

Multi-track dispositions

Barbara Vetter, HU Berlin
barbara.vetter@hu-berlin.de

This is the penultimate draft of a paper published in *The Philosophical Quarterly* 2013, 63, 330-352. Please cite the published version only.

Abstract

It is a familiar point that many ordinary dispositions are multi-track, that is, not fully and adequately characterizable by a single conditional. In this paper, I argue that both the extent and the implications of this point have been severely underestimated. First, I provide new arguments to show that every disposition whose stimulus condition is a determinable quantity must be infinitely multi-track. Second, I argue that this result should incline us to move away from the standard assumption that dispositions are in some way importantly linked to conditionals, as presupposed by the debate about various versions of the ‘conditional analysis’ of dispositions. I introduce an alternative conception of dispositionality, which is motivated by linguistic observations about dispositional adjectives and links dispositions to possibility instead of conditionals. I argue that, because of the multi-track nature of dispositions, the possibility-based conception of dispositions is to be preferred.

1 Introduction

Dispositional properties have traditionally been characterized in terms of conditionals. Thus a fragile vase is one that would break if it were struck, an irascible person is one who would get angry if she were provoked, and a water-soluble substance is one that would dissolve if it were immersed in water. On the simple conditional analysis, the conditional is all we need to provide a complete analysis of such a disposition. Thus the simple conditional analysis of fragility goes as follows:

(SCA) x is fragile iff, were x struck, x would break.

It is now agreed by most that the simple conditional analysis is false. A vase’s fragility may be *finked* by a sorcerer who would immediately remove its fragility should it be struck, or *masked* by anti-deformation packaging which would prevent it from breaking even if it were struck. In both cases the left-hand side of the biconditional in (SCA) is true, but its right-hand side is false. Moreover, a non-fragile concrete block may have a *reverse fink* if a sorcerer were to make it fragile as soon as it were struck, or it may *mimic* fragility if it is attached to an explosive that would go off and shatter the block

if it were struck. In these latter two cases, the right-hand side of the biconditional is true although the left-hand side is false.¹

But there is another kind of problem with (SCA) altogether, which David Manley and Ryan Wasserman have done much to establish². These are *structural* problems with (SCA) and any other attempt at analysing a disposition in terms of a single conditional. A first structural problem is that dispositions are *gradable*: a champagne glass is more fragile than an ordinary tumbler, and some of us are more irascible than others. This is a problem because (SCA) does not afford the materials for an account of the gradability of ‘fragile’: unlike fragility, the truth of a conditional is an all-or-nothing matter. Second, and relatedly, it appears that disposition terms are context-sensitive: what counts as fragile in the context of aeronautics may not count as fragile in the context of ordinary life. It is appealing to account for that context-sensitivity in terms of the degrees of fragility: *whether* something counts as fragile in a given context is a matter of *how* fragile it is. (SCA), again, cannot deliver such an account of context-sensitivity, because it cannot deliver an account of gradability.

To solve these structural problems, Manley and Wasserman propose that (SCA)’s mistake lies in its focus on a single conditional. Instead, they suggest, we need to think of a disposition as correlated with a large (in fact, a non-denumerably infinite) number of conditionals, each of which specifies in its antecedent a ‘fully specific scenario that settles everything causally relevant to the manifestation of the disposition’ (‘Gradable Dispositions’, p. 72). We can then say that x is more fragile than y just in case *more* of these specific conditionals are true of x than of y, and having established this comparative ordering, context-sensitivity can work by setting a threshold on it.

The problems are structural because they depend on a very general structural feature of (SCA): the correspondence between a disposition and single conditional. The problems arise whatever we think is the right response to the problems of finks and masks. Various responses to these problems have been suggested, of which Lewis’s ‘reformed conditional analysis’ may be the best known. Typically, however, these have shared the structural feature of (SCA) which gives rise to the structural problems pointed out by Manley and Wasserman: they correlate one disposition to one conditional.

Take, for instance, Lewis’s ‘reformed conditional analysis’ (p. 157):

¹C. B. Martin, ‘Dispositions and Conditionals’, *The Philosophical Quarterly* 44 (1994), pp. 1–8; D. Lewis, ‘Finkish Dispositions’, *The Philosophical Quarterly* 47 (1997), pp. 143–158; M. Johnston, ‘How to Speak of the Colors’, *Philosophical Studies* 68 (1992), pp. 221–263; A. Bird, ‘Dispositions and Antidotes’, *The Philosophical Quarterly* 48 (1998), pp. 227–234. Bird speaks of ‘antidotes’ where I, following Johnston, speak of ‘masks’.

²D. Manley and R. Wasserman, ‘A Gradable Approach to Dispositions’, *The Philosophical Quarterly* 57 (2007), pp. 68–75; ‘On Linking Dispositions and Conditionals’, *Mind* 117 (2008), pp. 59–84.

(RCA) Something x is disposed at time t to give response r to stimulus s iff, for some intrinsic property B that x has at t , for some time t' after t , if x were to undergo stimulus s at time t and retain property B until t' , s and x 's having of B would jointly be an x -complete cause of x 's giving response r .

The conditional is considerably more complicated than that on (SCA)'s right-hand side, but it is a single conditional no less. Note, however, that Lewis switches (as have most subsequent authors) from ordinary disposition ascriptions to *overt* disposition ascriptions of the form ' x is disposed to ... if ...'. The aim is to avoid masks by incorporating their absence into the stimulus condition. But it may be observed that he thereby also evades the structural problems, insofar as they concern fragility. (RCA) carries no commitment to a correspondence between fragility and a single conditional. It does, however, carry a commitment to a correspondence between an overtly ascribed disposition and a single conditional. The structural problems are thereby not avoided, they are merely transferred to the overtly ascribed dispositions. The same goes for more recent attempts to salvage a conditional analysis.

Sungho Choi³ defends the simple conditional analysis for *overt* disposition ascriptions (p. 374, slightly adapted and simplified):

(SCA*) x is disposed to M if S iff, if x were S , x would M .

(SCA) fails, according to Choi, not because (SCA*) fails but because it specifies the wrong stimulus and manifestation for fragility. The correct specification would rely on something like *ideal* or *normal* conditions⁴, which exclude masks and finks.

Jesse Steinberg⁵ has offered a different emendation of the simple conditional analysis. He holds that, while finks and masks provide counter-examples to (SCA) and (SCA*) alike, they can be accommodated by a simple *ceteris paribus* clause, yielding (p. 340):

An object is disposed to Φ when conditions C obtain if and only if, if conditions C were to obtain, then the object would Φ *ceteris paribus*.

Steinberg's analysis, like Lewis's and Choi's, tacitly requires that for any values of Φ and C there is a single (suitably qualified) conditional to characterize the disposition to Φ if C .

These responses to the problems of finking and masking, thus, do not address the structural problems with (SCA), they merely transfer them to overt disposition ascriptions.

³S. Choi, 'The Simple vs. Reformed Conditional Analysis of Dispositions', *Synthese* 148 (2006), pp. 369–379.

⁴Choi (p. 378, fn. 3) refers to S. Mumford, *Dispositions* (Oxford University Press, 1998), and to W. Malzkorn, 'Realism, Functionalism and the Conditional Analysis of Dispositions', *The Philosophical Quarterly* 50 (2000), pp. 452–469.

⁵J.R. Steinberg, 'Dispositions and Subjunctives', *Philosophical Studies* 148 (2010), pp. 323–341.

A different reaction to problems of finks and masks has been that of *dispositionalists* such as C.B. Martin and George Molnar⁶, who take these problems to show that dispositions are unanalysable, and indeed metaphysically irreducible. Insofar as they take (SCA) as their starting point, however, even dispositionalists are subject to the structural problems. In a sense, the structural problems are prior to the question of whether dispositions can be analysed or reduced. For that question arises only once we have a firm grasp of what the analysans, or the terms of the reduction, are supposed to be. If (SCA) fails for structural reasons, then the question whether it or anything that is structurally like it is an analysis or a reduction of dispositions is moot.

In this paper, I am going to add another structural problem to those which Manley and Wasserman have pointed out, and I will argue that the structural problems have wide-ranging consequences. In a first part of the paper (sections 2-3), I provide a novel argument for the conclusion that any but the most contrived dispositions are infinitely multi-track, in a sense to be precisified in section 2. I first consider the relatively uncontroversial case of fragility (section 2), and then argue that the same considerations which make it multi-track also apply to overtly ascribed dispositions, such as the disposition to break if struck (section 3).

The conclusion I draw here is not entirely new – it has been drawn by Manley and Wasserman⁷ – but my argument is novel. In fact, my structural argument is stronger than Manley and Wasserman's. They show that analyses of (SCA)'s form fail to capture certain features of dispositions: gradability and context-sensitivity. It is open, though undesirable, for the defender of such analyses to reject the need for those features to be captured. My argument shows that such analyses not only fail to capture certain features, but fail as an account of disposition ascriptions altogether. My considerations, like Manley and Wasserman's, will be purely structural, and entirely independent of issues of finks and masks. I will occasionally adopt Steinberg's *ceteris paribus* clause to mark this fact.

In the second part of the paper (sections 4-5), I draw some surprising conclusions from the multi-track nature of (almost) all dispositions. Section 4 will draw the general picture of dispositions that results from the first part of the paper in conjunction with the widespread assumption that conditionals capture the very nature of a dispositional property. Section 5 will introduce an independently motivated alternative conception of dispositions, which links them to possibility rather than to conditionals, and provides a simpler and more unified picture of ordinary dispositions. Together, the two sections establish a strong *prima facie* case for favouring the alternative conception over the standard assumption.

⁶Martin, 'Dispositions and Conditionals'; G. Molnar. *Powers. A Study in Metaphysics* (Oxford University Press, 2003).

⁷And by a few others: D. H. Mellor, 'The Semantics and Ontology of Dispositions', *Mind* 109 (2000), pp. 757–780; T. Cross, 'What is a Disposition?', *Synthese* 144 (2005), pp. 321–341.

2 The problem of qualitative diversity

A multi-track disposition, as I shall use the term, is a disposition which *cannot* be characterized adequately in terms of a single conditional. What is it for a disposition to be ‘characterized adequately’ in terms of a single conditional? I take it that it is for (at least) the following two conditions to hold:

(ST) Where D is a disposition and C a conditional of the form: if x were S, then x would be M, D is adequately characterized by C (and C alone) only if

- (1) For all objects x that have D, if x were S, then *ceteris paribus* x would be M.
- (2) For all objects x that have D, if x manifests D at t, then x is M at t and x is S at or before t.

Single-track dispositions are those that are characterized adequately by some one conditional. Hence if D is a single-track disposition, then there is a conditional C that fulfils conditions (1) and (2). If there is no conditional that fulfils the two conditions, and hence no conditional that alone adequately characterizes D, I shall call D a multi-track disposition.

(1) is simply the condition that a conditional which alone characterizes a disposition must be true, *ceteris paribus* (i.e. finking and masking aside), of all those objects which possess the disposition. This seems uncontroversial; for how else should objects count as possessing the disposition in question? Condition (2) is slightly less familiar but should nonetheless be uncontroversial. It is based on the idea, first, that if a disposition is characterized by a single conditional, it is characterized by a single pair of stimulus condition and manifestation; and second, that if the disposition is so characterized, it is only manifested in yielding the manifestation upon being subjected to the stimulus condition. Again, this seems utterly uncontroversial, for how else would D be manifested? If D were manifested, say, in an object’s being M upon being S* instead of S, then C alone would not adequately characterize D; a further conditional with S* in its antecedent would have to be added.

It is easy to see that most of our everyday dispositions are multi-track. A fragile glass may manifest its fragility in breaking upon being hit with a spoon, being dropped onto the floor, being sung to by a soprano, or being subjected to pressure over a period of time. Fragile parchments break upon being merely touched, and a fragile old wooden chair may split when transferred into a different temperature. Irrascible people may manifest their irascibility by becoming angry upon being yelled at, being politely told to wait, or being disagreed with. At first look, what we have here are qualitatively different stimulus conditions, and accordingly several different conditionals: x would break if it were dropped onto the floor, x would break if it were touched, x would break if it were sung to by a soprano, and so forth. It is

worth spelling out in some detail why this leads to fragility being a multi-track disposition according to (ST) above.

Suppose (for reductio) that fragility was to be treated as a single-track disposition. Given the different stimulus conditions that I have cited, how might we come by a single conditional to characterize that putative single-track disposition in a way concordant with (ST)? There are two things we might do.

(A) We may *pick one stimulus* from our list and focus on the resulting conditional. There may be worries about clause (1): given the variety among fragile objects, is there really one stimulus condition (and hence one conditional) that works for them all? Let us suppose, however, that there is. The conditional

(C1) If x were hit with a hammer, x would break

has a good chance of being true, *ceteris paribus*, of all fragile objects. The deeper problem lies with clause (2). By (2), any fragile object that manifested its fragility would have to do so upon being hit with a hammer. Breaking upon being hit with a spoon, being dropped onto the floor or being transferred into a different temperature will no longer count as a manifesting an object's fragility. (C1) will classify all the right objects as fragile. But it will not classify all the right breakings as manifesting fragility: only those breakings which happen upon being struck with a hammer will count as manifesting fragility. (C1), then, falsifies clause (2).

(B) The obvious solution is to *generalize the stimulus* by providing a condition which captures all the conditions under which fragility is ever manifested. Define *being stressed* as the disjunction of all the properties that may trigger a fragile object to break: to be stressed is to be hit with a spoon *or* dropped onto the floor *or* sung to by a soprano *or* subjected to pressure over a period of time *or* transferred into a different temperature *or* ... The conditional to characterize fragility is now

(C2) If x were stressed, x would break.

A disjunctive property such as being stressed is *multiply realizable*: one object may be stressed by being hit with a spoon, another by being transferred into a different temperature, and so on. Thus a glass's breaking upon being hit with a spoon will count as the object's breaking upon being stressed, as will a chair's breaking upon being transferred into a different temperature. Clause (2) is no longer violated.

The problem is that this strategy now falsifies (1). Suppose I uttered (C2), replacing 'x' by 'this vase'. You may take the vase, transfer it to a different temperature, and point out that the vase has now been stressed but did not break. It may well be true that the vase would break if it were hit. But (C2)

does nothing to privilege any particular one of the disjuncts that were used to define stressing. If it did, it would simply collapse back into strategy (A) and (C1). The very feature that made it superior to (C1), its generality, makes (C2) more susceptible to refutation than it should be.

Of course, clause (1) required only that a conditional which characterizes fragility be true of all fragile objects *ceteris paribus*. But that is of little help in the present context. The *ceteris paribus* clause was inserted to take care of finks, masks etc. But what we have here is not a case of finking or masking. Finks and masks interfere once the right stimulus condition has been applied. But our case is not one of interference. It is, rather, a case in which the right stimulus has not been applied.

It will help to spell this out in terms of the standard semantics for the counterfactual conditional. According to the orthodox Lewisian semantics, a counterfactual $A \Box \rightarrow C$ is true (at our world) just in case all the closest worlds in which A holds are also worlds in which C holds. (This is a simplification, but good enough for our purposes.) Note that, where A is equivalent to a disjunction $A_1 \vee A_2 \vee \dots \vee A_n$, the truth of $A \Box \rightarrow C$ does not follow from the truth of $A_i \Box \rightarrow C$ for any one disjunct A_i . Even if all the closest A_i -worlds are C -worlds, there may always be a non- A_i -world among the closest A -worlds, say an A_j -world, in which C is not true.

Accordingly, the truth of (C2) does not follow from the truth of ‘If x were hit, x would break’, although the antecedent of (C2) is equivalent to a disjunction one of whose disjuncts is that x is hit. This is so because there is no guarantee that the closest stressing-worlds include only hitting-worlds. It is easy to imagine that there is a world in which the vase is not hit, but transferred to a different temperature, and that world is at least as close as the closest worlds in which the vase is being hit. In fact, this may be so precisely because we know that the vase would break if it were hit and consequently take greater care not to hit it than we do not to transfer it to a different temperature. But if this is so, then (C2) is false, despite the fact that the vase is fragile, and clause (1) is violated.

This concludes part (B) of my argument.

We now have a dilemma: in formulating a single conditional to characterize fragility, we must either *pick one stimulus* condition to the exclusion of others, and hence violate (2) because that one stimulus condition may not be involved in all the manifestations of fragility; or we must *generalize the stimulus*, and violate (1) because there are too many ways for the generalized conditional to be false even of a fragile object. Either way, fragility is not a single-track disposition. The argument did not rely on any features that are specific to fragility; parallel arguments can be run for irascibility and any other disposition that has multiple qualitatively different stimulus conditions.

I conclude that fragility and irascibility are multi-track dispositions. I do not expect much controversy concerning this result. It is one of the reasons why philosophers have preferred to discuss the better-behaved, 'overt' disposition ascriptions such as 'x is disposed to break if hit' or 'x is disposed to break if heated'. Those, it is often thought, ascribe single-track dispositions. After all, they specify a single stimulus condition. They must therefore escape the problems that I have pointed out.

I am now going to argue that this thought is mistaken: the problems with (1) and (2) repeat themselves even with such qualitatively uniform dispositions.

3 The Problem of quantitative diversity

Not only dispositions, but also their typical stimulus conditions come in degrees. A glass can be struck with a greater or lesser force, a vase can be dropped from a greater or lesser height, a person can be yelled at more or less loudly. These properties are *quantities*: determinable properties with a range of determinates ordered by a relation such as that of *being greater than*. Striking, for instance, if understood as the exertion of mechanical force, has determinates such as *striking with a force of 8.35N*. Typically, the disposition's degree is (inversely) correlated with that of its stimulus. Thus, in general, a glass that is more fragile will break if struck with a lesser force, and one which is less fragile will break only if struck with a greater force.

In the remainder of this section, I will focus exclusively on the disposition to break if struck. For convenience, I will abbreviate 'the disposition to break if struck' to 'fragility*', and say that objects with this particular disposition are fragile*. When, in what follows, I appeal to intuitions about fragility* and fragile* objects, readers are invited to examine either their intuitions regarding the disposition to break if struck, or their intuitions concerning fragility but with the proviso that all stimulus conditions except striking are irrelevant. I will assume for the sake of the argument that *being struck* has the qualitative homogeneity that forestalls the argument of section 2, but will argue that its quantitative nature leads to problems very similar to those that faced the qualitatively different stimulus conditions in section 2.

We may begin by noting that the determinable property of *being struck* relates to its determinates in much the same way as the disjunctive property of *being stressed*, discussed in the previous section, related to the properties that formed its disjuncts. An object may be stressed by being struck, heated, dropped, sung to, or subjected to pressure. Similarly, an object may be struck by being struck with 1.01N, being struck with 8.35N, or being struck with 142.56N, and so forth. For any instantiation of *being stressed*, there is a particular disjunct of the property of being stressed that is instantiated; and

similarly, for any instantiation of *being struck*, there is a particular determinate of the determinable property of being struck that is instantiated.

Where we had, in the case of fragility, a multiplicity of *qualitatively* diverse stimulus conditions, we now have, in the case of fragility*, a multiplicity of qualitatively homogeneous but *quantitatively* diverse stimulus conditions. If we want to formulate a single conditional to characterize fragility*, we again have the following two options in dealing with this multiplicity.

(A*) We may *pick one stimulus* and consider only one of the many determinates, say being struck with 8.35N, as the stimulus condition of fragility*. However, the corresponding conditional

(C3) If x were struck with 8.35N, x would break

would falsify condition (2) of the characterization of a single-track disposition in (ST) above. Clearly, not all manifestations of fragility* are preceded by the fragile object's being struck with exactly 8.35N.

(B*) Again, the obvious alternative is to *generalize the stimulus*. To avoid privileging any particular determinate stimulus condition, we should incorporate the determinable and characterize fragility by the conditional

(C4) If x were struck, x would break

However, this strategy now violates clause (1) of (ST). Take a fragile* vase which would break, *ceteris paribus*, if struck with a force of 5N or more, but not if struck with any lesser force. Call this vase *Ming*. Now strike Ming very lightly, say with 2N. It will not break, and (C4) will therefore not be true of it. Once again, this is not a case of finking or masking: finks and masks interfere once the right stimulus has been applied, but the right stimulus has precisely not been applied in our case. In terms of the formal semantics, (C4) is true just in case x breaks in all the closest worlds where it is struck. But a fragile* object may be such that (i) it breaks in all the closest worlds where it is struck with 5N or more, yet (ii) among the closest worlds in which it is struck there are worlds in which it is struck with less than 5N and it does not break. In fact, Ming is – by stipulation – just such an object. To guarantee the truth of (C4), we would need to make sure that some of the at-least-5N worlds count as closer than any less-than-5N worlds; we would have to privilege the former worlds over the latter. But (C4), like (C2), provides no foothold for such a principled privilege.

So much for the parallel between fragility and fragility*. But perhaps the privilege is easier to come by in the present case. The determinable quantity of being struck is not entirely analogous to the disjunctive property of being stressed. Unlike the disjuncts of being stressed, the determinates of being struck are *ordered* on a continuous scale of magnitudes, from the very slight to ever greater forces. This

ordering allows for more than just the picking out of one determinate (as in (A*)) or the inclusion of all of them (as in (B*)). It allows, in particular, for the setting of a threshold and consideration of everything above or below that threshold. The problem with Ming in (B*) was precisely that worlds below the 5N threshold were allowed to count as relevant. What we need to do is exclude such worlds.

The idea, then, is to incorporate into (C4) a threshold value. This connects happily with the observation that ‘fragile’, as well as ‘fragile*’, is context-sensitive. Each context will determine a threshold value such that objects which would break, *ceteris paribus*, if struck with at least that force count as fragile* while others do not. As I said above, degrees of fragility* appear inversely correlated to the degree of the stimulus that is required for its manifestation: the more fragile* an object is, the smaller the force it takes to break it. By setting a threshold on the scale of forces, we determine *how* fragile an object has to be to count as fragile* within a given context.

For each context C, then, the property ascribed in that context by ‘x is fragile*’ (or by the unabbreviated ‘x is disposed to break if struck’), is adequately characterized by a conditional of the form:

(C5) x would break if x were struck with a force of at least n_C ,

where n_C is the threshold value determined by C.

It may be objected that the proposal is oversimplified. In particular, it may be claimed that a context does not simply set one threshold value for all kinds of objects. Perhaps the threshold value varies, within a given context, with the kind of object to which fragility is ascribed.⁸ In order to avoid commitment on this question, I will continue with the simplification but restrict my examples to only one kind of object, vases. I am not convinced that thresholds are kind-relative, but this will not affect my argument.

However this question is resolved, (C5) will not do either. The reason, in general, is that *being struck with at least n_C* still has the structure of a determinable quantity.

Take a context C, and suppose that the minimal force required to break those objects which count as fragile* in C ranges from 0.1N to 10N. Anything that would break, *ceteris paribus*, only if struck with a force greater than 10N does not count as fragile* in C (though it may in a different context). Clearly, C is a context in which Ming counts as fragile*. After all, it would break if it were struck with at least 5N. Now, where in the interval from 0.1N to 10N is n_C to be set?

We have three options here. (I will occasionally drop explicit relativization to C in what follows, but all ascriptions of fragility* are to be understood as true in the context C unless stated otherwise.)

⁸Thanks to an anonymous referee for bringing this to my attention.

First option: set the threshold at 0.1N (or below), with the most fragile* objects. The conditional that characterizes all objects that count as fragile* in C is: if x were struck with at least 0.1N, then x would break.

But that conditional is not true of any but the most fragile* objects. *Being struck with at least 0.1N* exhibits the by now familiar structure of a determinable property that we have seen in the property of being struck. An object may be struck with at least 0.1N by being struck with exactly 0.1N, with 2.0N, or with 8.3N, and so forth. Accordingly, our new conditional can be refuted in just the same way. Take Ming and strike it with a force of exactly 2N; it will not break. Since being struck with a force of exactly 2N is a way of being struck with at least 0.1N, the conditional at issue is not true of Ming although it is fragile*. Thus condition (1) in (ST) is violated once again.

Second option: set the threshold at 10N (or above), with the least fragile* objects. The conditional that holds of any object x which qualifies as fragile* is now: if x were struck with at least 10N, then x would break.

This conditional, to be sure, is true of all the fragile* objects. It does not, however, adequately capture their shared fragility*. Take Ming again and strike it with a force of 8N; nothing devious is going on; Ming breaks. Has it thereby manifested its fragility*? Surely we should say that it has. But on the present proposal, we cannot say that: for fragility*, on the present proposal, is the disposition to break if struck with at least 10N, and Ming has not been struck with at least 10N. Given clause (2) in (ST), if Ming is to manifest its fragility*, it must be struck with at least 10N. Indeed, on the present proposal, it is precisely the wrong breakings that qualify as manifestations of fragility*: breaking upon being struck with 10.1N, 11N or 100N. Breakings upon being struck with any force less than 10N are excluded from being a manifestation of fragility*; breakings upon being struck with any force of 10N or more may be included among the manifestations of fragility*. Given clause (2) of (ST), the conditional cannot adequately characterize fragility*.

The first option, then, violated condition (1) by not classifying the right objects as fragile*; the second option violated condition (2) by not classifying the right breakings as manifestations of fragility*.

A third option is to set the threshold somewhere in between. Since we are concerned with Ming, 5N would be a natural candidate for a threshold value. The conditional that would characterizes fragility* in C would be: if x were struck with at least 5N, then x would break.

But rather than solving the problem of the two earlier options, this solution will incur the problem of both. As far as Ming is concerned, the third option gets everything right. But that is not enough. If Ming counts as fragile* in C, then so must a more delicate vase, Meissen, that would break if struck

with at least 1N. (In fact, we can ascribe fragility* to both of them within a single sentence: Ming and Meissen are both fragile*.) But now strike Meissen with 2N; it will break; but its breaking will not count as manifesting its fragility*. Thus condition (2) is violated once again. Moreover, Ming's qualifying as fragile* in C should not exclude a slightly less fragile vase, call it Amphora, from also counting as fragile* in C. (In fact, we can ascribe fragility* to both of them within a single sentence: Ming and Amphora are both fragile*.) But since Amphora is, by stipulation, less fragile than Ming, it will not be true of Amphora that it would break, *ceteris paribus*, if struck with at least 5N; for it would not break if struck with exactly 5N. Hence condition (1) is violated.

Whatever value we fix for n_C , then, (C5) will not yield an adequate characterization of fragility*.

Perhaps (C5) made a mistake in its treatment of the threshold value. After all, what sets the fragile* things apart from the non-fragile* ones is not that they break if struck with greater force, but that they – and only they – would break if struck with slighter force, a force *below* a certain threshold. On this new proposal, fragility* is adequately characterized by the conditional

(C6) x would break if x were struck with a force of *at most* n_C ,

n_C being, again, the threshold value determined by the context C.

But (C6) founders on condition (1) wherever the threshold value is set in context C.

Suppose that n_C is set at 0.1N (or below). It is not true of any, not even the most fragile* objects, that they would break, *ceteris paribus*, if struck with 0.1N or less; so clause (1) is violated. Setting n_C to be at 10N (or above) will not help; the conditional will once again be too easily refuted, by simply striking our fragile* vase Ming with a force of, say, 2N. The vase will have been struck with at most 10N, but it does not break. So the conditional is false of it, against clause (1). And, again, setting the threshold anywhere between 0.1N and 10N will incur the problems of both previous options.

Are there any other options? We have tried using the threshold value as the lower bound of the interval in (C5) (the upper bound being infinity) and as the upper bound of the interval in (C6) (the lower bound being 0). A final option is to let our two main candidates for threshold values – 0.1N and 10N – specify both the lower and upper bound, yielding

(C7) x would break if x were struck with a force between n_C and m_C ,

n_C and m_C being the lower and upper threshold values determined by the context C.

By now it should be easy to see why (C7) does not work. Let $n_C=0.1N$ and $m_C=10N$. Then (C7) violates (1) because it will not be true of any but the most fragile* objects that they would break if struck

with any force between these two values. Setting n_C and m_C at values outside or inside the interval from 0.1N to 10N will not change the situation.

I conclude that the considered strategy fails. Setting a threshold value for the stimulus condition may help with context-sensitivity, but it does not solve the problem at hand: it does not provide us with one conditional that is both true, *ceteris paribus*, of all objects that count as fragile* in a given context, and that captures the right breaking events as manifestations of fragility*. If there is no one such conditional, then fragility* is not a single-track disposition in the sense outlined in section 2; it is a multi-track disposition. Further, the argument relied on no specific features of *being struck* other than its being a determinable quantity. The same is true of other stimulus conditions for fragility, or so I have suggested earlier. The same argument, then, applies to other qualitatively homogeneous ‘tracks’ of fragility – and indeed to any disposition that has a quantitative stimulus condition.

4 What to do with multi-track dispositions

The argument of section 2 showed the inadequacy of a simple conditional analysis of ‘fragile’, which takes the form

(SCA’) x is fragile iff, if x were S , then x would M .

We have seen that there were too many candidates to fill the place of S , and that (SCA’) failed whether we picked only one of them (strategy (A)) or tried to generalize enough to cover them all (strategy (B)). Section 3 showed, further, the inadequacy of a simple conditional analysis of the ‘disposed to’ locution along the lines of

(SCA*) x is disposed to M if S iff, if x were S , then x would M .

(SCA*) will fail, for the reasons outlined in section 3, whenever S is a determinable property that has a range of different determinates. The argument applies not only to fragility and the disposition to break if struck, but to any disposition with a stimulus condition that admits of degrees. I have yet to see an ordinary disposition that does not have such a stimulus condition. Take water-solubility: an object can be immersed in more or less water for a longer or shorter period of time. Or take inflammability, with (one of) its stimulus condition(s): being heated, clearly another quantity. An irascible person may be provoked more or less fiercely. And so on.

Unlike (SCA’), (SCA*) has some instances that escape the arguments I have adduced: those instances, that is, where S is replaced by a fully determinate property. For all that I have said, an analysis

along the lines of (SCA*) may succeed for such very specific dispositions as the disposition to break if struck with exactly 8.35N, call it fragility_{8.35}. If we are looking for single-track dispositions, we need to look here.

Note, again, that finks, masks, and so forth have played no role in my argument; we have assumed that no such factors were present in any of the cases considered. Whatever is done to accommodate such interfering factors – be it Lewis’s reference to a causal basis that is retained throughout the process, Choi’s incorporation of ideal conditions into the stimulus condition, or Steinberg’s *ceteris paribus* clause – will have no force against the arguments of sections 2 and 3.

What is the upshot of all this? That depends on how deep we think the connection is between dispositions and conditionals. The standard assumption in the contemporary literature is that the connection is very deep indeed, though its precise nature depends on other views about the metaphysics of dispositions.

Categoricalists hold that the world, at bottom, does not contain dispositional properties. All the sparse, the perfectly natural properties are categorical. Of course, we use dispositional idioms, and we say true things in using them. The task is to provide truth-conditions for sentences containing dispositional idioms. The simple conditional analysis was a first try, though most of us are now convinced that we will need something more complex.

Typically, categoricalists will take a similar stance to modality in general: the world is, at bottom, not modal. Of course, we use modal idioms such as ‘can’, ‘must’, and the counterfactual conditional, and we say true things with them. The task is to provide truth-conditions for sentences containing modal idioms. Possible-worlds semantics is likely to do the job, though we may disagree about what is the right conception of possible worlds.

The categoricalists that I have described constitute the vast majority of categoricalists, inspired by something like David Lewis’s credo of ‘Humean supervenience’⁹. They aim to analyse disposition ascriptions in terms of counterfactual conditionals, and counterfactuals in terms of possible worlds. Ultimately, what we say about disposition ascriptions is going to be phrased in terms of possible worlds, plus some further ingredients that depend on the precise nature of the analysis that is put forward: the intrinsic (categorical) properties of the objects in question, causation, and/or laws of nature. The deep connection between dispositions and conditionals is a matter of how we get there: via the counterfactual conditional(s) that characterize(s) a given disposition.

⁹See the introduction in D. Lewis, *Philosophical Papers II* (Oxford University Press, 1986); and D. Lewis, ‘Humean Supervenience Debugged’, *Mind* 103 (1994), pp. 473–490.

The result of section 2-3 is that the way via counterfactual conditionals is a very complex one, for any disposition except the very specific ones that we rarely, if ever, ascribe in actual discourse, such as fragility_{8.35}. For a disposition such as fragility, the truth-conditions will be couched in terms of an infinity of conditionals with qualitatively and quantitatively varied antecedents, or ultimately in terms of an infinity of clauses of the form: in the closest possible worlds in which x is subject to stimulus S_i (and some further conditions hold), x yields the manifestation M. Complexity is not a decisive reason to reject any analysis. But note that the complexity of the current analysis is an artefact of its going via conditionals. If a simpler analysis is to be had which achieves the same goal of reducing disposition ascriptions, ultimately, to unobjectionable statements about categorical properties and possible worlds, then in as much as it simpler the alternative analysis would give us reason to prefer it. In section 5, I will offer such an alternative analysis, which reaches the ultimate analysans for disposition ascriptions not via conditionals, but via possibility claims. But before I do so, let us look at an approach to the metaphysics of dispositions that is diametrically opposed to categoricism.

Dispositionalists hold that the world, at bottom, contains dispositional properties. Some, or perhaps all, of the fundamental or perfectly natural properties are dispositional. Dispositionalists may take the failure of the simple conditional analysis, and the problems that beset its less simple successors, as confirmation of the irreducibility of dispositional properties.¹⁰ Alternatively, they may accept that there is a correspondence between dispositions and conditionals but turn around the order of explanation, claiming that the dispositions, so far from being analysed away in terms of conditionals, are what provides for the truth of those conditionals.¹¹ In either case, the dispositionalist too assumes that the connection between dispositions and conditionals is deep indeed; even the failure of a conditional analysis can go towards showing the irreducibility of dispositions only if it is already settled that a reductive analysis would have to be a conditional one. The common core in both versions of dispositionalism is this: the (irreducible) nature of a dispositional property is best or adequately characterized – though non-reductively – by a counterfactual conditional.

For dispositionalism, too, the assumption of a deep connection between dispositions and conditionals combines with the results of sections 2-3 to yield great complexity for ordinary dispositions such as fragility. If the nature of fragility is best or adequately characterized by counterfactuals, then it will be infinitely complex, for it requires an infinity of conditionals. The nature of fragility_{8.35}, on the other hand, is as simple as a disposition can possibly be. In fact, it looks as though fragility is in some way

¹⁰Martin, 'Dispositions and Conditionals'; Molnar, *Powers*.

¹¹A. Bird, *Natures Metaphysics. Laws and Properties* (Oxford University Press, 2007); J. Jacobs, 'A powers theory of modality: or, how I learned to stop worrying and reject possible worlds', *Philosophical Studies* 151 (2010), pp. 227–248.

built up from such single-track dispositions, incorporating their characterizing conditionals within its complex nature. The simple building-blocks are, in whatever area, more fundamental than their complex compounds; hence fragility_{8,35} is more fundamental than fragility.

An *austere dispositionalist* may draw from this the conclusion that fragility is no disposition at all. We need to distinguish between dispositional properties and dispositional predicates, or so it may be said. As far as dispositional properties are concerned, we need accept no more than the simple, single-track ones. Dispositional predicates, such as ‘fragile’, are merely umbrella terms that we use to ascribe a bunch of them.¹² A more *liberal dispositionalist* will conclude instead that there is a priority ordering among dispositions, and that single-track dispositions are prior to, or more natural than, the more familiar multi-track dispositions. At bottom the explanatory work is done by an infinity of simple, single-track dispositions.¹³

Whichever way we go, our ordinary disposition ascriptions, if true, are underwritten ultimately by an infinity of dispositional properties. For the dispositionalist metaphysics, as for the categoricist semantics, an ordinary disposition such as fragility becomes a beast of infinitely many heads. Note that the multiplicity of dispositional properties cannot simply be subsumed under a familiar schema, such as that of a determinable and its determinates. Different determinates of a single determinable are mutually exclusive; nothing can be both crimson and scarlet (all over, at the same time, and so on). But the different single-track dispositions corresponding to a multi-track disposition can be, and often are, co-instantiated. We must accept the multitude of dispositions as a phenomenon in its own right.

Again, complexity is not a decisive reason to reject a view. But it is a comparative drawback, if there is a simpler alternative.

And indeed there is. Note, again, that the key premise in producing the complexity both in dispositionalist metaphysics and in categoricist semantics was the idea that there is a deep connection between dispositions and conditionals. If there was no such connection, then the categoricist semantics of disposition ascriptions would not be required to take the long and winded road via an infinity of conditionals. Nor would the dispositionalist have to take the multitude of conditionals that are needed to characterize a disposition as proof of the property’s complexity. From the mere fact that a property can be given a complex, conjunctive or disjunctive, characterization we cannot infer that the property itself is accordingly complex. The property of being green can be characterized as a disjunction of being grue and observed before a time *t* or being bleen and observed thereafter. This does not make *being green* a

¹²Brian Ellis seems to be an austere dispositionalist: see his *Scientific Essentialism* (Cambridge University Press, 2001).

¹³Alexander Bird is a liberal dispositionalist: see *Nature’s Metaphysics*.

more complex property than *being grue*, because we have no reason to believe that the disjunctive characterization is what gives the nature of *being green*. Similarly, the fact that fragility can be characterized by a multitude of conditionals is evidence for its complexity *only if* that characterization is what gives the nature of fragility. That it does, follows from the assumption that the nature of a disposition is quite generally to be characterized by conditionals. If we reject that assumption, we thus become free to give a more unified characterization of ordinary dispositions such as fragility.

5 An alternative approach

How are we to characterize dispositions, if not in terms of conditionals? This final section will make a suggestion, and compare it favourably to the standard approach.

We may begin by taking a hint from lexicography.

Fragile things, we are told in the *Oxford English Dictionary*, are those that are ‘liable to break or be broken; ... easily destroyed’. ‘Soluble’ means ‘capable of being melted or dissolved’, and to be ‘inflammable’ is to be ‘[c]apable of being inflamed or set on fire’. Two points are striking about these dictionary definitions: first, they cite only the disposition’s manifestation, no stimulus conditions; second, and relatedly, they seem to relate the disposition terms to claims of possibility of capability, rather than to conditionals of any kind.

Dictionary definitions, of course, cannot be a guide to truth in metaphysics. But the definitions I have cited (and the many others like them) suggest a conception of dispositionality that provides the simplicity and unity which, I have argued, is lacking when we subscribe to a more standard, conditional-based, conception of dispositions. What unifies fragile objects is the simple fact that they *can* break *easily*. It is this that underwrites the complexity of the infinitely many conditionals that we have seen to be associated with fragility. In general, what characterizes the nature of a disposition, on the alternative conception, is primarily what the objects so disposed *can* do and not, or only derivatively, what they *would* do under specific circumstances.

The alternative approach will thus replace (SCA) and its descendants, as a first step towards understanding dispositions, with something along the lines of a *possibility analysis*:

(PA) x is fragile iff x can break (easily).

In other words, the alternative approach entirely disposes of the deep connection between dispositions and conditionals, and substitutes for it a deep connection between dispositions and (easy) possibil-

ity. Exactly how that deep connection is spelled out will depend, again, on the more general metaphysical framework.

If we are *categoricalists*, then our main task is to provide a semantics of disposition ascription. In that framework, the alternative approach amounts to the claim that the truth-conditions of disposition ascriptions are to be formulated not in terms of a counterfactual conditional, but in terms of possibility. Simple possibility will not, perhaps, be enough: not anything that *can* break is fragile. The dictionary definition suggests, instead, that we think of the relevant possibility as graded: different things break more or less easily, and the more fragile an object is, the more easily it can break. For an object to count as fragile in a given context, it has to be sufficiently fragile. In the suggested framework, this amounts to the requirement that it can break sufficiently easily. In further developing the alternative approach, the task for the categoricalist is, first, to provide a better understanding of the graded possibility that is at issue; and second, to show how the proposal accommodates masked and mimicked dispositions.

There are two main models for understanding graded possibility. One is *closeness*: it is *more possible* that *p* than it is that *q* just in case *p* holds in a world that is closer to some contextually determined ideal than any world in which *q* holds.¹⁴ The second option is *proportion*: it is *more possible* that *p* than it is that *q* just in case there are more (relevant) worlds where *p* holds than there are (relevant) worlds where *q* holds.¹⁵

Whichever way we go, the alternative conception of dispositions as it stands is subject to masking and mimicking. A vase that stands in front of a bulldozer can break more easily than one which is wrapped in styrofoam, yet both vases may be equally fragile. Like (SCA), (PA) requires modification. The first steps of such modification are not hard to come by: the relevant worlds must hold fixed the intrinsic features of the object in question at the time of the disposition ascription, while being maximally varied with respect to features that are extrinsic to the object. Thus we look at worlds where the vase retains its intrinsic basis for fragility, but give no preference to bulldozer-worlds over styrofoam-worlds or vice versa.

A *dispositionalist* will have a different take on the alternative approach. She will insist, rather, that the right ‘can’ statements characterize the dispositional essence of a given property. She may continue by saying that the dispositional property provides for the truth of the relevant ‘can’ statement: the glass

¹⁴For a useful discussion of the various types of ideals, or ‘ordering sources’, see A. Kratzer, ‘The Notional Category of Modality’, In H. J. Eikmeyer and H. Rieser (eds.), *Words, Worlds, and Contexts. New Approaches in Word Semantics* (Berlin/New York: de Gruyter, 1981), pp. 38-74; now reprinted and updated in A. Kratzer, *Modals and Conditionals: New and Revised Perspectives* (Oxford University Press, 2012).

¹⁵Manley and Wasserman (‘On Linking Dispositions and Conditionals’) support a version of this model for dispositions. They also discuss the difficult technical question how such proportions are to be construed (pp. 79–82).

can break easily *because* it is fragile, and not vice versa. Or she may insist that ‘can’ statements are no more than imperfect approximations to the corresponding disposition.

Whichever of these attitudes the dispositional realists adopts, the alternative approach that I have proposed has the consequence that a disposition is individuated by its manifestation alone. Stimulus conditions play no part in individuating, or in giving the essence of, a disposition. A plausible realist picture of dispositions will start with a general conception of powers or capabilities to do or to be such-and-such. It will hold, further, that these powers come in degrees. For a disposition ascription such as ‘Ming is fragile’ to be true, Ming’s power will have to have a certain, contextually fixed, minimum degree. Powers and their degrees will remain primitives on such a view. But that is as it should be. After all, the project is designed to be anti-reductionist.

This is not the place to develop the alternative conception of dispositionality in more detail or defend it against objections¹⁶, but I will argue now that it is preferable to the standard, conditional-based conception because of the results of sections 2-4. To see why, we must note first that the alternative approach escapes the kind of argument that I have put forward in sections 2-3.

First, there is in general more qualitative unity to a disposition’s manifestation than there is to its putative stimulus conditions. The fragility of a glass, a vase, a wooden chair or an old parchment is manifested in *breaking* even if these manifestations are triggered by such diverse factors as the glass being hit by a spoon, the vase being sung to by a soprano, the chair being transferred into a different temperature and the parchment being merely touched. Of course, the manifestations, just like the putative stimulus conditions, will typically exhibit the quantitative diversity of determinable properties in general: things may break into two, three or any number of pieces, their breaking may be slower or quicker, and so forth.

But this is no problem, for *second*, possibility claims interact with disjunctions rather differently from counterfactual conditionals. The problems sketched of sections 2 and 3 arose largely from a structural feature of counterfactual conditionals: the universal quantifier that is, on the standard semantics, implicit in them. For it to be true that x would M if it were S, x has to M in *all* the closest worlds where it is S. If being S is a disjunctive property (such as being stressed in section 2) or a determinable property (such as being struck in section 3), those closest worlds may include worlds where any or all of the disjuncts or determinables of S are true of the object x. But, I argued above, that is too much to ask for a property such as fragility. The same problem does not arise with a possibility-based conception of dispositionality simply because there is no universal quantifier implicit in a possibility claim. According

¹⁶I do so elsewhere: B. Vetter, *Dispositions without Conditionals*, unpublished manuscript.

to that conception, when we say that *x* is fragile, we say that it breaks in *some* (contextually relevant) world or worlds. This is equally true of an object which breaks into two pieces in some such world, of an object that breaks into seven pieces in some such world, and of an object that breaks into a thousand pieces in some such world.

Adopting (PA) as a starting point in understanding dispositions, then, avoids the argument of sections 2–3, and thereby the consequences of that argument as drawn out in section 4. I concluded section 4 rather cautiously, stating that the great complexity engendered by the assumption of a deep connection with conditionals was a comparative drawback. Now that the simpler alternative is on the table, I want to strengthen that conclusion.

In comparing two types of approach to a phenomenon – in this case, the conditional and the possibility conception of dispositions – there are various criteria that must be weighed against each other. One is theoretical elegance or simplicity. On this account, as we have seen, the possibility conception does better than the conditional conception. But there are other criteria. Theoretical simplicity is sometimes in tension with a different virtue, which we may call *intuitiveness*. Intuitiveness is a matter of closeness to our natural understanding of a phenomenon, to our intuitions and ordinary dealings with it. How do the two conceptions compare with respect to intuitiveness?

It may be said that the conditional conception is the most intuitive one; that there is a pretheoretic link between dispositions and conditionals. However, a pretheoretic link may be claimed, with at least equal right, between dispositions and easy possibility. This latter link is confirmed by lexicography and linguistics. (Indeed, in her seminal work on the semantics of modal language, linguist Angelika Kratzer, without further ado, treats dispositional adjectives as expressing possibility.¹⁷ The former link is confirmed chiefly by philosophers' introspection. In the absence of further evidence, I am inclined to trust lexicography and mistrust the philosophers' introspection: the latter is likely to be informed by philosophical theory already. In any case, no decisive reason in favour of the conditional conception has emerged.

In fact, it seems to be the conditional conception that, in preferring such contrived dispositions as fragility_{8.35} over ordinary dispositions such as fragility, has moved away from our natural understanding of dispositions. This preference takes different forms. It is weakest with the categoricist, for whom ascriptions of the former merely have simpler truth-conditions. It is stronger with the liberal dispositionalist, for whom the ordinary dispositions are grounded in the contrived ones. And it is strongest in the austere dispositionalist, who gives up entirely on the ordinary dispositions. Moving away from our nat-

¹⁷Kratzer, 'The Notional Category of Modality'.

ural understanding of a phenomenon is a cost, albeit one that may be superseded. Theoretical simplicity may outweigh it; but we have seen theoretical simplicity to be on the side of the possibility conception too.

Science often prompts us to move away from our intuitive conception of a phenomenon. Is the conditional conception, perhaps, in better accord with the scientific image of the world? I doubt that it is.

Take the dispositionalist view first. The conditional conception has prompted it to accord a greater degree of naturalness to the highly specific, single-track dispositions, than to the more familiar multi-track ones. (The austere dispositionalist then rejects the less natural dispositions.) This would lead us to expect that the perfectly natural dispositions are highly specific too. At the fundamental level, then, we do not find properties such as charge (a general disposition to repel like-charged objects and attract differently charged ones with a force proportional to their charge and distance), but such highly specific disposition as the disposition to repel objects of charge e^- at a distance of exactly $5.3 \times 10^{-11}m$. But there are no independent reasons, apart from the conditional conception, to form such an expectation. We tend to think that the more natural properties figure in the more fundamental laws; and that the laws and properties dealt with in physics are more natural than those dealt with, say, in sociology. But physics is better understood to be dealing with charge and other equally general dispositions, than with the highly specific disposition to repel objects of charge e^- at a distance of exactly $5.3 \times 10^{-11}m$. It is the general dispositions, not the highly specific ones, that correspond to the laws of physics. Coulomb's law states that for *any* charges Q and q and distance r , the force exerted by a particle of charge Q when at distance r from a particle with charge q is $F = \epsilon \frac{Qq}{r^2}$. Instances of this law – concerning, for instance, the force exerted by a particle of charge e^- when at exactly $5.3 \times 10^{-11}m$ from another particle of charge e^- – follow from the general law and should not be prior to it, on pain of raising the question why there is such mathematical uniformity among them. Accordingly, charge itself has a better claim to be a (nearly) perfectly natural disposition than the specific dispositions described. Physics, then, warrants no preference for the specific dispositions over the more general ones. If anything, it advises an ordering that takes such general dispositions as charge to be prior to the highly specific dispositions, and by extension treats fragility as prior to fragility_{8.35}. (Some dispositionalists – the dispositional essentialists¹⁸ – have appealed to the conditional conception to explain laws of nature in terms of dispositions. For reasons

¹⁸Ellis, *Scientific Essentialism*; Bird, *Nature's Metaphysics*.

that cannot be given here, I believe that conditionals do not serve the dispositional essentialist's purposes either.¹⁹

For the categoricist, appeal to the scientific image has little force in the present context. Physics does not deal with dispositions, and the reductive task is not committed to a naturalness ordering between dispositions. The main consideration, for the categoricist project, remains theoretical simplicity.

I have considered three criteria that may plausibly be invoked to choose between the standard conditional conception of disposition, which starts with (SCA), and the alternative, possibility-based approach that takes (PA) as its starting point. Each criterion, though each in a different way, provides defeasible support to the alternative approach rather than the standard one. Taken together, they provide excellent reason to set aside the conditional conception, and to explore the alternative conception based on possibility.²⁰

¹⁹See B. Vetter, 'Dispositional Essentialism and the Laws of Nature', in A. Bird, B. Ellis, and H. Sankey (eds.), *Properties, Powers and Structures. Issues in the Metaphysics of Realism* (Routledge, 2012).

²⁰I would like to thank Romy Jaster, Beate Krickel, and Thomas Kroedel for their comments on a draft of this paper, and Simona Aimar, Natalja Deng, Antony Eagle, Daan Evers, John Hawthorne, Timothy Williamson, and Alastair Wilson for helpful conversations at earlier stages.