Decisions, Diachronic Autonomy & the Division of Deliberative Labor

Luca Ferrero

University of Wisconsin at Milwaukee

1. Introduction

§1.1

A distinctive feature of our agency is the ability to bind our future conduct by making future-directed decisions. The binding of decisions is not one of mere physical constraint. A decision is not the trigger of some mechanism that takes control of the agent at the future time f and physically forces her to ϕ. When the agent ϕs out of her past decision to do so, she is in rational control of her conduct at the time of action.[1] Decisions appear to have *rational authority* over the agent's future conduct. When the time of action comes, the agent is normally guided by no other rational consideration but her past decision. She is guided, not goaded, by it.[2] Unlike manipulative forms of distal self-control such as precommitments, decisions do not seem to alter the future situation of choice by introducing features extraneous to the original merits of the case.[3] Decisions appear nonetheless to make some kind of difference at the time of action. Were it not so, they would not be effective at influencing future conduct. A successful theory of future-directed decisions must account for the distinctive rational guidance of decisions and show how they can be effective without being manipulative.[4] A theory of this kind does not deny that decisions might play a causal role in the agent's psychology and that their effectiveness is, in part, a causal matter. But such a theory rejects the suggestion that genuine future-directed decisions operate as mere time-delay devices such as lit-fuses, that is, by way of mechanisms of brute, non-rational causality that bring about the inception of the action at f by bypassing the agent's contemporaneous exercise of rational governance.[5]

1. See Velleman (1997: 45–46), Anderson (1996: 542), and Hinchman (2003: 40).

2. For the contrast between goading and guiding, see Falk (1953).

3. On non-manipulation, see Pink (1996: 6, 114, 269).

4. For a statement of the combination of the requirements of effectiveness and non manipulation, see Velleman (1997: 47–48).

5. For a further discussion of decisions and psychological mechanisms, see §3.11 below.

§1.2

An obvious suggestion to account for the rational authority of future-directed decisions is to take them as sources of a particular kind of reasons — "decision-based reasons" — that are usually strong enough to move the agent at the time of action to act as originally decided. This is a promising start, but the two standard accounts of these decision-based reasons prove unsatisfactory. The standard views focus either on strategic uses of decisions or their roles as tie-breakers. As I will show in §2, the views that focus on the strategic uses violate the desideratum of non-manipulation. This is not a problem with the views that focus on the tie-breaking role of decisions (henceforth, unless otherwise noted, I will use 'decisions' to refer to future-directed decisions). These views, however, at most offer a partial account of the contribution of decisions to diachronic agency. These difficulties might suggest doing away with the idea of decision-based reasons altogether and explaining the rational influence of decisions in terms of normative requirements.[6] I will argue, however, that there are problems with an analysis in terms of normative commitments and that there might still be a role for an account in terms of decision-based reasons. Starting with §3, I will develop a view of the authority of decisions based on their role in the transtemporal intrapersonal division of deliberative labor. Roughly stated, the view consists of two major claims. First, a decision gives rise at the time of action to a particular kind of exclusionary reason, a second-order reason to ignore the balance of first-order reasons and simply act as previously decided, which saves us the costs of a novel deliberation. Second, the grounds for the validation of this exclusionary reason are 'evidential'. The intuitive idea, to be subjected to refinement in due course, is that the agent at the time of action is justified in acting directly out of the exclusionary decision-based reason only insofar as she is warranted in believing that, were she to consider the matter anew at that time (that is, independently of her past decision), she would come to the same conclusion (§§3.4–3.11).

6. See Broome (2001).

Despite some intuitive appeal and the pioneering work of Joseph Raz, an account of the authority of future-directed decisions in terms of exclusionary reasons has been surprisingly neglected in the literature.[7] In this work, I will show how an account of this sort can be developed to meet the desiderata of non-manipulation and effectiveness while respecting the agent's autonomy over time (see §4). In §5 and §7, I will compare decision-based reasons with other exclusionary reasons with which they can be easily confused, a comparison that helps sharpen the claims of the view defended in this paper. In §6, I will discuss the contribution of choices between equally choiceworthy options to the distal authority of decisions. I will close by suggesting a direction for further work on the relation between decisions and the temporal stability of reasons for action.

2. Strategic Manipulation, Selection, and Normative Requirements

§2.1

Standard accounts of the authority of future-directed decisions in terms of decision-based reasons fall into two categories: the "strategic" and the "selection" accounts, as I will call them. According to the strategic accounts, the agent at f has a decision-based reason to ϕ because ϕ-ing satisfies a decision-related preference, such as the desire to preserve her reputation of steadfastness[8] or the desire to avoid the costs associated with the undoing of the investments and the preparatory arrangements that the agent made under the expectation of her future ϕ-ing.[9]

7. The basic idea of decisions as sources of exclusionary reasons is presented by Raz as part of a general theory of authority (Raz 1975, 1979). Although Raz's work has been extensively debated in many quarters, it has received at best only a passing mention in the literature on future-directed decisions and intentions. Among recent works, Rovane (1998: 144ff.) and Hinchman (2003) offer accounts of intentions that bear some similarities to Raz's view, but they do not seem aware of his pioneering contribution.

8. See Sobel (1994: 249–250), Ainslie (2001).

9. See Pink (1996: 130). Bratman (1987: 82) calls this phenomenon the "snowball effect". Another kind of strategic account is McClennen (1990) and (1998). For discussion of strategic accounts, see Ferrero (2006) and Ferrero (2009).

These are "strategic" accounts because they see decisions primarily as tools for the strategic distal management of one's reluctant or irrational future selves. As such, they describe common and effective techniques of self-control. The problem is that these techniques have more in common with the manipulative devices of precommitment than with the operation of genuine future-directed decisions. In particular, the strategic accounts fail to account for the fact that the agent who ϕs at f in virtue of a prior decision to ϕ is not normally *induced* to ϕ indirectly, *i.e.*, by way of some effect of the prior decision, an effect that gives rise to a sufficiently powerful incentive to overcome the agent's initial resistance to ϕ. Normally, one ϕs *directly out of* one's decision to ϕ rather than as an *indirect result* of it.

§2.2

Strategic accounts violate the desideratum of non-manipulation. In order to counteract the agent's failure at f to see that she is to ϕ, decisions have to introduce into the situation of choice considerations *extraneous* to the original merits of the case. This is a source of a further problem. A future-directed decision would have, to use Bratman's terminology, an unacceptable "bootstrapping effect".[10] If a decision to ϕ were to operate as a kind of strategic inducement in favor of ϕ-ing, the agent at f would still have a reason to ϕ (and often a decisive one) even if, at that time, she would not decide to f were she to consider the matter independently of the effects of the strategic inducement. This is because, once an effective manipulative mechanism is set in motion, the agent at f is either unable to counteract its effect or, if she is, she can only do so via some costly tampering.

   The bootstrapping effect is required of any manipulative mechanism that is meant to correct for an expected future irrational change of mind. Only by offering an additional — and usually decisive — reason to φ can the agent's irrational resistance to φ be overcome at f. However, decisions do not have a bootstrapping effect. By default, they

are immediately revocable upon the agent's change of mind about the advisability of her ϕ-ing. As such, decisions are instruments for the correction of failures of temporal rationality.

§2.3

These problems do not arise for "selection" accounts. According to such accounts, decisions have rational authority in virtue of their role in the prospective *selection* or *adoption* of a course of action among equally choiceworthy options.[11] A decision as selection generates a novel reason to act as originally selected, but this reason is neither manipulative nor bootstrapping. First of all, the selection is meant to break a tie, not to counteract the reluctance of irrational future selves. Secondly, the selection is in principle always revocable. The troubles with selection accounts are of a different sort. First, it is unclear whether a *future-directed* selection can be truly effective in the mode of rational guidance. The fact that a present-directed selection is usually (and immediately) effective does not guarantee that selections can have a similar power over an extended temporal interval. The very feature that makes a selection non-manipulative — the fact that it introduces no truly *substantial* difference in the situation of choice, but only an arbitrary tie-breaker — seems to jeopardize the temporal stability of the selection-based reason. Given that the original selection was arbitrary, why should one be bound by it at a later time? At that time, it is neither impossible nor irrational to make a novel selection regardless of the prior one. Revoking a past arbitrary selection is not like repudiating one's original assessment of the merits of the case. This is not to deny that, for purposes of transtemporal coordination, we might have to select in advance among equally choiceworthy future options. But, as I will show in §6, either the selection is not truly directed to the future,

---

10. See Bratman (1987: 25).

11. Versions of the selection account can be found in Pink (1996) and Goetz (1998: 212). On selection-based reasons, see Scanlon (1998: 46, 70) and Mintoff (2001). Hints of a selection account can be found in Velleman's suggestion that the role of a future-directed decision is to tilt the balance among first-order *motives* (rather than reasons); see Velleman (2000: 22) and Velleman (2007: 18).

or the effectiveness of the selection is due to psychological propensities whose operation, although possibly rationally sanctioned, do not constitute a manifestation of rational governance.

A further trouble with selection accounts is that they do not offer a sufficiently general theory of decisions. Not all choices involve selections between equally choiceworthy options. Moreover, the contribution of decisions as selection does not seem to offer the starting point for a complete theory of decisions. It seems more plausible to begin by accounting for the authority of decisions that involve no selection. Only when this account is available should we consider what selection might add to the basic rational authority of non-selective decisions.

§2.4

The difficulties with the selection and strategic accounts might suggest giving up the very idea that decisions guide via decision-based reasons. John Broome, for instance, has argued that the relation between decision and action is better understood in terms of a "normative requirement", that is, a normative relation that is strict, wide-scope, and non-detaching.[12] To illustrate, let's imagine that q can be inferred from p by an immediate valid inference. There is then a normative requirement to the effect that one ought to (believe $p \supset$ believe q). One is not as one ought to be if one happens to believe p but not to believe q. However, simply believing that p is not a reason to believe that q. From the belief that p one cannot detach the claim that one ought to believe that q. After all, p might be false — in which case, one would do better by rejecting the belief that p. Normative requirements impose a symmetrical constraint. In the present example, the agent can satisfy the requirement either by believing that q or by rejecting the belief that p. The normative requirement is a requirement of consistency in one's attitudes. The idea of normative requirement seems to apply straightforwardly to *present*-directed decisions. If an agent has decided to ϕ *now*, rationality demands that she is to ϕ *right now*, but the decision

12. Broome (1999: §3).

generates no extra reason in support of ϕ-ing. This reason would be one too many, with a troubling bootstrapping effect in case one ought not to ϕ.[13] The grounds of this normative requirement appear to lie in the very nature of a present-directed decision and the primitive inconsistency involved in the failure to act on such a decision.

Let's grant to Broome that the normative relation between a present-directed decision and action is best accounted for in terms of a normative requirement. The question that interests us here is whether this solution can be extended to *future*-directed decisions. Unfortunately, there seems to be no straightforward extension. There seems to be no inconsistency — or at least no basic one — in a failure to carry out a decision that has been made well in advance of the time of action. The passage of time seems to loosen up the normative connection between decision and action. For one, the passage of time allows for changes of mind, which are not necessarily irrational. Broome explicitly acknowledges this. This is why he claims that the normative requirement that applies to future-directed decisions is a weaker one: A future-directed decision is normatively required to be carried out only if *it is not repudiated*, and no irrationality is involved in repudiating a decision.[14]

§2.5

If, however, repudiating a decision is always rationally permissible, how could there be a genuine *requirement* to carry out a decision? Broome thinks that this worry is unjustified because to repudiate is not just to stop having an intention. It is rather something that the agent must do "deliberately" in that she "must at least think about it for a moment". The repudiation, as he writes, "requires you to distance yourself from the intention — set yourself apart from it in some way".[15] In other words, Broome is claiming that, although *by default* one is normatively

13. Broome (2001: §1).

14. See Broome (2001) and Mintoff (1999: 271).

15. Broome (2001: §7).

required to act on a prior decision, one can always repudiate a prior decision without any irrationality. For argument's sake, let's assume that Broome is right in thinking that repudiations involve no irrationality.[16] But let's consider the nature of the inertia of the demand imposed by a decision. Why does it take an actual repudiation to cancel this demand? If one were to insist that this is simply a matter of a "normative requirement", this seems only a restatement of the *explanandum*, *i.e.*, of the fact that decisions have some sort of default authority. But it neither explains the source of this authority nor dispels the worry that the ease of repudiation might ultimately undermine the authority of decisions.

Broome does not appear guilty of offering a mere restatement of the *explanandum*. He offers a more informative account of the inertia. In his discussion of future-directed intentions, the idea of normative requirement is explicitly glossed in causal terms. What is "normatively sanctioned as rational" is the causal process that leads from the acquisition of the intention to the action.[17] The inertia is causal. Broome writes, "As a causal matter, we usually carry out our intentions; once you have an intention, you usually retain it until you carry it out. Without this tendency, you would never be able to complete any course of action that takes time."[18] This explains not only the inertia of decisions,

---

16. On the issue whether repudiation is never irrational, Broome's presentation is ambiguous. In some passages, he links the acceptability of repudiation to those specific decisions that are made for no particular reason, as it happens in selections ("You may have acquired the intention for no reason and *consequently* need no reason to give it up" §7 italics mine; "[Abraham] has no reason not to repudiate [his intention] *because* he had no reason to form this particular intention in the first place", §8 italics mine). But his official statement of the normative requirement as applied to future-directed decisions does not qualify the repudiation accordingly. If he did, the requirement would read something like: One is required to ϕ as previously decided unless the decision was based on no reason, in which case one would have to ϕ only if one had not repudiated the decision — which one is rationally permitted to do *as long as* the decision was based on no reason. The objections to Broome that I press in the main text, however, still apply to this revised formulation.

17. See Broome (2001: §4).

18. Broome (2001: 113).

but also why the repudiation takes the deliberate (albeit minimal) effort of "distancing oneself from the decision". This is because to repudiate a decision one must tamper with the causal mechanism that the decision has already set in motion.

Broome also writes that "the causal process that usually brings us to carry out our intentions is a rational one. It is normatively sanctioned". As I read this passage, Broome is claiming that we normatively sanction the causal process in the sense that we approve of it (say, because it makes our conduct more stable). But the fact that a causal process is rational in the sense of "being normatively sanctioned" does not necessarily mean that that process is tantamount to an exercise of the agent's rational governance. To this extent, according to Broome's account, decisions might just operate like lit-fuses or delaying devices. This reading is confirmed by Broome's claim that the causal processes initiated by a decision, if they are not deliberately interrupted by repudiation, "bring us to carry out our intentions". Hence, Broome's view ultimately denies that future directed decisions exercise rational authority by *guiding* our conduct *at the time of action* via the agent's *contemporaneous* and *direct* exercise of her rational governance. Once the decision is made, the exercise of rational governance is only called for if the agent wants to deliberately interrupt the causal tendency, that is, if she is going to repudiate the decision. The trouble with this account is that it gives up the intuitive desideratum that acting out of a prior decision is an exercise of agential governance at the time of action (§1.1). This is a radical suggestion, one that should not be pursued unless the desideratum is first proven impossible to satisfy. But there is no need to take this route. As I am going to argue in this paper, there is a particular kind of decision-based reason that explains how decisions exert authority *in* the exercise of agential governance at the time of action.

§2.6

There is nonetheless room for a different interpretation of the operation of normative requirements. According to such an interpretation,

standard compliance with a normative requirement amounts to a direct exercise of rational governance. When the agent responds to a perceived inconsistency in her attitudes, she is not aiming primarily at fixing some distinct causal process that has gone awry. She is rather trying to remove the inconsistency *as such* by operating directly on her attitudes, *i.e.*, by either rejecting or suspending some of them. In this sense, removing the inconsistency is an immediate manifestation of the agent's contemporaneous rational governance. This is also a causal operation, of course, in that the agent's psychology is realized or constituted by causal processes. What is normatively sanctioned as rational, however, is the consistency of the attitudes *in the exercise* of agential governance. The sanction only indirectly targets the causal processes that realize or constitute this agential governance.[19] If we apply this reading to future-directed decisions, we do not have to give up the idea that in acting out of her prior decision the agent exercises her contemporaneous rational governance. At the time of action, the agent is faced with the threat of inconsistency if she does not act as originally decided. This is meant to explain why she usually carries out her decision (unless she decides to repudiate it, which is another way to secure her consistency). Responding to the threat of inconsistency is a manifestation of rational governance at the time of action. To sum up, according to this reading, past decisions can indeed exercise genuine rational authority at the time of action. The problem with this account, however, is that it stops at this point. Its proponents still need to say more about the inconsistency that is allegedly involved in the failure to act on a prior decision. As I said above, to simply claim that, by default, one is normatively required to act on prior decisions is to restate that prior decisions have rational authority over future conduct. We still need an account of the nature and source of this authority.

19. It seems to me that this is the most common interpretation of the idea of normative requirements in the literature. Broome himself might favor this reading, or at least allow for it, in many discussions of normative requirements. However, at least in his explicit application of the idea of normative requirement to future-directed intentions, he explicitly embraces the causal interpretation, as indicated in the previous section.

At this point one might protest, on behalf of normative requirements, that even if we still lack an explanation of the source of the inconsistency, the idea of normative requirements offers a superior account of the structure of the authority of decisions. In virtue of the non-detachable character of normative requirements, the account avoids the objectionable bootstrapping that is allegedly produced by decision-based reasons. But this is not a conclusive consideration against decision-based reasons. It rather shows that any adequate account of the authority of decisions must make sure that this authority is not bootstrapping. As I am going to argue, this desideratum can be met by a particular kind of decision-based reasons. In any event, other features of normative requirements are not as attractive. In particular, the symmetry of the requirement seems to create troubles with the *default* nature of the authority of decisions once we give up the unappealing causal-inertial proposal discussed above. To sum up, it seems that the notion of normative requirements has not yet been sufficiently developed to offer a convincing alternative to decision-based reasons in explaining the authority of *future*-directed decisions.[20] What

20. Notice that I am not denying that there might be a primitive inconsistency in a failure to act immediately on a *presently*-formed decision directed at a *present* action. I have already granted that there might be a normative requirement that applies to presently-formed present-directed decisions. If we grant this, the question about the rational authority of future-directed decisions can be stated thus: How is it that, once the time of action comes, a future-directed decision demands by default that the agent put herself under the normative requirement to ϕ as if she had just decided to ϕ presently? When the time of action comes, the agent does not find herself automatically "saddled" with the normative requirement to ϕ given that she did not elect to interrupt the inertial causal mechanism set in motion by the earlier decision (see Buss 1999: 405). True, we normally keep up with the passage of time by the automatic updating of the contents of our attitudes. When the time of action f comes, the decision is no longer to ϕ at a future time f, but to ϕ at $f=now$. But this updating is not the brute transformation of a past decision to ϕ at f into a *presently-formed* decision to ϕ now. The updating does not put the agent under the normative requirement that applies to the latter decision. When the agent updates her past decision and finds out that f is now, she is not thereby in the same position as if she had just decided to ϕ now. She is rather under a *rational demand* to put herself in that position. Responding to this demand is part and parcel of the agent's exercise of rational governance at the time of action. The question I raise in this paper concerns the nature and source of

still needs to be shown is whether there are any kinds of decision-based reasons that can meet all the desiderata laid out thus far. This is the task for the rest of this paper.

3. The DDL view

§3.1

The rational authority of future-directed decisions derives, I shall argue, from their contribution to the *transtemporal division of deliberative labor* (DDL, hereafter). It is obvious that stable and effective decisions contribute to DDL in that they spare the agents both the costs of repeated deliberation about the same subject matter and the expenses of contingency planning. This contribution offers an important rationale for the general reliance on stable and effective future-directed decisions. But the DDL view goes beyond the uncontroversial statement of this rationale. The DDL view claims that the role of decisions as tools for the division of deliberative labor explains both the *nature* and the *source* of their *rational authority* on future conduct. It explains how they normally manage to shape the conduct of rational agents without encroaching upon their diachronic rational governance.[21]

§3.2

The DDL view holds that, under normal circumstances, the agent at the time of action (t_{act}, hereafter) has a distinctive kind of reason to act as originally decided. She is normally warranted in taking the memory of her past decision (or any other reliable record of it) as a *stand-in* for the decision that she would make at t_{act} if she were to engage in a full deliberation at that time independently of her prior decision. This is because, under normal circumstances, the agent is warranted in believing both that the record of her past decision is accurate and that her

past self at the time of decision (t_{dec}, hereafter) was not a biased, misinformed, or incompetent judge about her future situation of choice at t_{act}. The combination of these warrants justifies the agent's taking the memory of her past decision "at face value" and behaving as if she had delegated the labor of deliberation to her earlier self. The agent at t_{act} is thereby spared the costs of a novel deliberation; and she is going to act out of her earlier decision without further ado, *i.e.*, without consulting and balancing any consideration other than the memory of the earlier decision. In so doing, she exercises her contemporaneous rational governance at t_{act}. This exercise normally takes the form of the default acceptance of the conclusion of her past deliberation. Barring paralysis or akrasia, this acceptance usually leads to her ϕ-ing at t_{act} as if she had just decided to ϕ for the first time.[22]

There are two basic components to the DDL view. First, an account, in terms of the notion of exclusionary reasons, of the role that decision-based reasons play in the agent's exercise of rational governance at the time of action. Second, an account of what validates these exclusionary reasons. I will argue that the validation is of an epistemic and evidential kind since, at least as a first approximation, it is based on the expected convergence of the conclusion of the agent's deliberation at t_{dec} with the deliberation she would engage in at t_{act} if she were to do without her prior decision.[23]

---

this demand, a question to which the appeal to normative requirements does not seem to offer a satisfactory answer.

21. For a discussion of different views of the rationale of the capacity to act out of future-directed decisions, see Ferrero (2009).

22. If one accepts Broome's claim about the normative requirement that applies to presently-formed present-directed decisions, the acceptance amounts to the agent's putting herself under such requirement at t_{act}. At that point, a failure to ϕ would be in violation of that requirement.

23. The combination of the exclusionary character of these reasons and the epistemic/evidential nature of their validation is a distinctive contribution of the DDL view. The two elements have been separately endorsed in the literature. A sustained defense of the exclusionary character (under the notion of "pre-emptive reasons") is offered by Hinchman (2003). Hinchman explicitly rejects the evidential validation of these reasons (see fn27, below). The evidential character of the validation is briefly defended in Rovane (1998) and mentioned in passing by Joyce (1999: 60, fn16). Joyce talks of the effect of the evidential decision-based reason as an increase in the likelihood that one would act as originally decided. This is, however, too weak an effect. It does not capture the full extent of the rational authority of decisions. On the increase of likelihood, see also Pink (1996). Hartogh (2004) offers a sustained

§3.3

Let's begin with the claim that decision-based reasons (or better decision-cum-memory-based reasons) are particular kinds of exclusionary reasons — what Joseph Raz calls "protected reasons".[24] An exclusionary reason is a second-order reason to exclude some first-order reasons from deliberation. Exclusionary reasons neither override nor conflict with first-order reasons. Rather, they determine which considerations are to be excluded from the calculation of the balance of first-order reasons, even if they might in principle come to bear on such calculations. A *protected reason* to ϕ at t_{act} is the combination of (i) a first-order reason to ϕ at t_{act}, and (ii) an exclusionary reason to disregard some of the first-order reasons that bear on the choice at t_{act}. The larger the exclusion imposed by (ii), the larger the protection enjoyed by (i). In a *maximally* protected reason, (ii) demands that *all* first-order reasons be disregarded with the exception of (i). As a result, the balancing of first-order reasons reduces to the degenerate case of taking into account one reason only, namely, (i).

The DDL view claims that a decision made at t_{dec} to ϕ at t_{act} normally gives rise to a *maximally protected* reason to ϕ at t_{act}, a reason based on the agent's assessment at t_{dec} of what she is to do at t_{act}. Let's call this

---

discussion of the *primacy* of the epistemic nature of the authority of decisions over their coordinative or strategic effects, but he does not argue directly for the epistemic validation. He simply claims that the epistemic validation is an "obvious" matter (Hartogh 2004: 7). This is too optimistic a statement. Not only has the view received limited support in the literature, but, as I show in this paper, much work is still needed to establish it.

24. See Raz (1975: 37–39; 1978, 1979: 18; 1990: 191). Raz's original discussion of exclusionary reasons is primarily addressed to accounting for the nature of authority and norms in general, not for the specific authority of future-directed decisions. Raz's presentation emphasizes the idea that exclusionary reasons justify a *departure* from the merits of the case, as it indeed happens in certain sorts of exclusionary reasons, such as those based on commands, coordination, policies, and selection (see §§5–7 below). When he turns to the discussion of decisions, Raz mentions evidential considerations among the kinds of possible validation of decision-based protected reasons, but he does not single them out as the distinctive source of the authority of decisions. He rather bundles them together with other sources of validation (including coordination-based and "end-of-deliberation" ones, see §7).

reason an *assessment-based maximally protected reason* (APR, hereafter). Normally, an agent at t_{act} who remembers her past decision that she is to ϕ at t_{act} is warranted in taking the APR to ϕ at t_{act} as valid. In addition, she normally accepts it by default. Given the protected structure of this reason, once she accepts it, there is nothing for her to do but to ϕ. The acceptance of the APR puts the agent at t_{act} in the same position as if she had just decided to ϕ at that very time (including subjection to the normative requirement that might apply to the latter decision, see §2.4). As happens for a presently-acquired and present-directed decision, the acceptance of the APR normally leads immediately to the ϕ-ing (that is, barring those interferences and impediments — such as paralysis and akrasia — that might interpose between any *present*-directed decision and action).

§3.4

That a decision is usually taken to give rise to a valid maximally protected reason explains why the agent at t_{act} can act directly on the record of the earlier decision without consulting any other consideration. What does it take for such a decision-based reason to be valid? The structure of protection by itself is open to different kinds of validation, many of which are irrelevant to the division of deliberative labor and to the distinctive authority of decisions. The dictates of a sovereign, for instance, might generate valid protected reasons to act as commanded, but the legitimacy of his authority and the validity of his commands may have nothing to do with the transtemporal division of the deliberative labor of his subjects (for a discussion of different kinds of protected reasons, see §5 below). The intuitive idea behind the DDL view is that an agent at t_{act} would not be justified in acting out of a past decision unless she deemed the decision to be true to the merits of the case as she would see the matter for herself and from her own practical standpoint at the *time of action* (*i.e.*, from the set of the basic cares, concerns, values, and preferences she has at t_{act}).

As a first approximation, the agent at t_{act} is warranted in delegating deliberative work to a past self only when she believes that her past

self could correctly deliberate about the situation at t_{act} from her point of view at t_{act}. The deliberating self at t_{dec} must not only have sufficient expertise and information about the situation of choice at t_{act} but also must adjudicate the matter from the point of view of the later self. Were it not so, the later self could not take the previous conclusion to speak for her. In order to act on the decision of the past self, the agent at t_{act} is to assume that there would be a *transtemporal convergence of verdicts*. She is to assume that, were she to engage in full deliberation at t_{act} regardless of her past decision, she would reach the same conclusion that was actually reached by her past self. The APR could thus be said to be validated on evidential grounds. The agent is justified in accepting an APR if she takes it to give her sufficient evidence for the decision she would make if she were to deliberate at t_{act} independently of her prior decision.

§3.5

If we set aside for the time being cases in which the merits of the case are at least partially under-determined (which I discuss at length in §6 below), the account of validation in terms of convergence works fine in two kinds of cases. First, the convergence is obviously to be expected when the agent at t_{act} is in the same deliberative situation as she was at t_{dec}, that is, when at both times she enjoys the same deliberative resources (including time and information), she is not irrational, and she occupies the same practical standpoint. Second, the idea of validation in terms of convergence explains why the agent at t_{act} should set the APR aside if she knows either that she is under overall better deliberative circumstances than at t_{dec} or that her practical standpoint has undergone a substantive change. If either is true, a convergence is no longer guaranteed and the agent is not to rely on the APR.

§3.6

A problem with the validation in terms of convergence arises when the agent's deliberative circumstances worsen over time. In such cases, the agent at t_{act} can no longer be guaranteed that the conclusion of her

present deliberation would match her earlier one. This does not necessarily mean, however, that the agent is to reject the APR. It is rather the opposite. When the deliberative circumstances worsen over time, there is a more pressing need for the division of deliberative labor. Should we therefore give up the intuitive idea that the validation is based on the evidence of a transtemporal convergence of verdicts? Not really. What we need is a qualification of the initial interpretation of this convergence. What is to be expected is not a convergence between her actual past conclusion and the conclusion she would reach under her current deliberative circumstances. The match is rather with the conclusion that agent at t_{act} would reach now *under suitably idealized deliberative conditions*. More precisely, the agent at t_{act} is warranted in relying on the APR if she deems (i) that her deliberative conditions have not improved over time,[25] and (ii) that, *if she were now in the same deliberative conditions as her past self* and she engaged in full deliberation independently of her past decision, then she would now reach the same conclusion as her past self.

Notice that at t_{act} the agent is not to be concerned with the decision that would be made under deliberative conditions that improve on those at t_{dec} (let alone perfect ones) but that cannot be accessed prior to the time of action. Knowing that under better circumstances a better decision could be made is of no help to the agent if those circumstances are outside of the agent's reach, given the pressing need to solve her practical problems. What is helpful, instead, is the knowledge that there is a reliable record of the decision made by a rational deliberator under actual deliberative conditions that were both good-enough (at least for limited beings like us) and at least as good, if not better, than one's current deliberative conditions. The idealization built into the idea of convergence is meant to secure that the agent at t_{act} acts out

---

25. More precisely, I should say that they are not *relevantly* better, since the agent might believe that, although her deliberative circumstances have improved in principle (say, she has more time for deliberation and access to more information), these improvements are not going to impact the conclusion of her deliberation and are thus irrelevant. For simplicity's sake, I am setting aside this complication in the following discussion.

of a prior decision on which she cannot improve, given her current position, and with which she can identify. The identification is in part a matter of the agent's projection at t_{act} into the better deliberative conditions, including the absence of whatever defects of rationality from which she might suffer at t_{act}. The agent at t_{act}[26] might be aware, for instance, that she is currently too nervous or unfocused to carry out a satisfactorily complete deliberation, even if she were given all the time, resources, and information required. In this sense, the validation of APRs depends on the convergence between the actual past decision at t_{dec} and the one that the agent at t_{act} believes would be made by her better (although not necessarily perfect) self if this self were now under the better (although not necessarily perfect) deliberative conditions she already enjoyed at t_{dec}.

§3.7

To sum up, an APR is valid if and only if (a) the agent at t_{dec} did not suffer from any disqualifying defects such as irrationality and akrasia; (b) the agent's deliberative conditions at t_{act} have not improved over the good-enough conditions at t_{dec}; (c) the deliberation at t_{dec} was conducted from the practical standpoint of the agent at t_{act}; and (d) there is a reasonable expectation that she would reach the same conclusion that she did at t_{dec} if she were to engage in deliberation under conditions as good as those she enjoyed at t_{dec} and in the absence of any defect of rationality from which she might suffer at t_{act}. This is what convergence amounts to when we move away from the original suggestion that the earlier decision is the same one that would be made by the agent at t_{act} if she were to engage in full deliberation under her *actual* circumstances at t_{act}.[27]

26. Although the crucial practical standpoint is that of the agent at the time of action, it is not supposed to be the standpoint of a momentary agent. The standpoint is presumably shared across time as the standpoint of a temporally extended agent.

27. The convergence in verdicts via the idealized scenario avoids the objection to the naïve evidential view that is correctly criticized by Hinchman (2003: 41). An APR is not validated exclusively on the basis of inductive evidence about

§3.8

What I have just presented is a reflective articulation of the grounds of valid APRs. Ordinarily, we do not act out of decisions on the basis of such a sophisticated, explicit articulation (nor out of an articulated understanding of the concept of protected reasons). Nevertheless, we normally have some implicit and coarse sense of the nature and source of the authority of decisions that I have articulated in this paper. We are normally sensitive to at least the most blatant defeaters of the authority of earlier decisions (see §4.5). As a first stab toward such articulation, the agent might claim that she is justified in acting out of her prior decision because she takes her earlier decision to be correct.[28] The more elaborate account offered above is meant to articulate along two dimensions this intuitive understanding of the authority of decisions. First, we need to articulate the grounds of the agent's warrant. The agent at t_{act} has to have some sense of the extent of changes in her deliberative circumstances, if any, and of the credentials of the earlier self as an expert deliberator, as spelled out in clauses (a) and (b) above. Second, the question arises as to what kind of correctness is at stake. As a first approximation, the correctness is a matter of the decision's

the decision that one would take if one were to engage in full deliberation at t_{act} in exactly the same deliberative situation (including one's imperfections and disqualifying defects, if any). As an alternative to the naïve evidential view, Hinchman (2003: 34) suggests that we look at what he calls the "deeper level" of *trust*. For him, decision-based reasons are based on the agent's trust on the decision of a trustworthy past self (Hinchman 2003: 41). I endorse this claim but I am concerned that appeal to the idea of trust only addresses the structure of rational authority — the protected status of reasons — but leaves unspecified the source and nature of their validation. To trust the verdicts of the earlier self is a matter of having an exclusionary reason, possibly a protected one, to act on this verdict. But what makes the earlier self trust*worthy*? The idea of trust by itself does not go sufficiently deep in answering this question. It seems to me that there are different possible grounds of trustworthiness, not all of which make the earlier self's verdicts authoritative in the mode of decisions (see §5 and §7). But if I am wrong about this and the validation characteristic of decisions is implicit in the idea of trustworthiness, it is still true that we need an explicit account of this validation and thus we need to go even "deeper" than trust.

28. By "taking" the decision to be correct, I mean that the agent either believes or "accepts" — in Bratman (1999)'s sense — that the decision is correct.

being true to the original merits of the case, modulo the limitations due to the agent's finite rationality and limited deliberative resources. This might seem a trivial point, but it is not an uncontroversial ground for the authority of decisions. Both the strategic and the selection accounts reject it. The correspondence to the original merits of the case is not what makes a strategically induced reason to ϕ compelling at t_{act}; as for a selection-based reason, this is supposed to fill in gaps in the merits of the case. Moreover, simple talk of correspondence with the merits of the case still leaves open the issue of the point of view from which the merits of the case are to be appreciated. Agents who are equally rational and have the same deliberative resources are not guaranteed to agree on the assessment of the case since they might differ in their practical standpoints (including the temporal horizons of their application, such as the rate and shape of the temporal discounting). Moreover, the agent at t_{act} acts out of a decision — rather than as a result of it — only if she sees the decision *as her own*. The decision is not paternalistically *imposed* by her past self on to her current self. It has to respect the agent's autonomy at the time of action and, as such, reflect her practical standpoint at that time. The agent does not need to have an articulate understanding of her standpoint at t_{act} in order to validate the APR (she might actually rely on the division of deliberative labor to compensate for this inarticulacy), but she can grant authority to her past decisions only insofar as she has some sense that they are in keeping with her present practical standpoint, however inarticulate this standpoint might be. This last set of considerations explains the need for clauses (c) and (d). It explains why the validation is articulated from the point of view of the agent at the time of action, and why it is articulated in terms of a convergence with her verdicts, rather than in terms of an unqualified correspondence with the merits of the case.

§3.9

In the last few sections, I have shown how the naïve interpretation of convergence is to be modified to account for the validation of APRs. I think that this refinement is still in the spirit of the original suggestion,

but it makes clear that the evidential grounds of the validation were never meant to be understood purely in terms of theoretical-inductive evidence, or in terms of the *prediction* of the decisions that would be made at t_{act}. The expected convergence is rather a matter of the agreement in the conclusions reached by the agent in taking a first-personal *deliberative* stance over the practical problems that she is facing.[29] In taking such a stance, the agent is trying to figure out what the merits of the case demand of her. To this extent, the authority of APRs has an epistemic component; its validation is a matter of practical "expertise", the capacity to provide the right answers to one's practical questions. But it is expertise from the practical standpoint of the agent at t_{act}, whence the evidential component to the validation. The agent is warranted in putting herself under the authority of an APR only insofar as she is warranted in expecting that the APR would give her evidence of the conclusion she would reach from her own, first-personal, deliberative engagement with the practical question. The idealization built into the refined version of the convergence articulates an idealization that is already built into the very idea of a deliberative stance. The agent is not interested in the mere anticipation of the conclusion of her deliberative processes, even if under idealized conditions, unless at t_{act} she can project herself into these processes as the work of her first-personal, deliberative stance at its best — compatibly with the general limitations of rationality and the particular constraints imposed by the specific features of her current predicament.

§3.10

What is the relation between deliberation, assessments, and decisions? Although I have been speaking of the deliberation of the agent at t_{dec}, I do not want to suggest that a decision is always the product of an actual deliberation, let alone of an explicit one. Nevertheless, I maintain that decisions are in principle responsive to demands for rational justification. This is the sense in which they are supposed to be

29. For a contrast between the theoretical and the deliberative stance, see Moran (2001).

based on the assessments of the merits of the case. The psychological processes by which decisions are formed do not necessarily involve either explicit or even actual assessments. But a decision is responsive to rational considerations regardless of the actual process by which it is first acquired. It is subjected to rational criticism and it is supposed to be given up if shown to be unjustified.[30]

The relation between assessment and decision is particularly evident in the case of future directed decisions. A paradigmatic *present*-directed decision amounts to the executive transformation of an assessment into action (modulo the additional contribution of selection in cases of under-determination discussed in §6). Barring the interference of akrasia or paralysis, the decision marks the closing of the deliberation and the immediate transition to action.[31] In the absence of an actual deliberation and an assessment of the merits of the case, a present-directed decision just marks the inception of the action, but it still remains rationally accountable in terms of the correctness of the assessment on which the decision should have been based if one had actually engaged in deliberation. What happens to this relation when the decision is future-directed? Future-directed decisions are not *executive* in the way of present-directed ones (if they were, they would either exercise action-at-distance or trigger mere causal, time-delaying mechanisms). The temporal distance between the decision and the action takes away the basic executive contribution of the decision. This is why future-directed decisions are effective by way of rational authority. This brings to the forefront the assessment that is supposed to be the basis for the justification of a decision. This is the only basis for the decision's distal authority.[32] This is why the protected

reason that transmits the authority of the decision cannot but be, as I have claimed, an *assessment*-based one. The validation of the APR, in turn, depends on the expected convergence, under suitably idealized conditions, of the agent's decisions as the manifestations reflective of the assessments of the merits of the case that justify the decisions (although they do not necessarily precede it).

§3.11

The DDL view does not aim at explaining all the ways in which our diachronic agency is structured and influenced by decisions. It is rather meant to account for the *paradigmatic* influence of genuine future-directed decisions in the mode of an effective, non-manipulative rational guidance. The particular combination of protection and validation suggested by the DDL view is the reflective articulation of a *regulative standard* that is implicit in our everyday future-directed decisions. The DDL view does not aim at describing all of the actual psychological operations underlying the effectiveness of future-directed decisions. It rather articulates the model to which these operations are supposed to conform. Therefore, it should not come as a surprise that many of our decisions fail to live up to this standard. We often grant rational authority by default to decisions that are not backed up by a correct assessment and would not stand critical scrutiny. On the other hand, the regulative standard also allows for the acceptability of surrogates. Our conduct is often the product of psychological mechanisms that bypass our rational governance at the time of action but are nonetheless rationally sanctioned because their outcomes offer reliable surrogates for the paradigmatic operation of genuine future-directed decisions. The operation of these mechanisms is rationally acceptable, however, only in those cases in which we do not care that our conduct be the direct manifestation of our *contemporaneous* rational governance — they would not be acceptable, for instance, when signing a contract or say-

---

30. This is true even for those decisions that incorporate an arbitrary selection. An arbitrary choice is beyond rational criticism, but one can be criticized for thinking incorrectly that the situation calls for a selection, or for mistaking the extent of the under-determination. The basis for this criticism lies in the assessment of the existence and extent of the background under-determination over which the selection is supposed to operate (see §6.3).

31. See Rundle (1997: 202).

32. One might object that this overlooks two important functions of present- and

---

future-directed decisions: that they select over equivalent options and they bring a deliberation to a close. However, neither of these functions is responsible for the distal authority of decisions as I argue in §6 and §7 below.

ing "I do" in a marriage ceremony.[33] These mechanisms might play a considerable role in giving shape to our diachronic agency, but we rationally sanction their non-rational mode of influence only insofar as they approximate the model of the rational guidance of genuine future directed decisions.

4. Decisions, Non-Manipulation, and Effectiveness

§4.1

The influence exerted by decisions via APRs is of a rational kind. When the agent ϕs at t_{act} on the basis of an APR she is in rational control over her conduct *at that time*. This is so even if her rational governance at t_{act} is limited to the *default acceptance* of her past assessment to which she is normally entitled. This acceptance is usually an unobtrusive episode in the agent's mental life. The seamless psychological transition from past decision to present action explains why we might be tempted to think of past decisions as exerting *direct* control on future action, as if they bypassed the exercise of rational governance at t_{act}. In actuality, in the paradigmatic instances of acting out a past decision, there is no relinquishing of active and rational control at t_{act}, although this control is usually omissive. By default, one accepts the APR by refraining from calling it into question.[34] This kind of acceptance accounts for the "inertia of decisions", or the fact that normally a decision continues to exert rational influence until it is carried out, repudiated, or forgotten.[35] But the psychological seamlessness of this default omissive acceptance is not to be confused with the passive inertia of non-agential mechanisms. As Bratman correctly remarks, "[F]ollowing

33. For a discussion of the importance of the relation between intentional diachronic agency and the manifestation of contemporaneous rational governance, see Ferrero (2006).

34. On this point I disagree with Bratman (1987: 60). He claims that "the non-reflective non-reconsideration" of a past decision is the *absence* of an action rather than an instance of refraining. For a discussion of this difference, see Ferrero (2006).

35. See Bratman (1987: 16–17, 27).

through with one's plan is not, after all, like following through with one's tennis swing."[36] Acting out of an APR is not like being on autopilot. At the time of action, the agent is responsive at least to one reason, the APR — which tells her to ϕ without any further ado. This is not to suggest that the agent is explicitly consulting the APR. Nor is it to suggest that she is reflectively aware either of the structure of the APR as a protected reason or of the nature of the APR's validation. The DDL view does not mean to offer an over-intellectualized picture of diachronic agency. The account of reason-responsiveness given by the DDL view must be taken with the standard caveat of any philosophical account of rational activities: Responsiveness to reasons is a personal-level phenomenon, although one that is, for the most part, tacit, swift, and beyond of the focus of attention.

§4.2

The DDL view satisfies both the desiderata of non-manipulation and that of effectiveness. An APR is not manipulative, bootstrapping, or — for limited beings like us — redundant. It makes a difference to the situation of choice at t_{act}, but it does not alter the original merits of the case. The APR is, after all, only a *stand-in* for a prior assessment. The APR is a sort of anaphoric device in the order of justification. The agent who acts out of an APR is ultimately acting out of the original and unadulterated assessment of the merits of the case. The agent at t_{act} might be unable to offer any explicit reason to ϕ except for the APR, but she can always defer to the time and place of the original decision as the locus of the original, possibly articulated acknowledgment of the case for ϕ-ing. The APR is in principle *transparent*; it is as if the agent could see through it and look at the original assessment as the primary source of the justification of her conduct. But the agent does not have to see through the APR to be moved to ϕ. Transparency is not invisibility. After all, the point of relying on the APR is to make the agent ϕ at t_{act} by responding solely and directly to the APR. Hence, the difference

36. Bratman (1999: 60).

that the APR makes at t_{act} does not bear on the content of her choice but on how she comes to it. It follows that the APR is not manipulative. It guides the agent at t_{act} by easing the burden of her deliberation rather than goading her by generating extraneous considerations.[37]

§4.3

Is the guidance offered by APRs compatible with their efficacy at influencing future conduct? It might not be, if we consider the following intuitive test of effectiveness: A decision is effective only if it makes the self at t_{act} "buy into a choice that she would not otherwise have made", to put it as Velleman does.[38] This test is easily met by the strategic and selection accounts since they claim that decision-based reasons introduce an actual difference in the situation of choice and thereby induce the agent to do something that she would not have done otherwise, if not by accident. APRs operate differently. The basic idea behind the validation of the APRs is that by relying on them the agent at t_{act} chooses to do what she would choose if she were to do *without* them. The point of APRs is to spare the repetition of full deliberation at t_{act} while securing that the agent does something she can approve of from her *contemporaneous* practical standpoint. The fundamental difference that APRs make is neither to what the agent does at t_{act}, nor to why she does it. It is rather a difference in *how* she comes to realize what she is to do at t_{act} with respect to her autonomous rational governance at that time.

---

37. Because of transparency, an agent who acts on an APR that she takes to be valid makes herself liable to two distinct but related criticisms. First, she can be criticized for acting on an invalid APR. In this case, her action might still be correct but only accidentally so. Second, she can be criticized for the incorrectness of the original decision. This criticism is indirect given that the agent does not engage in full deliberation at t_{act}. It is not a fault attributable to her exercise of rational governance at t_{act}. But by accepting a transparent APR she makes herself accountable for the original decision. It might happen that an agent is subjectively justified in accepting an APR in that she cannot be blamed for thinking that the APR is valid, but at the same time she might be subject to criticism for the objective incorrectness of the original decision. This should not be surprising. It is the sort of risk that comes with any kind of delegation and deferral, including the intrapersonal division of deliberative labor.

38. See Velleman (1997: 47–48).

In this sense, if all that is at stake is the division of deliberative labor, in principle an agent could do without APRs, that is, without future-directed decisions altogether. At least, this is so when she is under perfect deliberative conditions both at t_{dec} and at t_{act}. When so, the agent does not need to be under the authority of a future-directed decision to sustain her project to ϕ at t_{act}. She could simply rely on her expectation that whenever the issue of what to do at t_{act} will arise she will be able to engage in a full deliberation and continue to reach the same conclusion. Notice that agents under these ideal conditions need not be time-slice agents, nor need they be concerned only with the immediate effects of their present actions. They can indeed be planning agents. They can embark on temporally extended activities and coordinate in advance with their future conduct. They can do this even better than we, as finite beings, can. What they lack is only the need for DDL and thus for effective *future*-directed decisions.[39] They can manage with effective *present*-directed decisions (which make them do whatever is momentarily required of them to sustain the present progress of their temporally extended projects) and stable expectations about their future present-directed decisions.

Under more realistic deliberative conditions, however, future-directed decisions are neither dispensable nor idle. This is why the APRs are effective in the sense of making a real difference regarding *what* an agent does at t_{act}. If the deliberative conditions at t_{act} are not as good as those at t_{dec}, in acting out of an APR, the agent at t_{act} chooses to do something other than what she would have chosen otherwise. Nevertheless, this effectiveness has a different flavor than in the case of the strategic management of reluctant future selves. The contribution of APRs is welcomed by the agent in the sense that she takes them *not* to make a difference to what she would choose if she were under suitably idealized conditions. In this sense, both the actual effectiveness of decisions via APRs and its justification depend on the dispensability of APRs under suitably idealized conditions. Hence, the suggestion that the primary difference that they make is in regard to *how* the agent

---

39. See Ferrero (2006).

comes to realize what she is to do rather than to what she actually does.

§4.4

APRs are not bootstrapping. They offer no resistance to being voided if the agent at t_{act} suspects that he might not be entitled to them. Like other exclusionary reasons, APRs are defeasible. Although they cannot be overridden by the first-order considerations that they exclude, they can be rejected if one suspects that the grounds for the exclusion are invalid; that is, if one suspects that the memory of the past decision is inaccurate or that the past decision is either incorrect or unjustified. The agent therefore is not bootstrapped into ϕ-ing by an APR. She can always set an APR aside and engage in a full, on-the-spot deliberation at t_{act}, whereas the influence of a manipulative mechanism can be counteracted, if at all, only by tampering with it.

This does not guarantee that APRs are automatically rejected when ϕ-ing at t_{act} is no longer choiceworthy. The agent might fail to notice that there are grounds to suspect that an APR is invalid and thus end up getting stuck with it. This danger is the inevitable price of the mechanisms for the division of deliberative labor. It could be avoided only if the agent were to make DDL pointless by *constantly* either keeping direct track of the choiceworthiness of ϕ-ing or checking the credentials of the source of APRs and the reliability of her memory of them.

§4.5

The real danger with APRs is that of *de facto* bootstrapping. This might happen if they are systematically accepted in an uncritical way. A rational agent should always be ready and willing to reopen a settled matter if reasonable doubt is cast on her entitlement to an APR. Hence, even if a decision settles the question of what to do at t_{act}, an agent is still required to be alert for at least the most obvious and manifest defeaters of her APR. Moreover, if properly challenged to justify her decision to ϕ, the agent is supposed to offer more than an uncritically reaffirmation of the APR, that is, more than a statement like "I am to ϕ

because I've decided to do so in the past". In response to a proper challenge, the agent should attempt either to show that she is entitled to the APR or to offer an articulate justification of her ϕ-ing in terms of the original merits of the case, thereby ignoring the APR.

An utterly uncritical acceptance of past assessments is a degeneration of DDL. APRs are meant to ease deliberative labor by delegating this labor to one's past self. They do not encourage a blind and unreflective acquiescence in the past assessment, which would end up *alienating* the control of one's conduct to the past self. Becoming an instrument of a past self might be desirable at times (*e.g.*, when one needs to counteract temporary irrationality), but it is not the correct model for the standard operation of future-directed decisions. As I have said before, one does not normally ϕ *as an indirect consequence* or *as a result of* one's past decision (as it happens in the case of strategic devices), but *out of* one's decision, although through the transparent mediation of an APR.[40]

Only in the latter case is the agent's "diachronic autonomy" respected. A self-directed manipulation or the uncritical acceptance of a past decision respect the agent's diachronic "rational governance" in that the agent exercises contemporaneous rational control over her conduct both at t_{dec} and t_{act}. But they do not respect her autonomy at the later time, since she is either acting as a result of an adulterated situation of choice, or she is uncritically submitting to the dictates of the prior self, with no guarantee that they might make her do what she would autonomously choose if she were to decide for herself at t_{act}. By contrast, the standard operation of APRs aims to respect the agent's diachronic autonomy. It is meant to secure that the agent at t_{act} exercises her contemporaneous rational governance in response to

40. The difference between acting *out of* a decision and acting *as a result of* a decision is apparent in what the agent is to do if she has to counter a decision-based reason. In the latter case, the efficacy of the decision-based reasons is matter of its influence on the balance of first-order reasons. These reasons work by introducing *substantial* modifications of the situation of choice. Hence, they can be countered only by neutralizing their effects, by a counter-modification. By contrast, an APR can be neutralized simply by setting it aside. There is no need to counteract its substantive effects since it has none.

her contemporaneous and autonomous appraisal of the case for ϕ-ing, although one that is transparently mediated by the APR (see §4.2).[41] In *directly* responding to the APRs, the agent at t_{act} is supposed to be *ultimately* responding to the original merits of the case[42] from her *contemporaneous* practical standpoint.[43]

---

41. I am using the term "diachronic autonomy" in a sense narrower than Velleman (1997: 46). He uses it to refer to what I call "diachronic rational governance".

42. More precisely, diachronic autonomy requires the agent to be responsive at t_{act} to reasons that are either identical to those she acknowledged at t_{dec} or partly different but still in the spirit of the original assessment in that they are the outcome of a "rational development" of that assessment. This development includes the augmentation, reinforcement, refinement, articulation, and specification of the justification for ϕ-ing as a result of the information and deliberative skills that the agent might have acquired after t_{dec}. On this development, see (Wilson 2000: 14). The development includes the reaffirmation of intention discussed in Bratman (1987: 96).

43. Bratman (1987) presents several objections to the idea that the binding force of decisions is due to decision-based reasons. None of these objections affect the DDL view. Bratman first objects that decision-based reasons are too weak if they work as standard first-order considerations. This is not a problem for APRs given that they are protected. Second, Bratman (1987: 24) objects that decision-based reasons are too strong because they are bootstrapping. But this is not a problem for APRs (see §4.2). Finally, Bratman (1987: 68) is worried that a decision-based reason could induce undesirable stubbornness by standing against all possible reconsiderations. This is not a problem with APRs since they are defeasible (see §4.4). Another set of objections, partially inspired by Bratman, is presented by Mintoff (2002: 349–50). Mintoff argues against the idea that decisions can have an indirect, second-order epistemic relevance, in the sense that the agent sees her past decision as evidence that ϕ-ing is favored by the balance of her reasons *and* she has a second-order desire to do what she believes she most desires, which "gives normative force" to the decision. This is not, however, the sense in which the DDL view appeals to second-order reasons. Protected reasons are second-order in the sense that they exclude certain first-order considerations from the balance of reasons. To this extent, Mintoff's concern that the indirect second-order account fails to account for the role of intentions as "predominant motives" does not affect the DDL view. Mintoff also objects that the view does not explain how a decision can make an action rational in cases of selection. But as I will argue in §6, decisions as selections do *not* carry authority over time.

5. Other Kinds of Protected Reasons

§5.1

APRs are not the only kind of exclusionary reasons generated by decisions or by similar phenomena. For instance, there are exclusionary and protected reasons generated by commands, policies, coordinating choices, and selections. Because of the shared structure of exclusion, it is easy to get confused about the nature of the reasons generated by genuine future-directed decisions.

Consider a command to ϕ at t_{act}. This command gives rise to a protected reason whose validity depends on the legitimacy of the issuing authority. This legitimacy, however, does not need to depend on the expectation that the commands correspond to what the subject would choose to do if she were to choose independently of the authority's issuances, not even under suitably idealized deliberative conditions. Therefore, command-based reasons are not necessarily epistemic/evidential in the mode of APRs. As such, they do not provide a good model for decisions. The common temptation to think of decisions as self-directed commands can be explained by the similar structure of protection, but the analogy is only partial in that it ignores the crucial difference in the origin of the validation of this protection.[44]

§5.2

Many of our decisions, including possibly some of the most basic ones,[45] are decisions to adopt general policies, rather than to perform specific actions. This might create some confusion about the effect of these decisions, given that policies are another source of protected reasons. In adopting the policy P to ϕ when circumstances C obtain, the agent is given a reason to refrain from balancing first-order reasons when C obtains. A policy, like a decision, eases the agent's deliberative labor.

---

44. For the suggestion that decisions are a kind of self-directed commands, see Kenny (1963: 216–220), Castañeda (1975: 42–43, 155), Velleman (1989: 99) For a criticism, see Rundle (1997: 189) and O'Shaughnessy (1980: ii, 342).

45. For the central role of policies in the constitution of the agent's diachronic identity, see Bratman (2007).

Could policy-based reasons offer an alternative or at least complementary account of the rational authority of decisions? This suggestion might arise by noticing that the validity of a policy-based reason is not necessarily of an epistemic/evidential character. Unlike a valid APR, a valid policy-based reason justifies the agent's ϕ-ing even if it is not true that the agent would decide to ϕ if at t_{act} she were to engage in full deliberation under suitably idealized conditions. The policy might be the result of an acceptable trade-off between the correctness of individual actions and the costs of full individual deliberations.

Should these considerations suggest a revision of the DDL view? Not really. If C obtains at t_{act} and the agent ϕs by applying P, she is not acting solely out of the policy-based protected reason. Her ϕ ing is also due to an APR. The policy-based reason concerns the application of the policy; it tells the agent at t_{act} what she is to do at that time given that C obtains. The APR instead tells her that she is to *hold onto* P (and thereby apply it) out of her prior decision to adopt it. The APR preserves the transtemporal authority of the decision to adopt the policy rather than the present authority of the policy. The policy-based reason only kicks in (at least notionally) when the agent at t_{act} is set to consider what to do then on the basis of her having *already* set, thanks to the APR, to continue to hold onto P.[46]

46. An interesting case is that of the exclusionary reasons associated with the policies for the reconsideration of decisions. Let's imagine that under circumstances C, the potential defeaters of a decision usually turn out to be false positives so that one adopts a policy Q to ignore them under C. When C holds and the agent applies this policy, she might end up dismissing a potential defeater of her APR that is not really a false positive. Despite the awareness that she is running this risk, the agent is justified in taking her APR to be valid. She is justified in thinking that there would be a convergence in verdicts even if she knows of a potential defeater. The agent is not of two minds, however. The policy Q and the APR operate at different levels. By dismissing the relevance of the potential defeater, the policy Q authorizes the agent in assuming the validity of her APR as if no such defeater existed. From that moment on, this APR operates in the standard way. To defeat this APR, the agent would have first to reject the policy Q, which is not defeated by the existence of those potential defeaters of the APR that are already factored in the considerations that supported the adoption of the policy.

6. Selection and Decision

§6.1

Among the possible grounds for exclusionary and protected reasons are the demands of coordination. The need to solve interpersonal coordination problems is a fairly common ground for accepting the dictates of an authority invested with the task of selecting among equivalent patterns of coordination (*e.g.*, choosing on which side of the road to drive). The dictates of the coordinating authority are binding even if they are arbitrary. Structurally similar situations arise in *intra*personal coordination. Several of our present- and future-directed decisions seem to addressing coordination problems via arbitrary selections. If so, the DDL view might not be the whole story of decisions given that the arbitrariness of a selection offers no basis to expect the convergence of verdicts required by valid APRs. In this section, I will argue that this is not a real threat to the DDL view.

Consider a case in which the balance of reasons under-determines my future choice, but I need to coordinate with my choice in advance. Perhaps, there are two equally choiceworthy ways of driving to a party, and I have offered to give a ride to a friend on my way there. My friend is indifferent between the two routes, but he needs to be told in advance where I will pick him up. This situation seems to require that I first make a future-directed selection and then coordinate my present and subsequent actions accordingly. This advanced coordination works only if the selection normally makes a difference to the future conduct. The idea of a selection-based reason is supposed to account for this difference by analogy with the interpersonal case of coordination-based reasons. The fact that we succeed at this kind of intrapersonal transtemporal coordination, however, does not show that selections exert distal rational authority. Consider what happens once I am already on my way to the previously selected meeting point. At that point, the prior arbitrary selection is only of historical interest to me in the sense that, had I made a different selection earlier, I would now be on a different route to the party. But there is no point in appealing to the selection as a source of a protected reason to stick to

the present route. Rather, what matters is the *fact* that I am *already* on that path and it is now too late or too expensive to turn back. Had I not already told my friend where to pick him up or had I not already been driving along the selected route, I would have no reason to stick to my selection. It seems that the future-directed selection does not make a difference to future conduct; rather, this is the job of the immediately executive action that effectively breaks the tie between the alternative routes. We should think of this situation in terms of the selection, not of my future action of ϕ-ing in isolation from what I do prior to it, but of a temporally extended course of action that culminates with my ϕ-ing and begins only when I take the first discriminating step in coordinating with the intended culmination. In this sense, there is no need for the distal authority of future-directed selections. One only needs effective selections that occur when one takes the *first actual step* that discriminates between the equally choiceworthy courses of action. Prior to calling my friend and telling him where to meet, or prior to taking the junction where the two alternative routes diverge, there really is no effective selection. I might tell myself in advance that I will take route A rather than B, say, but this selection gives me no reason to act on it. Up to the moment of the first discriminating step, I can forget about my selection or keep changing my mind about it without being irrational or jeopardizing the success of my project of going to the party and picking up my friend. On the other hand, once I am on a specific route, the selection makes no rational difference. What matters is only that I am already on route B and that I have already told my friend that I would be on B, not my selection of it over route A. This is true even if I had originally selected A. As long as I am not already on A, I am neither mistaken nor irrational if I tell my friend that I will meet him along route B. I would only be irrational if I were to take A after telling him otherwise. Advanced coordination with a future under-determined conduct, therefore, does not show that arbitrary selections can exert a distal authority of a different kind than APRs. Selections do not have this kind of authority. Nor is this authority required for successful transtemporal intrapersonal coordination.

§6.2

An analogous conclusion holds of those future-directed selections that are "bare" in that they are not meant to contribute to any advanced coordination. At no point in time do these selections exercise rational authority. There is no extended course of action on which the agent is to embark on the basis of such a selection. Hence, once the time of action comes, there is no reason to act as originally selected. There is nothing irrational or problematic with an agent who systematically forgets, mischaracterizes, or repudiates her past bare selections. The real trouble, as Buridan's ass teaches us, arises only for agents who are unable to make *presently executive* bare selections.

This is not to deny that a bare future-directed selection might affect future conduct. The selection to ϕ might instill a psychological propensity that reliably inclines, but does not physically constrain, the agent to ϕ at t_{act}. The agent can easily defuse this propensity, and she is not criticizable if at t_{act} she picks any of the other equally choiceworthy options. (The same is true of coordinating selections prior to any actual discriminating step since, up to that point, they are just like bare ones.) The operation of selection-based psychological propensities carries no rational authority. The psychological effects of bare selections do not compete with the DDL view. They do not account for an alternative ground for the validation of the authority of future-directed decisions. Rather, they account for a different kind of influence on future conduct altogether. This is not to deny that we might approve of letting ourselves be under the influence of these propensities. But rationality does not demand that we do so. In any event, the rational approval of the psychological effects of bare selections is not to be confused with granting any rational authority to alleged selection-based reasons.[47]

47. Agents who are unable to acquire psychological propensities as a result of their bare selections are not missing any fundamental rational ability, nor are they crippled in their diachronic agency. The only difference is that they have to find alternative ways of giving assurance about their future conduct when the merits of the case are under-determined. But this is not an impossible task. They can make themselves predictable, for instance, by making side-bets or promises. Real problems arise only for agents who do not understand, even if only implicitly, the role of *non-bare* selections in transtemporal

§6.3

The effect of future-directed selections is not only different in kind from that of paradigmatic future-directed decisions but also subordinated to the reliance on valid APRs. A selection is supposed to operate in the space left open by the under-determination of the merits of the case. Prior to making a selection at t_{dec}, the agent is to assess whether there is any under-determination and, if so, its extent. At the later time t_{act}, the selection is rationally acceptable only if both the merits of the case are still under-determined and the selected option is still among the equally choiceworthy ones. Normally at t_{act} the agent does not have to rehearse once again the full assessment of the background under-determination. There is a division of deliberative labor here as well given that figuring out the extent of the under-determination is often far from obvious. But the division of labor concerns the background of selection rather than the selection itself. There is no need to go beyond the DDL view to account for this phenomenon. APRs carry the authority of the assessments of background under-determination over time. This kind of protected reason is different from the one associated with decisions that involve no selection in that it does not tell the agent to ϕ at t_{act}. Rather, it rather tells the agent that there is nothing else for her to do now but make an arbitrary selection. In case a selection has already been made against the same background, the protected reason can be taken to tell the agent that she can let herself be under the influence of the psychological propensity instilled by the prior selection. Like standard APRs, these protected reasons are validated on the basis of the expected convergence (under suitably idealized conditions) of verdicts about the merits of the case, although verdicts concerning the existence and extent of the under-determination rather than of the arbitrary selection (over which no convergence could ever be guaranteed).[48] To sum up, a decision-cum-bare-selection

---

intrapersonal coordination. For this understanding is required to appreciate the structure of temporally extended activities, including those that involve no selection.

48. That APRs operate in the background of selection is shown by the fact that

might exert distal influence by a combination of the rational authority of the assessment of background under-determination *and* a rationally approved psychological propensity. But this propensity only operates on top of, and in the space left open by, the rational authority of an assessment whose distal influence is mediated by APRs, which have the structure and validation suggested by the DDL view.

7. End-of-Deliberation Reasons

§7.1

Oftentimes we have to bring a deliberation to a close[49] before all the relevant considerations have been taken into account. This happens, for instance, when the action is urgent. But even when the action lies in the distant future, we might have to come to a conclusion at once if we do not expect any future improvement in our deliberative circumstances. When a decision closes a deliberation before its ideal resting point, *i.e.*, before all relevant reasons have been duly collected and combined, the decision embodies what I call an "end-of-deliberation protected reason" (EPR, hereafter). A decision of this kind comes together with a protection from the disruption of potentially countervailing considerations that have not been taken into account in the actual deliberation. The fact that these considerations might exist or become available in the future does not affect the force of the decision, since the matter has been settled by the decision. The protection offered by the EPR is what stands behind the idea of a decision as settling a matter by bringing a deliberation to a close. In this sense, decisions taken

---

a challenge to the alleged authority of a future-directed selection is really a challenge to the authority of the background APR. Selection-based reasons per se are not defeasible, whereas background APRs are always defeasible. And when background APRs are put into doubt, a bare selection becomes *ipso facto* irrelevant.

49. Notice that the choice of when exactly to terminate a deliberation might involve some arbitrariness. But the conclusion that is reached need not be based on an arbitrary selection between equivalent options. The conclusion might actually be univocally dictated by the considerations that one has been able to take into account before closing the deliberation.

under the special conditions that allow a deliberation to reach its ideal resting point do not come with an EPR. These decisions do not have to exclude any potentially disruptive considerations since, by definition, all relevant considerations have been duly taken into account in the ideal situation.

§7.2

The protection secured by EPRs is not specifically meant to exclude *future* reconsideration. The concern with potentially disruptive evidence might arise even for a present-directed decision. The agent might know of potentially countervailing evidence at the time of decision but lack the time and resources to take it into account. Under these conditions, to decide *is* to *exclude* this evidence. This exclusion carries over time. This is however the work of APRS. EPRs are not in the business of securing the distal authority of decisions, including their role as the closings of deliberations. EPRs are part of what is transmitted by APRS (*i.e.*, EPRs are included in the full case for ϕ-ing, in the set of considerations to which a transparent APR defers if we need to articulate the justification of our decision). If the agent at t*act* does not expect to improve on previous, good enough deliberation that was brought to a close prior to its ideal resting point, she has a valid APR to act on that conclusion. That conclusion embodies the protection of the EPR.

These two kinds of reasons address separate, but related, issues. APRS concern the question "Why act as previously decided rather than engage in full deliberation at t_{act}". EPRS concern the question "Why act as decided given the existence of possible countervailing considerations". If their structure and validation were made fully explicit, EPRS would protect decisions by telling the agent both at t_{dec} and t_{act} that, once a decision is made, the matter is closed and it is thus too late to take any novel considerations into account; APRS would protect instead by telling the agent at t_{act} that all relevant evidence has *already* been taken into account at t_{dec} and that she can not do any better. Therefore, the availability of unheeded considerations does not automatically defeat either APRS or EPRS. It does not automatically defeat the APR because it

is not sufficient to show that the agent could do better at t_{act} if she were to engage in a novel deliberation even under suitably idealized conditions. It only offers a case against the APR if the APR is supposed to carry the authority of a decision that does not embody an EPR because it was made under deliberative conditions that were considered to be ideal. As for EPRS, novel considerations do not automatically defeat them, since their availability might have already been factored in the agent's determination that she was to bring the deliberation to a close at t_{dec}. The novel considerations defeat the EPR only if their availability constitutes an overall improvement on the original circumstances of deliberation.

Both APRS and EPRS help us deal with the scarcity of deliberative resources, and contribute to DDL. APRS make it possible to avoid the costly repetitions of deliberation over time; EPRS make it possible to avoid the continuation of the original deliberation past the point of decreasing marginal returns. Only APRS, however, account for the effect of the passage of time on the binding force of future-directed decisions. As such, APRS also operate in those circumstances in which the original decision embodies no EPR since it was made under perfect deliberative conditions.

§7.3

When the authority of an EPR is transmitted over time via an APR, it might seem that there is overlap and overdetermination in the structure of protection. Both reasons dictate that the agent is to ϕ by paying heed to no consideration other than her prior decision to ϕ. Nevertheless, the protections are at least notionally distinct, as shown by their different defeaters. The phenomenology of everyday agency might not register the complexity of the structure of protection. This complexity, therefore, does not normally get in the way of the standard seamless psychological transition from decision to action (see §4.1). Nevertheless, we normally rely on an implicit, albeit rough, sense of the changes of our deliberative conditions over time. Under normal conditions, we tend to accept APRS and EPRS by default but we are not blind to

their most blatant defeaters such as discoveries of the most egregious errors in the original deliberation, exceptional improvements in our deliberative conditions, or the occurrence of events, such as a major catastrophe, that would impose massive changes in our priorities and practical standpoints.[50]

§7.4

Let's take stock. In the last few sections, I have discussed various kinds of exclusionary and protected reasons that are generated by decisions but differ from APRs in terms of the source of their validation. These reasons can be grounded on legitimate authority, policies, coordination needs, and the necessity to terminate deliberation prior to its ideal resting point. The DDL view does not deny the existence of these reasons, nor does it deny the psychological effects of bare selections. It claims, however, that only APRs account for the basic rational authority of future-directed decisions. This is most evident when deliberative circumstances are ideal at t_{dec} and the merits of the case are not under-determined. A future directed decision made under these circumstances generates no other protected reason but an APR. This is the source of the basic kind of distal rational guidance of decisions. Situations of this sort are nonetheless rare. Usually, a decision comes with a bundle of protected reasons. Here is a typical scenario: (i) the original deliberation is conducted under less-than-ideal and not-improving circumstances, hence it is brought to a close before its ideal resting point, thereby embodying an EPR; (ii) the merits of the case are partly under-determined, hence the need for an arbitrary selection; (iii) what is at stake is the adoption of a policy rather than a single action. In this case, once the time of action comes, the agent is under

---

50. When we explicitly investigate the status of potential defeaters, we are not yet reopening the issue. We are rather considering whether to reopen it. At most, we might suspend the force of the exclusionary reason while we investigate the potential defeater. Hence, if at the end of the investigation we determine that the protected reason to ϕ is undefeated, our ϕ-ing still counts as a case of acting directly out of the original decision, no matter how circuitous and lengthy our investigation on the possible defeaters might have been.

an APR, an EPR, a policy-based reason, and the psychological effects of the selection. On top of them, there are also considerations about the potential defeaters of APR, EPR, and policies, which in turn might make appeal to prior decisions, policies, and selections. The complexity of these ordinary scenarios contributes to the difficulty of sorting out the specific contribution of APRs to the transtemporal rational authority of decisions. But if we consider the basic contribution of future-directed decisions to diachronic agency, APRs stand out as the exclusive source of the paradigmatic way in which future directed decisions exert distal rational authority. The other protected reasons are only part of what is transmitted by the APRs over time.

8. Conclusion

§8.1

In closing, I would like to address one residual worry about the DDL view. Some might be concerned that the DDL view fails to account for the power of decisions to *settle* in advance what the agent is going to do. If a decision is effective via the acceptance of an APR at t_{act}, it seems that what ultimately settles the matter is the later acceptance, not the earlier decision. This is true, but only in the sense that the agent exercises her rational governance at both times. In deciding to ϕ, rather than setting up some other mechanism of manipulative distal self-control, the agent exposes herself to the risk of a future repudiation of her decision. But this is the price that must be paid to make diachronic autonomy possible. At the same time, APRs contribute to settling what the agent is to do in the sense that, under normal conditions, it is reasonable to expect that the agent is going to accept the APR and act on it.

§8.2

An important element of these "normal conditions" is the assumption of the *background stability* both in the agent's rationality and in her reasons for ϕ-ing. A rational agent is not going to entrust the success of her pursuit to the future self's acceptance of an APR unless she is

entitled to the expectation that her future self will be both correctly exercising her rational capacities and in agreement with her prior assessment of the situation of choice. The latter agreement is ultimately grounded on the expectation that (i) the reasons in support of ϕ-ing stay stable over time and that (ii) one continues to acknowledge this stability. The agent must have, as George Wilson correctly writes, "a continued sense of the intrinsic force and authority of the reasons for action which have been decisive for him in his practical assessment of his position".[51] This sense is something that—I want to add—the agent can continue to have even if she becomes unable to articulate the reasons for her ϕ-ing. When that happens, the agent can defer to her prior deliberation if an articulation is ever called for. But it is this continued sense of the stability of these reasons that justifies both the agent's deferral to her past self and underpins her reliance on APRs at the time of action. If the agent at t_{act} were not to expect this stability, the APR would be defeated. Furthermore, at the prior time of decision, the expectation of the stability is what relieves the agent from the need to manipulate her future self to secure that she ϕs at t_{act}. The background stability of rationality and reasons is thus crucial for any sort of diachronic agency that respects diachronic autonomy. It is a mistake to think that the contribution of future-directed decisions is to correct or compensate for instability of reasons and rationality. Future-directed decisions are not to be confused with the tools of "egonomics"—as Schelling calls the art of managing one's own recalcitrant and irrational future selves.[52] This is not to say that this stability can be taken for granted as a trivial condition. Rather, more work needs to be done to understand when this stability can be assumed and what is required to secure it in a non manipulative way,[53] but this is a topic for another occasion.[54]

51. Wilson (2000: 18).

52. Schelling (1978: 290).

53. For discussion of the relation between stability, diachronic autonomy, and the agent's identity, see Ferrero (2009).

54. I want to thank Dan Rosenberg for help with proofreading, and Carla

References

Ainslie, George. (1992). *Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person*. Cambridge University Press.

Ainslie, George. (2001). *The Breakdown of Will*. Cambridge University Press.

Anderson, Elizabeth. (1996). "Reasons, Attitudes, and Values: Replies to Sturgeon and Piper". *Ethics*, 106.

Bratman, Michael. (1987). *Intention, Plans and Practical Reason*. Harvard University Press.

Bratman, Michael. (1999). *Faces of Intentions*. Cambridge University Press.

Bratman, Michael. (2007). *Structures of Agency*. Oxford University Press.

Broome, John. (1999). "Normative Requirements". *Ratio*, 12: 398–419.

Broome, John. (2001). "Are intentions reasons? And how should we cope with incommensurable values?" in C. Morris & A. Ripstein (eds.), *Practical Rationality and Preference*. Cambridge University Press: 98–120.

Burge, Tyler. (1993). "Content Preservation". *The Philosophical Review*, 102.4: 457–488.

Buss, Sarah. (1999) "What Practical Reasoning Must Be If We Act For Our Reasons". *Australasian Journal of Philosophy*, 77: 399–421.

Castañeda, Hector-Neri. (1975). *Thinking and Doing: the philosophical foundations of institutions*. Reidel.

Davidson, Donald. (1980). *Essays on Actions and Events*. Oxford University Press.

Elster, John. (1979). *Ulysses and the Sirens*. Cambridge University Press.

Falk, W. David. (1953). "Goading and Guiding". *Mind*, 62: 145–171.

Ferrero, Luca. (2006). "Three Ways of Spilling Ink Tomorrow". In E.

Baccarini & S. Prijic-Samarzija (eds.), *Rationality in Belief and Action*, Rijeka, Filozofski faulted: 95–127

Ferrero, Luca. (2009). "What Good is a Diachronic Will?" *Philosophical Studies*, 144.3: 403–430.

Goetz, Stewart. (1998). "Reasons for Forming an Intention: A Reply to Pink". *Mind*, 107: 205–213.

Hartogh, den Govert. (2004). "The Authority of Intention". *Ethics*, 115: 6–34.

Hinchman, Edward. (2003). "Trust and Diachronic Agency". *Noûs*, 37.1: 25–51.

Joyce, James M. (1999). *The Foundations of Causal Decision Theory*. Cambridge University Press.

Kenny, Anthony. (1963). *Action, Emotion and Will*. Routledge.

McClennen, Edward F. (1990). *Rationality and Dynamic Choice*. Cambridge University Press.

McClennen, Edward F. (1998). "Rationality and Rules". In P.A. Danielson (ed.), *Modeling Rationality, Morality, and Evolution*. Oxford University Press. (1998).

Mintoff, Joe. (1999). "Are decisions motive-perpetuating?" *Analysis*, 59: 266–275.

Mintoff, Joe. (2001). "Buridan's Ass and Reducible Intentions". *Journal of Philosophical Research*, 26: 207–221.

Mintoff, Joe. (2002). "How Can Intentions Make Actions Rational?" *Canadian Journal of Philosophy*, 32.3: 331–354.

Moran, Richard. (2001). Authority and Estrangement: An Essay on Self-Knowledge. Princeton University Press.

O'Shaughnessy, Brian. (1980). *The Will*. Cambridge University Press.

Pink, Thomas. (1996). *The Psychology of Freedom*. Cambridge University Press.

Raz, Joseph. (1975). *Practical Reasons and Norms*. Oxford University Press.

Raz, Joseph. (1978). "Reasons for Action, Decisions and Norms". In J. Raz (ed.), *Practical Reasoning*. Oxford University Press: 128–143.

Raz, Joseph. (1979). *The Authority of Law*. Oxford University Press.

Raz, Joseph. (1990). *Practical Reason and Norms*, 2nd edition. Princeton University Press.

Rovane, Carol. (1998). *The Bounds of Agency*. Princeton University Press.

Rundle, Bede. (1997). *Mind in Action*. Oxford University Press.

Scanlon, Thomas. (1998). *What We Owe To Each Other*. Harvard University Press.

Schelling, Thomas. (1978). "Egonomics, or the Art of Self-Management". *The American Economic Review*, 68, 2: 290–294.

Sobel, J. Howard. (1994). *Taking Chances*. Cambridge University Press.

Velleman, J. David. (1989). *Practical Reflection*. Princeton University Press.

Velleman, J. David. (1997). "Deciding How to Decide". In G. Cullity & B. Gaut (eds.), *Ethics and Practical Reason*. Oxford University Press: 29–52.

Velleman, J. David. (2000). *The Possibility of Practical Reason*. Oxford University Press.

Velleman, J. David. (2007). "What Good is a Will?" in A. Leist & H. Baumann (eds.), *Action in Context*, de Gruyter/Mouton.

Wilson, George. (2000). "Proximal Practical Foresight". *Philosophical Studies*, 99: 3–19.