

Inconsistency Without Contradiction

Achille C. Varzi

Department of Philosophy, Columbia University, New York

(Published in *Notre Dame Journal of Formal Logic*, 38:4 (1997), 621–638.)

1. Genuine worlds, ersatz worlds

David Lewis has argued that impossible worlds are nonsense.¹ If there were such worlds, one would have to distinguish between the truths about their contradictory goings-on and contradictory falsehoods about them. One would have to distinguish between, say, the alleged truth

- (1a) In so-and-so world pigs can fly; and in that world pigs also cannot fly,

and the contradictory falsehood

- (1b) In so-and-so world pigs can fly; and it is not the case that, in that world, pigs can fly.

But this—Lewis argues—is preposterous: there is no such distinction to be drawn. (1a) and (1b) are equivalent.

Of course, Lewis does distinguish between the following:²

- (2a) In the world of Sherlock Holmes, Watson limps; and in that world Watson also does not limp.
(2b) In the world of Sherlock Holmes, Watson limps; and it is not the case that, in that world, Watson limps.

(2b) is contradictory. But (2a) is, in a sense, true: there is a discrepancy in Conan Doyle's writings. In one of the Holmes stories, Watson limps because

¹ See, e.g., Lewis 1983a: 21.

² See, e.g., Lewis 1983b: 279 (my example is slightly different).

of an old war wound in the leg; in other stories Watson's wound is located in his shoulder, and he does not limp.

According to Lewis, the difference between the two cases is to be explained in terms of the different nature of the operators 'in so-and-so world' and 'in the world of Sherlock Holmes'.³ 'In so-and-so world', if we take it seriously, is a restricting modifier: it limits the domain of quantification of modal discourse to the domain of the world in question (for a modal realist, that amounts to limiting the domain of quantification to a certain part of all that there is) but has no effect on the truth-functional connectives. Thus,

- (3a) In so-and-so world, it is not the case that p
- (3b) It is not the case that, in so-and-so world, p

are equivalent, and so are

- (4a) In so-and-so world, both p and q
- (4b) In so-and-so world p , and in so-and-so world q .

This grants the equivalence of (1a) and (1b), as well as their equivalence with

- (1c) In so-and-so world, pigs both can and cannot fly.

In this sense, 'in so-and-so world' behaves like 'in Australia' or 'last Tuesday'. By contrast, 'in the world of Sherlock Holmes' is on a par with 'in the Holmes stories' or 'Arthur wrote that': these are not restricting modifiers and do not pass through the truth-functional connectives. Hence (2a) and (2b) are not equivalent, and neither is equivalent to

- (2c) In the world of Sherlock Holmes, Watson both does and does not limp.

Conan Doyle might have contradicted himself about Watson's limp (2a), but he has never written a contradiction (2c), and we certainly need not use a contradiction (2b) to report that fact.

There is, no doubt, a boundary to be drawn somewhere. Not every sentence modifier passes through the truth-functional connectives, and some do. Yet there is room for philosophical disagreement on where exactly the bound-

³ See, e.g., Lewis 1986: 7n.

ary should be drawn. Lewis is no ersatz; for him worlds are not like stories or story-tellers. *Therefore* ‘in so-and-so world’ is not on a par with ‘in the world of Sherlock Holmes’, and therefore there is no room for genuine impossible worlds. But that would hardly count as a conclusive argument. Ersatz modal realists would simply be unmoved by it. And *therefore*, ersatz modal realists might in principle find room for impossible worlds.

My purpose in this paper is not to defend or attack the ersatz conception—the view that unactualized worlds are merely linguistic descriptions of the ways this world of ours might have been. Rather, I am interested in that conception insofar as it pertains to the challenge raised by Lewis. Suppose we do *not* treat ‘in so-and-so world’ as a restricting modifier. Suppose we take Conan Doyle to be describing a *bona fide* world, though a world that is impossible because of some unfortunate discrepancies. What are we to make of the truths and falsehoods in such a world? Can we distinguish its contradictory truths from contradictory falsehoods about it? I want to say that we can. If an impossible world is one in which there are discrepancies of the sort illustrated by the Watson example—a world in which certain facts both do and do not obtain—then we can keep such worlds under logical control exactly as we can keep the Holmes stories under control. There are discrepancies, but these discrepancies are local and do not force incoherence upon our discourse about them because they can be explained away by reference to an underlying set of coherent goings-on.

One may also view this as a concern about the prospects of defining semantic and logical notions with regard to stories rather than genuine worlds, though this conception would have to be refined in a number of ways which I shall not consider here. In fact, in the following I shall mostly speak of models rather than worlds or stories. A model may be inconsistent in that it may involve some discrepancies. And the line I wish to consider is that the truths and falsehoods of an inconsistent model are determined, in some way to be made precise, by the truths and falsehoods of the family of its consistent fragments.

2. Quarantining Inconsistencies

Lewis has his own way of presenting this line of reasoning, so long as the model is merely a story or a fictional world.⁴ It proceeds from equating discrepancy (inconsistency) with *ambiguity*. Cassical models are perfectly unam-

⁴ The suggestion is detailed in Lewis 1982; see also 1983b.

biguous: each sentence can be evaluated in just one way. By contrast, inconsistent models are ambiguous and may support contrasting evaluations. Some sentences may come out true on some disambiguations and false on others, and therefore they may be regarded as both true and false in the model. However, this is not to say that anything goes: other sentences may come out unambiguously true, or unambiguously false. So inconsistencies need not lead to logical chaos.

The world of Sherlock Holmes is ambiguously defined in this sense: it would be more correct to speak of the worlds of Sherlock Holmes, in the plural. Each of the sentences

(5a) Watson limps

(5b) Watson does not limp

is both true and false if these worlds are taken collectively, because each is true on some ways of sorting out the ambiguities in the Holmes stories and false on others. Yet in no story will you find the plain contradiction

(6) Watson both does and does not limp

which is therefore unambiguously false. The world of Scherlock Holmes is no logical chaos.

This way of dealing with inconsistencies is of course reminiscent of Stanislaw Jaskowski's discussive logic: the propositions that hold in a piece of discourse are those that are put forward by at least one of the participants, who may contradict one another while being perfectly self-consistent.⁵ Other authors have proposed similar accounts, typically in connection with a conception of truth-value gluts that is epistemic or doxastic or information-theoretic rather than ontological.⁶ In general, such accounts agree that a model in which contradictions are true is nonsense, but show how one can make sense of a model with contradictory truths. For another example, Nuel Belnap considered this view in relation to the problem of dealing with inconsistent data banks (or explaining how a computer should think, if one conceives of a data bank as an epistemic state of the computer).⁷ Sam and Elisabeth enter the data. Each is

⁵ See Jaškowski 1948.

⁶ See *inter alia* Grant 1975, Rescher & Brandom 1980, Jennings & Schotch 1984, Schotch & Jennings 1989.

⁷ See Belnap 1977.

trustworthy, and the computer is so programmed as to answer Yes to a query if the answer follows from what one of its informers have entered in the data bank. If Sam enters (5a) while Elisabeth enters (5b), there is a discrepancy. It does not follow, however, that the computer should answer Yes to (6).

My own account of this way of quarantining inconsistencies (Lewis's phrase) comes from a different route, and I like to see it as embodying the spirit of Bas van Fraassen's supervaluationary approach.⁸ Generally speaking, a supervaluation is just a way of evaluating sentences on a non-standard model M on the basis of the values determined by a suitable class of standard models, namely those models that represent the various ways in which M can be sharpened. Traditionally this notion has been applied to models whose non-standardness lies in a form of incompleteness rather than inconsistency, i.e., models that are consistent but not fully determined with respect to the semantic value of some atomic sentences. (For instance, a data bank is typically incomplete in this sense, since it does not contain *all* the information about the objects in its domain.) If M is such a model, then its sharpenings are essentially its complete expansions, i.e., those models that correspond to some way or other of arbitrarily filling in the gaps in M , and the supervaluation is naturally defined as the function that registers the pattern of agreement among the valuations induced by such sharpenings:

- (7a) A sentence A is true (false) in an incomplete model M if and only if A is true (false) in every complete expansion of M .

The intuition is that if we get the same outcome no matter how we fill in the gaps, then the gaps don't matter. For example, the world of Sherlock Holmes is incomplete with regard to Watson's feelings about broccoli: maybe he likes broccoli, maybe he does not—we are not told. Clearly, this gap will be relevant when it comes to evaluating a sentence like

- (8a) Holmes smokes the pipe and Watson likes broccoli,

which in fact will be neither true nor false by (7a). However, the gap is irrelevant when it comes to a sentence like

⁸ See Varzi 1994, 1995 and the application in Hyde 1997; similar perspectives are briefly considered in Visser 1984 and in Anderson *et al.* 1992: 523 (with reference to a suggestion by Anil Gupta). As for van Fraassen's supervaluationary semantics, it dates back to his 1966a, 1966b, though similar ideas can already be found in Mehlberg 1958, §29.

(8b) Either Holmes smokes the pipe or Watson likes broccoli,

which in fact will come out true: the determinate truth of the first disjunct is enough to make the disjunction true no matter how Watson feels about broccoli.

In short, supervaluations reduce the problem of evaluating a sentence on a gappy model to that of evaluating it on the model's gapless sharpenings. Now, if M is *inconsistent* rather than incomplete, i.e., if it yields opposite semantic values for some atomic sentences, then a similar insight applies. Only, in this case we have to look for a different sort of sharpening—we have to look for those models that correspond to some way or other of weeding out, as it were, the gluts in M (an operation that dualizes the gap-filling operation on an incomplete model). Thus, the sharpenings will have to be found, not among the expansions of M , but among its contractions, specifically its consistent contractions.⁹ And rather than taking the logical product of the values induced by such sharpenings, in the case of an inconsistent model we may take their logical sum, since each sharpening delivers only part of the truth about the whole model. This leads to the following:

(7b) A sentence A is true (false) in an inconsistent model M if and only if A is true (false) in some consistent contraction of M .

For instance, in the world of Sherlock Holmes a conjunction such as

(9a) Holmes smokes the pipe and Watson limps

will both true and false, because of the glut concerning Watson's limp. However, when it comes to the disjunction

(9b) Either Holmes smokes the pipe or Watson limps,

the glut will be irrelevant: the truth of the first disjunct is enough to prevent the whole disjunction from being false whether or not Watson limps.

We may call an assignment of truth values conforming to (7b) a *subvaluation*, to stress the duality with (7a). Going back to (5a) and (5b), then,

⁹ The terms 'contraction' and 'expansion' are meant to suggest a connection with related work in belief revision theory (see Gärdenfors 1988 and Levi 1980, 1991). However, I will not elaborate upon this connection here.

my point is that the process whereby both those sentences come out true (and false), whilst their conjunction (6) comes out false (and not true), is precisely a subvaluational process of this sort: there is no consistent contraction (no single story) that verifies the conjunction (6), though some such contractions (stories) verify (5a) and others verify (5b).

The point is best appreciated if we consider that, on this understanding, the relationship between (5a)–(5b) and their conjunction (6) is just the dual of the relationship between, say,

(10a) Watson likes broccoli

(10b) Watson does not like broccoli

(neither of which has a definite truth value in the Holmes stories), and their disjunction

(11) Watson either does or does not like broccoli

(which is nevertheless true). This is a well-known fact about supervaluational semantics. If the expansions of an incomplete model M do not all yield the same truth-value, then A is neither true nor false in M (a semantic gap) but the tautologous disjunction ‘ A or not A ’ is true nonetheless. In other words, supervaluational semantics violates the Principle of Bivalence while preserving the Law of Excluded Middle.¹⁰ Dually, then, if M is inconsistent, then A will be both true and false (a semantic glut) when its contractions yield different truth-values, but the contradictory conjunction ‘ A and not A ’ is bound to be false nevertheless.

Of course, as they stand conditions (7a) and (7b) are in need of various refinements, since the gap-filling and glut-weeding operations that they presuppose have only been given an approximate and intuitive characterization. Even so, it should already be clear in what sense this account is similar to the other accounts mentioned above. In each case, the problem of evaluating sentences on an inconsistent (or incomplete) model is reduced to that of evaluating them on a number of consistent (complete) models: as long as we know how to do *that* (and we may well suppose that standard semantics will serve the purpose), we know how to keep our non-standard models under control. This is why it becomes possible to formulate the account in general terms, as in (7a)–

¹⁰ Van Fraassen emphasized the distinction in 1966a, §8; see also McCall 1966.

(7b). And this is what differentiates the account from the restricting-modifier outlook endorsed by other, more familiar treatments, such as Kleene’s three-valued semantics and its four-valued generalizations.¹¹

The formulation given in (7a) and (7b) is actually more general than Jaskowski’s, though it is still rather close to Lewis’s. It does not require that we actually *construe* an inconsistent model as a composition of consistent models (Nicholas Rescher and Robert Brandom have introduced the term “superposition” to indicate this mode of composition¹²), just as one typically does not construe an incomplete model as the result of combining (by “schematization”¹³) a certain collection of complete models. Simply, an inconsistent model is a model with truth-value gluts and an incomplete model is one with truth-value gaps—two opposite ways of generalizing the classical notion of a model in which every atomic sentence is exactly either true or false. Now, classical models induce a unique valuation function assigning a unique truth-value to each sentence of the language. If the model is incomplete (but consistent) there is no complete valuation to be performed. However there are several such valuations that we could tentatively perform, one for each way of completing the model. None of these valuations could legitimately qualify as the valuation determined by the model because none of them reflects *only* what goes on in the model. But their intersection does, and this is exactly what the supervaluationary account (7a) amounts to. Likewise, if a model is inconsistent (but complete), then there are several such valuations that we could perform, one for each way of consistently refining the model. None of these valuations could legitimately qualify as *the* valuation determined by the model because none of them reflects *all* the goings-on of the model. However, their union does, and this is exactly what (7b) amounts to.

What about the case of a model that is both inconsistent and incomplete? (The model of the Holmes stories is like that, since it has at least one glut and very many gaps.) It is obvious that the two strategies can be combined to account for the goings-on in such models as well. In fact, (7a) and (7b) already tell us what to do: if M is both inconsistent and incomplete, neither the subvaluation nor the supervaluation are uniquely defined in terms of classical

¹¹ See e.g. Fitting 1992 and Gupta & Belnap 1993 for reviews and developments. Various semantics for paraconsistent logics could also be viewed in this light: see Priest & Routley 1989 for a survey.

¹² In Rescher & Brandom 1980.

¹³ Lewis 1978 uses ‘superposition’ to indicate this mode of composition.

truth-value assignments; but one can compute the supervaluation on the basis of the admissible subvaluations, or the subvaluation on the basis of the admissible supervaluations. In other words, if incomplete models are allowed to be inconsistent, then (7a) amounts to the following:

- (12a) A sentence A is true (false) in a model M if and only if A is true (false) in some consistent contraction of every complete expansion of M ;

and if inconsistent models are allowed to be incomplete, then (7b) amounts to:

- (12b) A sentence A is true (false) in a model M if and only if A is true (false) in every complete expansion of some consistent contraction of M .

(On the other hand, in the case of models that are both consistent and complete, these conditions reduce to the classical conditions: we are of course to understand contractions and expansions in such a way that every model counts as a (vacuous) contraction and expansion of itself.)

There is but one complication: (12a) and (12b) do not coincide. This can be intuitively verified by evaluating a biconditional such as

- (13) Watson limps if and only if he likes broccoli

with respect to the model of the Holmes stories (where the left-hand side involves a glut and the right-hand side a gap). According to (12a), the biconditional is both true and false. For whether you consider an expansion where Watson likes broccoli or an expansion where he does not, you can always come up with two sorts of contraction: contractions where Watson limps (making (13) true) and contractions when he does not limp (making (13) false). On the other hand, (12b) will evaluate (13) as neither true nor false; for whether you consider a contraction where Watson limps or a contraction where he does not limp, you can always come up with an expansion in which the right-hand side of the biconditional has the opposite value of the left-hand side.

This asymmetry between (12a) and (12b) is disturbing. For if the two schemas do not coincide, then we must choose one or the other on pain of contradiction; yet either choice would seem arbitrary. At the same time, the asymmetry should come as no surprise. It simply reflects the impossibility of performing simultaneous gap-filling and glut-weeding sharpenings. So when the sentence to be evaluated expresses an *equivalence* between gluts and gaps

(as in (13)) the difference must become apparent in some way: on one policy gaps go first, so gluts may prevail; on the other policy it is gluts that go first, and gaps may prevail. (When it comes to sentences that do *not* express an equivalence between gaps and gluts, the two policies would seem to be in agreement, though I will not elaborate on this conjecture here.¹⁴) One could also consider mixed policies intermediate between (12a) and (12b), where some sentences get super-evaluated first while the others get sub-evaluated. I don't know how to choose among all of these slightly different options. But in the present context I would like to leave that issue open. Our question was whether one can find a way of distinguishing between truth and falsehood in an impossible world. And so far, the answer to that question is that we seem to have many ways of doing the job.

3. Ways of sharpening

There are indeed many other, more or less interesting variations of the strategies described above. However I shall not aim at a complete account here, not even a survey of the main options.¹⁵ Rather, I now want to take a closer look at the basic apparatus. In particular, I want to focus on the intuitive idea that to any model there corresponds a class of consistent and complete sharpenings. This idea hides various presuppositions and simplifications, some of which are decisive when it comes to cashing out the logic of a language with *bona fide* inconsistent or incomplete models.

3.1. Existence The first important presupposition is existential: the notions of expansion and contraction are defined only if the class of all models is partially ordered in terms of definiteness. This is intuitively clear, and in some simple cases there is no question as to what the ordering should look like. For instance, suppose we are dealing with propositional models, i.e., assignments of truth-values to unanalyzed (atomic) sentences. In this case, a model is inconsistent or incomplete depending on whether it assigns both values (true and false) or no value at all to some sentences. Formally this means that a model can be any relation between atomic sentences and truth values (and not necessarily a classical *total function* mapping each atomic sentence to a unique truth-

¹⁴ The conjecture exploits a remark of Visser 1984: 194.

¹⁵ I have examined some options in Varzi 1995. In that context I also elaborate on the disagreement between (12a) and (12b).

value). Taking relations to be sets of ordered pairs, the relevant ordering \sqsubseteq is therefore neither more nor less than ordinary set-theoretic inclusion:

$$(14) \quad M \sqsubseteq M' \text{ if and only if } M \subseteq M'.$$

Thus, if M is incomplete (but consistent) its sharpenings are simply its complete supersets; if it is inconsistent (but complete) its sharpenings are consistent subsets; and if it is incomplete and inconsistent, its sharpenings are complete supersets of consistent subsets (or consistent subsets of complete supersets).

When it comes to other sorts of models, however, things are less straightforward. Consider for instance the models of a first-order language. In that case we certainly want to say that the extension and counterextension of a predicate constant in a model's expansion (contraction) must include (be included in) the extension and counterextension, respectively, of that predicate constant in the model itself. If we think of an n -ary predicate as standing for some way of assigning a truth-value to each n -tuple of objects in the domain, then this is indeed the obvious generalization of (14):

$$(15a) \quad M \sqsubseteq M' \text{ only if } P^M \subseteq P^{M'} \text{ for all predicates } P.$$

(I write ' P^M ' to indicate the interpretation of P in M : if M is classical, then P^M is a total function assigning a truth value to each n -tuple of objects in the domain $|M|$; otherwise, more generally, P^M will be a relation between n -tuples and truth values.) This much is clear. But what about the domain of quantification itself? Certainly we want

$$(15b) \quad M \sqsubseteq M' \text{ only if } |M| \subseteq |M'|$$

Shall we allow for proper inclusion? Shall we allow a contraction to drop objects from the domain, or an expansion to add objects?

It would seem so. The model may be inconsistent precisely in that it contains a discrepancy concerning the existence of certain individuals. Suppose the model says that Holmes has a brother, Mycroft, but also that he was an only child. This means that two different sorts of consistent contractions are possible: those in which

$$(16a) \quad \text{Holmes has a brother}$$

is true, and those in which it isn't. And there seems to exist no *prima facie* reason to assume that the latter type of models should all have the same domain as

the former. We can make (16a) false simply by changing the blood tie between Holmes and Mycroft; but we can also falsify (16a) by getting rid of Mycroft altogether. Likewise, if the model is incomplete in that it says nothing about Watson's having a sister, there are two sorts of expansions in which

(16b) Watson has a sister

is true: those in which the sister is some woman already mentioned in the stories *and* those in which she is a totally new character, Lynda.

Now, whether the domain can indeed change as we go from a model to its sharpenings is a question on which I wish to remain neutral here. But *if* we allow for that possibility, then we face an option. Either we evaluate quantified sentences with reference to the domain of the sharpening, or we evaluate them with reference to the domain of the given model (for that, after all, is given). The first option amounts to construing the valuation on a non-standard model on the basis of the *standard* valuations induced by its sharpenings, as tacitly suggested above. This would imply that, in the envisaged situation, (16a) is both true and false and (16b) neither true nor false, respectively. This is intuitive. But it would also follow that

(17a) There exists an x such that $x = \text{Mycroft}$

(17b) There exists an x such that $x = \text{Lynda}$

are both true and false and neither true nor false, respectively. More generally, it would follow that any existential statement $\exists x P x$ may come out (true and) false on an inconsistent model by virtue of positive instances lost in the contraction (for instance, if Mycroft is the only individual with the property P); and it may fail to be (true or) false in an incomplete model by virtue of positive instances acquired in the expansion (for instance, if nobody has the property P except for the added sister). This means that the resulting semantics would not be monotonic, i.e., would not satisfy the conditional:

(18) if $M \sqsubseteq M'$, then every sentence that is true (false) in M is also true (false) in M' .

And this seems to be a high price to pay, for after all monotonicity is implicit in the intended reading of \sqsubseteq : what is already definitely true or false should remain so upon sharpening. The second option amounts to construing the valuation on a non-standard model on the basis of valuations that are not exactly standard,

and which do not therefore yield standard truth conditions (a sort of valuations “from the point of view” of the given model, as in Ermanno Bencivenga’s semantics for free logic¹⁶). This would satisfy monotonicity and would yield the same values as above for (16a) and (16b) while making (17a) definitely true and (17b) definitely false. So this can do admirably, though of course the standard quantification laws will have to be revised. In any case, the moral is that both options will eventually result in a non-classical logic due to the falsehood of (17a) and (17b), respectively. This is relevant, since one could otherwise be tempted to conclude that a semantics based on (12a) or (12b) is bound to yield a classical logic, at least as far as theoremhood goes.

3.2. Maximality and minimality Obviously things get even more complex when it comes to models with a significantly richer fabric than propositional assignments or first-order structures, but we need not go into that. Let us now assume a partial ordering \sqsubseteq to have been fixed on the relevant class of all models, along with a relevant strategy for evaluating sentences on the expansions and contractions of any given model. (We may generally think of \sqsubseteq as an approximation relation in Dana Scott’s sense, in which case it is most natural to assume it induces a complete lattice ordering.¹⁷) This means that (12a) is to be understood, broadly speaking, along the following lines:

(12a') A sentence A is true (false) in a model M if and only if A is relevantly true (false) in some consistent \sqsubseteq -contraction of every complete \sqsubseteq -expansion of M ,

and likewise for (12b). (Henceforth I shall focus on (12a), since nothing significant will hinge on the difference between the two policies.) The next important clarification concerns the scope of the sharpening process: How far can we go when we expand or contract a model M for the purpose of evaluating sentences on M ? Shall we consider every consistent contraction, or only those that are \sqsubseteq -maximal? Every complete expansion, or only those that are \sqsubseteq -minimal?

There is a sense in which the answer to these questions is constrained: if the only glut in the world of Sherlock Holmes is the discrepancy about Wat-

¹⁶ See e.g. Bencivenga 1981.

¹⁷ See Scott 1972. This conception is advocated by Belnap 1977. Of course there is no reason why one should not consider two distinct, possibly non-symmetric orderings to move along the two directions of the sharpening process, but I shall ignore that too.

son's wound and consequently his limp, there is no need to consider contractions that result from deleting other parts of the stories, and it would actually be misguided to consider contractions in which

(19) Watson was wounded

is not true. Likewise, it would be misguided to consider expansions that result from adding inconsistencies in those parts where the stories are perfectly consistent. There is, however, a sense in which a sharpening need not be the result of adding or deleting the bare minimum that will yield completeness and consistency. For instance, when we consider the contractions that emerge from the Holmes stories by deleting the fact that Watson was wounded in the leg, we may also want consider the result of deleting a number of other facts that would otherwise make little sense in spite of being perfectly consistent with the rest: in a contraction where Watson was wounded in the shoulder, Watson need not have a limp, need not own a collection of walking canes, etc. I am not sure that a dual example can be given for expansions, when these apply only to consistent models (or to consistent contractions of inconsistent models). Why should one fill in the gaps of an incomplete model to the extent of making it inconsistent? Besides, the whole idea is to contract and then expand so as to examine the various consistent and complete sharpenings of the given model M , *because* we are assuming that nothing is controversial about the goings-on of such sharpenings. If the sharpenings are themselves inconsistent or incomplete, then the entire semantic machinery is stalled. This is an argument to consider only the \sqsubseteq -minimally complete expansions even if contractions need not be \sqsubseteq -maximally consistent, at least if we base our evaluation policy on (12a). If, on the other hand, the sharpening process proceeded in the direction defined in (12b), by first expanding and then contracting, then of course the situation is reversed and one may find reasons to add more than just the bare minimum necessary to fill in the gaps. One would then consider the \sqsubseteq -maximally consistent contractions of expansions that need not be \sqsubseteq -minimally complete.

Be it as it may, all this suggests that the general evaluation rules for non-classical models should be made more precise by providing criteria for selecting the relevant class of sharpenings. We may speak of *adequately consistent* contractions and *adequately complete* expansions to indicate the result of such a selection. (Of course we assume every consistent model to count as an adequately consistent contraction of itself, and every complete model to count as an adequately complete expansion of itself.) Then (12a') can be amended as follows:

(12a") A sentence A is true (false) in a model M if and only if A is relevantly true (false) in some adequately consistent \sqsubseteq -contraction of every adequately complete \sqsubseteq -expansion of M .

3.3. Admissibility There are other, independent reasons for amending the evaluation rules in terms of selected contractions and expansions, as indicated in (12a"). For, generally speaking, there is no reason why the goings-on of an inconsistent or incomplete model should be evaluated by considering *every* possible sharpening. Generally speaking, one should only consider contractions and expansions that are *admissible* in some relevant sense. For instance, if the model is meant to represent the world of Sherlock Holmes, then not every way of filling in the gaps will do, for the fiction is meant to be read against a background of implicit facts (as Lewis pointed out¹⁸). The stories do not specify Watson's date of birth. But an expansion which says that

(20a) Watson was born in 1750

would be too far-fetched to be included in the admissible sharpenings of the world of Sherlock Holmes. Likewise, some complete expansions should be left out on account of certain "penumbral connections" (in the terminology of Kit Fine¹⁹) determined by the given (incomplete) model. For example, both Watson and Lestrade's dates of birth are left undefined in the model. However, suppose the model says that

(20b) Watson is older than Lestrade.

Then a complete expansion in which Watson was born in 1850 and Lestrade in 1849 should not be admissible. Analogous considerations apply of course to glut-deleting sharpenings. If both Mycroft and Lestrade are modelled as tall and also as short, and if Mycroft is unquestionably taller than Lestrade, then a consistent contraction in which Mycroft is short and Lestrade tall is simply not admissible.

This introduces a complication. For, in general, there is no guarantee that every model comes with a non-empty class of complete and consistent sharpenings of the relevant sort. It follows that the limit of a chain of contractions (expansions) is itself a contraction (expansion). We may even suppose that

¹⁸ This is the point of Lewis 1978.

¹⁹ See Fine 1975. A similar point is made in Kamp 1975.

such a limit is adequately consistent (complete). But of course it does not follow that the limit is admissible in the more general sense that we are now considering.²⁰ What, then, if there are no admissible sharpenings?

As it stands, (12a") delivers the following responses. (i) If M is a complete but inconsistent model, and if it has no admissible consistent contractions, then no sentence can be assigned a truth-value in M . (ii) If M is a consistent but incomplete model, and if it has no admissible complete expansions, then every sentence will be assigned both truth-values in M (vacuously). (iii) In general, if M is neither consistent nor complete, then every sentence will be neither true nor false if M has no admissible consistent contractions, and every sentence will be both true and false (vacuously) as long as M has a consistent contraction with no complete expansions. Now, this can't be right. If an inconsistency is so bad as to be ineradicable, one might perhaps find reasons to accept (i): one might say that it is hopeless to ask for an assignment of truth-values in such circumstances. But then, what parallel motivation could justify (ii)? One can hardly think of any reasons to evaluate a sentence as both true and false *because* of an unfillable gap. Even worse, one can think of no reasons to make a sentence true and false because of an unfillable gap in a consistent contractions (as (iii) has it), especially if other contractions work fine.

One way of resolving the issue would be to go precisely in the opposite direction, turning (i) and (ii) around. An incurable inconsistency then results in an abundance of truth-value gluts, and an incurable incompleteness results in an abundance of truth-value gaps. (iii) would then go along the same lines: a lack of admissible consistent contractions would yield an abundance of truth-value gluts, and a contraction with no admissible complete expansions would yield no truth-values and could therefore be ignored. I do not intend to pursue this suggestion here. After all, when there are no admissible sharpenings the whole rationale behind (12a") founders, and everything is up for grabs. It may even be that this is where the boundary should be drawn between reductionist semantics of the sort under discussion, where all truth-value assignments are classical at a remove, and a full-blooded paraconsistent semantics. (Or perhaps this is where the boundary should be drawn between manageable and unmanageable inconsistencies.) Be it as it may, what I want to stress here is simply

²⁰ One form of unsharpenability is examined in Collins & Varzi 1997. Fodor and Lepore's criticism of supervaluationism in 1996 may also be regarded as embodying a form of radical skepticism concerning sharpenability: what is unsettled in the actual world is unsettled in every world. See also Priest 1998 for some examples of unsharpenable impossible worlds.

that one must make room for some way of resolving the issue. From the present perspective the case of unsharpenable models is *sui generis*, and the basic evaluation rule should take this fact into account.

To this end, let us make explicit the presupposition that the admissible contractions and expansions (including M itself if it is either consistent or complete) be properly sharpenable. Let us say that a consistent contraction of a given model M counts as admissible only if it has some adequately complete expansion; and let us say that a complete expansion of M counts as admissible only if it has some adequately consistent contraction. Then the point is that (12a'') must be amended along the following lines:

(12a''') A sentence A is normally true (false) in a model M if and only if A is relevantly true (false) in some admissible, adequately consistent \sqsubseteq -contraction of every admissible, adequately complete \sqsubseteq -expansion of M .

Non-normal truth and falsehood can then be handled according to one's favorite views about the goings-on of an unsharpenable unsharp model. But this must be done independently of (12a''').

3.4. Relativity One last complication. Suppose that the model is unsharpenable—suppose there is no admissible, adequately consistent contraction of the world of Sherlock Holmes due to some deep discrepancies about Watson's old war wound. Why should this have any effect on the attribution of a truth-value to an innocent sentence such as

(21) Holmes smokes the pipe ?

Why should this sentence fail to be normally true simply because of those discrepancies about Watson? After all, everything said so far embodied the intuition (familiar from paraconsistent logics) that inconsistencies need not metastabilize throughout the entire language. Condition (12a''') (and its preliminary versions) capture this intuition in connection with sharpenable inconsistencies. Why should things be so radically different when it comes to indelible inconsistencies?

Another way of putting the same question is more cognitively oriented. As it is, (12a''') requires that in order to evaluate a sentence such as

(5a) Watson limps

we consider models that are complete and consistent relative to the entire language. But there is no obvious reason why such a task should require a radical operation like that—an operation that in some cases may even be impossible. A more natural account would be to consider sharpenings that are complete and consistent *relative* to (5a): models in which all the relevant facts are specified so as to make (5a) true or false, as the case may be. Similar considerations apply of course in case of a gap. Why should the task of evaluating a simple sentence like

(10a) Watson likes broccoli

require a global, simultaneous grasping of the indefinitely many sharpenings of the world of Sherlock Holmes? It should suffice to consider models that are complete and consistent *relative* to (10a).

The general point is that even when a model as a whole is unsharpenable, some (perhaps even all) of its proper fragments may be individually sharpenable, and for many purposes that is all that matters. One should be able to proceed as far as possible in the attribution of normal truth-values before giving up the approach, but (12a''') makes this impossible. To overcome this limitation, let us say that a model is *A*-consistent or *A*-complete (where *A* is any sentence) if and only if it is consistent or complete, respectively, relative to that fragment of the language that is explicitly involved in *A*.²¹ (In familiar cases, e.g., sentential or first-order languages, the relevant fragment is simply the set of symbols occurring in the construction tree of *A*.) Let us then suppose that the relevant notion of sharpening is redefined accordingly in terms of *A*-consistent contractions and *A*-complete expansions. In particular, an *A*-consistent contraction counts as admissible only if it has some adequately *A*-complete expansion, and an *A*-complete expansion counts as admissible only if it has some adequately *A*-consistent contraction. Then (12a''') can be amended along the following lines:

(12a''') A sentence *A* is normally true (false) in a model *M* if and only if *A* is relevantly true (false) in some admissible, adequately *A*-consistent \sqsubseteq -contraction of every admissible, adequately *A*-complete \sqsubseteq -expansion of *M*.

²¹ The notion of a model that is both *A*-consistent and *A*-complete is similar to Bencivenga's notion of an *A*-world. See e.g. Bencivenga 1978.

Of course, if M does not have any admissible A -sharpenings, then (12a''') will deliver the same unfortunate output as (12a''): A cannot be normally evaluated. But in general (12a''') will be much more efficient than (12a'').

3.5. Relevance There is still room for discussion. For instance, suppose the Holmes stories contain some such recalcitrant sentence A : there is no way to make the stories consistent (say) with respect to A . Then (12a''') will fail to handle, not only A , but also sentences such as

(22a) Holmes smokes the pipe, and also A .

(22b) Holmes smokes the pipe, or else A .

Now, this seems fair in the case of (22a): we must be able to handle A in order to handle the conjunction. However, (22b) seems different: here the first disjunct is normally true (and only true) in the world of Sherlock Holmes, and that seems sufficient for the purpose of evaluating the disjunction as true. The problematic status of A is arguably irrelevant.

There are various ways of dealing with such cases. However, I think the best option is indeed to accept this consequence of (12a''') and deny special status to (22b) in spite of the appearances. After all, the sense in which a disjunction with a true disjunct should come out true regardless of the value of the other disjunct is supervaluational: the disjunct is true because it *would* be true (truth-functionally) no matter what. But this motivation cannot serve its purpose here. (22b) cannot be normally true in this sense, for there simply is no admissible way in which it *could* be true (truth-functionally). One can still stipulate that a disjunction with a true disjunct is to be true (and only true), or that a conjunction with a false conjunct is to be false (and only false); but that would indeed be a stipulation from the present perspective. Therefore it should be handled as a case of non-normal truth, leaving (12a''') as it stands.

4. Sharpenability and logic

The upshot of this discussion is that the basic apparatus needed to make sense of inconsistent (and incomplete) models in the spirit of §2 is bound to be rather intricate.²² Now, this is not to say that the approach is unworkable. To the contrary, (12a''') is indicative of a certain flexibility. As long as the relevant

²² Kyburg 1997 has recently pointed out a number of related subtleties.

gluts and gaps are sharpenable, we can be assured that their admission will not bring logical disaster in its wake. Inconsistency need not yield contradictions, and tautologies do not imply completeness.

Formally, this general fact is reflected in the abnormal behavior of the relation of logical consequence, which will not in general satisfy the adjunctive and disjunctive laws. In particular, the following will fail:

$$(23a) \quad A, \neg A \vDash A \quad \neg A$$

$$(23b) \quad A \quad \neg A \vDash A, \neg A$$

There are in fact three distinct ways of defining the entailment relation, and each of them violates the above classical principles:

$$(24a) \quad \vDash_1 \quad \text{if and only if, in every model in which every element of } \Sigma \text{ is true, some element of } \Sigma \text{ is also true.}$$

$$(24b) \quad \vDash_2 \quad \text{if and only if, in every model in which every element of } \Sigma \text{ is false, some element of } \Sigma \text{ is also false.}$$

$$(24c) \quad \vDash_3 \quad \text{if and only if } \vDash_1 \text{ and } \vDash_2 \text{ .}$$

It is in any of these senses that the inferences from (5a)–(5b) to (6) and from (11) to (10a)–(10b) are blocked. On the other hand, the failure of (23a)–(23b) goes hand in hand with the fact that all tautologies remain true—and all contradictions false—in every model, or at least in every model where they can be adequately sharpened (and assuming that the strategy for evaluating sentences on a model’s sharpenings is classical enough to satisfy the usual truth-functional conditions). More generally, if we focus on such models, contradictions imply everything and tautologies are implied by anything:²³

$$(25a) \quad A \quad \neg A \vDash$$

$$(25b) \quad \vDash A \quad \neg A$$

Of course, if we end up with unsharpenable gaps and gluts, and if our treatment of such cases makes (some) contradictions non-normally true, or (some) tautologies non-normally false, then (25a)–(25b) do not hold either, and everything is up for grabs.

²³ See Varzi 1994, 1995 for proofs and related results.

Now, this gives a measure of the sort of departure from classical logic called for by (12a'''). One could say that the resulting logic is paraconsistent, but it is only half-heartedly paraconsistent (in the terminology of Priest and Routley²⁴) unless one allows for models that are truly unsharpenable. For in the absence of such models we lose classical entailment (more precisely, classical multi-premiss inferences) but we stick to the classics as far as tautologies and contradictions are concerned.

Two issues arise at this point. The first is how these constraints on sharpenability relate to our initial concern. Perhaps it is precisely these constraints that make the difference between ersatz worlds and genuine, Lewisian worlds—we can sharpen a way of representing a world but not the worlds themselves, not those “huge things” in one of which we have our being. If so, impossible worlds may well be nonsense, as Lewis has it, unless one goes paraconsistent all the way. On the other hand, for an ersatzer such constraints need not be drastic. There may be some unsharpenable goings-on, but one can still make sense of the rest as long as one does not close under classical logical implication. In this regard, (12a''') serves well its purpose.

The second issue concerns the relationships with classical logic. Exactly what classic patterns of validity are lost besides (23a)–(23b)? And how do such facts as (25a)–(25b) extend to non-truth-functional (e.g., quantificational) logical principles? As far as I know, both questions are very hard even in relatively simple contexts. For one thing, the first question seems to eschew a general answer already at the level of propositional logic, even assuming complete sharpenability and universal admissibility for all models.²⁵ This is so because in that case the notion of logical consequence becomes extremely language sensitive in the presence of gluts or gaps. It is true, for instance, that all instances of adjunction of the form (23a) and all instances of disjunction of the form (23b) fail. But if ‘*p*’ and ‘*q*’ are sentence symbols, the following instances are perfectly valid unless a restricted relation of admissibility is used:

$$(26a) \quad p, q \vDash p \quad q$$

$$(26b) \quad p \quad q \vDash p, q$$

Among other things, this means that the law of substitution will not hold any longer: propositional symbols cease to behave as variables, as it were.

²⁴ Priest & Routley 1989, p. 160.

²⁵ This is detailed in Varzi 1994, pp. 83ff.

On the other hand, the second question seems to call for different answers even with regard to the same language. It is here that the subtle complications involved in (12a''') become relevant. We noted in § 3.1 that as soon as we move to a language with quantifiers, the process whereby a sentence is evaluated on a model's sharpenings may not conform to the standards of classical logic. Thus, for instance, if we take the process to be performed "from the point of view" of the given model (so that quantified sentences are evaluated with reference to the domain of the given model), then we may go in the direction of a free logic. This means that not only classical entailments, but also classical validities are lost. In particular, (25a) and (25b) cannot be extended to their natural first-order analogues:

$$(27a) \quad (\exists x)A(x) \quad \neg A(a) \vDash$$

$$(27b) \quad \vDash (\exists x)A(x) \quad \neg A(a).$$

(27a) may fail because the premise can be true when the model is $A(a)$ -inconsistent (if we read entailment as \vDash_1 or as \vDash_3) or because the premise can fail to be false if the model is $A(a)$ -incomplete (if we read entailment as \vDash_2 or \vDash_3). Dually for (27b). By contrast, if we did not allow for contractions and expansions that involve a change of composition in the domain of quantification, then it is apparent that (27a) and (27b) retain their classical status exactly like (25a) and (25b).

These considerations are enough to suggest that the model theory behind (12a''') can be a messy business. This is fair, however, since we are talking about messy models after all. Whether genuine possible worlds are like that (if not worse), or whether the only genuine worlds are those that match the standards of perfectly sharp models, those are very difficult questions that I think we are entitled to postpone.²⁶

References

- Anderson, A. R., Belnap, N. D. Jr., and Dunn, J. M., 1992, *Entailment. The Logic of Relevance and Necessity. Volume II*, Princeton: Princeton University Press.
- Belnap, N. D. Jr., 1977, 'A Useful Four-Valued Logic', in J. M. Dunn and G. Epstein (eds.), *Modern Uses of Multiple-Valued Logics*, Dordrecht: Reidel, pp. 8-37.
- Bencivenga, E., 1978, 'Free Semantics for Indefinite Descriptions', *Journal of Philosophical Logic* 7, 389-405.

²⁶ Thanks to Roberto Casati and Graham Priest for helpful comments on an earlier draft.

- 1981, 'Free Semantics', *Boston Studies in the Philosophy of Science* 47, 31-48.
- Collins, J., and Varzi, A. C., 1997, 'Unsharpenable Vagueness', ms.
- Fine, K., 1975, 'Vagueness, Truth, and Logic', *Synthese* 30, 265-300.
- Fitting, M., 1992, 'Kleene's Logic, Generalized', *Journal of Logic and Computation* 18, 225-56.
- Fodor, J. A., and Lepore, E., 1996, 'What Cannot Be Evaluated Cannot Be Evaluated, and It Cannot Be Supervalued Either', *The Journal of Philosophy* 93, 516-35.
- Gärdenfors, P., 1988, *Knowledge in Flux*, Cambridge, MA, and London: MIT Press.
- Grant, J., 1975, 'Inconsistent and Incomplete Logics', *Mathematics Magazine* 48, 154-9.
- Gupta, A., and Belnap, N. D. Jr., 1993, *The Revision Theory of Truth*, Cambridge, MA, and London: MIT Press
- Hyde, D., 1997, 'From Heaps and Gaps to Heaps and Gluts', *Mind* 106, 641-60.
- Jaskowski, S., 1948, 'Rachunek zdań dla systemów dedukcyjnych sprzecznych', *Studia Societatis Scientiarum Torunensis, Sectio A* 1, 8, 55-77 (Eng. trans. as 'Propositional Calculus for Contradictory Deductive Systems', *Studia Logica* 24, 143-57).
- Jennings, R. E., and Schotch, P. K., 1984, 'The Preservation of Coherence', *Studia Logica* 43, 89-106.
- Kamp, H., 1975, 'Two Theories about Adjectives', in E. L. Keenan (ed.), *Formal Semantics and Natural Language*, Cambridge: Cambridge University Press, pp. 45-89.
- Kyburg, H. E., Jr., 1997, 'The Rule of Adjunction and Reasonable Inference', *The Journal of Philosophy* 94, 109-25.
- Levi, I., 1980, *The Enterprise of Knowledge*, Cambridge, MA, and London: MIT Press.
- 1991, *The Fixation of Belief and Its Undoing*, Cambridge: Cambridge University Press.
- Lewis, D., 1978, 'Truth in Fiction', *American Philosophical Quarterly* 15, 37-46.
- 1982, 'Logic for Equivocators', *Noûs* 16, 431-41.
- 1983a, 'Postscripts to "Anselm and Actuality"', in *Philosophical Papers, Volume I*, New York and Oxford: Oxford University Press, pp. 21-25.
- 1983b, 'Postscripts to "Truth in Fiction"', in *Philosophical Papers, Volume I*, New York and Oxford: Oxford University Press, pp. 276-80.
- 1986, *On The Plurality of Worlds*, Oxford, UK, and Cambridge, MA: Blackwell.
- McCall, S., 1966, 'Excluded Middle, Bivalence, and Fatalism', *Inquiry* 4, 384-86.
- Mehlberg, H., 1958, *The Reach of Science*, Toronto: Toronto University Press.
- Priest, G., 1998, 'Sylvan's Box', *Notre Dame Journal of Formal Logic*, this issue.
- Priest, G., and Routley, R., 1989, 'Systems of Paraconsistent Logic', in G. Priest, R. Routley, and J. Norman (eds.), *Paraconsistent Logic. Essays on the Inconsistent*, Munich: Philosophia Verlag, pp. 151-86.
- Rescher, N., and Brandom, R., 1980, *The Logic of Inconsistency. A Study in Non-Standard Possible-World Semantics and Ontology*, Oxford: Basic Blackwell.
- Schotch, P. K., and Jennings, R. E., 1989, 'On Detonating', in G. Priest, R. Routley, and J. Norman (eds.), *Paraconsistent Logic. Essays on the Inconsistent*, Munich: Philosophia Verlag, pp. 306-27.
- Scott, D., 1972, 'Continuous Lattices', in F. W. Lawvere (ed.), *Toposes, Algebraic Geometry and Logic*, Berlin/Heidelberg: Springer-Verlag, pp. 97-136.
- Van Fraassen, B. C., 1966a, 'Singular Terms, Truth-Value Gaps, and Free Logic', *The Journal of Philosophy* 63, 481-95.

- 1966b, 'The Completeness of Free Logic', *Zeitschrift für mathematische Logik und Grundlagen der Mathematik* 12, 219-34.
- Varzi, A. C., 1991, 'Truth, Falsehood and Beyond', in L. Albertazzi and R. Poli (eds.), *Topics in Philosophy and Artificial Intelligence*, Bozen: Mitteleuropäisches Kulturinstitut, pp. 39–50.
- 1994, *Universal Semantics*, PhD Thesis, University of Toronto.
- 1995, 'Super-Duper Supervaluationism', in T. Childers and O. Majer (eds.), *Logica '94. Proceedings of the 8th International Symposium*, Prague: Filosofia, pp. 17-40.
- Visser, A., 1984, 'Four-Valued Semantics and the Liar', *Journal of Philosophical Logic* 12, 181-212.