Routledge
Taylor & Francis Group

Intention inertia and the plasticity of planning

Piotr Makowski

Department of Social Sciences, Institute of Philosophy, Adam Mickiewicz University, Poznań, Poland

ABSTRACT

In this article, I examine Michael Bratman's account of stability in his planning theory of intention (PTI). Future-directed intentions should be stable, or appropriately resistant to change, over time. Bratman claims that the norm of stability governs both intentions and plans. The aim of this article is to critically enrich Bratman's account of stability by introducing plasticity as an additional norm of planning. I construct plasticity as a kind of stability of intentions which supplements Bratman's notion of "reasonable stability." Unlike the latter, plasticity applies mainly to cases in which plan states are abandoned without reconsideration. I focus on the intra-theoretical problems of PTI and elucidate: (1) the distinction between future-directed intentions and plans, (2) the conceptual difference between stability and inertia, which is only implicit in PTI, and (3) the role of the environment of the planner, which has a vestigial role in Bratman's work. I also defend my incorporation of plasticity against one possible objection and support it in the context of Bratman's later works. Although critical, my proposal is in moderate harmony with PTI.

1. Introduction

In his early work, Michael Bratman, following Gilbert Harman (1976), introduces the concept of intention *inertia*. The term has been used as a synonym of intention *stability*. Future-directed intentions are stable, meaning, roughly, that they are "defeasible defaults" (Bratman, 1987, p. 16). This stability over time should involve an appropriate resistance to change (Bratman, 2009, p. 153; 2013, pp. 51–52). In the so-called *planning theory of intention* ("PTI"), stability is one of the pivotal norms of the rationality of diachronic agency.

Bratman (1987, chapter 5, 1992) offers a subtle and detailed analysis of stability. Intention stability and the problem of its revision over time has attracted considerable interest, both in the form of refinement and critique (DeHelian & McClennen, 1993; Ferrero, 2010; Harman, 1976; Holton, 2004; Kolodny, 2008; Mintoff, 2004). Interestingly, in response to the criticism and feedback received, Bratman (1999a, 1999b, 2007) has been remarkably précising and changing his views. Nevertheless, a few aspects of his account seem to remain underdeveloped or only implicit. What exactly is the rationale for treating 'inertia' as a synonym of 'stability'? Are the conditions for the stability or inertia of an intention always the same as the conditions for the stability of a plan? If this is not the case, how does it affect Bratman's account of the overall psychological stability of agents?

In this paper, I try to answer these questions regarding the intra-theoretical aspects of Bratmanian philosophy. First, I will address Bratman's identification of intentions with plans and connect it to

the conceptual difference between inertia and stability. Second, and most importantly, I will expand Bratman's account of stable agency to incorporate cases in which intentions are abandoned without being reconsidered. Such cases are purposely ignored by Bratman. I will argue that any theory of planning in the philosophy or psychology of agency cannot be complete without considering the agent's environment and, drawing inspiration from the work of Tadeusz Kotarbinski (1961, p. 195) of the Lvov-Warsaw School, I introduce the norm of planning *plasticity*. I shall defend plasticity as a norm that is peculiar to plans (but not to single intentions), especially in unstable agential ecosystems in which agents frequently abandon their intentions without reconsideration. The banal but philosophically non-trivial idea behind my account is that we are agents who always plan our actions in an environment, as embodied agents immersed in a natural and social world, and this seemingly innocent factor determines the nature of the stability of our agency.

This paper is a critical addendum to Bratman's theory—a corrective supplement in moderate harmony with its overall spirit. In the last part of this work, I will defend my account of planning plasticity against one criticism that might be raised from the perspective of Bratman's later work, in which stability is connected with self-governance.

2. Interpreting PTI: Future-directed intentions and plans

Future-directed intention, the intention at t_1 to F at t_2, has been studied by many authors, and Bratman's is the most influential account. Why is future-directed intention interesting? In Bratman's words, "Why bother with intentions for the future? Why don't we just cross our bridges when we come to them?" (1992, p. 2).

Bratman cites two pragmatic reasons. First, we are resource-limited agents. Our deliberation takes time, so it is not a surprise that our prior intentions shape, as default options, our future actions. Second, intentions directed at the future support the organization of our actions both *intra*personally and *inter*personally (Bratman, 1987, p. 18; 1992, p. 2). When, for example, I intend now to go to Japan next month, I do not merely express my wish or desire to do so. I practically commit myself, in a particular way, to the course of the intended action. This is what Scanlon calls the "predictive significance of intent" (2008, p. 13). When I intend something, this changes what the future is likely to hold. But this commitment is not irrevocable. "Given new information, or a change in what I want, I may well reopen the question and reconsider" (Bratman, 1987, p. 16). The commitment of intention is not irrevocable, but as a facilitator of the control of future conduct, it resists instant change and reconsideration:

> If I now intend to go to Tanner later today, I normally see … the question of whether to go as settled and continue so to intend until the time of action. My intention resists reconsideration: it has a characteristic *stability* or *inertia*. (Bratman, 1987, p. 16)

Stability allows intentions to serve as an entry point for further reasoning. Prior intentions prompt reasoning about and the formation of further intentions (Bratman, 1987, p. 17). So, we have three basic and interconnected features of future-directed intentions, understood as "pro-attitudes" that carry commitment: (1) *conduct-control*, (2) *inertia* or *stability*, and (3) *further-reasoning-input*.

The PTI posits plans, in the sense of "having a plan," as constructs that help to explain the character and role of future-directed intentions. Consider my intention to go to Japan. To fulfill this intention, I need to carry out a number of actions. I need to check my schedule, buy tickets, book a hotel, and the like. I also need to know that checking my schedule precedes buying tickets and not the reverse. Therefore, I need to *plan* my actions, and this planning seems to share the features of future-directed intention—action-control, inertia, and further-reasoning-input. According to Bratman, plans have two more properties. First, plans are always *hierarchically structured*. My goal to visit Japan embeds the appropriate means to achieve that goal. The overall plan must contain subplans as stages of the process of reaching the goal. Second, I do not have to settle my plan all at once, as though it were a complete and full set of steps. Plans are generally *partial*. I do not need to know, at the time of planning, when

and where to buy the tickets, which hotel to book, and so on. Such choices can be left for later consideration when I am filling out the details of my plan. I may also have a very general plan—something like a life strategy (such as having children or becoming a lawyer)—which cannot be planned all at once (Bratman, 1987, pp. 29–30).

Future-directed intentions are quite similar to plans. The main difference between them is the complexity of the latter. "Intentions are … the building blocks of such plans; and plans are intentions writ large" (Bratman, 1987, p. 8). This juxtaposition does two things: future-directed intentions explain the basic structure and function of plans, and vice versa. Plans, when understood as built from intentions, gain theoretical attire. The constraints on plans help us to better understand intentions (Bratman, 1983, p. 272; 1987, p. 32). To summarize these points by appealing to a Kantian distinction, future-directed intentions are *rationes essendi* of plans and plans are *rationes cognoscendi* of future-directed intentions.

Bratman himself does not always keep to this distinction, for he sometimes seems to treat intentions and plans as equivalent, and he is followed in this by others (DeHelian & McClennen, 1993, p. 319; Holton, 2008). My intention becomes a plan when I realize that I need to form other future-directed intentions to reach my goal. General plans contain more precise sub-plans, which are our intentions. This is why there is a strong rationale behind PTI. Still, there are reasons to prefer Bratman's official stance, on which "future-directed intentions are the building blocks of plans … [and] … intentions are plan states" (2010, p. 9). Plans are thus differentiated from future-directed intentions. A plan involves a hierarchy of embedded intentions and an intention does not. My intention to go to Japan is obviously a plan, but my intention to have a glass of Californian cabernet sauvignon can be understood as a plan only if it requires a more complex net of future-directed intentions as its subplans, and this is not obviously the case.[1] So we should take the difference between future-directed intentions and plans seriously. There are important theoretical insights that will be lost unless we introduce the distinction.

3. Making it explicit: Inertia and stability

First, let us consider the term 'intention inertia'. Bratman (1987, p. 16) uses 'inertia' as a synonym of 'stability', defined as a psychological disposition not to reconsider prior intentions. The concept of inertia is present only in his early work (Bratman, 1987, pp. 16, 27, 30, 37) and is later replaced by the concept of stability. Inertia is not only a façon de parler borrowed from Harman (1976, pp. 62–63), but also initially shares its descriptive character. Both Harman and Bratman ascribe inertia to intentions, but on closer inspection, inertia is something that intentions inherit from the structures of our agential psychology. Bratman says:

AQ2

> If my future-directed intention *manages to survive until the time of action* [emphasis added], and I see that time has arrived and nothing interferes, it will control my action then. … Lacking new considerations I will normally simply retain my intention up to the time of action. Retention of my prior intention and nonreconsideration is, so to speak, the "default option." (1987, pp. 16–17)

Inertia of intention, described in this way, reveals an innocent ambiguity. Bratman tries to explain a particular feature of intention, its tendency to last over time, by its "defaultness," which is a result of our psychological make-up, a set of deeper propensities and dispositions not to reconsider our prior intentions. So stability, understood as inertia, is a primary, and in this sense natural, property of intention. It is also taken as something positive. It reveals the hidden pragmatic and rational character of our agential psychology. We tend to retain our future-directed intentions, and this generally supports the fulfillment of our desires. I think that this is all there is to the inertia of our diachronic psychology, but it is not the whole story regarding its *stability*.

Bratman abandons the term 'inertia' for a reason. One of the crucial aspects of his account of an intention's stability is its "non-reflective (non)reconsideration"—the *reasonable stability* of an intended action in all cases in which we can assess its rationality, that is, when we can reasonably retain an intention over time (Bratman, 1987, chapter 5). Intention stability is not only a thesis about the default character of prior intention, it is also a thesis about the *rationality* of retaining or reconsidering intentions in a particular case, as ascribed from an external perspective. This aspect of stability is connected

with what Bratman calls a "two-tier approach to nonreflective (non)reconsideration" (1987, pp. 64–71). It is a model of the rationality of an agent. We can describe it as follows. The higher tier concerns the rationality of our pragmatic habit not to reconsider prior intentions, and the lower tier concerns the assessment of the rationality of particular actions. According to the interpretation which I propose, the rationality of habits considered on the higher tier must be anchored in the inertial aspects of our psychology to be assessed then as stability at the level of a given action.[2] This suggests an important conceptual difference. If inertia of intention can be described roughly as its "managing to survive until the time of action," then reasonable stability cannot be so described, as it also embraces cases of intention reconsideration. Stability is different from inertia. Before I discuss the former, let me say something about the reasons for this distinction.

Inertia is a term derived from physics. It describes an inherent tendency of an object to maintain its state unless acted upon by an external or additional factor. *Stability* means different things in different contexts. It may refer to a type of property distribution, the equilibrium of a static or dynamic system with or without respect to its environment, and so on. Thus, there are various types of stability.[3] Inertia is one of them—the type for which the general tendency of a system to preserve its state is the "default," resisting modification. Inertia is not a synonym of stability. It is obvious that an object or a system may be stable without being inertial. It can reach its state because of internal or external processes, or in response to internal or external factors. Inertia always has some degree of passivity, whereas stability may be reached via active processes.

Now consider the stability of intention as explained by the two-tier model. On closer examination, it appears to be the hybrid product of a mixture of inertia and reasonable stability. This shows that the stability of future-directed intentions involves both active and passive aspects of our psychology—a default tendency toward nonreconsideration combined with a readiness to reassess intentions when needed, as in the face of new information or a change in belief, making intentions *defeasible*. In other words, the reasonable stability of future-directed intentions is a *non-inertial stability*. It is context-dependent, and the degree to which each of these two aspects influences action will depend on the specific circumstances. A reasonable agent will always have an available disposition to reconsider or to refrain from reconsidering, creating a sensible equilibrium.

It is clear why Bratman needs to leave inertia out of his norms of planning. It is hard to understand inertia in a normative sense. But the ambiguity remains. Bratman's account, both early (1987, 1992) and late (2010, 2012, 2013), seems to give some priority to the inertial aspects of the stability of intentions. Consider, as Bratman calls it, the *three-pronged support* ("TPS") for stability:

(1) Snowball effect of intentions acting on an intention changes the world: these changes may make it increasingly sensible to continue to act on that intention. In acting on an intention one gets closer to its target and further from the target of a competing plan. There can also be a social snowball effect (social costs in the abandonment of prior intentions).

(2) Costs of reconsideration by resource-limited agents: reconsidering an intention takes time, uses various mental resources, and may require re-thinking other courses of action on which one had earlier settled as part of a package that includes the intention being reconsidered. So there is frequently reason not to reconsider.

(3) Propensities favoring non-reconsideration given our limited mental resources, we frequently depend on general, non-deliberative habits and strategies about when to reconsider. These propensities *to some extent favor non-reconsideration*: they are conducive to the overall effectiveness of our resource-limited agency. (see Bratman, 2010, pp. 12–13)

Given these aspects of our agential psychology, it seems quite clear that despite the overall sensibility and balance of the two-tier model of practical rationality, Bratman has definitely construed stability as inertia-laden. On my interpretation of the place of TPS in PTI, the default aspect of the stability of intentions appears to be more central to the overall account of the stability of agential psychology than those aspects of stability that allow for the reconsideration of intentions. Otherwise, TPS would not highlight the costs of reconsideration.

4. Making it explicit: Stability of plans

So far, I have interpreted Bratman's account of reasonable stability as it applies to future-directed intentions. Given what I have said about Bratman's ambiguous characterization of the distinction between future-directed intentions and plans, it is not a surprise that when proposing his subtle account of nonreflective (non)reconsideration, reinforced by TPS, Bratman moves the problem of stability to the level of plans. He discusses the stability of intentions and plans simultaneously, switching almost undetectably between them:

> It will be useful to introduce the notion of the stability of a plan, a notion that coincides roughly with the commonsense notion of the firmness of an intention. An agent's habits and dispositions concerning the reconsideration or nonreconsideration of a prior intention or plan determine the stability of that intention or plan. Such dispositions may vary in a wide variety of ways. My disposition to refrain from reconsidering my prior plan may be rather minimal: I might be inclined to reconsider it given only a slight divergence between the way I find the world when I come to act and the way I expected it to be when I first settled on my plan. Or my disposition may involve substantial rigidity, as when I would only reconsider it in the face of some extreme divergence from my expectations—an earthquake, for example. (Bratman, 1987, p. 65)

Still, according to the strategy endorsed here, plans outrank future-directed intentions "because of their increased complexity (as compared with relatively simple intentions)" (Bratman, 1987, p. 29). It would be natural to expect a two-step account of stability, one for future-directed intentions and one for plans, suited to their level of complexity and to the contexts in which they play their typical roles. But Bratman, for the reasons presented in section 2, gives no such account and applies his account of stability to both. In the above quote, Bratman appears to offer the same account of stability that he provides for intention in his "two-tier" model. Of course, it may be the case that the stability of simpler intentions and the stability of plans, despite their difference in complexity, have one and the same pattern, as Bratman seems to claim. Both are "determined by general, underlying dispositions." But this does not resolve the problem of the increased complexity of plans, as compared to intentions, in regard to *reasonable stability*. It is, therefore, only a stipulative claim that the conditions for the stability of complex plans are the same as those for the stability of simple intentions—a stipulation which Bratman accepts neither explicitly nor implicitly.

This supports my modified Bratmanian account of plans. Plans, or plan states, should be understood as complex, more or less coherent clusters or chains of intentions. I think that the two-tier model, though promising, needs correction, but I also think that Bratman's own strategies can provide it. We do not need important departures from the PTI. My corrected version should be taken as an enrichment. There are two reasons for this, and both are needed to explain diachronic agency. The first is connected with the *complexity* of plans and the second with the *specificity* of planning—retaining, reconsidering and abandoning intentions, as plan states, over time according to the demands of the agential environment. Both of these issues arise in the context of cases which have been deliberately ignored by Bratman—intentions that are abandoned without reconsideration. Before elucidating this problem, I will introduce the core idea that promises to rectify the limitations of Bratman's account of plan stability.

5. Introducing plasticity

There are strong reasons to suspect that the norm of reasonable stability for plans is not exhausted by the notion of stability provided by TPS. That is why I propose to introduce another norm—*plasticity*. An initial, and rough, definition follows:

Planning plasticity (1). Plans should be flexible. That is, the complex chains of intentions which constitute a plan should be considered alongside the demands of the environment of the planning agent.

The main rationale for positing plasticity is that plasticity helps to retain the stable character of our agential psychology with respect to the environment—its (in)stability. This raises the question of whether plan stability is adequately described by the defeasible default model or the two-tier model.

The resulting account of plasticity as a norm of planning will require (1) the presentation of its conditions, function, and importance; (2) support for plasticity via a weakening of the TPS for stability and the proposal of a competitive "TPS" for plasticity; and (3) a defense of plasticity as reasonable compared with the model of plan stability provided by the framework of the stability of intentions. In the next section, I will sketch my initial account in terms of these desiderata.[4]

I will start with Bratman's view of the role of the agent's environment:

> If I am [the] sort of person who is constantly on the alert for *dangers in my environment* [emphasis added], my plans are likely to have a kind of instability they would not have if I were a person too engrossed in my projects to be very sensitive to such dangers. I might be more likely than most to reconsider my plan to go to Tanner Library upon hearing of a smog alert in the area. … Suppose … that I have a *high sensitivity to certain environmental hazards*, a sensitivity associated with certain kinds of *plan instability* [emphasis added]. These kinds of instability may in fact conduce to a kind of useful preventive activity. But they may also lead to a kind of apprehensive and fearful consciousness that will undermine my long-term interest. (1987, pp. 65-66)

Bratman appears to consider the environment an important factor in theorizing about the stability of planning, but a few remarks are in order.

First, although this account of the environment may be taken in a narrow sense, as the natural human habitat, it may equally be understood in a broad sense, as embracing the socio-cultural setting as well. This setting may itself be *stable* or *unstable.* This distinction has some import for an adequate account of agential planning.[5]

Second, the account of a planning agent in her environment reveals that the way she responds to environmental hazards is a matter of her psychology. It is this way of responding that makes her planning stable or unstable. This is only part of the story. We can easily imagine an agent who is sensitive to the changes in agential environment. She abandons her plan states with *no* reconsideration, but without losing her overall psychological stability. Changes in our planning may be made sensibly according to the environmental conditions within which we act. Suppose that the environment of a planning agent is highly unstable. It may mean that it is dangerous, but it may also mean that the environment is dynamic and unpredictable. This supposition is not just a theoretical possibility. It embraces the everyday situations one come across in travel, in business, or in politics. Bratman's picture of the stability of agential psychology does not say much about such cases.

Third, the Bratmanian image of the planner suggests that the only problem with plan stability is the psychological setting of the planning agent (Bratman, 1987, pp. 64–65). Again, this seems only partially apt. The way we respond to certain aspects of the environment depends, to a large extent, on how stable it is. The planning (in)stability depends also on the (in)stability of the environment. The socio-cultural ecosystem of our plans may be generally stable, but it may be generally unstable as well. The range of application of a feasible account of rational agency cannot be restricted to the cases that are not affected by this variable. **AQ8**

So far, we have considered three interlocking aspects of planning agency that merge in the ecosystem relevant to that agency: (1) natural and socio-cultural conditions, (2) sensible responses of a planner to these conditions, and (3) the stability or instability of the environment viewed as an additional factor, which, as an independent variable, enriches our understanding of stable agential planning. This account is richer than the one that can be extracted from Bratman's work. His account of stability disregards the nuances of the environment as an additional element that influences planning. For this reason, Bratman's account does not embrace all of the necessary conditions for a stable psychology, especially when the environment is unstable.

Why is understanding the unstable environment important for understanding planning stability? The more unstable the environment, the more likely it is that agents will have to change their plans. This is what we are looking for—situations in which the standard account of the reasonable stability of intentions cannot do justice to our intelligent functioning as planners. The standard account fails in these situations because it has been deliberately narrowed down to cases in which we change or retain our plans via (non)reconsideration.

> I will be making the simplifying assumption that an agent abandons a prior intention only as a result of some form of reconsideration of that prior intention. … In an earthquake I might just abandon my prior intention to play bridge tonight without engaging in anything that amounts to reconsideration of that intention. (Bratman, 1987, p. 183, note 1)

The two-tier model cannot embrace all the cases of planning in which we have to abandon our plan states without reconsideration because of environmental fluctuations, as in an earthquake. Nevertheless, we need such an account. If we accept the idea that there are many such big or small "earthquakes" in the lives of planning agents, it should become clear that Bratman's account of reasonable stability does not give an adequate picture of the psychological stability of planning in unstable environments.

This is why PTI needs another form of stability, namely *sensible plasticity*. This idea was originally developed by Tadeusz Kotarbinski (1961, 1983). I adapt his conception to the context of PTI.

If my plan is being carried out in an unstable environment, I need at least two things for my actions to be reasonable and effective. My planning has to be both *richer* and *humbler* than Bratman allows. To begin with the latter, my plan has to be intentionally *underdetermined* with respect to the means which allow me to realize it. In other words, it has to be partial (Kotarbinski, 1961, p. 195). Bratman does claim that plans are essentially partial, but the incompleteness of a plan on my account is *deliberate*. It is a sign of the agent's sensitivity to the possibility of change in the plan's environment. It is therefore reasonable. As for the richness of planning plasticity, planning as a complex chain of intentions requires imagination and knowledge regarding potential additional steps, change of means, or even reframing of the whole plan. We need be prepared to have alternative plan states, or even whole clusters of substitutions that may enter into our planning.

Planning plasticity (2). The readiness of an agent to modify a plan, according to the new information about the planning environment.   AQ9

The point is that we are, and should be, prepared to change our plans in order to keep our diachronic agency stable. While discussing the relationship between stability and inertia in section 3, I made the point that stability of intention is not the same as overall psychological stability, understood as a set of dispositions and propensities. It now becomes apparent that plasticity, as a norm governing our plans, is grounded in the latter and that this stabilizing aspect of our psychology is different from the one supported by the TPS for intentions. Our psychology reveals two different dispositions, depending on   AQ10 the context—the (inertial) tendency to stick to prior intentions and the readiness to flexibly abandon plan states when needed.

Consider two examples. First, recall my plan to travel to Japan. I usually book my flights with Lufthansa. Suppose that this company undergoes a long-term strike after I have planned my travel. No matter how much I like to fly with Lufthansa, I need to take account of the possibility of having to fly with another carrier at the time of planning. Suppose now that I intended to fly to Japan for a vacation, and that this was a part of larger plan to see Asia. I had been thinking of it for a long time. I had partially planned my trip—I had chosen my vacation time and the area to stay in, but nothing more. Before filling in my plan with more details, I receive information about a tsunami, which has flooded the area I had planned to visit. Undoubtedly, I may abandon my plan completely. However, my plan has been *structured*. The intention to visit Japan, being itself a plan, was a part of the larger intention to visit Asia. So I *reframe* my overall plan. I start to plan a visit to Singapore, say. Both of these scenarios are examples of plasticity in planning. They manifest the way in which we retain psychological stability despite environmental fluctuations.

Consider the opposite example. What happens if we are not plastic planners? Take a businessman, Mark, who is planning a merger of his company, A, with a similar company, B. Both firms have been preparing for this manoeuver for some time, since it promises huge profits. The plan looks perfect. Nevertheless, Mark is more committed to it than the management of B, and he has rigidly considered only the merger option. He has taken preliminary steps for administrative changes, prepared his employees, and invested in advertising. But, due to an unpredicted opportunity, company B suddenly retreats, pursuing a new, more promising option. Mark's planning turns out to have been highly ineffective. He not only has failed to achieve his goal, but he has also caused huge losses to company A.

CPHP 1213799
20 July 2016

Initial

CE: XX   QA: XX
Coll:XX   QC:XX

8   P. MAKOWSKI

There are plenty of such examples, showing that plasticity is a full-fledged standard of planning. Without it, we could not acheive anything in situations that require us to abandon our plan states. The instances of plasticity can be multiplied; our everyday life is full of them. Abandoning plan states, introducing changes in preparatory steps or means, and even reframing the goal of a plan are typical experiences of planning agents.

Now it should be plain that the need for overall stability in our agential psychology has dimensions beyond the simple stability of intentions. Sensible plasticity of planning is a way of retaining the stability of intentions over time without the retention or rational reconsideration of intentions. This is the reason to accept the claim that plans are complex nets of intentions, and that the norms of intention need be supplemented at the level of plans. However, this initial presentation of plasticity as an additional norm of planning now calls for support and defense.

6. Support for plasticity (the weakening of stability)

Recall the Bratmanian TPS for stability. We have (1) the snowball effect of intentions, (2) the costs of reconsideration for resource-limited agents, and (3) the propensities favoring non-reconsideration. There appears to be a tension, or even a contradiction, between the TPS for stability and the account of plasticity which I have sketched.[6] I think there is a tension, but not a contradiction. So instead of trying to refute the TPS for stability, I will construct a similar support for plasticity. Overall psychological stability requires a slightly richer account which makes space for plasticity in planning. One might call this strategy the weakening of (the Bratmanian picture of) stability.

Let us begin with the problem of the snowball effect. It is true that acting on an intention changes the world in a way that makes it reasonable to continue to act on prior intentions. This provides both individual and social benefits. Nonetheless, the snowball effect is ambivalent. Continuing to act on our prior intentions when our knowledge of the (in)stability of the plan's environment is insufficient as we fill in the details of that plan often has catastrophic effects, as in the case of the businessman Mark. To defend the snowball effect unconditionally is to commit to the undesirable effects of inertial planning. Plasticity better deals with this problem than does the two-tier account of stability. Plasticity is, and should be, the way we respond to the fact that the world may change before we change it.

Furthermore, the mental costs of changing our plans can be high. Hence, it is generally better to continue acting on a prior plan. "We are resource-limited agents" is a constant motif in Bratman's writings, almost an axiom of PTI, as it is rooted in the bounded rationality model. I am not going to question it in full. Instead, I want to point to a different dimension of this problem. Planning plasticity reveals those aspects of our agential psychology that allow us to face the fact that changes without reconsideration are sometimes part of our reasonable plans. The costs of replanning are included in them. An intelligent action always has to take account of relevant transformations in its environment. This is both rational and pragmatic. We tend to accept the cost of changes in our plan, especially when we know that these changes allow us to remain effective. The pragmatic rationale of planning is more connected with *effectiveness*, reaching our goals despite the additional costs of changes in planning, than with the "default" character of intentions. As Kolodny says, "intending X may make one more likely to take costly means to X" (2008, p. 369). This is how the plasticity of planning undermines the account of stability implied by the TPS for intentions. It appears that the latter works only for relatively simple intentions in relatively stable environments, while the former applies to complex plans, especially those in unstable ecosystems.

AQ11

AQ12

I have said that plasticity requires the exercise of imagination and of our capacity for prediction when dealing with demanding situations like Mark's. This poses a challenge for an account which treats resource-limitedness as given. Our mental resources are of course limited, but it does not mean that we should not be able to use our inventiveness and cognitive flexibility when dealing with demanding situations. It's true that some of us do it better and others do it worse. Nonetheless, our mental economy is generally adequate to enable us to respond to such difficulties. Otherwise, we would have to treat our plans as *traps* whenever the planning environment fluctuates in an unpredictable way. Therefore,

instead of celebrating our limits as a default, we should support and honor the propensities and skills which are responsible for our ability to be flexible—our ability to find new strategies and solutions.

The third problem is connected to the one above. We need to take a closer look at what is responsible for the psychological underpinnings of the "higher tier" of the rationality model. This is the set of psychological, nonreflective propensities and dispositions, which, according to Bratman, support stability and non-reconsideration. The task is not to challenge the existence of the propensities which stabilize our agential psychology, but to show how they operate in the context of complex plans in an unstable environment and how they support the plasticity of planning. I have suggested that we should understand sensible plasticity as a kind of meta-stability of intentions. This is the key to understanding the question at issue. Plasticity should apply to the contexts of agency which are complex and unstable. Such contexts impose on us the need for relevant changes in planning. So, if we are rational, our predisposition to simply retain our intentions should be overridden, albeit within limits. These predispositions lead us to retain, at least generally, the same course of action which was set in our initial plan, but with possible substitutes or surrogates in the preliminary steps, the means, or even the goals. Such substitutes and surrogates come under strain from the changes in our planning ecosystem and are supported by our skills for finding alternatives. Nevertheless, our overall inclination to retain our intentions is still in force, even if transformed, when the environment of our diachronic agency has changed.[7] This inclination still aids the overall effectiveness of planning after the plastic adjustments. Plasticity is a kind of "catalyst" for stability.

These observations also illustrate my initial point that intention inertia is a result of our general psychological setting. The same is true of plasticity. Planning plasticity is moored in our psychology, and it enters into diachronic agency when we need to abandon our plan states. Now we can see that the picture of our agential psychology should be richer. Our tendency toward stability is modified by our tendency to be flexible, within limits.

I have sketched a kind of TPS for plasticity. This modification is minor, since it does not contradict the TPS for stability which Bratman has given. Instead, it broadens our understanding of agential psychology and sensitizes us to its nuances.

7. Defending plasticity

I have painted the picture of plasticity with a broad brush, to borrow a phrase from Quine, without discussing its minute details. Although the present account of plasticity is not highly specified, it requires an initial defense against a critique which is implicit in Bratman's work. In what follows, I will argue against the challenge that plasticity is irrational when analyzed in terms of self-governance, as construed by Bratman. My strategy will be to defend plasticity and offer a moderate critique of Bratman's account of stability, along with some related ideas.

Bratman's later approach to planning concerns rationality in the context of diachronic autonomy. The highlighted role of stability, in combination with the assumption that it is rational to stick to our intentions over time, is taken to support autonomy. Consider the following definition:

> Agent S is (diachronically and locally) irrational throughout time t_1—t_2 iff S is intending *at* t_1 to F at t_2; throughout t_1—t_2 S is supporting this very intention; and yet at t_2 S is abandoning this intention to F at t_2. (see Bratman, 2012, p. 79)

AQ13

Given this principle, one might suppose that plasticity entails irrationality. According to Bratman, the above rule in combination with the TPS for stability supports diachronic self-governance (2012, p. 74). By 'diachronic self-governance', Bratman means that certain plan states pertain to what matters in our lives and link our thought and action in a relevant way. There are three conditions for diachronic self-governance: (1) self-governance "along the way," (2) semantic interconnectedness between intentions, and (3) the stability of intentions "in the absence of supposed conclusive reason for change" (Bratman, 2012, p. 83). We govern our lives as long as we retain the stability of our plans in which we realize "what matters." The idea seems feasible. However, in cases of diachronic agency

AQ14

in unstable environments, the last condition should often be modified. Plasticity concerns the cases in which we have overriding reasons for change, so there is a tension between plasticity and Bratman's account of self-governance. It is true that autonomy requires stability of intentions unless we have conclusive reason for change, but it should also embrace all cases of planning in which we remain plastic and abandon our plan states. Plastic agency in unstable environments cannot be excluded from the account of autonomy. This would be implausible. Self-governance cannot, ultimately, depend on the stability of planning as conceived by the standard framework of the two-tier model. This would restrict autonomy only to those environmental contexts which are stable. Cross-temporal autonomy cannot be something that depends on the environment. Instead, it should reflect the way in which we sensibly respond to the variations of our agential ecosystem.[8]

Now I will offer specific reasons to believe that plasticity corrects Bratman's account of autonomy. Diachronic rationality is the key here. Plastic modifications in planning, which allow for changes, are possible only if semantic interrelations between plan states are not rigidly determined (see condition 2). So we have another approximate account of plasticity:

> **Semantic underdetermination of plans.** The connections between intentions are semantically underdetermined with respect to the partial character of plans and the responses of a planner to possible fluctuations in the plan environment.[9]

The example of Mark the businessman shows the pragmatic effects of neglecting this underdetermination. Bratman suggests a role for this type of underdetermination in situations "*when* it is rational *to stick with* a prior intention [emphasis added]" (1999b, p. 86). Intentions are not simply "stuck" to other intentions, determining their meaning in diachronic plan structures. It is *we* who may, with respect to our desires and beliefs, stick to them when appropriate, depending on the circumstances. Paradoxically, this suggests that there is no such a thing as "diachronic glue" which prevents semantic interconnections between future-directed intentions from being underdetermined (Bratman, 2012, p. 83). Plasticity would not be possible if these connections were tight. It should always be possible to have diverse strategies or planning alternatives that would work in relevant practical scenarios. Without suitably "lax alternativity" of steps or pro-attitudes (Kotarbinski, 1961, p. 195), we would not be able to navigate all of the problematic situations which appear in our planning. Plasticity, therefore, does not hamper our diachronic autonomy—it supports it.

8. Conclusion

I have tried, in this paper, to map the most important issues connected to sensible plasticity, understood as a specific norm of planning and as a dimension of our psychology. Plasticity has been construed as a standard which supplements the norm of the stability of intention for complex plans. To show its importance, I have made heuristic distinctions between future-directed intentions and plans and between stability and inertia. Both distinctions touch the core of PTI. They help deepen our understanding of plans as chains of intentions and of the non-inertial stability of plans. Both issues are especially relevant to the conditions that explain the rationality of changing our plans, and particularly the rational abandonment of our intentions without reconsideration. This is the aspect of planning which Bratman has paid no attention to, instead focusing on nonreflective (non)reconsideration.

I have tried to show that the changes which are forced by new information that frustrates our plans are also important to the specific structure of our psychology in action and to our rationality. To this end, I have introduced the factor of a plan's environment. Plasticity is a norm of planning which applies especially in unstable environments, but its psychological roots are no less deep than the roots of intention stability presented by Bratman in his insightful account of PTI. Our planning psychology not only is oriented toward a simple, default stability, but is also plastic enough to embrace the potential nullification of our projects. Furthermore, this plasticity is moored in the same stabilizing dimension of our psychology.

My aim has been to enrich and supplement the picture of stability drawn by Bratman. Stability has been overstressed, but not overrated, in his PTI. He has shown a very important aspect of stability, but this is only one aspect of a broader and more complicated issue.

Notes

1. Future-directed intentions probably *always* enter, sooner or later, some planning structure in which they make sense. See Bratman (1987, p. 32).
2. Holton, whose account is close to mine, notices that there is an "obvious analogy" between the two-tier approach and rule utilitarianism (Holton, 2004, p. 510).
3. See types of stability in engineering, mathematics, and the natural sciences.
4. A more specific account is the subject of my further work.
5. In AI, the factor of the environment is normally included in modelling. AI theorists have drawn inspiration from PTI, but it is high time to reverse the direction of inspiration. Elijah Millgram (2014) has made some independent contributions to the subject.
6. This criticism was proposed by David Copp in discussion.
7. Here, Bratman's (2007) discussion of the default dimension of stability explicitly supports plasticity.
8. My thinking here is quite close to Ferrero's (2010). Roughly, every decision resulting from changes in intentions, either via reconsideration or via simple abandonment, should support our autonomy.
9. Plans are underdetermined in a number of ways. Here I point out only semantic underdetermination, which I understand broadly in Recanati's (2002) sense as dependence on pragmatic factors.

Disclosure statement

No potential conflict of interest was reported by the author.   AQ15

References

Bratman, M. E. (1983). Taking plans seriously. *Social Theory and Practice, 9*, 272–287.
Bratman, M. E. (1987). *Intention, plans, and practical reason*. Cambridge, MA: Harvard University Press.
Bratman, M. E. (1992). Planning and the stability of intention. *Mind and Machines, 2*, 1–16.
Bratman, M. E. (1999a). Planning and temptation. In *Faces of intention: Selected essays on intention and agency* (pp. 36–57). Cambridge: Cambridge University Press.
Bratman, M. E. (1999b). Toxin, temptation, and the stability of intention. In *Faces of intention: Selected essays on intention and agency* (pp. 58–90). Cambridge: Cambridge University Press.
Bratman, M. E. (2007). Temptation revisited. In *Structures of agency* (pp. 257–282). New York, NY: Oxford University Press.
Bratman, M. E. (2009). Modest sociality and the distinctiveness of intention. *Philosophical Studies, 144*, 149–165.
Bratman, M. E. (2010). Agency, time, and sociality. *Proceedings and Addresses of the American Philosophical Association, 84*, 7–26.
Bratman, M. E. (2012). Time, rationality, and self-governance. *Philosophical Issues, 22*, 73–88.
Bratman, M. E. (2013). The fecundity of planning agency. In D. Shoemaker (Ed.), *Oxford studies in agency and responsibility* (pp. 47–69). Oxford: Oxford University Press.
DeHelian, L., & McClennen, E. (1993). Planning and the stability of intention: A comment. *Mind and Machines, 3*, 319–333.
Ferrero, L. (2010). Decisions, diachronic autonomy, and the division of deliberative labor. *Philosophers' Imprint, 10*, 1–23.
Harman, G. (1976). Practical reasoning. *Review of Metaphysics, 29*, 431–463.
Holton, R. (2004). Rational resolve. *Philosophical Review, 113*, 507–535.
Holton, R. (2008). Partial belief, partial intention. *Mind, 117*, 27–58.
Kolodny, N. (2008). The myth of practical consistency. *European Journal of Philosophy, 16*, 366–402.
Kotarbinski, T. (1961). The property of a good plan. *Methodos, 13*, 189–201.
Kotarbinski, T. (1983). The general concept of plan. *Dialectics and Humanism, 10*, 15–31.
Millgram, E. (2014). Segmented agency. In M. Vargas & G. Yaffe (Eds.), *Rational and social agency. The philosophy of Michael Bratman* (pp. 152–189). New York: Oxford University Press.
Mintoff, J. (2004). Rule worship and the stability of intention. *Philosophia, 31*, 401–426.
Pollack, M. E. (1992). The uses of plans. *Artificial Intelligence, 57*, 43–68.
Recanati, F. (2002). Does linguistic communication rest on inference? *Mind & Language, 17*, 105–127.
Scanlon, T. M. (2008). *Moral dimensions: Permissibility, meaning, blame*. Belknap: Cambridge, MA.

AQ16