

Conditional Intentions

LUCA FERRERO

University of Wisconsin, Milwaukee

“[Don] Giovanni’s own end and the end of the opera both affirm that there is no actual human life that can be lived as unconditionally as his, but, at the same time, they recognize that any vitality that other people have must sustain the dream of being as free from conditions as his was.”

Bernard Williams (2006, p. 22)

1. Introduction

Many, if not most, of our avowals of intention are linguistically unconditional. However, when we say ‘I will φ ’ or ‘I intend to φ ’ we rarely mean to profess an intention to φ *no matter what*—‘*Fiat φ , pereat mundus.*’ Often even an emphatic unconditional profession of intention such as ‘I intend to φ at all costs’ or ‘I will φ no matter what’ is not to be taken literally given that one is not really willing to φ at the price of losing one’s life or making Heavens fall. Most intentions appear to be conditional in their ‘deep structure’ even when the conditions are not explicitly stated. This should not be surprising. For agents like us, who have a plurality of ends and live in a complex world, truly unconditional intentions directed at particular kinds of actions are often recipes for disaster. They are likely to impose too rigid a constraint on our conduct by demanding that we pursue them at all costs. This seems to be too high a price to pay, unless we would like to live in the unconditional mode of Don Giovanni, whose “single-minded determination to live at the fullest energy, at the extreme edge of desire—as Bernard Williams writes—neglects consequences to himself as much as to others,” and leads him to an inevitable end, facing the natural consequences of his recklessness.¹

In spite of the ubiquity of conditional intentions, the contemporary debate in the philosophy of action has only occasionally taken up the issue of

their nature and distinctive contribution to our agency.² This neglect could be justified if one could easily carry over to intentions the outcomes of the lively philosophical debates about the nature of indicative conditionals and conditional obligations.³ But this is not so. Because of the special role that conditional intentions play in our agency, especially in its diachronic dimension, conditional intentions raise a distinctive set of questions and puzzles that deserve a separate—and long overdue—investigation.

In this paper, I will discuss the various ways in which intentions can be said to be conditional, with particular attention to the internal conditions on the intentions' content. I will first consider what it takes to carry out a conditional intention. I will then discuss how the distinctive norms of intention apply to conditional intentions and whether conditional intentions are a weaker sort of commitments than the unconditional ones. This discussion will lead to the idea of what I call the 'deep structure' of intentions. Roughly, this is the idea that the conditional nature of our intentions is only partially made explicit in the expressions we use to communicate our intentions and in the explicit form of our thinking about and reasoning with them. Most conditions that qualify our intentions are part of a deep functional structure that can be evinced by observing the actual psychological functioning of intentions and by considering the rational requirements that they engage. I will argue that the deep structure of intentions is characteristically conditional. Genuinely unconditional intentions are only limiting instances of conditional intentions and their contribution to agency can only be understood in light of this fact. I will conclude by showing that the characteristic conditional structure of intentions is intimately related to distinctive features of human agency, especially to its unity over time.⁴

2. Internally Restrictive Conditions

Conditional statements of intentions are often ambiguous between an external and an internal interpretation of the condition.⁵ In describing the meeting between the Commendatore's statue and Don Giovanni in the churchyard scene in Act II of Mozart's *Don Giovanni*, one might say, 'If Don Giovanni invites the statue to dinner, the statue will dine with him.' Externally interpreted, the invitation is a condition on the Commendatore's *acquisition* of the intention to dine with Don Giovanni. In this sense, the occurrence of the invitation either causes the Commendatore's brute acquisition of the intention to dine with Don Giovanni or provides a sufficient reason for him to adopt such intention. In either case, the conditional clause does not qualify *what* the agent intends to do. When the condition obtains, if ever, the agent is supposed to acquire the intention to φ *simpliciter*. By contrast, when interpreted *internally*, Don Giovanni's invitation qualifies the *content* of the Commendatore's intention. What the Commendatore intends to do is not 'to dine with Don Giovanni *simpliciter*,' but 'to dine with

Don Giovanni if invited to do so.’ The internal reading implies that the agent has already acquired the intention and thus he is already subject to the distinctive rational pressures imposed by it. The externally conditional statement, ‘If C, x will φ ,’ instead, suggests that the agent does not yet intend to φ and thus he is not yet subject to its rational pressures.⁶ If we consider a first person conditional profession of intention, ‘I will φ if C,’ the internal reading stands for an avowal of intention, ‘I hereby undertake the intention to: φ if C;’ the external reading stands for a prediction of one’s future undertakings, ‘I predict that, if C, I will undertake the intention to φ *simpliciter*.’⁷

Externally conditional statements of intention are kinds of ordinary indicative conditionals. As such, the philosophical questions that might arise about them are not specific to the philosophy of action. Internally conditional statements of intention, instead, concern genuinely conditional *intentions*. This paper is only concerned with the latter.

There are two basic kinds of internal conditions: enabling and restrictive.⁸ Both kinds of conditions are expressed by the distressed Donna Elvira when she first enters the stage and sings, ‘Ah, se ritrovo l’empio / E a me non torna ancor, / Vò farne orrendo scempio / Gli vò cavare il cor’ (*Don Giovanni*, Act I.5)—‘Oh! If I find the traitor / and if he does not return to me, / I’ll inflict havoc on him / and tear his heart out.’ The first antecedent, ‘If I find the traitor,’ states an *enabling* condition on Elvira’s intended action. She intends to slay the traitor, Don Giovanni, but she is able to do so only if she first finds him. Let’s call all the circumstances that must obtain for an agent x to be able to φ the *preconditions* P of x ’s φ -ing. The preconditions include the universal enablers of agency (e.g., Elvira is to be alive and not paralyzed), the general enablers of the action *type* φ (e.g., Elvira’s having the tools required to tear someone’s heart out), and the specific enablers required for x ’s performance of a token φ -ing in x ’s particular situation (e.g., that Elvira finds Don Giovanni).

The second condition stated by Elvira, ‘If he does not return to me,’ is *restrictive*. She does not intend to slay Don Giovanni in all circumstances in which she has the ability and opportunity to do so. She is giving him another chance; she is not to tear his heart out if he returns to her. A restrictive condition C is any contingent circumstance that delimits what it takes for the agent to *carry out* her qualified intention to φ , that is, to *succeed* in that specific undertaking.⁹ Restrictive conditions restrict the range of situations in which the agent is required to φ in light of her qualified intention to do so.¹⁰ If an agent intends to φ unconditionally, in the sense that she imposes no restrictive conditions on her φ -ing, as long as she has not given up the intention, she is under a rational demand to φ under all circumstances; she is to φ *no matter what*. By contrast, if the agent intends to ‘ φ if C,’ she is under a rational demand to φ only when she *takes* C to obtain that is, only when she either believes or accepts that C is the case. (Throughout this paper, I will

use ‘acceptance’ in Michael Bratman’s sense as the context-sensitive attitude of taking a certain proposition for granted in practical deliberation.¹¹)

The notion of restrictive conditions is the one usually associated with the idea of conditional intentions. Unless otherwise noted, hereafter, I will use ‘conditions’ and ‘conditional’ to refer to restrictive conditions and restrictively conditional intentions. I will use the terms ‘preconditions’ and ‘preconditional’ to refer to enabling conditions. For reasons soon to be made clear, I will use the formula, ‘to intend to (φ IF C),’ to refer to the intention to φ qualified by internal restrictive conditions C and the formula, ‘to intend to [PRE P] φ ,’ for the intention to φ qualified by internal preconditions P.

It might be tempting to think that restrictive conditions could be characterized as standing for *reasons* in support of the intended action. But this is not sufficient to differentiate them from both preconditions and external conditions. First, we often cite the obtaining of an enabling condition as a reason for φ -ing. Second, the obtaining of an external condition might be a decisive reason for the adoption of the intention to φ (e.g., Don Giovanni’s invitation might be a decisive reason for the Commendatore to decide to dine with him). What distinguishes restrictive conditions is not that they are reasons in support of φ -ing, but their role in determining what it takes to carry out the intention that they qualify. They qualify the content of the intention and thus determine under which circumstances the intention is fulfilled. This is not to deny that restrictive conditions bear on the advisability of φ -ing. If an agent intends to (φ IF C) rather than to φ unconditionally or to φ under a different sets of conditions, it is reasonable to expect that, in the agent’s own lights, there is something about C that makes φ -ing advisable when C obtains but not necessarily otherwise.¹² Hence, certain features of C are going to be among the agent’s subjective reasons for her φ -ing under those circumstances (for more on advisability, see §8).¹³

It is obvious that, if C obtains by the time of the intended action and the agent is unable to make C false, the only way to carry out the intention to (φ IF C) is to φ . But what about those cases in which, by the time of action, C can no longer obtain (whether because C was never possible or because it is now too late for it to obtain)?¹⁴ Is the intention satisfied by the falsity of the antecedent? And what about cases in which the agent has the power to make C false? Can the agent carry out the intention by falsifying its antecedent? If satisfying a conditional intention were a matter of making true the material conditional ($C \supset \varphi$), the answers would be positive. When C can no longer obtain, the agent would automatically be guaranteed the satisfaction of her conditional intention, no matter whether she is going to φ or not. And if the agent had some control over C, she could carry out the conditional intention by making sure that, by the time of action, C does not obtain. For instance, a nurse who avows the internally conditional intention, ‘If the patient is still alive tomorrow, I will change his dressing,’ might carry out her intention by killing the patient.

Arguably, there is something odd with this way of satisfying the conditional intention. In the literature, the oddity of cases like the nurse's scenario is sometimes presented as the main case against the truth-functional reading.¹⁵ This is, however, far from conclusive. Considerations independent of the conditional structure of the intention—say, the nurse's professional obligations—might explain why it is unacceptable, in her particular situation, that she satisfies the conditional intention by falsifying the antecedent. A defender of the truth-conditional interpretation might thus reply that, as far as the conditional structure of the intention is concerned, it is in principle always possible to satisfy that intention by falsifying the antecedent.

A stronger case against the truth-functional reading requires a deeper understanding of its distinctive features. If satisfying a conditional intention to $(\varphi \text{ IF } C)$ were a matter of satisfying a the material conditional $(C \supset \varphi)$, the intention is carried out not only by φ -ing when C obtains but also as an immediate consequence of the fact that C can no longer obtain. The material conditional does not discriminate between the two alternative ways of satisfying the intention, formally described. This is especially apparent when the agent has power over the occurrence of C since the conditional intention does not put any more pressure on being satisfied by trying to φ when C obtains rather than by trying to make C false. According to the truth-functional reading, *in case C can no longer obtain*, the agent's conduct is still in compliance with the rational pressures imposed by her conditional intention, the agent is still under the jurisdiction of the conditional intention even if, by that point, she satisfies it no matter what she is going to do. The conditional intention gives the agent permission to either φ or not- φ . The agent is still under the intention's guidance, even if this guidance takes the form of a permission to do what the agent wants about her φ -ing.

The problem with the truth-functional reading is that it does not acknowledge the special status of the action of φ -ing in the satisfaction of a conditional intention. The only way in which a conditional intention can be genuinely *carried out* is by φ -ing in those circumstances when C obtains by the time of the intended action. There is nothing that counts as carrying out a conditional intention when, by the time of action, C can no longer obtain. When C can no longer obtain, the agent does not thereby *succeed* at her conditional pursuit. Carrying out an intention is a matter of a practical achievement. But discharging a conditional intention when C can no longer obtain does not involve the exercise of any practical skill, effort, or willpower, i.e., of any of the practical abilities associated with the idea of an accomplishment.¹⁶ Hence, when C is no longer possible, success at $(\varphi$ -ing IF $C)$ is no longer possible either.

Could it be that these considerations apply only when the agent has no control over C ? Perhaps, but even so, this would show that there is a problem with the truth-functional reading. For the truth-functional reading is supposed to apply to all conditional intentions, whether or not the agent

has control over C. In any event, even when the agent has control, the fact that preventing C might be an accomplishment and satisfy the agent's intention to make C false does not entail that succeeding at preventing C amounts to carrying out the intention to (φ IF C). Rather, by making C false, the agent makes the conditional intention *pointless*. It makes the satisfaction of the intention impossible, since the agent can no longer φ when C obtains, that is, he can no longer φ 'against the backdrop' of C, so to say.

When C no longer obtains, the agent is no longer under any of the rational pressures generated by the conditional intention because they no longer apply. The intention does not allow or permit the agent to do whatever she prefers to do about φ , as the truth-functional reading would have it. Rather, the agent is no longer under the intention's jurisdiction. For the intention has turned out to be moot. This implies that the conditional intention offers no more guidance to the agent about her future conduct. To the extent that the intention still says something to the agent, it is not telling her, 'You are permitted to do whatever you want about φ -ing,' but, 'Go back to the drawing board and figure out what to do, now that C is no longer to obtain.' (In some cases, it might not be necessary to go back to the drawing board because the agent might already have the intention to (ψ IF \neg C); but having a complementary conditional intention of this sort is neither a necessary nor a usual feature of conditional intentions.)

The fundamental problem with the truth-functional reading is that it misses the distinctive role of C in 'setting the stage' for the performance of φ . It is because of this stage-setting role, that a conditional intention can continue to guide the agent and, eventually, be satisfied only as long as C might still obtain. When intending to (φ IF C), the agent is not aiming at making true that ($C \supset \varphi$). The material conditional is not the content of a conditional intention.¹⁷ Rather, the agent is aiming at φ -ing against the backdrop of the obtaining of the restrictive conditions C. Hence, when C can no longer obtain, the conditional intention can no longer be carried out, it becomes *moot*.¹⁸ This is the characteristic structure of internally conditional intentions, which I represent by the use of IF.¹⁹

Notice that the claim that a conditional intention turns moot when its condition can no longer obtain is not peculiar to conditional intentions. As some philosophers have argued, the same holds true of conditional imperatives and prescriptions.²⁰ In addition, even an unconditional intention to ψ can become moot. This happens whenever ψ -ing turns out to be impossible (including those situations where—through no fault of her own—the agent loses either the ability or the opportunity to ψ). What the conditional structure adds is only a further way in which an intention might turn out to be inapplicable: the case where the condition C can *no longer* obtain.²¹

My considerations against the truth-functional reading are centered on the idea that an intention cannot be genuinely *carried out* by the fact that C can no longer obtain. Carrying out an intention by the performance of

the intended action is the basic sense in which we talk of the satisfaction of an intention. There is, however, another sense in which an agent can be said to be successful with respect to her intentions. Over a certain period of time, an agent can succeed at *living up to* the rational demands imposed by an intention even if, eventually, she does not carry it out. For instance, having done everything that is in her control and having complied with all the rational pressures of the intention to φ , the agent might still fail to φ and thus to satisfy the intention because the world does not collaborate with her efforts. Alternatively, the agent might successfully live up to the demands of a particular intention up to the time when she revokes it or when the intention turns out moot (as it happens when φ -ing turns out to be either impossible or, in the case of a conditional intention whose conditions can no longer obtain, pointless).²² Given that the rational pressures are in place as soon as the intention is adopted, continuous compliance with them is a kind of ongoing success with respect to that intention, although one that is not guaranteed to issue in the ‘perfecting’ of the intention, that is, in the accomplishment of a successful performance of the intended action. Last but not least, success at carrying out a conditional intention does not necessarily mean that the very preferences the agent had in adopting that plan are entirely satisfied. This is the familiar predicament of the pursuit of what Bratman calls ‘precautionary plans,’ plans aimed at dealing with the unfortunate occurrence of undesirable conditions. For instance, Masetto intentions’ to ‘beat Don Giovanni IF he tries to seduce Zerlina’ is a precautionary plan, given that Masetto prefers Don Giovanni never to try to seduce his fiancée.²³ In spite of his preference for C never to obtain, Masetto is still supposed to live up to the demands of his precautionary plan. As long as Don Giovanni’s seduction is still a live possibility, Masetto wants to be able and ready to beat him. It also follows that, while pursuing the precautionary plan, Masetto might also intend to prevent the obtaining of C. That is, there might be no inconsistency when Masetto intends to make the conditional intention moot while also trying to live up to its demands. For Masetto wants to be ready to beat Don Giovanni in case he fails to prevent his seduction. But contrary to what the truth-functional reading suggests, if Masetto succeeds in preventing the seduction, he has not thereby carried out the conditional intention. He has rather succeeded in his other (and primary) pursuit: depriving the precautionary plan of its point, making the conditional intention moot.²⁴

3. Suspended Intentions

The suggestion that a conditional intention is moot when C can no longer obtain appears similar to Dorothy Edgington’s account of indicative conditionals. She claims that affirming a conditional ‘If C, then p’ amounts to asserting that p *on C being true*, asserting *nothing* otherwise. Edgington suggests that such an analysis should be extended to other kinds of propositional

attitudes and speech acts, such as desires, commands, questions, promises, and apologies.²⁵ She does not offer a sustained discussion of this extension but for a few remarks on conditional desires²⁶ and for the suggestion that a conditional command “can be construed as having the force of a command of the consequent, conditional upon the antecedent being true.”²⁷ In other words, the condition obtaining activates the command, which prior to that moment is not yet in force. Nothing is commanded *unless* and, if so, *only when* the condition obtains. A conditional command is a ‘suspended’ command, so to speak. In a similar fashion, conditional promises, bets, wishes, apologies, etc. are ‘suspended’ until, if ever, the condition obtains. Prior to that moment, no one is expected to comply with the demands generated by the conditional attitude or speech act, given that at that time the attitude or speech act is not yet in place. And if the condition can no longer obtain, the agent has commanded/wished/promised/etc. *nothing* (although she might have done something significant in performing the conditional speech act). For instance, if Don Giovanni were to say to Donna Elvira, ‘If I have hurt your feelings, I apologize,’ he would not have apologized if he had never hurt her feelings (although he would have still done something significant in saying so).²⁸ Using my terminology, Edgington’s view is that, if the condition can no longer obtain, a conditional command/promise/apology/etc. is moot.

An account à la Edgington, however, cannot be extended to internally conditional intentions. An intention to (φ IF C) is in full force even when C does not yet obtain. A conditional intention is not a *suspended* intention. As long as the agent has that intention and the intention has not turned out to be moot, the agent is expected to *live up to* the demands imposed by her conditional intention. She is expected to do so *from the moment* she acquires the conditional intention. The obtaining of C does not activate the force of the conditional intention. When the agent comes for the first time to believe or accept that C,²⁹ she is already subject to the characteristic norms of intending (such as means-end coherence, intention-belief consistency, and agglomerativity, see §5). The only difference is that, in light of the obtaining of C, these norms impose more specific demands. For instance, prior to coming to believe that C, means-end coherence requires that the agent take the means she believes necessary in order to be *ready to* φ in case C obtains. Once the agent takes C to obtain, however, she is required to take the means she believes necessary to φ . In addition, once she takes C to obtain, the agent’s plans that are contingent on \neg C become moot.

Conditional intentions are not ‘suspended intentions.’³⁰ In some special situations, however, a conditional intention might mimic the operation of a suspended attitude. This happens for those intentions that, prior to the obtaining of C, are ‘practically idle’ since there is nothing that the agent is supposed to do in order to prepare for φ -ing, avoid threats to its future execution, and secure the agglomeration of this intention with his other

plans. In these special cases, the conditional intention operate as if it were suspended, since prior to the occurrence of C there is nothing that the agent is actually to do to live up to the distinctive norms of intention. However, this is not because the demands are suspended. It is only because living up to them comes for free.³¹

Conditional intentions should not be confused with another kind of suspended states, namely, the undischarged conditional statements in incomplete episodes of practical reasoning. Let's imagine that, Donna Anna, at some point in her deliberation about how to avenge her father's death, assents to the conditional, 'If Don Giovanni killed my father, I'll kill Don Giovanni.' The antecedent states a restrictive condition on her killing of Don Giovanni. If Anna were to close the deliberation at that point, she should adopt the conditional intention to (φ IF A), to 'kill Don Giovanni IF Don Giovanni killed my father'. However, she might not yet be ready to close her deliberation, i.e., to assent to the all-out judgment that she is to (φ IF A). Although she thinks that, on the basis of the reasons she has taken into account until now, φ -ing is the best she can do if A obtains, she still does not want to commit to the conditional intention. She can suspend her deliberation for the time being. She might simply want to sleep on it before drawing such a momentous conclusion. Or she might want to gain some time to collect more relevant information, review the evidence already available, or double-check her reasoning. Having suspended the deliberation, she has not yet acquired a conditional intention. Hence, she is under none of the rational pressures generated by this intention. She has not suspended her intending but her reasoning. As such, she might still be under the rational pressures generated by her assent to the conditional in her deliberation. Depending on where she was in her reasoning, she might even be under a rational pressure to acquire the intention to (φ IF A). This might happen, for instance, if she acquires conclusive evidence that Don Giovanni is the murderer. But when this happens, it is not because the obtaining of A activates the force of a suspended intention. For prior to that moment, Anna has never adopted the conditional intention; she has never drawn that conclusion and thus never put herself under any requirement to φ in case A obtains. As for the commitments that she adopts or incurs in her reasoning as a result of her assent to 'If A, I'll φ ,' they are not suspended either. They are in force from the moment of her assent onward, but they are distinct from the requirement of a genuine conditional intention.³²

4. Unconditional Intentions

A genuinely *unconditional* intention is an intention to φ *no matter what*.³³ This intention puts no restrictions on the pursuit of φ -ing. As soon as the time of the intended action arrives, and the agent has the ability and opportunity to φ , the agent is supposed to φ , period. Pure unconditional intentions of

this kind are much less than common than it might appear at first. Many linguistically unconditional expressions of intention do not really express pure unconditional intentions. For instance, when Leporello says, ‘Io deggio ad ogni patto per sempre abandonar questo bel matto’ (Act I.15)—‘*Whatever the consequences may be, I must leave this madman forever,*’ it is unlikely that he is avowing a pure unconditional intention. He is not really willing to leave his master at the price of bringing about his death or the end of the world (after all, four gold pieces are sufficient to convince him to keep working for Don Giovanni, see Act II.1). The problem is not that Leporello does not understand what an unconditional intention commits him to. In many cases, unconditional expressions of intention, even those that emphasize their alleged unconditional structure, are not to be interpreted *literally*. They do not express *pure* unconditional intentions. They are rather linguistic devices for signaling that one’s intention to φ is qualified by a much smaller set of restrictive conditions than it would be normally expected.

What happens to a conditional intention to (φ IF C) when the agent comes to take that C is going to obtain on time for her φ -ing? It would be a mistake to think that his intention is thereby transformed into a pure unconditional one. If Masetto comes to believe that Don Giovanni attempted to seduce Zerlina, he is not thereby committed to beat Don Giovanni *no matter what*. Masetto is committed to doing exactly the same thing as he was committed before, to beat Don Giovanni *in the eventuality that* Don Giovanni attempted to seduce Zerlina. The only difference is that now Masetto believes that eventuality to obtain. Masetto intends to φ *given that he now takes C to obtain*. Masetto goes from intending to (φ IF C) to intending to (GIVEN C; φ). The condition C is not erased from the content of the intention; it is rather pushed in what might be called the ‘cognitive background’ of the agent’s plans and projects. This ‘cognitive background’ is comprised of the set of all the circumstances relevant to the possibility and reasonability of an agent’s intentions that the agent takes to obtain either now or at any time prior to the completion of her plans. Once a condition C is moved into the cognitive background, there is usually no longer a need for the agent to entertain the condition explicitly in his practical reasoning. The agent might thus proceed to reason as if the intention were unconditional. In a similar fashion, there is often no need to make the condition explicit in the expression of one’s intention, at least in so far as one’s audience is expected to believe that C obtains as well.

Clearly, once the agent takes C to obtain, he no longer needs to direct the audience’s attention to the issue whether C is going to obtain. He also acknowledges that the rational pressures generated by the norms of intending now take a particular shape (see §3, p. 707). But the conditional intention never becomes a pure unconditional one. If, prior to completing the beating of Don Giovanni, Masetto realizes that he is mistaken about Don Giovanni’s attempted seduction and the condition C does not obtain, he must revert to

his original conditional intention in the ‘to (φ IF C)’ form. He must push C out of the cognitive background and reconsider what is to do in order to carry out this intention, since it is no longer the case that he is required to beat Don Giovanni. By contrast, if a conditional intention were literally *to turn into a* pure unconditional one once the agent thinks that C obtains, Masetto would be stuck with the requirement to φ *regardless* of whether he later finds out that he is mistaken about C. For a pure unconditional intention is literally under literally under no restrictive conditions at all, whether implicit or explicit, whether given in the cognitive background or not.

Although the obtaining of C does not transform a conditional intention into a pure unconditional one, it can make it into a *circumstantially* unconditional one. This happens when all the restrictive conditions on the φ -ing (explicit and implicit alike) are taken as given; when they are all moved into the agent’s cognitive background. When this occurs, the agent has removed any uncertainty about the circumstances that bear on her φ -ing. She can now drop any related contingency plan and focus her efforts making sure that she φ ’s. Superficially, a circumstantially unconditional intention behaves like a pure unconditional one. As long as the agent is right that all restrictive conditions obtain, what she has to do to carry out her intention to φ IF C is no different from what she would have to do if she were to carry out a pure unconditional intention to φ no matter what. But this convergence in the demands imposed by the two distinct kinds of intention is merely accidental. The convergence only concerns the implementation of the intended action. A circumstantially unconditional intention to φ imposes different demands on the articulation, justification, and explanation of one’s intentions. For instance, the two kinds of intentions are expected to yield different answers to the question of what the agent is committed to under the counterfactual situation in which C does not obtain. This difference bears on the agent’s understanding of the nature of her intention and on what she is supposed to do to articulate, justify and explain it.

Raising questions about these counterfactual situations helps us make explicit, at least in part, what I will call the ‘deep structure’ of the agent’s intention. This structure includes all the restrictive conditions that the agent attaches to his φ -ing independently of their present epistemic status. In this sense, the deep structure of the circumstantially unconditional intention to (GIVEN C; φ) is still conditional. The agent’s intention is still to (φ IF C) rather than to intend to φ *no matter what*. But since the agent now takes C to obtain, C has now been pushed in the agent’s cognitive background (more on the deep structure at pp. 714–717).

Notice that if the agent is not challenged to partially articulate her intention, the ordinary expressions of her circumstantially unconditional intentions are linguistically *unconditional*. It might actually be misleading for the agent to report a circumstantially unconditional intention by a conditional

indicative. The use of the conditional indicative might suggest that she is still uncertain about the obtaining of the explicitly mentioned conditions.³⁴ It might suggest to her audience that she does not think that her intention is circumstantially unconditional. As previously remarked, linguistically unconditional expressions of intention do not necessarily reflect the deep structure of the intention.

A circumstantially unconditional intention to φ is not necessarily an *executively* unconditional intention, that is, an intention on which the agent is to φ *hic* and *nunc*. First, the agent might take some conditions as given only prospectively; she expects them to obtain in the future although on time for her φ -ing. Second, even when all restrictive conditions already obtain, the agent might still be unable to φ (that is, some of the necessary *preconditions* of her φ -ing do not yet obtain). If Elvira believes that Don Giovanni is never going to return to her but she has not yet found him, her intention to slay him is circumstantially but not yet executively unconditional. She still has to wait for the time of action to turn ripe. When it does, the agent is required to φ immediately. This is not because she is now to φ *no matter what*, but because there is *no more matter* she still *has to wait for*.

5. Conditional Intentions and the Norms of Intending

Does the conditional structure of an intention make a difference to the application of the characteristic norms of intending? Is an agent who intends to (φ IF C) under a different set of rational requirements than the agent who intends to φ *no matter what*? Is a conditional intention a less stable and weaker commitment than an unconditional one? (Hereafter, unless otherwise noted, I will use ‘unconditional intention’ to refer only to *pure* unconditional intentions.) Let’s first consider how the norms of intending apply to a conditional intention to (φ IF C). I will focus here on three norms that according to Bratman are characteristic of intending: means-end coherence, intention-belief consistency, and agglomerativity.³⁵ Consider means-end coherence. An agent who intends to (φ IF C) is subject to the demand to take the means she believes necessary in order to φ if C no less than when she intends to φ *no matter what*. The only difference is that some of the things necessary to pursue an unconditional φ -ing are not required by the conditional pursuit of φ (given that in the latter case one is to see to it that one φ ’s under fewer circumstances). This is a difference in the individuation of the necessary means not in the norm engaged by the intention. Moreover, the norm begins to apply from the moment the agent acquires the intention to (φ IF C), even if she does not yet take C to obtain (see p. 707). From that moment on, the agent is to get ready for her φ -ing *just in case* C is to obtain. Similar considerations can be made about intention-belief consistency, i.e., the rational demand that—as Bratman puts it—“one’s intentions, taken together with one’s beliefs, fit together into a consistent model of one’s future.”³⁶ The

irrationality of intending something that is believed to be impossible follows from this norm. This norm applies to conditional intentions as much as to unconditional ones. In particular, the agent is not to intend to (φ IF C) if she takes that, were C to obtain, it would be impossible to φ . Finally, the norm of agglomerativity requires that if the agent intends to φ and she also intends to ψ , then she is to intend to (φ & ψ). An intention to (φ IF C) is to be agglomerated, no less than an unconditional one, with the agent's other intentions, whether conditional or not. This is so even prior to her discovery whether C obtains.³⁷ To sum up, conditional intentions are under exactly the same requirements as unconditional intentions. To this extent, a conditional intention, qua intention, is no less of a commitment than an unconditional one.³⁸

Are conditional intentions less stable than unconditional ones? Are they more prone to reconsideration? Is the agent less settled on a conditional intention than on an unconditional one? I see no reason to think that conditional intentions are any less stable, in any of the above senses. There might still be a concern, however, about the relation between stability and the role of intentions in transtemporal coordination. Prior to knowing whether C obtains, a conditional intention does not offer a fixed point for future coordination. The agent cannot yet rely on her future φ -ing as she could if she intended to φ unconditionally. This might seem a troubling limitation, given that a major rationale for adopting intentions is their contribution to transtemporal coordination. But we should not exaggerate the difference between conditional and unconditional intentions on this point. First, a conditional intention might already offer some simplification over the possible future scenarios that the agent is to consider for transtemporal coordination. The conditions C that the agent imposes over her φ -ing need not be ultra-specific. They often cover a relatively large set of possible circumstances. In this sense, they might eliminate the need for some contingency planning. The agent can already anticipate this much: under the *set* of circumstances C, she is going to φ . Her conditional intention is not highly sensitive to the minutiae of the actual circumstances in which C occurs.³⁹ We do not conditionalize intentions in the mode of ideal deliberators with unlimited time, resources, and information. Moreover, even prior to determining whether C obtains, the agent can rely on some expectations about her future conduct. In particular, she can expect that, at least for the time being, she is going to do what she is required to do by the norms that apply to her conditional intention even prior to her knowing about C, such as that she is expected to take the means that she already knows to be necessary to be ready to φ *just in case* C is going to obtain.

This does not mean that the agent should ignore the scenarios in which C is not going to obtain. She might even care about getting ready for them. But contemplating these scenarios and preparing for them does not necessarily undermine the firmness of the agent's purpose and the stability of

her resolve to (φ IF C). The norms of intending are violated only when contemplating and preparing for the alternative scenarios is incompatible with contemplating and preparing for the situations in which C obtains. When there is no incompatibility, the agent cannot be accused of being irrational for contemplating and preparing for the case in which C does not obtain. Doing so, however, might be unreasonable. This depends on the agent's expectation about the obtaining of C, what is at stake in φ -ing, how burdensome the preparation for φ -ing is, and how risk-averse the agent is. In certain contexts, the most reasonable conduct is to accept that C and thus behave as if the intention were circumstantially unconditional. At the opposite extreme, C might be so remote a possibility and φ -ing so unimportant that the most reasonable strategy is to do only what is *minimally* required to keep the possibility of her φ -ing alive in case C obtains, while devoting all other available resources to planning and preparing for $\neg C$.⁴⁰

It might be tempting to think that a conditional intention involves a weaker degree of commitment.⁴¹ This cannot mean, however, that the weakness is in the attitude itself. As we have just seen, conditional and unconditional intentions are subject to the same characteristic norms and they are equally stable. Nor can it mean that agent is less firm in her purpose, since a restrictive condition is an *internal* qualification of the content of her attitude.⁴² (Some conditional linguistic expressions of intention might suggest weaker resolve, but they do not impose internal restrictive conditions, see p. 725.) The only sense in which a conditional intention might be said to be weaker is a matter of its *content*. The agent who intends to (φ IF C) is set on φ -ing under a *smaller* set of circumstances than an agent who intends to φ no matter what. This restriction affects the specific *application* of the characteristic norms of intending. Attending to these differences might give us evidence about the particular restrictive conditions, if any, that the agent imposes on her φ -ing. This, in turns, bears on the individuation of what she is up to. But these differences in the application of the norms do not suggest that the conditional intention is a weaker kind of commitment. They rather help us appreciate what the agent thinks about the merits of her φ -ing. Given that the agent could φ even when $\neg C$, we can ask why she is imposing those restrictions on her φ -ing. What is it about C that, in the agent's own eyes, makes φ -ing to be pursued under those circumstances rather than different ones? Why does she not intend to φ unconditionally? Or not to φ at all? The answers to these questions offer glimpses into the agent's beliefs, practical standpoint and character. These considerations, in turn, bear on our evaluations of the agent and of her conduct, including attributions of *mens rea*,⁴³ blame, and responsibility.⁴⁴ It is only in this sense that we might legitimately speak of 'degrees of commitment.' The variation reflects the agent's attitude toward the relative merits of φ -ing, rather than the nature of the attitude she acquires and the strength of the commitment

she undertakes once she settles on φ -ing, whether conditionally or otherwise.

6. Necessary Preconditions

When Elvira avows, ‘If I find the traitor, I’ll inflict havoc on him,’ the conditional clause, ‘If I find the traitor,’ plays a different role than would the restrictive qualification, ‘If he does not return to me’ (pp. 702–703). The former clause specifies an enabling condition or precondition P on her slaying of Don Giovanni. The clause refers to the circumstances under which it is possible for Elvira to inflict havoc on Don Giovanni rather than to those under which she would *elect* to do so.⁴⁵ Whereas the conditions of satisfaction of the intention to (φ IF C) are different from those of the intention to φ *no matter what*, the conditions of satisfaction of the unconditional intention are the same as those of the intention to ‘ φ if P.’ This is shown by the fact that appending the clause ‘If P’ to the avowal of the intention to φ does not add any interesting information about the content of the intention, at least for those who already know about that P is a precondition of φ -ing. The explicit mention of P might be a matter of the pragmatics of the communication of intentions rather than of their structure. This is the same conclusion that Donald Davidson reached in his discussion of conditional intentions. For him, there are conditional clauses that do not express *bona fide* conditions. They only warn the audience that the agent is more likely to fail than they might expect given that a precondition P is more unlikely to obtain than the audience might expect.⁴⁶ Using my terminology, what Davidson argues is that the logical form of a statement like, ‘I intend to tear Don Giovanni’s heart out if I find him,’ is not, ‘I intend to (tear Don Giovanni’s heart out IF I find him),’ but a combination of an *avowal* of the unconditional intention to tear Don Giovanni’s heart out and a *warning* about the likelihood of the precondition P of finding him. The precondition P does not qualify the content of the intention. What Elvira intends to do is to tear Don Giovanni’s heart out *simpliciter*: ‘I intend to tear his heart out *simpliciter*. . . but be warned that my finding him—a necessary precondition of my intended action—is less likely to obtain than you might reasonably expect.’

Davidson is right in calling attention to the distinction between preconditions and restrictive conditions. But he is wrong both in claiming that preconditions are not *bona fide* conditions at all, and that they can only be mentioned to cancel conversational implicatures. As I am about to argue, preconditions play a distinctive and central role in the structure of our planning attitudes. Let’s adopt the formula, ‘to intend to [PRE P] φ ,’ to stand for, ‘to intend to φ on the necessary precondition P,’ or—as I will sometimes say, ‘to intend to φ *preconditionally* on P.’ Although it is internal to the content of the intention, a precondition does not qualify the intention restrictively.

Preconditions and restrictive conditions can appear side by side but they are independent of each other. An agent x might intend, on the precondition that P , to φ on the restrictive condition that C , i.e., ‘ x intends to [PRE P] (φ IF C).’ And one can intend to φ in a pure unconditional way, meaning with no restrictive qualification, while intending to φ preconditionally on P , ‘ x intends to [PRE P] φ no matter what.’ In fact, one must always intend to φ on its necessary preconditions P since one cannot intend to do something that is impossible. Hence, whereas restrictive conditions might be elective, the agent has no choice over preconditions.

Although one cannot but intend to φ on its preconditions, we can speak—by analogy with the case of restrictive conditions—of pure, circumstantially, and executively *unpreconditional* intentions (pp. 708–711). First, there might be *pure* unpreconditional intentions. In principle there could be actions that can be intended in a pure unpreconditional way. These are actions that the agent, as long as she is alive, has the opportunity and ability to perform under any imaginable circumstances. Since by definition the agent always has the ability and opportunity to perform them, these are very unusual actions—like the thinking of the ‘cogito.’ A more useful notion is that of a *circumstantially* unpreconditional intention: an intention to φ *all* of whose preconditions the agent takes to obtain *at the appropriate time* for her φ -ing.⁴⁷ By analogy with the case of circumstantial unconditional intentions (p. 710), the preconditions that are taken as given are moved in the agent’s cognitive background but they are not erased from the intention’s deep structure. The intention ‘to [PRE P] φ ’ is transformed in the intention ‘to [GIVEN P] φ .’ There are two main reasons why preconditions that are taken as given are still part of the deep structure of intention. First, the agent might later find out that she is mistaken in thinking that P obtains. When so, she has to revert to the preconditional intention to φ and, in case she discovers that P is never to obtain, she is to drop her intention to φ because φ -ing has turned out to be impossible. The second reason concerns the agent’s need to keep track of the compatibility of her intentions (see p. 718).

That an intention is circumstantially unpreconditional does not mean that the agent is to act on it without further ado. First, some of the preconditions taken as given might not *yet* be in place. Second, even if they are in place, some restrictive conditions might not. An intention can be executive only when all of its conditions and preconditions are taken to obtain not just on time for her φ -ing but also *at the present time*. It is only at this particular point that the time of action is ripe.

To understand the role of preconditions P in the structure of intentions, let us consider how preconditions relate to the conditions of satisfaction of intentions, the norms of intending, and transtemporal coordination. First consider conditions of satisfaction. When P does not obtain, the intention is moot because φ -ing is impossible. This mootness is more momentous than the one induced by the lack of a restrictive condition C (for in the latter case,

the agent is still able to φ but she is under no requirement to do so). One important consequence is that, if the agent has power over P, she cannot use it to prevent P without violating means-end coherence. By contrast, means-end coherence allows the agent to try to falsify C while intending to (φ IF C), (see p. 717). The status of preconditions bears immediately on the compliance with the norm of intention-belief consistency. It is irrational, after all, to intend to do something that one believes to be impossible, that is, whose preconditions can no longer obtain. However, if one is uncertain about P, there is inconsistency in intending to φ , as long as one is ready to give up this intention immediately upon learning, if ever, that P can no longer obtain.

Finally, preconditions are crucial for securing the agglomerativity of intentions. Restrictive conditions by themselves do not tell us whether two intentions can be agglomerated, unless one intends to (φ IF C) and one also intends to ($\neg\varphi$ IF D) where condition D overlaps with C. In all other cases, to determine whether two intentions can be agglomerated, we must rule out any incompatibility in their preconditions. Such an incompatibility would imply that, sooner or later, the pursuit of one plan would make the pursuit of the other one impossible.

Don Giovanni is throwing a party for Masetto and Zerlina. Let's imagine that Donna Anna intends to go to the party if and only if Don Ottavio is going to be there. Anna's intention is the combination of two conditional plans: the intention to (go to the party IF Ottavio is going to the party) and the intention to (not go to party IF Ottavio is not going to the party). An agent can have a genuine biconditional intention only if its two components agglomerate. That is, only if the two sides of the biconditional can be pursued at the same time without interfering with each other up to the moment when the agent finds out whether the condition obtains or not (at that point one of the two intentions becomes moot and the agent is to make true only one side of the conditional). For instance, if Ottavio is supposed to come and pick up Anna on the way to the party, there is nothing that Anna is to do to prepare for her going to the party that might interfere with the success of the other side of the biconditional, with her intention *not* to go to the party in case Ottavio is not going either. In this scenario, there is no problem with the agglomeration. Let's now imagine that Ottavio is supposed to go to the party independently of Anna and that he might not be able to get there before Anna does. In this case, Anna's two intentions might at some point conflict with each other. If, by the time Anna is supposed to be at the party, she does not yet know whether Ottavio is going to show up, she has to decide which intention to give up. She cannot wait and see which side of the biconditional turns out to be moot. By that time it might be too later for her to fulfill the intention that is left standing (say, if she waits home too long, by the time she learns that Ottavio is going to the party, it might be too late for her to leave and get to the party on time). She is rather to take a chance, say, to go to the party in the hope that Ottavio goes as well. Thus she would

give up the intention to stay home in case Ottavio is not going to the party. Once the time to go to the party comes, if Anna does not know whether Ottavio would be there, the two intentions can no longer be agglomerated, given that some of their preconditions have become incompatible. A similar story can be told about any pair of intentions of the agent, even those that, taken together, do not amount to a biconditional intention.

If two intentions cannot be agglomerated, there are several strategies that the agent can adopt to try to meet the norm of agglomeration. The most obvious but drastic one is to give up one of the intentions. Alternatively, she might hold on to both by trying to add further restrictive conditions so as to eliminate the conflict in their preconditions.⁴⁸ Another solution is to give priority to one project and subordinate the other to the fate of the primary one: in case of an *actual* conflict in preconditions, the primary one takes precedence at the expenses of making the secondary one impossible.⁴⁹ Finally, if no subordination is possible or desirable, one might *demote* the projects from full-fledged intentions to the lesser status of mere ‘endeavorings’ or ‘settled objectives,’ which are attitudes that do not engage the full range of norms of intention, especially agglomerativity (see p. 711).⁵⁰ To sum up, the important conclusion is that, no matter which strategy is adopted to secure agglomerativity, the strategy must ultimately aim at avoiding the conflict of preconditions. Preconditions have the last word on agglomerativity.⁵¹

Finally, preconditions play a role in transtemporal coordination. One cannot expect to φ unless one expects that its necessary preconditions obtain. The extent to which one can rely, for coordination purposes, on the intention ‘to [PRE P] (φ IF C)’ is a function of one’s expectations about *both* P and C. As a result, it might happen that one is more confident about one’s future φ -ing on the basis of one’s conditional intention to φ than about one’s future ψ -ing on the basis of one’s unconditional intention to ψ . This happens if ψ -ing is preconditional on circumstances that are less likely than the combination of both the conditions and the preconditions of φ -ing. For instance, my intention (to go to Italy in 2010 IF I am awarded a sabbatical) offers me more assurance about my future conduct than my intention to (go to Mars in 2012 *no matter what*). The latter intention, although unconditional, is preconditional on circumstances that are much more unlikely than the combination of the conditions and preconditions of the former intention. Therefore, although a pure unconditional intention to φ is a better guide to coordination than the conditional intention to (φ IF C), a pure unconditional intention to ψ is not necessarily a better guide than a conditional intention to (φ IF C).⁵²

What is the agent required to do about the preconditions that are under her control? Here lies an important difference with restrictive conditions. In principle, the intention to (φ IF C) allows the agent to do whatever she wants about C. As long as the agent does not jeopardize the possibility of φ -ing in case C turns out true, the agent is allowed to try to prevent C.

For instance, Masetto's precautionary plan to (beat Don Giovanni IF Don Giovanni attempts to seduce Zerlina) does not require Masetto to be idle; it is perfectly reasonable that he should do his best to prevent Don Giovanni's seduction (see p. 706).⁵³

By comparison, controllable preconditions impose stricter constraints on what the agent can rationally do. Elvira has some control over the precondition P on her intention to slay Don Giovanni. She has some control on whether she is going to find him since she might decide not to actively seek for him. Given that she might run into Don Giovanni even if she does not look for him, she is not required to be actively looking for him. But means-end coherence requires Elvira not to *preclude* the possibility that P might obtain, i.e., that she might find him. If Elvira were to deliberately withdraw from the world to live in seclusion, she would thereby repudiate her intention to slay him.

A particularly interesting kind of repudiation occurs when the agent gives up an intention that she already has so as to make room for novel plans that, because of conflicting preconditions, cannot be agglomerated with the older ones. Imagine that Ottavio is settled on spending a quiet evening at home reading Molière's *Dom Juan*, when he unexpectedly gets invited to Don Giovanni's party. Having sworn to Donna Anna that he would avenge her, he accepts the invitation in order to be able to come face to face with Don Giovanni. His new intention to avenge Anna is incompatible with the older one to read *Dom Juan*. The precondition of reading—i.e., staying at home—is incompatible with the precondition of avenging Anna—i.e., going to the party. Up until he receives the invitation to the party, Ottavio takes his staying at home as given. But this does not prevent him from contemplating the possibility of going to the party. This is because he does not take the precondition of his reading *Dom Juan* as unchangeable. He realizes, however, that were he to go to the party, agglomerativity requires him to give up the intention to read Molière's comedy.

Keeping track of the preconditions of one's intentions, keeping track of each intention's deep structure (including those portions that are already taken as given), is crucial to one's ability to appreciate the revisions that might be rationally required in order to embark in novel undertakings.

Finally, although means-end coherence forbids the agent to attempt making a precondition P false, the agent might qualify her intention by imposing restrictive conditions about her control over P so that she can be allowed to try to prevent P without any incoherence. For instance, Don Giovanni cannot marry Zerlina unless she consents to it. If he proposes by saying, 'I will marry you if you consent,' he makes explicit a precondition P on his intended action but he imposes no restrictive condition on it. His proposal is, 'I'll marry you, preconditional on your consent,' i.e., 'I intend to [PRE P] φ .' However, if Don Giovanni says, 'I will marry you if you consent by *next week*,' he introduces a restrictive condition on the intended action by

qualifying a necessary precondition of his marrying her. The form of his new proposal is, 'I intend to [PRE P] (φ IF C_P),' where C_P stands for a restrictive condition C on the precondition P. The intention is both *preconditional* on Zerlina's accepting the proposal at some unspecified future time, and *restrictively conditional* on her accepting by next week. If the agent happens to have some control over C_P , he is under no rational pressure to use this power one way or the other. Don Giovanni might make a sincere proposal to marry Zerlina IF she accepts by next week, even if at the same time he is trying to prevent Zerlina from accepting the proposal by the deadline. This case might seem bizarre, but it violates none of the norms of intending (see p. 717). Don Giovanni does not violate any norms of intending as long as his attempt at preventing Zerlina's timely acceptance does not jeopardize his ability to marry her in case she happens to accept by the deadline, that is, before he succeeds at preventing it.⁵⁴

7. The Deep Structure of Intentions

Elvira's conditional intention to slay Don Giovanni explicitly mentions the precondition of finding him and the restrictive condition that he is not going to return to her (p. 702). Many other qualifications are left implicit. Most obviously, Elvira is not listing all the preconditions of her action (for instance, she does not intend to slay Don Giovanni if she is paralyzed). But it is also plausible that she is imposing additional, non-act-essential, restrictive conditions on her φ -ing. Some of these conditions are generic; they apply to large classes of her pursuits. Most of our ordinary projects are usually conditional on the non occurrence of massively disruptive events such as natural disasters, terrorist attacks, medical emergencies, and personal tragedies. Events of this kind do not necessarily make our ordinary pursuits impossible, but they often make them unadvisable. Many of our intentions, therefore, are implicitly qualified by clauses such as, 'If I am not the victim of an earthquake,' 'If my partner does not die in a car crash,' etc. These generic conditions are not the only ones to be left implicit. Usually, our intentions are qualified by many *specific* restrictive conditions that we not normally bother to make explicit. For instance, Elvira's intention to slay Don Giovanni might not be made moot by Don Giovanni's return *simpliciter*. Even if she does not say so, she might mean that his return is to be sincere and lasting. Moreover, even if Elvira is willing to forgive Don Giovanni for all the seductions meticulously recorded in Leporello's catalogue (Act I.5), she might not be willing to take Don Giovanni in if, say, he were to turn out to be a child molester or a serial killer. She might neither mention nor think about these conditions because she takes them to be either too obvious or too unlikely. Nonetheless, these conditions have a place in the psychological economy of her planning attitudes and behavior. This role is made manifest not only by the rational agent's behavior in case the implicit conditions do

obtain, but also by her responses to the request to contemplate counterfactual or hypothetical scenarios involving these conditions (a request that is often prompted by the demand to either justify or articulate one's plans). These implicit conditions, both generic and specific, are part of what I call the 'deep structure' of her intention to φ .

The deep structure of an intention to φ includes *all* the preconditions P and the conditions C that actually qualify its content, whether generic or specific, whether explicit or implicit. Let us schematize it as follows,

(DS) x intends to [PRE $P_1, P_2 \dots P_m \dots$] (φ IF $C_1, C_2 \dots C_n \dots$)

Appeal to the deep structure is required for a complete explanation of the functioning of a particular intention and the appreciation of the specific rational pressures to which it is subjected. This appeal, however, is most effective when we factor in the status of these qualifications in the agent's cognitive background, i.e., when we take into account which qualifications the agent takes as given (p. 709).⁵⁵ The 'epistemic' version of the deep structure reads as follows,

(DS-E) x intends to [GIVEN $P_1, P_2 \dots$; PRE $P_m, P_{m+1} \dots$]
(GIVEN $C_1, C_2 \dots$; φ IF $C_n, C_{n+1} \dots$)

Although the deep structure DS of the intention to φ is not supposed to change over time, its epistemic version might constantly change to reflect the modification in the agent's beliefs and acceptances concerning the preconditions and conditions of her φ -ing.⁵⁶

As previously discussed, once a qualification is taken as given, it is not erased from the content of the intention. If the agent discovers that a condition is not to be taken as given, she is not thereby saddled with an unconditional intention (p. 709; p. 715). In any event, qualifications that are taken as given play important roles in the economy of planning independently of any revision in the cognitive backgrounds of intentions. First, the agent must keep track of these qualifications to determine whether her plans are going to aggregate (p. 718). Second, the extent of what the agent takes as given bears on her contingency planning. It does so in that it determines how she is to allocate time and resources in preparing and monitoring for various contingencies. Last but not least, the agent might have to appeal to the 'given' qualifications when she is challenged to justify or articulate her projects.

The deep structure of an intention includes all the qualifications that play a role in the psychological economy of our planning attitudes. A qualification plays this role even at those times when it is not explicitly entertained or expressed. The qualification has a role in the overall economy of planning even if it is in the background of attention or lies dormant in the recesses of the agent's mind. As long as it is ready to contribute, when appropriately

prompted, to the shaping of the agent's conduct and practical thinking, any qualification could play this role.

In several cases, the remoteness of a condition from the focus of attention does not correlate with the time it takes for the agent to respond to the appropriate trigger. For instance, if one is asked to contemplate (let alone actually faced with) the effects of natural disasters and personal tragedies on the intention to φ , often it takes very little thought to acknowledge whether these occurrences would make the intention moot or not.

In other cases, however, delving into the deep structure might require more time and thought. It might not be immediately clear to the agent whether certain circumstances count as a condition or precondition of her plan to φ . But this does not entail that these conditions do not play a role in her psychological economy. Figuring out the full extent of the deep structure of an intention is often a work in progress.

This is not simply because intentions are, as Bratman says, 'partial plans' (that is, plans that need to be filled-in as we proceed in their implementation). Even when considered in their 'partial' form, our intentions are partly shaped by self-directed interpretation. Unearthing their deep structure is often a matter of *articulation* in Charles Taylor's sense rather than of the description of a fully independent object.⁵⁷

Furthermore, intentions are not freestanding attitudes. They are supposed to fit with our view of the world, our practical standpoint, and the full network of our cognitive and practical commitments. To articulate the deep structure of an intention is often a matter of making sure that the intention fits with our other attitudes. This task might not always be accomplished once and for all. For discoveries and changes in other parts of our mind might always call for further digging into the deep structure of any particular intention to determine its compatibility with the novel configuration of our mind.

Finally, our intentions are responsive to characteristic rational pressures. Hence, their deep structure is supposed to be partly open to the specifications and changes required to meet the demands imposed by the distinctive norms of planning agency. In this sense, not all the qualifications of an intention need to be acknowledged and fully articulated prior to actual demands to demonstrate or secure compliance with these rational pressures.

Intentions whose deep structure gets *fully* articulated might be quite rare. Such full articulation is rather something to which we aspire when we strive to secure that our intentions are a proper fit both for our world and for our mind. The notion of a fully articulated deep structure of intention works as a regulative ideal of planning agency. Hence, the full articulation of the deep structure of any particular intention is in principle always an ongoing job. As such, we must allow for some open-endedness in the listing of the qualification of each particular intention. This is the reason why I have used ellipses in listing of conditions in the two formulas (DS) and (DS-E).

This open-endedness is part and parcel of the nature of our intentions as attitudes that are both responsive to rational pressures and accountable to the regulative ideal of full articulation. This open-endedness is a feature of the actual psychological *modus operandi* of intentions. In this indirect sense, it is thus possible to claim a psychological role and reality for the yet unarticulated qualifications of an intention and consider them among the components of its deep structure.⁵⁸

Only few qualifications, if any, of the deep structure of our intentions are usually made explicit both in communicating our plans and in thinking about them. Which conditions are usually made explicit? And which conditions should be made explicit? It is not the job of this paper to offer either descriptions or prescriptions about the *salience* of particular conditions in everyday practical thinking and communication.⁵⁹ I will only remark that the relevance of a particular qualification depends on a complicated combination of a variety of contextual factors. To mention a few: how likely that condition is taken to be; how specific it is to the intention in question, rather than a more general qualifiers of a larger class of intentions; whether the agent has the power to alter that condition; and how obvious the relation between the qualification and the intended action is; etc. These are some of the considerations that affect the relevance of particular conditions in practical reasoning (including how much attention and articulation they deserve) and in the explicit professions of our intentions. Because of the limitations in the resources we can devote to deliberation and communication, it is not surprising that usually we explicitly deal with only a small portion of the deep structure of our intentions.⁶⁰ We are usually quite proficient at determining which conditions should get the privileged treatment. Being a *reasonable* planning agent is partly a matter of having a sense of salience that is properly tweaked to respond, in each distinct context, to the complex interaction of the various factors of relevance so as to determine which conditions are to be made explicit both in practical thinking and in communication.⁶¹

8. Generic Placeholders

Although we are usually unable to offer a complete and explicit articulation of all the conditions and preconditions that qualify our intentions, we are not blind to the complexity and open-endedness of their deep structure. As planning agents, we operate with (at least) the implicit sense that, even in the case of a seemingly unqualified intention to φ , we are set on φ -ing only in so far as we think that our φ -ing is both feasible and advisable. Our planning attitudes are supposed to be responsive both to their advisability (as partially reflected in the articulation of the specific restrictive conditions of each individual plan) and to their feasibility (as reflected in the articulation of each individual plan's specific preconditions). I suggest to represent our implicit sense of the open ended articulation of intentions by inserting in the

(DS) schema two generic placeholders—FEASIBLE and ADVISABLE—in lieu of the ellipses in the lists of preconditions and conditions.

(DS*) x intends to [PRE P_1 – P_m & FEASIBLE] (φ IF C_1 – C_n & ADVISABLE)

The two generic qualifiers stand in for the specific Ps and Cs that, although not immediately operative in her psychology, could be elicited when the agent either faces some unusual and unexpected scenarios, or is challenged to justify and articulate her plans (oftentimes in response to hypothetical and counterfactual scenarios). Unlike the qualifications for which they stand in, the placeholders are psychological operative. As such, they manifest the agent's persistent sense that the content of her intentions is in principle always subjected to further articulation. The psychological operation of the placeholders is a matter of the agent's readiness to partially replace them with more specific qualifications when a more detailed and explicit articulation of her plans and projects is called for. (By analogy with the previous remarks about the role of specific qualifications, the placeholders' psychological role is in part a matter of the regulative ideal of intentions (p. 721). As rational planning agents, we are required to spell out yet unarticulated qualifications of the deep structure of our plans in response to relevant circumstances and appropriate challenges. Hence, in our planning we are expected to be operate with a sense that our intentions come with the implicit generic placeholders, which we are ready to fill if called upon to do so.)⁶²

By default, when the generic placeholders are taken as given, they are in the agent's cognitive background. Thus, the epistemic version of (DS*) reads as follows.

(DS-E*) x intends to [GIVEN P_1 – P_m & FEASIBLE; PRE P_{m+1} – P_{m+n}]
(GIVEN C_1 – C_q & ADVISABLE; φ IF C_{q+1} – C_{q+r})

If the agent comes to doubt whether φ -ing is either feasible or advisable under certain circumstances, she is prompted to find out which hitherto unarticulated specific qualifications would bear on φ 's feasibility or advisability under those circumstances. Once articulated, these qualifications no longer fall under the extension of the generic placeholders. Some of them might also move out of the cognitive background, if the agent can no longer take them as given. The articulation, therefore, might not allay the agent's concern about the feasibility or advisability of her φ -ing, but it would give it a more definite and manageable form. For the agent would come to know something more definite about the preconditions and conditions that must obtain if she is to go ahead with her φ -ing. On the basis of this knowledge, she can then take concrete steps in coordinating this intention with her other plans (including her epistemic activities directed at finding out whether the preconditions and conditions obtain).

If the placeholders were not by default under the scope of the GIVEN, the agent's intentions would be marred by a persistent but unspecific doubt about their feasibility or advisability. This would undermine their role as planning attitudes. Because the agent would be left without any clear guidance about what to expect about her future conduct and no determinate framework for the coordination of her projects. She could only rely on the useless expectation that she is going to φ *were it not* for the obtaining of *yet unspecified* eventualities that would make her φ -ing either impossible or inadvisable. Hence, the fact that by default the generic placeholders are taken as given does not mean that we are blind to the possibility that φ -ing might turn out to be either impossible or unacceptable if certain yet unarticulated conditions do not obtain. It is precisely because we are open to this possibility that we accept the feasibility and advisability of φ -ing in the form of *generic placeholders*. These placeholders stand for our readiness, on the face of legitimate suspicions, to articulate more specific qualifications and to determine their epistemic status.

These considerations might not be sufficient to allay all worries about the role of generic placeholders. According to my analysis, it seems that generic qualifications such as, 'If I can,' 'If there is still reason to,' 'If I do not change my mind,' are necessary constituents of the deep structure of intentions. But appending these qualifications to a statement of intention might be problematic. For instance, qualifications such as, 'If there is still reason to,' 'If I do not change my mind,' etc. are troublesome if they license whimsical changes of mind. This is not, however, the correct reading of the 'if ADVISABLE' placeholder. The placeholder stands for the agent's acknowledgement that her intentions come with yet unarticulated conditions that, if taken to obtain, would give her a *good enough* reason to φ .⁶³ It does not suggest that the advisability of the action is ultimately a matter of the agent's arbitrary election, nor that there is any vacillation in the agent's resolve.⁶⁴

What about the explicit uses of clauses such as 'If advisable,' 'If I do not change my mind,' 'If nothing alters,' 'If I can,' etc. in everyday communication of intentions? Depending on the context, the qualifiers perform a variety of different functions. Usually, they are not meant to make explicit the generic placeholders in the deep structure of intention. Rather, they might be used to signal the existence of specific conditions or preconditions that the agent does not yet take to obtain but whose specific nature she does not want to divulge; to convey the agent's hesitation about the future advisability or feasibility of the intended action; to cancel standard conversational implicatures of the avowals of intention (e.g., to weaken the assurance normally generated by avowals of intention); or to remind the audience of one's fallibility. Because of the variety of functions that these if-clauses can play in overt expressions of intention, we should thus be cautious about relying on their ordinary usage to cast light on the deep structure of intention (in particular, we should not assume that these if-clauses

necessarily play the role of the generic placeholders ‘If FEASIBLE’ and ‘If ADVISABLE’).⁶⁵

The uses of the generic if-clauses in ordinary conversation I have just illustrated are unproblematic. But there are cases in which appending a generic if-clause of feasibility or advisability seems to undermine one’s avowal of intention. First-personal expressions of intention such as, ‘I’ll φ unless I forget,’ ‘I’ll φ unless I change my mind,’ or ‘I’ll φ unless I am weak-willed,’ are troublesome if the agent uses the if-clauses to qualify her avowal of intention by relieving herself, in the very act of avowing, of the responsibility to meet the demands for the stability of her intention—a responsibility that she is supposed to undertake in the very act of adopting an intention. One cannot sincerely and genuinely avow the intention to φ if one is not committed to making earnest efforts to stick to the intended course of action in the face of the hurdles and interferences produced by temptation, forgetfulness, akrasia, and fickleness; the hurdles and interferences produced by the ‘flighty mind’ of the human animal, as Nietzsche would put it.⁶⁶ Hence, if-clauses such as, ‘unless I forget,’ ‘unless I change my mind,’ ‘unless I am weak-willed,’ do not generate genuinely conditional intentions. If anything, by attempting to qualify the avowal of intention by putting restrictions on the agent’s commitments, they undermine the very adoption of the intention.⁶⁷

Notice that the problem arises only if these qualifications are entered ‘in the very same breath’ as the first-personal avowal of intention. There is no problem if they are meant to issue warnings about the intentions that are distinct from the avowals of them and that are not meant to weaken the agent’s commitments. One might use them as reminders of the limitations and frailty of the human will in general, with no suggestion of any specific threat to one’s resolve and no rejection of full responsibility for one’s undertaking (compare the hedging of assertions with generic disclaimers such as ‘If I am not wrong’). In other cases, one might use an if-clause to warn the audience about some determinate but unusual threat to the stability of one’s intention, in spite of one’s best efforts to thwart them and to live up to one’s responsibility for it. For instance, if I am unusually prone to forgetfulness, I might hedge my expression of the intention to φ with the clause ‘If I do not forget.’ In doing so, I am not qualifying the content of my intention, nor am I undoing the avowal of it.⁶⁸ Rather, in uttering, ‘I’ll φ if I do not forget,’ I am combining the avowal of the unqualified intention to φ —which is offered from the first-personal stance of commitment—with a warning about the likelihood of some specific failures of intention—which is offered from the third-personal stance of prediction.⁶⁹

9. Conditional Intentions and Agency

It is no accident that the deep structure of our ordinary intentions is conditional. We hardly ever intend to do something literally ‘no matter what.’⁷⁰

Why is this so? And what does this tell us about the nature of our agency? Consider preconditions first. It should not be surprising that by default the deep structure of intentions makes room for preconditions. Only an omnipotent agent could have intentions that are not qualified by enabling conditions. Our intentions can be without any enabling condition only at the time of action when they become *executively* unconditional (see p. 715). This is however only a temporary (and non guaranteed) stage in the unfolding of plans, whose deep structure remains preconditional. What about restrictive conditions? Adopting a strictly unconditional intention to φ is to be set on φ -ing *no matter what*, on φ -ing under all circumstances, provided that φ -ing is possible. A pure unconditional intention is utterly unresponsive to changes in the agent's situation. This inflexibility is potentially quite troublesome. For the agent is set on φ -ing—literally—at *any price*. An unconditional intention might thus require the agent to engage in the most obstinate and reckless conduct.

This danger is compounded if the agent's entire existence is organized around pure unconditional pursuits. Consider Don Giovanni. He appears to be set on the pure unconditional pursuit of the pleasures of food, wine, and seduction. This explains the problems that Bernard Williams finds with Don Giovanni's conduct, with the single-minded determination and reckless neglect of the effects of his conduct on himself and others. In this sense, Don Giovanni can be said to be living an unconditional life (see p. 700).

This kind of existence faces two major problems. First, its inflexibility exposes the agent to the risk of paying too high a price for his pursuits, including the loss of his own life or the fall of Heavens. Second, if the agent is set on a multiplicity of unconditional pursuits, how is he going to handle conflicts between them? Because of their unconditional structure, unconditional plans offer no resources to handle these conflicts. Each plan demands, with equal force, to be pursued at the expenses of any other one. Pure unconditional plans constantly expose agents like Don Giovanni to the dangers of either ineffectiveness or disintegration. These plans do not engage the norm of aggregativity. This means that, strictly speaking, Don Giovanni does not act out of full-fledged *intentions* but rather out of unconditional 'settled objectives' (p. 717). That is, he acts out of conative attitudes that, unlike genuine intentions, are not responsive to rational pressures for consistency and agglomeration.⁷¹

The pursuit of unconditional settled objectives resembles the relentless operation of a natural force. This relentlessness seems to explain the attraction that—as Williams suggests in the opening quote—we feel for the vitality of Don Giovanni's life. But it also shows that this vitality can be achieved only at the price of Don Giovanni's reckless obstinacy. As Williams writes, "Those who survive Giovanni—not only the other characters, but, on each occasion that we have seen the opera, ourselves—are both more and less than he is: more since the conditions *on* humanity, which we accept, are also the

conditions *of* humanity, and less, since one thing vitality needs is to sustain the dream of being as free from conditions as he is.”⁷²

The limitations and risks inherent in living unconditionally, Don Giovanni's style, show why the deep structure of full-fledged intentions by default makes room for restrictive conditions. Pure *unconditional* intentions suit only agents who never face the need to coordinate and integrate their conduct both at a time and over time. These agents can be said to have genuine intentions but only in a trivial sense: their planning attitudes are guaranteed to engage the norm of aggregativity given that these agents never encounter circumstances that demand actual aggregation. Pure unconditionality suits only a thoroughly ‘atomistic’ agency; one in which each single project is necessarily isolated from any other project, so that the different pursuits never impinge on each other's feasibility and advisability.⁷³ But we are not atomistic agents.

This atomism might be exemplified by agents with extremely simple motives and/or a very narrow temporal horizon, such as time-slice agents. We are not agents of this kind. Against the background of a complex, interconnected, unpredictable, and often hostile world, we engage in a variety of temporally extended projects that are isolatable from each other only under special and temporary circumstances. By default, we face coordination and integration problems both at a time and over time. Our plans have to be adaptable and opportunistic so that we can both take advantage of the changing circumstances of action and avoid conflicting pursuits. For agents like us, a systematic and sustained engagement in a multiplicity of pure unconditional pursuits carries the risk of massive frustration and ineffectiveness, not to mention the attendant danger of causing serious harms both to ourselves and to others. Hence, full-fledged intentional agency requires intentions with a deep structure like (DS-E*). That is, a deep structure that paradigmatically includes not only preconditions and restrictive conditions but also the generic placeholders, so as to secure both the sensitivity to the demands of aggregation and the flexibility of potentially open-ended articulation. These are the features that make our life as intentional agents profoundly different from (and possibly preferable to) a Don Giovanni-style unconditional existence.⁷⁴

These considerations do not prove that Don Giovanni is utterly unable to negotiate his way in this world by relying only on unconditional settled objectives. Nevertheless, given that his agency has both a complex motivational structure and an extended temporal horizon, Don Giovanni needs a mechanism for securing some kind of integration and flexibility in his pursuits. He needs some surrogate for the work that, for agents like us, is performed by the *internal* qualifications in the deep structure of our genuine intentions.

This surrogate might be offered by a system of checks and balances *external* to the content and structure of Don Giovanni's unconditional settled objectives. This system would work by switching his unconditional pursuits

on and off in response to appropriate triggers. For instance, each individual unconditional pursuit might be activated or inhibited by the success or failure of certain other pursuits. Alternatively, each pursuit might lose out to 'stronger' ones in the competition to gain access to the agent's limited focus of attention or to the scarce resources for its implementation. Perhaps each pursuit might simply be dropped or forgotten if it is not completed within a given timeframe.

A psychology that combines mechanisms of this kind might become quite complex (for instance, it might include both intersecting and hierarchical layers of activation and inhibition) and it might be able to curtail many conflicts among unconditional settled objectives and to avoid the most serious dangers of their dogged pursuits. This should not be surprising. A psychological architecture of this kind might underpin the agency of non-human animals and secure the kinds of stability, integration, and temporal extension distinctive of animal agency.

Nonetheless, this architecture can only secure the *external* unification of agency. The flexibility and the integration of agents *à la* Don Giovanni is not produced by features internal to the very structure of their practical attitudes.⁷⁵ The difference between internal and external integration is reflected in the handling of conflicting pursuits. Incompatible settled objectives have to be eliminated or deactivated; they are to be defeated, overpowered, preempted, dropped, inhibited, or forgotten. The mechanisms for the elimination, deactivation, or reactivation of the pursuits work outside of the agent's practical reasoning and planning. The agent does not have to take them into account in determining what she is to do. These mechanisms are simply *triggered* by the appropriate circumstances or by actual conflicts between unconditional settled objectives. Don Giovanni's agency can achieve a certain degree of flexibility and integration by yielding to *brute psychological* pressures that his conative attitudes might exert over each other.

By contrast, genuine intentions respond to *rational* pressures for consistency and agglomeration that operate from within their own content and structure. When a particular intention turns out to be either inadvisable or incompatible with other plans, this intention does not have to be *counter-acted*. Rather, the intention either stops to exert its influence because the agent acknowledges that it is moot, or is accommodated by the uncovering of yet unarticulated conditions. In either case, this is something that occurs within the agent's practical thinking and planning: the intention yields to the demands of overall integration as part of its own operation *qua* the kind of rational attitude that it is.

The internal integration does not depend simply on the existence of specific qualifications in the deep structure of intentions. An indispensable contribution is given by the generic placeholders and their role in the articulation of the content of the intention and the responsiveness to the demands of agglomeration. When one adopts the intention to (φ IF C), one is not aiming at

φ -ing when C obtains *and no matter what else*. Don Giovanni's conduct could still be dangerously single-minded and recklessly inflexible, and thus eschew any internal integration, if he were simply to qualify each settled objective 'to pursue φ ' so that it becomes 'to pursue φ IF C, *percat mundus* otherwise'. A partially but inflexibly conditional existence of this sort does not seem much of an improvement over a purely unconditional one. This existence still fails to engage the demands for aggregation. It still suffers from structural deficiencies and the attendant threats to the effectiveness and the unity of the subject's agency. Its success and unity still depend on the brute psychological operation of mechanisms of external integration. This is because this kind of existence lacks the contribution of the potentially open-ended articulation of the internal conditions of the deep structure of intention.

Because of the possibility of open-ended articulation, an important part of our practical thinking is devoted to the uncovering of specific qualifications within the deep structure of our intentions in response to challenges to the advisability, feasibility, and mutual compatibility of our plans. Much of our planning activity is devoted to avoiding seeming conflicts among our projects by articulating their restrictive conditions and taking advantage of their flexibility in securing mutual compatibility.⁷⁶ In the rational order of planning, giving up an intention altogether, as common a phenomenon as it is, might only come as last resort.⁷⁷ Even so, when the agent gives up the intention, she engages with it as a judgment-sensitive attitude rather than as a brute psychological phenomenon that needs to be defused or counteracted. She gives up the intention by (and in) her acknowledgment that the intention is either impossible, inadvisable, or non-agglomerable.

The mechanisms of external integration might secure some remarkable degrees of unified and effective conduct (as exhibited by non-human animals, for instance) but they cannot aspire to the same unity secured by internal integration. First, they cannot secure the kind of flexible integration made possible by two features of internal integration: the direct engagement with the norms of aggregativity and the possibility of open-ended articulation. Second, the unity produced by external integration is ultimately adventitious. It is a mere by-product (although one that might be either designed or selected for) of the combined operations of otherwise isolated settled objectives. This is different from the kind of unity to which the agent is directly sensitive (although usually unreflectively so) in the very exercise of her intentional agency because of its constitutively necessary engagement with the demands of agglomerativity.

If, as I have argued, the deep structure of intentions is paradigmatically conditional, what role is left for unconditional plans? To begin with, there are situations in which we behave like Don Giovanni. Situations in which we are unconditionally set on φ -ing as a matter of having an unconditional settled objective to φ . This might be a good thing as long as the mechanisms of external integration provided by our animal agency secure an acceptable

surrogate of genuine internal integration. But these behaviors are not instances of genuine unconditional intentional agency. To find genuine pure unconditional intentions we must rather look at those situations that allow for the pursuit of individual plans in utter isolation from each other. That is, for situations that allow for the exercise of a sort of atomistic agency. These circumstances, however, are rather unusual, and at best only temporary and local. This is why most genuinely unconditional pursuits occupy only a marginal role in our agency.

Pure unconditional intentions are not to be confused with plans that are pursued 'at all costs' in the sense that one is willing to die in order to succeed at them. A plan of this sort need not be literally 'no matter what.' For the agent might not be willing to φ if this were to cause, say, the loss of other lives or the fall of Heavens. The deep structure of a plan that demands extreme self-sacrifice is often a conditional one. Nevertheless, there is a sense in which a plan of this kind can be said to be 'categorical.' It is similar to what Bernard Williams calls a 'categorical desire,' a desire that—as he writes—does not 'operate conditionally on [the agent's] being alive, since it itself resolves the question whether he is going to be alive.'⁷⁸ A 'Williams-categorical' intention to φ does not operate conditionally on the expectation of the agent's being alive as a result of his successful φ -ing, but it itself resolves the question whether he is going to be alive as a result of that pursuit.

This is not to deny that, even outside of temporary pockets of atomistic agency, an agent might intend to φ literally *at all costs*, including the end of the world. Such an unconditional intention, however, is in order only for an *ultimate* end. It is no surprise that one is to pursue such an end at all costs. For it is by reference to it that the agent determines the highest possible price that she is ever willing to pay in her acting, namely, anything but her failure to φ . Unconditional intentions, however, are not in order for the more specific plans which are subordinated to the pursuit of the ultimate end and face the demands for mutual internal integration. If we set aside ultimate ends, the paradigmatic structure of our ordinary plans is therefore conditional by default.

In any event, it does not seem to be a requirement of agency that there always be a single ultimate substantive end to be pursued unconditionally. For instance, for a value pluralist the pursuit of each distinct value is likely to be unsubordinated to the pursuit of any other value and thus, in some sense, ultimate or categorical. Nonetheless, no pursuit of this kind is strictly speaking unconditional given the possibility of conflicts between incommensurable values. The agent might thus have a multiplicity of ultimate ends, each to be promoted under all circumstances although only *ceteris paribus*.

Even if there is no single ultimate substantive end there is always room for a particular kind of pure unconditional intentions. This is because the constitutive standards of agency (whatever these might turn out to be) can only be engaged unconditionally.⁷⁹ No agent can intend to violate these

standards in the very exercise of her intentional agency. (For instance, if we assume that acting under the guise of the good is constitutive of intentional agency, then an agent cannot knowingly intend to act—at least presently—under the guise of the bad, not even conditionally so.) This is not to deny that one might intend to bring about the future demise, whether temporary or permanent, of one's own agency. But this plan must still be carried out by continuing to engage the constitutive standards of agency as long as necessary to complete the project. In this sense, the standard of agency are *inescapable*, and any agent, in her present exercise of full-fledged intentional conduct, cannot but engage these standards in a purely unconditional manner.⁸⁰

The deep structure of our ordinary intentions is conditional. Pure unconditional intentions are rare. They are called for either under those special circumstances that allow for temporary and local atomism of agency or in connection with the most general characterizations of our agency, in terms of either a substantive ultimate end or the constitutive features of agency. The latter kind of intentions raise a distinct set of philosophical issues that goes beyond the scope of this paper. But rest assured that I intend to investigate them—if *feasible and advisable*—on another occasion.⁸¹

Notes

¹ Williams (2006, p. 41); on the inevitability of Don Giovanni's end, see Williams (2006, p. 39).

² The only extended discussions of conditional intentions of which I am aware are Bratman (1979), Cartwright (1990), Yaffe (2004), and Klass (2009). Davidson (1978/1980, pp. 92–95) made some brief but important remarks about conditional intentions, later criticized in Velleman (1989, pp. 117–121), Bratman (1985/1999), and Velleman (1997, p. 46 fn. 27). Brief remarks about conditional intentions can be found in Bach (1984, p. 50). See Roth (2004: sec. 4) for a discussion of conditional intentions and shared agency.

³ For an overview of recent work on conditional obligations, see Hilpinen (2001); on indicative conditionals, see Edgington (2001; 2006).

⁴ The promise of these insights on the nature of our agency provides a sufficient reason to embark on a detailed investigation on conditional intentions. But even apart from this promise, there are at least two current debates in the philosophy of action that are likely to benefit from a better understanding of conditional intentions. First, there is the issue of cognitivism about intentions and the norms of practical reason. (For a recent discussion of cognitivism, see Bratman 2009b.) The cognitivists claim either that intentions are ultimately a special kind of belief or, more modestly, that the characteristic norms of intentions are grounded on norms of theoretical reason. The discussion of cognitivism is usually cast in terms of unconditional intentions. Can cognitivism offer an adequate account of the nature and norms of *conditional* intentions? If, as I will claim, the deep structure of intention is characteristically conditional, the cognitivist cannot elude this question. This might turn out to be an intriguing challenge for the cognitivist, especially if the conditional structure of intentions turns out to reflect what at least *prima facie* appear as distinctive features of intentions as *practical* attitudes. (See for instance the brief but interesting remarks in Cartwright (1990, pp. 252 ff.) about the apparent difference between conditional intentions and the corresponding conditional beliefs.) Second, there is the issue of joint agency. To the extent that joint agency constitutively involves conditional intentions such as 'I'll ϕ if you will' (see Velleman 1997, Roth 2004), it is imperative to make clear how the

qualification is to be understood. This is because there are many possible ways of interpreting it, all of which mean substantial differences to the kind of joint agency underpinned by conditional intentions. Finally, the discussion of conditional intentions bears directly is the questions about conditional intent and its relation to *mens rea*. This question figures prominently in the two more extended treatments of conditional intentions in the literature: Cartwright (1990) and especially Yaffe (2004).

⁵ For the distinction between external and internal conditions, see Cartwright (1990, p. 235). External conditions in Cartwright's sense are distinct from external conditions in the sense discussed by Gilbert (1993, p. 633) in connection with promises. Gilbert's external conditions correspond to the conditions of what I call 'suspended intentions' (see §3).

⁶ In the case of external conditions, from 'If C, x intends to φ ' one can infer the contrapositive 'If x does not intend to φ , $\neg C$.' Not so for the internal condition in 'x intends to: φ if C', see Austin (1956/1961, p. 214).

⁷ A further kind of conditional statement might qualify the avowals of intentions—the 'giving of the undertaking', as Austin says (1956/1961, p. 214). Consider the expressions, 'If you don't mind, I will pick up the bill', 'I will put the biscuits on the sideboard in case you want them later' (see Austin 1956/1961, p. 210). In these cases, the qualification does not concern the content of the intention or the acquisition of the intention. My intention is to pay the bill, whether or not you mind it. The 'IF' is one of courtesy, so to speak. Likewise, my intention is to put the biscuits in the cupboard, whether or not you are ever going to want them. The fact that you might want them later is a consideration for my adopting the intention to put the biscuits in the cupboard, but your actually wanting them is not a condition on the acquisition of that intention (as it would be in 'If you want the biscuits, I will put them in the cupboard') or an internal qualification of it (as in 'I'll put the biscuits in the cupboard IF you want them').

⁸ A third kind of internal conditions, which limitations of space prevent me from discussing in detail, are *subordination* conditions (see note 49 below).

⁹ Necessary truths do not qualify intentions restrictively, although they might appear as antecedents in the derivation of intentions in practical reasoning. Notice that the restrictive conditions C are not to be confused with the conditional clauses, if any, that might be used to analyze the structure of the intended *action* of φ -ing. For instance, Don Giovanni unconditionally intends to seduce women ('Purché porti la gonnella', act I.5—'As long as she wears a skirt'). The intended action could be analyzed as 'For all x's (if x is a woman, seduce x).' But being a woman does not seem a *restrictive* condition on Don Giovanni's seduction in the way in which a qualification such as 'if she is not married' would. Having called attention to this distinction, I must confess that I do not know where to draw it exactly. If there were no such distinction, one could end up with all intentions being conditional in the following trivial sense: what one intends is always to 'act' in a most generic sense, qualified by the 'if C,' where C gives the specification of what this acting consists in. However, it is not clear at which point the qualification counts as being part of the nature of the intended action, of the φ -ing, rather than a restriction on the content of the intention to φ . Consider the timing of the intended action. If I intend to eat an apple tomorrow, is the time specification part of the intended action or a restrictive condition on the time-generic intention to eat an apple?

¹⁰ In this paper, I assume that having the intention to φ (whether conditional or not) puts the agent under some rational demand to φ . I want this assumption to be as neutral as possible between different accounts of the exact strength and ultimate source of this requirement. All that I assume is that an agent who intends to φ is—everything else being equal and as long as she continues to intend to φ —under some rational pressure to φ .

¹¹ See Bratman (1992).

¹² Oftentimes the agent who professes the intention to φ if C thinks that φ -ing is inadvisable if $\neg C$, i.e., his intention is really to ' *φ only if C.*'

¹³ Davidson (1978, p. 94) claims that "bona fide conditions are reasons for acting that are *contemporary with the intention*" (my italics). This qualification is supposed to draw a distinction

between what I call restrictive and enabling conditions. But I do not think that it works. Davidson claims that “not being able to go home early is not a reason for or against [the agent’s] going home early.” It might be true that, strictly speaking, there is no such thing as his going home early, given that he is unable to do so. But this does not show that the agent cannot intend to go home early and that he is given a decisive reason against so intending when he learns that he cannot go home. Compare the case of belief. I cannot believe something that, by my own lights, is false. Prior to learning that p is false, I believe that p . Imagine that certain conditions obtain that give me conclusive evidence against p . These conditions are thus a good enough reason not to believe that p . The conditions that make it impossible for me to φ appear to play exactly the same role in my rejecting the intention to φ , as the conditions that give me reason not to believe that p . In any event, the notion of a ‘reason contemporary with the intention’ does not really do any substantive work in Davidson’s theory in a positive characterization of restrictive conditions. It actually gets Davidson in trouble, since he ends up denying that preconditions are conditions at all (they are ‘bogus’ in his terminology), which gets his view into some serious troubles down the line (see note 51 below). Thus, I agree with much of Velleman’s criticism of Davidson (Velleman 1989, pp. 117–121). I do not think, however, that Davidson was ‘deeply confused,’ as Velleman (1997, p. 46 fn. 27) puts it, in trying to distinguish between what I call preconditions and restrictive condition. His mistake is rather that of claiming that preconditions are ‘bogus.’

¹⁴ The interpretation of the clause ‘C can no longer obtain’ varies depending on the different ways in which the timing of C might qualify the intention to φ —for instance, C is to obtain at a specific time t , by a specific time t , for a certain interval at/by a specific time t , at some indefinite time, etc. (notice that the φ -ing can be temporally qualified in similar ways and that all combinations are in principle possible; for instance, I might intend to φ at some indefinite future time if C obtains for a certain interval by a specific time t). It might also be possible to intend to φ at a time *earlier* than the obtaining of C. In this case, however, one is not really intending to φ conditional on the obtaining of C but on one’s *prediction* that C is going to obtain at a later time. The question then becomes what the conditional intention demands of the agent if by the time of the intended action she is not warranted in predicting that C is going to obtain in the future.

¹⁵ See Edgington (2001, p. 409).

¹⁶ This is why it would be ludicrous to suggest that one might want to undertake projects conditional on most unlikely circumstances in order to improve one’s ability to carry out long-term plans and to strengthen one’s resolve and steadfastness.

¹⁷ This is not to deny that in some cases one might intend to make a material conditional true, i.e., to intend that $(C \supset \varphi)$. (For instance, if one intends to be meet static rational requirements of consistency among one’s attitudes, the object of one’s intention is to satisfy a material conditional.) The intention to make $(C \supset \varphi)$ true, however, is not a conditional intention. Although one might intend conditionally to satisfy a material conditional, that is, one might intend to (make $(C \supset \varphi)$ true IF D). The ability to distinguish between these two kinds of intentions would be lost under the truth-functional interpretation of conditional intentions.

¹⁸ Cartwright (1990, p. 239) seems to have this in mind with his talk of the ‘cancellation’ of the intention. Notice that if the intention is moot but the agent does not know it (since she does not know that C can no longer obtain), she might still be warranted in having that intention. If so, until she realizes that the intention is moot, the agent is still under the rational requirements that apply to her in virtue of her having that intention. The case is the analog to the agent’s being under the rational requirements imposed by her believing that p even if, unbeknownst to her, p is false.

¹⁹ A further problem with the truth-functional interpretation is that it gives an implausible account of what happens when φ -ing turns out to be impossible. Consider Don Giovanni’s intention to ‘wear Leporello’s clothes IF he is going to try his luck with Elvira’s maid’ (see Act II.1, ‘Vò tentar la mia sorte’). Imagine that, contrary to what happens in the Mozart’s opera,

Don Giovanni does not run into Leporello at the beginning of Act II and, thus, he is unable to wear Leporello's clothes. (This is a variation on the example of a conditional imperative in Vranas (2008, note 21).) According to the truth-functional reading, once it becomes impossible for Don Giovanni to wear Leporello's clothes, the only way in which Don Giovanni might carry out his intention is by giving up on his attempt at seducing Elvira's maid. But this is counterintuitive. If Don Giovanni cannot find Leporello's clothes, it is more natural to say that the conditional intention is no longer applicable, given that the action of φ -ing—his wearing his servant's clothes—has turned out to be impossible. At that point, the conditional intention has nothing to say about Don Giovanni's planned seduction. Don Giovanni is to go back to the drawing board. If attempting the seduction of Elvira's maid is still a reasonable idea in spite of his inability to arouse her appetite by presenting himself in Leporello's clothes (see Act II.1), Don Giovanni violates no rational demand if he decides to attempt to seduce her while wearing his own hat and cloak. His decision amounts to the acquisition of a new intention that is not in conflict with his original conditional intention, since the latter has turned out to be moot.

²⁰ See Vranas (2008; cf. note 10 for extended references to the literature).

²¹ The intention is not moot when C does not *yet* obtain. As long as the agent takes that there is a possibility, however remote, that C might obtain, the agent is under all the requirements the conditional intention imposes on her. However, if the possibility is truly remote, the agent might *accept* in Bratman's sense of the term (see note 11) that C can no longer obtain. When so, the intention has become moot for her. In a sense, by accepting that C is no longer to obtain, the agent has made the intention moot but not as if she had actually made C false.

²² Once the agent has revoked an intention, she is no longer under any of the rational pressures imposed by that intention since the intention has disappeared. However, if the revocation is in violation of stability requirements, the agent might be under a rational demand to re-adopt the revoked intention and thereby subject herself, once again, to the intention's distinctive rational pressures.

²³ The expression 'precautionary plans' is taken from Bratman (1979, p. 280).

²⁴ For a criticism of my rejection of the truth-functional interpretation, see Klass (2009).

²⁵ See Edgington (1995, p. 288; 2001, pp. 409–411), see also Bennett (2003, pp. 124–125).

²⁶ Edgington (1995) argues that to have a conditional desire for 'p if C' is a matter of having a preference for p&C over \neg p&C; a preference that is nonetheless *silent* about the ranking of \neg C over \neg p&C. This claim accords with my view of conditional intentions. Having a conditional intention to (φ IF C) is in part a matter of having a conditional desire for ' φ -ing if C' in Edgington's sense rather than a desire for the material conditional (φ -ing \supset C). That the conditional intention is silent about the ranking of \neg C over ($\neg\varphi$ -ing & C) is compatible with the fact that the agent might intend to make C false while sincerely holding an intention to (φ IF C), see p. 717.

²⁷ Edgington (2001, p. 409).

²⁸ See Bennett (2003, p. 125).

²⁹ In a terminology to be introduced in §0, this is a case of intending to (GIVEN C, φ).

³⁰ A suspended intention would be an *unconditional* intention to φ that does not exert any rational pressure prior to the obtaining of C. Although the idea of a suspended intention is not inconsistent, suspended intentions seem to be an anomaly, if anything, because it is unclear what their rationale is. To see why consider the following. The content of suspended intentions is not internally conditional. The condition C is external. This is not a condition on the *acquisition* of the intention (as for external conditions see p. 701), but on the *activation* of its force. Once the agent acquires the suspended-until-C intention to φ *simpliciter*, the agent is not yet under the rational pressures generated by the intention to φ , such as the demand to take the known necessary means to φ -ing. She might have prudential reasons to take these means prior to C in mere *anticipation* of the rational pressures to which she is going to be *automatically* subject once C obtains. What happens if, prior to C, the agent does not take a known necessary means for her future φ -ing? There is no irrationality at that time. What about at the time when C obtains?

At that point, given that φ -ing is now impossible, it seems inappropriate to deem the agent irrational. At most, she could be chastised for her imprudence. This is, however, a much weaker rational pressure than the one exerted by an already active intention. Moreover, it is unclear what difference a suspended intention makes with respect to the case in which the agent, prior to C, simply anticipates that she is going to intend to φ , rather than actually adopting a suspended intention to do so. (Similar worries of instability might arise about what Gilbert (1993, p. 633) calls externally conditional promises, first-person statements such as “On condition that C, I promise to φ .” These promises appear to be ‘suspended’ in that, as she says, they generate appropriate promissory obligation only when, if ever, C obtains. What makes them unstable is that it is not clear what kind of assurance they can offer. Prior to C, the promisor is under no obligation to make sure that she is able to φ once C obtains, and it is not clear in which sense, once C obtains, she can be criticized for failing to secure her ability to φ at the earlier time. If I am the promisee, I’d rather get from you an immediately active internally conditional promise to (φ IF C) rather than an externally, suspended-until-C promise to φ . And if you are unwilling to bind to the internally conditional promise, there seems to be no reason why you should give more than a *prediction* that, when C obtains, you are going to promise me to φ .)

³¹ For instance, after I buy a lottery ticket, I might start planning what I will do if I win the multi-million jackpot. Rather than just fantasizing about it, I might genuinely lay out a plan and commit to it. But prior to the announcement of the winner, there might literally be nothing that I can do to promote the actions that are conditional on my winning. I might not be able to afford these actions now and it might even be impossible or unreasonable for me to prepare for them. Likewise, none of my other current plans are in principle incompatible with the intentions conditional on my becoming filthy rich. For if I win, I will be able to pay whatever penalties and cost might be associated with giving up my present intentions.

³² A full analysis of the relation between assenting to the claim ‘If A, I’ll φ ’ in practical reasoning and adopting the internally conditional intention to (φ IF A) must be left for another occasion. One of the interesting questions is the relation between the conditional connective in reasoning and the ‘If’ in the content of the intention.

³³ See Cartwright (1990, pp. 237 ff.).

³⁴ See Davidson (1978), cf. Velleman (1979, pp. 117–121).

³⁵ The norms are first formulated in Bratman (1987). For a more recent discussion, see Bratman (2009b).

³⁶ Bratman (2009b).

³⁷ Although it is in principle impossible to intend both to φ and to $\neg\varphi$ unconditionally at the same time, the agent might be able to agglomerate the intention to (φ IF C) with the intention to ($\neg\varphi$ IF \neg C). She might be able to do so as long as the necessary preparatory work for the first plan does not interfere with the preparatory work for the other. This is just a special case of the possibility to agglomerate the intention to (φ if C) with the intention to (ψ if \neg C) in case the preparatory work of the first intention does not interfere with that of the latter (see p. 716 for more on agglomerativity and preconditions).

³⁸ Because of the restrictive qualification, conditional intentions might be subject to a norm concerning the evidence for C. It might be the case that if I intend to (φ IF C) I cannot deliberately shield myself from being exposed to evidence about C, especially if the evidence is readily available and not to expensive to acquire, even if I might not be required to be actively investigating about C. (When properly qualified, this requirement might be turns out to be rather weak. Consider the following scenario, suggested by Alec Walen (in conversation). I intend to give you \$5,000 if your business goes under. I promised you. But meanwhile, our friendship has gone sour. I am no longer so happy to give you the money. I am not ready to break my promise, but I want to make it hard for you to reach me just in case. So I change my email address and don’t give you the new one. If you find me and tell me that your business failed, I’ll help you as I promised. But why should I not make it at least a bit hard for you to find me and tell me?)

³⁹ This is not to deny that there are going to be small and obvious differences in the details of *implementation* of φ -ing because of the need to respond to the details of the circumstances at the time of action.

⁴⁰ These considerations concern the *reasonable* allocation of scarce resources in planning and preparing for various contingencies. One can be most unreasonable in this allocation without violating the characteristic norms of intending. For instance, an agent might be wholeheartedly, single-minded, and even obsessively devoted to the conditional intention to (φ IF C), being utterly indifferent about what she is to do in case C does not obtain even if she knows that C is most unlikely (consider agents who are waiting for the unlikely return of a loved one, such as a soldier missing in action or a patient in a deep coma, as in Luigi Pirandello's short story *La Camera in Attesa*; thanks to Ermanno Bencivenga for the reference).

⁴¹ For instance, see the talk of 'degrees of commitment' in the discussion of conditional intent in Yaffe (2004).

⁴² See Cartwright (1990, pp. 245–246).

⁴³ On the issue of conditional intent and *mens rea*, see Yaffe (2004).

⁴⁴ In a similar fashion, if the agent has some power over C, her attitude toward it shows something about her stance toward φ -ing. For instance, if Masetto intends to (beat Don Giovanni IF he tries to seduce Zerlina), there is a significant difference between this being a non-precautionary project, so that Masetto actually tries to make sure that the seduction occurs, rather than a precautionary plan, so that Masetto hopes that the seduction never takes place and tries to prevent it.

⁴⁵ Cf. Klass's use of the expression 'elective conditions' to refer to what I call 'restrictive conditions,' Klass (2009).

⁴⁶ See Davidson (1978/1980, p. 93).

⁴⁷ Preconditions cannot be temporally qualified in the same ways as a restrictive condition. A precondition P is to obtain no later than necessary in order to enable the φ -ing at the intended time of action. The timing of a restrictive condition is, instead, an elective matter. When imposing a restrictive condition C, the agent is also to determine, among a range of options, the timing of its obtaining (see note 14).

⁴⁸ For instance, in the example of the biconditional intention discussed in the previous section, Donna Anna might secure the agglomerativity of her plans by reformulating them as follows: I intend to (be at the party IF, by the time I have to leave home, to the best of my knowledge, Ottavio is going to be at the party) and I intend to (not be at the party IF, by the time I have to leave home, to the best of my knowledge, Ottavio is not going to be at the party). In this case, one of two intentions is going to be moot before Anna is going to take any preparatory step that forces her to repudiate one of the intentions. But there is no guaranteed that a similar strategy is always available.

⁴⁹ When φ -ing is subordinated to ψ -ing, we give *lexical priority* to the latter, primary pursuit. As far as the φ -ing is concerned, we are expected to comply with the norms of intention only in the space left, if any, once we have taken care of the demands of ψ -ing. In many cases, it might be that we never get around doing anything to promote the subordinated pursuit, since all our resources are absorbed by the primary plans. Once the fate of the primary plan is settled, we might be required to pursue the subordinated plan, *provided* that this is still possible. (Whether we are so required depends on the nature of the subordination. In some cases, the φ -ing is subordinated to the *failure* of the primary plan; φ -ing is the 'plan-B.'). Subordinated intentions are still intentions; they are not mere endeavorings or settled objectives, in Bratman's sense (1987, ch. 9) and (2009a). They are still sensitive to agglomerativity; they do so *by* their subordination. And they are sensitive to this demand from the time they are adopted, although this might mean that no other steps need to be taken until S is fulfilled—which makes them superficially resemble alleged 'suspended intentions' (see §3). The conditions of subordination are internal to the intention. They look like restrictive conditions C but they play a different role in the deep structure of intentions. This structure should thus be symbolized

as follows, where S stands for the subordination conditions: x intends to [SUB S][PRE P](φ IF C); which makes clear that a restrictively conditional intention is not necessarily a subordinated one. Conditional intentions can play the role of primary plans, although all subordinated plans are internally conditional (at least in that they are qualified by S). Finally, sometimes conditional expressions such as ‘I will call you if my appointment is cancelled’ only express external conditions on intentions that we are not yet adopting. They only indicate that we are *disposed* to φ in case the condition obtains. Likewise, sometimes we talk of having contingency ‘plans’, which however are neither full-fledged intentions nor merely subordinated ones. In these cases, we only mean that we have readily available *recipes* that we expect to follow if the contingency obtains. (Notice that many precautionary plans (note 23) are also subordinated plans, subordinated to our irremediable failure to avoid C. I say ‘irremediable’ since we might not be committed to φ -ing when C first obtains, but only when C cannot be permanently avoided.)

⁵⁰ See Bratman (1987, ch. 9) and Bratman (2009a). The expression ‘endeavoring’ is introduced by Bratman (1987), ‘settled objectives’ by McCann (1991).

⁵¹ The role of preconditions to secure agglomerativity is what underlies Bratman’s sound criticism of Davidson’s account of conditional intentions (Bratman 1985/1999, pp. 217 ff.). Bratman argues that, given Davidson’s claim that only restrictive conditions are genuine, it follows that for Davidson an agent might happen at the same time to have two unconditional intentions to pursue two courses of action that the agent already knows to be incompatible. This outcome violates agglomerativity. Let me point out that one of Davidson’s premises is correct. In determining whether φ -ing is choiceworthy, one is to consider the scenarios in which φ -ing turns out to be possible. This is however only a preliminary step, which establishes only a negative point. If φ -ing is not choiceworthy in case possible, there is no point to try to pursue it if one does not yet know whether it will turn out to be possible. From this one should not derive an unconditional intention to φ (or better an *unpreconditional* one, in my terminology) by assuming that P holds. This is, however, what seems to follow from Davidson’s account, given that he fails to recognize that preconditions are genuine internal qualifications of the intention, although of a different sort from the restrictive ones.

⁵² An intention ‘to [PRE P] φ ’ does not tell the agent what to do in case P can no longer obtain. Unless the agent knows that that P holds, she might be under some rational pressure to consider and possibly prepare for scenarios in which φ -ing is impossible. The extent of this pressure varies from case to case (depending on the likelihood of P, the importance of φ -ing, and the costs of contemplating and preparing for the different options). Under this respect, preconditions behave very much like restrictive conditions (see p. 712).

⁵³ It seems to me that if a plan is not understood to be a precautionary one, often there might be implicit restrictions on the agent’s use of her control over C. In these cases, to commit to φ -ing IF C is also to commit not to exercise one’s control over C, or at least to commit to φ -ing *once* C obtains, even if it would still be in one’s power to make C false. When these restrictions are in place, the agent might have to make explicit that she intends to exercise her control over C, in case she wants to, otherwise by default she would be committed not to do so. However, these restrictions on the agent’s control over C do not derive from the conditional nature of the intention *per se* but from contextual features of that particular intention (as in the nurse scenario at p. 703).

⁵⁴ A very interesting kind of qualification on controllable preconditions P concerns the extent of the agent’s actual exercise of her power over P. Again, consider marriage proposals. Don Giovanni cannot marry Zerlina unless she freely consents. But he wants her to consent out of pure love for him rather than of any other consideration. As it happens, he could offer her a huge financial incentive to accept his proposal. But he wants to exclude the exercise of this power from the possible ways of securing her consent. In proposing to her, therefore, he does not intend—preconditional on Zerlina’s free consent—to marry her *no matter what it takes to convince her to consent to his proposal*. He rather intends—preconditional on Zerlina’s free

consent—to (marry Zerlina IF she accepts exclusively out of love for him). Likewise, he might qualify a controllable precondition to make sure that it is his own achievement, i.e., that it does not obtain spontaneously or by someone else's doing.

⁵⁵ For simplicity's sake, I am ignoring the differences that degrees of belief might make to the deep structure. I am also ignoring the distinction between taking C to obtain now or at some future moment but on time for φ -ing (see note 14).

⁵⁶ I think that the deep structure of intentions is akin to what Bratman (1979) was trying to account for in terms of his notions of 'core-intention' and 'intention map', although he did not make a distinction between enabling and restrictive conditions.

⁵⁷ See Taylor (1985: 36–37).

⁵⁸ My talk of the 'psychological reality' of the deep structure of intentions is not supposed to commit my view to any specific mental architecture. Whether this structure is isomorphic with the structure of actual causal mechanisms and thus qualifies as psychologically real in some stronger sense is a matter on which my view is meant to stay neutral.

⁵⁹ For a more extended discussion, see Klass (2009).

⁶⁰ Because of the limitations in our resources, information, and abilities, we adopt several strategies to simplify our dealings with the complexities of planning and the exactness of the ideal articulation of the deep structure of intentions. First, there is the attitude of acceptance, in Bratman's sense of the word (see note 11), which allows us to cut through fine distinctions in our cognitive attitudes (including degrees of belief) and proceed by taking certain conditions for granted as if we were certain of their occurrence. Second, in the interest of viable categorization and classification, we individuate actions and their circumstances in coarse-grained ways. Third, we rely on rules and policies, trading off responsiveness to the nuances of individual cases for computationally manageable satisficing procedures. Fourth, we reason on the basis of presumptions, defaults, and generics (see Gabbay and Woods, 2005) Last but not least, we benefit from the stabilizing effects of the framework that intentions themselves provide for practical reasoning (see Bratman 1987).

⁶¹ Because of the complex interaction in the factors of relevance and their contextual variability, we should be cautious about drawing conclusions about the deep structure of intentions by some unsystematic observations about ordinary uses of conditional expressions of intentions. (This seems to be a problem, for instance, with Davidson 1978.)

⁶² Many of the circumstances that fall under the unarticulated generic placeholders are so remote that there is hardly ever any point in bring them out explicitly. But they are still part of the deep structure of the intention, even if it is safe to expect that they do so only by falling under the generic placeholders. Our most fundamental convictions about the world, however, might not fall under any placeholder. If some of these convictions were shaken, it is not clear to me that we would think that our intentions turned out to be moot because most unlikely circumstances (that we never explicitly entertained in the content of our intentions) just happened to obtain, making our intentions either unfeasible or inadvisable. If our worldview is fundamentally shaken, we might simply be at a loss. We really have to start from scratch in figuring out what to do. Convictions that are so basic are part of a different background of intentions than what I called the cognitive background of intentions. These convictions do not figure in the content of our intentions. They are rather part of the background against which we go about figuring out what to do. Our intentions can be said to be *conditioned* on these convictions, in the sense that they owe their very existence to this fixed background. Had we had a different worldview, we would not have been thinking of acquiring those intentions in the first place. *Pace* Cartwright (1990, p. 236), these convictions are not just external conditions on the acquisition of the intention, in the sense of being circumstances that offer considerations for the adoption of some particular intention within a fixed worldview (p. 701 above). All that said, it is hard to draw a clear line between the ordinary cognitive background that is included within the content of the intention under the GIVEN and the deep convictions that form the worldview on which an intention is conditioned.

⁶³ This is not to say that there is room for some arbitrariness in the adoption of one's intentions. But this arbitrariness is warranted whenever the merits of the case are underdetermined. Whimsical changes of mind, however, are changes of mind that are utterly insensitive to the merits of the case, even when they are not underdetermined, on this issue see Ferrero (forthcoming). On a related issue, cf. Korsgaard's (2007, pp. 61–62) on the universalizability of commitments (an agent is committed to φ -ing in all relevantly similar circumstances, that is, 'unless there is a good reason why not', which is not to be confused with 'in the face of any consideration whatever').

⁶⁴ A further worry concerns the predictability of future conduct. Although the agent is not whimsical, it seems that in many cases, once the time of action comes, it might not be difficult for the agent to find some unanticipated feature of the situation that she could use to claim that her φ -ing is inadvisable. True, but notice that we often appreciate the importance of giving reliable assurance about our future conduct. What makes our conduct advisable at the time of action is in part a function of this concern. In many situations it is better to stick to φ -ing even if, had one earlier considered the possibility of present circumstances, one would have explicitly qualified one's intention to φ to exclude them thereby being under no demand and expectation to φ if they occurred. But since that explicit exclusion was not entered and made explicit the agent, the φ -ing might still turn out to be all-things-considered the advisable course of action. (The question of when exactly one should to stick to one's φ -ing despite unfavorable turns of events on account of the need to make oneself predictable can be quite contentious. It could be argued that the practice of promising as separate from the simple communication of intentions is needed to secure clearer requirements and expectations about one's future conduct even on the face of unfavorable turns of events. This is not to say that all contentious matters are thereby avoided, since there are still issues about the circumstances that relieve one of the promissory obligations.)

⁶⁵ The pragmatics of expressions of intention is further complicated by the following phenomena. First, even if we explicitly entertain several qualifications, we usually have only the time, resources and audience's attention for one or two explicit qualifications. Second, there are contexts in which a profession of intention such as 'I plan to go if I am not struck by lightning' does not convey any uncertainty about one's success but rather reassures the audience about its likelihood and the firmness of one's purpose (see Hunter 1978, p. 50). Third, the range of expressions available for communicating our intentions ('I shall', 'I will', 'I am going to', 'I plan to', 'I intend to'—with possible adverbial qualifications such as 'certainly') can convey different degrees of confidence or cautiousness even without explicit if-clauses but with a similar ambiguity as to the exact nature of the qualification (Reminders of generic qualifications? Warnings about likely impediments? Warnings about possible laziness, forgetfulness, or indolence? Hedging of fallibility? Want of firmness? Reluctance? Cancellations of implied assurance?); see Hunter (1978, pp. 38–40; 50–51), Rundle (1997, pp. 130–132).

⁶⁶ See Nietzsche, *The Genealogy of Morals*, I.3.

⁶⁷ These statements have the ring of Moore's paradox: the agent appears to be undoing the adoption of the intention in the very avowal of it (while the third-person analogues are not problematic).

⁶⁸ Possible failures of resolve might appear in genuine restrictive conditions but only when the agent sets a limit to the effort she is willing to put in securing the success of her φ -ing on the face of these threats. An agent might intend to (φ IF the steps required to avoid forgetting about φ -ing, resisting the temptation not to φ , etc. are not too expensive or strenuous) (cf. Hunter 1978, p. 51). The agent is not undermining one's commitment as long as there is still a range of circumstances over which she takes herself to be required to resist possible failures (cf. Korsgaard 2007, pp. 61–62). This restriction amounts to a kind of half-heartedness in intention, but it is internal to its content. The agent is not half-hearted in her attitude. It is not a lack of firmness in her undertaking. It is not a matter of ambivalence about one's decision in Frankfurt's sense (that is, of being of two minds about it). Nor is one hesitant about it because overly concerned

about one's generic fallibility—"I am fully behind my decision but I worry that it might not be the right one"—a concern that might lead to vacillation.

⁶⁹ The trouble with switching between the two stances is that it might undermine the long-term effectiveness of intentions. If I avow intentions together with the third-person warning, I might end up being more prone to cave in when obstacles and temptations arise. My expectation of or concern about a future failure (which in turn might be based on evidence from prior failures) might undermine my confidence in my ability to overcome obstacles and temptations, thereby weakening my resolve when they actually arise (see Ainslie, 2001). Thus, some amount of self-deception about the strength and direction of one's motivation might be helpful to build strength of will. This does not make the disclosure of one's weakness and hesitation paradoxical, although it might advise against the disclosure on account of its detrimental psychological effects.

⁷⁰ See Yaffe (2004: 277).

⁷¹ On the limited norms engaged by settled objectives, see Bratman (2009a).

⁷² Williams (2006: 41).

⁷³ Cf. Bratman's concerns with the atomism of agency that he takes to be implied by Velleman's cognitivism, see Bratman (2009b).

⁷⁴ To this extent, Don Giovanni is similar to the 'tyrannical person' discussed by Korsgaard (2009: 172–3): "Who is willing to do certain things *whatever the consequences*, which makes him such an unsettling parody of the just person."

⁷⁵ As I argue in Ferrero (2009a), the internal character of the temporal integration affects the nature of what we can pursue and the very structure of our temporal existence.

⁷⁶ For instance, we avoid *prima facie* conflicts by such techniques as restrictive scheduling, the qualifying of preconditions, or the subordination of plans (see p. 718 and note 49).

⁷⁷ Echoing a suggestion by Baier (1979), we can say that differences in the distinctive ways of changing minds reflect basic differences in the kinds of minds.

⁷⁸ Williams (1973, p. 86).

⁷⁹ See the discussion of the impossibility of a conditional commitment to justice as the formal principle of deliberative action and of the unity of agency in Korsgaard (2009: 178ff).

⁸⁰ On the issue of the inescapability of the standards of agency, see Ferrero (2009b).

⁸¹ Earlier versions of this paper were presented at Stanford University, UC-Riverside, and UC-Irvine. I thank the audiences for their comments and criticism, especially Ermanno Bencivenga, Facundo Alonso, Margaret Gilbert, Aaron James, Krista Lawlor, Michael Nelson, and Gary Watson. I also benefited from conversations and correspondence with Eljah Millgram, Niko Kolodny, Bruno Verbeek, Alec Walen, R.Jay Wallace, and Gideon Yaffe. Thanks to Dan Rosenberg for proofreading. I am grateful to the anonymous reviewers for prompting me to clarify some important issues in the paper. The final version of this paper has greatly benefitted from the comments and criticisms that Gregory Klass made on an earlier version, some of which will appear in Klass (2009). A special thank to Michael Bratman for his encouragements, suggestions, and criticisms throughout the development of this project.

References

- Ainslie, George. (2001) *Breakdown of Will*. Cambridge: Cambridge University Press.
- Austin, John. (1956) "Ifs and Cans," in John Austin, *Philosophical Papers*. Oxford: Oxford University Press, 1961: 205–232.
- Bach, Kent. (1984) "Default Reasoning: Jumping to Conclusions and Knowing When to Think Twice," *Pacific Philosophical Quarterly* **65**: 37–58.
- Baier, Annette. (1979) "Mind and Change of Mind." *Midwest Studies in Philosophy* **4**: 157–176.
- Bennett, Jonathan. (2003) *A Philosophical Guide to Conditionals*. Oxford: Oxford University Press.
- Bratman, Michael. (1979) "Simple Intentions." *Philosophical Studies* **36**: 245–259.

- Bratman, Michael. (1985) "Davidson's Theory of Intention," in Michael Bratman, *Faces of Intention*, 1999. Cambridge: Cambridge University Press: 209–224.
- Bratman, Michael. (1987) *Intentions, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Bratman, Michael. (1992) "Practical Reasoning and Acceptance in a Context," in Michael Bratman, *Faces of Intention*, 1999. Cambridge: Cambridge University Press: 15–34.
- Bratman, Michael. (2009a) "Intention, Belief, and Instrumental Rationality," in David Sobel and Steven Wall (eds.), *Reasons for Action*, Cambridge: Cambridge University Press.
- Bratman, Michael. (2009b) "Intention, Belief, Practical, Theoretical," in Simon Robertson (ed.) *Spheres of Reason*, Oxford: Oxford University Press.
- Cartwright, J.P.W. (1990) "Conditional Intention." *Philosophical Studies* **60**: 233–255.
- Davidson, Donald. (1978) "Intending," in *Essays in Actions and Events*, 1980. Oxford: 83–102.
- Edgington, Dorothy. (1995) "On Conditionals." *Mind* **104**: 235–329.
- Edgington, Dorothy. (2001) "Conditionals," in Lou Goble (ed.), *The Blackwell Guide to Philosophical Logic*. Oxford: Blackwell: 385–414.
- Edgington, Dorothy. (2006) "Conditionals," in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <<http://plato.stanford.edu/entries/conditionals/>>
- Ferrero, Luca. (2009a) "What Good is a Diachronic Will?" *Philosophical Studies* **144**: 403–430.
- Ferrero, Luca. (2009b) "Constitutivism and the Inescapability of Agency." *Oxford Studies in Metaethics* **4**: 303–333.
- Ferrero, Luca. (forthcoming) "Decisions, Diachronic Autonomy, and the Division of Deliberative Labor." *Philosophers' Imprint*.
- Gabbay, Dov and John Woods. (2005) "The Practical Turn in Logic." *Handbook of Philosophical Logic* **13**: 15–122.
- Gilbert, Margaret. (1993) "Is an Agreement an Exchange of Promises?" *The Journal of Philosophy* **90**: 627–649.
- Hilpinen, Risto. (2001) "Deontic Logic" in Lou Goble (ed.), *The Blackwell Guide to Philosophical Logic*. Oxford: Blackwell: 159–182.
- Hunter, John. F.M. (1978) *Intending*. Halifax, N.S.: Dalhousie University Press.
- Klass, Gregory. (2009) "A Conditional Intent to Perform." *Legal Theory* **15**: 107–147.
- Korsgaard, Christine. (2007) "Morality and the Logic of Caring," in Harry Frankfurt, *Taking Ourselves Seriously and Getting it Right*. Stanford: Stanford University Press.
- Korsgaard, Christine. (2009) *Self-Constitution*. Oxford University Press.
- McCann, Hugh. (1991) "Settled Objectives and Rational Constraints." *American Philosophical Quarterly* **28**: 25–36.
- Roth, Abe S. (2004) "Shared Agency and Contralateral Commitments," *Philosophical Review* **113**: 359–410.
- Rundle, Bede. (1997) *Mind in Action*. Oxford: Oxford University Press.
- Taylor, Charles. (1985). "What is Human Agency?" in *Human Agency and Language*, Cambridge: Cambridge University Press: 15–44.
- Velleman, David. (1989) *Practical Reflection*. Princeton: Princeton University Press.
- Velleman, David. (1997) "How to Share an Intention." *Philosophy and Phenomenological Research* **57**: 29–50.
- Vranas, Peter B.M. (2008) "New Foundations for Imperative Logic I: Logical Connectives, Consistency, and Quantifiers." *Noûs* **42**: 529–572.
- Williams, Bernard. (1973) *Problems of the Self*. Cambridge: Cambridge University Press.
- Williams, Bernard. (2006) *On Opera*. Yale: Yale University Press
- Yaffe, Gideon. (2004) "Conditional Intent and *Mens Rea*." *Legal Theory* **10**: 273–310.