Author information

Title	Dr. phil.
Name	Jens Gillessen
Institutional affiliation	Universität Leipzig
Contact address	Universität Leipzig
	Institut für Philosophie
	Beethovenstr. 15
	04107 Leipzig
	Germany
e-mail	jens.gillessen@uni-leipzig.de

REASONING WITH UNCONDITIONAL INTENTION

Jens Gillessen

Abstract.
Suppose that you intend to go to the theater. Are you therein intending the unconditional proposition that you go to the theater? That would seem to be deeply irrational; after all, you surely do not intend to go if, for instance, in the next instant an earthquake is going to devastate the city. What we intend we do not intend 'no matter what', it is often said. But if so—how can anyone ever rationally intend simply to perform an action of a certain kind? In response to the puzzle, a 'conditionality' view of intention has emerged: The contents of everyday intentions are claimed to be fraught with hidden conditional clauses. The paper argues that such claims are radically unmotivated: Even unconditional intentions have only limited inferential import and hence contrast sharply with a 'no matter what' stance. The point is established by examining relevant patterns of reasoning from unconditional to conditional intentions.

I.   A puzzle about flat intention

Recent philosophical accounts of intention tend to ascribe to speakers internally conditional attitudes on the basis of even the simplest reports of intention. Suppose someone—call him Schmidt—states, in response to a colleague's question what he is up to in the evening:

(1)   I intend to go to the theater at 7 p.m.

Being unconditional in grammar, the statement surely invites ascription to Schmidt of an intention that he himself goes to the theater at 7 p.m., period; the content not involving any hidden conditional clauses. It thus seems most natural to ascribe to Schmidt an intention relating him to the proposition that he himself will go to the theater at 7; which 'flat' intention would be unequivocally captured by:

(2)   Schmidt intends to (go at 7)

Lately, the number of those who balk at 'flat', face-value accounts of reports of intention as well as desire have been on the rise. In the case at hand, they would argue that (2) misdescribes Schmidt as a kind of fanatic with a crazy intention: an intention to go to the theater at 7 'come what may', or 'no matter what'. And for sure, his assertion of (1) cannot be taken to entail:

(3)   Schmidt intends to (go at 7 if by then, an earthquake will have devastated the city).

If adopting flat intention (2) meant to commit to ludicrous intentions such as (3), taking statement (1) at face value would be a folly. Still, it is quite hard to see what a flat intention to A could be other than an intention *to A in any circumstances.*

The proponents of what I would like to call the *Conditionality View* have concluded that our everyday intentions cannot be of the flat sort that (2) exemplifies. Instead, they have contended, the contents of ordinary intentions must be systematically and deeply conditional.[1] What the present paper is going to show is that this contention is unmotivated, for a simple reason: A flat intention to A is not an intention to A 'no matter what'. My argument is straightforward.

1. If intending that (A) were intending that A 'no matter what', it would have to be rationally permissible to base on one's intention that (A) an intention that (A if C), *for arbitrary choice of circumstances C.[2]*
2. There is no such permission.
3. Therefore, intending that (A) is not intending that A 'no matter what'.

My defense of the first premise will be brief. On the broadly inferentialist picture I will be assuming, what it amounts to to bear a certain attitude is reflected by the inferences the attitude bearer is entitled to draw by virtue of her having that very attitude (against the background of her other attitudes). Thus, the state of believing that p is partially, but nonetheless essentially characterized by a (conditional) normative fact: It entitles its bearer to form the belief that q, conditional on her simultaneously believing that if p then q. Whatever state does *not* support such modus ponens reasoning cannot be counted as a belief that p. Analogously, I am suggesting, an intention to go at 7 cannot be counted as an

intention to do so 'no matter what' unless it rationally entitles its bearer to adopt conditional intentions such as, say, an intention to go at 7 *if* a quake will have devastated the city by then.[3]

That is all I have to say about the first premise. Vindicating the second premise will take a lot more reflection, and this is how I shall go about it: After a clarification of the concept of conditional intention and some preliminaries of reasoning toward them (section II), I will first point out why reasoning toward conditional intentions is particularly difficult to assess. For one we are dealing with *practical* reasoning, which calls for an apt standard of assessment (section III). For another (section IV:), the inference pattern in question elicits intuitions that call for *some* acknowledgment, however limited, of inferences from intentions such as (2) to intentions such as (3). In section V I am going to argue that reasoning *without further premises* from an unconditional intention to one of its conditional counterparts—in accordance with a rule I label (SCON)—is proven incorrect by striking counter-examples. This point will be established by critical examination of a number of specious responses to those counter-examples. Section VI then reveals a *correct* pattern of reasoning from unconditional to conditional intentions. Reflection on the necessary complexities of the pertaining inference rule will make it clear, however, that there is a world of difference between flat and fanatic intentions. At that point then, the truth of premise 2 will become obvious.

II.   Conditional Intention and Conditionalization

First of all, let me clarify what will be meant by "conditional intention". Virtually every intention is conditional in the sense that its being formed, entertained or abandoned depends on environmental, neural, mental and other conditions external to the intention itself. These need to be distinguished from the *internal* conditions with which I will be concerned. The point needs emphasizing because locutions of the form 'I intend that A if C'[4] admit of both an external and an internal reading. On the external reading, the speaker merely predicts that if C obtains, or is going to obtain, she will form an intention that A. On the internal reading, she expresses an intention that A if C which at the time of utterance she has already formed—an internally conditional intention.[5] In order to disambiguate, the latter will be noted down in the form 'N intends that (A if C)'. This notation is meant to be compatible with, but not to imply the contentious view that the agent therein intends the *proposition that* (A if C).[6] There may be no such things as conditional propositions, and my arguments will not presuppose a particular stance on the matter. What they will presuppose is the robustness of the distinction of internal and external conditions. Among internally conditional intentions, my focus will be on intentions with internal *sufficient* conditions exclusively.[7]

By *conditionalization* I shall mean a subject's passage by inference from a state of intending that A to a state of intending that (A if C). Sometimes, the formation of a conditional intention will not be inferential at all. Scared by long and deep cracks in the ceiling above you, you might decide to rush to a particular emergency exit as fast as you can in case they should start to widen. Obviously, deciding so makes sense in the absence of any intention to rush that could serve as a premise. On other occasions, however, it seems as though you will not rationally come to intend that (A if C) unless you already intend that A. On these occasions, you reason, or could reason, toward a conditional intention.

A good means to focus on these cases is to imagine playing the *conditionalization game*. The game involves two players, and the rules are as follows. Player 1 reports one of

her intentions, in the form 'I intend to A' or 'I intend that A'. Next, player 2 asks player 1 a question of the form 'Do you intend that A if C?', picking up the A-term while choosing C totally arbitrarily. Player 1 is then to give a sincere yes-no response. It will occasionally be helpful to recall this setup.

In dealing with conditionalization it needs to be acknowledged that there are two well-investigated correct sorts of reasoning which at least sometimes lead to conditional intentions as conclusions:

- *Instrumental Reasoning:* You adopt a conditional intention because you believe doing so to be a suitable means to achieve one of your goals.[8]
- *Enkratic reasoning:* You believe you ought to A if C, and therefore adopt an intention to (A if C).[9]

Almost needless to say, neither of these paths could be exploited to reason from an intention such as (2) to a crazy intention, such as (3). Were intention (3), or its adoption, recommended by either Schmidt's reasons or as a means to his ends, it would not count as relevantly crazy after all. But this does of course not prove the Conditionality camp wrong; one can conceive of further, more straightforward ways of conditionalizing. Let me summarize these under the label:

- *Genuine Conditionalization:* You adopt an intention to (A if C) because you intend that A without availing of either instrumental or enkratic reasoning.

The shortest and most natural-seeming route of Genuine Conditionalization would be *Simple Conditionalization*:

(SCON) N intends that (A)
∴ N intends that (A if C)

What we need to ask then in order to decide on premise 2 from above is whether *correct* reasoning along (SCON), or some relevantly similar such path, is possible.

III.   Assessing Practical Reasoning for Correctness

The task of judging schemes of inference such as (SCON) is complicated by the fact that we are dealing with practical reasoning. Assessing *theoretical* inferences can seem comparatively unproblematic: Truth-preservation in passing from premises to conclusions is commonly taken to be a reliable guide; at least when operating under idealizing abstraction from the limits of reasoners' cognitive resources. In the realm of practical inference, however, the lens of truth-preservation turns blind. We need an alternative standard of assessment. In explaining mine, I will draw heavily on John Broome's theory of active reasoning (Broome 2013: 177-91, 221-87).

Reasoning, I suggest in agreement with Broome, is a causal process by which a subject comes to have an attitude X because of other attitudes $Y_1, …, Y_n$, where this process satisfies at least one further condition: The conclusion attitude must be constructed from the premise attitudes by application of a rule. A particular piece of reasoning is correct just in case the therein applied rule is correct. Correctness of a rule, I suggest, should in turn be assessed by means of the following criterion:[10]

(COR) An inference rule of the form 'to derive attitude X as conclusion from attitudes Y_1, ..., Y_n as premises'[11] is correct iff it is rationally permissible to base X on $\{Y_1, ..., Y_n\}$.

On the account I am suggesting, to reason is to actively base attitudes on other attitudes. Thus, (COR) identifies correctness of reasoning with its rational permissibility. At the same time the criterion offers a deeper understanding of what reasoning is—at least in those cases in which reasoning leads one to adopt a new attitude: To *base* an attitude X as a conclusion on attitudes Y_1, ..., Y_n as premises is to construct X by operating on Y_1, ..., Y_n following a rule, therein adopting X.[12]

How do we know what rationality permits? Broome thinks that basing permissions must, in the last instance, derive from pertaining requirements of rationality (Broome 2013: 248, 256-59). Whether this is so need not be answered here. I take it that we have (fallible, but nonetheless) immediate epistemic access to what rationality permits when we carefully study exemplary instances of reasoning and therein come to have intuitions about their (ir-)rationality.

IV.   A Conflict of Intuitions

Following intuitions blindly, however, is out of the question. As in other philosophical domains, too, a reflective equilibrium needs to be established between considered intuitive judgments and a theory that explains them (Harman 1986: 9). The importance of such cautioning becomes evident when considering the confusingly wide range of intuitions (SCON) can evoke.

To the adherents of the Conditionality View, the pattern has apparently seemed almost trivially *correct*. As will be seen, there is quite a number of ways in which one could try to reconcile this favorable judgment with the apparent oddity of conclusions such as (3). Sometimes recalcitrant intuitions can be defused by tracing them to systematic biases of judgment or cognitive illusions. One such error theory[13] could claim, for instance, that in judging (3) a crazy state, what actually strikes one as crazy is not the *internal* reading:

(3-i)   Schmidt intends to (go at 7 if a quake occurs)

    ... but rather the *external* reading:

(3-e)   If a quake occurs, then Schmidt intends to (go at 7)

The impression of craziness being explained by the salience of reading (3-e), nothing seems to stand in the way of judging (3-i) a reasonable intention to adopt.

Against this suggestion, two rejoinders are in place. First, it is not at all clear that the planning state attributed to Schmidt by (3-i) is any less crazy than the disposition attributed to him by (3-e). In order to properly isolate intuitions about (3-i), consider a brief story. Suppose Schmidt in all seriousness deliberated as follows: "I'm going to the theater at 7 ... There could be a devastating quake tonight, though ... Looters would be all over the place... Might be better to guard the house then ... Anyhow, even if a quake will strike—I'll go!" Suppose further that as a result of so deliberating he acquires a number of dispositions that can be considered indicative of (3-i):

• a disposition to reply with "yes" to the question whether he will go to the theater in case of a quake;

- a disposition not to form any plans incompatible with going to the theater in case of a quake; and
- a disposition not to deliberate further about the matter.

Finally, let us suppose that he *fails* to acquire disposition (3-e): Around 7, the earth begins to shake violently, and in horror Schmidt forthwith lets go of his intention.

In the story, (3-i) holds true while explicit denial of (3-e) should prevent the alleged disturbance of one's judgment. Still, Schmidt's intention seems crazy—contrary to what the error theory predicts.

Second, advancing an error theory makes sense only if there is something to be said in favor of (SCON) in the first place. Now, as far as I can see intuitive support for (SCON) could originate from two sources.

1.) Friendly-looking instances of (SCON).
2.) Specious similarity of (SCON) with a certain correct path of reasoning toward conditional *belief.*

Let me comment on these one at a time. A correct path of belief reasoning is certainly:

(BC⊃)   I believe that (P)
   ∴ I believe that (C⊃P)

The pattern is rationally unassailable; after all, the content of the premise entails the content of the conclusion, and belief is 'closed under entailment'.[14] It is natural to conjecture then that analogous reasoning with intention is likewise correct:

(IC⊃)   I intend that (P)
   ∴ I intend that (C⊃P)

Yet, the analogy is anything but perfect. The correctness of (BC⊃) is obvious for the sole reason that belief is governed by a 'closure under entailment' norm of ideal rationality. No such norm governs intention, though, as can be seen from a counterexample: A rich philanthropist could fully rationally intend that (he exterminates poverty in the world) without intending that (there is poverty)—even though the former content clearly entails the latter. The absence of an underlying closure principle for intention does not prove (IC⊃) incorrect, to be sure; the scheme could still be correct in its own right. Even so, the absence of any deeper justification and the revealed disanalogy between intention and belief are bound to shatter much of its initial appeal.

Real support for (SCON), I take it, is provided by friendly-looking instances of the pattern. If Schmidt were asked: "Do you intend to go to the theater at 7 if the sun will already have set?", or "… if 2 plus 2 equals 4?", or "… if Florence is in Tuscany?", he might each time *affirm*, and this would not strike anyone as odd behavior. Conditionalizing toward circumstances in which the intended action would be reasonable generally looks harmless. This counts in favor of (SCON).

Still, these harmless instances contrast sharply with crazy-seeming moves such as the inference from (2) to (3). What we are facing is a *conflict* of intuitions, where both kinds of intuition will have to be accounted for. Acknowledgement of (SCON) without a working strategy to deal with bad-seeming moves will not do. Rejecting conditionalization altogether, by contrast, would account for the bad-seeming instances, to be sure. But this solution would come at the price of forgoing a quite appealing explanation of the intuition

that many, if not most instances of conditionalizing seem to be almost trivially permissible. Eventually I shall argue that (SCON) is in need of severe regimentation, but beforehand the rule cannot be considered either plainly correct or plainly fallacious.

It might be objected that the price of a clear-cut rejection of (SCON) may not be that high after all. Could the good-seeming instances of conditionalization on closer inspection not turn out to be instrumental or enkratic inferences in disguise? After all, if you intend to A, you will usually so intend either because you believe that so intending will promote your goals, or because you believe you ought to A. And if the latter, it seems as though the very reasons that support your ought-belief will also support an intention to (A if C); at least when C does not undermine those reasons.

The objection fails because there are seemingly good conditionalizations the permissibility of which is explained neither by instrumental nor by enkratic reasoning. To see this, it suffices to attend to constellations reminding of Buridan's Ass, in which the agent intends something she has no specific reason to intend, and which she also does not believe to be specifically conducive to any other goal of hers.

Suppose a saleswoman—call her Meyer—intends to go by car from Palo Alto to San Francisco. She is already on the road, and before reaching a certain intersection she will have to choose whether to go by Highway 101 or by Interstate 280. Being unfamiliar with the area, she does not have any beliefs about the two routes that would justify favoring one over the other. After a while, she plumps for 101. Another minute later, still traveling toward the intersection, she picks up a hitchhiker. On learning that Meyer intends to go by 101, the hitchhiker asks her whether she intends to go by 101 even *if it is going to rain*. Meyer does not mind driving in rain, and if it will rain on 101, she believes, it will be raining on 280 either. She replies with "yes"—and rationally so, as it seems.

Let us take stock. In the example, Meyer exhibits, at some point or other, two intentions:

(M1)   Meyer intends to (take 101)
(M2)   Meyer intends to (take 101 if it will be raining)

At no point does she have any ought-belief from which she could possibly have inferred (M2) enkratically. Arguably, she intends to (go to San Francisco) because she believes she ought to go to San Francisco; also, she believes that she ought to take *either* 101 *or* 280. But none of these beliefs is suitable for obtaining (M1) by way of enkratic reasoning— let alone for enkratically obtaining (M2). In the absence of any reason to take 101 in particular, she cannot rationally come to believe she ought to (take 101)—let alone to (take 101 if it will be raining).

However Meyer has arrived at (M2)—she cannot have reasoned her way enkratically. Nor could instrumental reasoning have yielded (M2); after all, no goal of hers would be served by adopting (M2) over and above (M1). Whether it will rain or not, (M1) suffices to take her to San Francisco. And still, she would not have adopted (M2), had she not adopted (M1) in the first place. There must be a third correct path of reasoning toward conditional intentions, apart from instrumental and enkratic reasoning then; and for all that has been said up to this point, this path could well be *Simple Conditionalization.*

V.   Simple Conditionalization in spite of malign instances?

If (SCON) as it stands were correct, both benign and malign instances of conditionalization would have to be good reasoning. Not only would Schmidt be permitted to reason

from (2) to an intention to go to the theater at 7 if by then the sun will have set; he would even be able to reason correctly toward crazy intention (3). Hence the proposal immediately runs into trouble.

That alone barely justifies its abandonment, however, because the force of the counter-examples is not entirely beyond doubt. In defending (SCON) against the allegation of licensing crazy inferences, two general strategies can be pursued.

1.) It might be held that bad-seeming conditionalizations, although good reasoning in themselves, in reality *never occur* (maybe with some very special exceptions) because despite appearances, agents generally do not have intentions as simple as (2) from which they could reason toward crazy intentions. This defense has at least implicitly been endeavored by Luca Ferrero, the main proponent of the Conditionality View.

2.) On the contrary assumption that agents do have intentions as simple as (2), appearances could still deceive in another way. In seemingly crazy moves such as from (2) to (3) the craziness of the conclusion can be (and has in fact been) credited to violation of rationality principles other than (those arguably underlying) Genuine Conditionalization. The candidates I shall discuss are intention consistency, the instrumental principle, Enkrasia, and a principle about the connection between intending and believing.

A. The Defense given by the Conditionality View

The Conditionality View of intention exploits (SCON) to argue that intentions are crazy states *unless* their content is regimented by the presence of strong implicit conditional clauses. On one simplistic, but instructive version of the view, the deep structure of (2) amounts to:[15]

 (4)   Schmidt intends that (he goes at 7 if he will then find it advisable to do so)

Call an intention with a built-in highly general evaluative clause a 'deep' intention. Now, conditionalizing toward the earthquake circumstance, such as in the step from (2) to (3), would lead to:

(5)   Schmidt intends that ((he goes at 7 if he will then find it advisable to do so) if by 7, an earthquake will have devastated the city)

It is not hard to see how the advisability clause saves (5) from craziness. On the assumption that the earthquake really will take place, Schmidt will at 7 hardly find it advisable to go; the innermost conditional will turn out vacuously true, and the intention will be satisfied whether or not Schmidt goes to the theater.

Adherents of the Conditionality View readily admit that (SCON), when applied to flat intentions, does support moves toward crazy conclusions; the inference from (2) to (3) is claimed (or at any rate implied) to be *correct*. Yet, since no sane person has intentions as flat as (2) to reason from,[16] examples of crazy inferences are claimed to be *unrealistic;* whereas the results of applying (SCON) to realistic premises are claimed to be sane thanks to their built-in advisability condition.

The Conditionality View thus offers a way to hold on to (SCON) even in the face of seemingly bad instances, but at the same time it raises a bunch of new issues. Let me confine myself to a problem about akratic intention. Are akratic intentions supposed to be deep or flat? Suppose George, a reckless undergraduate, against his own better judgment gives in to an urge of curiosity and decides to take heroin at 9 p.m. when he will be back in his room. On the deep construal, George has therein formed:

(6)   George intends that (he takes heroin at 9 if he will at 9 find it advisable to do so)

George does not *now* find it advisable to take heroin; that is what makes his intention akratic, if anything. Unless in the course of the day he at some point comes to find drug abuse advisable, (6) will at 9 fail to initiate drug-taking on his behalf; the intention will not be productive of akratic action. Yet this is what akratic intention often does. So akratic intention must be flat:

(7)   George intends that (he takes heroin at 9)

But if the Conditionality View must admit that akratic intention is flat, it cannot solve the problem of bad-seeming conditionalizations *for akratic intentions.* Just consider the inference from (7) to:

(8)   George intends that (he takes heroin at 9 if his bourgeois parents will then be giving him a surprise party in his own room)

Intending to take heroin while being watched by his parents will be a crazy thing for George to intend, even by his own lights (we can assume). The move from (7) to (8) would be a crazy instance the Conditionality View cannot accommodate. It will at any rate not do to claim that intentions such as (7) in reality never occur. Denying the possibility of akrasia seems too high a price for holding on to (SCON). And notice that, if (SCON) is incorrect, so are most likely the inferential moves which are supposed to provide the initial motivation for assuming the Conditionality View—moves such as from (2) to (3).

In some respects, I have been simplifying. Ferrero, who has put forward the best elaborated version of the view yet, actually suggests a more complicated 'deep structure' than displayed by (4).[17] Yet, even on his elaborate account it is hard to see how akratic intention could be dealt with satisfactorily. I suggest to explore other defenses of (SCON).

B.   The Inconsistency Defense

Here is the idea of a seemingly simple defense: Whenever (SCON) leads from an intention to a crazy intention, the craziness of the progression is claimed to be due not to (SCON) but to an alleged inconsistency with one of the agent's other intentions. More precisely, the Inconsistency Defense can be put forward in either of two forms.

(ID1)   The agent may be claimed to *have a background intention* that is inconsistent with intending to (A if C). If an agent has a 'blocking intention', as I shall say, this will explain sufficiently why she cannot fully rationally adopt the conclusion of (SCON).

(ID2)   Alternatively, it could be claimed that agents spontaneously *form* suitable blocking intentions whenever confronted with stimuli such as, say, the sort of questions asked by player 2 as part of the conditionalization game.

For these to be viable defenses of (SCON), we need to be able to plausibly ascribe suitable blocking intentions. Which intention of Schmidt's, for instance, could possibly block an inference from (2) to (3)? These candidates may come to mind:

(9)   Schmidt intends (never to end up in trouble)
(10)   Schmidt intends (not to go if a calamity occurs)
(11)   Schmidt intends (not to go if a quake occurs)

Actually, (9) will not do. Even on a wide construal of inconsistency, (9) conflicts with (3) only under non-trivial additional assumptions; e.g. that Schmidt believes that going in a quake would take him into trouble. Intentions with highly general content such as (9) are better seen as reflecting goals which call for adoption of suitable means. I will get back to blocking by way of instrumental reasoning below. (10) and (11), by contrast, are indeed inconsistent with (3). (11) is plainly so; and (10) on a reasonably broad construal of consistency which also takes a subset of the agent's beliefs into account.[18] For the sake of simplicity, let me focus on (11).

Could (11) be part of Schmidt's background of intentions, as a proponent of (ID1) might claim? It could, but that does not seem very likely. It should of course be admitted that agents are generally disposed to respond adequately to surprising events, rather than in foolish ways. And if Schmidt is like most of us, he will most likely all the while have dispositions to take quite specific measures in response to a quake; say, to immediately suspend his plans; to stay at home; maybe to turn on the radio, etc. Yet, being disposed to A in response to C usually falls short of *intending* to (A if C). An agent may be so disposed simply in virtue of a disposition to *form,* in response to C, an intention that A. Even without going into the details about when a disposition to A in response to C can be said to (co-)constitute a distal intention,[19] it is plain enough that the presence of an intention formation disposition must not be confused with the presence of a background intention. Schmidt is disposed not to go if a quake occurs, but unless he is prompted to take a stance on the scenario, he most likely does not harbor any quake-related intentions; not even tacitly. The point generalizes; proponents of (SCON) had better make do without appeals to agents' background intentions.

How about (ID2)—the claim that players of the conditionalization game *quickly form* blocking intentions? A simple temporal consideration will reveal a flaw in this line of defense, too. Suppose at t_0 you form an intention to go to the theater that night. At t_1 you are asked whether you intend to go if a quake will occur. Let it be granted for the sake of argument that the question triggers cognitive processes of yours that issue in the blocking intention (*not* to go if a quake will occur). Under the assumptions made, you will, however, *not* have had this blocking intention between t_0 and t_1. Hence, throughout the interval $[t_0, t_1]$ it would have been *rational* of you to avail of (SCON) and infer the crazy intention (to go if a quake occurs). This still seems utterly wrong. Apparently, defense (ID2) posits blocking intentions at too late a point in time to achieve its goal.

C. The Instrumental Defense

The Instrumental Defense holds that in bad-seeming conditionalization moves the agent violates either the consistency requirement on intention, or the instrumental requirement, or both. Thus, it may well be thought that what keeps Schmidt from rationally adopting (3) on the basis of (2) is that in adopting (3), he would flout the instrumental requirement as follows. Schmidt may tacitly be intending to X, and (part of) the best means to achieve X may be that if an earthquake occurs, he does *not* go to the theater. If a suitable X-term could be pointed out, he would be unable to fully rationally adopt (3) *even if* rationality should permit him to base (3) on (2).

What proposition could play the part of X in this story? Maybe Schmidt intends, however implicitly, not to get stuck in debris; or, to return to the somewhat more credible proposal from above, he may be intending *never to end up in trouble.* The impression that (3) is a crazy conclusion might then be said to stem from some such fuzzy background policy.

On this account, Schmidt would be entitled (and unless he drops his goal, at some point required) to reason:[20]

(I-1)   I intend that I never end up in trouble
         I believe forming an intention that (I do not go if a quake occurs) to be (part of) the best means never to end up in trouble
         ∴ I intend that I do not go if a quake occurs

And for sure, were he to adopt and hold on to this conclusion, he could not consistently adopt (3).

The problem with this defense is that the belief premises it needs to invoke are often too flimsy. In many cases, the instrumental belief will be false, which makes it implausible to see it as part of the agent's background of tacit beliefs. As it happens, Schmidt's is one of these cases. A minimal condition for m to be (part of) a means to e should be that the occurrence of m (jointly with the other parts) raises by some non-negligible amount the probability that e will obtain. Now, (3) will only take Schmidt into trouble if a quake will actually occur. This is extremely unlikely, though. Schmidt will almost certainly avoid trouble anyway, *whether or not* he will go if a quake occurs, and *whether or not* he *intends* not to go if a quake occurs. Consequently, not going if a quake occurs contributes so little to the likelihood of his not ending up in trouble that forming an intention not to go if a quake occurs cannot seriously be considered (part of) a *means* of trouble-avoidance; let alone (part of) the *best* means. The belief premise is simply false, and will always be so when reasoning toward intentions with unlikely internal conditions. Since the Instrumental Defense fails to block a broad range of bad-seeming conditionalizations, it will not do.

D. The Enkratic Defense

The Enkratic Defense is the attempt to exploit enkratic reasoning toward negative conditional intentions in order to block malign (SCON)-instances. Notice first of all that Schmidt is quite likely to have certain tacit ought-beliefs. In particular, he may well believe that all things considered, he *ought not* to go at 7 if a quake occurs. Now consider the following (conditional) Enkratic rule of inference:[21]

(I-2)   I believe that (I ought that A if C)
         I believe that (it is up to me whether or not A if C)
         ∴ I intend that (A if C)

According to the permission of rationality underlying (I-2), Schmidt is permitted to base on his ought-belief (and an additional belief that going would be 'up to him') an intention *not* to go if a quake occurs. And if Schmidt intends *not* to go if a quake occurs, he cannot consistently adopt an intention to *go* if a quake occurs.

But why think that Schmidt actually adopts the blocking intention? The idea is of course that enkratic reasoning in accordance with (I-2) will make him so intend. Now, (I-2) might be rooted in a *mere permission;* there might not be any corresponding rational *requirement.* The problem with a mere permission in the present context would be that you can avail of it, or not avail of it; in either case you are (in so far) perfectly rational. Absent a corresponding enkratic requirement, Schmidt could rationally choose simply not to avail of (I-2). If so, no opposing intention will be there to conflict with an intention to (go if a quake takes place), should he choose to infer it by way of (SCON) from (2). (2) would still look crazy, and the blame would rest with (SCON). To get the Enkratic Defense off the ground, we crucially need to appeal to an Enkratic *requirement*. So consider:[22]

(ECR)   Rationality requires of N that, if
    (1)   N believes at t that she herself ought that A if C, and if
    (2)   N believes at t that it is up to herself then whether or not she A-s if C, then
    (3)   N intends at t that A if C

At first glance, this may look promising. Since Schmidt satisfies sub-clause (2), there is just one way in which he can both have his ought-belief and satisfy (ECR): By not intending to (go at 7 if a quake occurs). (SCON), to be sure, would enable Schmidt to reason correctly from his intention to (go at 7) toward intending to (go at 7 if a quake occurs); but Schmidt's ought-belief would block him from rationally adopting the conditional intention.

The problem with the Enkratic Defense is that (ECR) is not really plausible. Suppose that Schmidt is asked whether he intends to go if (C:) a quake occurs. Suppose furthermore that Schmidt, like most of us, believes it to be extremely unlikely that on the day in question a quake will in fact occur. He may then well refuse to clutter his mind with a plan concerning what to do if C—and rationally so. Under limitations of time and memory, it is often simply not rational to form a plan with respect to circumstances which almost certainly will not come about. Adopting such plans may be *permissible,* but that rationality *requires* it is just as implausible as the view that agents are required to adopt believed logical consequences of what they believe, no matter whether or not they care about those consequences (Harman 1986: 12, 55-57; Broome 2013: 157f.).

To overcome the objection, (ECR) would need to be propped up with an additional sub-clause; maybe:

N believes at t that C

But then, the Enkratic Defense will be unable to block numerous bad-seeming conditionalizations toward circumstances which the agent does not all-out believe to (be going to) obtain. One of these unblocked moves would, once more, be the move from (2) to (3). Apparently then, the Enkratic Defense does not succeed either.

E.   The Intention-Belief Defense

A defense of (SCON) that turns on the coherence of intentions with the agent's beliefs about the future can be drawn from Donald Davidson's "Intending".[23] As is well known, Davidson's account identifies future-directed intention to Φ with an evaluative judgment of some sort that future Φ-ing is 'all-out', or unconditionally desirable.[24] In the case of present-directed intention, such a judgment can assume the simple form: "This action is desirable" (Davidson 1980: 98). But how can one reasonably judge *future* action desirable? This is the problem Davidson struggles with on the last pages of his essay. Suppose I affirm:

(12)   I intend to eat a candy tonight.

Can my assertion be understood as a judgment that eating a candy tonight will be unconditionally desirable? Since eating a *poisoned* candy would be highly undesirable, it seems that not. Therefore, Davidson concludes, "it would be *mad* to hold that any action of mine in the immediate future which is the eating of something sweet would be desirable" (Davidson 1980: 99; my italics).

How is this a (bad-seeming) instance of *conditionalization?* Although Davidson in the 'poisonous candy' case approaches the problem as an issue about specifying rather than conditionalizing one's intention, a relevant variation is definitely within range: It would

be equally mad to judge-desirable that I eat a candy *if it is poisoned.* The relevance of his considerations to the question at issue then become obvious.

In dealing with the problem, Davidson was well aware of the kind of solution nowadays offered by the Conditionality View. If statements such as (12) involve implicit conditions, these may block eating a poisonous candy from satisfying the reported intention's content. The conditional intention behind (12) would then be located somewhere along a spectrum of generality such as:

> I intend to eat a candy tonight *if and only if …*
> … it will not be poisoned.
> … it will be a good candy.
> … doing so will then be advisable.

But, as Davidson remarks, once propped up with conditions sufficiently comprehensive to block mad desirability entailments altogether, the statement will tell us almost nothing about what the agent intends (Davidson 1980: 93f.); so that, "if this is the road I must travel, I will never get my intentions right" (ibid.: 99). Thus having rejected the Conditionality View,[25] Davidson has to make sense of unconditional intention. Here is what I take to be his suggestion.[26] According to Davidson, the judgment an intention can be identified with is not concerned with the whole range of actions that would satisfy the intention, but only with those *compatible with the agent's beliefs.* This restriction is an essential part of his reductive definition of intention,[27] which can be stated thus:

(INT$_D$) N at t_1 intends that (A at t_2) iff N at t_1 judges that any instance of (A at t_2) *compatible with N's beliefs at t_1 about the present and the future* is unconditionally desirable.

How does the belief clause help to overcome the madness problem? Suppose:

(13)   I believe that I will not eat a poisonous candy.

In intending to eat a candy tonight, I judge that any instance of doing so that is compatible with my beliefs is desirable. But do I, against the background of (13), therein judge desirable eating a *poisonous* candy? The content of (13) entails that there will be no such actions. In this sense, the mad actions are not compatible with (13), so the answer is no. Given that an agent has the required beliefs, (INT$_D$) provides the means to exclude whatsoever crazy action from what is judged desirable.

Although Davidson's account of future-directed intention has been convincingly refuted (see Bratman 1985), it might still be thought that the gist of his 'belief filter defense' can be saved by incorporating it in a theory of rationality, the idea being that it is *irrational* to intend that (A if C) while believing all-out that in C-type circumstances one will not A. So consider the 'Belief Filter Requirement':[28]

(BFR)   Rationality requires ((N believes that ¬(C∧A))⊃¬(N intends that (A if C))

At least tacitly, Schmidt surely believes that *not both* an earthquake will occur (C) *and* he will go to the theater (A). As long as he holds on to this belief, he cannot come to intend to go if an earthquake occurs ('A if C') without violating (BFR).

The problem with the suggestion is that it inherits not only the power, but also the shortcomings of Davidson's defense. Let me confine my critique to his original suggestion. In order to block all mad desirability entailments with Davidson's belief filter, agents' beliefs would have to be either numerous or abstract enough to deny whatever mad-making

circumstance the agent could conceive of. Now it has of course to be admitted that theories of tacit belief indeed do ascribe surprisingly dense patterns of beliefs. Consider as an example Mark Crimmins' account: One "at-least-tacitly believes p just in case it is *as if* [one] has an explicit belief in p" (Crimmins 1992: 248). On this account, I tacitly believe that I never ate a bicycle: Although I never even considered the possibility, I do think and act *as if I explicitly* believed that I never ate a bicycle. Likewise, Schmidt is predicted to tacitly believe that nothing will interfere with his theater-going; after all, he plans on going just as he would if he had an explicit belief that nothing will interfere with his plan. Accounts of tacit belief can indeed make it seem as if our beliefs were virtually all-encompassing.

Yet, even our tacit beliefs are gappy. Suppose I plan to apply for an open job position at my home university, which has not been advertised. I know that at present, there is exactly one applicant, about whom I know nothing further. I also believe no one else is going to apply, and have no clue how the employer will make the decision. Being asked about my chances to get the job, I therefore estimate them at 0.5. Now consider the following possible circumstance:

(14)   The employer will give the job to my rival.

Judging it desirable to (apply for the job *if it will be given to my rival)* would certainly be foolish. Can the belief filter defense block this purported upshot of my plan? It would have to be contended that I at-least-tacitly disbelieve (14). But I lack the dispositions characteristic of having a (negative) belief. It is not as if I had an explicit belief that (14) is false. I am not in any way disposed to deny (14) when asked; nor do I in my behavior or deliberation in any way rely on my rival's failure. I do not happily tell my best friend that I am soon going to have a better job; I do not move into a larger apartment, and so on. All this makes it utterly implausible to suppose I tacitly disbelieve (14). And then, there will be no belief in my mental stock that could block the mad judgment.

The belief filter defense fails because it does not block all crazy judgments (or intentions, respectively). It would work, to be sure, if doxastic indifference were impossible; but if the only reason for assuming such an extreme view of belief is the wish to save (SCON), we should rather admit its failure.

VI.   Restricted Conditionalization (RCON)

The most natural defenses of (SCON) having failed, there may still be further requirements of rationality which give rise to defenses I have not considered. Despite this uncertainty, I suggest to radically change the strategy. By any reasonable standard, (SCON) should be considered untenable and be abandoned.

Recall, however, that we still need to make sense of those conditionalizations that seemed to be almost trivially permissible: Since Schmidt intends to (go at 7), he is permitted to adopt an intention to (go at 7 if Florence is in Tuscany). And recall also that enkratic reasoning can account for many, but not all of these harmless-seeming moves. If we discard (SCON), we had better replace it with some alternative scheme of Genuine Conditionalization. By modestly complicating (SCON), we may still obtain an inference scheme that does not produce bad-looking consequences in the first place. To that end, we will need to introduce additional schematic premises into the rule. We will thus obtain rules of *Restricted Conditionalization*—RCON, for short.

A. The Ought-Restriction and the Enkratic Restriction

Where to take suitable restrictive premises from? Even though enkratic reasoning has been proven insufficient to account for all of the data, the correctness of conditionalizing toward circumstances C does seem to be somehow *restricted* by what ought, or ought not, to be done if C were to obtain. Let us see if we can make sense of this idea. First consider a rule of 'Ought-Restricted' Conditionalization:

(I-3)   I intend that (A)
           I ought that A if C
           ∴ I intend that (A if C)

This natural suggestion will not do. (I-3) is either not yet fully explicit, or not an inference rule anyone could follow. Surely can no one avail of (I-3) without taking some *attitudinal stance on* whether she ought that A if C.[29] Hence, the second premise in (I-3) cannot yet be fully explicit. Yet, on making it explicit we obtain a rule that fatally reminds of enkratic reasoning:

(I-4)   I intend that (A)
           I *believe that* I ought that A if C
           ∴ I intend that (A if C)

With (I-4) we would once more buy into the limitations of enkratic reasoning while forgoing most of its virtues. The rule would still not support conditionalizing in the case of Meyer and the Hitchhiker from above. Further trouble is to be expected from its omission of Broome's 'up to me'-belief premise (Broome 2013: 159-63, 170f.). And most importantly, it seems simply irrational to base the conclusion on the set comprising of both the ought-belief *and the intention premise:* The ought-belief by itself sufficiently rationalizes adopting an intention to (A if C)—at any rate when A-ing if C is 'up to the agent'. The intention premise is dispensable, and would therefore have to be eliminated from the rule. Ought-Restricted Conditionalization thus collapses into a conditional version of enkratic reasoning.

B. The Disbelief Restriction

From this initial failure an important lesson should be drawn, to which I shall get back below in greater detail. The sought license to conditionalize cannot depend on an ought-belief providing *additional rational support*. Instead, the required sort of premise will have to act as a *constraint on* the rationality of conditionalizing one's intention; a premise that is negative in nature. Consider:

(I-5)   I intend that (A)
           I believe that ¬(I ought that ¬(A) if C)
           ∴ I intend that (A if C)

The second premise introduces what I shall call the disbelief restriction. Even though the rule cannot be adopted as it stands, it points the way ahead to a tenable version of RCON.

The motivation for the proposal should be clear by now. The Enkratic Defense was based on the fact that an instance of the following scheme suffices to block rational adoption of an intention to (A if C):

(15)   I believe that I ought ¬(A) if C

Even though absence of such ought-beliefs does not pose an *external* constraint on conditionalizing, one of the negations of (15) could still *internally* restrict the conditionalization rule itself. Now, there are three such negations:

I believe that I ought that ¬(¬(A)) if C
¬(I believe that I ought that ¬(A) if C)
I believe that ¬(I ought that ¬(A) if C)

The first option yields the discarded rule (I-4). The option I suggest to elaborate is the third, which is built into (I-5). As far as I can see, the second option would (after analogous elaboration) cater to the data just as well. I shall put it aside because it raises worries about reasoning with non-beliefs that need not be dealt with here.

In accounting for good-seeming and avoiding bad-seeming conditionalizations, the disbelief restriction takes us a big step forward. To show this, let me put (I-5) to work, first in explaining some good-seeming moves. (Assume for the moment that agents are fully rational.)

Example #1. Suppose that Schmidt is asked whether he intends to go to the theater at 7 if by then the sun will have set. On recalling his reasons for going, Schmidt will come to believe that in that case, he ought to go (as can be assumed). His ought-belief will then cause him[30] to tacitly *disbelieve* that he ought *not* to go if the sun will have set. Thus, he will be in the position to reason correctly toward intending to (go if the sun will have set), and that is precisely what the account needs to predict.

Example #2. Confronted with the hitchhiker's question whether she intends to (take 101 if it is about to rain), Meyer fails to see any reason *not* to take 101, even if it is about to rain. Since the decision to take 101 was underdetermined by Meyer's reasons in the first place, she thus comes to disbelieve that she ought not to take 101 if it is about to rain, and is thus in the position to infer by way of (I-5) an intention to (take 101 if it is about to rain). This is an important result because it shows that Genuine Conditionalization can yield conclusions beyond the reach of both instrumental and enkratic reasoning.

The next point to be noted is that thanks to the disbelief restriction, (I-5) steers clear of most bad-seeming conclusions. Assume as the basic scenario underlying the cases to follow that it is noon, that (2:) Schmidt intends to (go to the theater at 7), and that he intends (2) for the sole reason that he believes the announced play will please him. His present attitudes then defeasibly entail that, all things considered, he ought to (go at 7), and this is what he tacitly believes.

Example #3. Is Schmidt permitted to infer from (2) an intention (3:) to (go at 7 *if a quake will have devastated the city by then)?* It is natural to suppose that he tacitly believes he ought *not* to go if a quake occurs, which rules out ascribing to him the disbelief required to avail of (I-5). The inference is not permitted, and this is the desired outcome.

Example #4. Is Schmidt permitted to reason toward an intention to (go *if traffic jams will force him to depart hours before the play)?* In examples of this kind, it is not entirely clear what ought-beliefs the agent holds, or will come to form in cognitively processing the respective question. In order to apply (I-5) properly, we need to know the agent's relevant reasons, and how they relate to one another by her own lights. In addition, we need a formal account that tells us, given those reasons and their interactions, which (conditional) ought-beliefs it will be rational of the agent to form or have. The topic is too intricate to be treated here, so that an *ad hoc* suggestion seems appropriate. Conditional ought clauses can be construed by analogy with suppositional accounts of conditional assertions and beliefs. They can then be ascribed truth-conditions in the following, counterfactual manner, which links conditional ought with defeasible logic:

(OGT)   *N ought that A if C* iff, were the belief that C added to agent N's mental stock, N's revised set of attitudes would defeasibly entail[31] that N ought that A.

For the present purpose, this nutshell account makes it sufficiently clear what ought-beliefs agents could rationally arrive at, were they to deliberate about the matter. With respect to Schmidt, we need to ask what would happen if the belief that traffic jams will occur were added to his mental stock. Obviously, he would be provided, in addition to his reason in favor of going to the theater, with a reason *not* to go. Which of these reasons would rebut the other would then depend on their respective weights or (if you prefer) the priority ordering among Schmidt's reasons. If by his own lights, his passion for the play outweighs any worries about traffic jams he may have, his pleasure reason will pre-vail. In that case, (I-5) will enable him to reason toward an intention to (go at 7 *if traffic jams occur).* Notice that, given such ardent passion for theater, the move would not be crazy. Or else assume Schmidt would rather forgo the pleasure than have the hassle. Then his pleasure reason would be rebutted, he would come to *believe* he ought *not* to go if traffic jams occur. Consequently, he would not be in the position to conditionalize—rightly so, because in that case it would be crazy.

Example #5. Consider, by contrast, conditionalization toward an intention to (go at 7 *if by then, the theater will have collapsed).* Adding the respective belief to Schmidt's mental stock and revising his attitudes would result in a belief that he will not be able to go to the theater (in the sense of attending a performance). Since an agent cannot rationally intend an action she all-out believes she will not perform, Schmidt's ought-(dis-)beliefs are not even relevant to the case. What blocks him from being fully rational in coming to intend that (A if C) here is the intention-belief inconsistency he would incur (see Bratman 1987: 37f.).

These results are encouraging. The rest of the paper will have to deal with two seemingly problematic sorts of upshots, though.

C. Suppressed considerations

As it stands, (I-5) produces counter-intuitive results when reasoning with irrationally held intentions as premises.

Example #6. Consider once more George's case of distal akrasia. George has formed his intention (7:) to (take heroin at 9) for reasons he would by his own lights not consider decisive, were he willing to properly take them into account. This time, though, consider a good-seeming move, namely from (7) to:

(16)   George intends that (he takes heroin at 9 if by then the sun will have set)

The Conditionality View was unable to block bad-seeming inferences based on an akratic premise; the present account fails to permit an akratically based good-seeming inference. For George certainly *does* believe that he ought *not* to take heroin, whether the sun will have set or not. Unless he is outright insane (and I am suggesting he is not), this belief will block him from acquiring the disbelief required by (I-5) to take him to (16).

The counterexample features a broadly Davidsonian construal of akratic irrationality, which suggests a remedy as well. Agents sometimes base an intention on a proper subset of what, by their own lights, are reasons for or against adopting that very intention; they *suppress* some of their reasons, as I shall say. As the example shows, the sort of ought-belief that constrains conditionalization cannot require that on hypothesis C, the practical

conclusion ¬(A) would be defeasibly entailed by *the entirety* of the agent's reasons. Instead, we need to focus on those 'unsuppressed' reasons which the agent herself is willing properly to take into account in practical deliberation.

Technically, this is achieved in two steps. Let M be agent N's mental stock, and let S be the subset of M which comprises of the internal reasons suppressed by N. Then we can first define a partial ought operator by stating:[32]

(OGT')   N ought, *all unsuppressed reasons considered,* to X if C iff (M∪{N believes that C})\S after attitude revision defeasibly entails the practical conclusion to X.

Second, we can amend (I-5) so as to avail of the new partial operator:

(I-6)   I intend that (A)
I believe that ¬(I ought, *all unsuppressed reasons considered,* that ¬(A) if C)
∴ I intend that (A if C)

For all those cases from which reason-suppression is absent, (I-6) will produce the same result as (I-5). When it comes to examples such as #5, though, the weakened ought makes a decisive difference. George can plausibly be attributed the at least tacit belief that, all *unsuppressed* reasons (and only these) considered, he *ought* to take heroin at 9—whether or not the sun will have set. So he will at least tacitly *disbelieve* that he ought (unsuppressed reasons considered) *not* to take heroin if the sun will have set. He can thus plausibly be attributed the disbelief required to avail of (I-6), and this is the intuitively right result.

(I-6) is my final suggestion. And since applying the rule amounts to basing the conclusion state on the set comprising of all and only the premise states, we are now in a position to conjecture that the rational permission which underlies RCON is:

(RCP)   Rationality permits
(N intends that A
∧   N believes that ¬(she ought, all unsuppressed reasons considered, that ¬(A) if C)
∧   N intends that (A if C)
∧   N bases her intention to (A if C) on her intention that A and her disbelief)

D. The rationale of Genuine Conditionalization

A satisfactory account of Genuine Conditionalization has been reached. Some of its upshots, however, can still seem counter-intuitive. But rather than threatening the achieved result, these upshots can lead us to a deeper understanding of *why* RCON is correct reasoning; or so I am going to argue.

Example #7. Schmidt plans to go to the theater because he believes the play will please him. Next, he plays the conditionalization game, and player 2 confronts him with the following shocking figment. "Just when about to go to the theater, you will receive a call from inside the theater building. It will be your child, reporting in distress that a major earthquake has just destroyed the roof structure during the afternoon performance for kids. You will thereby gather two crucial pieces of information: There will be no play (¬P), and your child is trapped under the debris inside the theater (T). Do you intend to go to the theater if this story will come true?"

On the present account, Schmidt *is* entitled to avail of (I-6) and infer:

(17)   I intend to (go at 7 if (¬P∧T))

Here is why: No suppressed considerations being involved, Schmidt is permitted to draw conclusion (17) just in case he disbelieves he ought not to go if (¬P∧T). And that he at least tacitly entertains this disbelief is indeed warranted. Granted: In the fictitious circumstances, his actual reason for going—the expectation to see a play—would not hold any more; there would be no play. Yet, the second part of the conjunction would give him a substitute reason for going: He would need to save his child. Hence he will come to believe that he ought to go if (¬P∧T), and this in turn will make him *disbelieve* that he ought not to go if (¬P∧T). The present account therefore predicts the rational answer to player 2's question to be "yes".

Notice that this answer is not in itself wrong. If Schmidt redeliberated from scratch whether to go if (¬P∧T), he could not end up with anything but (17). It is not as though (17) were a crazy intention. What can make the example look troublesome is *the inferential path* on which Schmidt arrives at (17). It can seem crazy of him to infer (17) from (a premise set that includes) his priorly formed intention to go to the theater. After all, the reasons on which his intention is actually based would in the envisaged circumstances be annihilated without residue; and it seems as though Schmidt would not intend to go *for the right reasons,* were he to intend to in consequence of a purely hedonic decision.

The example thus displays what can be called *reason disconnect.* Here is a brief analysis. Call those reasons which an agent actually takes to support A-ing the 'A-reasons', and call those reasons which *would* support A-ing if C came about the 'C-hypothetical reasons for A-ing'. Then reason disconnect takes place just in case an agent's A-reasons do not intersect with her C-hypothetical reasons for A-ing. This is what happens in the example, and it gives rise to a quite natural worry: How can it in such cases be rationally permissible to base an intention to (A if C) on one's intention to (A)?

It might be thought that the problem can be coped with by further tinkering with the inference rule. I suggest to resist the pull and reflect instead on why precisely reason disconnect looks troubling. It is tempting to think of RCON as correct in so far it reflects the occasional robustness of A-reasons in the face of C-hypothetical reasons against A-ing. Yet, it is crucial to see that such robustness considerations, however well they may go with enkratic reasoning, are clearly out of line with the very idea of Restricted Conditionalization. If RCON is to fare better than enkratic reasoning, it must not draw on the force of agents' A-reasons. As Meyer's Buridan case illustrates, there will often be no such reasons at all; even so, conditionalizing basically is a plausible move in such cases.

I contend that RCON is sanctioned instead by the normative import of an agent's very intention to A, *as opposed to* the reasons in support of it. As Bratman has argued, an intention to A, once formed, has a normative import over and above the import of the reasons for its adoption. This is most conspicuously revealed in situations in which the scarcity of deliberative resources makes it irrational to reconsider one's intention even though the reasons which originally rationalized its adoption have, as a matter of fact, been overturned. In such constellations, it can be rational to hold on to intentions which, were the decision to be made from scratch, one ought by one's own lights not to adopt.[33] Intentions thus do not only physically tend to persist; to a certain degree, there is also rational pressure to retain them once they are formed, and to abstain from reconsidering them.[34] RCON, properly understood, reveals just another facet of this normative import of intending as such.

In this perspective, example #7 does not look threatening any more. That Schmidt's inference toward (17) would display *reason disconnect* does not threaten its correctness because he need not therein rely on the robustness of his pleasure reason to go. If he avails

of RCON (rather than redeliberating from scratch, thereby bypassing his theater-intention), he will therein be respecting the normative import of his actual intention to go. He will be fully rational in doing so because he believes that this normative import is not overturned by his C-hypothetical reasons. That these reasons include an *even stronger* reason for A-ing—his child being in distress—than his actual hedonistic reason, is simply a different matter.

VII.   Conclusion

As it has turned out, there is no support for identifying flat intention with the mostly crazy state of intending 'no matter what'. For the identification to hold true, Simple Conditionalization (SCON) would have to be good reasoning. But (SCON) faces obvious counterexamples which cannot be defused or deflected by any natural means. Restricted Conditionalization (RCON), by contrast, has turned out to be on much more solid footing. Here, an additional premise makes sure that in coming to intend that (A if C) the agent does not deem intending to A a bad idea should C obtain. In availing of ought-beliefs, RCON resembles enkratic reasoning, but there remains a crucial difference: While enkratic reasoning mobilizes practical reasons to rationalize its conclusions, RCON relies on the normative import of intending as such; practical reasons are invoked, but as a constraint only. As a consequence, RCON yields more conclusions than (both instrumental and) enkratic reasoning. This result is welcome because some conditionalization moves are rationally unassailable and yet dictated by neither reasons nor goals.

On the suggested account, flatly intending to A is intending to A in those circumstances which, loosely speaking, *would not make it a bad idea to A.* But then, to intend flatly to A cannot mean to intend to A *in any circumstances*; flat intention cannot be a 'no matter what' stance. The Conditionality View's key motive for systematically reading conditional clauses into ordinary intentions thus evaporates. Its adherents were not totally mistaken, to be sure. To my mind, Ferrero and others have started out with an accurate observation: Our intentions do not normatively commit us in any way to actions that would be foolish or immoral by our own lights. Where they have gone wrong is in explaining this observation by appeal to restrictive conditions within the content of intention. I have, by contrast, been suggesting to locate the sought restriction within the content of a *basing permission of rationality;* a permission that shapes the inferential role characteristic of intentions, thus reflecting a feature of the state of intending as such. As a consequence, a simpler picture of agency regains credibility; a picture according to which everyday intentions are in themselves just as simple, wieldy and unsophisticated as they present themselves to us when we form, revise, or execute them, or communicate them to others.[35]

Notes

[1]   The view is implicit in Meiland 1970: 18f. and Grice 1971. For an explicit elaboration, see Bratman 1979 and, more recently, Ferrero 2009 and Klass 2009: 111, 124. Analogous claims have been made with regard to desire, see McDaniel/Bradley 2008: 280, 282; Lycan 2012: 209, and Fara 2013. Parfit 1987: 151 has advanced the somewhat related view that most desires are implicitly conditional on their own persistence; see McDaniel/Bradley 2008: 270.

[2]   I.e. as long as C is compossible with A.

[3]   The proponents of the Conditionality View have not been outspoken about what precisely, on their view, makes unconditional intentions crazy; usually the claim is put forward with direct appeal to intuitions.

But since they propose a conditional view of intention as the remedy, my reading of the 'no matter what' clause should be acceptable to them as well. Cf. Klass 2009: 107; Ferrero 2009: 700.

4   In what follows, A and C stand proxy for full-fledged, self-standing propositions. The A-term will usually involve a self-representation of the intending subject as herself.

5   The external/internal distinction is drawn by Cartwright 1990: 235; see also Ferrero 2009: 701f. and, for desire, McDaniel/Bradley 2008: 272, who discern (read: internally) conditional from (read: externally) hypothetical desire. As the latter authors point out, a conditional desire need not be conditional on its own persistence (cf. Parfit 1987: 151); the same holds for intention.

6   There are various reasons to be careful here. McDaniel/Bradley 2008 have argued that conditional desires and intentions are attitudes toward two propositions rather than one; and in linguistics, Angelika Kratzer's 'restrictor view' offers a highly unified account of if-clauses in general by suggesting that they serve to restrict the modal base of some (at times covert) modal operator; see ch. 4 of Kratzer 2012.

7   On issues about intentions with internal necessary conditions, see Klass 2009: 120-24.

8   For the details, see Broome 2013: 159-170, 262, 264.

9   Cf. Broome 2013: 159-63, 170f., 290.

10 Cf. Broome 2013: 246-48, 255. Broome defends a first-order account of reasoning according to which what is usually operated on in active reasoning are not attitudes, but pairs of propositions and attitude markers, jointly constituting 'marked contents'; see ibid.: 251f. For the sake of greater convenience, (COR) operates on attitudes, but I do not intend to take a stance on whether reasoning is 'first' or 'higher-order'.

11 '$Y_1 \dots Y_n \therefore X$' will serve as an alternative rule notation.

12 Cf. Broome 2013: 225. Notice that basing permissions on this construal of 'basing' are permissions to perform acts of attitude-construction. Broome's basing permissions, by contrast, are permissions *to be based* rather than *to base*, as is clear from the temporal pattern displayed by their contents. Cf. ibid.: 187.

13 I am grateful to an anonymous referee for having brought the possibility of such an account to my attention.

14 Notice that $P \supset (C \supset P)$ is a tautology, and will therefore validate trivially in Hintikka-style belief semantics as presented e.g. in Hendricks/Symons, 'Epistemic Logic', sec. I. On closure principles for belief, more will be said in section V.D. below.

15 Cf. the two more sophisticated versions in Ferrero 2009: 723.

16 Ferrero admits certain exceptions, such as fanatics and agents pursuing an absolutely praiseworthy goal; see Ferrero 2009: 727, 730. These can be put aside here.

17 See Ferrero 2009: 720, 723.

18 It seems to be a requirement of rationality that ¬((N intends that X) ∧ (N intends that Y) ∧ (N believes that with at least metaphysical necessity, if X then not Y)). This goes beyond Broome's respective requirement of intention consistency; cf. Broome 2013: 156.

19 For an account of the difference, see Mele 2007.

20 For an elaborated account of instrumental reasoning, see Broome 2013: 159-170, 262, 264. I am simplifying matters here; against the background of Broome's account, (I-1) must be considered enthymematic.

21 For an accurate non-conditional version that operates on 'marked contents', see Broome 2013: 290. According to Broome, m is 'up to me' iff were I myself not to intend m, because of that, m would not be so. For the details, see Broome 2013: 159-63, 170f.—Why appeal to a *conditional* version of Broome's Enkratic inference rule? If we construed both conditional intentions and conditional ought-beliefs as attitudes towards conditional propositions, we could make do with Broome's rule, the first premise of which runs: "I believe that I ought that P". We could then substitute a conditional-proposition signifier "(A if C)" for "P", which would give us the sought inferential route towards a conditional intention. Yet, it is far from evident that in reports of conditional intention, the intention operator takes a conditional proposition as its argument, and the same worry could of course be raised about conditional ought statements. The difficulty is circumvented, however, when we leave the scope of 'ought' as ambiguous as found in natural language, stick with our convention to read "N intends that (A if C)" as not committing us to a particular stance on the conditional-proposition issue, and state a new, self-standing, conditional version of the enkratic rule.

22 This is an adaptation of the simplified enkratic requirement of rationality in Broome 2013: 171.

23 See Davidson 1980: 92-102. I shall untie his account of unconditional intention from most of its original context: i.e. the question of whether intentions are a kind of belief. For a discussion of this aspect see Velleman 1989: 114-24; see also Ferrero 2009: 732f., n. 13.

24 See Davidson 1980: 100. One might think the view has a problem to account for genuinely conditional intention—a sort of intention that Davidson clearly wants to admit (see ibid.: 94). But there is no problem here. A conditional intention is an intention to (A if C), and his proposal maps this smoothly onto judging that (A *if C)* is unconditionally desirable.

25 Later on, he seems to have changed his view on the matter: „I cannot imagine someone intending, even at the start, to buy *any* car. Even if it were wired to explode? Surely there are many things I could easily specify in advance, and would if I were ordering, say, from a catalogue. However, there is never a point at which I could completely specify the content of my intention […]", Davidson 1999: 498.

26 The following interpretation has been suggested by Bratman 1999: 213-19.

27 "[…] there is nothing absurd in my judging that any action of mine in the immediate future that is the eating of something sweet would be desirable *given the rest of what I believe about the immediate future",* Davidson 1980: 99. "To intend to perform an action is, on my account, to hold that it is desirable to perform an action of a certain sort in the light of what one believes is and will be the case", ibid.: 100.

28 Notice that (BFR), although closely related to, is *not* a substitution instance of Bratman's intention-belief consistency constraint, which can, for the sake of comparison, be rendered as: Rationality requires ((N believes that ¬P)⊃¬(N intends that P)). Cf. Bratman 1987: 37f. What prevents the derivation of a relevant norm from the latter is that because of the first negation sign, substituting "(A if C)" for the first occurrence of P makes sense only if '(A if C)' denotes a proposition.

29 Expressivists will disagree with this. If ought-utterances are treated as expressing some e.g. volitional attitude, my main objection against (I-4) applies directly to (I-3): If I ought that A if C, this by itself sufficiently rationalizes adopting an intention that (A if C).

30 On sub-personal processes which ensure consistency by automatically erasing beliefs, see e.g. Broome 2013: 77f., 189, 268. Such revision dispositions belong to the foundations of our rational capacity.

31 John Horty's theory of defeasible reasoning is well suited for fleshing out this proposal because he identifies conclusions of practical defeasible reasoning generally with unconditional oughts; see Horty 2012: 18, 65ff. His logic thus meshes nicely with my account of conditional ought. I take it, however, that practical reasoning (defeasible or not) can also yield intentions themselves as its proper conclusions; see Broome 2013: 250ff.

32 'A\B' denotes the set difference of A and B, and 'A∪B' their union. I take it that conditional ought is a four place operator OUGHT(N, X, C, T) where T stands for the 'things considered', but a definition of partial ought in general will not be needed.

33 For an example, see the second Mondale case in Bratman 1987: 74f. See also Bratman 2012: 74, 79f.

34 Bratman has outlined this rational aspect of the 'inertia' of intentions in terms of a system of rationality constraints on reconsidering one's intentions; see Bratman 1987: 16, 60-110. The nature of these pressures is still under debate; see Bratman 2012; Ferrero 2012, 2014. Problems stem from the fact that dropping one's intention before its execution is often perfectly rational. I take it, however, that there are genuinely diachronic rational pressures toward retaining one's intentions—even though I agree with Ferrero that these pressures never reach more than minimally into the future. See Ferrero 2014: 332; Gillessen 2015.

35

References

Bratman, Michael. 1979. "Simple Intention." *Philosophical Studies* 36: 245–59.

———. 1987. *Intention, Plans, and Practical Reason.* Cambridge, MA: Harvard University Press.

———. 1999. "Davidson's Theory of Intention." Reprinted in *Faces of Intention. Selected Essays on Intention and Agency*. Cambridge, New York: Cambridge University Press, 209-224.

———. 2012. "Time, rationality, and self-governance." *Philosophical Issues* 22: 73–88.

Broome, John. 2013. *Rationality through reasoning.* Malden, MA: Wiley-Blackwell.

Cartwright, J.P.W. 1990. "Conditional Intention." *Philosophical Studies* 60: 233–55.

Crimmins, Mark. 1992. "Tacitness and Virtual Beliefs". *Mind & Language* 7: 240–63.

Davidson, Donald. 1980. "Intending." In *Essays on Actions and Events.* Oxford: Clarendon Press, 83–102.

———. 1999. "Reply to Edna Ullmann-Margalit." In *The philosophy of Donald Davidson*, ed. L. E. Hahn. Chicago: Open Court, 497–99.

Fara, Delia Graff. 2013. "Specifying Desires." *Noûs* 47: 250–72.

Ferrero, Luca. 2009. "Conditional Intentions." *Noûs* 43: 700–41.

———. 2012. "Diachronic Constraints of Practical Rationality." *Philosophical Issues* 22: 144–64.

———. 2014. "Diachronic Structural Rationality." *Inquiry* 57: 311–36.

Gillessen, Jens. 2015. "Do Intentions Set Up Rational Defaults? Commitments, Reasons, and the Diachronic Dimension of Rationality." *Pacific Philosophical Quarterly.*
DOI: 10.1111/papq.12147

Grice, Herbert Paul. 1971. "Intention and Uncertainty." *Proceedings of the British Academy* 57: 263–79.

Harman, Gilbert. 1986. *Change in View. Principles of Reasoning*. Cambridge, MA; London, England: MIT Press.

Hendricks, Vincent, and John Symons. 2015. "Epistemic Logic." In *Stanford Encyclopedia of Philosophy*, ed. E. Zalta. Fall 2015 Edition.
http://plato.stanford.edu/archives/fall2015/entries/logic-epistemic/

Horty, John. 2012. *Reasons as defaults.* New York: Oxford University Press.

Klass, Gregory. 2009. "A Conditional Intent to Perform." *Legal Theory* 15: 107–47.

Kratzer, Angelika. 2012. *Modals and conditionals*. Oxford, New York: Oxford University Press.

Lycan, William G. 2012. "Desire Considered as a Propositional Attitude." *Philosophical Perspectives* 26: 201–15.

McDaniel, Kris, and Ben Bradley. 2008. "Desires." *Mind* 117: 267–302.

Meiland, J.W. 1970. *The Nature of Intention.* London: Methuen.

Mele, Alfred R. 2007. "Persisting Intentions". *Noûs* 41: 735–757.

Parfit, Derek. 1987. *Reasons and Persons*. Oxford: Clarendon Press.

Velleman, David. 1989. *Practical Reflection.* Princeton, NJ: Princeton University Press.