

Indeterminate Identities and Semantic Indeterminacy

Achille C. Varzi

Department of Philosophy, Columbia University

1. Event Identity and Indeterminacy

Consider the following familiar scenario. There was a duel at dawn between X and Y. X was faster and shot first; Y died. The background conditions are such that it is right to say, now, not only that X shot Y but also that X killed Y. So:

(1) *Data:* X shot Y at dawn. In fact, he killed Y.

Question: Is X's shooting of Y the same as X's killing of Y?

This familiar scenario illustrates a puzzle about which philosophers have conflicting views. Some are inclined to answer the question in the negative, by appeal to the fact that the events in question seem to have distinct properties. For example, suppose Y lingered on until dusk before dying of his bullet wound. Then at noon on the fatal day Y was still alive, so he had not been killed yet. But he had been shot. Hence the shooting can't be the same as killing. Other philosophers are inclined to say that the events in question are one and the same—that X performed one and only one action at dawn—though our linguistic resources are such that we can pick out that event by means of linguistic descriptions or action sentences, not all of which have the same semantic properties. For example, if indeed Y died at dusk, then at noon we can only describe what happened by saying that there was a shooting, not a killing. (Or: at noon we can truthfully say that X shot Y, not that X killed Y.) But this is not to say that there were two events. It simply means that what happened can be described in different, non-equivalent ways at different times, just as one and the same person (say, George W. Bush) can be described in different and even incompatible ways at different times (say, as the son of the US President, or as the US president himself).

How do we choose between these two views? Of course, if we insist that 'the killing' refers to an event that extends over a longer period of time than 'the

shooting', up to the death of the victim, then the question has an obvious answer: The killing and the shooting must be distinct. But suppose we insist that 'the killing' refers to an event that took place at the same time as the shooting. Or suppose that Y's death was instantaneous, so that the temporal factor is irrelevant. Shall we still insist on these events being distinct on account of their having different *modal* properties? (At noon on the fatal day Y *could* have been alive: so it *could* have been the case that, at noon, Y had been shot without having been killed yet.) Shall we insist instead on their being the same event in spite of there being significant modal differences in the statements and descriptions by which we can talk about what happened?

In his early work on events, Terence Parsons argued that identity questions such as these can, in some cases, be settled by the theory of events that one endorses. More precisely, in some cases they can be settled by the event-based theory that displays the logical forms of the statements in question, and hence of the event-referring descriptions that can be extracted from those forms, as long as that theory is sophisticated enough to take into account all the relevant linguistic data. (Standard Davidsonian analyses are too unsophisticated for this purpose, and therefore leave room for battles of intuitions. But Parsons's theory is more fine-grained.) In other cases, however, the data may not be enlightening enough and the theory may not deliver any definite answer. The theory may be compatible with the view according to which the events in question are identical as well as with the view according to which the events are distinct, and it may even be compatible with the view that the relevant identity statement lacks a determinate truth-value.

For example, consider the following two scenarios, both of which are also familiar from the literature on event identity.

- (2) *Data:* A person, P, is singing. In fact, she is singing loudly.
Question: Is P's singing the same as her singing loudly?
- (3) *Data:* A metal sphere, S, is spinning. In fact, it is also heating up.
Question: Is S's spinning the same as its heating up?

These are controversial cases. Nonetheless, Parsons's theory delivers a definite answer to both questions. In the first case, the theory immediately delivers a positive answer (*contra* Kim). For the theory tells us this:

- (2') P's singing = (the e)[Singing(e) & Agent(e, P)]
P's singing loudly = (the e)[Singing(e) & Agent(e, P) & Loud(e)]

And if both descriptions refer, then they must refer to the same event by ordinary principles of predicate logic. In the second case, the theory delivers a negative answer (*contra* Quine). At least, it delivers a negative answer if we make certain plausible assumptions—for instance, if we assume that all spinnings are rapid whereas our sphere heated up slowly. For in that case the theory tells us this:

- (3') S's spinning = (the *e*)[Spinning(*e*) & Agent(*e*, S)]
- S's fast spinning = (the *e*)[Spinning(*e*) & Agent(*e*, S) & Rapid(*e*)]
- S's heating up = (the *e*)[Heating-up(*e*) & Agent(*e*, S)]
- S's slow heating up = (the *e*)[Heating-up(*e*) & Agent(*e*, S) & Slow(*e*)]

The first two descriptions must be co-referential, if they refer at all, and so must the third and fourth descriptions. But if all descriptions refer, the common referent of the former must differ from the common referent of the latter by ordinary principles of predicate logic.

So far nothing new. Davidson's theory delivers the same answer. But Parsons's theory is more sophisticated than Davidson's and, as a result, delivers more answers. Consider, for instance, the following:

- (4) *Data*: Mary played the violin. In fact, she played a Bach sonata.
Question: Is Mary's violin playing the same as her sonata playing?

Here a standard Davidsonian account leaves the question unsettled (though one might want to distinguish the two events by trading on a difference in meaning between the two uses of 'play'). On Parsons's theory, however, the relevant logical forms are more detailed:

- (4') Mary's violin playing = (the *e*)[Playing(*e*) & Agent(*e*, Mary) & Theme(*e*, Mary's violin)]
- Mary's sonata playing = (the *e*)[Playing(*e*) & Agent(*e*, Mary) & Theme(*e*, Bach's sonata)]

And these descriptions cannot be coreferential, given that Mary's violin and Bach's sonata are distinct. So the events must be distinct (even if the two uses of 'play' are regarded as synonymous.)

So the theory goes further than Davidson's with regard to identity issues. It does not, however, go all the way, and there is no reason to think that it could. One obvious case where the theory does not and cannot deliver an answer involves vagueness: there may be indeterminacy as to whether John's balding and

his losing hair are one and the same process, due to the soritical conundrums that go with our ordinary use of ‘balding’. But vagueness is not the only sort of case where a theory such as Parsons’s leaves an identity question unresolved. Here is another case:

- (5) *Data*: X killed Y; he murdered him.
Question: Is X’s killing of Y the same as his murdering of Y?

In this case the theory yields the following logical forms:

- (5') the killing = (the e)[Killing(e) & Agent(e , X) & Theme(e , Y)]
the murdering = (the e)[Murdering(e) & Agent(e , X) & Theme(e , Y)]

And even assuming that the relevant predicates have precise boundaries of application, there is little we can say about whether these two descriptions are co-referential. If we make the additional assumption that every murdering is a killing, then of course the second description must be co-referential with the first. But without appeal to such an assumption the question is up for grabs. The theory does not entail a definite answer. Likewise, to finally go back to the initial scenario, Parsons’s theory does not automatically entail any answer to the question in (1). All the theory says is this:

- (1') the shooting = (the e)[Shooting(e) & Agent(e , X) & Theme(e , Y)]
the killing = (the e)[Killing(e) & Agent(e , X) & Theme(e , Y)]

And because there is no reason to suppose that every killing is a shooting, or that every shooting is a killing, there is no way to infer an answer to (1) directly from (1'). (In fact, Parsons has opinions about what the answer should be. He is inclined to say that the two descriptions in (1') refer to the same event. But such opinions are independent of the theory and are correspondingly controversial.)

It might be thought that the indeterminacy here is only apparent. After all, the predicates that we use to pick out our events are only partially informative. As cases (2) and (3) illustrate, we can make our event descriptions more precise by adding relevant information—information that at the level of logical form gets cashed out in terms of additional conjuncts; and once all the details are filled in we might hope that the answers follow as a matter of ordinary predicate logic. Still, this may be right in some cases, but not always. For the expressive limits of our language might be such that the events in question satisfy all the same predicates while having different properties. And even assuming that we have means for describing every property, we might simply not be in a position to achieve

the necessary descriptive completeness unless we already have an answer to our problem. Such is the limit of Leibniz's laws. For instance, to the top description in (1') we may add as a conjunct the statement that *e* is a killing, or the statement that *e* is not a killing. But clearly enough, the choice between these options involves the same difficulty that is involved in the question we are trying to answer. The shooting is *a* killing if and only if it is *the* killing.

2. Taking Indeterminacy Seriously

Now, what are we to say in those cases where the theory delivers no definite answer to a question of identity? There seem to be two options. One is to insist that *there is* an answer, whether or not we can figure it out. The other option is to say that there is no answer at all—that we are confronted with identity statements whose truth-values are truly indeterminate, and cannot be determined in any way.

In his recent work Parsons has gone a long way towards a clarification of the second option—a clarification of what is involved in the claim that a certain identity statement is indeterminate. There are, in fact, many other situations where we may find ourselves in the same sort of predicament, so one could hardly take the indeterminacy exhibited by the examples above to provide an indirect *reductio* of those theories that posit events in the first place. Here are some examples, of the sort that Parsons discusses in his works:

- (6) *Data:* The dog is in the garage. The cat is in the living room.
Question: Is the dog in the same house as the cat?
- (7) *Data:* There is exactly one person in the room, and exactly one human body in the room.
Question: Is the person in the room the same as the body in the room?
- (8) *Data:* Exactly one ship, A, left port, but as a result of a familiar repair/assembly process, two ships, B and C, docked (one with new parts, one consisting of the old parts reassembled).
Question: Is A identical with B, with C, or with neither?

In cases such as these, Parsons urges us to regard the identity statements in question to be indeterminate. Some might think that they have a definite but unknown truth-value, but this is not what Parsons means by 'indeterminate'. For Parsons the identities in question are indeterminate precisely insofar as they lack

a truth-value altogether. Of course, one might hope that a theory about the entities in question might justify a different attitude, just as a theory about events might resolve some of the *prima facie* indeterminacies that affect our event talk. A rigorous theory of houses will presumably be such as to specify whether a house with a garage (a garage that is separate from the rest of the building) includes the garage as a proper part, hence whether the question in (6) should be answered in the affirmative or in the negative. A materialist about persons will hold a view that, when properly articulated, will imply an affirmative answer to the question in (7), whereas a dualist will hold a view that implies a negative answer. And so on. But these would not be theories in the sense in which Parsons's is a theory of events. They are either stipulative theories (about houses, for instance) or genuinely metaphysical theories (about persons), whereas Parsons's theory of events is a semantic theory—a theory about language that has explicit ontological implications but few metaphysical commitments. (It does assume that events are spatio-temporal particulars, but nothing besides that.) If we can get help from a semantic theory of this sort, then it's good news. But if we have to resort to *ad hoc* stipulations or pull out a full-fledged metaphysical account, then chances are that the answers afforded by our theory will only be accepted by those who already share the same views, and no significant progress is made. Since no semantic theory of the desired sort seems available to do the job, the thought that we are confronted with indeterminate identity statements must be taken seriously.

There is an important methodological point to be stressed here, one on which Parsons has been much clearer than most philosophers and linguists who have been puzzled by such issues. Parsons puts it in terms of Peircean guidelines. One begins with ordinary beliefs, and one rejects them only if some reason is found to challenge them. *Ad hoc* stipulations or metaphysical theories that contradict such ordinary beliefs and judgements must be averted, and one should rather look for theories that “preserve the data”. Parsons's theory of events (or Davidson's, for that matter) is a theory that preserves the data. But stipulative theories about houses or metaphysical theories about persons are generally revisionary. In fact, to the extent that such theories force upon us an answer to the questions illustrated in (6) and (7), they are bound to be revisionary, for the apparent data is that in cases such as these we have no grounds at all for answering the questions in a determinate way. Such theories explain how and why we should *change* our beliefs, and how the puzzles are resolved if we do so. I don't know whether Parsons believes that revisionary theories of this sort—which are very popular and constitute a major chapter of contemporary analytic philoso-

phy—are bad philosophy. I myself am inclined to think that sooner or later that is what we have to do if we want to do philosophy at all. But never mind that. If we want to get clear about the prospects of indeterminate identity, and if we want to do so while preserving the data, we can hardly dismiss Parsons's methodology at the outset.

3. Ontic Versus Semantic Indeterminacy

So, of the two choices mentioned above with respect to the issue of event identity, Parsons favors the second. But as recent literature has emphasized, there is still room for disagreement here—a disagreement that arises even within the methodological constraint just mentioned. One can maintain that the indeterminacy in question is purely semantic (pertaining exclusively to the link between the words we use when we talk about the world and the world itself) or one can maintain that it is a matter of ontology (that there is no worldly fact of the matter as to whether the referents of our singular terms are identical or distinct). Most philosophers favor the first option, and I confess that I belong to this traditional crowd. Parsons favors the latter. And his book *Indeterminate Identity* is the best attempt ever made to make sense of this position.

Why go for an ontological account of indeterminate identity? Parsons gives us two sorts of reasons. On the one hand, he argues that the account is perfectly coherent, let alone plausible, in the face of various purported refutations that have been leveled in the literature (such as Evans's celebrated argument and the refinements that followed). On the other hand, Parsons argues that the account is actually better than its competitor *vis-à-vis* a number of identity puzzles of the sort mentioned above, at least from a broadly Peircean methodological perspective. It provides a better diagnosis of the puzzles and it yields a systematic account, whereas the view that the indeterminacy of identity lies entirely in our language, or in the system of concepts embodied in our language, fails to do so. It fails—Parsons says—because the view has never been fully articulated. Of course some cases of indeterminate identity are easily explained in terms of semantic indeterminacy, that much is uncontroversial. But the puzzles that worry Parsons the most cannot, on his reckoning, be properly handled by reconstruing them exclusively as puzzles about our language, not in such a way that the solution to one puzzle carries over to the other puzzles. At least, this is what Parsons concludes after taking a close look at the semantic theories of indeterminate identity put forward so far: such theories are not compelling in their present form, Parsons says, though there may be room for improvements.

I want to focus, here, on this second part of Parsons's work. I will grant the point of the first part entirely, not only for the sake of the argument but because I think that Parsons has indeed succeeded in making the point, providing us with a plausible account of what it means to say that the world might be genuinely indeterminate (more so in the book than in his earlier attacks on the problem). So I want to focus on some of the challenges issued by Parsons against the semantic conception of indeterminacy. They are serious challenges, indeed—challenges that as far as I know have never been seriously considered by the defenders of the semantic conception of indeterminacy. But I want to argue that they can be met. In fact, I want to argue that Parsons gives us all the necessary ingredients to come up with a formulation of the semantic conception that does the job, and that does the job in a methodologically acceptable way (by Parsons's own light).

In doing so, I will follow Parsons in taking the semantic conception to be, at bottom, supervaluational. On this conception, the apparent indeterminacy of some identity claims stems from the inexactness of our language, which may prevent us from making determinate judgements in spite of the world's being perfectly determinate. And to say that our language is inexact is to say that in some cases the referential pattern of our words is not fixed with sufficient precision. For example, in some cases the names and descriptions that we use to talk about the entities in the world lack a precise semantics, in that they have more than one potential referent. If so, then an identity statement involving such terms is essentially ambiguous, in that it has no unique reading that could yield a definite truth-value. The statement would be true according to some ways of refining (or "precisifying") the semantics of our names and descriptions, but it would be false according to some other ways of refining their semantics. And since there is no reason to prefer some refinements to others—since our linguistic practices are compatible with both sorts of refinements—there is no way to evaluate our statement as true, or as false. In other words, our statement admits of different interpretations; and although in some cases such an embarrassment of riches is ultimately irrelevant (a statement can be true under every admissible interpretation), when it comes to identity statements the situation is hopeless and the semantic variation in the potential reference of our terms results in a truth-value gap. Truth is super-truth, i.e., truth under every admissible semantic refinement—says the supervaluationist. And super-truth is not bivalent.

In other cases, the semantic imprecision affects other parts of our language. For example, it may affect our general terms (or predicates) rather than the singular terms by which we aim to pick out entities in the world. Or it may affect the concepts that are expressed in our general terms. In such cases the semantic

diagnosis is slightly different but the account is similar, and the multiplicity of the admissible refinements of our predicates and/or concepts leads naturally to a form of supervaluationism. If a statement turns out to have the same truth-value under every refinement, then we may conclude that our imprecision makes no difference to the truth-value of the statement. But if different refinements result in different truth-value assignments, then our imprecision does make a difference and the only safe option is to leave our statement truth-valueless.

The supervaluational account is not, of course, the only possible option for a friend of semantic indeterminacy. Other theories have been proposed, some of which emphasize other features of our linguistic and cognitive practices that supervaluationism must abandon (such as the truth-functionality of the logical connectives, or the disquotational property of the truth predicate). But I think Parsons is right in taking supervaluationism to be the better account, all things considered. I think supervaluationism is actually the *only* reasonable account one can give of the semantic indeterminacy thesis. So here I will just focus on Parsons's challenges to this theory, ignoring other options until the last section.

4. The Supervaluationist Account (and Its Limits)

Let us briefly review how supervaluationism works in some concrete cases. Consider again (6). We have a house with a separate garage. There is a dog in the garage, and there is a cat in the living room. Is the dog in the same house as the cat? Here we may introduce two names, 'A' and 'B', to designate the house where the dog is and the house where the cat is, and we may rephrase our question as a question about the truth-value of the identity statement 'A = B'. Since our two names have not been introduced with sufficient precision, our question cannot receive a definite answer. There are two candidate referents for 'A': the garage and the complex that includes the garage as a proper part. There are two candidate referents for 'B': the main building and the complex that includes the garage as a proper part. So there are four possible ways of interpreting the statement 'A = B'. On three of them the statement is false; on one of them (when both 'A' and 'B' denote the complex) the statement is true. So, on the face of it our lack of precision prevents us from settling the issue, and if truth is super-truth that means that the identity statement 'A = B' is indeterminate. (By contrast, many other statements involving 'A' or 'B' turn out to have a definite truth-value in spite of the imprecision. For example, if we say that the dog is in A, our statement is super-true, for it is true under both ways of choosing from among the candidate referents of 'A'. This is quite natural and the outcome "pre-

serves the data”, too. For we do want to say that the dog is in the house where the dog is, no matter how exactly we interpret that. In this sense, supervaluationism yields a parsimonious amount of truth-value gaps. It yields truth-value gaps only in those cases where, intuitively, we feel that there is no sufficient semantic information.)

This sort of account applies naturally to many other cases. For instance, it applies in cases where the relevant indeterminacy is a sign of vagueness. By vagueness I mean the sort of indeterminacy that yields the so-called sorites paradox, versions of which can be given not only to bring out the vagueness of some predicates but also the vagueness of some singular terms of our language—including proper names. For a familiar example, take Mount Everest. Let $A_1 \dots A_n$ be a long sequence of increasingly larger regions stretching from a small area around the highest point of Himalaya, at $27^\circ 59' \text{ N } 86^\circ 56' \text{ E}$, to a large area encompassing the whole of Himalaya.

- (9) *Data:* Mount Everest is part of Himalaya; it is different from A_1 and from A_n , but its peak is at $27^\circ 59' \text{ N } 86^\circ 56' \text{ E}$.

Question: Is Everest identical with A_i (for some intermediate i)?

Here one could bite the bullet and deny the data. But if we want to preserve the data, then it appears that for some i the question cannot be given a definite answer. (Whether there is a first such i is itself a difficult question, but let us put that aside.) And the supervaluational explanation of this fact is straightforward. The name ‘Everest’ has a vague semantics. There are several different ways of tracing the geographic limits of Mount Everest, all perfectly compatible with the way the name is used in ordinary circumstances. Indeed, when the members of the Geodetic Office of India baptized the mountain after the name of their British founder, they simply did not specify exactly which parcel of land they were referring to. The referent of their term was vaguely fixed. Each one of a large variety of slightly distinct regions from among the intermediate elements of the sequence $A_1 \dots A_n$ has an equal claim to being the referent of that name (each such region being perfectly determinate). So for any such region, A_i , the claim that Everest is identical with A_i is true on some (in fact, exactly one) admissible refinement of the semantics of ‘Everest’, but false on the others. So the claim is neither super-true nor super-false. It is indeterminate. (By contrast, other claims involving ‘Everest’ may receive a determinate truth-value in spite of the vagueness of this term. This is true, in particular, of the statements that fix the data. No matter how we choose a referent for ‘Everest’ from among the many admissible options, that referent will be part of Himalaya, so the claim that Everest is

part of Himalaya will be super-true. Ditto for the claims that Everest is different from A_1 and from A_n , since neither of these regions is among the admissible referents.)

Now, the important question that we have to ask is whether this sort of account can equally be applied to the other identity puzzles mentioned above. If it can, then the only reasons to reject supervenientism as a plausible account of the relevant indeterminacy would have to be reasons of a general sort, reasons that do not have anything specific to do with the issue of indeterminate identity. And there are such reasons, according to some critics. But Parsons is not worried about that. His challenges are more direct and concern specifically the adequacy of supervenientism *vis-à-vis* the critical identity puzzles. Fine with such cases as (6) or (9). But what about the rest?

Consider the person/body puzzle (7). There is exactly one person in the room, and exactly one human body in the room, but it is indeterminate whether the person in the room is identical with the body in the room. What is the supervenient explanation of this indeterminacy? Following the account sketched above, the explanation goes like this. ‘The person in the room’ and ‘the body in the room’ (henceforth: ‘the person’ and ‘the body’ for short) have multiple potential referents. One potential referent of ‘the person’, a , is also a potential referent of ‘the body’. But there is also a potential referent of ‘the person’, b , that is not a potential referent of ‘the body’, and there is a potential referent of ‘the body’, g , that is not a potential referent of ‘the person’. So the statement ‘the person = the body’ turns out true on one semantic refinement (the one where both terms are interpreted as denoting a) but false on the others (those in which ‘the person’ is interpreted as denoting b while ‘the body’ denotes either a or g and those in which ‘the body’ is interpreted as denoting g while ‘the person’ denotes either a or b). So the statement is neither super-true nor super-false, hence it is indeterminate, as desired.

What is wrong with this account? On the face of it, there are two problems. Parsons put his finger on one, which is bad enough, but there is a second one that I would like to mention. The problem stressed by Parsons is one of metaphysics. For Parsons the account is formally adequate, but it clashes with the intuition that there are only two things in the room. We are supposed to consider three potential referents overall, a , b , and g . But—Parsons asks—where are these entities? When we scrutinize the data and consider the situation, there are at most two entities that are relevant to the puzzle: a person and a body. Perhaps these are not *two* entities but only one, that’s the puzzle. But there surely are not three candidates to choose from. So the account may be formally ade-

quate, but it requires us to make explicit ontological assumptions that go beyond the data and that run against intuition. In this sense, the situation under examination is different from the house case, and also from the Everest case. There we have a multiplicity of potential referents, some of which are part of others, and the relevant indeterminacy is explained very naturally in terms of semantic imprecision. But here we do not have a comparable multiplicity of referents. The data tell us that there are at most two entities, and we want to preserve the data. In short, the analysis afforded by supervenience in the early cases does not generalize to the case at issue. (It might be objected that this is too quick. After all, if we have at least two entities, then we have at least three, since the mereological fusion of two disjoint things counts as a third thing. So here it could be objected that one of the entities in question, namely *a*, is the fusion of the others, *b* and *g*. And, as some like to say, a fusion is nothing over and above the parts that constitute it. But this wouldn't do. If you are a dualist about the mind/body problem, then fine: *b* is a mind, *g* is a body, and since it is unclear whether 'person' applies only to minds or to mind+body combinations, there is indeterminacy as to whether the referent of 'the person' here is *b* or *a*—whence the indeterminacy of the relevant identity statement. But our methodology prevents us from assuming a dualist position here. And if we do not make that assumption, then there is no way out of the charge that the supervenient account brings in too many creatures to explain the relevant indeterminacy. That is, there is no way out unless we allow for ontological indeterminacy, making it an indeterminate matter whether *b* and *g* are actually distinct potential referents.)

To make things worse, here is the second problem. We want the account to explain why the question in (7) lacks a determinate answer. But we also want the account to preserve the data. And unlike the previous cases (the house and the Everest cases), here the account does not preserve the data. Hence it is not formally adequate after all. Here is why the account does not preserve the data. When we say that 'the person' does not have a fixed referent, we implicitly say that 'person' does not have a fixed extension either—and similarly for 'body'. Otherwise *a*, which is supposed to be a potential referent for both terms, would have to be in the extension of 'person' as well as in the extension of 'body'; and given the relevant logical forms,

- (7') the person = (the *x*)[Person(*x*) & In(*x*, the room)]
 the body = (the *x*)[Body(*x*) & In(*x*, the room)],

this would contradict the proposed analysis as a matter of ordinary predicate logic. (If the descriptions in (7') denote at all, they must both denote *a*.) It is rea-

sonable to assume, instead, that the extension of ‘person’ varies with the referent of ‘the person’—and similarly for ‘body’. Here we have two options, depending on the theory of descriptions that we are assuming in the background. One option is to require that the predicate ‘person’ be always refined so as to include in its extension exactly one entity in the room, be it *a* or *b*, for we want the definite description ‘the person’ to have a unique referent under every refinement. Ditto for the predicate ‘body’. This would preserve some of the data. It would guarantee that the statements ‘The person is a person’ and ‘The body is a body’ are super-true. And it would imply, plausibly, that the statements ‘The person is a body’ and ‘The body is a person’ are truth-valueless. But it would also imply that the statement ‘Everything in the room is either a person or a body’ (with the quantifier restricted in the obvious way) is super-false. And this contradicts the implicit data, or so it seems to me. The other option is to allow the predicates ‘person’ and ‘body’ to include more than one object in their extension, which is to say either one or two objects. Since we want ‘The person is a person’ and ‘The body is a body’ to be super-true, this must be done so that the extension of ‘person’ is a non-empty subset of $\{a, b\}$ and the extension of ‘body’ is a non-empty subset of $\{a, g\}$, and as a consequence the statement ‘Everything in the room is either a person or a body’ would come out truth-valueless. This is better. But then the statements ‘There is exactly one person in the room’ and ‘There is exactly one body in the room’ would come out truth-valueless, too. This is so because some refinements of our predicates will yield an extension containing one object in the room, whereas others will yield an extension containing two objects. And to say that these statements are truth-valueless is to contradict the data explicitly.

So these are the problems. And it is easy to see that these problems arise also in other identity puzzles, including cases where no philosophical position comparable to a mind/body dualism could help. Take for instance the event identity puzzles with which we began. (I was surprised not to see any reference to such puzzles in Parsons’s recent work on indeterminate identity.) Consider the shooting/killing case in (1). If we apply the supervaluational procedure illustrated above, then the explanation is perfectly parallel to the person/body puzzle. We have two event descriptions, ‘the shooting’ and ‘the killing’ for short, and these have multiple potential referents. One potential referent of ‘the shooting’, *a*, is also a potential referent of ‘the killing’. But there is also a potential referent of ‘the shooting’ that is not a potential referent of ‘the killing’, and vice versa. So the claim that the shooting is identical with the killing comes out true on one semantic refinement but false on the others, and therefore counts as

supervaluationally indeterminate. This is good. But this account, like the person/body account, forces us to assume three entities—in fact, three events—to be the potential referents of our two event descriptions. And this is bad. For the issue was whether the events in question are one or two, and that should not entail that these events must be chosen out of *three* events. To the extent that the account entails that, it is open exactly to the same objections raised above.

5. More on Supervaluationism

Parsons is right in pointing his finger on these problems. For much recent literature on supervaluationism yields an approximate picture at best, and the approximate picture looks exactly like the picture we have just reviewed and is therefore open to the objections just raised. There is, however, an oversimplification in this picture. And since I believe that the oversimplification is playing a role in yielding the problems at issue, it is important to take a closer look.

Broadly speaking, supervaluationism tells us two things. The first is that the semantics of our language is not fully determinate, and that statements in this language are open to a variety of interpretations each of which is compatible with our ordinary linguistic practices. The second thing is that when the multiplicity of interpretations turns out to be irrelevant, we should ignore it. If what we say is true under all the admissible interpretations of our words, then there is no need to bother being more precise. In ordinary circumstances this is indeed the case, and that's why our language is still very imprecise. In some cases—particularly in the case of philosophically significant statements, such as identity claims—the multiplicity of interpretations does turn out to be relevant, and that's why in such cases what we say will suffer from a truth-value gap. Now, this second aspect is what truly captures the spirit of supervaluationism (and even the word). If you don't like this aspect, then supervaluationism is not even in the right ballpark for an explanation. But as I said, this is not a worry for Parsons (at least not the main worry), and it certainly is not a worry for me. I think the idea that truth is super-truth is both cognitively and linguistically compelling, in spite of the many criticisms that have been leveled against it, and so I am not going to worry about that here. The worry concerns the first aspect—the idea that the indeterminacy of our language amounts to a multiplicity of potential interpretations. And it is with respect to this part of the story that we have to be more specific if we want to avoid the objections raised above.

Here is one important point where we must be specific. When we say that our language admits of a variety of interpretations, we can mean either of two

things. We can mean that the semantics is indeterminate while the ontology is fixed—that is, that we have settled on our domain of quantification but not on a unique interpretation function, so that our singular and general terms may have multiple potential referents and multiple potential extensions in that domain. Or we can mean that the semantics as well as the ontology is up for grabs, in that the domain of quantification is also not completely fixed. (It is partially fixed by our commitments, but we may not have settled on the exact composition of the whole domain.) Formally, this means that in the first case we construct our supervaluation as the logical product of the valuations induced by the models that we get by taking the given domain together with one of the admissible interpretation functions defined on that domain, whereas in the second case we construct the supervaluation as the logical product of the valuations induced by the models that we get by taking a domain that is *compatible* with our ontological commitments, together with a corresponding interpretation function.

Now, it seems to me that if we want to stick to the broadly Peircean methodology recommended by Parsons, we have to go with the second option. There is no obvious reason to suppose that the indeterminacy of our semantic practices should entail complete clarity or agreement on what there is, so no reason to suppose that the domain of quantification should be fixed once and for all. There is one and only one world, but we may have conflicting models of it. However, the supervaluational procedure with which we have addressed the identity puzzles in the previous section implicitly assumed the first option instead. We implicitly assumed that the candidate referents of ‘the person’ and ‘the body’, for instance, must all belong to the same domain of quantification, just as the candidate referents of ‘Everest’ or ‘the dog’s house’ must all belong to the same domain. This is not a justified assumption from the present methodological perspective. In the case of ‘Everest’ it seems plausible to suppose that the potential referents for the name are all there, so to say; we just have not been accurate enough to select a unique one of them. But in the person/body case this assumption does not seem plausible at all. On the contrary, precisely insofar as we have not settled on whether a person coincides with her body, we have not settled on how many objects must be included in our domain of quantification. It seems plausible to include two for each person, namely a person and a body, and it seems equally plausible to include just one object for each person, identifying every person with her body. In other words, there is indeterminacy as to what the right model of the world looks like—not only in that the referents of our singular terms and the extensions of our predicates are not uniquely defined, but also insofar as the relevant portion of the domain of quantification has not been

uniquely determined. To bypass the issue by assuming the domain to be fixed is to surreptitiously build into the picture a metaphysical view that our methodology does not warrant. Ditto for the shooting/killing case and the like.

Now, if this is right—if the correct way to construe the indeterminacy of our semantics is to allow for some indeterminacy concerning the underlying ontology—then the supervaluational account looks very different from what we made it look like. In the Everest case, (9), as well as in the house case, (6), the ontology is locally clear, so we can proceed as indicated: the multiplicity of models associated with our language agree on the relevant denizens of reality. But in general we must proceed by associating our language with a multiplicity of models that may disagree on that, too. Thus, in the person/body case, (7), we should say this. (The shooting/killing case, (1), is perfectly similar.) We agree that there is exactly one person in the room, and we agree that there is exactly one human body in the room. So there are two main types of possible models for our language:

(7'') *Type 1*: The domain of quantification includes one object located in the room, **a**, which is in the extension of ‘person’ as well as in the extension of ‘body, and which counts as the referent of both ‘the person’ and ‘the body’.

Type 2: The domain of quantification includes two objects located in the room, **b** and **g**. One of these objects is in the extension of ‘person’ and counts as the referent of ‘the person’; the other is in the extension of ‘body’ and counts as the referent of ‘the body’.

This seems to me the correct way to describe the picture according to the first guideline of supervaluational semantics. The second guideline now tells us to calculate the truth-values of our statements by calculating the logical product of the valuations induced by the admissible models of our language. So let’s see. On models of the first type, the two terms have the same referent, hence the identity statement ‘the person = the body’ is true; on models of the second type, the terms have different referents and the statement is false. So our supervaluation tells us that the statement is indeterminate; it is neither true nor false. Likewise, it is easy to see that the supervaluation yields the following results:

(7''')	‘The person is a person’	<i>true</i>
	‘The body is a body’	<i>true</i>
	‘The person is a body’	<i>indeterminate</i>
	‘The body is a person’	<i>indeterminate</i>

This is precisely what we expected and coincides with the answer delivered by the early account as well. But now, unlike before, we don't run into the unexpected. For how many things are there in the room (with 'thing' understood restrictedly in the obvious way)? Certainly not three, except in the trivial sense that whenever we have got two things we also have their mereological fusion. So no metaphysical problems here. In fact, we have one thing in the first type of models and two things in the second type, so the supervaluation will give us the following results :

(7''') 'There is exactly one thing in the room'	<i>indeterminate</i>
'There are exactly two things in the room'	<i>indeterminate</i>
'There is at least one thing in the room'	<i>true</i>
'There are at most two things in the room'	<i>true</i>
'There are either one or two things in the room'	<i>true</i>
'There is exactly one person in the room'	<i>true</i>
'There is exactly one body in the room'	<i>true</i>
'Everything in the room is either a person or a body'	<i>true</i>
'Everything in the room is both a person and a body'	<i>indeterminate</i>
'Something in the room is both a person and a body'	<i>indeterminate</i>

All this shows that the account fully preserves the data. So it is formally adequate and metaphysically neutral, as desired.

As I said, the shooting/killing case is perfectly similar, so there is no need to go through the details. But for the sake of clarity let me briefly outline the picture in the ship case, (8). I have not discussed this case in the previous section, but as one might expect (and as Parsons pointed out in his book) this is another identity puzzle where supervaluationism would yield the wrong results if construed by reference to models with the same domain of quantification. For example, we would end up saying that there are at least five ships to choose from, which is absurd. If the domains are allowed to vary, however (as I am suggesting), then it is easy to see that the supervaluational account will fulfill all the intuitive desiderata. To recall: the data tell us that exactly one ship, A, left port, but as a result of a familiar repair/assembly process, two ships, B and C, docked (one with new parts, one consisting of the old parts reassembled). So here are our possible models:

- (8') *Type 1*: The domain of quantification includes three distinct ships, a, b, and c. Ship a left port and is the referent of 'A'. Ships b and c docked, and are the referents of 'B' and 'C', respectively.

Type 2: The domain of quantification includes two distinct ships, *a* and *g*. Ship *a* left port and docked, and is the referent of both ‘A’ and ‘B’. Ship *g* docked and is the referent of ‘C’.

Type 3: The domain of quantification includes two distinct ships, *a* and *b*. Ship *a* left port and docked, and is the referent of both ‘A’ and ‘C’. Ship *b* docked and is the referent of ‘B’.

At this point we can construct our supervaluation and we find that the only identity statement that has a definite truth-value (apart from self-equalities) is ‘B = C’, which turns out to be false. All other identity claims, as well as their negations, are indeterminate. It is indeterminate whether A is identical with B, with C, or with neither B nor C. This is what we expected. In addition, the supervaluation will yield the right sort of answer to cardinality questions. For example, it is easy to verify that the given scenario yields the following truth-value assignments:

(8") ‘There are exactly two ships’	<i>indeterminate</i>
‘There are exactly three ships’	<i>indeterminate</i>
‘There are at least two ships’	<i>true</i>
‘There are at most three ships’	<i>true</i>
‘There are either two or three ships’	<i>true</i>
‘Exactly one ship left port’	<i>true</i>
‘Exactly two ships docked’	<i>true</i>
‘Every ship that left port docked’	<i>indeterminate</i>
‘Some ship that left port docked’	<i>indeterminate</i>

(And so on.) Isn’t this exactly what we want? Isn’t this exactly what Parsons wants? I think it is. On the face of it, when properly spelled out the supervaluational account yields exactly the answers that we intuitively expect, and that preserve the data in the relevant sense—in this case as well as in the other identity puzzles.

6. Is This Just Superresolutionism?

If we agree with this way of doing supervaluationism, then, the above arguments against the semantic conception of indeterminacy founder. This is not to say that we have a complete vindication of the conception. Parsons, for one, has other misgivings about supervaluationism, specifically about its non-truth-functional character (on the one hand) and about the complications involved in specifying

exactly what counts as an admissible model in cases where this requires refining the extensions of our predicates (on the other). I'll come back to the first sort of misgiving in a moment. As for the second, there is no question that a fully articulated supervaluational semantics is bound to be extremely complex. There are very complex patterns of penumbral connection (in Kit Fine's terminology) that set constraints on the admissible ways of refining the predicates of our language, and it is not easy to see how the task can be fully accomplished. (In this respect, the examples reviewed above are relatively unproblematic and make things look easier than they may generally be.) However, this strikes me a general problem for semantics *tout court*, i.e., a problem for *any* attempt to lay down the semantics of a specific language, not just for supervaluationism. And there is no obvious reason why we should think that this general problem is insurmountable as a matter of principle. In his book Parsons goes some way towards showing how deep the difficulties lie. But his arguments are not conclusive, nor are they meant to be, so from the present perspective it seems fair to me to say that we do have a vindication of supervaluationism, at least with regard to the specific problems raised by the identity puzzles, and at least in principle.

There is, rather, a different sort of question that strikes me as important at this point. It can be briefly put as follows. How much ontological indeterminacy are we letting in by construing our semantic refinements as models with variable domains of quantification? To say that the domain may vary is to say that we have not settled on what there is. But isn't this a step away from a purely semantic account, a step into the territory of ontological indeterminacy?

In my view the answer is a straightforward no. Ontological indeterminacy, as I understand it (and as Parsons understands it, on my reckoning), occurs when there is no fact of the matter regarding whether a certain state of affairs obtains. There is no fact of the matter "because of the way the world is (or because of the way the world is not)". In particular, we can speak of ontological indeterminacy *vis-à-vis* existence if there is no fact of the matter regarding whether a certain putative entity, *a*, is a citizen of the world. And we can speak of ontological indeterminacy *vis-à-vis* identity if there is no fact of the matter regarding whether a certain entity *a*, a citizen of the world, is the same as a certain entity *b*, also a citizen of the world. That supervaluationism involves no ontological indeterminacy *vis-à-vis* identity should now be obvious. But neither does it involve ontological indeterminacy *vis-à-vis* existence. To say that we have not settled on what there is is not to say that the world leaves it open whether certain entities exist. More modestly, it amounts to saying that our models of the world—the models of the world that we construct by attaching a semantics to our lan-

guage—leave that open. In a way, this boils down to saying that we are not clear about the semantics of the quantifiers. Alternatively, this may be construed as a claim to the effect that we are not clear about the semantics of the existence predicate, or of whatever bits of language do the job of singling out the entities of which we can truly say that they exist. Either way, there is no implication that the relevant indeterminacy is due to the way the world is, or to the way the world is not. The indeterminacy is thoroughly semantic even if it affects the semantics of ontologically significant portions of our language. Or so I submit.

Still, it might be thought that once supervaluationism is construed as a semantic theory allowing for variable domains of quantification, supervaluationism turns out to be merely a variant of what Parsons calls superresolutionism, which unquestionably qualifies as a theory of ontological indeterminacy. Superresolutionism says this. Suppose there is genuine worldly indeterminacy. Now consider the various ways in which this indeterminacy might be *resolved*, i.e., ways in which the actual world might become determinate by determining (one way or the other) every indeterminate state of affairs. For example, if it is not determinate whether a certain property (say, personhood) applies to a given object (say, the human body in the room), one resolution will make the property determinately apply to that object, and another resolution will make the property determinately not apply to that object. Every such resolution would amount to a fully determinate world and would therefore classify each sentence of the language as either true or false. And if a sentence is true in every such resolution, we may say that the sentence is also true (in some plausible sense) in the actual world. Ditto for falsity. This is what Parsons calls *superresolved* truth (or *superresolved* falsity, respectively). And clearly enough, superresolved truth and superresolved falsity are very much like super-truth and super-falsity. Indeed, if supervaluationism is construed as indicated in the previous section, the two accounts coincide on every sentence.

Supervaluationism and superresolutionism are indeed very similar, up to coincidence of results. Yet this similarity falls short of identity. For even though the two accounts yield the same answers in every case, they are arrived at quite differently—superresolutionism by making *the world* more determinate, supervaluationism by making *our models of* the world more determinate. And this difference is all that matters when it comes to comparing the ontological conception of indeterminacy with the semantic conception. In fact, from the supervaluationist's perspective the similarity between the two accounts is welcome. For it shows that the semantic account of indeterminacy preserves essentially the same data as the ontological account, if this is construed superresolutionally.

So the question is not whether supervaluationism can be distinguished from superresolutionism, for it obviously can. The question is whether superresolutionism is a good way of cashing out the ontological conception of indeterminacy. Now, Parsons does not think so. He does not offer superresolutionism as a general theory of truth, because superresolutionism is not truth-functional (it can make a disjunction true even though both disjuncts are indeterminate, for instance), and for Parsons this amounts to making the assessment of a statement depend only partially on how its parts relate to the world. It makes semantics depend on the world only *indirectly*. As I mentioned, this is one of the reasons why Parsons does not like supervaluationism in the first place. Rather, Parsons says that something like superresolutionism is needed to explain why we are sometimes inclined to accept as true statements that apparently contradict the theory of indeterminate identity, which by itself is more naturally associated with a truth-functional three-valued semantics of a familiar sort. (*Which* semantics? Parsons favors Lukasiewicz's for its "naturalness". Yet other options are available and one may wonder how to justify one's final choice. I think this embarrassment of riches is deeply problematic, but never mind that.) The sort of cases where we may be inclined to accept as true statements that apparently contradict the theory of indeterminate identity are cases involving the counting of objects. In the ship scenario, for example, we are inclined to say that although it is indeterminate whether there are two ships or three, still there are *either two or three* ships all told. A truth-functional treatment of the disjunction connective prevents us from saying this, but a superresolutational account (just like the supervaluational account—see (8")) will deliver exactly that answer, since the disjunction will be true in every resolution of the indeterminate world corresponding to the envisaged scenario. If things are so, however, then the original motivation for objecting to the semantic treatment of indeterminacy because of its violation of truth-functionality is lost. Either we stick to one way of doing things, or we leave it to the context to decide which way to go. If we stick to one way, then a superresolutational account seems better because it delivers the right answers when it comes to counting, and so the supervaluational account should also be fine. If we leave it to the context to decide which way is better, then we can do that without buying into ontological indeterminacy. We can adopt a supervaluational account in those cases where we are inclined to abandon truth-functionality (and to read our statements "non-literally", in Parsons's terminology), and we can adopt a different, truth-functional partial semantics—e.g., a Lukasiewicz-style semantics—in those cases where we are inclined to favor truth-functionality (reading our statements "literally"). Supervaluationism is a

good way of dealing with semantic indeterminacy; but it is by no means the only possible way.

Besides, the superresolutional account appears to fare better also in cases where truth-functionality is not the issue. We want to say that exactly one ship left port, although it is indeterminate whether it is identical with one of those that docked. This is part of the data. But unless we rely on superresolutionism, the theory of indeterminate identity (which Parsons illustrates in detail with the help of Venn-like diagrams) will tell us that the number of ships that left port is indeterminate, since it is indeterminate whether A, that left, is identical with one of the ships that docked, B and C, and also whether either of these two ships left port. In cases such as these, Parsons reckons that sticking to the literal meaning of the statements that report our intuitions is misleading. In general, there is an ambiguity between counting the x 's such that $f(x)$ and counting the x 's such that *determinately* $f(x)$. And although in most circumstances this ambiguity is irrelevant, when it comes to counting it becomes relevant. If we are counting the ships that determinately left port, then we come up with exactly one; but if we are counting the ships that left port, with no assumption of determinacy, then we come up with no answer at all. Since the data tell us that exactly one ship left, we must resolve the ambiguity in the first way—what Parsons calls the “adorned” reading. And the semantics of the adorned reading is the one that we get by applying superresolutionism, at least under some plausible constraints on the admissible ways in which the indeterminacy of the world could be resolved. (We want to rule out, for instance, a resolution in which there are two ships to begin with—a possibility that the formal treatment of superresolutionism outlined by Parsons does not formally rule out.)

So superresolutionism may not do justice to the requirement of truth-functionality; it makes the assessment of a statement less directly a matter of how its parts relate to the world than a fully truth-functional account does. But on the face of it, there are plenty of cases where the disagreement between a superresolutional account and its alternatives favor the former, and I am not sure about whether there are cases where the disagreement favors the latter. So, on the face of it, superresolutionism appears to be a good if not the best way to cash out the truth conditions of our natural language statements on the assumption that the world is ontologically indeterminate. And since supervaluationism delivers exactly the same truth-conditions, one can hardly reject it as a way to articulate the semantic conception of indeterminacy. By contrast, if the friend of ontological indeterminacy is willing to switch from truth-functionality to superresolutionism depending on the context, I see no reason why a friend of semantic

indeterminacy should not be free to make a corresponding switch from a truth-functional semantics to a supervaluational semantics depending on the context. (I do not think that is necessary, but never mind that.) Either way, I conclude that the choice between ontological accounts and semantic accounts of indeterminacy is open.