


This article was downloaded by: [171.67.216.21]

On: 04 May 2014, At: 17:02

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954

Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Inquiry: An Interdisciplinary Journal of Philosophy

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/sinq20>

Diachronic Structural Rationality

Luca Ferrero^a

^a Department of Philosophy, University of Wisconsin-Milwaukee, USA

Published online: 08 Apr 2014.

To cite this article: Luca Ferrero (2014) Diachronic Structural Rationality, *Inquiry: An Interdisciplinary Journal of Philosophy*, 57:3, 311-336, DOI: [10.1080/0020174X.2014.894272](https://doi.org/10.1080/0020174X.2014.894272)

To link to this article: <http://dx.doi.org/10.1080/0020174X.2014.894272>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form

to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

Diachronic Structural Rationality

LUCA FERRERO

Department of Philosophy, University of Wisconsin-Milwaukee, USA

(Received 15 January 2014)

ABSTRACT *In this paper I investigate whether there are genuine and irreducible pressures of diachronic rationality grounded on the structure of the subject rather than on substantive considerations, such as pragmatic ones. I argue that structural pressures of diachronic rationality have a limited scope. The most important pressure only tells against arbitrary interference with the mechanisms for the retention of attitudes over time. I then argue that in the practical case, a substantial account in terms of the agent's temporal identity appears more promising than a purely structural one, but in the end it still leaves many questions about diachronic practical rationality underdetermined.*

I. Introduction

I.i.

Does rationality impose pressures and constraints on *cross-temporal* combinations of psychological attitudes? It is fairly uncontroversial that rationality governs combinations of *contemporaneous* attitudes—for instance, rationality seems to require that one does not at the same time both believe that p and believe that not- p . It is less clear whether there are also genuinely diachronic rational pressures, which range over cross-temporal combinations of attitudes and are irreducible to merely synchronic pressures.

Consider future-directed intentions. Is there a genuinely diachronic constraint in support of their stability over time? It is obvious that temporal stability is a diachronic phenomenon: it concerns the *persistence* of an intention to x at future time f . However, this pressure might ultimately result from the repeated operation of a synchronic pressure that governs the intention at each particular moment until f . At each moment m until f included, there might be a pressure to intend to do-at- f what at m one

Correspondence Address: Luca Ferrero, Department of Philosophy, University of Wisconsin-Milwaukee, PO Box 413, Milwaukee, WI 53201–6413, USA. Email: luca.ferrero@gmail.com

takes to have most reason to do-at-*f*. If what one takes to have most reason to do-at-*f* remains stable over time, then the *repeated* operation at each moment of the synchronic pressure will result in a stable intention to *x*-at-*f*. The stability is not the result of an independent pressure for an intention to be stable, regardless of or in addition to one's continuous belief in a stable case for the intended action.¹

I.ii.

Even if there are genuinely and irreducible diachronic rational pressures, there is still the question whether they are a matter of *structural* rather than *substantial* rationality.

Claims of structural rationality concern the relations among the psychological attitudes, as governed by the norms constitutive of the attitudes' functional roles and their attributability to a single and unified subject. Pressures toward consistency are typical examples of claims of structural rationality.

Substantial claims concern non-structural considerations that favor or disfavor particular attitudes.² Substantial considerations need not apply to the attitudes in isolation from each other and they often support general policies that apply to sets of attitudes. Even so, their authority does not have the same source as the structural pressures. Unlike the latter, a violation of substantial pressures does not pose a threat to the attributability of the attitudes to a single and unified subject.

In the diachronic case, pragmatic considerations are a typical example of substantial rationality: certain temporal patterns of attitudes are sanctioned as rational on account of their contribution to the maximizing (or at least satisficing) of the satisfaction of the agent's preferences or desires. For instance, the stability of intentions might be supported by its beneficial long-term effects on the agent's success in the pursuit of her goals, rather than in terms of its constitutive role in structuring the psychology of subjects with the ability to engage in temporally extended activities.

I.iii.

In this paper, I will explore the prospects for structural accounts of genuinely diachronic rational pressures. I will assume that there are genuine structural *synchronic* pressures, but my discussion will not depend on the specific shape of these pressures.³ I will use the term 'pressure' to refer generically to

¹For a more detailed discussion of the prospects of a genuinely irreducible pressure for the stability of intentions, see Bratman, 'Time, Rationality, and Self-Governance', and Ferrero, 'Diachronic Constraints'.

²For the distinction between structural and substantial claims, see Scanlon, 'Structural Irrationality'.

³This assumption is denied, for instance, by Kolodny, 'Myth of Practical Consistency'. For a rejoinder, see Bratman, 'Intention, Practical Rationality'.

principles, norms, constraints, requirements, demands, or pressures of rationality.

Although there appear to be powerful substantial considerations, especially pragmatic, in support of diachronic pressures, it is important to investigate whether structural accounts are available.⁴ For structural accounts seem able to address with relative ease skepticism about the pressures' authority, by appealing to their constitutive role: a failure to be governed by these pressures amounts to the metaphysical impossibility of attributing the relevant attitudes to a single and unified subject. This makes it impossible to question the pressures' authority, as long as one is indeed trying to attribute the attitudes in question to a single and unified subject.

The failure of attribution can take two forms. If one has independent reasons to take for granted that a unified and single subject exists, at issue is only whether the attitudes can be attributed to her. However, the existence of a single and unified subject is at least partially dependent on which attitudes can be attributed to her. A violation of the structural pressures can thus ultimately put at stake the very existence of a single and unified subject, rather than the mere attribution of the attitudes.

I.iv.

A structural story can also account for the seeming *peremptory* authority of the pressures. The pressures appear to apply with the same force in each and every instance in which the relevant attitudes are at play. The structural story can account for the peremptoriness by relying once again on the constitutive role of the pressures. In their role as constitutive of the attitudes of a single and unified subject, these pressures *necessarily* govern each and every instance of the attitude for each and every unified subject.

A substantial account cannot respond with the same strength to challenges to the authority of the pressures and secure its peremptoriness since the force of substantial pressures depends on the particular circumstances under which they operate. There might be cases in which there are no substantial grounds and the substantial pressures lack any force. On the substantial account, the authority of the pressures can only be contingent, rather than peremptory and categorical.

I.v.

A structural account does not rule out the existence of substantial considerations in support of particular pressures, at least under a normal range of circumstances. For instance, even if the stability of intentions can be

⁴For an example of a defense of stability in the substantial rather than structural mode, see Paul, 'Diachronic Incontinence'.

grounded on structural considerations, it can still generate pragmatic benefits. Nonetheless, structural accounts, if available, take precedence over the substantial ones because of their universal scope and categorical strength.

Finally, diachronic structural accounts appear plausible by analogy with the synchronic case. Structural accounts are paradigmatically applied to synchronic pressures. At first blush, it seems surprising that a move in the diachronic dimension by itself would make a structural account disappear.

I.vi.

In this paper, I will consider the prospects for genuine and irreducible structural pressures of rationality, with a special emphasis on the practical case. I will argue that although there are some irreducible diachronic pressures, they have only a limited scope. First, there are pressures for the immediate psychological transitions necessary for the processes of securing that one is in the correct synchronic state. Second, there is a pressure against arbitrary interference with the mechanisms of defeasible psychological retention. This pressure is supposed to prevent the temporal isolation of subjects but it is unable to support extended psychological stability and more complex cross-temporal patterns of attitudes.

I will argue that although there are important differences in how to model the diachronic rationality of cognitive and conative attitudes, respectively, in both cases there are no structural pressures that go beyond the defensive demand against arbitrary interference with retention. I will then consider the general prospects for a substantial account of practical diachronic pressures by appeal to the temporal identity of agents as extended beings. I will argue that the substantial account is more promising than the structural one, but it still leaves many crucial questions about diachronic practical rationality largely underdetermined.

I.vii.

Notice that my discussion will not be cast in terms of any specific rational pressure. Here I am only interested in more schematic considerations about the availability of structural accounts in the diachronic dimension. For the same reason, I am not going to talk about any specific synchronic structure. I will assume that there is some basic synchronic structure common to subjects like us and that this structure grounds some synchronic rational pressures. Here, I am only interested in the prospects of a structural account once we take into account the structure of the subject as a temporally extended being.

Before moving on, let me warn the reader about the tone of this paper: it is rather programmatic in nature and many of its claims are conjectural. I wish it could be otherwise, but this is the best that I have been able to do once I have come to realize that my high hopes for a much stronger structural story would have to be dashed.

II. Rationality and Minimal Persistence

II.i.

A time-slice agent, who exists entirely and solely in an instant, can be the object of a structural *assessment*. Her contemporaneous states could be evaluated in terms of their synchronic structural integrity. But she is not under any rational constraint as a *pressure*. The instantaneous existence makes it impossible for the subject to take any step to make her psychology come into accord with the demands of rationality, including making corrections and revisions in response to any violation. For these steps take time.

Any subject who is under a rational constraint as a *pressure* must thus be able to persist throughout the minimal interval necessary to take these steps, when they are called for. The persistence over this interval requires at least the persistence of a core of the subject's psychological states throughout this interval. This default persistence is not governed by structural pressures. It is rather a *precondition* for the operation of rational capacities and the application of rational pressures. In its absence, there would be no persistent subject who could be the target of any pressure whatsoever.

Hence, whether the persistence of any particular attitude counts as conforming to a pressure is something to be determined only against the background of a *minimally* persistent subject. This basic persistence is not something that rationality calls for. It is rather the necessary background for its operation. This is only a very basic and minimal persistence, the one that makes possible for any subject to be more than a mere static and instantaneous time-slice item.

In the rest of this paper, I will only be discussing subjects who are at least minimally persistent, even if only momentarily so. Unless otherwise noted, I will use 'time-slice' to refer to the minimal *interval* of persistence rather than to a static point—to refer to a 'dynamic' rather than a 'frozen' time-slice, so to say.

II.ii.

To look for genuine diachronic pressures, we should now turn our attention to those attitudes that, once acquired, demand the acquisition of additional attitudes. For instance, rationality might demand that when one comes to judge to have conclusive evidence that *p*, one ought to acquire the belief that *p*. If forming that judgment takes time (and if the judgment is indeed distinct from the belief), *at the very moment* when the judgment is first reached, one is not irrational for *yet* lacking the corresponding belief. What rationality demands is the successive step. For instance, if the demand is narrow scope, the subject is to acquire the belief that *p* without further ado.

II.iii.

Pressures for these *immediate transitions* from one attitude to another are structural; they only concern the pattern of the attitudes rather than the specific grounds supporting the individual attitudes. These pressures are diachronic: they govern *processes* of attitudes' acquisition or rejection. But they are *shallow* diachronic: they only govern the immediate steps that are supposed to take place in the minimal interval of basic psychological persistence. They do not support diachronic patterns of attitudes that extend beyond this interval, even if the subject does persist for much longer. For instance, they are not pressures toward long-term stability or any more complex combinations of attitudes over longer intervals.

A subject who never violates diachronic but shallow pressures can persist over a very long interval and yet have a psychology that is highly unstable or with a very narrow temporal horizon, in which combinations of attitudes directly relate to each other only over very short intervals. For instance, a subject who is unable to retain evidence might be perfectly rational in having, at any particular time, the beliefs that she judges to have conclusive evidence for at that moment, and yet constantly be changing her beliefs because of her inability to retain relevant information. If her beliefs are stable, this is only by accident. She sticks to them only because she happens to constantly *re-acquire* from scratch the evidence supporting them, rather than being responsive to a structural pressure for their stability.

III. Long-Term Retention**III.i.**

Consider now a subject with the capacity to *retain* psychological states. Not only does she persist, but her psychology is *cumulative*. The subject need not constantly acquire her psychological attitudes from scratch at each and every moment. Thanks to her power of retention, she saves on scarce psychological resources. She can also *learn*: she can acquire new attitudes by building on those that she has already acquired and retained.

In its basic form, retention is the capacity to make a previously acquired attitude *directly available* to the subject at the present time *as if* she had just acquired it. The attitude is directly available in the sense that it is on a par with the other present attitudes as far as (a) its potential contribution to shaping the subject's thought and conduct at that time and (b) its membership in the overall combination of the attitudes currently subjected to synchronic constraints and to pressures for immediate transitions.

However, the fact that an attitude has been *retained* does not directly bear on its structural rationality. Whether it is rational to have that attitude is a matter of its fit with the subject's other contemporaneous attitudes. Consider someone who has previously acquired and then retained the belief that

p. Whether this belief is rational for her to have now is a matter of its fit with her other contemporaneous attitudes. And this fit is first of all determined by the belief's content, rather than by how the belief has been acquired and retained. In this sense, a retained attitude is by default operating *at face value*, that is, in a transparent manner.

III.ii.

Diachronic structural rationality comes into play when we consider the operation of *defeasible retention*: the power to either reject (or suspend) a retained attitude *prior* to including it in the set of contemporaneous attitudes subjected to synchronic pressures.

Having room for rejecting the deliverances of the retentive faculty does not imply that one is always required actively to endorse the retained attitudes before making them fully operative. The need for constant active acceptance would seriously hinder the work of retention. Under normal circumstances, retention operates in the mode of default acquiescence in its deliverances.

Defeasibility allows the direct management of the retentive faculty and its deliverances. One can reject retained attitudes if one suspects or believes either that (a) the retaining mechanism is not working properly (which affects the accuracy of the retained attitude) or that (b) the accurately retained attitude might be incorrect because of concerns with the subject's expertise and information when the attitude was originally acquired.

How one handles the defeasibility of particular deliverances is not a matter of structural rationality. For it depends on the substantial considerations that bear on whether one ought to worry in those particular cases about either (a) or (b). But it is structurally irrational to abuse the defeasibility by an *arbitrary* rejection of the retained attitude, that is, by rejecting it without even a suspicion that either (a) or (b) is the case.

III.iii.

The arbitrary rejection of a retained attitude is a structural failure. It jeopardizes the possibility of a psychology that is not confined to the present moment. A psychological structure without retention might extend over time but it would do so only by the mere concatenation of time-slice combinations of attitudes, that is, as a result of minimal persistence and its responsiveness to pressures for immediate transitions without any cumulative and learning ability.

At issue in the arbitrary rejection of retained attitudes is not the fate of the attitudes which happen to be rejected, but the threat to the existence of a kind

of psychological structure. Namely, a structure that can non-accidentally support cross-temporal patterns of attitudes and that allows for the working of those attitudes whose characteristic operations unfold over time.⁵

III.iv.

The structural irrationality of arbitrary rejection is genuinely diachronic: the rejection only concerns *retained* attitudes as retained. As such, this irrationality is not reducible to the violation of a synchronic pressure. The pressure against arbitrary rejection targets unjustifiable interference with the default working of retention (henceforth, I will usually refer to it as the ‘non-interference pressure’). It is a *defensive* pressure rather than a constructive one. It does not directly contribute to securing cross-temporal patterns. Rather, it protects the positive but indirect effects of retention mechanisms in their role as providers of the components of other cross-temporal patterns.

III.v.

Retention makes possible two basic patterns—stability and accumulation. First, thanks to retention one can have stable attitudes without paying the price of re-acquiring them from scratch at each and every time. This price is often so high that stable attitudes in the absence of retention would usually be only a lucky accident.

Second, retention makes cumulative patterns possible. By having more available resources for the acquisition of new attitudes rather than the re-acquisition of old ones, the subject can *accumulate* additional (and valuable) attitudes: in particular, the subject can learn. Retention does not simply allow for a more efficient deployment of scarce resources, but it also makes possible to build on top of the attitudes that have already been acquired. For instance, by retaining conclusions previously reached, the subject not only saves on scarce resources that could be used elsewhere but she can make *progress* in the same line of reasoning by building on the previous conclusions.

The extent of stability and cumulation varies depending on the circumstances and the strength of the retention mechanism. Retention can cover

⁵Notice that to forget is not the same as actively to reject a retained attitude. A forgotten attitude is one that has not been retained long enough, rather than one that is rejected *once* retained. Forgetfulness is an even more basic psychological failure than arbitrary rejection but it is not an instance of irrationality. It is rather a failure of the capacity for basic retention, whose proper functioning is a precondition for the application of rational pressures. As Williamson, *Knowledge and its Limits*, 219, writes: ‘forgetting is not irrational; it is just unfortunate’. Which is not to say that agents cannot take responsibility for failures of their faculty of retention if they have the ability and opportunity to monitor for and correct any of its malfunctionings.

temporal intervals of various lengths and range over different domains. What is crucial is that for subjects with limited psychological resources, retention is required for any non accidental stability and cumulation. The pressure against arbitrary interference therefore protects any psychological structure which, thanks to its retentive capacity, has the power to secure stable and cumulative patterns of attitudes.

III.vi.

Unfortunately, retention is not sufficient to secure more complex cross-temporal patterns: those that embrace attitudes over longer temporal intervals and are sensitive to the distribution and relative position of the attitudes within these patterns. Two important kinds of more complex patterns are temporal *sequences* and temporal *unities*. In a sequence, what matters is the temporal order and relative position of the attitudes within an interval, as it happens in a line of reasoning. Sequential patterns allow for things such as progress, timing, and rhythm in combinations of attitudes. In a temporal unity, the pattern is a matter of the overall configuration of various attitudes at different times, which are 'embraced' as a unit over and above their sequential arrangements. A line of reasoning also exemplifies a pattern of temporal unity since it is not simply a matter of the temporal arrangement of the various steps but also an overall arrangement in which the various steps all hang together as a unitary process, on the basis of interweaving cross-temporal ties between both earlier and later stages that go beyond simple lines of (directional) continuity.

Both sequential and unitary patterns are crucial to the psychological life of subjects like us. Retention, although necessary, is not sufficient to make them possible. A subject who retains attitudes over long time intervals might still operate, at any particular moment, within a very restricted temporal horizon, ignoring whatever complex patterns her attitudes might form outside of that horizon. For instance, a subject might have the capacity to learn from her past experiences and update her attitudes accordingly, but operate only in light of immediate concerns and with no interest in securing any sequential or unitary connection between her past, present, and future attitudes.

A subject of this kind would relate to her past attitudes only as parameters within which her current psychological structure is to operate, rather than as elements of an extended psychological makeup. For this subject, there is no problem if the diachronic dimension of structural rationality goes no further than the non-interference pressure.

However, even for a subject with a long-term horizon in the configuration of her psychology, the diachronic dimension of structural rationality might still stop at the non-interference pressure. There might still be pressures that support more complex diachronic patterns of attitudes, but these pressures

might be substantial rather than structural. The structural synchronic pressures together with the substantial demands imposed on her current attitudes might be sufficient to sustain more complex diachronic patterns. These diachronic patterns might be appreciated as such by the subject, so it is no mere accident that her psychology exhibits them, but they would not be grounded on her psychological *structure*. The structural rational pressures that support these patterns are synchronic except for the pressure against interfering with retention.

IV. Cognitive Attitudes and Diachronic Rationality

IV.i.

Consider a subject whose only concern is demonstrating mathematical theorems and who relentlessly and unwaveringly engages in this activity. The activity takes time and she can only take a few steps at any one moment. This subject is not omniscient and her logical abilities and relevant information are limited. But working within those limitations, she can try to make some progress. At any particular time, she has various kinds of beliefs about (a) which mathematical statements have already been shown to be true or false; (b) the history of prior attempts, both successful and unsuccessful, at mathematical proofs; and (c) the extent of her abilities and opportunities for making further progress. In light of this information, she attempts to make some progress in one or more proofs and she immediately updates her set of relevant beliefs accordingly and in response to whatever additional relevant information she might have just received in the meantime.

Such a subject is under standard synchronic pressures, including those that govern the immediate updating of her doxastic states. She is also under the diachronic pressure of noninterference, since she cannot allow to isolate herself from the outcomes of at least some past mathematical investigations.

IV.ii.

Additionally, this subject is under the forward-looking counterpart of the pressure of non-interference. When she acquires and retains beliefs, she is to do so with a claim to their continuous bearing on her future cognitive activity, in the form of a direct although defeasible import on this activity. In other words, she operates on the assumption that, as long as the retention mechanisms are properly working in the background, many of her current beliefs will continue to be available to her future selves as part of the sets of their current attitudes, which will be under synchronic pressures at those times. The import is *direct* since by default the beliefs are supposed to be available in the future ‘at face value’ rather than as mere *indirect* evidence of the past outcomes of prior doxastic activity.

The prospective counterpart of non-interference is a rational pressure against any arbitrary refusal to grant one's current doxastic states a defeasible claim to direct import on one's future doxastic activity. That is, the subject would be structurally irrational if she were to take her beliefs immediately to *expire*, thereby making them unavailable for future retention and use.

Rationality does not require, however, the active retention of *all* beliefs. It is fine if one lets go of some beliefs as soon as they are acquired, especially those deemed to be irrelevant in the future. What is irrational is to stand arbitrarily in the way of beliefs to be made available for retention or to operate with them now ignoring that they should in principle have a direct import in the future. Luckily, it seems quite easy to comply with this rational pressure. All that the agent has to do is to let one's properly functioning retention mechanisms do their work. This prospective non-interference might appear trivial in the case of beliefs but, as we will see later, might be less so in the case of executive attitudes such as intentions.

IV.iii.

Let's return to the diachronic patterns. The subject of our example might be under pressures to engage in temporally extended cognitive activities with complex temporal patterns—including the sequential and unitary patterns of mathematical demonstrations. But these pressures do not come directly from her structure. They are rather substantial demands imposed by the nature of the activity itself. These demands might engage structural rationality in that, at each particular moment, the subject could be under a structural pressure to acquire certain attitudes or to act on account of her *current* appreciation of the case for that activity. These demands, however, are synchronic even if their cumulative effect is to bring about complex cross-temporal patterns.

In order for this cumulative effect not to be accidental, the subject might need to (a) appreciate the extended temporal structure of the activity, (b) have a desire to engage in the activity as temporally extended, and (c) see how her current attitudes and actions help in making the activity progress. But all these diachronic elements are reflected in the *content* of her current attitudes and in the substantial case for them, to which the subject is supposed to be currently responsive on account of the contemporaneous demand of synchronic structural rationality. The scope of structural rationality is thus left unaffected: the subject under discussion can do perfectly fine by relying on synchronic pressures plus the minimal diachronic pressure for non-interference (both in its backward- and forward-looking dimensions).

IV.iv.

At this point, someone might be tempted to suggest that the *diachronic* structure of the subject as a temporally extended being might play a role in

grounding at least the non-interference pressure. But this suggestion turns out to be incorrect. The point of retention is to avoid the subject's cognitive isolation from the past, but not necessarily on account of its being temporally extended. Even time-slice subjects benefit from retention mechanisms, which spare them the costs of starting from scratch at each and every moment. But for time-slice agents, the retained attitudes cannot come from their own past, since they do not have any. What a retention mechanism is supposed to do is to preserve the warrant of an attitude once acquired. Hence, if this warrant can be made in principle directly available to other subjects, there is no reason to deny that there could be mechanisms for the preservation and transmission of warrant across distinct subjects, that there could be *interpersonal* retention.

As Burge persuasively argues, testimony might work in a similar fashion as retention, i.e. as a mechanism for the default interpersonal transmission of warrant.⁶ Given this deep similarity between testimony and memory, I think it is in principle possible for a retention mechanism to work interpersonally but in a way that is more 'intimate' than testimony: interpersonal retention presents the retained attitude with the same psychological and phenomenological immediacy as a standard memory, even if it was originally acquired by a different subject.

This is the kind of retention that could be available to time-slice subjects to help them overcome the risk of isolation from the fruits of past cognitive activities of other agents. Time-slice subjects would thus be irrational as much as temporally extended ones if they were to interfere arbitrarily with reliable (interpersonal) retention mechanisms.

If this suggestion turns out to be correct, the structure of an extended subject would not provide the grounds for pressures of retention. For these pressures would be already in place for time-slice subjects, and they would look the same even for temporally extended subjects. The pressure is for the arbitrary non-interference with whatever retention mechanisms are deemed to preserve warrant, regardless of whether they operate intra- or interpersonally. The temporal extension of the subject would not be necessary to provide the connection with the past for the purposes of pure cognition.⁷

IV.v.

The possibility of interpersonal retention shows that the enterprise of mathematical demonstration can extend over time in a fully interpersonal

⁶See Burge, 'Content Preservation', and Burge, 'Memory and Persons'.

⁷This is not to deny that for each individual subject the intrapersonal mechanisms might often work better, being more reliable, accurate, more readily available, and less costly to access and maintain. But these are only differences in degrees of functionality, which do not show that interpersonal retention is inconceivable.

and thereby impersonal mode. Different subjects can at different times contribute to the activity entirely on the basis of the conclusions of prior steps taken by other subjects, conclusions that are offered to them not just by testimony but also by interpersonal retention. Even time-slice agents can thus take part in this temporally extended cognitive investigation with instantaneous one-off contributions. These contributions are undertaken in light of an appreciation of their role in the extended enterprise but they are not governed by diachronic pressures. They are only governed by synchronic pressures over attitudes whose content is about the extended enterprise.

Hence, even a temporally extended agent, when engaged into such an extended enterprise, is not under genuinely diachronic structural pressures. Each of her momentary steps is governed by synchronic pressures *as if* she were a time-slice subject. Her extended identity is never at stake in responding to the pressures that govern this activity.⁸

For the kind of subject we have discussed thus far, who is single-mindedly set on a specific cognitive enterprise, there is no need (nor room) for structural diachronic pressures that go beyond those that regulate the immediate transitions and the non-interference with retention. There can be diachronic elements in the content of her attitudes (such as her care for contributing to the extended cognitive endeavor), but these elements do not impose genuine and irreducible diachronic structural demands on her. These diachronic elements are part of the substantial matters to which the subject must attend at any particular time by complying with the synchronic pressures which, at that time, govern her contemporaneous psychology and conduct.

IV.vi.

How does this conclusion bear on subjects like us, with more complex cognitive and conative sets of attitudes? Let's first expand the target of the cognitive enterprise in which our subject is engaged, so that it now includes various empirical investigations. Does this broadening in the scope of the investigation introduces any new diachronic pressure? My conjecture is that it doesn't.

True, the progress of the cognitive investigation becomes more dependent on the subject's spatio-temporal location. The preservation of warrant

⁸A subject might engage in the enterprise in a stronger personal mode, i.e. only in so far as it is her own enterprise—even going as far as trying to thwart any impersonal version. But this would bring in her temporal identity as a substantial element of that specific form of the enterprise, not as part of the structural constraints that govern it. (For a discussion of the different forms in which one's own identity might enter into the content rather than the structure of one's intentional activity, see Ferrero, 'Can I Only Intend'.)

might reflect this—most likely by giving prominence to intrapersonal mechanisms since they are the ones that are guaranteed to keep track of the subject’s trajectory in space-time and might more easily update the spatio-temporal indexes of the attitudes’ contents. But this increased reliance on one’s own retention mechanisms does not introduce new structural pressures. After all, it is in principle still possible to rely on the preservation of warrant coming from other subjects. And one’s present contribution to the progress of that particular cognitive enterprise can be entirely regulated by synchronic pressures. Hence, it does not appear that by broadening the scope of cognition, one brings into play the structure of the subject as extended over time. But if this structure is not brought into play, then there is no basis available to try to ground more complex diachronic structural pressures.

IV.vii.

This conjecture is supported by the regulative ideal of pure cognition. The goal of pure cognition is the formation of as accurate and objective as possible a conception of the *one* world we all inhabit. Hence, there is a pull toward a convergence in the cognitive states of different subjects, against the background of a common intention to pursue the cognitive enterprise. Actual cognition falls short of this regulative ideal for a variety of reasons—including its impure character, i.e. its being pursued together with (and often at the service of) non-cognitive aims. Nonetheless, in their most general formulation, the rational pressures on cognitive attitudes still reflect the aspirational features of the pure form. This is why I suggested that the pressure for retention takes an impersonal form that allows for interpersonal retention. And this goes together with the possibility of an extended cognitive enterprise that takes place by a sequence of momentary contributions of distinct time-slice rational subjects.

IV.viii.

The main concern in dealing with the past is to avoid being cut off from the fruits of prior good cognitive labor, regardless of whom performed it. This is because it is impossible for limited beings like us to make progress by starting from scratch at each and every time. But once reliable retention mechanisms are in place, the cognitive job of the subject can be carried out at any particular time in terms of the momentary operations on her current combination of cognitive attitudes alone *as if* she were a time-slice cognitive agent.

IV.ix.

The basic pressures of cognitive rationality can thus be said to be those of a ‘time-slice epistemology’, in two senses: first, the pressures are *synchronic* and their grounds lie entirely in the subject’s current attitudes as if she were a time-slice subject. To the extent that past and future attitudes have any influence, this is only indirect and always mediated by the subject’s present states.

Second, the pressures in their minimal diachronic extension are *impersonal* as shown by the demand against arbitrary interference with retention. This demand applies equally both to intra- and interpersonal retention mechanisms, in the same manner in which it would apply to time-slice subjects. It follows that, intrapersonal stability and complex cross-temporal patterns of cognitive attitudes are, from the structural point of view, no more than by-products of compliance with a time-slice epistemology, even for temporally extended subjects like us. At the very least, this seems to be so when we are concerned with pure cognition and those rational pressures that are modeled on the regulative ideals of such cognition.⁹

V. Executive Attitudes and Diachronic Rationality**V.i.**

What happens once we take into consideration the conative side of complex psychological structures? Could the structural dimension of *practical* rationality be exhausted—in analogy with the doxastic case—by a combination of (a) synchronic pressures, (b) pressures for immediate transitions, and (c) pressures against arbitrary interference with (potentially interpersonal) retention mechanisms of conative attitudes? In other words, could we be satisfied with a time-slice structural practical rationality, by analogy with a time-slice epistemology, even for subjects like us, who are extended in time and have a more complex psychology?

V.ii.

To address this question, we must start by considering whether there is even room for a non-interference pressure with the retention of *conative* attitudes. The kind of retention in question is not just the memory of the fact that one has priorly acquired conative attitudes. This kind of memory is just an instance of cognitive retention, which is subjected to the non-interference pressure we have already discussed. What we are interested in is retention of the conative attitudes as having default direct import in their role as conative attitudes.

⁹For a programmatic defense of time-slice epistemology, see Moss, ‘Time-Slice Epistemology’.

V.iii.

Not all conative attitudes seem fit for a rational pressure of this sort. Although we might welcome stable desires, it does not seem that a desire as such is under any rational structural pressure to be stable. At any particular moment, the agent might be under a synchronic pressure to have a certain desire, say, because of a pressure for consistency with her other cognitive and conative attitudes at that time. But whether she is still under such pressure at a later time is a matter of whether there are pressures for the retention of these other attitudes, rather than for the desire as such.

The conative attitudes whose retention is paradigmatically fit for a retention pressure are *future-directed executive* attitudes, such as future-directed intentions. These are attitudes whose primary role is the direct control of the agent's conduct (which makes them executive) and whose satisfaction lies at least in part in the future (which makes them future-directed). Because of the latter, the attitudes call for retention. If not retained, the attitude is guaranteed to fail unless the agent has the ability and resources to reacquire it from scratch, which is unlikely to happen for limited agents like us. Hence, to the extent that there are indeed pressures of conative retention, it seems that they would have their characteristic application over future-directed executive attitudes.

V.iv.

Let's then consider how retention is supposed to operate on future-directed executive attitudes (henceforth, unless otherwise noted, I will use 'executive attitudes' to refer only to the future-directed ones). By analogy with the cognitive case, one might suggest that there is a structural pressure against arbitrary interference with retention mechanisms that are deemed to preserve the executive attitude's warrant, regardless of the identity of the subject who first acquired the attitude. But this suggestion immediately raises some serious concerns.

First, it seems that one's autonomy is violated if one were to allow the executive attitudes formed by others to have a direct import on oneself.

Second, one might have qualms about granting default direct import even to executive attitudes formed by one's past self. For two reasons: (a) the pervasive instability of many motives and (b) the effects of the mere passage of time even over stable motives. The agent's movement in time might automatically affect the temporal horizon of her concerns because of the constant movement of the origin of her temporal perspective, including the shape of the curve and the origin of her future discounting, which might change in non-linear ways.¹⁰

¹⁰For an introduction to 'present-based preferences' and effects on diachronic rationality, see Andreou, 'Dynamic Choice'.

Last but not least, there is always the possibility of estrangement from one's past because of the need to respect the subject's *current* autonomy. The agent at the present moment is the locus of the basic exercise and affirmation of her self-determination. This might be sufficient to generate a concern with the suggestion of relinquishing control even to one's own past self, on account on nothing other than his being in the past.

All of these considerations raise questions about the existence of a pressure of non-interference with the retention of executive attitudes.¹¹

V.v.

In the cognitive case, the demand for non-interference comes from the worry that a subject might be cut off from the past, becoming unable to enjoy the fruits of prior cognitive labor. Because of the regulative ideal of cognition, these fruits are in principle shareable over time, both intra- and interpersonally.

By contrast, in the conative case, there is at least an initial push in the opposite direction, toward the isolation of the subject at the present time both from other subjects and from her other temporal selves. In the conative case, there does not seem to be a regulative ideal that pulls in the direction of the interpersonal sharing of executive attitudes. Although the agents' lives ultimately unfold in one and the same world, each agent needs to carve herself as an individual and separate portion of that world, a portion that can take many different shapes and forms. There is thus a basic metaphysical push within agency toward isolation and separation, toward securing one's individual space as a distinct agent. And this push points even in the direction of the temporal isolation of the agent at the *present* moment, given that the agent's own present time is the primary locus for the actual exercise of direct and unmediated self-determination. This is reflected in the initial push toward the unshareability of executive attitudes (both inter- and intra-personally) since these are among the attitudes which are primarily responsible for giving shape to each individual agent as a distinct locus of self-determination.

The direction of this push is not absolute and irreversible. Ultimately, cognition and conation are two aspects of an integrated agent. The need to integrate them generates its own pressures, which might reverse or limit the opposing drives of cognition and conation. But in exploring the nature of structural rationality it is crucial to be sensitive to the basic metaphysical

¹¹The same considerations also raise a concern with the idea of a default direct import of executive attitudes on the agent's own future. These attitudes can still be directed at future conduct but the agent might no longer be justified in seeing them as able to control her future conduct by direct future import, that is, as operating immediately in the future as if they had just been acquired at that time. This means that at most the agent can take her future-directed executive attitudes to be able to influence her future conduct *indirectly*—say, by working as a kind of pre-commitment.

pulls exercised by the different aspects of the subject, since they are going to be reflected in the basic structural pressures.

V.vi.

The isolationist push within agency does not stand in the way of the agent's enjoyment of the products of the cognitive and deliberative labor of other subjects or selves. It would be irrational for an agent arbitrarily to interfere with retention of those states that would help her save on deliberative resources and avoid starting from scratch. But for this retention to be acceptable, it has both to secure the subject's current autonomy and to provide states whose content is justifiable from the agent's current *practical standpoint*, that is, the current combination of her preferences, cares, concerns, and values.

The first matter is procedural. The agent at the present time has veto power on allowing the retained states to operate in her; she has the option to either suspend or reject them. The presence of this capacity and the opportunity to exercise it might be sufficient to respect her current autonomy, even if the power is not actively exercised and the agent does not actively and explicitly endorse the retained attitudes. If the capacity and the opportunity to reject the attitudes were all that it would take to respect autonomy, in principle it could be possible to retain even the executive attitudes coming from other subjects, provided that they were deemed justifiable from one's current practical standpoint.

Whether the retained executive attitudes would be justifiable is, instead, a substantial matter. In order for the agent to take a retained executive attitude as warranted, she must deem it endorsable from her *current* practical standpoint in light of her view of her current circumstances. The standpoint need not be restricted to the present moment—the agent might have long-term cares and concerns—but it is *centered* at the present moment. It is the standpoint that the agent has *now*.¹²

V.vii.

A conative standpoint is *personal* in a way in which a cognitive standpoint need not be. By cognitive standpoint, I mean the complete set of the subject's cognitive states at a particular time. This standpoint is a function of the subject's current spatio-temporal location and of the path taken to get there. With respect to some cognitive states, such as perceptual beliefs, this standpoint is perspectival: the agent has these beliefs only in virtue of the position she occupies. In keeping with the regulative ideal of pure cognition, I

¹²Likewise, what she considers acceptable might depend on her expectations about her future circumstances but only to the extent that her *present* standpoint makes her to have any concern about those circumstances.

maintain that a cognitive standpoint is shareable with other subjects. Other agents could in principle occupy the same cognitive standpoint—including cognitive states with indexical content that is centered on the subject's present location, such as the belief 'I am at such and such a location now'.

The only exception are states whose content concerns one's own temporal identity in a first-personal mode of presentation. Consider the cognitive position of subject *S* at spatio-temporal location *l*. A different subject *T* might come to occupy the same location *l* and take on the same cognitive standpoint, but he would not thereby be able to share in *S*'s belief 'I am now the same as *R* at time *t*' (although *T* might come to share the same belief in its third-personal form as '*S* now is the same as *R* at *t*'). If my conjecture about the regulative ideal of cognition is correct and all we need is a time-slice epistemology, this limitation would not matter for structural rationality. For a time-slice epistemology is impersonal and thus the identity of the subject in the first-person mode is immaterial to the application of basic pressures of cognitive rationality. In other words, the subject's cognitive standpoint is neutral with respect to the subject's own temporal identity: different subjects might thus come to share the same cognitive standpoint without necessarily losing their identity.

V.viii.

By contrast, a practical standpoint includes elements (such as cares, concerns, values, and their temporal horizons) that play a deeper and more fundamental role in the determination of the agent's own identity. They contribute to giving her a distinctive shape as that particular item in (and portion of) the world that she strives to be in her activity of self-determination. This puts a limit on the shareability of the practical standpoint. Coming to actually occupy a different practical standpoint might amount to a loss of identity. Different agents might imagine what occupying that standpoint might be like for an agent, but they might be unable to imagine *themselves* occupying it as the self-same agents.

Whereas the ability to occupy the same cognitive standpoint is essential to the regulative ideal of cognition, the ability to *fully* occupy the same practical standpoint is accidental from the point of view of agency.¹³

¹³Let me stress that at issue here is the impossibility of 'fully' occupying the same practical standpoint and how this impossibility relates to the determination of the agent's identity. I am not denying that important parts of the standpoint might be shareable across distinct agents and that they might actually provide the grounds for important normative demands, such as those of morality (see for instance the discussion of the shareable reasons in the work of Christine Korsgaard). There might actually be a push within the regulative ideal of agency for this much sharing. But this sharing would operate in the *interpersonal* dealings among distinct individuals on account of the very recognition of their being separate loci of autonomous agency; a separation that, if I am correct, is in large part due to the unshareability of their *full* practical standpoint.

The elements of the cognitive standpoint offer both the parameters and the material for the operation of cognitive activity, against a shared practical concern for it, at least in its pure and idealized form. The elements of the practical standpoint, instead, determine what projects are worthwhile for the agent, what she is up to as the particular agent that she is. And this determination takes place at the present time, the one time when the agent is able to exercise, if ever, the power of direct self-determination. This is why, in the practical sphere, even the transtemporal sharing of a standpoint within oneself is not to be taken for granted.

V.ix.

How does the isolationist push of agency affect the prospects of diachronic practical rationality of the structural kind? To begin with, the push has a restrictive influence on the pressure for non-interference with retention of executive attitudes. Unlike the cognitive case, it is much more demanding for executive attitudes to meet the conditions for the preservation of warrant. This warrant is ultimately a function of the practical standpoint of each particular agent at her present time, a standpoint that, unlike the purely cognitive one, is under a pull to isolate each agent at that time as a distinct center of momentary self-determination.

The second problem is that the pressure against non-arbitrary interference is necessary but not sufficient to support more complex cross-temporal patterns of attitudes. The pressure only allows the current self to take advantage of the division of deliberative labor but it is silent on the diachronic features of the current self's practical standpoint and their possible effects in support of more complex diachronic pressures.

V.x.

There are two aspects of the practical standpoint that matter for the prospects of diachronic structural rationality. First, the agent's present standpoint should look favorably on some extended projects to which she might presently contribute. But even if it does, this would not yet be sufficient to secure that the agent supports the progress of an extended endeavor via the direct import of executive attitudes priorly acquired. For that to happen, a relatively stable practical standpoint is required, one that is shared by the subject's temporal selves throughout the endeavor.¹⁴ This is more demanding than the mere division of deliberative labor. In this division of labor, it is sufficient that the earlier self, acting as a kind of advisor, prepares and

¹⁴I take this suggestion to be in the spirit of the 'sharing of normative perspective' defended by Morton, 'Deliberating for Our Far Future'.

transmits the executive attitude from the standpoint of the recipient—the present self in an advisee’s role. But there is no demand that the advisor should share the same standpoint as her own. A discrepancy between the actual practical standpoint of the advisor and that of the advisee need not weaken or cancel the warrant in the retained attitude, assuming that the advice is issued from the point of view of the advisee.¹⁵

The discrepancy in standpoint between an earlier and a later self is problematic, however, when both selves are supposed to partake in the same extended activity. The discrepancy makes impossible to engage in the activity in the form of *internal unity*, that is, by the non-manipulated collaboration of the selves out of a shared appreciation of and desire for the extended activity as extended. If the standpoints do not agree with each other in the relevant ways, the agent might still be able to pursue the extended activity but only in the mode of *external* temporal unity. That is, this activity would be unified only via the manipulation of the later selves, on account of their lack of direct support for it, or simply by the accidental convergence on the required momentary steps, which are taken by later selves out of independent motives (including the pursuit of different activities which just happen to demand the same steps at those times).¹⁶

VI. Temporal Identity and the Substantial Account

VI.i.

Given the need for a shared practical standpoint as the necessary background for extended activities that exhibit complex temporal patterns, rational pressures in favor of these more complex patterns can be structural only if they can be grounded on a psychological structure that supports a stable practical standpoint.

Could our transtemporal identity offer the required psychological structure and be a sufficient basis for a stable practical standpoint? Unfortunately not. Under any plausible criterion of personal identity over time, the preservation of identity over a sufficiently long interval is compatible with massive changes in practical standpoint. The possibility of these changes stands in the way of establishing a default direct import of executive attitudes at least over sufficiently long intervals, since retention mechanisms cannot guarantee the default preservation of warrant in the executive attitudes over those longer intervals.

Nevertheless, over much shorter intervals, identity might still seem to secure the basis for some conative retention, given that a practical standpoint

¹⁵For further discussion of the division of deliberative labor, see Ferrero, ‘Decisions, Diachronic Autonomy’.

¹⁶For a more extensive discussion of temporal unity, see Ferrero, ‘What Good Is a Diachronic Will’.

is less likely to change over a shorter interval. Even so, this pressure would still fall short of securing long-term stability in executive attitudes, let alone more complex cross-temporal patterns. This is because the pressure for retention could only span rather short intervals and be compatible with the rapid accumulation of large and possibly radical changes in practical standpoint. Not to mention that in many cases a single and immediate change in one's motives and inclinations might be sufficient to threaten the warrant in the retained attitudes, as it happens when the agent is under the strong pull of temptation as she perceives an immediate opportunity to indulge in it.

To sum up, the agent's transtemporal identity might provide the basis for a structural pressure against the most extreme form of conative isolation since it opposes arbitrary interference with immediate conative retention. But this identity is unable to provide the basis for structural pressures in favor of long-term stability and more complex cross-temporal patterns of executive attitudes.

VI.ii.

The appeal to the agent's temporal identity might still help with grounding stability and complex patterns in the *substantial* mode. The condition of being one and the same agent over time might help by setting the proper *unit* for the assessment of the costs and benefits of having stable and complex patterns of executive attitudes. Even radical changes in the practical standpoint are compatible with the preservation of identity. But what matters from the substantial point of view is how the agent handles these changes so as to gain long-term benefits over one's life as a whole, rather than focusing on the benefits that only accrue to shorter portions (if not even just to the present moment).

The substantial pressure does not require that the different stages of the agent start by sharing a practical standpoint. What it requires is rather the agent's own appreciation of her being one and the same agent over a temporally extended life. The rational pressure can then operate at any particular time to retain, revise, and stabilize—if necessary—the elements of the practical standpoint at that time. The agent might realize that if she were to change some of her preferences, cares, concerns, or their temporal horizons, she would put in jeopardy her ability to attain the benefits of stable and complex patterns of attitudes and related conduct. Hence, the substantial pressure if successful can help bring about a shared and stable practical standpoint against the assumption of the agent's identity over time. This is an important difference with the structural account, which instead finds the ultimate grounds of authority of the rational pressures in the threat that a failure to be governed by them poses to the agent's own continued existence.

VI.iii.

In contrasting the substantial with the structural account, we need not specify the form taken by the latter. For instance, it does not matter whether the substantial account is cast in terms of maximization rather than satisficing. Likewise, for present purposes, it does not matter how the substantial account specifies the currency in which the costs and benefits are calculated or what is to be counted as a genuine cost or benefit. What we are focusing on is only the suggestion that the substantial account might rely on the agent's temporal identity as setting the unit within which to calculate the benefits of stable executive attitudes and complex cross-temporal patterns.

VI.iv.

The existence of substantive benefits of this kind for beings like us is fairly uncontroversial. What is at issue, however, is whether these benefits provide the basic support for stability and complex patterns *in lieu of* a structural account. The problem in giving up the structural account is that we are unable to show that there are genuine diachronic pressures of the categorical and unconditional form (but for the merely *defensive* pressure of non-interference with retention).

In the substantial story, the support for stability and cross-temporal complex patterns ultimately depends on the circumstances of the agent. There is no guarantee of a necessary rational pressure toward stability and complexity. The pressure exists only as long as (and to the extent that) one is justified in expecting a net gain over one's life. This conclusion prompts us to explore whether there might still be some other avenue to support a structural account of diachronic practical rationality.

VI.v.

It might seem that the obvious place to look is the planning character of future-directed executive attitudes. According to Bratman, planning attitudes are subjected to pressures of stability.¹⁷ Hence a failure to abide by these pressures is by definition a *diachronic* failure. But is it also a structural failure? At first, it might seem so: the pressures of stability are constitutive of planning attitudes, hence these attitudes cannot be attributed to a subject who is not responsive to these pressures.

These considerations, however, are not sufficient to support the structural account. The structure that must be at stake is that of the *subject*, not just of her attitudes. The question to ask is whether the instability in the purported planning attitudes would threaten the subject's own structure. It does not

¹⁷See Bratman, *Intention, Plans and Practical Reason*.

seem so. True, without stability the agent's life might not fare as well, but she might still be recognizable as having a sufficiently coherent and consistent psychology even in the face of the impossibility of being attributed (stable) planning attitudes. The subject might still be criticizable as structurally irrational on account of this impossibility. But this is so only insofar as the subject has other attitudes that call for her having, at the same time, planning attitudes as well.

For instance, if the agent presently cares about being a planning agent, she might be synchronically irrational for not having, at this time, the required planning attitudes. The irrationality concerns the misfit between her *present* care for having planning attitudes and her failure to meet *now* the demands imposed by the constitutive standards of planning attitudes. This is the misfit that threatens the subject's structure. Notice that meeting the attitudes's structural demand is not a *categorical* demand on the agent as such. In this example, it is conditional on her caring about being a planning agent. What is unconditional is only the synchronic pressure to meet the constitutive demand of the planning attitudes, *given* one's care for being a planning agent.¹⁸

VI.vi.

Even if structural rationality does not categorically require that we have planning attitudes, there might be a very robust and strong substantial case for them. Given our nature as limited temporal agents, it seems that—barring some very special circumstances—it is always in our best interest to develop, maintain, and exercise a capacity for planning agency as a general-purpose good; a good that greatly benefits temporally extended lives considered as a whole. This support is conditional on our normal circumstances, the general features of our lives as temporally extended, our basic care for our lives to fare well, and our taking our lives as the standard unit for the determination of the overall benefits that accrue to us. However, these appear to be among the most general and common features of human lives. Hence the substantial support for stable and complex patterns of executive attitude, although conditional, might turn out to be quite *robust*.

If so, to concede that we might be unable to ground the purported peremptory character of the diachronic pressures might not be too troubling. For the substantial account might still be able to deliver something quite close in scope and strength to what the structural account was supposed to give us.

¹⁸I am leaving open whether the synchronic pressure is wide or narrow scope, since this distinction is immaterial to the present discussion.

VI.vii.

Nonetheless, there is still a worry with the reliance of the substantial account on the agent's identity over time. Is this identity the proper unit for assessment of the benefits of planning? The benefits are long term, so we have substantial reasons for some stable and complex cross-temporal patterns. But why think that the proper unit should embrace all of one's life? Why couldn't the beneficiaries be just some extended portions of one's life? As long as they are sufficiently extended, they seem to be able to gain the benefits of planning. Once this option is on the table, however, we face a host of difficult issues: Which portions matter? How are we going to handle conflicts among them, which are very likely to arise? Which rate and shape of discounts should we use, if any, in addressing these conflicts?

These are just some of examples of the questions that we would be facing. And in addressing them the appeal to one's identity over time no longer helps: The portions in question are all portions of the same life, after all. The substantial account must thus find additional basis to address all the complexities of diachronic rationality. Just appealing to the general-purpose good of planning is not going to get it very far. It only tells us that we'd better have a psychology with some stability and cross-temporal complexity, but it otherwise leaves the substantial account largely underdetermined.

This might not be a fatal blow to the substantial account. After all, a structural one, if available, would be unlikely to fare any better, given that it could only appeal to the structure of the agent in its most general and generic form, a structure that could be equally exhibited in that form by the different conflicting portions of the agent's own life. But it is important to sound a note of caution about the limitations of the substantial account, and about how little we might accomplish if all that we can say about substantial diachronic practical pressures must ultimately be grounded on nothing other than the agent's identity over time. Figuring out where we should look next to find additional and more promising grounds, however, is something to be left for another occasion.

Acknowledgements

I thank the audience and the participants at the Center for the Study of the Mind in Nature workshop in Oslo (including Chrisoula Andreou, Michael Bratman, Sarah Paul, and Wlodeck Rabinowick). I also thank Sergio Tenenbaum for organizing it and for editing this volume, and Olav Gjelsvik for his hospitality. A special thank you goes to Jennifer Morton for a stimulating discussion that prompted me to start thinking harder about the issues in this paper.

References

- Andreou, Chrisoula. 'Dynamic Choice'. In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Fall 2012 edition. 2012. <http://plato.stanford.edu/archives/fall2012/entries/dynamic-choice>
- Bratman, Michael E. *Intention, Plans, and Practical Reason*. London: Harvard University Press, 1987.
- Bratman, Michael E. 'Intention, Practical Rationality, and Self Governance'. *Ethics* 119, no. 3 (2009): 411–43.
- Bratman, Michael E. 'Time, Rationality, and Self-Governance'. *Philosophical Issues* 22, no. 1 (2012): 73–88.
- Burge, Tyler. 'Content Preservation'. *The Philosophical Review* 102, no. 4 (1993): 457–88.
- Burge, Tyler. 'Memory and Persons'. *Philosophical Review* 112, no. 3 (2004): 289–337.
- Ferrero, Luca. 'What Good is a Diachronic Will?' *Philosophical Studies* 144, no. 3 (2009): 403–30.
- Ferrero, Luca. 'Decisions, Diachronic Autonomy, and the Division of Deliberative Labor'. *Philosophers' Imprint* 10, no. 2 (2010): 1–23.
- Ferrero, Luca. 'Diachronic Constraints of Practical Rationality'. *Philosophical Issues* 22, no. 1 (2012): 144–64.
- Ferrero, Luca. 'Can I Only Intend My Own Actions?'. In *Intentions and the Own Action Condition*, ed. David Shoemaker, 70–94. Oxford: Oxford Studies in Action and Responsibility, 2013.
- Kolodny, Niko. 'The Myth of Practical Consistency'. *European Journal of Philosophy* 16, no. 3 (2008): 366–402.
- Morton, Jennifer M. 'Deliberating for Our Far Future Selves'. *Ethical Theory and Moral Practice* 16, no. 4 (2013): 809–28.
- Moss, Sarah. 'Time-Slice Epistemology and Action Under Indeterminacy'. *Oxford Studies in Epistemology*, forthcoming 2014.
- Paul, Sarah. 'Diachronic Incontinence is a Problem in Moral Philosophy'. *Inquiry* 57, no. 3 (2014): <http://dx.doi.org/10.1080/0020174X.2014.894273>.
- Scanlon, T.M. 'Structural Irrationality'. In *Common Minds: Essays in Honor of Philip Pettit*, eds. Geoffrey Brennan, Robert Goodin, Frank Jackson, and Michael Smith, 84–103. Oxford: Oxford University Press, 2007.
- Williamson, Timothy. *Knowledge and its Limits*. Oxford: Oxford University Press, 2002.