

DIACHRONIC CONSTRAINTS OF PRACTICAL RATIONALITY

Luca Ferrero
University of Wisconsin-Milwaukee

1. Introduction

1.1. The Diachronic Practical Predicament

Rationality imposes constraints on combinations of psychological attitudes. For instance, rationality demands that, at any particular time, if one both intends to ϕ and believes that ψ -ing is a necessary means to ϕ -ing, then one ought to intend to ψ . In spite of disagreements about the nature, source, and scope of these rational constraints, it seems uncontroversial that the stock examples concern combinations of contemporaneous attitudes.¹ It is not immediately apparent, however, whether there are genuinely *diachronic* rational constraints, i.e., rational pressures on cross-temporal combinations of attitudes.

In this paper, I want to consider the more specific question of whether there might be genuinely diachronic constraints on the rationality of intentions. I maintain that if these constraints exist, they arise in response to the distinctive predicament of agents who have to act over time. Temporal agents are forced to pursue temporally extended activities using resources and abilities that are usually scarce, unequally distributed over time, and unstable.

There are three distinct features of the *diachronic predicament of temporal agents*:

- (1) Executive resources—including opportunities for action and the relevant abilities—can be scarce both at a time and over time. Agents need to manage the piecemeal progress in their activities by taking partial steps at appropriate times and in the proper sequence.

- (2) Deliberative resources are scarce and unequally distributed. Deliberation usually takes time, uses up scarce computational resources, and is based on partial information. Deliberative conditions are usually unstable (often in unpredictable ways); they can get better or worse as time goes by.
- (3) The agent's practical standpoint might be unstable. She can change her inclinations, wants, preferences, cares, concerns, and practical principles (and their respective temporal horizons).

The combination of these three factors provides the normal background against which ordinary agents engage in temporally extended activities.

1.2. Outline

In this paper, I will argue that there are distinct rational constraints that deal with the different dimensions of the diachronic predicament. I will introduce two constraints: one about the structure of temporally extended activities, another about the preservation of intentions over time. These two constraints have essential diachronic features, but they might still be reducible to synchronic rational constraints. As such, they might fail to meet the strictest criterion for diachronicity. By contrast, Bratman (2012; 2010) proposes a diachronic constraint that is genuinely irreducible. But he does so in a form that entails 'practical conservatism' in cases of normative underdetermination. In my view, this is too high a price to pay. We should rather settle for diachronic constraints that—even if reducible to synchronic combinations of attitudes—are diachronic in the weaker sense that they regulate attitudes with essential diachronic dimensions (either in their content or in their causal history). Or so I will argue.

1.3. Bratman's Desiderata

To isolate the different constraints of diachronic rationality, I proceed by a strategy of *reversed* Gricean creature construction:² I begin by discussing the pressures that apply to idealized agents, who face a milder form of diachronic predicament. I then proceed by relaxing the idealization and consider the rational demands that apply to agents closer and closer to us.

I frame my discussion in terms of a set of desiderata that can be distilled from Bratman's groundbreaking contribution in Bratman (2012) and Bratman (2010).

A genuinely diachronic rational constraint on intentions should have the following features:

- (1) Be a pressure for a *cross-temporal pattern of attitudes*;
- (2) Be a pressure of a *distinctive* and *direct* diachronic sort, rather than the mere outcome of the ‘snowball effect’ of earlier actions and attitudes on present intentions;
- (3) Avoid bootstrapping;
- (4) Conforming to it should matter in each *particular* instance rather than just on account of some general benefits.³

In addition to Bratman’s desiderata, it is also plausible to expect that:

- (5) A genuinely diachronic constraint should be *irreducible* to synchronic constraints on patterns of contemporaneous attitudes.⁴

2. Temporal Unity and Frictionless Deliberators

2.1. Temporal Unity

In discussing diachronic rationality, I focus on what I believe to be the paradigmatic form of diachronic intentional agency: *temporally extended intentional activities*. That is, activities that take time to complete and need to be intentionally sustained throughout their unfolding. Contrast them with extended activities that are not pursued out of a continuous sense of their overall temporal structure. That is, activities that the agent continues to pursue either because she is responding to local cues only in light of proximal concerns or because she is cajoled or forced by a pre-commitment. Let’s call the paradigmatic extended intentional activities ‘temporally unified’ activities.

As I argue more extensively in Ferrero (2009), throughout the unfolding of a temporally unified activity, the agent takes each momentary step:

- (1) in light of her appreciation and approval of the activity as extended and unified over time;
- (2) autonomously—in the sense that she is not manipulated or goaded by the effects of the steps she has already taken;
- (3) with the expectation that the future steps will be taken in a similar fashion (that is, as a result of her contemporaneous autonomous appreciation and approval of the unified activity as such).

Because of the expected stability in the autonomous appreciation and approval, at any particular time in the unfolding of the activity, the agent can rely on her future cooperation in that pursuit. Hence, at any particular time, she only needs to contribute what she takes to be necessary for the activity’s

progress *at that time*, without worrying about setting up devices for goading or manipulating herself in the future.⁵

The intentions that guide temporally unified activities are *future-directed* or *prospective*, not in the sense that the activity's inception lies in the future, but in the more fundamental sense that the activity's unfolding and eventual completion do.

2.2. *A Constraint of Temporal Unity*

Let's begin by considering idealized agents who suffer only from limitations in executive resources (the first dimension of the diachronic predicament). These idealized agents are 'frictionless deliberators'—to use Bratman (1987: 28)'s expression: At any moment, they can deliberate instantaneously in light of full information. In addition, they can rely on a stable practical standpoint. But they are not omnipotent. They have to deal with the scarcity and instability of opportunities and abilities for action. Therefore, in order to achieve some of their distal goals, they have to engage in temporally unified activities.

When a frictionless deliberator *A* has the intention to engage in a temporally unified activity ϕ over a certain interval t_{1-10} , she is under a distinctive rational pressure that derives from the structure of temporal unity. Not only is she required to believe that her ϕ -ing is in principle possible and to intend to take the known present necessary means to ϕ -ing, she is also required to believe that she will continue to autonomously appreciate and approve of the unified activity through its completion.

More formally, *A* is under the following rational constraint:

(R1) An agent *A* ought to (if she intends at t_n ($1 < n < 10$) to ϕ in the interval t_{1-10} , then she believes at t_n that—assuming she continues to have the ability and opportunity to ϕ —she will continue to autonomously appreciate and approve of her ϕ -ing until t_{10}).

(R1) is wide scope. One can satisfy it by giving up the intention to ϕ when one lacks the belief in one's future autonomous cooperation in ϕ -ing.⁶

2.3. *Is (R1) Diachronic?*

(R1) derives from the structure of temporally unified activities. As such, it has a distinctive diachronic aspect. But the demand is only addressed to the agent's *contemporaneous* attitudes, to her present intention about extended activity and her present belief about her future autonomous cooperation. The constraint does not directly govern a cross-temporal pattern of attitudes. This

pattern is necessary for the actual existence of a genuinely temporally unified activity, but it appears in (R1) only in the *content* of some of the present attitudes that fall under (R1) (e.g., in the belief about her future cooperation). (R1) could be said to be the ‘projection’ onto the agent’s psychology at *a particular time* of the constitutive features of temporally unified activities. To this extent, (R1) fails to meet the desideratum of irreducibility. After all, it is a synchronic rational constraint. But it is a synchronic constraint with a distinctive, necessary, and irreducible diachronic *element*. As such, it qualifies as diachronic in a weaker sense.

2.4. (R1) and the Stability of Intentions

A directly diachronic constraint would make an unmediated demand on the cross-temporal pattern of attitudes. An example of a constraint of this kind would be a simple pressure to stick with one’s prior intention:

- (i) If at t_1 A intends to ϕ at f , then at t_2 A ought to intend to ϕ at f (for $f \geq t_2$).

This constraint is troublesome because it makes sticking to prior intentions rational regardless of A ’s belief at t_2 in the existence of adequate grounds for ϕ -ing at f . To avoid this bootstrapping, one might add a clause about the continuous belief in adequate grounds. But in so doing, one might make the prior intention rationally irrelevant. The present belief in the adequate grounds for ϕ -ing is all that the agent needs to be rational in her intending to ϕ at the present time, *whether or not* she had a prior intention to ϕ . The stability of the intention would only be a by-product of a synchronic rational constraint like:

- (ii) A at t_n ought to intend to do one of the actions that at that time she takes to be supported by adequate grounds.

This is not to deny that the case for the intended action ϕ might depend on A ’s prior intention. For on the basis of that intention, A might have taken steps that are either necessary or sufficient for the present existence of a case for ϕ . (For instance, A might have taken means to ϕ -ing that were necessary at the earlier time). In addition, the case for a temporally unified activity always makes essential reference to cross-temporal patterns of attitudes, given that they are constitutive of that activity. But as far as rationality is concerned, the diachronic elements are a matter of the *substantive* rather than *structural* considerations in support of ϕ -ing. As far as structural rationality is concerned—the rationality of the relations between the agent’s attitudes—a constraint like (ii) is synchronic.⁷

2.5. Does (R1) Matter?

Let's return to the frictionless deliberators. As long as they intend to engage in temporally unified activities, they are bound by (R1). This is the basic rational demand that regulates their intentional pursuits over time, given that they suffer from limitations in executive resources.

Their diachronic predicament stops there. They are not concerned with the rationality of retaining intentions over time. In the presence of a stable case for the extended activities, they can always figure out, at any particular time, what they are supposed to do at that time, whether or not they remember their prior intentions. They do not need to *settle* practical questions in advance. They do not need to secure that future deliberation about the same matter be defeasibly closed by default. As long as the underlying case of ϕ is stable, neither the stability of their conduct, nor their rationality are jeopardized by this 'unsettledness'. In virtue of their unlimited deliberative resources, they can continue, at each and every time, directly to respond to their contemporaneous and full cognition of the stable case for ϕ -ing.

(R1), however, is not limited in its application to idealized, frictionless deliberators. I have introduced it in connection with them only in order to isolate its contribution from other features of diachronic agency. *All* temporal agents, idealized or not, are subjected to (R1) when they are pursuing temporally unified activities. This is because this constraint derives from the basic structure of temporally unified activities and addresses the basic predicament of temporal agents: the diachronic scarcity of executive powers (which is, at bottom, what makes all of them *temporal* agents).

In ordinary agents, this constraint operates alongside other rational pressures that derive from the other dimensions of the diachronic predicament. Nevertheless, some philosophical issues of diachronic rationality already arise, for all agents, in connection with (R1) independently of other constraints.⁸

3. Intentions as Summary Attitudes

3.1. The Preservation of Attitudes

Let's consider the second aspect of the diachronic predicament: the scarcity of deliberative resources. Because of limitations and uneven temporal distribution of information and computational resources, deliberation is usually costly and time-consuming, forcing trade-offs between speed and accuracy. Ordinary agents are not 'frictionless deliberators'. For them, it is usually dangerously inefficient (if at all possible) to recapitulate deliberation after the intention has been acquired. They have to deal with the temporal

management of both execution and deliberation. With respect to latter, they need to figure out—given the often unpredictable and changing deliberative circumstances—(a) when to engage in deliberation and for how long, (b) when to rely on their prior conclusions, and (c) when to re-open it, if at all.

To overcome the scarcity in deliberative resources, agents need a psychology endowed with memory, i.e., with a *preservative* faculty that makes them retain and retrieve attitudes over time. Preservative memory makes previously acquired attitudes available to the subject at a later time as ‘ready-made’ for their standard deployment.⁹

By default, agents like us take their preserved attitudes to be warranted, even if the grounds in their support might not be immediately (if at all) accessible. Our cognition of the grounds of these attitudes oftentimes is not full and direct, but partial and indirect; it is *oblique*, as Tenenbaum (2007) puts it. A preserved attitude can be said to carry the warrant secured by the process of its original acquisition. To the extent that it was acquired by some piece of reasoning, the preserved attitude plays the role of a ‘summary’ of the conclusion of this past reasoning. By extension of this role, I suggest to use the expression ‘summary attitudes’ for judgment-sensitive attitudes—such as beliefs and intentions—that are preserved over time with a default expectation that their warrant has been preserved over time.

3.2. *Psychologically Stable Intentions*

For agents with a preservative memory, intentions usually persist over time as summary attitudes. At a later time, they play their characteristic roles—such as framing further deliberation and controlling conduct—even if the agent has no direct cognition of their grounds.

The intention-as-summary to ϕ stands as a proxy for the conclusion of a prior deliberation (whether actual or potential). The agent is thereby *settled* on ϕ -ing, given that (further) deliberation about that matter is closed by default. (This is the kind of settledness that frictionless deliberators might have the leisure to forgo; see section 2.5.)

3.3. *A Constraint on the Preservation of Intentions*

The default reliance on a preserved intention depends on two assumptions: (a) that the force and content of the attitude have been properly preserved; (b) that the attitude is still warranted. Under scarcity of deliberative resources, if these assumptions hold, it is irrational to abandon an intention-as-summary (either by reopening deliberation or giving up the intention altogether).

Here is the distinctive rational constraint on intentions *as preserved summary attitudes*:

(R2) At t_2 , A ought

(if

- (1) A takes her deliberative resources to be scarce, and
- (2) A has the intention to ϕ as preserved from an earlier time t_1 (intention-as-summary), and
- (3) A takes the intention-as-summary to be correctly preserved, and
- (4) A takes the intention-as-summary to be still warranted,

then

A retains the intention-as-summary to ϕ (*unless she acquires the intention to ψ , which she takes to be as adequately supported as the intention to ϕ*).

(For the time being, the reader should ignore the italicized clause about the intention to ψ . I will return to it in section 5.2.)

Among the four antecedents, the least secure is the last one, which concerns the stable warrant. This is because, ordinarily, agents can safely assume the correctness of their memory of prior intentions. The belief in a stable warrant is not usually determined by checking directly for the correctness of the conclusion of the original deliberation. This would require engaging in an actual deliberation at t_2 , which would defeat the point of relying on the intention-as-summary.

Rather, the belief in the stability of the warrant usually depends on the agent's belief about the comparative changes in her deliberative circumstances (including changes in the available computational resources and relevant information). If the agent takes the deliberative circumstances at t_2 to be no better than those at t_1 , then she takes the intention's warrant to be stable. For she does not expect that, were she to reconsider the matter anew, she could improve on her prior conclusion.

Nevertheless, an improvement in deliberative circumstances does not necessarily translate in a loss of warrant. Reconsideration might still be unreasonable because of the *entrenching effect* of the intention's prior acquisition: the potential gain in the revision of the intention might be offset by the costs of either the new deliberation or the undoing of the steps already taken. Because of this entrenchment, the belief that one would choose differently if one were to deliberate *for the first time* at t_2 might not be sufficient to justify the re-opening of deliberation at t_2 .

Determining the stability of the warrant by comparing deliberative circumstances could be complicated. It depends on a variety of general

norms, policies, and rules of thumb.¹⁰ Nonetheless, the agent usually does not engage in an explicit and articulate comparison of the changes in circumstances. Rather, she usually relies on a general sense of whether major changes in deliberative circumstances have occurred, against the background operation of psychological habits and propensities of non-reconsideration (which are often adequate by themselves to secure the efficient use of scarce deliberative resources).

3.4. (R2) and the Stability of Intentions

Unlike (R1), (R2) is a rational constraint that relates attitudes over time. Specifically, it presides over the rationality of the stability of intentions as summary attitudes. (R2) avoids the problems with the earlier attempts to formulate a constraint of intention-stability. Unlike (i) in section 2.4, (R2) is not bootstrapping. The prior intention makes a rational difference at the later time only via the belief in the stability of its warrant, not simply because it has been priorly acquired.¹¹ Unlike (ii), (R2) is not redundant, since the stability of the warrant is determined by comparing deliberative circumstances rather than by direct assessment of the case for ϕ -ing.

3.5. Is (R2) Diachronic?

Is (R2) genuinely diachronic? Although the constraint ties together the intention at t_1 with the intention at t_2 , it does so only via the presence of the intention-as-summary at t_2 . At issue for (R2) is whether the intention-as-summary, which is given to the agent at t_2 by her retentive faculty, is to be left undisturbed, so to say, and perform its default job *at that time*.

If the intention is forgotten prior to t_2 , (R2) does not require the intention to be re-adopted. The rational demand only concerns the *full retention* of the intention that is already available to the agent at that time as a preserved attitude. By ‘full retention’, I mean that the agent is not simply acknowledging that she has the intention-as-summary, but she is also deploying it rather than suspending or abandoning it. (R2) puts a pressure on a combination of contemporaneous attitudes, even if one of these attitudes comes from the past via memory. Hence, (R2) fails to qualify as *irreducibly* diachronic (see section 1.3).

A failure of memory might be a failure of rationality. But if so, it would be because of a different rational constraint (most likely, one that applies to all judgment-sensitive attitudes and whose success is a precondition for the application of (R2)). A rational constraint against forgetfulness would be irreducibly diachronic. For it would directly secure a cross-temporal pattern, although one that relates two temporal stages of the same attitude: the

original attitude at the time when first acquired and the same attitude as preserved at a later time. This is a pattern of a different kind than the one normally associated with the idea of structural rationality, which relates two or more distinct attitudes. This suggests that the rational criticism of forgetfulness should not be bundled together with the standard kinds of structural irrationality.

A systematic failure to retain judgment-sensitive attitudes—like beliefs and intentions—is a failure in securing the necessary background for the proper functioning of the rational psychology of temporally extended agents like us. But it does not seem to me that, for each particular judgment-sensitive attitude that one might have at any particular time t_1 , one is under a rational constraint to preserve it at a later time t_2 . The preservation of the relevant attitudes is, however, a precondition for the application of many standard constraints of structural rationality. For instance, one cannot be deemed structurally irrational if at t_1 one intends to ϕ and at the later time t_2 one does not intend the known necessary means to ϕ -ing but one has, in the meantime, forgotten about the original intention and has not independently re-acquired it (even if one might be irrational for not adopting it at t_2 if one believes at that time that there are decisive ground in support of ϕ -ing).

I suspect that similar considerations can be made for all constraints of structural rationality. If so, it follows that these constraints all necessarily operate over combinations of contemporaneous attitudes—even if some of them are present as preserved attitudes. Setting aside the pressures against forgetfulness, therefore, there might be no *irreducibly* diachronic constraints. The rational cross-temporal patterns of attitudes would always have a synchronic counterpart, over which a synchronic constraint would first have to be applied. This is not to deny that there are diachronic dimensions to practical rationality. But diachronic structural concerns seem always to require a sort of ‘projection’ onto the lower-dimensional surface of the present moment, where synchronic constraints apply in the first place. Forgetting might make the agent immune from criticisms of structural irrationality because it blocks the projection of cross-temporal patterns into combinations of contemporaneous attitudes.

3.6. *Comparing (R1) and (R2)*

The two constraints (R1) and (R2) are complementary but independent. They address different features of intentions and diachronic agency. (R1) is concerned with the structure of diachronic agency in the mode of temporal unity (regardless of the contribution of intentions as preserved summary attitudes). (R2) concerns the effects of the temporal preservation of intentions as summary attitudes (even in cases where the intentions support activities that lack the complexity of temporal unity and, thereby, do not engage (R1)).

In a sense, (R1) is more closely related to diachronic agency than (R2). (R1) reflects the constitutive features of temporally unified activities; hence, it lies at the core of diachronic rationality for *all* kinds of temporal agents—whether idealized or not—as long as they are set on pursuing a temporally unified activity.

(R2), instead, is not distinctive of diachronic agency. As formulated, it applies only to intentions as summary attitudes. But this is just an instance of a more general constraint that regulates the preservation of all judgment-sensitive attitudes. For instance, there is a cognitive counterpart to (R2) which holds of beliefs as summary attitudes:

(R2-b) At t_2 , A ought

(if

- (1) A takes her deliberative resources to be scarce, and
- (2) A has the belief that p as preserved from an earlier time t_1 (intention-as-summary), and
- (3) A takes the belief-as-summary to be correctly preserved, and
- (4) A takes the belief-as-summary to be still warranted,

then

A retains the belief-as-summary that p (unless she acquired the belief that q , which she takes to be as adequately supported as the belief that p)).

Whereas (R1) is constitutively necessary for securing temporal unity, (R2) is only instrumental to it, and only to the extent to which deliberative resources are scarce. This is because constraints like (R2) and (R2-b) are primarily concerned with the stability of judgment-sensitive attitudes required for the efficient use of scarce deliberative resources.

4. Unstable Practical Standpoint

4.1. *A Spectrum of Instabilities*

The third source of the predicament of diachronic agency is the instability of practical standpoint: changes in preferences, cares, concerns, principles, and the temporal horizons of their scopes.

These changes, whether temporary or permanent, affect the ability of the agent to engage in temporary extended activities: the agent at the future time might give up the activity since she might no longer find it choiceworthy in light of her future practical standpoint.

How does diachronic rationality deal with these cases? There is a spectrum of cases. At one extreme, consider a persistent, irreversible, and massive change, like in Parfit (1984)'s Russian Nobleman. In this case, one's future conduct is beyond the direct reach of one's prior intentions. One can only deal with one's future conduct as if it were that of a different agent. Hence, this case falls outside of the scope of rational constraints on combinations of attitudes.

The opposite (and far more common) extreme is that of temporary, reversible, and limited changes in practical standpoint. These changes can be dealt with relatively easily by timely reminders, prompts, or nudges, which help set the agent back on the intended course of action. These devices are not manipulative. They do not violate the agent's autonomy over time; they are rather 'tricks' that the agent welcomes at the very time when they operate. In a similar fashion, progress in ϕ -ing might at times be simply 'inertial', ϕ -ing can progress without the agent's attending to it. This often happens thanks to various psychological habits and propensities that might either block reconsideration of prior decisions or make some actions automatic or habitual. Last but not least, many pursuits may be temporarily left 'dormant', without detriment to their future success, while the agent is busy with more urgent projects.

In all of these cases, the agent can make progress in her intentional ϕ -ing, even if it is not strictly true that she does so out of the contemporaneous full appreciation and approval of her momentary contribution (whether because she is not attending to it or just gently tricked into it). It would be unreasonable to deny that this kind of progress is compatible with genuine diachronic intentionality, even if the agent's conduct does not meet the strict standards of temporal unity articulated in (R1).

The solution, I suggest, is to relax the conditions of temporal unity by taking some of its demands as *regulative* ideals, which ordinary agents strive to approximate, oftentimes by relying on psychological habits, propensities, and various 'tricks' for on-the-fly adjustments.

Although the demand of continuous, full, and autonomous future cooperation is to be relaxed for ordinary agents, I still maintain that a temporally unified activity requires that, during its unfolding, there is a sufficient number of times when the agent carries out the activity out of her contemporaneous appreciation and autonomous approval of it—even if only in an implicit and inarticulate way. What is incompatible with genuine temporal unity is progress that extensively relies on manipulative devices directed at breaking the resistance of reluctant future selves.

These considerations affect the status of (R1) for ordinary agents who are subjected to the diachronic predicament under all of its dimensions. For these agents, the demand of future autonomous cooperation takes on the character of a regulative ideal. But this does not change that, however one might reformulate (R1), a rational constraint based on the structure of temporal unity is still a fundamental feature of intentional diachronic agency.

The milder forms of instability in the practical standpoint, therefore, do not require the introduction of new rational constraints but an adjustment in the status of the constraints that derive from the structure of temporal unity.

4.2. *Temptation*

In the intermediate cases, a practical standpoint is unstable because of changes that—although temporary, reversible, and limited in scope—produce some momentary strong resistance to ϕ -ing. Typical cases are those of powerful temptations. These temptations cannot be resisted by the nudging strategies I have just illustrated. More successful strategies involve self-directed manipulation, but only at the price of violating one's autonomy. When engaged in self-manipulation, at issue are not structural but substantive questions about the instrumental efficacy of the specific strategies of self-manipulation. This is not to rule out that considerations of structural rationality might be able to deal with some kinds of temptations in a non-manipulative way. The matter is still very much debated in the literature.¹² Further discussion must thus be left for another occasion. But this leaves open the possibility that we might have to countenance additional structural rational diachronic constraints on intentions in order to deal with some instabilities in practical standpoint.

5. Bratman on Diachronic Constraints

5.1. *Bratman's Diachronic Rational Constraint (D)*

As indicated at the outset, Bratman is to be credited for raising the issue of the existence, nature, and importance of diachronic practical constraints. I framed my previous discussion mostly in terms of his desiderata. Let's now consider how my conclusions relate to his positive view. Bratman has argued that the distinctive diachronic constraint on intentions is as follows:

(D) It is always locally irrational: Intending at t_1 to ϕ at t_2 ; throughout t_1 – t_2 confidently taking one's relevant grounds adequately to support this very intention; and yet at t_2 newly abandoning this intention to ϕ at t_2 .¹³

(D) seems closely related to (R2). Both constraints concern the retention of intentions (Bratman does not discuss any constraint about the structure of diachronic activities comparable to (R1)). But I will argue that (R2) is to be preferred to (D) on at least two counts. First, (R2) avoids an ambiguity in Bratman's formulation that might lead to a problematic reading of (D). Second and most importantly, (R2) does not entail what I take to be an untenable form of *practical conservatism*.

The ambiguity in (D) concerns the expression ‘taking one’s relevant grounds adequately to support intending’ to ϕ . Should the agent be interpreted as having direct or oblique cognition of these grounds (for the distinction, see section 3.1)?

Consider those scenarios where the agent believes that there is a *decisive* case for ϕ -ing at t_2 . If the agent has a *direct* cognition of these grounds, the rationality of her intending to ϕ is entirely determined by a purely synchronic constraint that says that one is to intend to do what one takes *at that time* to be supported by decisive grounds. Whether one has a prior intention to do it is entirely irrelevant. Hence, under the direct-cognition reading, and in cases of normative determination, (D) would be utterly redundant and have no genuine diachronic import—it would be a version of the constraint (ii), which I rejected in section 2.4.

The redundancy problem does not arise for the oblique reading. In the oblique reading, (D) is much closer to (R2). In order to continue to take her grounds to support ϕ -ing, the agent at t_2 only needs the belief that the deliberative circumstances have not sufficiently improved (if at all) to justify the re-opening of deliberation. This is a belief that she can have at t_2 even if she has no direct access at t_2 to the grounds for ϕ -ing. In this case, (D) is not redundant. It would work on the stability of the preserved intention in a manner similar to (R2). And in this case the prior intention would make a difference (otherwise, the intention would not have been preserved in the first place). Nonetheless, under the oblique reading, (D) fails to qualify as *irreducibly* diachronic (for the same reasons that (R2) fails, see section 3.5).

It seems to me that the oblique reading of (D) is to be preferred. If so, (R2) would be very much in the spirit of Bratman’s solution. Nonetheless, I maintain that the *letter* of (R2) is to be preferred. For two reasons: first, (R2) avoids the ambiguity between the direct and oblique reading; second—and most importantly—(R2) makes explicit that the constraint governs the *preservation* of intentions as summary attitudes; as such, it offers a better articulation of the conditions that bear on the rationality of their retention than (D) does.

5.2. Underdetermination

Even if the ambiguity between direct and oblique is removed, there is still a major difference between (D) and (R2), which becomes clear when considering normative underdetermination.

(D) supports the rationality of *practical conservatism* (Bratman 2007: 21). According to (D), if the case for ϕ -ing at t_2 is underdetermined (in that the agent takes that there are equally adequate grounds for ψ -ing), it is by default irrational to abandon one’s prior intention to ϕ even if only in exchange for the intention to ψ .

By contrast, (R2) does not make it irrational to abandon the prior intention to ϕ when the agent takes the case to be underdetermined.¹⁴ This is the point of the last clause in the consequent of (R2), which was left unexplained when I introduced (R2) in section 3.3. According to (R2), it is irrational to abandon the intention to ϕ *only* when the case in its support is taken (via the belief in the stability of the warrant) to be decisive.

There are two reasons why one might want a practically conservative constraint. First, the constraint would meet the desideratum of irreducible diachronicity. Second, the constraint might help with securing self-governance over time (Bratman 2012: sec. 6).

Consider irreducibility first. When the case remains underdetermined over time, no synchronic constraint can support the *constancy* of intention. A synchronic constraint can only be sensitive to contemporaneous normative underdetermination and, thereby, it cannot but be indifferent about prior intentions. (D) here has a clear advantage over (R2). However, I think that practical conservatism has no independent appeal. If so, the defense of (D) only in light of its irreducibility would be *ad hoc*. Or so I am about to argue.

5.3. Diachronic Underdetermination

As far as time is concerned, there are two basic kinds of underdetermination: momentary and persistent.

If the underdetermination is *momentary*, this is because the acquisition of the intention to ϕ at t_1 (even if only by plumping) removes future underdetermination via entrenchment (via the combined effects of the framing of deliberation and the preparatory steps, see section 3.3).¹⁵ As a result, at t_2 the rational pressure favoring ϕ -ing no longer amounts to conservatism. The rational agent now takes it that the case has been tipped—however slightly—in favor of ϕ -ing. This scenario can be handled without any difficulty by a constraint like (R2), so it does not tell us anything interesting about diachronic irreducibility.

Consider now *persistent* underdetermination: a scenario where the acquisition of the intention to ϕ has no effect on the comparative merits of ϕ -ing over ψ -ing. This happens when either (a) the intention is utterly idle, or (b) the consequences of the deliberation framed by the expectation of ϕ -ing and of the preparatory steps toward ϕ are so limited that they still leave the case between ϕ and ψ underdetermined (even if the costs of changing one's mind about ϕ might have increased as a result of the intention to ϕ).

A practically conservative constraint like (D) makes an actual difference in these cases, but only because (and to the extent that) these scenarios do not engage any interesting feature of planning diachronic agency. After all, in these cases the framing of future deliberation and the making of preparatory steps do not remove the underdetermination. In these scenarios, the

comparison between the alternative options is immune from the entrenching effects of planning. Consider agents who, in the face of persistent underdetermination, delay acquiring the intention to ϕ or, once they acquire it, change their mind. What reason is there to deem them *structurally* irrational given that, by the time they acquire the intention to ϕ or change it in favor of ψ -ing, they intend to do something that they believe to be supported by adequate grounds? If agents are operating in entrenchment-immune circumstances, why should they bother about constancy of intentions?

One possible reason is a concern with the increase in the costs of changing one's mind. Even if the underdetermination persists, what is immune from entrenchment are only the comparative merits of ϕ vs. ψ . The costs of changing one's mind about any one particular action might still increase as a result of one's acquiring the intention to do it. So it seems to be unreasonable to pay these costs solely to acquire an intention for ψ -ing, which one takes to be supported by adequate grounds that are not better than those supporting one's ϕ -ing.

In response, notice that as long as the underdetermination persists, the costs of changing one's mind are not sufficient to tip the balance in favor of ψ -ing. In this sense, these costs are factored in the present comparison between the two actions. Sticking with the original plan only on account of the costs already incurred and in spite of the present underdetermination is an instance of the sunk-costs fallacy. Hence, it is not structurally irrational to change one's mind in response to persistent underdetermination.

This is not to deny that an agent who, in this sort of situations, tends to change her mind (possibly shuffling between ϕ and ψ) might suffer from troubling *overall* inefficiencies in her use of scarce deliberative and executive resources. Constraints of pragmatic rationality might thus support a policy of practical conservatism in cases of persistent underdetermination. But this is only as a matter of a *general* policy and on account of *long-term aggregate* inefficiency and inefficacy. It is not a *local* rational pressure that applies to *each particular case* on its merits alone. It is a matter of pragmatic rather than structural rationality. As such, it fails to meet Bratman's own desideratum that conformity to a diachronic constraint should matter directly in each particular instance rather than on account of its general benefits (see desideratum 4 in section 1.3).

The genuine diachronic import of (D) is restricted to scenarios of persistent underdetermination. But these are circumstances where the extended activities do not engage, at least when considered locally, the distinctive contributions of planning agency, as attested by the persistent character of the underdetermination in spite of the effects of planning. Determining how frequently these scenarios occur is an empirical matter. If they are rare and precious, very little, if any, seems to be lost if practical conservatism is rejected. On the other hand, more might be at stake if persistent underdetermination is pervasive. Even so, this pervasiveness would

not support a demand for *local* practical conservatism in the mode of (D). Rather, it would support a *non-local* conservatism grounded in the overall pragmatic rationality of some general habits, propensities, and policies against inefficiency, which would usually include allowances for a modicum of local exceptions.

As long as we focus on constraints of local and structural rationality, (R2) is still to be preferred to (D).¹⁶ If practical conservatism turns out to hold at the global level, I think it can do so only by virtue of rational constraints of a different kind.

Practical conservatism is unattractive if it is needed only to secure the irreducible diachronic import of a constraint like (D). If so, I think that we are better off giving up on the desideratum of irreducibility.

This conclusion does not yet take into account the role that (D) might play in securing diachronic self-governance. Bratman (2012) claims that the rational pressure for the stability of intentions in cases of permanent underdetermination protects against the undermining of diachronic self-governance by ‘brute shuffling’: “the lurching from one plan-like commitment to another incompatible plan-like commitment seen as equal or incomparable, in a way that involves abandoning one’s prior intentions” (Bratman 2012: sec. 6).

In response, I am going to claim that diachronic self-governance is not necessarily threatened by brute shuffling. As I have just argued, there might be cases in which the underdetermination persists *in spite* of the effects of intentions as planning attitudes. When so, the power of intentions to *rationaly* settle practical matters is defused (the mere psychological power might still operate but it would not be stable under rational scrutiny; see Ferrero (2010: sec. 6)). *Pace* Bratman (2012: sec. 6), I see no point in insisting that the settling capacity of intentions be made rationally relevant in scenarios of this kind. This is because none of the other features of intentions as planning attitudes can get ‘a grip’ on these ‘planning-insensitive’ scenarios.

Nonetheless, as I have just discussed, there are scenarios that are only globally sensitive to planning: in them, planning does not remove the local underdetermination, but only affects the global efficiency and efficacy in the overall use of scarce deliberative and executive resources.¹⁷ In these scenarios, it might be true that brute shuffling eventually threatens diachronic self-governance. If so, there should be rational pressures for some kind of intention constancy of a practically conservative sort. The conservatism would be necessary to overcome the limitations of synchronic constraints, which can only be responsive to local underdetermination.

These pressures, however, are not going to be in the form of constraints of *local* rationality that apply equally to each individual scenario in isolation from global concerns (this is ultimately the same problem that besets the local practical conservatism required to secure irreducible diachronicity). Hence, a constraint like (D), together with its local practical conservatism, cannot be supported out of worries for the actual but only global effects of brute

shuffling on diachronic self-governance. As far as local practical rationality is concerned, a constraint like (R2) is still our best choice—which is open to being supplemented by pressures of global practical rationality that entail some kind of non-local practical conservatism.

My proposal does not deny that practical constraints have irreducible diachronic features. Nor does it rule out that some kind of practical conservatism might be true. But it locates both features in different places than Bratman does. My constraints are diachronic not because they directly regulate cross-temporal patterns of attitudes, but because they regulate synchronic patterns of attitudes with irreducible diachronic features. My practical conservatism is induced by general pragmatic considerations of efficiency and efficacy rather than stemming directly from an irreducible diachronic constraint like (D).

6. Conclusion

In this paper, I have argued that the distinctive predicament of temporal agents gives rise to various demands of diachronic practical rationality. Executive limitations are related to the constitutive features of temporal unity, regulated by (R1). Limitations in deliberative resources call for memory, as a faculty for the preservation of attitudes, regulated by (R2).

The predicament of temporal agents is genuinely diachronic. Both in its executive and psychological dimensions, it is a predicament that affects only temporal agents as they engage in temporally extended activities. As such, it cannot be reduced to the practical predicament of momentary agents who only engage in instantaneous actions.

But the diachronic predicament of temporal agents gives rise to rational constraints that might not be irreducibly diachronic. (R1) and (R2) are ultimately synchronic rational pressures, they constrain combinations of contemporary attitudes. Nevertheless, both constraints are diachronic in the sense that there are irreducible diachronic elements in the attitudes that fall under both constraints (these elements lying either in the attitudes' content or in their provenance as preserved attitudes). Both constraints apply only to synchronic combinations of attitudes, but of attitudes that belong to *temporal* agents as they engage in *temporally extended* activities. To this extent, (R1) and (R2) are synchronic constraints of *diachronic* practical rationality.¹⁸

Notes

1. This is not to deny that conforming to the demands of rationality might take time: one might need to acquire or abandon an attitude in order to conform to the rational constraint. Hence, there might be rational pressures on the

psychological *process* of acquiring or abandoning an attitude. But these demands are parasitic on the synchronic constraints on combinations of contemporaneous attitudes, since the processes aim at securing these contemporaneous combinations.

2. For the standard strategy of Gricean creature construction, see Bratman (2007: ch. 3).
3. See Bratman (2012: sec. 1,5) and Bratman (2010: 10–11, 20–21).
4. The desideratum of irreducibility was explicitly offered in earlier drafts of Bratman (2012). Bratman no longer explicitly mentions it in the final version (although he does not explicitly reject it either), and thus I won't attribute it to him. But notice that his positive view still meets it.
5. For further discussion of temporal unity and diachronic autonomy, see Ferrero (2009) and Ferrero (2010).
6. If an agent lacks the belief in her future cooperation, she might *re-characterize* the object of her intention: she might only intend to prepare for a possible future ϕ -ing, or attempt it, or set up a device that manipulates her into ϕ -ing, etc. In any event, unless these alternative projects can be completed immediately, the agent is still under a pressure to believe in her continuous cooperation with the re-characterized project for *some* future time—for instance, until the preparation is over or the pre-commitment device is in place.
7. For the distinction between structural and substantive rationality, see Scanlon (2007).
8. For instance, Kavka (1983)'s toxin puzzle already arises for frictionless deliberators. The puzzle depends simply on the passage of time and the related changes in the agent's causal powers, against a background of a stable practical standpoint. It does not depend, in particular, on the extent of the agent's deliberative resources. This seems to be true also of some of the questions of diachronic rationality discussed in Andreou (2012).
9. See Burge (2004) and Ferrero (2006).
10. These norms and policies add more stability to intentions. They protect them against some of the challenges that come from an improved deliberative situation. Notice that this stability is not a matter of either the distinctive structure of temporally unified activities or the basic working of memory. It is neither ontological nor psychological stability. It is rather pragmatic: the product of the instrumental considerations that support the preservation of intentions as summary attitudes in the face of our limitations and imperfections as temporal agents and deliberators.
11. The earlier acquisition of the intention could make a difference, but only by affecting the stability of the warrant via the entrenching effect discussed in section 3.3. This entrenchment is not an unacceptable bootstrapping. It rather reinforces the case for the intended ϕ -ing via the substantive effects of the intention acquisition on subsequent deliberation and conduct. It is not a matter of adding support for ϕ -ing on the basis of the *bare* fact of the acquisition of the intention to ϕ .
12. See Bratman (1999: ch. 34), Bratman (2007: ch. 12), Holton (2009: ch. 7), Ferrero (2012).
13. Bratman (2012: sec. 4) and Bratman (2010: 20).

14. The belief in the persistent underdetermination might be produced by preservative mechanisms under the guidance of (R2-b)—the cognitive counterpart of (R2) (see section 3.6). Having determined at t_1 that the case was underdetermined and would continue to be so until t_2 , the agent at t_1 might acquire a summary-belief about the underdetermination, a belief that she would be irrational to give up at t_2 unless she took that her deliberative conditions had improved so much as to warrant the re-opening of deliberation.
15. In cases of momentary underdetermination, the agent at the earlier time has the power to break the tie and impose the ϕ -ing on her future self. The agent at t_1 can thus ‘snowball’ herself into ϕ -ing at t_2 . Yet this is not normally a case of snowballing a future *reluctant* self. If the case was correctly seen as underdetermined at t_1 , at t_2 one should have no complaints about the choice made at the earlier time, given that, were she to choose for the first time at t_2 , ϕ -ing would still be one of the choiceworthy options.
16. (R2) has the additional merit of avoiding a differential treatment of the practical and theoretical domains. As suggested in section 3.6, (R2) is an instance of a more general constraint, which would therefore rule out practical conservatism in the case of beliefs as well. By contrast, the conservatism entailed by (D) is distinctively practical, which raises questions about the possible asymmetrical treatment of the two domains.
17. Notice that what is locally insensitive to planning is the ‘scenario’, that is, the comparison between the alternative pursuits. The pursuits themselves are usually sensitive to planning in their unfolding, even if the choice between them remains underdetermined.
18. Thanks for comments and criticisms to Chrisoula Andreou, Michael Bratman, Sarah Paul, the participants at the SOFIA Workshop in Huatulco 2010 and the audience at the second SLACRR conference in St. Louis 2011.

References

- Andreou, Chrisoula. 2012. Self-defeating self-governance. *Philosophical Issues*, 22(1):20–34.
- Bratman, Michael E. 1987. *Intentions, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA.
- Bratman, Michael E. 1999. *Faces of Intention*. Cambridge University Press, Cambridge.
- Bratman, Michael E. 2007. *Structures of Agency*. Oxford University Press, Oxford .
- Bratman, Michael E. Nov. 2010. Agency, time, and sociality. *Proceedings of the American Philosophical Association*, 84(2):7–26.
- Bratman, Michael E. 2012. Time, rationality, and self-governance. *Philosophical Issues*, 22(1): 73–88.
- Burge, Tyler. 2004. Memory and persons. *Philosophical Review*, 112(3): 289–337.
- Ferrero, Luca. 2006 Three ways of spilling ink tomorrow. In S. Prijic-Samarzija E. Baccarini, editor, *Rationality in Belief and Action*: 95–127. Rijeka.
- Ferrero, Luca. 2009. What good is a diachronic will? *Philosophical Studies*, 144(3):403–430.
- Ferrero, Luca. 2010. Decisions, diachronic autonomy, and the division of deliberative labor. *Philosophers Imprint*, 10(2):1–23.
- Ferrero, Luca. 2012. Review of Holton’s *Willing, Wanting, Waiting*. *Journal of Moral Philosophy*, 9(3):1–5.
- Holton, Richard. 2009. *Willing, Wanting, Waiting*. Oxford University Press.

- Kavka, Gregory. 1983. The toxin puzzle. *Analysis*, 43(1):33–36 .
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford University Press, Oxford.
- Scanlon, Thomas M. 2007. Structural irrationality. In G. Brennan, R. Goodin, and Michael Smith, editors, *Common Minds: Themes from the Philosophy of Philip Pettit*. Oxford University Press, Oxford.
- Tenenbaum, Sergio. 2007. *Appearances of the Good: An Essay on the Nature of Practical Reason*. Cambridge University Press.