Time-Slice Rationality and Filling in Plans[*]

Justin Snedegar, University of St Andrews

Forthcoming in *Analysis Reviews*, for a symposium on Brian Hedden's *Reasons Without Persons*.

1 Introduction

In *Reasons Without Persons*, Brian Hedden argues for a view he calls 'Time-Slice Rationality'. This view holds that there are no diachronic norms of rationality; all rational norms are synchronic norms about how an agent ought to be at a given time. The theory is also impartial: the fact that some attitude in the past or future is *yours* cannot, by itself, give it any special significance for determining what your attitudes ought to be now. If an attitude you had in the past, for example, plays a role in determining what your attitudes ought to be now, that role cannot be different in kind from the role played by the attitudes of others in the past (or present, or future).

In Chapter 10, Hedden considers the objection that a theory of rationality should provide norms of reasoning, which are diachronic, and may be personal rather than impartial. After all, reasoning takes place across time, and involves moving from some attitudes that you have to some new attitudes in a rule-governed way. Reasoning is a paradigm of rational activity, so it seems that a theory of rationality should tell us something about how one ought to reason—how one ought to move from some attitudes at one time to new attitudes at a later time.

Hedden argues that a theory of rationality in fact need not provide norms for reasoning. It is enough for the theory to say what states an agent ought to be in at a given time, given facts about that agent's current evidence, preferences, and so on. In Kolodny's (2007) terminology, Hedden thinks a theory of rationality only needs to provide *state* requirements, and not *process* requirements. The theory should tell us what states to be in at a given time, not how to go about getting into them.

What I take to be the central argument here is as follows.[1]

(1) A theory of rationality should be a theory of ideal rationality—it should not take into account our own contingent limitations.
(2) Norms of reasoning would really only apply to limited agents like us. Ideal agents could form beliefs directly in response to the evidence, updating automatically, without going through reasoning processes.
(3) So a theory of rationality need not provide norms of reasoning.

[*] Thanks to Brian Hedden for discussion.

[1] See pp. 182-187. This thought about ideal agents comes in most directly at the end of this discussion, on pp. 186-187, but seems to me to really be the thrust of the argument. Hedden also provides a more concessive response (187) that even if a theory of rationality must provide diachronic norms for reasoning, these norms are derivative on synchronic, impartial norms. I do not have space to evaluate this more concessive response, but note that even if it is "in the spirit of Time-Slice Rationality" (187), it does require giving up the central claim that all rational norms are synchronic and impartial.

Hedden's focus in Chapter 10, as we can see from (2), is on epistemic norms. But as he makes clear throughout the book, the theory is intended to cover both epistemic and practical rationality. I focus on an apparently diachronic, personal norm of reasoning from the practical domain. Importantly, this norm applies even to ideal agents. The norm I will focus on is the norm, roughly, to fill in your plans and intentions. This norm is discussed at length by Bratman (1987), and plays an important role in his planning theory of intention. Bratman's influence on what follows will be obvious.

After arguing that the norm to fill in your plans puts pressure on Time-Slice Rationality, I conclude by sketching and evaluating a way to understand this apparently diachronic norm synchronically.

2 Theories of rationality and norms for reasoning

Hedden notes that reasoning is a *tool*: its central use is to bring about new attitudes (or drop existing attitudes) that, we hope, are rational. It is natural to think that a theory of rationality should provide norms for how to use this tool—how to reason, to get ourselves into rational states. After all, some ways of reasoning seem to be clearly *irrational*, and to push us away from rational states.

Hedden begins his argument against the claim that a theory of rationality must provide norms for reasoning by noting that in fact, most of our beliefs are not formed as a result of reasoning—at least if we conceive of reasoning as a conscious process of moving through premises to a conclusion. Perceptual beliefs, for example, are not the result of reasoning. Many other beliefs seem to be analogous to perceptual beliefs in that they are more naturally thought of as direct, automatic responses to the evidence. Since these beliefs are nevertheless evaluable as rational or irrational, the strong claim that a theory of rationality should *only* be concerned with reasoning is false.[2] Our theory of rationality should give us norms for evaluating beliefs that are *not* the result of reasoning.

This observation provides the thin end of a wedge for Hedden. He claims that it might have been that *all* of our beliefs were the result of automatic, direct responses to the evidence, rather than reasoning. Even if *we* couldn't be that way, it seems possible that there could be some intellectually superior creatures who had no need for reasoning, and could instead respond directly to the evidence. These creatures' beliefs would still presumably be evaluable as rational or irrational, and so a theory of rationality should apply to them, as well.

Hedden combines this claim with the plausible claim that a theory of rationality should be necessary—apply to all possible agents—to draw the conclusion that a theory of rationality need not include norms for reasoning:

> In my view, we should say that the normative theory of rationality includes requirements of rationality which are necessary and hence do not vary depending on an agent's contingent psychological limitations. And this means that it is not a desideratum on a theory of rationality that it issue norms for

[2] Hedden cites Kolodny (2007) as an advocate of this strong position.

> reasoning, since it is only because of our contingent psychological limitations that we resort to reasoning[.] (pp. 186-187)

Since idealized agents are not under norms to reason in any particular way, since they don't need to reason at all, a theory of rationality—as a theory of *ideal* rationality—need not give norms for reasoning.

As an initial objection, even if we accept that the norms of rationality are necessary, in the sense that they do not vary depending on what the agents are like, it does not follow that a theory of rationality need not include norms for reasoning. It might instead be that some of the norms of rationality do not apply to all rational creatures, in the sense that they could not ever be in violation of them or guided by them. One way to make sense of this is to think of the norms as having various restrictions built in. For example, a norm for reasoning could have the form 'If you can't respond to the evidence bearing on p directly, then reason about p like this'.[3] There is an important difference between the claim that the content of the theory of rationality varies between creatures and the claim that different norms—all of which are part of the theory of rationality—apply to different creatures, due to differences between those creatures. Compare the moral case: if members of some society never make promises, then they're never guided by or in violation of a requirement to keep one's promises. On the other hand, members of our society are frequently under this requirement. But it does not follow that the true moral theory for the first society differs from the true moral theory for us.

Setting this aside, the remainder of this paper develops a different objection. This objection holds that there are some norms for practical reasoning that apply—in the more robust sense—even to idealized agents.

3 Ideal agents and planning

Before I develop the argument against Time-Slice Rationality, I should address one preliminary objection. The argument will appeal to norms for practical rationality that involve planning or intending. As we saw in the previous section, Hedden argues that a theory of rationality should be a theory of ideal rationality—it should provide norms for ideal agents that lack our contingent limitations. So a preliminary response to the argument I'll develop is that ideal agents will not need to form plans or intentions, and so norms meant to govern plans or intentions will not apply to them.[4] Showing why this is mistaken will help set up the discussion of the main objection I want to raise.

In the epistemic case, Hedden claims that an ideal agent would be able to evaluate the evidence at a time and respond to it directly, without going through a process of reasoning. The analogous thought here is that the only reason agents like us must form plans or intentions is that we are non-ideal in important respects.

[3] Alternatively, they could be wide-scope; for example, 'You ought to be such that either you don't need to reason about p or you reason about p like this'. Such norms *would* apply to all agents, but ideal agents would always satisfy them, since they would always satisfy the first disjunct.

[4] Perhaps idealized agents need to form what is often called intentions *in acting*. What this objection suggests is that they nevertheless do not need to form intentions or plans for the future.

As Bratman (1987) emphasizes, central function of plans and intentions is to let us decide now what to do later, so that we do not have to think about it later. This is useful since when the time for action comes, we may not have time to go through the potentially complicated task of weighing up the reasons for and against our options to come to a decision about what to do in the moment. The ability to plan for the future lets us distribute these costs—both in terms of time and cognitive resources—across time, and to think about what to do when we have the opportunity to do so, rather than at the moment at which a decision is needed.

The thought, then, is that we only need something to play this kind of role because we are non-ideal. If we were just smarter and could think more quickly, we wouldn't need to distribute the cognitive costs of thinking about what to do, since we would have no trouble coming to an immediate decision about what to do. Just as idealized agents can form rational beliefs directly in response to the evidence, idealized agents can make rational decisions about what to do directly in response to the reasons.

But there is an important disanalogy between theoretical and practical rationality, which has to do with the fact that the latter is practical. Our intentions and plans have to be carried out in the world. Whereas the rational belief is simply the belief that's best supported by the evidence, or (epistemic) reasons, there are other rational requirements bearing on intentions and plans. In particular, there are requirements springing from what Bratman (1987) calls the *coordinating* role of intention: intentions and plans let us coordinate both our own actions across time and our actions with those of others with whom we need to cooperate. For example, our intentions must be consistent with one another, otherwise we cannot carry them out. Thus, even if options A and B are equally well supported by the reasons, it is irrational to intend both, if they're incompatible. As the case of Buridan's ass illustrates, it can also be irrational to simply not make up your mind between these two equally desirable options. For another example, our intentions must have a certain degree of stability across time, in order to play their coordinating role.[5]

The most important upshot of this coordinating role for the current discussion is that there is a point to forming intentions, even if you're an ideal agent and could immediately see which action the reasons most strongly supported at every decision point. If you always waited until the last second to make your decision about what to do, coordination, especially with others, would be impossible. So practical rationality is about more than simply deciding to do what's most strongly supported by the reasons at every decision point; even ideal agents may be rationally required to form intentions and plans for the future, to allow for coordination.

[5] The requirement that our intentions be relatively stable across time already suggests that some requirements of practical rationality will be diachronic. Hedden discusses something close to the stability requirement on intention, but in connection with preferences (see the *Career Decisions* case on p. 3 and discussion of Utility Conditionalization in section 8.2). There are important differences between the stability requirement on intention and alleged stability requirements on preferences that I believe make the former more likely to be irreducibly diachronic (as well as making them more plausible). But I will leave this issue aside, in part because formulating a suitable stability requirement is tricky, since the requirement can be, at best, that intentions be *mostly* stable—changing your mind cannot *always* be irrational, after all. See Bratman (1987), Chapter 5.

4 Filling in plans

4.1 A diachronic norm

As Bratman (1987) observes, our plans and intentions are typically *partial* in a certain sense. They do not specify every detail of the action to be carried out; they are relatively *coarse-grained*.[6] This fact generates a norm to *fill in* these partial plans. This is because plans that are too partial or coarse-grained cannot lead to or guide action in the way that plans are supposed to. One obvious way in which we need to fill in our plans is by selecting *means* for carrying them out. But the norm is more general. We often need to decide on a *way* of carrying out a plan, where this is not naturally described as selecting a means.

Suppose my partner asks what I want to do for dinner tomorrow night. I form the plan to make dinner at home. If dinnertime tomorrow rolls around and I haven't revisited this plan and filled it in, then I will be in some trouble. I can't actually carry the plan out without filling it in. I need to decide at least *what* to make and *when* to make it (note that neither of these questions is naturally understood as asking for a *means* to making dinner). Since the primary function of plans is to get us to act, it is plausibly a norm that we fill in the partial plans that we form. It is irrational to fail to fill in your partial plans before the time for action.

There are a few important qualifications. First, we don't need to fill our plans in "all the way", to a maximally fine-grained level of detail (if there is such a level). For example, we can plan to get on the bus but don't need to fill this plan in by planning to get on the bus right-foot-first. I will say that the plan must be *sufficiently* filled in, leaving it open exactly what this amounts to. Second, different people need to fill in their plans to different degrees. Maybe someone with an ankle injury does need to fill in their plan to get on the bus by deciding which foot to use first. Finally, we can drop our partial plan instead of filling it in. If I change my mind about making dinner at home, then I'm obviously under no rational pressure to decide *what* to make. So the norm directing us to fill in our plans will be somewhat complicated. I will largely ignore these complications here, though, because they do not bear on whether the norm is synchronic or diachronic, or whether it's personal or impartial.

Keeping those qualifications in mind, but setting them aside for now, here is a simple version of the norm that highlights the diachronic, personal nature of the norm.

> **Fill In:** If at t1 you have a partial plan to A at t2, then by t2 you ought to have sufficiently filled in that plan.

[6] From now on I will talk mainly of plans, because there is another use of 'partial intention' in the literature from Holton (2008) that is very different from the sense I intend here. What I say about plans, I believe, can be carried over to intentions pretty straightforwardly, even though plans are not exactly identical to intentions, at least for Bratman (1987) (intentions are the "building blocks" of plans, while plans are "intentions writ large", p. 8). The examples of intentions/plans I use will not involve long-term plans or projects, which may be connoted by 'plan'. If you like, replace my talk of plans with talk of intentions, keeping in mind that 'partial intentions', in that case, are very different from what Holton (2008) has in mind.

Note that this principle is explicitly diachronic: it refers to the state you are in at one time, t1, and says that given that you are in that state at t1, you ought to be in a certain state at a different time, t2. It is also explicitly personal, since it refers to plans that *you* have at different times. So if a principle like this is a true principle of practical rationality, then there is at least one diachronic, personal principle of practical rationality, and Time-Slice Rationality is false.

4.2 Options as decisions

So far, this discussion of partial plans has been put in terms that Hedden may not accept. In Chapter 6, Hedden defends a somewhat surprising view about what an agent's options are, at least for the purposes of a theory of rationality: rather than being actions that the agent can perform, options are instead decisions that an agent can make. More precisely, options are propositions about decisions to perform an action, rather than the actions themselves, as I have been assuming.

On this view, a set of options is then a set of mutually exclusive and exhaustive propositions of the form *S decides at t to A*. These sets must be maximal, in the following sense: "there is no proposition of the form *S decides at t to A* which is not a member of the set but which is incompatible with each member of the set" (106). Though the sets of options must be maximal in this sense, it does not follow that the options themselves must be "maximally specific", in the sense that the content of the decision, A, must be a "maximally specific" action. The decisions may be relatively coarse-grained.

So though I have not put things in Hedden's preferred terms, treating options as decisions, the crucial point that options can be more or less partial, or coarse-grained, is compatible with his view. Filling in a partial plan to make dinner, on Hedden's view, would be regimented as follows: I first select an option like *I decide at t1 to make dinner*, and then fill this in by selecting an option like *I decide at t2 to make pizza for dinner*. This does not affect the objection to Time-Slice Rationality. For the remainder of this section, then, I will continue to write in (what I take to be) the more natural way, as if options were the actions themselves. But as we'll see in the next section, the view that options are decisions may give Hedden a way to reply to the objection I've presented.

4.3 Reply: Ideal agents and partial plans

In the previous section I argued that even ideal agents ought to form plans. But we may still doubt whether ideal agents ought to form *partial* plans. Perhaps ideal agents could form maximally (or at least sufficiently) fine-grained plans from the start, and thus would not be under any rational norm to fill in their plans.

In fact, according to Bratman (1987), one of the main reasons we form partial plans, rather than maximally specific ones, has to do with our cognitive limitations. Settling on maximally specific plans (or even plans that are sufficiently specific for action) takes lots of cognitive resources. Since it is our forming partial plans in the first place that makes us subject to the (alleged) norm to fill them in, Hedden may simply reject the claim that a satisfactory theory of (ideal) rationality will include

anything like **Fill In**. Ideal agents, who lack our cognitive limitations, have no need for partial plans, and so do not need to fill them in.

But there are reasons to form partial plans not deriving from our contingent limitations. One such reason, also discussed by Bratman and mentioned earlier, has to do with the importance of *stability* of plans. Since both the agent herself and those with whom she must coordinate need to use her plans as fixed points in reasoning about what to do, it is important that plans are stable, and not constantly being reconsidered or abandoned. Highly specific plans are more likely to need to be abandoned in the face of new information. For example, a plan to *drive the car to the store at 8:35am, taking the scenic route* is more likely to need to be abandoned in the face of new information (about the hours of the store, the route, the availability of the car, etc.) than a plan to *go to the store*. Partial plans are more flexible, and therefore more stable. Relatedly, partial plans are more likely to be compatible with our other plans and with the plans of others with whom we need to coordinate.

One potential objection here is that ideal agents would not have this reason to adopt partial plans, since it is based partly in the fact that we don't always have full information, about the world or about the plans of others, which is one of our contingent limitations. But constructing a theory of rationality for omniscient agents is surely idealizing too much.[7] Moreover, Hedden discusses rational norms about proportioning one's beliefs to the evidence, so it seems that he does not assume that ideal agents are omniscient.

It is worth noting here that planning theorists in the field of artificial intelligence have recognized these sorts of advantages of forming partial plans. They have thus developed planning systems that lack many of our contingent limitations but still form partial plans; this is often called "least commitment planning", highlighting the advantages that come from the flexibility of partial plans. Since these systems form partial plans, they must eventually fill them in. Of course, even these computer systems are not ideal agents, since they will have some processing limitations, but it is telling that even agents that are more ideal than us in important respects have a need for partial plans.[8]

So even ideal agents should sometimes adopt partial plans. This means that even ideal agents will be under rational norms to fill in those plans, since any agent that adopts partial plans is under such a norm. Crucially, the norm to fill in your plans is a diachronic and personal one, since it refers to your plans at one time, and how they put rational constraints on what you ought to do at a later time. Thus, even a theory of ideal rationality should include at least one diachronic, personal norm for reasoning, contrary to Time-Slice Rationality.

5 A synchronic norm?

In this section I will consider a way to defend Time-Slice Rationality by understanding the norm to **Fill In** your plans synchronically. The strategy is to

[7] It is common to assume that ideal agents are *logically* omniscient, but full omniscience of the kind that would be required to offer the current reply goes far beyond that.
[8] For a sampling, see Weld (1994), Sacerdoti (1977), and Bratman, Israel, and Pollack (1988).

appeal to two synchronic norms, one meant to capture the rationality of adopting partial plans and one meant to capture the rational pressure to have sufficiently filled in plans. I'll argue that this at least gives us the basis for a synchronic understanding of partial planning, and point out some notable features of the sort of deliberation it posits.

Here is a simple, plausible synchronic norm meant to explain the rational pressure to have sufficiently filled in plans.

> **Synchronic Filled In:** At t either do not plan to A at t or have a plan to A at t that is sufficiently filled in to guide you in A-ing at t.

When the time comes to A, if you plan to A, then your plan ought to be sufficiently filled in. This norm explains why it's irrational to leave your plans partial. But we have not yet explained why agents are sometimes rationally required to adopt partial plans in the first place.

Here is a very general and plausible norm (or at least a simplified version):

> **Planning:** At t, adopt the plan that is best supported by reasons.

For Hedden, options are decisions, so adopting a plan would be choosing an option of the form, *S decides at t to A*. As we saw in section 4.2, these options can be more or less fine-grained (e.g., deciding to make dinner tomorrow vs. deciding to make pepperoni pizza for dinner tomorrow at 7pm). We can exploit this to show why this simple **Planning** norm can capture the rationality of adopting partial plans. The crucial claim is that sometimes, a partial plan—or relatively coarse-grained decision—is the one that is best supported by reasons. It is sometimes assumed (at least implicitly) that the only reasons that could bear on the rationality of a decision to A are the reasons that highlight good or bad features of A or its alternatives. But one important thing that Bratman's (1987) discussion of the coordinating role of intentions and plans shows is that this is a mistake: as we've seen, there may be reasons that bear on the rationality of deciding to A that have to do not with the qualities of A itself, but rather with the good or bad features of deciding to A.[9] So in many cases, there will be coordination-based reasons in favor of adopting partial plans, rather than more detailed ones.

If this is correct, then this **Planning** norm—which again is a highly general synchronic norm—explains why it is sometimes rational to adopt a partial plan. The **Synchronic Filled In** norm above explains why it is irrational to leave these plans partial up through the time of action. Notice that it does not say anything about reasoning: there are no norms about how to get from the state required by **Planning** into the state required by **Synchronic Filled In**. This is very much in the spirit of Time-Slice Rationality.

For this proposal to work, it must be that sometimes, the agent's reasons best support a partial plan. So the agent's options must sometimes be relatively coarse-

[9] Many people would classify these as the 'wrong kind of reasons' for deciding to A, since they are 'state-given' rather than 'object-given'; see Parfit (2001), Piller (2001), and Hieronymi (2005), for example. But I agree with Schroeder (2012) that Bratman's observations about the coordinating role of intention (or decision) show that this is a mistake.

grained, or partial. As we saw before, Hedden's favored theory of options allows for this. He holds that a set of options for an agent at a time consists in all and only the *decisions* that the agent is able to make at that time. Importantly, agents are generally able to make both coarse- and fine-grained decisions. So, for example, a set of options might include all three of *S decides to make dinner*, *S decides to make pizza for dinner*, and *S decides to make pepperoni pizza for dinner at 6:30pm*.

One initially puzzling thing is that for all of these to be options in the set, they must be mutually exclusive—both because Hedden insists that options are mutually exclusive, and because, intuitively, it is hard to see how an agent could face a choice between compatible options. It is perhaps even harder to see how an agent could face a choice between deciding to do A and deciding to do B when doing A *constitutes* doing B, as in the pizza example above. In fact Hedden wants to allow for just this. He says that on his intended reading, *S decides to A* is incompatible with *S decides to A&B*; we can attempt to secure this reading by inserting 'only': *S only decides to A* and *S only decides to A&B* (106).

If *S decides to make dinner* and *S decides to make pizza for dinner* can each be options, then the proposal I've suggested in this section seems defensible. Given this set of options, it may be that the more coarse-grained decision is most strongly supported by the reasons. Crucially, the coordination-based reasons favoring *S (only) decides to make dinner* may *not* favor *S (only) decides to make pizza for dinner*, since the former is more flexible than the latter.

Such an approach to understanding partial plans and the corresponding process of filling them in does miss out on explaining some of the advantages emphasized by Bratman (1987). One of the most important advantages for agents like us is that forming partial plans lets us put off thinking about and evaluating more detailed ways of carrying them out until we have more time. If the coarse-grained decision and the many fine-grained ways of carrying it out are all options, then in order to determine that what we have most reason to do is make the coarse-grained decision, we will have to consider all of these more fine-grained decisions, as well. But this advantage is due to our contingent cognitive limitations. Ideal agents could consider even highly detailed plans, but could then nevertheless opt for more partial ones, to secure the coordination-based advantages.

This approach relies crucially on treating options as decisions. This is because, though it is at least plausible that deciding (only) to make dinner and deciding (only) to make pizza for dinner are mutually exclusive, it is not at all plausible that the actions themselves—making dinner and making pizza for dinner—are mutually exclusive. This brings out an important point: the kind of deliberation—where 'deliberation' just means weighing up the reasons for and against the options, which need not involve reasoning—that is at work in this kind of process, at least for ideal agents, is in an important way a kind of higher-order deliberation *with* or *about* decisions or plans.

For Hedden, all deliberation is most directly like this, since options are decisions.[10] The reasons we weigh up are reasons that support decisions themselves,

[10] This assumes that what reasons bear on are our options. It may be possible to defend a more indirect account here, such that reasons bear on the actions—the content of the decisions—most fundamentally, and only bear on the options (decisions) themselves indirectly.

rather than actions. This strikes me as a surprising, if not problematic, description of ordinary deliberation about what to do, which at least intuitively involves weighing up first-order reasons, bearing on the actions themselves. But the deliberation described above is very naturally described in this higher-order way. Put another way, in ordinary deliberation the reasons we weigh up seem to bear on the actions we can perform, while the deliberation involved in partial planning, for ideal agents, at least some of the most relevant reasons seem to bear most directly on the plans or decisions themselves, and do not (necessarily) say anything in favor or against the actions which form the content of the plans.

This illustrates a divergence between the deliberation of ideal agents and the deliberation of agents like us, when it comes to partial planning. We are of course capable of engaging in higher-order deliberation, considering reasons bearing most directly on the attitudes we may adopt. For example, we may sometimes consider reasons to *withhold* intention, and in fact may withhold intention—not make a decision—for those reasons.[11] But this is not what ordinarily happens when we engage in partial planning. Typically, we consider reasons for and against the relatively coarse-grained actions themselves, and not the coordination-based reasons to adopt partial, rather than more detailed, plans. So ordinarily, the coordination-based advantages of engaging in partial planning do not enter into our deliberations. They are more like the advantages of being able to employ heuristics in reasoning—not something we think about explicitly, but which may enter into, for example, an evolutionary explanation of why we reason in a certain way. For ideal agents engaged in the sort of deliberation posited by the synchronic approach developed in this section, though, these coordination-based advantages do function as reasons favoring forming partial, rather than more detailed, plans.

The kind of approach I've suggested here provides the basis for a synchronic account of the rationality of the apparently diachronic phenomenon of filling in our partial plans, given Hedden's view that options are decisions. But it also brings out an instance of a general problem facing theories of ideal rationality like Hedden's, namely that it may give a picture of rational agency that is not very useful or familiar for agents like us.

Bibliography

Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Cambridge: Harvard University Press.

Bratman, M., Israel, D., and Pollack, M. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4: 349-355.

Hedden, B. (2015). *Reasons Without Persons: Rationality, Identity, and Time*. Oxford: Oxford University Press.

[11] See Schroeder (2012).

Hieronymi, P. (2005). The wrong kind of reason. *Journal of Philosophy*, 102(9): 437-457.

Holton, R. (2008). Partial belief, partial intention. *Mind*, 117: 27-58.

Kolodny, N. (2007). State or process requirements? *Mind*, 116: 371-385.

Piller, C. (2001). Normative practical reasoning. *Proceedings of the Aristotelian Society*, 25, suppl., 195-216.

Parfit, D. (2001). Rationality and reasons, in *Explaining Practical Philosophy*, eds. Egonsson, Jesefsson, Petersson, and Rønnow-Rasmussen. Burlington: Ashgate, 17-39.

Sacerdoti, E. (1977). *A Structure for Plans and Behavior*. New York: Elsevier.

Schroeder, M. (2012). The ubiquity of state-given reasons. *Ethics*, 122(3): 457-488.

Weld, D. (1994). An introduction to least commitment planning. *AI Magazine*, 15: 27-62.