that α is *just true*, he says nothing more nor less than that α is ttrue; however, his assertion carries the implicature that Max rejects $\neg\alpha$. *Just true* and *ttrue*, on this account, remain equivalent; it's just that an utterance of the former carries 'autobiographical' implicatures that the latter doesn't.

« *Parenthetical remark.* I return to the issue of acceptance and rejection—and, relatedly, probability theory—in Chapter 5.

I should also note that if truth were less than transparent (something other than *ttruth*), other options would emerge. Graham Priest (2006*b*), for example, agrees that *just truth*, like truth itself, is inconsistent; however, he rejects that truth and 'just truth' are equivalent. Unlike me, Priest takes truth to be an important (semantic) property, rather than a merely logical (see-through) device, and the resulting theory is a non-transparent theory of truth. On Priest's (non-transparent) theory, 'just true' is defined as $T(\ulcorner\alpha\urcorner) \wedge \neg T(\ulcorner\neg\alpha\urcorner)$, where $T(x)$ is Priest's (non-transparent) truth predicate. Since, on Priest's approach, $T(\ulcorner\neg\alpha\urcorner)$ fails to imply $\neg T(\ulcorner\alpha\urcorner)$, the equivalence between $T(\ulcorner\alpha\urcorner)$ and $T(\ulcorner\alpha\urcorner)\wedge\neg T(\ulcorner\neg\alpha\urcorner)$ breaks down. My focus, as throughout, is on *transparent truth* and its paradoxes; however, in Chapter 5 I discuss some of Priest's well-known (dialetheic) views. *End parenthetical.* »

3.4 Remarks on revenge

Some might think that the foregoing discussion of 'just true' misses the real issue. The real issue, one might press, is that dialetheic positions face a typical 'revenge' problem involving *just true*. The thought is that there's clearly a notion of *just true* that is essentially used in constructing the formal account, but one that—on pain of triviality—cannot be expressed in the overall theory. The aim of this section is to briefly touch on this issue.[4]

In the following subsections, I first discuss a broad, background topic to set the stage. (See §3.4.1.) I then turn to the topic of 'revenge', in particular 'just true' qua revenge problem. (See §3.4.2.) Finally, in §3.4.3, I give my reply to the alleged problem.

3.4.1 *Models and reality*

Like much in philosophical logic, constructing a formal account of truth is 'model building' in the ordinary 'paradigm' sense of 'model'. The point of such a model is to indicate how 'real truth' in our 'real language' can have the target (logical) features we take it to have—Release and Capture features or, in our case, full intersubstitutability of $Tr(\ulcorner\alpha\urcorner)$ and α. In that respect, formal accounts of truth are idealized models to be evaluated by their adequacy with respect to the 'real phenomena' they purport to model.

Formal accounts (or theories) of truth aim only indirectly at being accounts of truth. What we're doing in giving such an account is twofold.

[4]The following subsections, in addition to parts of this chapter above, are from Beall 2007*a*, which is the introductory essay to Beall 2007*b*.

1. We construct an artificial *model language*—one that's intended to serve as a heuristic, albeit idealized, model of our own 'real' language—and, in turn, give an account of how 'true' behaves in that language by constructing a precise account of *truth-in-that-language*.
2. We then claim that the behavior of 'true' in our language, at least in relevant, target respects, is like the behavior of the truth predicate in our model language.

By far the most dominant approach towards the first task—viz., constructing a model language—employs a classical set theory. (This is certainly the approach employed in this book, though it has largely been a background assumption.) One reason for doing so is that classical set theory is familiar, well-understood, and generally taken to be consistent. A related reason is that, in using a classical set theory, one's formal account of truth may serve as more than merely a heuristic picture; it can also serve as a 'model' in the technical sense of *establishing non-triviality* or, in non-paraconsistent settings, *consistency*.

That a classical set theory is used in constructing our artificial language serves to emphasize the heuristic, idealized nature of the construction. We know that, due to paradoxical sentences, there's no truth predicate in (and for) our 'real language' if our real language is (fully) classical.

« *Parenthetical remark.* The same applies if the truth predicate has an extension: the extension isn't really a classical set. Every classical set $\mathcal{S}$ is such that $x \in \mathcal{S}$ or $x \notin \mathcal{S}$, which, given paradoxical sentences, results in inconsistency. (The point is independent of 'size' issues. Classical *proper classes* are likewise such that $x \in \mathcal{C}$ or $x \notin \mathcal{C}$.) If $\mathcal{T}$ is the extension of $Tr(x)$ and $\mathcal{T}$ is a set, a sentence λ that 'says' $\ulcorner\lambda\urcorner \notin \mathcal{T}$ makes the point—assuming, as is plausible, suitable 'extension' versions of Capture and Release (e.g., $\alpha \Rightarrow \ulcorner\alpha\urcorner \in \mathcal{T}$, etc.). *End parenthetical.* »

The project, as above, is to show how we can have a truth predicate in our 'real language', despite the various paradoxical spandrels that arise from introducing our truth predicate (our see-through device). And the project, as above, is usually—if not always—carried out in a classical set theory. Does this mean that the project, as typically carried out, is inexorably doomed? Not at all. Just as in physics, where idealization is highly illuminating despite its distance from the real mess, so too in philosophical logic: the classical construction is illuminating and useful, despite its notable idealization. But it is idealized, and, pending argument, on the surface only heuristic. That's the upshot of using classical set theory.

One might think it odd that we can enjoy classical set theory or classical logic, but it is not uncommon among theories of *ttruth*. So long as classical logic is an extension of our (weaker) logic, we can enjoy classical set theory (or the like). Since classical logic is an extension of LP, which is the logic of our conditional-free $\mathcal{L}^+$ (i.e., base language plus 'ttrue'), you can stand squarely in an appropriate fragment of $\mathcal{L}^+$ and enjoy a perfectly classical theory—like classical set theory.

In effect, this is the situation that we enjoy when we are utilizing a classical 'metalanguage' in constructing our formal 'model language'.

3.4.2 *'Just true' qua revenge problem*

As above, in giving a formal theory of truth, one does not directly give a theory of *truth*; rather, one gives a theory of $\mathcal{L}_m$-truth, an account, for some formal 'model language' $\mathcal{L}_m$, of how $\mathcal{L}_m$'s truth predicate behaves, in particular, its logical behavior. By endorsing a formal theory of truth, one is endorsing that one's own truth predicate is relevantly like *that*, like the truth predicate in $\mathcal{L}_m$, at least with respect to various phenomena in question (e.g., logical behavior).

In what follows, I use '$\mathcal{L}_m$-truth' to abbreviate *the (logical) behavior—e.g., the logic of—$\mathcal{L}_m$'s truth predicate* (i.e., the logic of $\mathcal{L}_m$ predicate set up to be the $\mathcal{L}_m$'s truth predicate). As discussed below, this is generally very different from the metatheoretic, model-dependent notion of *truth in* $\mathcal{L}_m$, which we normally use to specify $\mathcal{L}_m$-truth (i.e., specify the logic of $\mathcal{L}_m$'s given truth predicate).

Revenge qua objection—revenger's revenge—is an *adequacy objection*. Typically, the revenger charges that a given 'model language' is inadequate due to expressive limitation. Let L be our 'real language', English or some such natural language, and let $\mathcal{L}_m$ be our heuristic model language. (Recall that '$\mathcal{L}_m$-truth' abbreviates 'the behavior of $\mathcal{L}_m$'s truth predicate'.) In broadest terms, the situation is this: we want our (heuristic) $\mathcal{L}_m$, and in particular $\mathcal{L}_m$-truth, to illuminate relevant features of our own truth predicate, to explain how, despite paradoxical sentences, our truth predicate achieves the features we take it to have. Revenge purports to show that $\mathcal{L}_m$ achieves its target features in virtue of lacking expressive features that L itself (our real language) appears to enjoy. But if $\mathcal{L}_m$ enjoys the target features only in virtue of lacking relevant features that our real L enjoys, then $\mathcal{L}_m$ is an inadequate model: it fails to show how L itself achieves its target features (e.g., consistency). That, in a nutshell, is one common shape of revenge.

In our case, the objection might run as follows. Let $\mathcal{L}_m$ be our formal 'model language' (sketched in Chapters 1 and 2). In constructing $\mathcal{L}_m$, we use—in our metalanguage—classical set theory, and we define *truth-in-$\mathcal{L}_m$*, which, with other similar notions, is used to discuss $\mathcal{L}_m$-truth (the behavior of $\mathcal{L}_m$'s truth predicate). In particular, *truth-in-$\mathcal{L}_m$* is defined as *designated*, which, in our case, amounts to being true at @ in a given model (i.e., $@ \models \alpha$ iff α is designated in the given model). (In the standard many-valued approach, discussed in Chapter 1 appendix, designation amounts to being assigned something in $\{1, \frac{1}{2}\}$.) Moreover, we can prove—in our metalanguage—that, despite paradoxical sentences, a sentence $Tr(\ulcorner \alpha \urcorner)$ is true-in-$\mathcal{L}_m$ exactly if α is true-in-$\mathcal{L}_m$. (Indeed, we get WT and, more generally, TP, as discussed in Chapter 2.)

The familiar 'revenge' charge at issue is that $\mathcal{L}_m$, so understood, is not an adequate model of our real language; it fails to illuminate how our own ttruth predicate, despite paradoxical sentences, achieves non-triviality. In particular, the revenger's charge is that $\mathcal{L}_m$-truth achieves its non-triviality in virtue of $\mathcal{L}_m$'s expressive poverty: $\mathcal{L}_m$ cannot, on pain of triviality, express certain notions *that our real language can express.*

Example. First, notice that we have a notion of *just false in* $\mathcal{L}_m$, which, like the other notions used to discuss $\mathcal{L}_m$-truth, is defined via our semantic values; in particular, α is just-false-in-$\mathcal{L}_m$ iff α is not-true-at-@ (in a given model, i.e., $@ \not\models \alpha$). Suppose, now, that $\mathcal{L}_m$ contains a predicate $JF(x)$ that defines $\{\alpha : \alpha \text{ is just false in } \mathcal{L}_m\}$. Given the resources of $\mathcal{L}_m$, there will be a sentence λ that says $JF(\ulcorner \lambda \urcorner)$. But, then, we can immediately prove—in the metalanguage—that λ is true-in-$\mathcal{L}_m$ iff $JF(\ulcorner \lambda \urcorner)$ is true-in-$\mathcal{L}_m$ iff λ is just-false-in-$\mathcal{L}_m$ iff λ is not true-in-$\mathcal{L}_m$. But this is impossible, since—by construction—no sentence of $\mathcal{L}_m$ is both true-in-$\mathcal{L}_m$ and not true-in-$\mathcal{L}_m$. Hence, we conclude that $\mathcal{L}_m$ cannot express *just false in* $\mathcal{L}_m$.

The same objection can be put more generally (and more concisely) by skipping the 'just false' notion and simply using *not true in* $\mathcal{L}_m$. Just suppose that $\mathcal{L}_m$ contains a predicate $NT(x)$ that defines $\{\alpha : \alpha \text{ is not true-in-}\mathcal{L}_m\}$, and the reasoning is the same. In turn, one concludes that $\mathcal{L}_m$ cannot express *is not true-in-$\mathcal{L}_m$*.

The revenger's charge, then, amounts to this: that our model language fails to be enough like our real language to explain at least one of the target phenomena, namely, *ttruth's non-triviality*. Our metalanguage is part of our 'real language', and we can—because we do—define $\{\alpha : \alpha \text{ is not true-in-}\mathcal{L}_m\}$ in our metalanguage. Since, as above, the given model language cannot similarly define $\{\alpha : \alpha \text{ is not true-in-}\mathcal{L}_m\}$, the given model language is inadequate: it fails to illuminate the target features of *ttruth*, in particular, its non-triviality.

3.4.3 *Remarks on revenge*

The foregoing 'revenge' objection turns on the claim that *is not true in* $\mathcal{L}_m$ (or, more narrowly, *is just false in* $\mathcal{L}_m$) cannot be expressed in $\mathcal{L}_m$. The relevance of such a result, if it in fact goes through (a point to which I return below), is not obvious. After all, even if the notion is not expressible in the model language, the given notion is a *classically constructed* notion; it is a 'model-dependent' notion—a notion that makes no sense apart from the given (classically constructed) models—defined entirely in a classical metalanguage. As such, the given notion is not one of the target (say, model-independent, or 'absolute') notions in our real language L that $\mathcal{L}_m$ is supposed to model. The question is the relevance of the revenger's result, if the result goes through at all.

At this stage, a bit more detail on the 'model language', and in particular its resources, is in order. For present purposes, I mention only two situations: classical arithmetic and classical set theory. In the appendices of Chapters 1 and 2, the former was assumed for simplicity, but the latter is equally available.[5] What is important to note is that if (classical) set theory is used as one's base

[5] As mentioned in Ch. 1's appendix, the non-triviality results go through where set theory (versus arithmetic) is the base language, but it is in various ways simpler—and certainly more standard—to use classical arithmetic as the base language.

'model language', in addition to its serving as one's metalanguage (in which one proceeds, as discussed above, to define *true-in-*$\mathcal{L}_m$ etc.), then the revenger's argument towards the inexpressibility of *true-in-*$\mathcal{L}_m$ clearly fails. After all, $\mathcal{L}_m$ enjoys classical set theory as a part, and hence can express any classical set-theoretic notion—a fortiori, *true-in-*$\mathcal{L}_m$. So, quite plainly, the revenger's arguments do not establish that such set-theoretic notions are not expressible in set theory. Instead of establishing that notions like *true-in-*$\mathcal{L}_m$ are not expressible in $\mathcal{L}_m$, the revenger's argument at best establishes that such notions do not play the role in $\mathcal{L}_m$ that $\mathcal{L}_m$'s truth predicate plays—namely, being a (transparent) truth predicate in $\mathcal{L}_m$. The relevance of such a result is not clear.

In the case in which classical arithmetic is the base model language, which case I have been assuming here (because of its use in earlier Chapters), the same question of relevance arises but for a slightly different reason. In particular, it is clear that set theory is not expressible in arithmetic, and so the revenger's argument(s) needn't be given to establish as much. What the arguments do establish—or, if need be, can be tweaked to do—is again that such model-dependent, set-theoretic notions do not play the role in $\mathcal{L}_m$ that $\mathcal{L}_m$'s truth predicate plays. This is correct, but, again, the relevance of the result is unclear.

Now, one might think that, regardless of which sort of model language is in play (arithmetic or set theory), the relevance of the revenger's argument(s) is plain. One might, for example, think that the semantics for $\mathcal{L}_m$ is intended to reflect the semantics of L, our real language. Since the semantics of the former essentially involves, for example, *not true-in-*$\mathcal{L}_m$, the semantics of our real language must involve something similar—at least if $\mathcal{L}_m$ is an adequate model of our real language. But, now, if there's nothing in $\mathcal{L}_m$ that plays exactly the role—and enjoys exactly the same (logical) behavior—in $\mathcal{L}_m$ as *true-in-*$\mathcal{L}_m$ plays in the formal semantics, we should conclude that $\mathcal{L}_m$ is an inadequate model of our real language L since our real language clearly does enjoy such a notion (viz., *true-in-*$\mathcal{L}_m$ itself).

Such an argument might serve to turn the revenger's result into a plainly relevant and powerful objection; however, the argument itself relies on various assumptions that involve quite complex issues. One conspicuous assumption is that the 'semantics' of $\mathcal{L}_m$ is intended to reflect the semantics of our real language L. This needn't be the case. For example, suppose that, with 'deflationists', one rejects that semantics—the semantics of our real language—is a matter of giving 'truth conditions' or otherwise involves some explanatory notion of truth. In the face of Liars (or other paradoxes), one still faces questions about one's truth predicate, and in particular its logical behavior. By way of answering such questions, one might proceed as above: construct a model language that purports to illuminate how one's real truth predicate enjoys its relevant features (e.g., Capture and Release) without collapsing from paradox. In constructing and, in turn, describing one's 'model language', one might give truth-conditional-like semantics for the model language by giving truth-in-a-model conditions for the language. If so, it is plain that the 'semantics' of the model language are not intended to reflect the 'real semantics' of one's real language; they may, in the end, be only tools used for illuminating the logic of our real language, versus illuminating the 'real semantics' of our real language. Indeed, this is the perspective that I endorse, namely, that our formal tools are just that: tools. The formal account of *ttruth* is given to elucidate the logical behavior of *ttruth*; it isn't intended to reflect the 'real semantics' of our real language, whatever such 'real semantics' might come to. So, a critical assumption in the argument above—the argument towards the relevance of the given inexpressibility results—requires argument.[6] Of course, pending a full theory of meaning—something I am not prepared to give—such an 'instrumentalist' reply is at best promissory, but there are some promising 'use-theoretic' approaches to meaning that might finish the job (e.g., Field 2001).

Meaning aside, I should note that Ross Brady (1989), in a closely related context, gestures at the sort of response to 'just true' qua revenge problem that I've advocated here. I do not know—and, in fact, somewhat doubt—whether Brady intends the 'instrumentalist' reply above, but his diagnosis of conflating model-dependent versus absolute (what he calls 'ordinary') is in agreement with the way I look at the matter. Brady writes (and here I take some liberty in paraphrasing in brackets):

> It seems to me that ... the [relevant notion] used in generating the [alleged problem] involves reference to the details of the model [used in our 'formal account']. That is, 'p takes the value t', or some equivalent, makes reference to the specific values of a model ... and thus goes beyond natural language expressions which just refer to truth and falsity. (Brady, 1989, p. 467)

The point, as I (perhaps idiosyncratically) read Brady, is along the lines that I've suggested. While we can—in our 'real', natural language—classically construct various notions such as 'just true in a model' or etc., we mustn't forget that these are classically constructed (i.e., bound by classical logic) and thereby not the non-model-dependent notions that we really care about. With respect to our formal, model-theoretic picture, we have our classically behaving 'just true' and so on; however, as Brady points out (and I endorse), this is all simply model-dependent. Our 'absolute' notion of *just true*, as I've said, is just *ttruth*.

3.5 Limited notions of 'just true'

I maintain, as above, that our notion of *just true* is just *ttruth*. The thought, perhaps based on taking the formal picture too seriously (see previous section),

[6] Hartry Field (2008), as I understand him, takes a very similar position. Indeed, though I've long endorsed such an 'instrumentalist' view of the formal account, I am indebted to Field's work for forcing me to be more forthright about the matter. Field's basic account of *ttruth* is sketched in the next chapter (although his position on similar revenge-like objections is left to his work). I should also emphasize that Field, unlike other (formal) truth theorists, is explicit in his use of set theory—versus arithmetic—as the base 'model' language. This affords responses to revenge-type arguments that I have not herein discussed, and for brevity leave to Field 2008.

that we need to recognize a notion of truth stronger than ttruth seems to me to be in need of argument. Still, what is worth noting is that there may be various notions of 'just true' available, should they be needed. In fact, there are clearly some such operators in the theory already advanced. By way of closing this chapter, I briefly discuss two natural options, one that's already available, one merely speculative.

To clarify the issue, suppose that we're after a notion of *just true* that is stronger than *ttruth*. In particular, far from having intersubstitutability (the chief see-through service of *ttruth*), the target 'just true' doesn't even satisfy Capture (either RC or CC). The thought is that while any α that is just true is ttrue, the converse, presumably, fails. The thought, more clearly, is that α may be ttrue without thereby being *just true*—on the target, stronger notion of 'just true'. At a minimum, then, our 'just true' operator $\mathbb{J}$, whatever it may be, satisfies Release (at least in rule form) but not Capture (in either form). What we want is at least (rule) Release for $\mathbb{J}$

» $\mathbb{J}$-Release: $\mathbb{J}(\alpha) \vdash \alpha$

but no Capture:

» No $\mathbb{J}$-Capture: $\alpha \nvdash \mathbb{J}(\alpha)$

In what follows, two options towards such an operator are discussed.

« *Parenthetical remark.* One might wonder why we want an *operator* (something that, syntactically, takes sentences and makes sentences) versus a *predicate* (which, syntactically, takes names or, generally, singular terms and yields sentences). As it turns out, when we enjoy, as we do, a fully transparent truth predicate, the distinction between operators and predicates diminishes in importance. Using *ttruth*, one can always define an operator in terms of a given predicate, and similarly vice versa (a predicate in terms of an operator). For ease of exposition, I simply concentrate on an operator. *End parenthetical.* »

3.5.1 *Available: many 'partial' operators*

One natural idea is to rely on the 'merely linguistic', classical base-language restriction of the current proposal and acknowledge a plurality of notions of *just true*, each tied to a different fragment of the language. For example, let $\mathcal{L}$ be our *ttruth*-free (and, generally, semantic-free) fragment of our language. We can then define a *just true* operator $\mathbb{J}_0$ as α *is ttrue and in the given fragment* $\mathcal{L}$. Since anything that's both ttrue and in $\mathcal{L}$ is ttrue, we have Release:

$$\mathbb{J}_0(\alpha) \vdash \alpha$$

Moreover, as desired, Capture fails:

$$\alpha \nvdash \mathbb{J}_0(\alpha)$$

That Capture fails is clear: just consider a ttrue α that is not in $\mathcal{L}$.

The given approach goes some way towards a non-vacuous account of 'just true', but it is at best partial. An obvious limitation is that, for all that's been said, $\mathbb{J}_0$ is not ttrue of any 'ttrue'-ful sentence, since such sentences are not in the given fragment.

There may be smoother fixes, but the current idea is that we simply use more and more such operators: $\mathbb{J}_1, \mathbb{J}_2$, and so on, each $\mathbb{J}_i$ being tied to the relevant fragment of our language. For example, let $\mathcal{L}_1$ comprise all base-language sentences and the sentences that result from substituting $Tr(\ulcorner\beta\urcorner)$ for any occurrence of β in α, where β and α are themselves base-language. Then $\mathbb{J}_1$ is tied to $\mathcal{L}_1$ along the lines of $\mathbb{J}_0$ and $\mathcal{L}$. In turn, we let $\mathbb{J}_2$ be tied to $\mathcal{L}_2$ (understood as per $\mathcal{L}_1$), and $\mathbb{J}_3$ be tied to $\mathcal{L}_3$, and so on.

This sort of approach goes only so far, but it yields the appropriate Release feature and avoids the relevant Capture feature. While I think that the approach is viable (notwithstanding details), I set further discussion aside. It is clear that, given the account advanced in Chapters 1 and 2, we have such (limited) notions of 'just true' if we need them. For now, I turn to one other (fairly speculative) approach.

3.5.2 *Speculative: another 'not'*

On my theory, negation is exhaustive but not exclusive; it satisfies E1 but not E2. If, unlike the suggestion in §3.5.1, we have a *single* notion of 'just true' that is stronger than ttruth, in the sense of satisfying Release but not Capture, it is unlikely to be definable in terms of negation and ttruth. Assuming—only for brief exploration—that we do have such a (single) notion, how might it be understood? That's the question.

A natural idea, I think, is to acknowledge a usage of 'not' distinct from negation, a usage that, in concert with ttruth, gives us our (supposed) single notion of *just true*. Let 'NoT' be the relevant term, although in English it is presumably spelled 'not'.

The role of *NoT*, on the going suggestion, is to serve as a 'commentary device' on *non-gluts*. The given usage is at work in our claims that α is a non-glut, that α is ttrue but NoT tfalse (or tfalse but NoT ttrue). Letting $\flat$ be *NoT*, its role is to serve as a non-vacuous way of expressing, for example, that α is a non-glutty ttruth (or, similarly, non-glutty tfalsehood).

$$\alpha \wedge \flat(\alpha \wedge \neg\alpha)$$

Clearly, negation does not do the NoT-ting trick. After all, if α is ttrue, then

$$\alpha \wedge \neg(\alpha \wedge \neg\alpha)$$

is ttrue, for *all* ttruths—gluts included. So, negation will not play the role of $\flat$; it will not serve as our supposed 'commentary device' on non-gluts. (Negation does not NoT, as one might say.)

How, then, is $\flat$ to be understood? Though I have no proof, my conjecture is that any purely 'extensional' operator will be inadequate. (I do have strong inductive evidence, but as yet no proof.) One might, as we do with negation, go 'intensional', invoking 'worlds' or the like, but a different approach might be more natural.

Instead of taking $\flat$ to be an ordinary operator, either extensional or intensional, we take $\flat$ to be a 'partially defined' connective. In particular, $\flat$ has only a *sufficient condition* for its application, in addition to a few other guiding principles. Indeed, as suggested below, the natural idea here is to follow the lead of da Costa and Alves (1977), where 'negation' is quite non-compositional.[7] The difference, on the present account, is that what we're specifying is *not* negation; rather, it is a usage of 'not' developed for one purpose—namely, to 'comment' on the non-gluts. Were it not for the given role, $\flat$ would be utterly superfluous, an otherwise terribly weak and uninteresting connective.

What we might want is a necessary and sufficient condition for the ttruth and, in turn, tfalsity of $\flat$-claims; however, it's not clear that we need as much—or, for that matter, can have as much in any simple fashion. The idea is modest. With da Costa and Alves, we recognize only a sufficient condition for $\flat$. In particular, where $\flat$ is our device, we give the following *sufficient* condition for the truth of $\flat$-claims.

N. If $w \not\models \alpha$ then $w \models \flat\alpha$

Notice that N is compatible with negation itself. Indeed, restricting to 'normal worlds' or, for simplicity, base worlds @, we already have it that if $@ \not\models \alpha$ then $@ \models \neg\alpha$, but, as with $\flat$, we do not have the converse. (See Chapter 1.) The difference is that *only* the minimal condition N governs $\flat$, leaving it very, very weak but adequate for its (supposed) 'commentary' role.

One clear desideratum of a commentary-on-non-gluts device is the failure of 'Non-Contradiction' in at least the following form.

$$\flat(\alpha \wedge \neg\alpha)$$

Call this 'LNC1' (only for a convenient tag). That LNC1 should fail is clear from the chief role of *NoT*, namely, to (non-vacuously) characterize α as being a *non-glut*, a sentence such that $\alpha \wedge \neg\alpha$ is NoT ttrue. If LNC1 held for all α, as its *negation* correlate does (i.e., replace $\flat$ with $\neg$), then *NoT* would fail to serve as a non-vacuous 'commentary device' on non-gluts. That LNC1 *does* fail is clear. In the formal picture, a relevant case is one in which $@ \models \alpha \wedge \neg\alpha$ but $@ \not\models \flat(\alpha \wedge \neg\alpha)$. Because $\flat$ is governed only by N, and N does not require anything of $\flat\alpha$ when α is true, such countermodels are ubiquitous.

NoT, so understood, is clearly very, very weak. Typical de Morgan principles fail, as does much (perhaps most) familiar 'negation-like' behavior, such as Double Negation features (in both directions). With respect to the latter, one could

[7] Actually, the *tradition* behind or around da Costa and Alves' work might be in tension with some of my proposal. The point is that the proposal joins da Costa and Alves in advancing only a relatively modest and, in effect, non-compositional account of the device.

follow da Costa and Alves and stipulate that if $w \models \flat\flat\alpha$ then so too $w \models \alpha$, thereby achieving at least Double Negation Elimation (DNE). (This still fails to ensure the converse, which, in fact, fails.) I see no obvious reason why DNE, let alone equivalence, should hold for the device, which, as above, is not our *negation* but rather a usage of 'not' that serves only to 'comment on non-gluts'. The given role of *NoT* requires a great deal of latitude and, ultimately, a very weak logic.

« *Parenthetical remark.* This isn't to say, of course, that, with further details in place, *NoT* couldn't collapse into negation over a suitable proper fragment of our language. Indeed, in previous work (Beall, 2005*a*) I have advocated a Diderik Batens-inspired 'adaptive' approach to the theory of *NoT*, wherein *NoT* is ultimately to be understood via a non-monotonic logic such that, for much of the language, *NoT* behaves in perfectly familiar ways but, when it comes to 'commentary work' (on non-gluts), *NoT*'s very weak logic takes over. I do not go into this approach here, mainly because I've since come to think that, as earlier in this chapter, we needn't recognize a notion of *just true* beyond *ttruth*, except perhaps for the ones—as in §3.5.1—already at hand. *End parenthetical.* »

Notice that, despite its weakness, *NoT* does exhibit some familiar behavior. In particular, as with negation, $\flat$ is exhaustive, that is, 'Excluded Middle' for $\flat$ holds.

$$\vdash \alpha \vee \flat\alpha$$

One notable difference is that, *unlike negation*, which is forced to take a 'holiday' in the face of gluts (i.e., negation can sometimes yield 'gaps' at abnormal worlds), *NoT* is exhaustive at *all* worlds—normal or abnormal. Let $w \in \mathcal{W}$. If $w \models \alpha$, then so too $w \models \alpha \vee \flat\alpha$. If $w \not\models \alpha$, principle N gives us that $w \models \flat\alpha$, and hence $w \models \alpha \vee \flat\alpha$.

That $\flat$, unlike negation, is 'everywhere exhaustive' (taking no holidays) requires that it not be 'explosive'. (Think of E1 and E2 discussed above in §3.2.) Not surprisingly, given its exhaustive behavior, $\flat$-ful spandrels of *ttruth*, like a sentence λ equivalent to $\flat Tr(\ulcorner\lambda\urcorner)$, wind up glutty, by familiar reasoning. As such, 'Explosion' fails for $\flat$.

$$\alpha, \flat\alpha \nvdash \beta$$

In the formal account, a model such that $@ \models \alpha$ and $@ \models \flat\alpha$ and $@ \models \neg\alpha$ but $@ \not\models \beta$ does the trick.

Spandrels aside, that *NoT* is exhaustive raises a question. One might wonder why we should recognize both negation and our 'loose' connective *NoT*. While I do not officially endorse *NoT*, a few comments on this issue are in order.

The main reason for rejecting that *NoT* is our only 'negation-like' device is that *negation* seems to enjoy a recognizably independent usage. For example, typical de Morgan behavior, or double-negation behavior seems to be at work in at least one common usage of 'not', namely, what I take to be *negation*. Admittedly, such appearances might arise from *NoT* behaving as such over some suitably proper fragment of our language, one on which our 'intuitions' of 'negation' are built. As such, the given appearances mightn't be ultimately decisive. In the end, it may be that the strongest reason to recognize negation (as I've construed it) as something different from NoT—if the latter is ultimately recognized—is that NoT is fairly wildly non-compositional. While I am very open to the idea of enjoying a non-compositional usage of 'not', I find it hard to accept that there's no compositional usage at all. For this reason, it strikes me that if NoT were recognized—invoked for purposes of cashing out a stronger notion of *just true* than what I've advanced—we should nonetheless see it as something other than negation. But I admit that the issue is not easy.

More might be said about the NoT approach to 'just true', but the basic idea is clear enough. The idea, in short, is that when we say that α is *just true*, we're using *ttruth* combined with a special (and especially weak) connective NoT. As above, one could very easily define a unary 'just true' connective $\mathbb{J}$ (it is just true that…) in terms of *ttruth* and NoT, namely,

$$\mathbb{J}\alpha \ =_{df} \ \alpha \wedge \flat(\alpha \wedge \neg\alpha)$$

This, in turn, will yield Release in at least rule form (RR) but, as desired, not Capture.

» $\mathbb{J}\alpha \vdash \alpha$. This follows from the basic features of Conjunction (namely, Simplification).
» $\alpha \nvdash \mathbb{J}\alpha$.

Whether the NoT approach is ultimately viable or, contrary to what I've suggested, even needed is for debate to tell. While I find the proposal to be interesting, I do not ultimately endorse the idea. The chief reason, as discussed above, is that I'm not convinced that the apparent notion of *just true* requires as much.

Summary

The apparent problem of *just true* is that it cannot be accommodated in a dialetheic framework. Whether this is correct depends on what *just true* amounts to. As in §3.1, it cannot be an operator that yields both E1 and E2, assuming, as I am, that we enjoy F1 and F2, namely, Capture and Release (in at least unrestricted rule form) and DP. What, then, is it?

In this chapter, I've argued that, aside from more limited (or speculative) notions discussed in §3.5, *just true* is just *ttruth*. This account respects the surface of 'just true', which appears to be either 'true and not false' or 'true and not a glut'. As discussed (see §3.3), both of these are equivalent to *ttruth*, at least if, as I assume, 'true' in 'just true' is just *ttruth*.

If one wants more from a notion of *just true* than *ttruth*, one at least has the various $\mathbb{J}_i$ operators discussed in §3.5.1, and perhaps—details notwithstanding—something along the NoT approach. But, again, pending some argument to the contrary, the need for more remains unclear.

CHAPTER THREE APPENDIX: ANOTHER APPROACH

This appendix briefly notes another approach to *just true* already available in BXTT (and, indeed, the basic star-simplified framework for BX). (I noticed the idea late, when this book was days away from press, at which point only two pages could be added.) Like the other avenues discussed in Chapter 3, this approach is not intended to capture *everything* that *everyone* wants;[8] however, it's arguably as much as one should reasonably demand.

Desiderata of a 'just true' operator $\mathbb{J}$

D1. Release (at least rule): $\mathbb{J}\alpha \vdash \alpha$.
Rationale. if α is *just true*, then surely it's true.

D2. No Capture: $\alpha \nvdash \mathbb{J}\alpha$.
Rationale. α may be a glut.

D3. Explosive: $\neg\alpha, \mathbb{J}\alpha \vdash \beta$.
Rationale. α's being *just true* should 'rule out' its being false.

D4. Not Exhaustive: $\nvdash \mathbb{J}\alpha \vee \mathbb{J}\neg\alpha$.
Rationale. we don't want that all α are just true or just false.

D5. Ensuring Detachment (etc.): $\mathbb{J}\alpha, \alpha \wedge (\alpha \supset \beta) \vdash \beta$.
Rationale. we want our 'consistent truth' or 'just truth' operator to do work validity-wise. (MMP is but one among many such 'inferences' at issue.)

D6. Non-triviality: not only do we want $\mathbb{J}$ to not yield triviality; we want a proof that it doesn't.

The idea

Our focus is BX and the star-simplified semantics (see Chapters 1–2), where truth in a model is truth at @ (with @ normal). To see the idea, let τ be any just-true sentence in an arbitrary (non-trivial) BX model: $@ \models \tau$ and $@^\star \models \tau$. I claim that

$$\tau \rightarrow \alpha$$

is a plausible *just true* operator in the D1–D6 sense.[9] Our key proposition is this:

Proposition *(Forcing Consistency: FC) Consider any non-trivial BX model, and let τ be a just-true sentence in the model. If $@ \models \tau \rightarrow \alpha$ then $@ \not\models \alpha \wedge \neg\alpha$.*

Proof of FC. Assume $@ \models \alpha \wedge \neg\alpha$, in which case $@ \models \alpha$ and $@ \models \neg\alpha$. Since $@ \models \neg\alpha$ we have $@^\star \not\models \alpha$. But $@^\star \models \tau$. Hence, there's a point (viz., $@^\star$) at which τ is true but α not. As @ is normal, we have: $@ \not\models \tau \rightarrow \alpha$.

[8] In fact, there's a sort of essential 'incompleteness' (of predictable sort) that remains.

[9] Once the idea is seen, it's obvious that many sentences will do the trick. I've picked the simplest.

The desiderata D1–D6

Note, first, that the trivial model is inessential to the logic BX.[10] Hence, we can restrict our class of models to the non-trivial ones, and we thereby ensure at least one just-true sentence—a sentence that plays the role of τ. Letting $\mathbb{J}\alpha$ be $\tau \rightarrow \alpha$, the desiderata, in turn, are achieved.

D1. $\mathbb{J}\alpha \vdash \alpha$. *Proof.* Since $@ \models \tau$ and $@ \models \tau \rightarrow \alpha$, we have $@ \models \alpha$.

D2. $\alpha \nvdash \mathbb{J}\alpha$. *Proof.* Let $@ \models \alpha \wedge \neg\alpha$, in which case FC gives $@ \not\models \tau \rightarrow \alpha$.

D3. $\neg\alpha, \mathbb{J}\alpha \vdash \beta$. *Proof.* This follows from D1 and FC.

D4. $\nvdash \mathbb{J}\alpha \vee \mathbb{J}\neg\alpha$. *Proof.* Let $@ \models \alpha \wedge \neg\alpha$, in which case FC gives $@ \not\models \tau \rightarrow \alpha$ *and*, letting α be $\neg\alpha$ in FC, $@ \not\models \tau \rightarrow \neg\alpha$.

D5. $\mathbb{J}\alpha, \alpha \wedge (\alpha \supset \beta) \vdash \beta$. *Proof.* Given $@ \models \mathbb{J}\alpha$, FC delivers $@ \not\models \neg\alpha$, in which case we have to have $@ \models \beta$ since we have $@ \models \neg\alpha \vee \beta$.

D6. Non-triviality? The non-triviality proof in Chapter 2 Appendix covers BX and, in particular, the theory BXTT. $\mathbb{J}$, as here given, adds nothing that isn't already available in the theory. [Indeed, we can make the point stronger by adding a constant $\top$ axiomatized $\alpha \vdash \top$ and $\neg\top \vdash \alpha$. This is covered by a non-triviality construction along the lines of Chapter 2 Appendix, wherein 'truth-value' constants are added. See Brady 1989. Chapter 5 contains related discussion.]

A few remarks

As Chapter 3 makes plain (see 'incoherent operators'), there are essential limits on any plausible candidate for a (coherent) 'just true' operator. Our $\mathbb{J}\alpha$, as $\tau \rightarrow \alpha$, is no different. One might think that if α is neither ttrue nor tfalse, then it ought count as just true or just false. Along these lines, one might want to impose

$$\neg(\alpha \wedge \neg\alpha) \vdash \mathbb{J}\alpha \vee \mathbb{J}\neg\alpha$$

as another desideratum; however, quick reflection shows that this is asking to transcend the limits—asking for incoherence. After all, according to BXTT, *every* sentence α is such that $\neg(\alpha \wedge \neg\alpha)$ is true. But now consider predictable sentences like some ζ equivalent to $\mathbb{J}\neg\zeta$. On the going account, ζ has the form $\tau \rightarrow \neg\zeta$, saying in effect 'I am just false'. Pushing the going desideratum clearly pushes to incoherence: since we have $\neg(\zeta \wedge \neg\zeta)$, we would get $\mathbb{J}\zeta \vee \mathbb{J}\neg\zeta$, which—as is straightforward to check—implies triviality.

The current account of *just true* need not be seen as replacing the other approaches discussed in Chapter 3. The point of this appendix is to note the rather surprising fact that, in effect, BX already contains a plausible 'just true' operator. Certainly, the features of $\tau \rightarrow \alpha$ canvassed here are available in BXTT, for which we enjoy a non-triviality proof.

[10] The trivial model cannot invalidate anything that isn't already invalidated. (Whether particular BX-ish *theories* need—for some reason—to acknowledge the trivial model is another matter. BXTT, as advanced in this book, shuns a trivial model.)