent – his guilt follows logically from his innocence and vice versa. Thus neither of the two possible adjudications that the judge might make has any more claim to correctness than the other, but nonetheless he is required to make a decision which will have important consequences for all concerned. Pity the man. (Perhaps he ought to wash his hands of the matter and call for an expert witness who claims to have found the standard model solution to the liar). We can see, then, that whilst an ad hoc decision to call any particular problematic sentence either true or false may be all that is possible, this fact may prove highly lamentable in certain unusual circumstances.

want, in the next chapter, to develop an observation which expands and, I hope, supports the alternative diagnosis outlined above.

Chapter 1

1.0

In chapter 0 I made a number of suggestions about the logic of various paradoxes; in this chapter, I shall take up, and expand upon, just one of those: the remark, made with reference to the SLC and Mrs Finline, that ‘all there is to say about the matter is that this is the rule and that these are its consequences’. This may not have been particularly conspicuous at the time, but it does, I think, point towards an important area which demands exploration.

1.1

So let’s return to take a harder look at the SLC paradox. Recapping: the SLC entry criterion does not yield a coherent pronouncement in the case of Mrs Finline. Therefore, because there is nothing more to SLC-eligibility than is given in the entry criterion, and nothing remaining for us to discover about the matter, it follows that there is no answer to the question ‘Is Mrs Finline eligible for the SLC?’. There is no fact of the matter of Mrs Finline’s SLC-eligibility, and so the appropriate response to the paradox is to cease asking the question ‘Is Mrs Finline eligible for the SLC?’.

It is crucial for everything I want to say in this chapter that this analysis is not mistaken for a denial of excluded middle (where I take excluded middle to be whatever is denied by allowing that some object may neither have nor not have some property). To say that there is no answer to the question of Mrs Finline’s eligibility is not to say that she is neither eligible nor not eligible. Indeed, a denial of excluded middle would constitute a standard model solution to the paradox: it would invalidate the move from ‘If Mrs Finline is eligible then she is ineligible’ to ‘Mrs Finline is ineligible’. We must take the full force of the fact that there is *nothing* to be said regarding Mrs Finline’s SLC-eligibility.

But there are, however, certain things that we can say about the surrounding situation. We might truly say, for instance, that the SLC criterion does not rule in favour of Mrs Finline, and also that it does not rule against her. Again, we might truly say that we cannot consistently call Mrs Finline either SLC-eligible or SLC-ineligible – that is, if we assume the one then the other follows in a certain immediate manner. Such claims are not (directly) concerned with Mrs Finline’s SLC-eligibility, but rather are descriptive of what happens

when we apply the rule to the case of Mrs Finline. I want to introduce some terminology to mark this distinction: An assertion about a person's SLC-eligibility is an *object* level assertion; an assertion about the application of the rule to a particular case is a *meta* level assertion. (When we make meta assertions we do not assert the conclusion of an argument in which we have used the entry criterion (an object level argument), rather, in many circumstances, what we report is the argument itself.) And I may now reformulate my diagnosis of the SLC paradox in this new jargon: Solving the paradox would consist in finding the correct object level assertion to make (a denial of excluded middle would be an object level claim); realising that there is no solution to the paradox consists in recognising that there is nothing to be said about Mrs Finline at the object level (there is no object level fact in her case), that all there is to say about her situation is located at the meta level.

To all this, I want also to add a further claim: that the sentence 'Mrs Finline is eligible for the SLC', like the sentence 'Mrs Finline is not eligible for the SLC', has no (standard) sense; SLC-eligibility cannot significantly be applied to or withheld from Mrs Finline. This additional claim does not follow immediately from (though it does lie easily with) what has gone before; I shall not justify the claim here, but rather in the following section 1.2.

Now unfortunately it is quite natural, when predicating some meta property or other of an object, to use a form of words which would normally indicate that an object level assertion is being made. It is quite natural, for instance, to use the words: 'Mrs Finline is neither eligible nor ineligible for the SLC' not to deny excluded middle, but rather to say that the entry criterion rules neither for nor against in her case. In an ideal world we would always add the rider 'in the sense that the entry criterion does not rule either way' to such an assertion, but the fact that we do not is apt to occasion a certain amount of confusion (even in the speaker – he may hesitate over what may be inferred from his assertion). In particular, we may fail to make a distinction between 'N is eligible for the club' and 'The club's entry criterion rules in favour of N'. Certainly, a person is eligible for the SLC if, and only if, the entry criterion rules in his or her favour, but this must be understood as meaning: 'If the one side is true then so is the other', for their negations are not interchangeable: it would be correct (*true*) to claim the second negation for Mrs Finline and incorrect (*senseless*) to claim the first. The predicate 'the SLC entry criterion pronounces in favour of' has a wider

domain of (sensical assertive) applicability than does the predicate ‘is SLC eligible’ (in its standard sense).^{13,14}

And it is important to highlight the fact that despite the possibility of making a meta claim using a form of words which would normally indicate an object level claim, very different things will follow from the different assertions. If I truly say: ‘Mrs Finline is not eligible for the SLC’ to mean that the entry criterion does not rule in her favour, it does not follow from what I have said that Mrs Finline is, after all, eligible for the SLC. The entry criterion: ‘A person is eligible for the SLC if, and only if, they are secretary to a club they are ineligible to join’ contains object uses of the predicate ‘eligible’, and so it licences inferences only between object level assertions. Meta assertions, no matter in what form of words they are couched, are not subject to the logic of the entry criterion.

Returning now to Grelling’s paradox, I concluded, in section 0.2, that:

There is, then, no such thing as a solution to the Grelling conundrum: the rules lead to contradiction if we ask a certain, slightly bloody-minded question, and that is all there is to the matter. It is senseless to continue to ask whether or not ‘heterological’ is heterological, as if there were some hidden fact out there waiting to be discovered.

We may now rephrase this thought as follows: Because the rules do not pronounce in the case of the adjective ‘heterological’, there is no fact of the matter of the heterologicality of that adjective; there is nothing to be said at the object level regarding the adjective ‘heterological’. But as with the SLC paradox, this diagnosis is not a denial of excluded middle – I am not saying that ‘heterological’ is neither heterological nor not heterological. The question ‘Is ‘heterological’ heterological or not?’ is to be rejected, not answered; it is not even to be answered with the word ‘Neither’ (indeed such a response would be a contradiction!).

¹³ Again, I have not yet justified the assertion that ‘Mrs Finline is not eligible for the SLC’ is senseless (see section 1.2 below), though I have justified the assertion that it is incorrect (see section 0.1 above).

¹⁴ This is not to say, of course, that the SLC entry criterion has an applicability restricted to those persons for whom the assertion ‘N is SLC-eligible’ has a (standard, object level) sense. On the contrary, the entry criterion is fully applicable to Mrs Finline. (In the entry criterion the phrase ‘Mrs Finline is SLC eligible’ appears only as one side of a biconditional).

Interestingly, in the case of the Grelling, there is not one problematic adjective but two: things also go awry if we look at the adjective ‘autological’. We may not be led, if we ask after the heterologicality of the adjective ‘autological’, to a contradiction, but then nor are we led to an answer. ‘Autological’ is heterological if, and only if, it has no property which it names, i.e. if, and only if, it is not autological, i.e. if, and only if, it is heterological. We can see, then, that applying the definition of heterologicality to the adjective ‘autological’ does not lead to a contradiction, but rather it leads to the partner in crime of a contradiction: a tautology. The rule may not tie itself up in quite the same type of contradictory knot for ‘autological’ as it does for ‘heterological’, but nonetheless something is not quite right. I want to suggest that, in just the same way as there is no fact of the matter of the heterologicality of ‘heterological’, there is no fact of the matter of the heterologicality of ‘autological’. The intuition here should be clear: the rules no more pronounce in the case of ‘autological’ than they do in the case of ‘heterological’. For the moment, this suggestion will be left without any more solid support; section 1.2 will explore, and hence provide a firmer basis for, the intuition.

Again paralleling the SLC, I want to make the further claim that we may not significantly apply or withhold the property of heterologicality to or from the adjectives ‘heterological’ and ‘autological’ (but again I defer a full justification of this to section 1.2). If there is no fact of the matter of X’s heterologicality, then the sentence ‘X is heterological’ has no (standard) sense.

Now despite the silence surrounding the heterologicality of both ‘heterological’ and ‘autological’, we may say a number of interesting things about them at the meta level – about how the rule (the definition of heterologicality) applies to the cases of ‘heterological’ and ‘autological’. Regarding ‘heterological’ we may truly say that it ‘has no object level’, and that it may not consistently be called either heterological or not heterological (meaning here that if we assume the one then the other follows in certain direct manner); regarding ‘autological’ we may truly say that it has no object level, and that it may consistently be called either heterological or not heterological (calling it the one will not, of itself, commit us to calling it the other). We may also, of course, make meta assertions in unproblematic cases: we may, if we like, say that the rules pronounce in favour of heterologicality in the case of the adjective ‘rude’. And as with the SLC, it is not unnatural to use a form of words which would most normally indicate an object assertion in order to make a meta assertion.

We might, for example, choose to use the words ‘‘Heterological’ is neither heterological nor not heterological’ to make the claim that there is no fact of ‘heterological’'s heterologicality; this is perhaps confusing, and it should be emphasised that such a meta-use of the words ‘is neither heterological nor not heterological’ would not constitute an answer to the object level question ‘Is ‘heterological’ heterological?’. Indeed, in such circumstances, the words ‘is neither heterological nor not heterological’ would be a fused idiom (if not fused, then a contradiction!), and would certainly not be subject to the logic of the definition.

A natural objection to the (as yet not fully supported) claim that we cannot significantly predicate heterologicality (SLC-eligibility) of ‘autological’ (Mrs Finline) might be expressed as follows: ‘There is, you maintain, a standard, object level sense of the predicate ‘is heterological’. Why, then, cannot one take the predicate in this standard sense and declare, significantly, that the adjective ‘autological’ is, in the standard sense, heterological?’

The response here is straightforward: one cannot magically glue a sense to a predicate, carrying it with us into whatever new sentence we please. It is as absurd to say: ‘You understand the word ‘heterological’, don’t you? – well then I am using it in just that sense which you understand’ as it is to say: ‘You understand the words ‘soul’ and ‘purple’, don’t you? – well then how is it that you don’t understand me when I say that my soul is purple, for I am using the words in quite their ordinary senses?’.¹⁵

But pressing the objection we might ask about the person who falsely believes that Mrs Finline is secretary not just to the SLC but also to another club for which she is ineligible, and who therefore asserts that Mrs Finline is eligible for the SLC? Or about the person who has no firm beliefs about Mrs Finline’s secretarial positions, but nevertheless asserts that she is SLC-eligible? Or even about the person who is mistaken about the SLC entry criterion and so is led to assert that Mrs Finline is eligible? Are not all these people’s claims precisely object level assertions about the eligibility of Mrs Finline? Well no, they are not. They are, in fact, an unusual type of nonsense – a nonsense uttered by a person under the false belief that they are talking sense.

¹⁵ See Wittgenstein 1963 §117

At first sight this may seem implausible – if a person means something by a sentence, then surely it must make sense.¹⁶ But what is this ‘means something by’?¹⁷ Consider the following example: ‘I have minus two books on my table’. Certainly, this is a nonsense; it cannot be understood. But what about: ‘The number of books on my table is equal to the greatest root of the equation $x^2+4x+4 = 0$ ’? Here it is not obvious that the sentence is nonsense, but nonsense it certainly is. The only difference is that we can see how someone might be mistaken in thinking that it made sense (they believe that the equation has a non-negative integer as a solution). The very same goes for the person who utters ‘Mrs Fineline is eligible for the SLC’ because they have inadequate or faulty beliefs about the situation: we can understand how they might be led to declare the sentence, but this does not mean that we can give the sentence itself any significance. The only response we can make to an utterance of this type is to try to repair the utterer’s epistemic failings. In these circumstances, we cannot give any sense at all, object or meta, to the sentence uttered.

1.2

The previous section is, I hope, reasonably intuitive. But, up to a certain point, that is how it remains: intuitive. In this section I aim to provide a more substantial basis for picture drawn in the previous.

What needs support in the first instance are the meta claims. It is all very well to claim that there is nothing to be said at the object level about Mrs Fineline, or about ‘heterological’ or ‘autological’ – that there is no answer to the question ‘Are these adjectives heterological or not?’ (‘Is this person SLC-eligible or not?’) – but it has not been made clear why exactly this is so. Is it possible that there are other adjectives or people ‘without an object level’? We need a more precise understanding of the meta claim ‘The rules do not pronounce in case X’. Likewise for the other meta claims such as ‘We may not consistently call Mrs Fineline either SLC-eligible or SLC-ineligible’, ‘We may not consistently call ‘heterological’ either heterological or not heterological’ and ‘We may consistently call ‘autological’ either heterological or not heterological’. These claims were passed off as meaning that a

¹⁶ See Wittgenstein 1981 §247

¹⁷ This ‘means something by’ means roughly (is best understood *in this context* as meaning), I suggest, ‘intends their words to be seen in the light of a certain practice, and believes that their linguistic action is a legitimate move in that practice’. But that one intends an action to be seen as a part of a certain practice and believes that it is a legitimate move in that practice does not, of course, guarantee that that action is in fact a legitimate move in that practice. See Wittgenstein 1981 §297.

particular object level sentence does or does not yield its contradictory ‘in a certain direct manner’, but what is this ‘certain direct manner’?

A central source of concern about these meta claims is that if we are working in the vicinity of an inconsistency, we may derive anything we like.¹⁸ Suppose Mrs Thickline is secretary to an elitist golf club which she is ineligible to join. Then she satisfies the SLC entry criterion, and so she is SLC-eligible. But we may also reason as follows: Mrs Finline is secretary only to the SLC, therefore Mrs Finline is SLC-eligible and Mrs Finline is not SLC-eligible (the paradox reasoning), therefore Mrs Thickline is ineligible to join the SLC (from a contradiction anything follows). Combining the two, we see that the SLC entry criterion does not yield a coherent ruling even for Mrs Thickline. Further, if we take ‘the rules are such that we may consistently say that p’ to mean ‘‘not-p’ does not follow logically from the rules and the worldly facts’, then we may not consistently say anything. (Removing the ‘and the worldly facts’ from the above phrase will not, of course, turn the trick – in the case of the SLC we would then report that we may consistently call Mrs Finline SLC-eligible (the paradox relies upon the worldly fact that she is secretary to no other club), and in the case of heterologicality we would still have to report that the rules oblige us to say everything and that we may not consistently say anything).

But Mrs Thickline *is* SLC-eligible, and so we do not want to say that that the entry criterion does not yield a coherent ruling in her case, that there is no object level fact here. Further, we want to say that one may not consistently declare a person to be (in)eligible only in unusual cases.

We want a precise sense of ‘rules pronouncing in favour of’ which is true of just those persons who are SLC-eligible. How, then, are we to clarify our meta-level assertions in a non-trivial manner which tallies both with the intuitive understanding which we thus far have of them, and also with the object level facts?

The key here is, I suggest, to look at *how we actually operate at the object level*; to look carefully at the circumstances under which, the reasons for which, we actually assert that a person is SLC-eligible or that an adjective is heterological. How do we actually answer the

¹⁸ See section 2.1 for a more detailed discussion of this principle.

question ‘Is X heterological?’¹⁹ Once we have a solid view of this, things will fall into place:

We answer the question ‘Is N SLC-eligible?’ by looking around to see if N is secretary to a club (s)he is ineligible to join, that is we look for an answer to the question ‘Is N secretary to a club (s)he is ineligible to join?’. And to do this we ask, of each club C for which N is a secretary, whether or not (s)he is eligible for that club. If we obtain a negative answer to any of these questions we answer the initial question ‘Is N SLC-eligible?’ positively, and if there are no such questions (if N is not secretary to any club), or if all such questions are answered positively we answer the initial question negatively. Now this procedure works perfectly well in the vast majority of cases, but it falls down in the case of Mrs Finline because, being secretary only to the SLC itself, the question to which we find we need an answer in order to answer the initial question ‘Is Mrs Finline SLC-eligible?’ is, in fact, the very same question: ‘Is Mrs Finline SLC-eligible?’. In order to give an answer to the question ‘Is Mrs Finline SLC-eligible?’ we already need to have an answer to it, and we are thrown round in a circle. As a result, we are unable to provide an answer.

A precisely parallel situation arises in the Grelling: In order to answer the question ‘Is X heterological’, we need to ask, of each property which X names, whether or not X has that property. Normally this is unproblematic, but in the case of the adjectives ‘heterological’ and ‘autological’ this leads us to ask: ‘Is ‘heterological’ heterological or not?’ and ‘Is ‘autological’ autological (not heterological) or not?’, i.e. to ask the very same questions to which we were initially seeking answers. Thus we need already to have an answer to the question ‘Is ‘heterological’ heterological?’ before we can give an answer to it, and so we are unable to say anything at all in response to the enquiry.

For both paradoxes, we may also see this very same phenomenon coming into play not when we seek answers to questions, but when we look to make assertions. The SLC entry criterion contains, implicitly, what we might call an assertion directive: You may assert that N is eligible for the SLC if, and only if, you have a right to assert that N is secretary to a club (s)he is ineligible to join, and you may assert that N is ineligible for the SLC if, and only if, you have a right to assert that N is not secretary to any club (s)he is ineligible to join. But

¹⁹ This methodology is, I think, important in answering (or removing the impulse to ask) questions in any number of philosophical areas.

this directive gives us no right to assert either way in the case of Mrs Fineline; in order to gain such a right, we need to already have it. And what is more, there is no way by which we can gain a right to assert that any person is eligible or ineligible for the SLC other than by this assertion directive. Similarly, the definition of heterologicality directs us to assert that an adjective A is heterological if, and only if, we have a right to assert that it does not have any property which it names, and to deny that A is heterological (assert that A is not heterological) if, and only if, we have a right to assert that it has some property of which it is a name. But, again, this rule can give us no right to assert either way in the cases of ‘heterological’ and ‘autological’, for we need already to be in a position to assert something about the heterologicality of these adjectives before the directive can tell us what to assert about their heterologicality, but there is no way of getting into such a position other than by following the directive. And so we assert neither that Mrs Fineline is SLC-eligible, nor that she is not SLC-eligible; we assert neither that ‘heterological’ (‘autological’) is heterological, nor that it is not heterological. Nor, of course, do we assert that Mrs Fineline is neither SLC-eligible nor SLC-ineligible, or that ‘heterological’ is neither heterological nor not heterological, rather: *we assert nothing at all in these cases.*

It cannot be stressed enough that these ‘assertion directives’ do represent the procedure by which we actually come to make our claims, the same (inverted) procedure by which we actually seek an answer to the question ‘Is X heterological/Is N SLC-eligible?’. Moreover, it is *the* procedure implicit in the definitions for coming to make assertions or to answer questions. When we (truly) assert that Mrs Thickline is eligible for the SLC, it is because we have a prior right to assert that she is secretary to a club which she is ineligible to join; when we (truly) assert that ‘obscure’ is heterological, it is because we have a prior right to assert that ‘obscure’ does not name any of its own properties.

Taking the assertion making/question answering procedure and the possibility of its failure as well enough identified, the thought will now be this: that we may explicate our meta phrase ‘the rules do not pronounce’ as meaning ‘the procedure fails’.

Shifting perspective now, it is notable that if we wanted to teach someone how to use the word ‘heterological’ without simply presenting him with the definitions, we would teach him precisely the assertion making/question answering procedure described above, and once he

had understood our instructions, had mastered the procedure, he would have a perfect command of the use of the word ‘heterological’. The question answering/assertion making procedure is at the heart of actual meaning of the predicate.

This points to a further suggestion I made in the last section, namely the suggestion that if the heterologicality rules do not pronounce in the case of X (now: if the procedure associated with those rules fails in the case of X), then the sentence ‘X is heterological’ will lack a standard, ‘object level’ sense. And we may now support this claim: it is *because* the standard mechanism by which an assertive use of the predicate ‘is heterological’ is justified fails to render a ground in the case of adjective X (*because* the rules do not pronounce) that the sentence ‘X is heterological’ has no standard sense.

Of course any one sentence may be asserted with the same sense on a variety different grounds (perhaps I assert that ‘pulchritudinous’ is heterological because I have been so informed by Trustworthy Tim, even though I do not understand the word); the point rather concerns the *canonical* means by which the application of a predicate is justified. There is a standard procedure for obtaining answers to the question ‘Is X heterological’, and I want to suggest that the sense of the predicate ‘heterological’ is tied to this procedure. This should not be thought a particularly contentious claim: the interest we have in a predicate, the use we make of it, the response we make to its assertive application and so on will all be linked to the method by which we understand its application to be justified. Indeed, we have noted that learning this justification method amounts, in the case of ‘heterological’ to learning the meaning of the word. (In the case of SLC-eligible, learning the procedure would not be sufficient for understanding – the pupil would also have to know that a person P for whom the sentence ‘P is SLC-eligible’ is assertible may be admitted to the SLC). I do not want to investigate exhaustively what should be said about the exact relationship between sense and canonical justification procedure, but I do want to assert the following: If the canonical method for answering the question of whether or not an object has a property named by a certain predicate fails, *as a matter of logic*, for some object, then the predicate cannot, in its standard sense, be applied to or withheld from that object.²⁰

²⁰ I do not mean to suggest that all predicates P have just one unique canonical procedure for answering ‘Is x P?’; some predicates may have more than one, and this may raise interesting issues. However, for heterologicality and SLC eligibility there is a unique procedure.

The sense of the predicate 'red' is linked to the fact that its assertive application is canonically justified via an act of seeing. It therefore makes no (standard) sense to say that a thought is red (or is not red). *What is novel in the cases of paradox, however, is that the canonical procedure may fail logically without the occurrence of a category error.* The domain of applicability of 'red' (the category to which it applies) is (roughly) the set of objects which are, in principle, visible. There is therefore no threat of the canonical procedure failing for an object within the domain of applicability. However, in the case of the Grelling (the SLC) the procedure fails for an object after whose eligibility we may quite standardly enquire, namely the adjective 'heterological' (the person Mrs Finline). (The definition of heterologicality applies to all adjectives, 'heterological' and 'autological' included; the SLC entry criterion applies to all people, Mrs Finline included). As a result, we may quite properly ask whether or not 'heterological' is heterological, but no answer will be given.

So we have cleared up our worries about the meta idea of 'the rules pronouncing'. The rules pronounce in case X if, and only if the assertion directive gives us a ground in this case, that is if, and only if, we are not sent around in circles by the question-answering procedure. We also have an explanation of how the absence of certain object level facts – the absence of a pronouncement from the rules – leads to the impossibility of significantly applying or withholding the predicate 'is heterological' in those certain cases. That is, we have a link between the claims that there is nothing (correct) to be said about Mrs Finline at the object level, and that nothing *can* be said about Mrs Finline at the object level.

We may also make more precise (explicate) our other meta idea of 'the rules allowing us to consistently call X Y': 'We may (may not) consistently call X heterological' means 'If we assume that we have the right to assert that X is heterological, then this will not (will) of itself provide us, following the assertion directive, with reason to assert that X not heterological'. And in the case of the SLC: 'We may not consistently call Mrs Finline ineligible (eligible)' means 'If we assume that we may assert that she is ineligible (eligible), then this will of itself provide us, following the assertion directive, with reason to assert that she is eligible (ineligible)'.

It is probably worth ending these slightly delicate last two sections with a paragraph emphasising that nothing I have said makes for a solution to the Grelling or SLC paradoxes; each of their premises is perfectly in order and each of their inferences is perfectly correct. It would be a misunderstanding of the remarks of this chapter so far to think either that the SLC entry criterion does not properly apply to Mrs Fineline, or that excluded middle does not hold for SLC-eligibility. The SLC entry criterion applies equally to all people, the definition of heterologicality applies equally to all adjectives, and there is no object which neither has nor does not have some property.

1.3

Returning to the liar, we can, predictably, make some parallel observations. As with the just considered cases of heterologicality and SLC-eligibility, I am going to make a number of suggestions centring around the idea that there is no answer to the question of whether or not the liar sentence is true.

Again, I want to call the absence of an answer to the question ‘Is the liar sentence true?’ the absence of an ‘object level’ fact. There is no fact of the matter of the truth or untruth of the liar sentence; there is nothing to be said here. But there are many things that can be said about the liar at a ‘meta level’, that is we can say interesting things about what happens when we apply the T-schema to the liar sentence. We may want to say, pending clarification, that we cannot consistently call the liar either true or not true, or we may want to say that there is no fact of the matter of the liar’s truth. And we may (perhaps) use the words: ‘The liar is neither true nor not true’ to say the former, or the words ‘The liar not true’ to express the latter, but it should not be forgotten that such meta uses of ‘not true’ and ‘neither true nor not true’ would not, like the object uses, be subject to the inferences of the T-schema.

Paralleling the previous paradoxes, I will refuse to say anything about the liar’s truth or untruth; in particular I will refuse to say that the liar is neither true nor not true (in the standard (object level) sense of the word ‘true’), that is I will refuse to deny excluded middle. I say nothing about the liar at the object level neither through omission and nor due to some bizarre ineffability – rather, it is that there simply is nothing to be said about the truth or otherwise of the liar.

Further, I shall suggest that we cannot significantly predicate the property of truth (or untruth) of the liar sentence. A person who says the words ‘The liar is not true’, intending his use of ‘is not true’ to be subject to the T-schema inferences, is talking nonsense.

Now the logic of the truth predicate is at once extremely simply and extremely complex. The T-schema itself is as simple as could be, but the situations which may develop around its inferences can be highly complicated. The Grelling paradox had just two problematic cases, namely the adjectives ‘heterological’ and ‘autological’, the SLC had just the one (Mrs Fineline); in the case of truth, however, there is an infinity of problematic sentences X where the question ‘Is X true?’ should be rejected as opposed to answered. (To see this, suppose that there are n people who, in their lifetime, say only one thing each: Person i ($1 \leq i \leq n-1$) says that person $i+1$ tells the truth, person n says that person 1 is a liar).

Interesting examples are not hard to come by; here is one which does not compare directly with any of the problematic cases of the SLC or Grelling: Bert says: ‘The next sentence Betty says will be true’. And then before Betty, who is sitting in her living room at the far side of town, says anything, Bert says again: ‘The next sentence Betty says will not be true’. At that point Betty makes the following claim: ‘The last two sentences Bert said were both true’. We may reason about this as follows: The last two sentences Bert said contradicted each other. Thus they cannot both be true. Thus Betty’s sentence is not true. And so Bert’s first sentence is not true, his second true. But, really, what are we to actually understand by these conclusions? Only, I suggest, this: that the only consistent way of assigning truth values to the three sentences is true to Bert’s second sentence and not true to the other two. But a person who asks whether or not Betty’s sentence is true is not asking how it may consistently be assigned a truth value, she is asking whether or not it is *true*. I shall contend that we should reject any enquiry after the truth of Betty’s sentence, that all we can do here is make the meta claim that the only consistent way of assigning a truth value to Betty’s sentence is to call it untrue.

Now we saw in the last section that we were able to circumscribe the problematic cases of the Grelling situation, the cases on which the rules did not pronounce, by considering whether or not a certain procedure, implicit in the heterologicality definition, provides a ground for asserting an object level sentence, a ground for answering the question ‘Is X

heterological?'. And I think that we may fruitfully look for the same kind of guidance in the case of truth and the liar. If we want to understand what is going on with the liar paradox, we should make a close examination of *how we actually use the predicate 'true'* – of the circumstances under which, the reasons for which, we actually apply it assertorically, and of the means by which we actually seek an answer to the question 'Is X true?'.

Implicit in the T-schema is the following assertion directive: Assert that a sentence is true if, and only if, you are entitled to assert the sentence itself; deny that a sentence is true (assert that it is not true) if, and only if, you are entitled to deny that sentence (assert its negation). When answering the question 'Is 'p' true?' we must look to see whether p or not p. Most often, this procedure works perfectly well, but in the case of the liar things are unfortunate. We see that we should assert that the liar (the sentence which says, of itself, that it is not true) is true if, and only if, we are entitled to assert that the liar is not true; in order to answer the question 'Is the liar true?' we are directed to look and see whether or not the liar is not true. So we are sent round in a circle and can gain no right to assert anything about the liar's truth unless we already have such a right; we may not answer the question 'Is the liar true?' unless we already have an answer. There is no possibility, then, of obtaining a right to assert anything about the truth of the liar, of being able to answer the question 'Is the liar true?'.

It might have been noticed that the directive 'assert that a sentence is (not) true if, and only if, you are entitled to assert that sentence (its negation)' does not tell us (directly) when to assert sentences of the form 'Some sentence with property P is true' or 'All sentences with property P are true'. But existential generalisation gives us a hand here: we may assert (deny) 'Some sentence with property P is true' if and only if we may assert the truth of some sentence with property P (assert the untruth of all sentences with property P), i.e. if, and only if, we may assert some sentence with property P (deny all sentences with property P). And we may assert (deny) 'All sentences with property P are true' if, and only if, we may assert, of each sentence with property P, that it is true (deny the truth of some sentence with property P), i.e. if, and only if we may assert each sentence with property P (deny some sentence with property P).

We might now try out this procedure on Betty and Bert: We may assert that Betty's sentence is true if, and only if, we may assert that both of Bert's two sentences are true. And we may assert that both of Bert's two sentences are true if, and only if, we may assert of

each of Bert's two sentences that it is true. Now we may assert that Bert's first sentence is true if, and only if, we may assert that Betty's sentence is true, and we may assert Bert's second sentence if, and only if we may assert that Betty's sentence is not true. But this throws us around in a circle, for it is precisely the right to assert something about the truth of Betty's sentence that we are seeking to gain.

Interestingly, not all problematic truth predications fail to gain an assertibility ground for reason of being sent around in a circle: the procedure may also fail by 'sending us off to infinity'. Consider the possibility of an infinity of sentences S_i ($i \geq 1$) where $S_i = 'S_{i+1}$ is true'. In this case we may assert that S_1 is true if, and only if, we may assert that S_2 is true. And we may assert that S_2 is true if, and only if, we may assert that S_3 is true. And so on. We thus never gain the right (there is no means of gaining the right) to assert anything about the truth or untruth of S_1 , despite the fact that the procedure does not, at any point, return us to the sentence S_1 itself.

Now it is useful, as it turns out, to look at the matter from a slightly adjusted perspective, that is to look not at any particular sentence X and see whether or not the procedure for answering the question 'Is X true?' leads us to a ground for an answer, but rather to start with the set of grounds and move on up to see where they may take us, what assertion rights they give us. So we start by considering those sentences which do not contain the truth predicate and divide them into the true and the false – into sets (T, F) . These are our initial grounds. Next, we use the assertion directive for truth to move to a new pair of sets (T_1, F_1) where T_1 is the set of sentences which may be asserted, given our prior right to assert any sentence of T and to deny any sentence of F , allowing just (at most) one step of the assertion procedure for the truth predicate, and where F_1 is the set of sentences which may be denied from grounds (T, F) allowing just one step of the assertion procedure. Thus if 'Grass is green' is in T then both 'Grass is green' and "'Grass is green' is true' will be in T_1 , "'Grass is green' is not true' will be in F_1 ; if 'Grass is orange' is in F , then both 'Grass is orange' and "'Grass is orange' is true' will be in F_1 , "'Grass is orange' is not true' will be in T_1 . And, supposing Pete has, at some point in his life, said both 'Grass is green' and 'Grass is orange', then 'Some sentence said by Pete is true' will be in T_1 , as will 'Some sentence said by Pete is not true'; 'Every sentence said by Pete is true' and 'Every sentence said by Pete is not

true' will both be in F_1 . Et cetera. Having thus sorted out sets T_1 and F_1 , we naturally repeat the step and move on to a new point (T_2, F_2) . And so on.

Now as this process proceeds, those sentences which end up in one or other of the expanding sets (T_n, F_n) are precisely those sentences for which we gain a right, following the assertion directive for truth, to either assert or deny – precisely those sentences X for whom we may, following the procedure implicit in the T-schema, obtain an answer to the question 'Is X true?'.²¹

(To see that the liar sentence will never enter $T_n \cup F_n$ we need only note that in order to be a member of $T_n \cup F_n$ the liar must first have been a member of $T_{n-1} \cup F_{n-1}$. But then it must have also been a member of $T_{n-2} \cup F_{n-2}$ and so on all the way down to $T \cup F$. But it contains the truth predicate and so it is not a member of $T \cup F$.)

The point of looking at the assertion procedure in this way is that it is under this aspect that Kripke has managed, in his article 'Outline of a Theory of Truth'²¹, to formalise matters. His work, given the potential complexities of the various examples, is something of a godsend, and it is well worth a brief look:

So we adopt Kleene's strong 3-valued logic and consider a first order language L interpreted over a domain D which includes (codes for) its own sentences. Primitive n -ary predicates of L are interpreted by (totally or partially defined) n -ary relations on D . Next, we extend L to a language $\Lambda(S_1, S_2)$ by adding a monadic predicate T whose interpretation is only partially defined as $(S_1, S_2) \in P(D)^2$ (i.e. (extension, anti-extension), $S_1 \cap S_2 = \emptyset$; $P(X)$ = the power set of X ; \emptyset = the empty set). Then we let S_1' be the set of (codes of) true sentences of $\Lambda(S_1, S_2)$, and S_2' be the set of elements of D which are either (codes of) false sentences of $\Lambda(S_1, S_2)$ or are not (codes of) sentences of $\Lambda(S_1, S_2)$.²² (Note that if $p \in S_1$ then $T(p) \in S_1'$).

Having set this up we now proceed as expected:

²¹ Kripke 1975; reprinted in Martin 1984.

²² For reasons that should be clear by now (I do not want to say that a table is not true), I would prefer not to add the second disjunct here. Everything follows through perfectly well without it; I have left it in as Kripke has it in his presentation.

We define the mapping $\varphi: (P(D))^2 \rightarrow (P(D))^2$ such that $\varphi(X,Y) = (X',Y')$ (defined only where $X \cap Y = 0$)

And put $\Lambda_0 = \Lambda(0,0)$;

If $\Lambda_n = \Lambda(X,Y)$, then $\Lambda_{n+1} = \Lambda(X',Y')$

We may also extend into the transfinite, simply taking the union below at limit ordinals.

Now φ is order preserving on $(P(D))^2, \leq$, that is:

$$(A,B) \leq (C,D) \Rightarrow \varphi(A,B) \leq \varphi(C,D) \quad \text{where} \quad (A,B) \leq (C,D) \Leftrightarrow_{\text{defn}} A \subseteq C \text{ and } B \subseteq D$$

Which is to say that if the interpretation of T is extended by giving it definite truth values where it was previously undefined, no truth value previously established changes or becomes undefined. This follows from the Kleene valuation rules. And it has the key consequence that $\varphi^n(0,0)$ is an increasing sequence: $\varphi(0,0)$ extends $(0,0)$ because everything extends $(0,0)$, and the above will serve as an inductive step. (Extension into the transfinite is also increasing). Truth values, once assigned, stay assigned.

Moreover, from any initial argument, φ reaches a fixed point, that is a point $(X,Y) \in (P(D))^2$ such that $\varphi(X,Y) = (X,Y)$; from $(0,0)$ we reach a minimal fixed point, $\varphi^\sigma(0,0)$, of which all other fixed points are extensions (see Kripke 1979 for proof). We thus have a language Λ_σ containing (something like) its own truth predicate. The following definitions are now possible: A sentence is *grounded* if, and only if, it has a truth value in Λ_σ . A sentence is *paradoxical* if, and only if, it has no truth value at any fixed point. A sentence *has an intrinsic truth value* if, and only if, it is assigned that value by at least one fixed point and is not assigned a conflicting truth value by any other fixed point. A sentence is *easy* if, and only if, it is assigned the value 'true' by at least one fixed point and the value 'not true' by at least one other.

Now this formal work is all very pleasing (and it is *very* pleasing), but Kripke is, as he himself freely admits, rather thin on the ground when it comes to philosophical interpretation.²³ In fact, it is far from clear exactly what Kripke's ambition was in writing the paper. He does not subscribe to the standard position assumption that there is a uniquely

²³ Kripke 1975 in Martin 1984 p63: 'I have not at the moment thought through a careful philosophical justification of the proposal, nor am I sure of the exact areas and limitation of its applicability.'

correct solution to the liar paradox²⁴, and he even seems to doubt the existence of a fully universal solution²⁵. Yet despite not working with a standard model background, he offers us no alternative view. Whatever his broader ambition, however, in the paper in question he appears to be doing little more than asking us to look at his construction.

And it is therefore without the slightest hesitation that I would like to incorporate his formal theory into my own diagnosis of the liar paradoxes. We saw in the section 1.2 that a consideration of the procedure by which we come to make claims of the form ‘X is heterological’ or ‘N is SLC-eligible’ is essential to understanding how and why certain questions fail to have answers, and to explaining why certain sentences do not have a (standard) significance. I want to suggest that an identical role may be played, in the case of truth, by the procedure by which we arrive at claims of the form ‘X is true’. And because the procedure for truth gives rise to far more complex situations than do the procedures for heterologicality and SLC-eligibility, it is useful to have a formal theory as an illustration of what is going on.

In what comes next, I do not want to suggest that there is a definitive way in which the precise notion of groundedness (/paradoxicality/intrinsicity) applying to sentences of Λ_σ should or does carry across to sentences of English. What we have is a precise formal picture which mirrors, to a useful extent, the kind of patterns which we are inclined to see developing in English when we ask after the truth of a problematic sentence. Nor even do I want to suggest that there is a uniquely preferred notion of groundedness applying to Λ_σ – we could easily replace Kleene’s 3-valued logic with one of its ‘competitors’ (so long as the monotonicity of Kripke’s construction is not lost). I shall talk of English sentences being ungrounded (paradoxical, intrinsically true etc.), but it is to be understood that in such a context these are vague predicates with borderline cases where intuition (either about the correctness of a translation from English to the formal language, or about the which formal logic to prefer) is unclear and on which we need not necessarily pronounce.

²⁴ ‘I do not regard any proposal, including the one to be advanced here, as definitive in the sense that it gives *the* interpretation of the ordinary use of ‘true’, or *the* solution to the semantic paradoxes’: Kripke 1975 in Martin 1984 p63.

²⁵ ‘The present approach certainly does not claim to give a universal language, and I doubt that such a goal can be achieved’: Kripke 1975 in Martin 1984 p79.

With qualification in mind, I want to make the suggestion that there is no fact of the matter of the truth of ungrounded (English) sentences; the procedure for answering the question ‘Is X true?’ does not allow for a ruling in these cases and so there simply is no answer to the question being asked – we should stop asking it.

But whilst there is nothing to be said about the truth of ungrounded sentences, there is, we noted above, much we can say about them at the meta level. About the liar, we wanted to say that it could not consistently be called either true or false; this can now be clarified to: ‘The liar is paradoxical’. About Betty’s sentence, we wanted to say that the only consistent way of assigning a truth value is to call it untrue; this can now be clarified as: ‘Betty’s sentence is intrinsically untrue’. And about the truth teller (‘This sentence is true’), we wanted to say that we may consistently call it either true or not true; this can now be clarified to: ‘The truth teller is easy’.

Further, and again paralleling what was said in section 1.2 about the Grelling and SLC paradoxes, I want to suggest that we cannot significantly predicate either truth or untruth of ungrounded sentences. A person who says: ‘The truth teller is true’ and intends his use of the predicate ‘true’ to be subject to the inferences of the T-schema (to be an object level use) is talking nonsense.

If we wanted to teach someone how to use the word ‘true’, what we would do, implicitly if not explicitly, would be to teach him the very assertion-making/question-answering procedure which Kripke has formalised, and once the pupil has mastered the procedure, he will have a complete understanding of the word ‘true’. This procedure is central to the very meaning of the word; central in such a way, I suggest, that if it falls apart in a certain case, then it is not clear what can be meant by an application of either the predicate or its negation to that case. The canonical procedure for justifying an assertion of ‘X is true’ fails if X is an ungrounded sentence, and I claim that as a result the sentence ‘X is true’ will have no standard sense in such a case.

The ungrounded sentences are like cogs spinning out of gear; they fail to engage in the standard manner with the world beyond and behind and giving point to our truth attributions. The cost of this lack of engagement is the loss of the possibility of making sensical assertions (or denials) of truth simpliciter.

1.4

Returning to Kripke's article, he seems a little displeased by the fact that:

Liar sentences are *not true* in the object language, in the sense that the inductive process never makes them true, but we are precluded from saying this in the object language... If we think of the minimal fixed point, say under the Kleene valuation, as giving a model of natural language, then the sense in which we can say, in natural language, that a Liar sentence is not true must be thought of as associated with some later stage in the development of natural language, one in which speakers reflect on the generation process leading to the minimal fixed point... The necessity to ascend to a metalanguage may be one of the weaknesses of the present theory. The ghost of the Tarski hierarchy is still with us.²⁶

What, though, is this 'ghost of the Tarski hierarchy'? In Kripke's theory, the predicates 'intrinsically true', 'paradoxical', 'grounded' and so on are all part of the language of the theory, not of the object language. This is thought to pressure the emergence of a Tarskian hierarchy. It is interesting to examine exactly how strong this pressure actually is, and what sort of hierarchy might be required in response.

It is certainly true that the meta predicates do not feature in Kripke's initial language L , and nor do they feature in the finished product Λ_σ . A strengthened liar, Kripke implies, would look like:

(SL): This sentence is not true

where the sense of this 'not true' is 'not made true in the inductive process' (comparable directly with the meta sense of 'SLC-ineligible': 'the entry criterion does not pronounce in favour of') – a meta sense which is also beyond the expressive reach of the object language Λ_σ . But note that this meta sense of '(not) true' is not the standard sense of '(not) true' (the sense which we are attempting to get a clearer picture of), and we would do well to highlight this by using a different word here in place of 'true', perhaps 'madetrue'. We may then construct the following strengthened sentences:

(M): This sentence is not madetrue

(G): This sentence is ungrounded

(I): This sentence is intrinsically untrue

²⁶ Kripke 1975 in Martin 1984 p80

(P): This sentence is paradoxical
all of which might seem to pose some sort of problem.

But even without these awkward sentences, we still want to say that some meta claims, such as 'The liar is ungrounded', are true, and we currently have no formal model for predicating truth of sentences containing meta predicates. With this in view, the obvious step to take, having run the procedure and determined the extensions of the meta predicates, is to extend the initial language L by adding these meta predicates (which will be defined only on object level sentences), and run the process once more. If this is done, the sentence 'The liar is ungrounded' will come out true (T_2 , perhaps (T for the extended language)) and the above problematic sentences will be nice and undefined (the meta predicates were added to L as defined only on object sentences). And we will, of course, gain a new set of (meta) meta predicates that might be labelled 'ungrounded₂', 'madetrue₂' etc., with which we can repeat the procedure to get a further model – the beginnings of something of a Tarskian hierarchy, perhaps. But see that whilst the nth language in this process will contain 'grounded_{n-1}', ..., 'grounded₁' (= 'grounded'), there will, in each language, be only one truth predicate T_n . This reflects the fact that 'The liar is ungrounded' is true in just the same sense of the word 'true' that 'Grass is green' is true. And because a predicate present in two different languages is given the same interpretation in each, it follows that if a sentence of the nth language is in the extension of T_n , then it's equivalent in the m(>n)th language will be in the extension of T_m . Likewise, a sentence in the counter-extension of T_n or undefined for T_n will be in the counter-extension of or undefined for T_m . All that will happen as we move from one language to the next is that the addition of the new predicates will mean there will be new sentences to be assigned (or not assigned) a place in T.

There seems to me to be no harm at all in this slightly complicating extension of the formal model; the meta-predicates are all relatively technical and reflective notions, the presence of subscripts is not inappropriate. We will, in the light of the extension, want to say that the sentence 'The liar is ungrounded (=ungrounded₁)' is true and grounded₂; there is nothing to be said about its groundedness (=groundedness₁).

It is not, as I see it, a significant problem that this extended model does not allow for quantification over the subscripts, that we have no model into which we can translate the sentence 'The liar is ungrounded_n for all positive integers n', for the aim is not to provide a

universal model for all possible predications of truth, it is simply to provide a formal illustration of the standard mechanism by which those predications may get a grip on the non-semantic world. The formal model is useful only as a tool to direct us towards what is the right thing to say about certain sentences; that we can construct sentences which lie outside the scope of the model may be a little unfortunate, but it is far from disastrous. For the model is, after all, simply an object of comparison. Clearly the right thing to say about the sentence ‘The liar is ungrounded_n for all n’ is that it is true and is in some sense grounded, but because such sentences are far removed from the simple phenomenon of the liar paradox, and are interesting only insofar as they bear some distant relation to that phenomenon, there is no very pressing reason to say anything about them at all.

To return to the strengthened liar, it is essential here to remind ourselves that even if we choose to use the words ‘true’ or ‘not true’ to predicate some meta property or other of a sentence, then those uses will not be subject to the logic of the T-schema. For we are not, in these circumstances, predicating *truth*. We can, in a reflective state of mind, use the words ‘The liar is not true’ to mean that the liar is not madetrue, but we are not here using a truth predicate of some metalanguage; the logic of ‘madetrue’ is not given by a ‘meta T-schema’. If X is an assertoric sentence then we have X is true if, and only if, X is madetrue, in the sense that if one side of the biconditional is true then so is the other, but in some cases ‘X is madetrue’ will be false and ‘X is true’ senseless (consider the liar sentence). Recall here the ‘madetrue’ equivalent in the case of the SLC: ‘the rules pronounce in favour of’; a person is SLC eligible just in case the rules pronounce in favour of them (the truth of the one side implies the truth of the other), but the falsity of ‘The rules do not pronounce in favour of N’ does not ensure the falsity of ‘N is SLC-eligible’ as is illustrated by the case of Mrs Fineline. Madetruth is no more truth than the-rules-pronouncing-in-favour-of-ness is eligibility.

Kripke’s worry that ‘Liar sentences are *not true* in the object language, in the sense that the inductive process never makes them true’ stems, I hesitatingly suggest, from the assumption that such a claim (‘The Liar is *not true*’) uses the word ‘true’ in quite its ordinary sense but employs a rather harder sense of the word ‘not’. But I would contend that the sense of the word ‘not’ here is quite unchanged from its use at the object level, rather it is that we are using a completely different (and non-standard) sense of the word ‘true’.

Still, the strengthened liar fanatics might continue to press. What, for instance, am I to say about the sentence:

(S): Either this sentence is not true or there is no matter of its truth

Well, we might translate it into the model (at level 2) as '(S) is not true or (S) is ungrounded₁'. And as (S) is not an object level sentence (it contains the meta predicate 'ungrounded') the predicate 'ungrounded₁' will be undefined for it and so, under the Kleene rules, it will behave exactly like the simple liar and come out ungrounded₂. We might, then, choose to say that either there is no matter of its truth (it is ungrounded₂), or even that there is no matter of whether or not there is a matter of its truth (it is undefined for ungrounded₁). I would tend to prefer the latter, but really what I say here is totally unimportant; the formal model is not particularly helpful (the translation is dodgy), and there is no very good reason why it should be. In fact, it would very probably be misleading to say anything at all here, for that would give the impression that there is something very definite and perhaps important to be said.

The strengthened liar is something which people who are in the business of blocking the liar reasoning must contend with. For if the means by which they block the liar argument itself generates the possibility of a higher level argument to a contradiction, then they have merely replaced one problem with another. But we must not lose sight of the fact that I am not attempting to block the liar argument, and so if the manner in which I diagnose the logic of the liar presents (suggests) the possibility of a strengthened liar, this is not something I need worry unduly about. Indeed, I need not say anything very definite about outlandish sentences such as (S) not just because whatever the correct thing to say is, it doesn't matter if it leads to contradiction, but also because, unlike the object predicate 'true' which has a very precise and simple logic (the T-schema), the behaviour some of the meta predicates (such as 'there is no matter of the truth of') in extraordinary circumstances is much less clear (our uses and intuitions, even when carefully examined, are much less precise), and so we should not suppose that there is a uniquely 'correct' thing to say about bizarre sentences containing meta predicates.²⁷

²⁷ Implicit, however, in the science of the standard model solution seeker, is the assumption that every proper term has its own definite and consistent logic. McGee 1990 p1: 'Each of the sciences aims to find out the truth about its subject matters. One of the subject matters of philosophy is truth. So one of the aims of philosophy is to find out the truth about truth.'

1.5

Now it will very probably not have escaped notice that, unlike that cases of the SLC and the Grelling, my diagnosis of the liar affects our view of certain properties of the object of investigation (the liar sentence (Mrs Fineline, ‘heterological’)) other than the property with which the paradox is directly concerned (truth (SLC-eligibility, heterologicality)). Precisely, the diagnosis bears on our view of the sensicality of the liar sentence. For I have claimed that we cannot sensically predicate truth (or untruth) of ungrounded sentences. But if we now look more closely at the ungrounded sentences, we see that this is precisely what they themselves involve: predications of truth of further sentences which are themselves ungrounded (or truth predications of the form ‘some sentence with property F is true’ where no sentence with property F is true and at least one such sentence is ungrounded, or the form ‘all sentences with property F are true’ where no sentences with property F are untrue and at least one is ungrounded)²⁸. And so we may conclude that ungrounded sentences, the liar included, are themselves senseless.^{29, 30}

Now this might be thought to occasion some concern for my position at this point in the proceedings, for asserting that the liar is senseless might seem to amount to a solution to the liar, whereas I am (or at least purport to be) in the business of challenging the assumption that the liar has a solution. Indeed, it is certainly true that some philosophers have, whilst attempting to solve the liar, made the claim that the liar is senseless (makes no statement, does not express a proposition etc.). But this noted, it is not, on the face of it, obvious how labelling the liar sentence as senseless blocks the paradox reasoning.

There seem to be two points at which the block might come. First is the thought that the T-schema applies only to sentences with a sense, and so if the liar is senseless then the standard

²⁸ This follows from the Kleene valuation rules used in Kripke’s construction. As I have noted before, I do not want to suggest that these rules have any more right to be considered the correct rules than their competitors; so long as the monotonicity of Kripke’s process is maintained, little hangs on the matter. Our idea of groundedness need not pronounce definitively on cases where intuition wavers.

²⁹ We might object that our diagnosis of Grelling paradox also affects our view on a property of ‘heterological’ other than its heterologicality, namely the property of being sensically applicable, in its standard sense, to all adjectives. This is true, but such a property does not seem so weighty (or potentially relevant to the paradox itself) as the property of sensicality is in the case of the liar.

³⁰ Kripke suggests that if we are going to end up saying that ungrounded sentences are meaningless, then we ought not to use the Kleene valuation, as under those rules sentences such as ‘The liar is true or grass is green’ are grounded, and the implication is that if ‘The liar is true’ is senseless, then so must a disjunction of ‘The liar is true’ with any other sentence. Presumably he has some thought regarding compositionality in the background here, but I see no good reason to follow his suggestion.

inferences of the truth predicate do not apply in this case. But why should we think this? Indeed, it is precisely because of what happens when we apply the T-schema to the liar sentence that we have been moved to assert that it is nonsensical. Are we really to suppose that a person who reasons with the T-schema about a sentence, unaware of the contingent fact that it is ungrounded, is doing something wrong? The person who asserts a contingently ungrounded sentence in the false expectation that his sentence is grounded fully intends his use of the predicate 'true' to be subject to the logic of the T-schema. How can it then be incorrect to apply the T-schema to his sentence? The T-schema is applicable in all cases where a standard (object) use is made of the predicate 'is true', and the use (object, meta etc.) made on any particular occasion is determined, in part, by the speaker's intentions. The use made is in no way be determined by remote contingent facts. Back with the simple liar, the T-schema instance 'The liar is true if, and only if, the liar is not true' is to be fully endorsed. To this it might be complained that one cannot endorse a biconditional both of whose sides are held to be nonsensical. But why not? We assert that Mrs Finline is eligible for the SLC if, and only if, she is secretary to a club she is ineligible to join, but the left hand side of that biconditional would, if asserted on its own, be senseless. And since Mrs Finline is secretary only to the SLC, it follows that she is secretary to a club she is ineligible to join if, and only if she is ineligible for the SLC – a perfectly acceptable biconditional with a 'senseless right hand side'. From here we may move without hesitation to assert 'Mrs Finline is eligible for the SLC if, and only if, she is not eligible for the SLC'. (Indeed, the result that the liar is true if, and only if, the liar is not true is a graphic illustration of the fact that the sentence '(L) is true' has no sense – what sense could we possibly make of a sentence which implies its negation and vice versa?).

Secondly, it might be thought that if (L) is senseless, then the move from 'If (L) is true then (L) is not true' to '(L) is not true' will not hold. But, as I have insisted throughout, there is no failure of excluded middle. I do not reply to the question 'Is the liar true?' with 'It is neither true nor not true; it is senseless'. The assertion that the liar is senseless is not an appropriate response to an enquiry after its truth value, for it is not that the liar has the truth value 'senseless' – rather it is that there is no matter of its truth. The move in question would fail to hold only if it were possible for a sentence to be neither true nor not true, but this very suggestion is a contradictory nonsense (can only be given a meta-sense irrelevant to the object level logic at hand). Certainly the inference is from an acceptable sentence to an

unacceptable one, but this does not mean that it is incorrect – such a step will occur in any solutionless paradox (the SLC, for instance).

1.6

Looking back at section 0.3, I claimed that an ad hoc decision to call any particular semantically problematic sentence either true or false would constitute a perfectly satisfactory solution to the liar difficulties, but without a means of identifying (characterising) those sentences which are ‘semantically problematic’, this remains a piecemeal response; the notion of groundedness promises the possibility of a more general approach³¹. Kripke’s work allows for a more general strategy for avoiding inconsistency. We might, for instance, adjust the T-schema to:

(T1): If X denotes the (assertoric) and grounded sentence ‘p’, then X is true if, and only if,
p

(T2): If X denotes an ungrounded sentence then X is not true

This, of course is no more a solution to the liar than an adjustment to the rules for SLC eligibility asserting that Mrs Fineline is eligible constitutes a solution to the SLC paradox, or a decision to let the traveller across is a solution to the bridge paradox. We might equally well have put:

(T2’): If X denotes an ungrounded sentence then X is true
or even

(T2’’) : If X denotes an ungrounded sentence then X is true if it contains an even number of words and not true if it contains an odd number of words
in place of (T2).

These ‘solutions’ identify the problematic sentences, refuse them the normal logic of truth, and make an arbitrary decision to call them either true or not true. And to repeat a point made several times in chapter 0, we need only ‘solve’ the liar in this way if, for some unusual reason, we have a practical need for a consistent truth predicate.³² Such ‘solutions’

³¹ I don’t want to suggest here that there is anything inherently unsatisfactory with a piecemeal response to the liar paradoxes, but only that there may be circumstances in which a more general solution would be desirable.

³² It might be thought that we would need a consistent truth predicate if we wanted to program a computer to reason with truth, but it should not be forgotten that the standard first order propositional calculus (with function letters and individual constants) is, along with the other more interesting first order systems such as ZF set theory and first order arithmetic, recursively undecidable.

will, of course, have slightly odd looking consequences (e.g., on (T2), ‘The liar is not true’ and ‘The liar is not true’ is not true’), but that is the price of consistency.³³

³³ It should not be assumed from the work so far that I regard the semantic paradoxes as mere quirks – it remains open to the person who accepts my diagnosis to link their occurrence with some deep issue in the relation between language and world. Or, indeed, to see them as interesting only as an example of the kind of way language can torment us (Wittgenstein 1978 III-13).

Chapter 2

2.0

There are, of course, objections that can be raised against my no solution standpoint on the liar. In this chapter, I shall consider the two objections that I have found most commonly raised and most stubbornly maintained.

2.1

First, then, is the worry, mentioned in section 1.2, presented by the principle of *ex contradictione quodlibet* (or ECQ for short). This is the principle that everything follows from a contradiction. If I accept that we can validly (*correctly*) derive contradictions, then am I to suggest that *everything* can be correctly derived simply by inferring whatever is desired from the contradiction?— surely not.

Given what has been said in the last two chapters, the most reasonable option available to me here might seem to be the claim that whilst we need not necessarily block arguments with contradictory conclusions (particularly: we should not fault the T-schema in any way in order to block the liar argument), we should take issue with the inference from ‘The liar is both true and not true’ to ‘Grass is orange’. We are allowing the derivation of certain contradictions, but we cannot allow the derivation of arbitrary sentences, and so we must deny ECQ.

But now the worry would be: is this denial of ECQ sufficient?— does it really cut deep enough into the logic of the paradoxes to prevent the derivation of arbitrary sentences? Perhaps we can use the T-schema to argue that grass is orange without passing through the ‘singularity’ of a contradiction. The following reasoning from the existence of sentence (S) (Curry’s paradox) indicates that the worry would be fully justified:

(S) If (S) is true then grass is orange

Now argue as follows:

Assume:

(a) (S) is true

Then substituting for (S) in (a) we obtain:

(b) 'If (S) is true then grass is orange' is true

Now deduce from the disquotational nature of truth that:

(c) If (S) is true then grass is orange

Next, from (a) and (c) we deduce:

(d) Grass is orange

So now we have established that (d) follows from (a). i.e. that:

(e) If (S) is true then grass is orange

So, re-quoting we have established that:

(f) 'If (S) is true then grass is orange' is true

i.e. that:

(g) (S) is true

And finally, from (e) and (g) we conclude:

(h) Grass is orange

It looks, then, as if I must do something more drastic than simply deny ECQ.

But must I? The pressure for positive action here is coming from the assumption that we cannot allow the derivation of arbitrary sentences, that we must always find fault with an argument which concludes that grass is orange, but a little reflection on the last two chapters should make clear that it is, in my opinion, just this assumption that is to be rejected. I neither deny ECQ nor fault the Curry reasoning.

In refusing to fault (solve) the liar (Grelling, SLC) paradox I am in no way suggesting that contradictions are somehow acceptable (or true!). My rejection of the standard model is the assertion that the acceptable may occasionally lead, via acceptable reasoning, to the unacceptable. It is emphatically not the assertion that contradictions are acceptable after all. Curry's paradox, and the argument that the liar is both true and not true and therefore that grass is orange, are merely further examples of the phenomenon of the acceptable leading to the unacceptable. Indeed, the possibility of inferring that grass is orange via the principle of ECQ or the Curry paradox is, from my perspective, simply an illustration highlighting the fact that we have a paradox here – they do not present any concern that has not been (at least implicitly) addressed in chapters 0 and 1. The only work these additions might be thought to do is to enable us to argue for contingently false conclusions as opposed to a contradictory

conclusion. (Though really we could have done this even without ECQ by treating the argument as a reductio of ‘Mrs Fineline is secretary only to the SLC’, or of a contingent premise in a contingent liar).

Still, the possibility of arguing for contingently false conclusions might be thought to be of some extra weight. It might be thought that whilst examples such as the paradox of the SLC may reassure us that not all arguments with contradictory conclusions need be faulted, not enough work is done by those examples to convince us that we may sometimes leave arguments with contingently false conclusions (and which do not proceed by reductio) unblocked. Perhaps the Curry is symptomatic of a deeper worry about the logic of truth than any present in the logic of SLC-eligibility – perhaps we ought, after all, to deny ECQ (so we don’t reason from ‘Mrs Fineline is both eligible and ineligible for the SLC’ to ‘Grass is orange’) and block the Curry (perhaps (but not necessarily) thereby also blocking the liar). In order to dispel this suspicion, consider the following example: A club called The Secretaries and Orange Believers Liberation Club is founded as a protest at once against social elitism and the persecution of people who believe that grass is orange, and has the following entry criterion: A person is eligible to join the SOBLC if, and only if, either they are secretary to a club they are ineligible to join, or they believe that grass is orange (or both). As in the case of the SLC, the SOBLC functions very happily, growing to be so large that they themselves require the services of a secretary. The lady they hire, Mrs Orangeline, happens not to be secretary to any other club. We may now reason about Mrs Orangeline as follows: Suppose for a moment that she is not eligible to join the SOBLC. Then because there is a club for which she is a secretary and which she is ineligible to join, she is, after all, eligible for the SOBLC. Therefore she is eligible for the SOBLC. But, given that she is eligible for the SOBLC, it follows that she believes that grass is orange, for, because she fails to satisfy the first disjunct of the entry criterion, she must satisfy the second. We thus have a perfectly acceptable argument that Mrs Orangeline believes grass to be orange, a contingent falsehood, which neither proceeds by reductio nor makes use of the principle of ECQ.

If we hold to the principle of ECQ, both the SLC and liar paradoxes can be augmented into arguments from contingently true premises (that there exists a liar sentence, that Mrs Fineline is secretary to, and only to, the SLC) to contingently false conclusions (that grass is orange). If we reject ECQ these arguments go only as far as the contradiction, which would

commonly be deemed necessarily false. Setting ECQ aside, the Curry and SOBLC paradoxes are arguments from contingently true premises (that sentence (S) exists, that Mrs Orangeline is secretary to, and only to, the SOBLC) to contingently false conclusions (that grass is orange, that Mrs Orangeline believes grass to be orange). And so under the assumption that the SLC and SOBLC paradoxes have no solution – that the reasoning in those arguments is not to be faulted – then there is, on the face of it, no compelling reason to fault either the liar or the Curry. Or, then, the principle of ECQ.

I should note (rather than hide) here that my position has a certain fairly radical consequence: accepting that we may faultlessly reason from truth to falsehood entails rejecting the identification of a correct (valid) argument/inference with one which is necessarily truth preserving. Grass is not orange, so whether I conclude that it is via the Curry reasoning or via the liar paradox and ECQ, I have an argument with a false conclusion. And if we understand a valid argument to be one whose inferences are necessarily truth preserving, we must, to have arrived at a false conclusion, either have started from a false premise, or have argued invalidly. The further standard model assumption is that there must be at least one specifically identifiable error (false premise or faulty inference) in the argument, and the task of the paradox solver is to find and correct that error. But the whole thrust of my work in the first two chapters has been to undermine the standard model assumptions, *and in tow here comes a questioning of the importance/interest of a notion of valid reasoning on which it is insisted that an argument with true premises and a false conclusion must be invalid.*

If we do insist upon regarding the possibility of a true premise's leading via a passage of reasoning to an unacceptable (perhaps false) conclusion as a sufficient criterion for considering the reasoning invalid, then the force of this essay is that we obtain an extremely thin and uninteresting notion of reasoning invalidity – a notion which does not entail the existence of an 'error' permitting the possibility of a discrimination between alternative strategies for avoiding the undesirable conclusion as 'correct' (identifying and rectifying the error) and 'incorrect' (misidentifying the source of the problem and so introducing a further error which, happily, compensates for the first in this particular argument). And if the possibility of this distinction is lost, one might well wonder what is achieved by calling the reasoning invalid. Certainly, we can choose one strategy or other in order to avoid the

undesirable conclusion, but there is no very good reason to do this unless we have a practical need for a more acceptable conclusion; choosing a strategy for avoiding the conclusion is not a matter of ‘correcting’ the reasoning, rather it is a matter of taking a pragmatic decision to adopt a certain means of solving a (practical) problem.

What is correct, in the cases of the SLC, the Grelling and the liar, is the reasoning which leads us to the contradiction; what is correct in the Curry and the SOBLC is the reasoning which leads us to the contingently false conclusion. The inferences made in these arguments are fully in line with the meaning of the predicates and logical particles in question.³⁴ But what we do not thereby obtain is a demonstration of the truth of a contradiction or of a contingent falsehood.

Leaving aside for the moment the falsity of the Curry paradox’s conclusion, a secondary concern may be with rationality. The following might be held: As rational beings we ought to analyse the logical consequences of our set of beliefs and choose either to accept them or to revise those of our beliefs from which they follow. It is irrational to hold certain beliefs in the knowledge that they have certain logical consequences but to refuse either to accept the consequences or to revise the beliefs.

But why so? Would I be counted rational if I argued that Mrs Finline is both eligible and ineligible for the SLC, that from this contradiction it follows that shooting oneself in the head is good for one’s health and so proceeded to shoot myself in the head? Clearly not. Thus, it seems, I must either reject ECQ or revise some of my beliefs regarding the SLC and Mrs Finline. Which is to say that, on pain of irrationality, I must either reject ECQ or solve the SLC paradox. So if we agree that the SLC paradox has no (standard model) solution, then the existence of Mrs Finline forces us to reject the validity of the principle of ECQ. Strange indeed that such an unassuming lady can have such a profound impact. And no less strange is that even after rejecting ECQ we may still derive the claim that Mrs Finline is both eligible and ineligible for the SLC, the norms of rationality therefore still forcing us to

³⁴ It would be well beyond the scope of this essay to explore thoroughly how we should understand the idea of a ‘correct’ inference other than as necessarily truth preserving (indeed, this is my proposed PhD topic), though the phrase ‘in line with the meaning of the words involved’ would seem to point in the right direction. (Equivalently, perhaps: a correct inference is one licensed by a necessary proposition).

believe that particular flat contradiction (not that I have the least idea what it might be to believe a contradiction!).

Something is clearly wrong. Perhaps the relationship between logical implication and rational belief revision is not so simple as was supposed – consider the following ‘Preface paradox’: A modest philosopher prefaces her book with the acknowledgement that not everything to be claimed therein is true. But being sincere she believes everything written in the book, preface included, and so (knowingly) has a logically inconsistent belief set without, it seems, being in any way irrational. But this possibility is precluded by the naïve norm of rationality suggested above. (Consider similarly the lottery paradox in which it is fully rational to believe, of each lottery ticket, that it will not win, but simultaneously to believe that there will be a winning ticket).

Whatever the correct epistemic aim might be, it does not entail that we should always believe all logical consequences of our beliefs. The SLC and Preface paradoxes provide compelling counterexamples to this supposition. We should not (and in fact do not) accept all logical consequences of our beliefs.

To close this section, it is worth reminding ourselves that no one ever seriously reasons using ECQ (i.e. derives a contradiction and then concludes whatever they like) or with sentences like (S). To do so for any purpose other than as a logical exercise is plainly ridiculous; neither the Curry paradox nor the SLC paradox plus ECQ will convince anyone that grass is orange. Nor will the SOBLC paradox convince anyone that Mrs Orangeline believes that grass is orange. In chapter 1 I diagnosed the liar phenomena as stemming from the ungroundedness of certain sentences; one possible strategy, then, for avoiding arguments in this area with undesirable conclusions would be to refuse to reason about the truth or untruth of ungrounded sentences (about the SLC- or SLBOC-eligibility of Mrs Fineline or Mrs Orangeline). And in fact this strategy is more or less in line with how we do actually behave.³⁵ We do not, in serious (non-philosophical) argument (i.e. in an everyday argument intended to establish the truth of a conclusion), reason about the truth of such sentences as

³⁵ Harman 1986 p16: ‘The rational response [to the liar] for most of us may simply be to recognise that our beliefs about truth are logically inconsistent, agree that this is undesirable, and try not to exploit this inconsistency in our inferences’

(S) or the liar. We do not, in a courtroom, declare ‘If this sentence is true then I am innocent’ and then proceed to deduce our innocence. Not that to do so would be to commit a logical error; it would be to commit an error of absurdity.

2.2

A second (and perhaps not unrelated) objection against a no solution stance on the liar or the Grelling is that such a position is ruled out by any of a number of semantic theories. Indeed, it has been suggested that my whole approach begs the question against certain semantic theories.

Broadly speaking, pressures for solutions arise within semantic theories for English in which predicates (‘is true’ and ‘is heterological’ included) are assigned a type of ‘semantic value’ which precludes them from being paradoxical. The most popular line regarding the semantic value of a predicate is that a predicate determines a general, object independent condition of application – a condition which is indifferent to both the identity of the individuals which meet it, and also to the number of such individuals³⁶. Gareth Evans is among many philosophers who have adopted this idea: ‘To treat an expression as a predicate is to associate with it a certain condition, upon whose satisfaction by objects depends the truth or falsity of the sentences in which the expression occurs’³⁷. Other possibilities for semantic values for predicates include (platonic (in some way)) properties and sets; I shall, however, look more closely only at the idea of application conditions, the corresponding arguments for sets and properties being more or less parallel. (Semantic ideas of this ilk tend to be inspired by the semantics Tarski provides for formal languages, central to which is the idea of predicates being satisfied by (ordered n-tuples of) objects).

As a first shot, the idea of a condition might bear upon the paradoxes in something like the following way:

(P1): Predicates are associated with conditions of application

(P2): Vagueness aside, the application condition of a predicate must, determinately, either be met or fail to be met by any object falling within the domain of applicability of that predicate

³⁶ See, for example, Larson and Segal 1995 p127

³⁷ ‘Identity and Predication’ in Evans 1985 p31

(P3): Simple object-predicate sentences are true just in case the object denoted meets the condition associated with the predicate used

(P4): ‘is heterological’ is, logically, a predicate; ‘‘rude’ is heterological’ is true

Therefore

(C): ‘heterological’ is determinately either heterological or not heterological; the Grelling paradox has a solution

This will not quite do, though, for it should be clear that predicates, conceived, at least, as syntactic items, will not always determine the same condition. ‘James is tall’ could be true on one occasion and false on another – perhaps he is tall for an adult male but short for a basketball player. Moreover, in the absence of any appropriate context, it is plausible to suggest that an utterance of ‘James is tall’ will not be rightly evaluable as either true or false. Perhaps, then, we should adopt the following picture:

Sentences may be used to express propositions, which are either true or false. On any particular occasion in which a simple object-predicate sentence expresses a proposition, the predicate determines a particular condition, and that sentence is true or false (via the proposition it expresses being true or false) depending on whether the object denoted does or does not meet that condition. The adjusted argument that the Grelling must be solved will now be slightly more complicated:

(P1’): Predicate tokens, when forming part of sentence tokens which express propositions (sentence tokens capable of truth or falsity), determine conditions of application

(P2’): Such conditions are, determinately, either met or not met by all objects falling within the domain of applicability of the predicate-type

(P3’): Simple object-predicate sentence-tokens which express propositions are true just in case the object denoted meets the condition determined by the predicate-token

(P4’): ‘is heterological’ is, logically, a predicate. It may be used to express subject-predicate propositions

Therefore

(C’): ‘heterological’ either meets or does not meet the condition determined by each token of the predicate ‘is heterological’ which forms part of a proposition-expressing sentence (a sentence which is true or false). A solution to the Grelling will consist in making clear *when* a token of ‘is heterological’ determines a condition, and, if a condition is determined, *which* condition is determined.

A solution will therefore give an account both of the various conditions (or perhaps the unique condition) which may be picked out by a token of 'is heterological', and an account of how context fixes the condition picked out by any one token (the proposition expressed). (Perhaps each token of 'is heterological' has a tacit numerical subscript attached in a similar way to that in which each token of 'is tall' (arguably) has a tacit (or even pronounced) 'for an X' attached).

(I don't want to suggest that the idea of a condition as the semantic value of a predicate *forces* any sort of 'contextualist' approach, or even that all those who have been attracted to contextualist approaches (Burge, Parsons, Simmons amongst others) would endorse the idea of a condition as the semantic value of a predicate. Rather, I want merely to indicate a plausible (and not atypical) way in which a philosopher might be pressured into solving the semantic paradoxes by his commitments regarding a semantic theory for English).

Of course it remains open to the semantic theorist, if we insist upon stipulating what is to count as correct use in such a way that leads inextricably to paradox (consider the community which introduces the word 'true' to their dialect by an explicit setting down of the T-schema), to deny (P4) (or (P4')). In such cases the syntactical predicate 'is true' or 'is heterological' will not be a logical predicate; sentences involving such predicates will not express propositions, they will be incapable truth or falsity. Utterances of these sentences will now be speech acts whose content is not available for logical analysis in the way that the contents of the majority of assertoric utterances are. This shows the limits of stipulation: we cannot go around dictating what is to count as correct use with complete freedom and expect that our dictations will always be sufficient to engage those correct uses with truth and falsity.

But if we do not engage with truth or falsity – if the predicate fails to determine a condition – does this make our sentences meaningless? I don't think that the theorist need be quite this unsympathetic; perhaps, so long as we find the predicate intelligible, so long as we find it useful in our interactions, he can allow that it has a sense even though we fail to express propositions by its use. It would be a brave but implausible theorist who grabbed the nettle and suggested that a predicate might be completely meaningless even though it is fully assimilated into our linguistic activities. (Although, on the other hand, if we can understand, believe etc. sentences containing predicates which fail to determine conditions then we will,

on the account developed, be forced to deny that the propositional attitudes are relations to propositions. Perhaps some might consider this worse than calling the non-proposition-expressing sentences meaningless (and so incapable of being believed, understood etc.).

It does, however, seem rather unlikely that the semantic theorist would allow that *our* truth predicate is not a logical predicate, that no token (utterance, inscription) of the word ‘true’ determines a condition and so the word may not form part of any proposition-expressing sentence³⁸. If this were so then a sentence such as ‘‘Grass is green’ is true’ could not itself be true – a rather unpalatable result. Therefore, the theorist will be forced to solve the paradox of the liar.

As I stressed above, the semantic ideas pencilled above by no means represent the only position which will force us into solving the liar; perhaps no philosopher would agree entirely with the outline I have suggested. What is more important is the acknowledgement that various semantic theories entail that the existence of a solution to the liar is a prerequisite for us speaking truly when we use the word ‘true’. I want now to examine how a philosopher with such a theory might respond to the thoughts of chapter 0.

So let’s consider again the paradox of the Secretaries’ Liberation Club. I have argued, and will take it as read, that this paradox has no (standard model) solution. The stipulated entry criterion does not determine a condition (property, set) which every person determinately does or does not satisfy (have, is or is not a member of). We seem, then, to be in the position considered above where our stipulations of what is to count as correct use (here correct use of the predicate ‘is eligible to join the SLC’) fail to determine a condition. Therefore it must be held (see footnote 5) by the semantic theorist that we cannot speak truly when we use sentences of the form ‘N is eligible to join the SLC’. This might seem more than a little strange: that the sentence ‘Mrs Thickline is eligible to join the SLC’ is not true even though she satisfies the entry criterion for the SLC, everyone accepts that she is eligible to join the SLC, and indeed she proceeds to join the SLC to everyone’s approval. In fact, it seems

³⁸ Of course, it would be possible for a theorist to deny that our truth predicate is, logically, a predicate without committing himself to the claim that no sentence containing the word ‘true’ may express a proposition – he could, perhaps, claim that the word has a logical role other than that of a standard predicate. I shall not investigate the possibility here; it doesn’t seem to be a popular option; the obvious choice for the logical form of the sentence ‘Sentence X is true’ is the subject predicate form.

impossible to deny that Mrs Thickline *is* eligible to join the SLC even if the sentence that says the same (or, at least, appears to say the same – perhaps it says nothing at all in some weighty sense of ‘say’!) is not true (or at least not appropriate for truth). How very strange.

We might next ask whether the SLC paradox infects just eligibility for the SLC, or whether the disease spreads to bar all our sentences containing the word ‘eligible’ from expressing propositions. Have we reason to suspect the latter? Perhaps we do, for it is the concept of eligibility which appears to lie at the heart of the paradox. If we take sentences like ‘N is eligible for C’ to have the logical form of a simple two place predicate relating two objects, then the predicate ‘eligible’ should determine a condition on ordered pairs of people and societies. But the SLC paradox demonstrates that it does not do this (consider <Mrs Fineline, the SLC>), thereby contaminating not just the one place predicate ‘is eligible for the SLC’ but also the two place predicate ‘is eligible for’. We thus seem to be driven into saying that no sentence of the form ‘N is eligible for C’ expresses a proposition and so no such sentence is ever, strictly speaking, true or false. (Or worse still, if the theorist makes determining a condition necessary for having meaning, all sentences of the form ‘N is eligible for C’ are meaningless).

Perhaps this is a pill the semantic theorist will elect to swallow (though surely not without a considerable grimace). But even if he does, he will not, I assumed above, swallow the suggestion that the same might go for the predicate ‘is true’. The assumption is that his semantic theory demands a solution to the liar, and it makes this demand as a result of the fact that we can speak truly using sentences in which we predicate truth of other sentences. He therefore owes an account of the deep semantic difference between eligibility and truth (or at the very least between SLC-eligibility and truth) which indicates why our uses of the word ‘true’ determine conditions but our uses of the word ‘eligible’ do not. But there seems to be no clear intuitive basis for any important distinction: it seems just as much a fact that ‘Mrs Thickline is eligible to join the SLC’ may be true or false as it is that ‘The first sentence on page 120 is true’ may be true or false. Pending the provision of a good reason to support the claim that the two sentences have a wildly different semantics, the semantic theorist is in something of a quandary. There are paradoxes of both truth and eligibility, why does one demonstrate that we have not yet fully discovered the condition determined by a predicate but the other demonstrate that a predicate fails to determine a condition?

In the light of a semantic theory which claims that the possibility of using a predicate to speak truly entails that the predicate must pick out a condition which in turn entails that any paradox surrounding the predicate has a solution, there seem to be four responses available to the problem posed to the theorist by the SLC and liar paradoxes:

- (1): Thump the table and insist that there is a relevant difference between truth and eligibility which the theorist will no doubt think of soon, a difference which will justify his opposing attitudes towards the two paradoxes (and hence towards the semantics of truth and eligibility)
- (2): Change his mind on the SLC paradox and insist that it too must have a solution
- (3): Change his mind on the liar paradox, allowing that it has no solution and thus commit himself to the consequence that no sentence predicating truth of another sentence can itself be either true or false
- (4): Adjust his semantic theory to allow that the existence of a solution to a paradox involving a certain predicate is not a necessary prerequisite for using that predicate to speak truly

Option (1) is rather dogmatic, and will remain so until the purported difference is brought to light. In the meantime, we have been given no good reason to regard the SLC and liar paradoxes differently, and so, assuming that the SLC has no solution, there is, for the moment, no compelling reason from semantic theory to think that the liar must have a solution. This assumption about the SLC can be denied, of course (option (2)), but I have already argued at some length against such a move (section 0.1 etc.); I shall not add to those arguments here. The third option entails denying the obvious (that we can speak truly using the word 'true'); I don't think it would be seriously countenanced by any philosopher. Which leaves option (4), the option I would advise. But if this option is preferred, we lose the very reason for insisting that the liar must have a solution which has been our object of investigation in this section. The idea that, whenever a truth is uttered using a predicate, the predicate determines a condition which all objects in its domain of applicability must determinately either satisfy or fail to satisfy will have to be rejected.

Certain semantic theories do entail, given reasonable assumptions about what we can truly say, that the liar must have a solution. But these theories do not, as I see it, provide a basis for an adequate response to my arguments that the liar need not have a solution, for they are

extremely uncomfortable with the case of the Secretaries Liberation Club – a case which, on the face of it, appears to provide some evidence against the theories (motivation for taking option (4) above). If, however, the semantic theorist sticks to his guns (taking options (1), (2) or (3)) he must respond to charges either of deep implausibility (options (2) or (3)), or of failure to provide sufficient basis for his position (option (1)).

2.3

There are, I don't doubt, other possible objections to a no solution position on the semantic paradoxes. But apart from the concern that an argument with true (acceptable) premises and a false (unacceptable) conclusion must contain an error, the concern which is the central motivation for the standard position (and also the concern highlighted by such arguments as the Curry paradox), all objections seem to be made off the back of some pre-developed philosophical position on, e.g., the nature of properties. My entire argument against the standard position (see, particularly, chapter 0) is aimed at allaying this first concern, but I am not much moved by responses of the second sort. The more theory needed to support an objection, the less need there is to answer it.

Chapter 3

3.0

In chapter 0 I argued that there is no good reason to suppose that the paradox of the liar must have a solution of a particular standard kind (a ‘standard model’ solution), indeed I suggested that solving the liar would be a purely practical matter. In chapter 1 I developed certain of the observations of chapter 0 into a more comprehensive alternative diagnosis of the phenomenon of semantic paradox. Subsequently, in chapter 2, I defended those previous chapters from two commonly raised objections – objections whose force is an insistence that the liar must after all have a standard model solution.

So far, then, I have presented my case almost entirely in opposition to the standard position. So much so, perhaps, that I might be taken to be suggesting that that position represents the only alternative to my approach. Certainly, the majority of the philosopher-logicians who have worked on the semantic paradoxes have, in their mostly very brief remarks on the philosophical basis of their work, indicated that they are working with standard position assumptions (or something very like them) in the background, but this rule is not without exception. Its most notable exception, a philosopher who takes trouble to explore not just the various ‘solutions’, but also what accepting a particular ‘solution’ would amount to, and indeed who rejects the standard position, is Quine. His view on the matter, as spelt out in ‘The Ways of Paradox’³⁹, provides an interesting alternative both to my position and to the standard position.

3.1

In this paper, he distinguishes three types of paradox: falsidical paradoxes, veridical paradoxes and antinomies. In each case we have a surprising conclusion arrived at by seemingly valid argument from plausible premises; the situation is veridical if, despite initial appearances, the conclusion is perfectly acceptable, it is falsidical if we can come to recognise an error in the premises or reasoning, and if we can neither accept the conclusion nor fault the argument, we have an antinomy:

A veridical paradox packs a surprise, but the surprise quickly dissipates itself as we ponder the proof. A falsidical paradox packs a surprise, but it is

³⁹ Essay I in Quine 1976

seen as a false alarm when we solve the underlying fallacy. An antinomy, however, packs a surprise that can be accommodated by nothing less than a repudiation of part of our conceptual heritage.⁴⁰

The liar is, for Quine, a paradigmatic case of antinomy. The conclusion, being a contradiction, is unacceptable, but, he suggests, we cannot fault the means by which we arrive there. The contradiction is inherent in our current concept of truth and a solution to the liar will not therefore be a matter of discovering and describing hidden structure in that current concept, as is supposed by the standard position, rather it will involve adjusting our current concepts to remove the flaw of inconsistency. Once the concepts have been revised and the dust has settled, that is once a solution has been generally adopted and become the new 'common sense', future generations (the date offered is AD 4000!) will consider the liar to be a straightforwardly falsidical paradox.

This view of the matter (and mine, in fact) is given some intuitive appeal by the 'Of course! – that's it' phenomenon. Comic mathematical 'proofs' that one equals zero are plentiful, and in some the fallacy (perhaps a division by zero) is well hidden, but when the sleight of hand is pointed out there is, perhaps, a characteristic 'Of course, that's the error' experience. Similarly for other falsidical paradoxes. But the fact that the liar has resisted solution for thousands of years coupled with the immense complexities of the current best solution attempts, make it reasonable to suppose that no solution exists which would produce this effect. And such an absence would lie very easily with the view that a solution would represent a conceptual shift and not a discovery of an error in the argument as it currently stands.

By means of comparison with the contemporary problems, Quine offers the example of Zeno's paradox of Achilles and the tortoise. Zeno pictures Achilles pursuing a tortoise and argues that whenever Achilles reaches a place where the tortoise once was, the tortoise has moved on, and so Achilles can never overtake the tortoise. Zeno's error, of course, is the assumption that an infinite succession of positive time intervals must add up to eternity. Now Quine suggests that for the ancients the argument did not represent a falsidical paradox but an antinomy - that the ancient mathematical concepts were such that an infinite sum of positive terms was always infinite. But as mathematics has developed, perhaps under

⁴⁰ Quine 1976 p9

pressure from just such examples, our concepts have shifted in such a way that a twenty first century sum of an infinity of twenty first century positive terms need not be infinite, and we can view Zeno's paradox as straightforwardly falsidical.

Another example Quine discusses is Russell's paradox which is presented as being in a state of limbo between antinomy and falsidical paradox. At the time of its discovery, the contemporary 'habits of thought' of the mathematical community contained an 'overwhelming presumption' of the validity of existence principles allowing the existence of the class of non-self-members, and so at that time Russell's paradox represented an antinomy. But we are not yet in a position to view the paradox as straightforwardly falsidical, for the dust thrown up by the dropping of the antinomy-bomb has yet to fully resettle on the conceptual landscape of set theory – the 'crisis in the evolution of thought' is not yet over.

These examples allow a slightly more careful analysis of Quine's position. Quine does not use the words 'true', 'false', 'correct' or 'incorrect' to describe the conceptual schemes of which he talks, but he does describe them as being 'discredited' by the presence of contradiction. Is there, then, any sense in which the paradoxes do have unique correct solutions – is there such a thing as the correct conceptual scheme? And if there is a correct solution, what is its source? (to what are our conceptual schemes responsible?)– does Quine subscribe to a realism which allows unconceptualised things-in-themselves to provide the source?

The following passage should make his position on set concepts more clear:

[Regarding a response to Russell's antinomy] We are driven to seeking optimum consistent combinations of existence assumptions, and consequently there is a great variety of proposals for the foundations of general set theory. Each proposed scheme is unnatural, because the natural scheme is the unrestricted one that the antinomies discredit; and each has advantages, in power or simplicity or in attractive consequences in special directions, that each of its rival lacks.⁴¹

Again:

⁴¹ Quine 1976 p16

Common sense is bankrupt, for it is wound up in contradiction. Deprived of his tradition, the logician has had to resort to mythmaking. That myth will be best that engenders a form of logic most convenient for mathematics and the sciences; and perhaps it will become the common sense of another generation.⁴²

Regarding set theory, then, Quine is a pragmatist as opposed to a realist, there being no uniquely correct set theory corresponding to the investigation-independent nature of what it is that sets *really* are. Further, the suggestion seems to be that we do have a certain amount of freedom in responding to Russell's antinomy, for no one solution is uniformly superior to all others. The pragmatic approach fails even to identify a best candidate which we could see as the (pragmatically) 'correct' solution. Are we therefore to assume that he takes a parallel line on the semantic paradoxes?

Three different responses to the liar are considered: Firstly, that we should cease to use the word 'true' and all related locutions, secondly that we should refuse to apply such locutions to expressions which themselves contain examples of the same locutions, and thirdly that we should introduce a Tarskian hierarchy by adding numerical subscripts to each such expression on each occasion of use. Quine prefers the third, *but he does so purely on grounds of 'cost'* – the first two methods of antinomy-elimination are both 'less economical' than the method of subscripts. Thus resolution even of the liar is simply a utility optimisation game: we want to adjust our semantic conceptual scheme in a way which eliminates the inconsistency but minimises the loss of reasoning we find useful and so would like to keep. In this light, then, there is no such thing as the correct solution to the semantic paradoxes (where 'correct solution' is to be understood here as 'solution which corrects the error'); at best there will be one concept-adjustment which stands out above all others as the least costly, according to whatever scales of cost we find important. The very idea of a 'correct' solution is appropriate only to falsidical paradox not to antinomy.

3.2

It is unfortunate (for me, at least) that Quine puts forward no argument in 'The Ways of Paradox' by way of support for his central paradox/antinomy distinction. One gets the

⁴² Quine 1995 p27: 'Whitehead and the Rise of Modern Logic'

distinct impression that he adopts the position he does not in response to any pressure local to the paradoxes, but rather because it is the one which best fits with his broader commitments. Which makes an investigation of his surrounding philosophy a prerequisite for an understanding any deeper than the one gained so far. Interpreting the Quinean metaphysics/epistemology is a hazardous occupation, but the following is a concise stab at the matter which I hope does not fly too far from the mark:

The totality of our beliefs – our ‘total science’ – is to be pictured as a web of logically interconnected statements whose ultimate purpose is to predict future experience in the light of past experience. Included in this network are not only the statements of everyday, of history, geography, natural science etc., but also those of mathematics and logic. Now this web is in more or less continual flux as we bring it into line with our ongoing experiences, certain of which may, when coupled with the demand for consistency, occasion widespread adjustments over the web. Importantly, no statement is immune from having its truth value revised, not even those statements we are inclined to call ‘necessary’. There is no principled distinction to be made between necessity and contingency – these ideas consist only in the adoption of a revision policy which shields certain areas (subjects, types of statement) at the expense of others.⁴³

Now with this empiricist picture comes a rejection of metaphysical realism – a denial of the idea of a unique correspondence between our linguistic expressions and things in themselves. And the temptation is to describe Quine’s alternative position as the claim that the notions of truth and falsity are internal to, relative to, our evolving theories (conceptual schemes, world sciences etc.), but this is not a description which he would allow, for despite rejecting the realism of an unconceptualised world which lies behind our evolving theories and to which those theories are in some way responsible, he remains a robust realist with respect to his own current science – a theory with a fully bivalent truth predicate which has a fully transparent disquotation schema. This relativism-realism tension is visible in several of his papers:

To say that the statement ‘Brutus killed Caesar’ is true, or that ‘The atomic weight of sodium is 23’ is true, is in effect simply to say that Brutus killed Caesar or that the atomic weight of sodium is 23. That the statements are about posited entities, are significant only in relation to a surrounding body

⁴³ See ‘Two Dogmas of Empiricism’ in Quine 1980

of theory, and are justifiable only by supplementing observation with scientific method, no longer matters; for the truth attributions are made from the point of view of the same surrounding body of theory, and are in the same boat.⁴⁴

But this does not amount to any sort of relativism about truth, because:

We continue to take seriously our own particular aggregate science, our own particular world-theory... Unlike Descartes, we own and use our beliefs of the moment, even in the midst of philosophising, until by what is vaguely called scientific method we change them here and there for the better. Within our own total evolving doctrine, we can judge truth as earnestly and absolutely as can be; subject to correction, but that goes without saying.⁴⁵

For Quine, then, there simply is no Olympian point of view from which he can see himself as a relativist; the phrase 'from my current perspective' is completely vacuous when it comes to assessing statements made from within divergent conceptual schemes (the ancient mathematical scheme, for example). But nonetheless, metaphysical realism remains rejected; when we adjust our conceptual scheme, we are directed not by 'how the world is' but rather, where rational, by purely pragmatic considerations.

As for the matter of whether or not there is such a thing as the uniquely correct world theory, a unique total systemisation of science 'conformable to the past, present and future nerve-hits of mankind'⁴⁶ towards which we are fumbling, Quine is quite frank in accepting that the constraints of experience will not single out any one theory. Global science is underdetermined by empirical evidence, and this shows that there are 'various defensible ways of conceiving of the world'⁴⁷. What is not defensible, then, is the insistence that the semantic antinomies have, from our current perspective (a vacuous qualification), uniquely correct solutions; our semantic concepts may evolve in response to the presence of paradox, but there is no uniquely correct direction in which they must (even pragmatically speaking) go.

⁴⁴ Quine 1960 p24

⁴⁵ Quine 1960 p24

⁴⁶ Quine 1960 p23

⁴⁷ Quine 1992 p102

3.3

A critique of Quine's wider position into which his view of antinomy fits is beyond the remit of the present chapter. Rather, I shall examine the key area where his response to paradox is importantly different from my own: the suggestion that concepts (or conceptual schemes) containing contradictions are 'bankrupt' and must be modified.

The working assumption at this point is the non-standard position that our current truth concept, as is reflected in the unrestricted T-schema, is inconsistent and may lead to contradiction. But how does this 'bankrupt' or 'discredit' the concept? The obvious suggestion given Quine's (and equally my) pragmatic attitude regarding the appropriate response to paradox, is that a concept is bankrupt if it is incapable of serving the purpose for which we want to use it. This is why the discovery of Russell's paradox presented a crisis for the set theory of the time: the inconsistency in our naive set concept made it incapable of serving the purposes it was hoped a set theory would serve, namely to provide some sort of foundation for mathematics (though whether or not such a hope was or is misguided is another question). Therefore, in order to proceed with the stated aim, a means of eliminating the paradox was essential. But for what purpose is it essential to solve the liar?

Quine's story of antinomy-induced conceptual revolutions may seem less convincing for the liar than it is for Russell's paradox. The liar has been a recognised part of our languages for nearly twenty five hundred years, where then is the 'crisis in the evolution of thought' that a discovery of antinomy is supposed to ignite? The very fact that we have not solved the liar, that we have not adjusted our concepts, indicates that for all our standard practical purposes our inconsistent concept serves us perfectly adequately. So, given Quine's pragmatism in the face of antinomy, why does he unquestioningly insist that the liar must be solved?

The answer to this should hopefully be clear from the previous section: Quine cannot be as pragmatic regarding *whether* to respond to antinomy as he is regarding *how* to respond to antinomy because our world theory is, at heart, a machine of prediction. The issue of resolving the antinomy must, despite its artificial and pragmatic nature, be addressed, because the presence of inconsistency reduces a theory to triviality and uselessness: an inconsistent theory predicts everything. In our current conceptual scheme, the existence of a liar sentence implies that grass is orange, and this is what bankrupts. Antinomies are therefore 'very much less frivolous' than they might at first appear.

(Of course, it is not set in concrete (for Quine, at least) that everything must follow from a contradiction – we might choose to adjust our logic in such a way as to isolate the contradictions, preventing them from infecting the entire system (no part of our theory, not even our logic, is immune to revision). But this is just another way of solving the liar antinomy (an interestingly different way in that future generations will view the liar as a veridical paradox as opposed to a falsidical one) which, given the centrality to our world theory of the logical principles we would be obliged to disavow, looks unlikely to represent the most economical (c.f. Priest and paraconsistent logics).)

3.4

How, then, might Quine respond to the central argument of this essay, an argument whose conclusion is that we may be as pragmatic regarding whether or not to respond to the semantic paradoxes as we may be regarding how to respond (if we so choose)? In particular, what is he to say about such examples as the Secretaries' Liberation Club?

There seem to be three options open to the Quinean at this point:

- (a): Find an important difference between the case of the SLC and that of the liar which allows us to accept the SLC paradox without reducing our world science to triviality in the way that an acceptance of the liar would.
- (b): Retract the assertion that the liar trivialises our world theory.
- (c): Bite the bullet and admit that, like the liar, the SLC paradox infects our world science with contradiction and therefore bankrupts our concept of eligibility. The SLC paradox is a full blown antinomy which must, like the liar, be resolved.

None seems an easy course to take. Anyone preferring option (a) would owe an account of the important difference they are postulating – and they would seem to have their work cut out, for whilst there may be no knockdown argument against the possibility of such a difference, there is certainly no obvious basis for a distinction which would do the work required. The possibility of the taking option (b) depends heavily on the details of the Quinean's view of the nature of her world theory – a matter it is not my intention to explore in any further detail – but it looks very much as if this option does not represent a move she can make without abandoning tenets central to her position. Which leaves possibility (c), the option I estimate Quine himself would take.

I have argued at some length against the supposition that there is a *standard model* solution to the liar paradox, and used the fact that the SLC paradox has no such solution as a weapon in my assault. The hard line Quinean of option (c) will agree with me that it is ridiculous to suppose that the SLC paradox has a standard model solution, but disagree with my suggestion that, so long as Mrs Fineline does not apply for membership, there is not the least reason to adjust the SLC entry criterion (the concept of SLC-eligibility). Rather, she will maintain, the paradox *does* provide good reason to adjust our concept of SLC-eligibility, but the reason is not local to matters facing the SLC admissions committee, it is identifiable with a much broader concern for consistency. And accordingly, I would hazard, the adjustment will most likely be made not directly to SLC-eligibility, but to the notion of eligibility simpliciter. (Perhaps we might devise a way of attaching subscripts to our uses of the word 'eligible' which removes the inconsistency (though this would be a far from simple matter).)

To this hard line Quinean position (the one which takes option (c) above), I have no reply (beyond the label of implausibility) local to the paradoxes. (Not that the label of implausibility is not totally without force – the Quinean's philosophical baggage forces her to maintain that the SLC paradox bankrupts our concept of eligibility, a concept which must therefore be adjusted in order to avoid a disaster in our mappings of the world!). If the idea of a theory of global science as a machine of experience prediction is developed to a point at which we are forced to solve the SLC paradox, then my only response, if I choose to respond, can be to question that notion of global science. Such a move would, as remarked above, be well beyond the scope of this essay.

The arguments of this essay will not, therefore, persuade the hard-line Quinean that we may solve the liar when and only when we have a particular need for a consistent truth predicate. But she is already persuaded of the essay's main conclusion – that the semantic paradoxes have no standard model solution – and so is my ally for the majority of the battle. I shall not, therefore, press the issue with her any further. It is of some interest to note that there is a position one can defensibly take which, despite involving a rejection of the standard position, does not coincide fully with my own.

References

Aristotle

1984 *The Complete Works of Aristotle*, Barnes, J. ed. Princeton; Princeton University Press

Austin, J. L.

1976 *How To Do Things With Words*. Oxford; Oxford University Press

Cervantes, Miguel de

1993 *Don Quixote de la Mancha*. Ware, UK; Wordsworth Editions Limited

Chihara, Charles

1979 'The Semantic Paradoxes: A Diagnostic Investigation' *The Philosophical Review* 88: 590-618

Evans, Gareth,

1985 *Collected Papers*. Oxford; Oxford University Press

Harman, Gilbert

1986 *Change in View*. Cambridge Mass.; MIT Press

Kripke, Saul

1975 'Outline of a Theory of Truth' *The Journal of Philosophy* 72: 690-716.
Republished in Martin 1984

Larson, R. and Segal, G.

1995 *Knowledge of Meaning*. Cambridge Mass.; MIT Press

Martin, R. L. ed.

1984 *Recent Essays on Truth and the Liar Paradox*. Oxford; Oxford University Press

McGee, Vann

1990 *Truth, Vagueness, and Paradox*. Indianapolis; Hackett Publishing Company

Popper, Karl

1954 'Self-Reference and Meaning in Ordinary Language' *Mind* 63: 162-169

Quine, W.V.

1960 *Word and Object*. Cambridge Mass.; MIT press

1976 *The Ways of Paradox and other essays*. Cambridge Mass; Harvard University Press

1980 *From a Logical Point of View*. Cambridge Mass; Harvard University Press

1992 *Pursuit of Truth*. Cambridge Mass; Harvard University Press

1995 *Selected Logic Papers*. Cambridge Mass; Harvard University Press

Sainsbury, R. M.

1988 *Paradoxes*. Cambridge; Cambridge University Press

Simmons, Keith

1993 *Universality and The Liar*. Cambridge; Cambridge University Press

Soames, Scott

1999 *Understanding Truth*. Oxford; Oxford University Press

Tarski, Alfred

1983 *Logic Semantics and Meta-Mathematics*. Indianapolis; Hackett Publishing Company

Wittgenstein, Ludwig

1963 *Philosophical Investigations*. Oxford; Blackwell

1978 *Remarks on the Foundations of Mathematics*. Oxford; Blackwell

1981 *Zettel*. Oxford; Blackwell