

Extended Agency and the Problem of Diachronic Autonomy*

Julia Nefsky and Sergio Tenenbaum

University of Toronto

1. Introduction

It seems to be a humdrum fact of human agency that we often act on intentions or decisions that we have made at an earlier time.¹ At breakfast, you look at the Taco Hut menu online and decide that later today you'll have one of their avocado burritos for lunch. You're at your desk and you hear the church bells ring the noon hour. You get up, walk to Taco Hut, and order the burrito as planned.

As mundane as this sort of scenario might seem to be, philosophers have raised a problem in understanding it. If you are simply abiding by this morning's decision, how are you acting autonomously? Your earlier self seems to be calling the shots; if you are just acting accordingly, without thinking through it or in some other way trying to ensure that the past decision conforms to

* For very helpful comments on earlier drafts of the paper, we would like to thank Michael Bratman, Luca Ferrero, Kim Frost, Amelia Hicks, Dana Howard, Hanti Lin, Berislav Marusic, Sarah Paul, Regina Rini, Tina Rulli, John Schwenkler, Itai Sher, Matthew Silverstein, Walter Sinnott-Armstrong, Karen Stohr, Larry Temkin, Mike Teitelbaum, and Stephen White. We would also like to thank audiences at Duke University, the 2015 Workshop in Ethics at Dartmouth University, the 2015 St. Louis Annual Conference on Reasons and Rationality, the Belief Rationality and Action Over Time Workshop in Madison, the 2017 Meeting of the Israeli Philosophical Association, the Varieties of Agency Workshop at Stanford University and departmental colloquia at Universidade de Campinas, Universidade Federal do Rio de Janeiro and Université Paris I (Panthéon-Sorbonne). We greatly profited from the discussions in all these venues.

¹ In this paper, we use "future-directed intention" and "future-directed decision" interchangeably. More precisely, our view is that a future-directed decision is a kind of future-directed intention, namely one that is formed by deliberation. There are other future-directed intentions that are not formed by deliberation, but this doesn't matter for our purposes here. So we can, for simplicity, use the two terms interchangeably.

your present standpoint, it is not clear how this amounts to an exercise of your present autonomous agency. It seems, rather, that your earlier self has succeeded in slaving you to her own purposes. She was the one who wanted (intended, judged it to be good, etc.) to have an avocado burrito; in simply following through, your current self seems to be just an automaton performing the commands left behind by your former self.

Of course, you might not allow yourself to be shackled by your earlier self. You might refuse to follow anything but your own present judgments: you will only go to Taco Hut if this is what you judge you should do right now, and once at Taco Hut will only eat the avocado burrito if that is what you want to eat once there. But if this is the way you generally operate, this seems to block your ability to make effective future-directed decisions. If you will always open up the question of what to do when the time comes, acting only on your present judgments, then you cannot successfully decide in advance what to do later. But, as we have said, this is an ability that we do seem to have and employ frequently.

The puzzle, then, is one of explaining how the future self can do the bidding of her past self without losing her autonomy. Or, as Luca Ferrero puts it, how can our future-directed decisions be “effective without being manipulative”²? Let’s call this “the Problem of Diachronic Autonomy.”

Philosophers raising the Problem of Diachronic Autonomy take it to show that there must be reasons or rational requirements to follow through with our past decisions. They claim that unless our decisions put rational pressure on us to follow through, we cannot solve the problem.

In this paper we examine these philosophers’ different proposals, and we argue that none of these views seem to correctly capture the autonomy that we have in acting on our past decisions. We then argue that what is going wrong with these views is that they are, effectively, trying to solve a

² Luca Ferrero, “Decisions, Diachronic Autonomy, and the Division of Deliberative Labor”, *Philosopher’s Imprint* 10: 2 (February 2010), 1-23, at p. 1.

problem that isn't there. In other words, there is no Problem of Diachronic Autonomy. There is no puzzling situation that needs explaining. Consequently, there is no need coming from this purported puzzle to think that our future-directed decisions generate reasons or rational requirements to follow through. In particular, we argue that the "future self" can do the bidding of the "past self" without giving up its autonomy because, very simply, the past self is the *same agent* as the future self. In following through with my decision, I am acting on my own freely formed intention; that is straightforwardly autonomous. There is nothing puzzling or difficult here.³ Indeed, we argue that the Problem of Diachronic Autonomy, as specified, is ultimately incoherent: it implicitly relies on two inconsistent conceptions of human agency. We then consider other possibilities for what people might have in mind in thinking that there is a special problem of diachronic autonomy, and argue that none of them work either. There is no relevant difference between the diachronic case, in which I make a future-directed decision and later act on it, and the synchronic case, in which I act immediately on a decision about what to do in the present.

2. Attempts to Solve The Alleged Problem

David Velleman argues that we can only solve the Problem of Diachronic Autonomy if in making a future-directed decision one gives one's future self reason to follow through with it:

The only way to control our future behavior without losing future control, I believe, is by making decisions that our future selves will be determined to execute of their own volition; and the only way to determine our future selves to do something of their own volition is by giving them reason to do it. Hence future-directed intentions or commitments must be capable of providing reasons to our future selves. Unless we can commit ourselves today in a

³ There is, of course, a lot of disagreement about what it takes for an action to be autonomous. This does not matter for our purposes. As long as you accept that for my act to be autonomous, I must at least act on my own intention (or reasons, or desires) then we have what we need for the arguments we want to make. While there will be other conditions for an act to count as autonomous beyond that basic one, it is that basic one that is at play in the Problem of Diachronic Autonomy: the concern is that if I act directly on my past decision, then I am not acting on my *current self's* intentions (or desires or reasons), and am instead being manipulated by my past self.

way that will generate reasons for us to act tomorrow, we shall have to regard our day-older selves either as beyond the control of today's decisions or as passive instruments of them.”⁴

Velleman recognizes that some people's normative intuitions oppose the idea that our decisions generate reason to follow through with them: they “think that abiding by a commitment for its own sake is foolish.”⁵ But, he says, whatever your normative intuitions are on the matter, the question of whether there are such reasons is settled by the need to solve the Problem of Diachronic Autonomy. A fact about human agency is that we have the ability to decide for our future selves without undermining our future autonomy, and – Velleman claims – we can only make sense of this fact *if* our decisions generate reasons to follow through.

How does making a future-directed decision generate reason to follow through with it?

Velleman does not give a full answer, but he does sketch an idea of at least one way this could work: namely, if we accept his view that autonomy is the constitutive aim of action. He writes:

[I]f autonomy is the constitutive goal of action, and hence the internal criterion of success for action, then reasons for acting will be considerations relevant to autonomy [...] And we can at least hope that reasons of this kind will be generated by future-directed decisions.⁶

Like Velleman, Sarah Paul claims that for a future-directed intention to serve its function of settling a deliberative question, “the agent must view the fact that she has formed that very intention as carrying a distinctive weight that is independent of the weight of the reasons favoring the

⁴ David Velleman, “Deciding How to Decide” in *Ethics and Practical Reason*, G. Cullity and B. Gaut, ed. (Oxford University Press, 1997), 29-52, at p. 46.

⁵ Velleman, op. cit., p. 49.

⁶ Velleman, op. cit., p. 49. Velleman explains that on a *maximize utility* conception of the aim of action, there is no room for a future-directed decision to ϕ to generate reason to ϕ . On such a conception, one only has reason to ϕ *if* doing so maximizes utility. So, if Velleman is right that our future-directed decisions must generate reason to follow through, then the maximizing-utility conception cannot be correct. He explains: “If an action were the sort of thing whose success or failure could be judged solely by utility-maximizing considerations, then it wouldn't be the sort of thing that we could decide on today in a way that would necessarily give us reason to perform it tomorrow... But action *is* that sort of thing--it *is* behaviour over which rational agents have diachronic autonomy.” (Velleman, p. 49)

intended action.”⁷ Paul’s story of how this works is different from Velleman’s. Paul focuses on cases in which one’s options are on par or incommensurable.⁸ So, take a case in which you decide to go to Taco Hut for lunch, but in which lunch at Taco Hut is, in your assessment, completely on par with lunch at Fajita Palace. In such a case, there is no direct cost to breaking with your decision when lunchtime comes and going to Fajita Palace instead. But, Paul claims, there is “an indirect cost” to not following through with your decision: it gives you evidence that your decisions do not determine what you will do, and this evidence erodes your ability to treat your decisions as authoritative when there are benefits to be gained from doing so – for instance, in cases in which you are using future-directed decisions to try to avoid temptations or to try to cooperate with others.⁹ So, in Paul’s view, our decisions generate reasons to follow through because following through with them helps maintain our capacity to treat our decisions as authoritative, and this is a useful capacity to have.

While the details of their proposals differ, Paul and Velleman agree that our future-directed decisions must give us reason to act accordingly because without this we cannot make sense of how future-directed decisions can decide deliberative questions effectively without being manipulative. So, on both of their views, I have reason to go to Taco Hut and order the avocado burrito coming from the fact that I decided in the morning to do so. Let r be the reason in question.¹⁰ The problem with both proposals is that now when I go to Taco Hut and order the burrito, the reason I do so is

⁷ Sarah Paul, “Diachronic Incontinence is a Problem in Moral Philosophy,” *Inquiry*, special issue on Choice Over Time, Sergio Tenenbaum, ed. 57: 3 (June 2014), 337-355, at pp. 349 - 350.

⁸ Such cases are useful in thinking about whether a decision to ϕ generates reasons to ϕ , since they rule out that other reasons to ϕ (e.g. that it’s easier to ϕ than to take the other options) are what are really doing the work.

⁹ Paul, op. cit., p. 351.

¹⁰ For Paul r is, roughly, that every time one follows through with a decision this helps maintain (or, at least, avoids eroding) one’s ability to understand one’s own decisions as authoritative, and this is an ability that we have an interest in maintaining. For Velleman r is something like: following through will preserve my autonomy (or, some other reason connected to the constitutive aim of action, autonomy.)

(in part) *r*. But *r* is, by definition, not the reason why I formed the intention to order the burrito in the first place. I formed the intention to order the burrito because it would be a delicious lunch. So, a view like Paul's or Velleman's only vindicates my autonomous acting on a past decision at the cost that I do not act on the same reasons for which I made the decision. I now eat my burrito to preserve my autonomy, or to preserve my capacity to make effective future-directed decisions, rather than (or, rather than just) because it would be a delicious lunch. A picture on which I conform with my decision but for different reasons from those for which I made the decision does not seem to be an accurate or acceptable picture of how our future-directed decisions guide our actions.¹¹

Furthermore, such a picture seems at risk of actually undermining, rather than providing an explanation of, diachronic autonomy. On this picture, it is by giving you this new reason that your decision exerts control over your future actions, and because it exerts control via this reason, this is supposed to make the control non-coercive.¹² But does this make it non-coercive? On this view, my

¹¹ Velleman might reply to this objection as follows: while the future-directed decision generates reason to follow through with it, one would only act *for* this new reason if one reconsidered before following through. If one does not reconsider, then one might simply follow through with the original decision, acting for the original reasons. If this is Velleman's view, this avoids the objection that on his account I must be eating the avocado burrito at least partly in order to preserve my autonomy. But even if Velleman can avoid that first objection in this way, this does not help with the issues we will raise in what follows. This is because even if I need not act *for* the new reason, it is still true – on his account – that the existence of this new reason is doing the work of explaining why my following through is autonomous. That is, whether or not I explicitly entertain it or act for it, his solution to the Problem of Diachronic Autonomy is that I have this new reason and that this new reason is (part of) what justifies my following through with the decision. The new reason is something I should take into account if I do reopen the question, and if I do not reopen it, the reason is still doing “behind-the-scenes” work in justifying my action and explaining it as autonomous. This is enough for the next objection we make, and it is enough for our general critique in section three.

¹² In this paper, we use the terms 'manipulation' and 'coercion' fairly loosely to capture the ways of interfering with someone's agency that are of concern in the Problem of Diachronic Autonomy. Interpersonal coercion often involves threats, but, of course, intrapersonal coercion could not be threats-based. But in the interpersonal case too, coercion need not involve threats: it can occur through acts that constrain the agent's choice situation. If I set up an automated electric shock system so that every Tuesday, if you do not come visit me, you will receive a very painful shock, I am coercing you into visiting me on Tuesdays, even though I have not issued any threat.

past self is doing something (deciding) that gives my present self a reason that it would not otherwise have had to act according to my past self's will. Controlling someone's behaviour in this way can be a way of coercing or manipulating them. As Ferrero writes, controlling your future behavior by "introducing features extraneous to the original merits of the case" is precisely what happens in "manipulative forms of distal self-control like pre-commitments."¹³

Velleman's and Paul's idea seems to be that the control my past decision exerts is not coercive because it goes *through* my rational deliberative capacities, appearing in the reasons that I entertain in deciding for myself now whether to go through with the decision, rather than bypassing my deliberation with brute causal force. But this argument does not work. My decision to give my wallet to the person who threatens, "your money or your life" also goes through my rational deliberations: their threat gives me reason to give up my wallet, a reason that outweighs my reasons to keep it. But that surely does not make this threat a non-coercive way of exerting control, nor does it make giving up my wallet a case of purely autonomous agency.

Perhaps, though, the specific sort of reason that one's past self gives one's future self on Velleman or Paul's account helps explain why the situation is non-coercive. After all, not all cases of giving someone a reason to act as you wish are cases of coercing them.¹⁴ But looking at the content of the reason does not seem to help here. The content is, in very rough and general terms, something like: *I* decided to ϕ and if you don't do as I – your past self – decided, then your agential

¹³ Ferrero, *op. cit.*, p. 1.

¹⁴ In particular, we often can give others positive incentives to act as we wish without coercing them. (E.g. "if you help me move, I'll buy you a nice dinner".) Interestingly, it is hard come up with a clean case of giving *yourself* a "positive" incentive. If I promise myself ice cream as an incentive to finish unpacking, it might seem as though I'm giving myself a positive incentive. But ice cream, in this example, must be something that I have access to unless I withhold it from myself. So, it seems more accurate (or, at least, equally accurate) to characterize this as a case of negative incentives: the incentive is that unless I finish the unpacking I will withhold ice cream from myself.

capacities will be compromised in such-and-such a way.¹⁵ Giving someone this sort of reason to φ certainly seems coercive.

There are two ways philosophers have avoided having the future self acting for different reasons from those that came into the decision, while preserving the idea that the decision controls future behaviour non-coercively by changing the rational deliberative scene for the future self. The first one, proposed by Bratman, is to argue for wide scope rational requirements governing the relation between future-directed intentions and their later execution. If the rational pressure that past decisions exert on our current selves is determined by wide scope requirements, we need not say that these decisions generate new reasons to act. A second possibility, proposed by Ferrero, is that decisions give us new exclusionary reasons – that is, second-order reasons to exclude certain kinds of reasons from our deliberations. This would also avoid the conclusion that the past decision creates a new reason for action. We will consider each of these views in turn.

Let's start with the proposal that we are diachronically autonomous when following through with a past intention insofar as and because doing so is in conformity with a wide scope rational requirement. Bratman's proposal for the wide scope rational requirement is:

(D) You ought (If you intended at t_1 to X at t_2 and throughout t_1 – t_2 confidently took your relevant grounds adequately to support this very intention, to X at t_2 .)¹⁶

Now, to be clear, Bratman himself does not present *D* as a direct response to the Problem of Diachronic Autonomy. Rather, Bratman argues in a different way that diachronic autonomy requires *D*. That is, he argues that in order to have diachronic autonomy (or, in his terminology “self-

¹⁵ This should not be read as a threat. It is, on the accounts in question, what you *make true* by deciding for your future self, rather than something you threaten to do to your future self. It looks, thus, more like the case of coercing by setting up an electric shock system (see footnote 12 above), rather than coercing via threats.

¹⁶ Michael Bratman, “Time, Rationality, and Self-Governance”, *Philosophical Issues* 22:1 (October 2012), 73-88, at p. 79; we have rephrased it into “You ought(If p, to q)” form.

governance over-time”¹⁷) we need *D*, but he does not argue for this by appeal to the puzzle about how our past decisions can control our behavior without undercutting our autonomy.¹⁸ In this section, we will look at *D* as a response to the Problem of Diachronic Autonomy (even though Bratman himself does not present it that way). But as we will explain later, at the end of section three, our main point carries over to Bratman’s actual argument as well.

If we put rational requirement *D* as a response to the Problem of Diachronic Autonomy, the idea would be as follows: since your past self made the decision to ϕ , this limits what it is rational for your current self to do: either your current self must change (if she hasn’t already) her assessment of the relevant grounds *or* she must follow through. Following through, then, is an exercise of your autonomous agency, just an exercise that is – rightly – responsive to the rational constraints that the past decision puts in place.

While invoking a rational requirement in this way might avoid positing new first-order reasons, it is not clear that this view does any better at capturing following through on one’s past decisions as autonomous. Whether one appeals to new reasons created by past decisions or to rational requirements whose antecedent is partly made true by a past decision, the past decision controls my future conduct by way of rational pressure on me to follow through. But, putting rational pressure on someone to act as you wish can be a way of manipulating or coercing them; it

¹⁷ Bratman notes that “self-governance over time” for him is the same idea as “diachronic autonomy” for Velleman. See Bratman, *op. cit.*, p. 87, n.26.

¹⁸ It’s worth noting that Bratman does raise the Problem of Diachronic Autonomy elsewhere. For instance, in “Toxin, Temptation and the Stability of Intention”, Bratman writes that a theory of instrumentally rational planning agency “needs to be responsive to a fundamental tension”: “On the one hand, a planning agent settles in advance what to do later. On the other hand, she is an agent who, whatever her prior plans, normally retains rational control over what she does when the time comes. Following through with one’s plan is not, after all, like following through with one’s tennis swing. We need to do justice to both these aspects of planning agency.” See his “Toxin, Temptation, and the Stability of Intention” in Michael Bratman, *Faces of Intention* (Cambridge University Press, 1999), p. 60.

does not seem to make a difference whether the pressure is coming from a first-order reason or from a rational requirement. Suppose I want you to do my laundry. My clothes are very dirty from gardening, and I feel lazy after all the hard work. I ask you if you would mind taking on the task. You refuse. You explain that you have more important things to do today than my laundry: you have big plans to thoroughly clean your house. In light of this information, I run into your house, scatter my dirty laundry on the floor, and blockade the doors and windows so that you cannot get rid of it. I know how meticulous you are when you clean: so meticulous that you wouldn't consider the task complete if there was a single article of dirty clothing on the premises. Since there is a rational requirement for means-end coherence, I have made it such that you are now rationally required to either abandon your cleaning project or to do my laundry. In this way, I have put rational pressure on you to do my laundry. This is, quite obviously, a case of coercing or manipulating someone. By manipulating the rational requirements that apply to you, I am interfering with your autonomy.

Ferrero's response to the Problem of Diachronic Autonomy seems to avoid this problem. According to Ferrero, rather than a new first-order reason or a rational requirement, my future-directed decision gives rise, under certain circumstances, to a second-order exclusionary reason to follow through with it. On his view, we make future-directed decisions because our finite capacities require that we have a "division of deliberative labour"; if I expect that I am in a better position to deliberate now than I will be at the time of the action, I should decide now what I will later do. But if this is true, and I decide now to ϕ later, it would be wasteful and "risky" to reopen deliberations when it comes time to ϕ . The fact that I deliberated in good conditions thus creates an exclusionary reason not to reconsider first order considerations in favour or against ϕ -ing, and to instead simply act on the basis of my knowledge of my previous decision. This exclusionary reason is – Ferrero says – "maximally protected", meaning that it cannot be defeated by first-order considerations. Such

exclusionary, maximally protected reasons are generated by an agent's past decision as long as (a) "the agent's deliberative conditions at [the time of the action] have not improved over the conditions at [the time of the decision]"; and (b) "there is a reasonable expectation that she would reach the same conclusion if she were to deliberate in conditions as good as those she enjoyed [when she made the decision]".¹⁹

So, Ferrero's solution to the Problem of Diachronic Autonomy is that, insofar as you are following through with your past decision when conditions (a) and (b) hold, you are exercising your autonomous agency in doing so, because in doing so you are responding to the exclusionary reason that your decision gives you – a reason that says to "φ without any further ado."²⁰ Now, like the other views, Ferrero's account tries to vindicate one's diachronic autonomy by saying that one's past decision puts a certain sort of rational pressure on you to follow through. But, unlike the other views, this sort of rational pressure does not seem at risk of being coercive. The "pressure" to follow through comes entirely from my current self's recognition that my past self deliberated better than (or at least, as well as) I could expect to deliberate now based on our shared practical perspective.

The first thing to notice is that on Ferrero's view acting on my past intention is exactly like acting on expert advice. I act on my past intention because I recognize that it would be better, or at least expectedly better, to act on a "borrowed opinion"; it's just that the borrowed opinion comes from my past self. But is this plausible? It seems to us that there is an important difference between following through with my past decision and deferring to the advice of an expert.²¹

¹⁹ Ferrero, *op. cit.*, p. 10. Ferrero imposes further conditions but they do not affect our argument. We cannot do justice here to all the aspects of Ferrero's interesting discussion of the division of deliberative labour.

²⁰ Ferrero, *op. cit.*, p. 13.

²¹ What this difference is will become clear in the next section. Note that we don't deny that there might be special cases in which following through on your past decisions is like following expert advice. Our point is that this is not the normal case of acting on a past decision. See fn. 33 for more on this.

A related objection is that if conditions (a) and (b) must be satisfied in order for acting on your past decision to count as autonomous then the scope of cases of diachronic autonomy would be very limited. Take our original example of acting on one's earlier decision to have lunch at Taco Hut. It seems unlikely that either (a) or (b) is satisfied in that case. I'm probably in a much better position at lunchtime to deliberate on what kind of meal I would enjoy for lunch. And it might also be true that there is no reasonable expectation that I would reach the same conclusion if I were to deliberate again when lunchtime comes (under equally good conditions to my original deliberations); after all, I know that I often change my mind when I reconsider my lunch options. But despite the fact that (a) and (b) are probably not satisfied, I still seem perfectly autonomous if I do not reconsider and simply act on my earlier intention and go to Taco Hut.²²

Perhaps – one might suggest – this means that Ferrero should drop or modify conditions (a) and (b), so that the maximally protected, exclusionary reason applies *even if* I am in a better position now to deliberate about the matter, and even if I would reach a different conclusion were I to deliberate now under equally good conditions. This way he could still capture the Taco Hut case (and others like it) as a case of diachronic autonomy. But the problem with Ferrero's account is not just that (a) and (b) are likely not satisfied; it's also that it is simply not plausible that there is a maximally protected, exclusionary reason not to re-check first order considerations in the Taco Hut case. There surely would be nothing wrong if, while walking to Taco Hut, I reconsidered my

²² You might propose that this is because where I am going to have lunch is a rather unimportant decision. Maybe the conditions for autonomy are not as strict in fairly trivial cases like this. But most of our actions on the basis of future-directed intentions are, in the grand scheme of things, rather minor. And it's not that trivial pursuits are immune to failures of autonomy. It would be very different, for instance, if my action was the result of the subliminal but infallible effect that Taco Hut commercials have on my mind. My following through with a decision I made this morning, even if the decision was made under less than ideal circumstances, seems to bear no resemblance to this kind of manipulation.

decision and chose to go to Rawesome instead on the basis of the fact that Rawesome has healthier options than Taco Hut. Ferrero might try to argue that there is something wrong with reconsidering: it would be a waste of my deliberative resources. But why think that? As I walk to Taco Hut, deliberating about lunch might be no worse of a use of my faculties than humming “The Wheels on the Bus” in my mind. So, Ferrero’s appeal to maximally protected, exclusionary reasons does not seem to work to explain why I am autonomous when I act on my decision to go to Taco Hut. Since there is nothing unusual about the Taco Hut example, the point will carry over to many more cases of acting on my past decisions.

3. It Was Me All Along

Despite the variations in their views, all four philosophers agree on the same general point: that our past decisions must give us reasons or rational requirements to abide by them, because without this we cannot make sense of our diachronic autonomy. We now will argue that this general idea is a mistake.

The Problem of Diachronic Autonomy asks: when I follow through with my decision to have an avocado burrito for lunch, how is it that I am acting autonomously? How is it that I am not being coerced or manipulated by my earlier self, if it was my earlier self who settled the question for me of what I will have for lunch? Our view is that there is nothing puzzling or difficult to explain here. The answer is simple: I am acting autonomously when I go to get the avocado burrito, because *I* was the one who decided to get the burrito – not someone else. I am not being coerced or manipulated by someone else (my earlier self) into executing their will, instead of my own. Rather, I am a temporally extended agent, and so it is the same self (me!) who intends and performs the action. This is the commonsense response to the Problem of Diachronic Autonomy. And this commonsense, simple answer turns out to be the correct one.

This answer relies on accepting that our agency is temporally extended.²³ It requires, in other words, denying a time-slice conception of our agency: a conception on which myself-at- t_n is a different agent from myself-at- t_m . But the philosophers who raise the Problem of Diachronic Autonomy recognize that we are temporally extended agents. They do not want to subscribe to a time-slice conception. Rather, they present the problem specifically as a problem in understanding the diachronic nature of our agency. The fact that we are temporally extended agents means that we can decide now and execute later, and they think that, in cases in which we do this, the temporal distance between the decision and the execution creates a puzzle. They think, in other words, that *given* that (\mathcal{A}) we are temporally extended agents, there is a problem in understanding how (\mathcal{G}) our future-directed decisions can guide our actions without being manipulative. Our claim is that once we accept \mathcal{A} there is, in fact, no problem at all in explaining how it is that \mathcal{G} .

In fact, more than that, it seems that in order to be able to formulate the supposed problem to begin with one must implicitly rely on both of these two mutually inconsistent conceptions of human agency – a time-slice conception and a temporally extended conception. First, to set up the problem, we need a temporally extended conception of our agency. On a time-slice conception, your relation to your past selves is of the *same type* as your relation to other people. So, on that conception, your past self can only decide – or, settle a practical question – for you in the sense that someone else can settle a practical question for you. Someone else can settle a practical question for you by making a convincing argument, or by offering you an incentive that you cannot resist, or by forcing

²³ We are not assuming any particular solution to the problem of personal identity over time, nor are we assuming that there is no such problem. We do assume, as we will presently explain, that human agents, or persons, persist through time such that – for instance – myself from this morning is (in the normal case) the same person as myself from this afternoon. But we do not assume or require any particular account of how this works. Whatever the correct account is of what makes it the case that myself from this morning is the same person as myself from this afternoon, this should be fine for our purposes here, and thus we can remain neutral on this difficult question.

your hand, or by some other causal shenanigans. But none of these are the same as forming an intention for yourself to act on at a later time. On a temporally extended conception of our agency, on the other hand, you – as a single agent – persist through time, and so you can form an intention now about what you will do later, and then, later, act on it. Since the Problem of Diachronic Autonomy is a problem about cases in which we form intentions for ourselves to act on at a later time, it thus requires a temporally extended rather than time-slice conception.²⁴ But second, in setting up the puzzle, we must implicitly assume a time-slice conception of our agency. It is only on a picture under which my earlier self is a *different agent* from my later self that simply acting on my earlier freely-formed decision could even look like a case of being coerced or manipulated.

It is worth noting that the time-slice conception of agency is often relied upon in the language the above-mentioned philosophers use in discussing the problem. While they are explicitly recognizing and studying the cross-temporal features of our agency, the time-slice conception seems to creep in. Paul, for instance, asks, “Finding myself with an intention formed last week to go to the dentist today, why should I allow these instructions from the past to guide *my* action – permit myself to be manipulated by the ‘dead hand of the past’, so to speak?”²⁵ Velleman writes that the only way of making sense of our diachronic autonomy is if we “make decisions that our future selves will be determined to execute of *their own* volition,”²⁶ and he talks about what it takes for a person’s intention from the past to become a “volition of his current self” – for the “volition to become *his own*.”²⁷ Ferrero, in discussing an agent who follows through with her decisions without sensitivity to the conditions required for the exclusionary reason to apply, writes:

²⁴ Now, the time-slice conception might have it’s own version of the Problem of Diachronic Autonomy. We will discuss this briefly in section five.

²⁵ Paul, op. cit., p. 339.

²⁶ Velleman, op. cit., p. 46, our emphasis.

²⁷ Velleman, op. cit., p. 47, our emphasis.

[The agent's] uncritical acceptance of a past decision ... [does not] respect her autonomy at a later time, since ... 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} .²⁸

All of these ways of talking treat the self who formed the intention as a different agent from the one who might act on it.²⁹

Once we consistently recognize the temporally extended nature of our agency, we see that the relation between future-directed intention and future doing is no different than the relation between current intention and current doing. That is, at least as far as autonomy is concerned, there is no difference between the diachronic and the synchronic case. In either case, if (a) you intended, (b) you acted, and (c) you acted because you intended, then you acted autonomously.³⁰ The passing of time between the initial formation of the intention and the execution does not change this. Of course, the passing of time does give you more of an opportunity to rethink the decision and change your mind. But often we don't rethink our decisions; often we are simply moved to act by the intentions we formed earlier. I might decide now to call my Grandma Doris when I finish writing this paragraph. Having formed that intention, when I get to the end of the paragraph, I might simply close my computer, pick up the phone and dial. That is, the future-directed intention that I formed might simply persist and directly cause my intentional activity of calling my grandmother, without any new exercising of my rational deliberative capacities. The philosophers in question think that if this is the way things go, then "our later selves would lack autonomy of their own, since they would

²⁸ Ferrero, *op. cit.*, p. 15, our emphasis.

²⁹ In "A Planning Agent's Self-Governance Over Time", forthcoming in *Planning, Time and Self-Governance*, Oxford University Press, Bratman is explicit that, at least for him, these ways of talking are metaphorical. He doesn't think there are literally multiple agents at play; it is just one agent acting at different times. The metaphor, though, he thinks is helpful. (p. 7) But the worry is that, without the metaphorical talk of more than one agent, the problem simply disappears. See footnote 38 for more detail.

³⁰ There will likely be other conditions that must be satisfied for the act to count as autonomous (e.g. that there are no deviant causal chains), but this does not make a difference to our point.

find their limbs being moved by the decisions of earlier selves.”³¹ But things are quite the opposite. If I formed the intention of my own accord, and this intention moves me to act, that is *paradigmatically* autonomous. To think that being moved to act directly by the intention I formed earlier would be a case of lacking autonomy is, again, to mistakenly treat my future self and past self as two different agents.

On our view, future-directed decisions are effective at settling deliberative questions in virtue of the fact that, in the absence of interference, we will act on them, exactly like in the synchronic case. They do not gain their ability to control our future behaviour via some pressure they put on our future deliberations. Rather, they can control our future behaviour because we can act on them.³²

In sum, the answer to the question of how I am acting autonomously when I follow through with my future-directed decision to ϕ is simple and commonsense. I am acting autonomously when I ϕ because *I* decided to do it: I am following through with my own decision. There is no need to invoke special reasons or rational requirements in order to explain this.³³ And for this very reason,

³¹ Velleman, op. cit., pp. 45 – 46.

³² It should be clear that we do not claim to have shown that future-directed decisions never generate reasons or rational requirements to follow through with them. What we are arguing is that there is no argument for such reasons or rational requirements *from* the Problem of Diachronic Autonomy: we do not need them in order to explain how it is that our future-directed decisions effectively control our behaviour or how it is that we are autonomous when we act on them. This does not rule out that there may be some totally other sort of argument for them.

³³ To clarify: this is our understanding of normal, typical cases of acting on past decisions. We don't want to deny that there could be unusual cases in which acting according to one's past decision is like following the advice of an expert. Suppose that I am walking in the park and see a group of people practicing Tai Chi. Tai Chi has never appealed to me before, but this time it suddenly does. I decide that I will sign up for Tai Chi classes. Immediately, after making the decision, a stray baseball hits me on the head and I am knocked unconscious. When I wake up, I have no persisting intention to sign up for Tai Chi, and indeed no memory of having made the decision. If I later remember having made the decision and, on the basis of this memory alone (without having any idea whatsoever what reasons I might have had to sign up for Tai Chi), sign up for classes, this would seem to be a case of treating my past self's decision as something like expert advice. But this is far from the normal case of acting on a future-directed decision. In the normal case, the decision either

this commonsense account of diachronic autonomy does better than the others in terms of how autonomous we turn out to be. That is, it does better than the others at capturing following through on one's past decision as autonomous (and this despite the fact that those other accounts put more effort into the task.) On all other accounts, my current self would not have been autonomous *but for* the fact that it is responding to rational pressure imposed on it by my earlier self. To the extent that this counts as a way of being autonomous at all, this is a rather distorted, minimal kind of autonomy.³⁴ On our account, on the other hand, I am autonomous when I follow through with my future-directed decision in the most basic and full sense: I am autonomous because I am acting on my own, non-coercively formed decision.³⁵

Before moving on, we want to come back to Bratman's argument. As we said above, Bratman does not himself argue for *D* (his rational requirement for intention stability over time) by appeal specifically to the Problem of Diachronic Autonomy. He argues for *D*, instead, by appeal to his account of what diachronic autonomy, or – in his terminology, “self-governance over time” – consists in. We now want to argue that our main point applies to Bratman's argument for *D* as well.

In arguing for (D), Bratman is particularly concerned with “Buridan's Ass” and incommensurability cases, in which our reasons support multiple incompatible courses of action. In such cases, in the absence of a rational requirement to stick with one's decisions, a rational agent could engage in “brute shuffling”; she could abandon a previously formed intention to ϕ in favour of an intention to ψ while having no reason to prefer ψ -ing over ϕ -ing (as long they are equally good or incommensurable). Bratman claims that brute shuffling in this way undermines your self-

is or gives rise to an intention that persists over time, and which I act on (so long as there is no interference, like a change of mind).

³⁴ Ferrero might be excepted from this charge. But still the kind of autonomous agency at work in his account – namely, the kind that we exercise when we defer to the advice of another – is hardly paradigmatic autonomy, and arguably is less than ideal.

³⁵ Assuming, of course, there is nothing else going on to take away from the significance of this fact.

governance over time. On Bratman's conception, "self-governance essentially involves ... guidance and control by attitudes that help constitute a sufficiently unified point of view, a point of view that constitutes the agent's ... practical standpoint."³⁶ And, Bratman thinks, brutally shuffling from one equal value or incommensurable option to another means that one does not have a unified point of view over time. Thus, conforming to the rational requirement to either change your assessment or follow through with your decision is required for being self-governing over time.

If this argument succeeds then, at least on Bratman's conception of diachronic autonomy, it would be true that diachronic autonomy requires conformity to a rational requirement to follow through with one's past decisions. So, contrary to what we claim, autonomy in the diachronic case would not be the same as autonomy in the synchronic case.

But this argument does not seem to us to succeed – at least, not if we accept that our agency is temporally extended. As Bratman himself agrees, a unified point of view over time must admit of change; a practical standpoint evolves over time and, at least as long as such changes are rational, not manipulated, and not in some other way undermining of the agent's self-governance, they do not threaten the agent's self-governance over time. But "shuffling" in incommensurability and equal value cases seems to clearly count as an unproblematic change in the agent's attitudes over time. If – while I shuffle – each time I form a new intention, I do so on the basis of an enduring "practical perspective" that provides sufficient grounds for it, then all my decisions and actions over this time frame are being guided by the same fundamental point of view. So, this is a clear instance of being guided by a "unified point of view". Thus, even if we adopt Bratman's conception of autonomy, violations of (D) do *not* seem to threaten our autonomy over time.³⁷

³⁶ Bratman, "Time, Rationality, and Self-Governance", p. 77.

³⁷ See Luca Ferrero, "Diachronic Constraints of Practical Rationality," *Philosophical Issues* 22 (2012): 144-164, p. 160 for a similar critique.

Now, if we assume a time-slice conception of agency, such that the self who decided to ϕ and the self who decided to ψ were different agents, then it might indeed be hard to think that there is a unified point of view. But this lack of unification would be a result of the fact that they were two different selves. Two different selves reaching different conclusions from the same “information” (namely, that the options are on par or incommensurable) might not seem to have a unified point of view. If we reject the time-slice conception, there is no difficulty in seeing the point of view from which the agent shuffles from her decision to ϕ to her decision to ψ as a unified one.³⁸

To further illustrate the point, think of a (quasi)-synchronic shuffling case: a case in which there are no time gaps between the shuffles. Suppose you are deciding between two shirts, a red and a blue one, and you reach the conclusion that the two choices are completely on par. You tell the salesperson: “I’ll take the red”. As he is about to pick up the red shirt, you say: “No, the blue!”, then “Actually the red”, and finally “No, definitely the blue”. You are certainly exhibiting indecisiveness, but you do not lack a unified point of view that is guiding and controlling your actions. Your point

³⁸ In his forthcoming “A Planning Agent’s Self-Governance Over Time”, Bratman tries to further back-up his argument that diachronic self-governance requires (D) by arguing that temporally extended planning agency is largely analogous to interpersonal shared agency. When one forms a future-directed intention for one’s future self, this is *analogous* to (even though not the same as) trying to act jointly with another person. To be self-governing over time, then – Bratman argues – one must coordinate with and be responsive to one’s past intentions in the same sorts of ways that agents acting together must coordinate and be responsive to each other’s intentions. It is not clear to us, though, that the analogy really works in the way that Bratman says it does, if we truly accept that your future self who acts is not a different agent from your present self who decides. (For instance, rather than – as Bratman’s analogy seems to require – following through on a future-directed intention to A involving having a new intention to A that interlock with one’s past intention to A, it can simply involve retaining one’s intention to A over time and acting on it. If that’s right, then the analogy does not hold.) But more importantly: even if the analogy does largely hold, it does not seem to us to support a rational requirement against shuffling, provided that we really are – as Bratman intends to – accepting that it is only an analogy, and there is really only one agent. If there is one agent with a single practical perspective that equally supports both ϕ -ing and ψ -ing all along, then while one would, of course, not be following through with one’s intention to ϕ if one shuffled to intending to ψ instead, there is no reason to think that this amounts to a breakdown of self-governance of any kind.

of view is that the two are on par, and that it is, thus, a toss-up which one to go with. There is nothing substantially new introduced if we add time gaps between the shuffles (e.g. if I decide to go with the red shirt, put it on hold and go for lunch, and then switch to the blue shirt when I come back to make the purchase.)

4. Isn't the Diachronic Case Special?

Our claim is that there is no special problem in understanding our autonomy in acting on past decisions: it is just like acting autonomously on decisions we make in the moment. But it is important to be clear that this is a broader and stronger claim than just that the Problem of Diachronic Autonomy, as specified in section 1, is not really a problem. Even if you grant that the Problem of Diachronic Autonomy, as specified, does not itself amount to a real problem, you might propose that there is another important difference between the diachronic and the synchronic case that explains why there is indeed a special problem in understanding how we can be autonomous in the diachronic case. In this section, we discuss three such arguments.

The first such argument appeals to the fact that in the diachronic case there is the opportunity to change one's mind. One can change one's mind over time, but one cannot (or at least should not) be of two minds *at a time*. Doesn't this show that there is a special problem in explaining why we are autonomous when we act on past decisions?

But why would it? If an agent changes her mind about her evaluations or overall preferences, but acts – despite this – on her earlier views, she is *not* acting autonomously. Such a case would be, at best, a case of *akrasia*. Few philosophers argue that an agent in such a situation would be manifesting diachronic autonomy. So, the case in which an agent changes her mind but nonetheless acts on her

past decision cannot be said to pose a special problem for explaining diachronic autonomy, since that wouldn't be a case of diachronic autonomy.³⁹

Perhaps the claim is not that actual cases of changing one's mind pose a problem for understanding diachronic autonomy, but rather that the *possibility* of changing one's mind in the diachronic case poses the problem. That is, even if you do not in fact change your mind, and you do act on your past decision, there is no guarantee that you would have stuck with that decision *had* you deliberated further in the time that elapsed between making the decision and acting. The fact that you could have changed your mind might be thought to pose a problem for understanding how you are autonomous when you simply follow through.

But once we move to merely possible changes of mind, there doesn't seem to be any interesting difference between the diachronic and synchronic case. In the synchronic case it is also often true that *had* you deliberated further you might have come to a different decision. Suppose I deliberate and decide to bake a cake "right away". It is true that further deliberation might have changed my mind and even revealed to me that I ignored some important preferences that I have or that I was too quick to conclude that I prefer my homemade cakes to the store bought ones. But deliberation must come to an end at some point, and the fact that I *could have* continued deliberating, and even the thought that were I to continue to deliberate I would probably end up deciding something else, does not threaten my (synchronic) autonomy when I proceed to bake the cake.

³⁹ Even those who think that agents ought to act on their past resolutions when facing potential judgment or preference shifts, do not think that agents should act contrary to their current judgments. Holton, for instance, argues that an agent should not reconsider her resolution, exactly because she knows that were she to reconsider she would abandon her resolution and, in that case, it would not be rational to act according to the resolution. See Richard Holton, "Rational Resolve", *Philosophical Review* 113:4 (October 2004), 507-535. Bratman considers a view on which rationally resisting temptation involves acting contrary to one's current judgment in "Toxin, Temptation, and the Stability of Intention". But in a later paper, he rejects that understanding of rationally resisting temptation. See Michael Bratman, "Temptation and the Agent's Standpoint" *Inquiry*, special issue on Choice Over Time, Sergio Tenenbaum, ed. 57: 3 (June 2014), 293-310.

The diachronic case is just an instance of this same general phenomenon. Sometimes deliberation ends because we are convinced that no further deliberation would change our mind. But sometimes deliberation ends because we ran out of time, were too tired, didn't think it was worth the effort, etc.⁴⁰ Of course, when making decisions in such circumstances, I might have second thoughts, or change my mind as I start acting on the decision. However, if I do not have second thoughts, and do not change my mind, this is no threat to my autonomy. Acting on the outcome of such deliberations is perfectly autonomous. There is no difference introduced by the diachronic case other than the time lag between the decision and the action. Of course, if we think that the time lag makes it the case that the person who decides and the person who acts are not, in the relevant sense, the same person, then the introduction of the time lag would be significant. But, as we said above, all of those concerned with diachronic autonomy are committed to (and indeed must be committed to) the identity of the agent through time.

So, the first argument fails: the possibility of changing my mind does not pose any special problem for understanding our diachronic autonomy. In neither the diachronic case nor the synchronic case does the fact that further deliberations might have resulted in a different decision threaten the autonomy of acting on the decisions that we do make.

Perhaps, though, the distinctive problem of diachronic autonomy comes from a particular sort of risk that one will change one's mind, namely the risk of temptation. This is the second argument. Suppose I know that tomorrow evening my cousin will call and ask if I want to go to a karaoke bar. However, I promised my aunt that I would have her quilt ready by the day after tomorrow, and if I don't spend all of tomorrow evening sewing, the quilt will not be ready on time. I think that I really should finish the quilt on time, but I suspect that tomorrow I might end up

⁴⁰ And, of course, sometimes we do not even deliberate.

finding a way to rationalize joining my cousin. I thus make a resolution, or at least form an intention, today that I will stay home sewing the quilt tomorrow evening. When the time comes to act on my intention, I recognize that I might be tempted to reconsider my decision and go to the karaoke bar instead. I thus make a special effort not to reopen the question about my evening plans; I keep chanting to myself “quilt for auntie!” so that I’ll not deliberate on this issue. I avoid deliberating exactly because I suspect that my current evaluative standpoint would lead me to make a different decision.

There seems to be a real challenge to our understanding of my autonomy here that finds no parallel in the synchronic case; it seems that we need to understand how acting on an intention that I formed exactly so as not to act on my current evaluative standpoint could be an expression of my autonomy. After all, at least on some conceptions of autonomy, in order to act autonomously, my actions must somehow express my evaluative standpoint or a motive that I (truly) endorse.

But is there really no “synchronic parallel” to such a case? In the karaoke case, I avoid reopening the deliberative question because I am aware that the proximity of a certain kind of enjoyment will warp my practical standpoint. Specifically, I am aware that, were I to deliberate now, my reasoning would be corrupted by the nefarious influence of the propinquity of the satisfaction afforded by karaoke singing. But similar impediments to my deliberative capacities can be present in synchronic temptation cases. Suppose my wedding ring falls into the bottom of a cold and deep pool. I immediately jump into the pool to try to recover it. As I dive into the cold water, I realize that I might be tempted to give up pursuing the ring in order to avoid the pain and discomfort I’m enduring. Fearing this possibility, I avoid thinking about the pros and cons of continuing in my search for the ring, and instead just keep chanting to myself “must get the ring” until I snatch it from the bottom of the pool. This is not a diachronic case of deciding in advance to do something that I suspect I will be tempted not to do later when the time comes. I never formed a prior

intention to save my ring in such a situation. But here, just as in the diachronic temptation case, I avoid deliberating, and so avoid acting according to deliberation from my current practical standpoint, because I fear that doing so would lead me astray.

We can see what the two cases have in common: in certain contexts, I realize that my best judgment is not the one that would be formed by careful deliberation at the time of action, because I realize that some part of my evaluative standpoint has been corrupted. I thus rely on judgments or motives that I form (or have formed) in other ways. In the diachronic case, I rely on a judgment that was made in the past. In the synchronic case, I simply act on my desire to save the ring without allowing myself to consider the merits of the other options. In both cases, I refuse to open deliberation due to my realization that it might lead me astray. In the diachronic case my action relies on prior deliberation, but given the assumption that we are temporally extended agents, why should this make any difference to the question of whether or how my action is autonomous? Whether the best judgment or motive happens to lie in the present or in the past, in both cases I act on an intention that expresses *my* best judgment or motive.⁴¹

Let's turn to the third argument. The third argument says that there is a crucial difference between the diachronic and synchronic case that we have overlooked, namely that it is possible to manipulate your future self. Even though the past self is the same agent as the future self, you can manipulate your future self, and this – the claim goes – shows that there is a special problem in understanding our autonomy when we act on past decisions.

It is true that it is possible to manipulate my future self. There are two ways in which I can do this. First, I can create “coercive” incentives for my future self, with the aim – for instance – of

⁴¹ Of course, our reply to this argument does not depend on this particular description of what makes these action autonomous; it does not even commit us to accepting that these actions are indeed autonomous. All that we need for our purposes is the claim that the synchronic and diachronic cases are perfectly parallel in the relevant respects.

avoiding temptation.⁴² For example, I can make sure that I do not have ice-cream for dessert by making it the case that when the time comes for dessert the “cost” of getting ice-cream will be artificially high: I can give away all the ice cream that I have in my freezer, so that the only way for me to have ice-cream for dessert would be to make a special trip to the grocery store across town. Second, I can try to bind my future self in a way that blocks my future self from having any choice on the matter when the time comes. For instance, I might hypnotize myself now, so that when I later hear the theme song of “Days of Our Lives”, I will walk to my office (because otherwise I would watch the show and get no work done).

But while I can indeed manipulate my future self in these ways, this does not give rise to any special problem in understanding diachronic autonomy. This is because there is nothing distinctive about the possibility of manipulating oneself in the diachronic case: both varieties of diachronic self-manipulation have synchronic counterparts.

Starting with the second kind of case, we can imagine a synchronic version of the hypnosis example. Imagine that *while* the theme songs of “Days of Our Lives” is playing, I hypnotize myself to walk to my office when I hear the theme song of “Days of Our Lives”. I thereby immediately (before the song is over) walk to my office and get to work. There is no temporal distance here between the hypnosis and walking to my office. But I am, nonetheless, clearly manipulating myself into walking to my office.

The first variety of self-manipulation (giving myself coercive incentives) is also possible in the synchronic case. Suppose I am eating the leftovers off my dinner guests’ plates as I clean the dishes. The leftover morsels of food are so tempting, but I also think to myself that I should not be eating little pieces of partially consumed food off other people’s plates. So, as I am driving the fork

⁴² See George Ainslie classic discussion of such cases in his *Breakdown of Will* (Cambridge University Press, 2001).

to my mouth with my right hand, I pour detergent onto the food with my left hand to make it significantly less desirable. Since I know I will not like the taste of detergent, I stop eating the leftovers.

As these examples of synchronic self-manipulation reveal, there is nothing distinctive about the possibility of manipulating yourself in the diachronic case. Technological obstacles make synchronic cases of self-manipulation less common than diachronic cases, but they are not impossible and do sometimes happen. Thus, it is a mistake to think that the possibility of self-manipulation poses a special challenge in understanding cases in which we do act autonomously on our past decisions.

Indeed, on the contrary, reflection on diachronic self-manipulation cases helps illustrate why acting directly on my past decision is a straightforward, paradigm case of acting autonomously. In the self-manipulative ice cream example, when I eat fruit for dessert instead of ice cream, the relationship between my past decision to have fruit for dessert and my current eating of fruit for dessert is deviant. Rather than the past decision directly causing me to eat fruit for dessert when the time comes, the past decision spurred me to set up incentives so that I would, when the time comes, abide by my decision. Indeed, when I eat the fruit for dessert, I am not *acting on* my past intention to eat fruit for dessert at all. Rather, I am acting on a new intention to eat fruit based on my current evaluation of the options. The case involves self-manipulation⁴³ because I interfered with the factors that I knew would come into my evaluation of the options.

Suppose now that I decide this the afternoon to have fruit for dessert this evening; I go about my other business for the rest of the day, and after dinner, I simply *act on* my earlier decision: I

⁴³ If it does. Some might want to deny that this case and the analogous synchronic case involve self-manipulation at all, since in both cases the agent is following her best judgment. This is fully compatible with our view; our claim here is that even if we grant that these cases involve self-manipulation, they do not show a disanalogy between the synchronic and diachronic case.

get out a bowl of strawberries from the fridge and eat it. In this case, the relationship between my past decision and the action that executes it is not deviant. I eat strawberries simply and directly because I decided to, and that's all there is to it. Nothing could be more straightforwardly autonomous than that.

5. Time-Slice Agency and Diachronic Autonomy

Our account of autonomy in acting on past decisions appeals to the fact that our agency is temporally extended; we claim that once we recognize this fact, there is no distinctive problem in understanding our diachronic autonomy. As we mentioned, some people dispute that we are temporally extended agents and subscribe, instead, to a time-slice conception of human agency. We think that the time-slice conception is mistaken, but more relevantly for our purposes in this paper, so do the philosophers concerned with the Problem of Diachronic Autonomy. Not only do they happen to subscribe to a temporally extended picture of our agency, they need to if they are interested in giving an account of future-directed decisions. In this section, though, we want to briefly consider the time-slice conception of agency, and look at a version of the Problem of Diachronic Autonomy that might be thought to arise on such a conception.

Here is the time-slice conception's version of the problem: while there is no such thing as one's past self literally deciding for one's future self, there is something like it that happens all the time. Our past selves very often, somehow or other, seem to control our future selves. Our past selves make what *look like* decisions for our future self (even though they cannot really be doing anything different in kind from advising or commanding). And our future selves often seem to simply follow those dictates, without question. It sounds then, that the life of the average person involves constantly being manipulated by other people (i.e., her past selves) – a terrifying prospect! But at least, in standard cases of my present time-slice following the dictates of my past time-slice, it doesn't *feel as though* I'm being manipulated; rather, I regard myself as acting autonomously. How can

this be? We can call this the “Time-Slice Problem of Diachronic Autonomy”. It isn’t a problem that concerns our diachronic agency, since our agency is not diachronic, on such a conception. But it is a problem that concerns how one time-slice can maintain its autonomy in these diachronic situations.

How to understand a time-slice view in this context is a complicated issue. What we will do here is just briefly explain why we think that the Time-Slice Problem of Diachronic Autonomy is not a distinctive problem for those with a time-slice conception. That is, we will argue that there is no special problem for the time-slice conception in explaining our autonomy in diachronic situations, and rather – much like we argued in the case of the cross-temporal conception – as far as autonomy is concerned, there is *no interesting difference* between diachronic and synchronic cases.

The difference between acting autonomously and not is not a difference in whether you are acting *in accordance with* your preferences, or judgments, etc. One can act in accordance with one’s current preferences without being autonomous at all; one can end up with one’s most preferred outcome due to deviant causes or fortuitous circumstances. Second, I can be fully autonomous, even though I have made a mistake of calculation, overlooked an important aspect of my situation, etc., such that I am not acting in accordance with my current preferences. The difference between acting autonomously and not is, rather, a difference in whether you are being *guided* in a certain sort of way by your own judgments, or evaluations, or some such thing. Autonomy requires that my action is properly guided by my judgments, evaluations, or preferences, or that my action is the outcome of my (possibly implicit) deliberations. But deliberation, or any process or event that will count as “being properly guided”, takes time. So, even in the “synchronic” case, in which I act immediately on a decision that I have just made (with no time gap) autonomy seems to require at least a causal relation between items that exist at different times or extend through time. It is, thus, not clear how a time-slice agent can be autonomous at all unless autonomy is a relation that the time-slice agent bears to its past selves. The process of deliberation or of being guided by one’s judgments,

preferences or evaluations spans more than a moment, and thus must operate through more than just the present time-slice agent. If this is right, then even on the time-slice conception, there is no special problem in understanding our autonomy in the diachronic case. Autonomy, even under that conception, must be by its very nature diachronic.

6. A Concluding Thought

If we reflect on the nature of action, it is not surprising that there is no special problem of diachronic autonomy. Any action takes time to execute; there is no such thing as action that occurs in a mere moment. So, there is no truly “synchronic” case. Because actions always take time, intentions must always operate over time. The intention that I have at the beginning of action must sustain my action throughout its duration. When I walk to the library, my intention needs to persist until I reach the library. While I could reconsider what I am doing part way through, most often I do not. Most commonly, as I make my way through the city blocks to my destination, I do not redeliberate, reconsider my reasons, or even rely on exclusionary reasons not to reopen the question; I simply continue to act on the same intention. The difference between this case and the “diachronic case” that was supposed to pose a special puzzle is only a matter of the length of the time that elapses between forming the intention and (completely) executing it.

Though it certainly does lend some support, noticing that all action is diachronic is not enough to establish that there is no special problem of diachronic autonomy concerning future-directed decisions. Even if all actions take time, it could be that once we are talking about decisions concerning the not-merely-immediate future, the puzzle arises. It could, in other words, be that once the time that elapses crosses some rough threshold, something additional is needed than in the more “synchronic-like” cases to make sense of how we are autonomous in following through with the original intention. Our argument shows that this is not so: it shows that a difference in degree does not make for a difference in kind in this case.

