

Deliberation and Pragmatic Belief

Brad Armendt

Arizona State University

June 13, 2018

How do our beliefs interact with our practical interests? Does the practical significance we attach to the truth or falsehood of p sometimes influence how strongly we believe it? If so, when are we more likely to notice the influence, and what can we say about how it happens?

When we deliberate about what to do, our beliefs and our interests interact. Suppose we evaluate our options in light of our expectations about their possible outcomes. Then our beliefs about those outcomes influence our evaluations, and so influence our choice and action.

In the other direction, our comparisons of the competing options influence our beliefs about what we will do, and they sometimes influence our beliefs about what our actions may accomplish. Sometimes a decision does not come quickly, and deliberation takes time. The interactions among all these influences may then take a deliberator through a succession of shifting evaluations and opinions. The theory of *deliberation dynamics* provides ways to model that.

Let us consider the questions raised at the beginning, as they apply to contexts of deliberation, when deliberation stretches over a period of time. The answer to the first question has many parts. Here I mean to pursue one of them, suggested by the second question, namely the possibility that beliefs are sometimes *stake-sensitive*. Might the strength of a deliberator's state

of belief vary, depending upon how it matters, upon what she takes to be at stake on the truth of the belief? I have elsewhere discussed at some length the idea that a belief state might be stake-sensitive in this way; here I will be more concise.¹

Stake-sensitive belief. You are in a conversation with a student, who asks you: in which part of his paper on truth and probability did Ramsey give that example about the unwholesome yellow toadstools? After a moment, you reply: in the last section of the paper, on the 'logic of truth.' Soon after the student leaves your office, the phone rings. A local radio station invites you to play their quiz game. The prize for a correct answer to their question is valuable, a fabulous overseas trip, and they offer you several subjects to choose among. You pick 'Probabilism,' and to your amazement, they ask the very same question. Now you think a little longer; it seems like it was the last section of the paper, but could it have been earlier when Ramsey was talking about beliefs and frequencies?

What, if anything, might be different for you in the two situations? What, if anything, is different about your belief that the yellow toadstool example is in the last section of the paper? The story is not perfect, but it is intended to bring us to the neighborhood of our topic.² To what extent do our beliefs, and how strongly we hold them, depend upon how they matter to us, on what we take to be at stake on them? Might we move from one context to another, remaining in the same doxastic state concerning p , yet believing that p more strongly in the one context than in the other? In order for that to happen, a doxastic state, a belief state, must have a certain sort of complexity, a context-sensitivity that yields, in the presence of one set of stakes, a belief of

one strength, and in the presence of different stakes, a belief of a different strength. So the question being asked is about the nature of belief states, as we understand them, or as we think they should be modeled in a theory about them.

If a belief that p is stake-sensitive, then in the absence of contextual stakes, its strength is not definite.³ It is a persistent doxastic state, let us assume, yet it exerts patterns of influence on choice and inference that vary, when what is understood to be at stake on p varies. Given its persistence, over time we might notice the belief's stake-sensitivity when there is a shift in the pragmatic significance of p . But a change in what is at stake could affect a belief in two ways, only one of which is what I have in mind. My question is whether the strength of a *given* state of belief is sensitive to what is at stake, not on whether *news* of what is at stake leads to *belief change*, from a state of belief in one context, to a different state of belief in another context. No one doubts that our beliefs sometimes change through learning or forgetting; it is less clear that they are sometimes stake-sensitive.

The question is not confined to beliefs that are rational; I am interested in stake-sensitive belief, whether rational or not. For what it's worth, and without argument here, I tend to think that sometimes beliefs are stake-sensitive (in the peculiar way I have described), while also thinking that the ways that may happen are often not rational.

Contexts of deliberation, however, might provide occasions where we find rational stake-sensitive beliefs. Periods of extended, perhaps intermittent, deliberation are in some ways

special contexts, to be sure, and my beliefs about what my current deliberation will lead me to do are in some ways special beliefs. But if we seek examples of rational stake-sensitive belief, then contexts of deliberation, special though they may be, are a natural place to look.

A belief is stake-sensitive when the strength with which you hold it varies in response to variation in how it matters to you, variation in what you take to be at stake on it. One sort of variation is variation in the *size* of the stakes. Sensitivity to the size of the stakes is often raised in pragmatic encroachment discussions, when we consider a believer or knower moving between low-stakes and high-stakes contexts, for example. In previous work on belief I highlighted sensitivity to the sizes of the stakes, as in the above story of the radio quiz game. Another sort of variation is variation in the *odds* of the stakes, or in the *shape* of the stakes. These variations may occur together. By variation in the odds, I mean variation in the comparative value the believer sees in the truth and in the falsehood of the belief. This might be taken to be the *ratio* of the perceived value of the belief's truth to the perceived value of its falsehood, on some particular value scale. Imagine that in one context you regard p as much better than $\sim p$, but in another context you do not, and in another $\sim p$ is better than p . If that is the source of shifts in how strongly you believe that p in each context, then your belief is stake-sensitive.⁴

The question of whether beliefs are stake-sensitive can be raised for full beliefs, for categorical beliefs, and for degrees of belief. Suppose we are talking about action-guiding degrees of belief; a stake-sensitive belief in p could be characterized by a *set* of fair betting quotients, rather than a single one. Which of the fair betting quotients accurately measures the belief could depend on

the context, and in particular, on what you then take to be at stake on p . In other words, a stake-sensitive degree of belief in p could have a strength that depends on context, and in particular, on the stakes that you take to actually be in play, as given by valuations you currently attribute to p and to $\sim p$. The current valuations reflect the part that p plays in your current pragmatic interests; the degrees of belief reflect your willingness, in various contexts, to take on new interests and commitments to which p is relevant.⁵ In the discussion that follows, let us take beliefs to be degrees of belief; I will assume that rational degrees of belief are subjective probabilities.

If p is one of many members of a partition, it is natural to think of the shape of the stakes on p as comparisons of the value of p to the values of the other members. If you are deliberating among many different options, the shape of the stakes on A involves the comparison of its value to the values of the other options B, C, \dots ⁶ Consider, then, in the midst of my deliberation about whether to do A or do B , my present belief that I will do A . This is a belief about the future, perhaps the near future or perhaps a more distant one.⁷ The value I attach to the truth of the belief that I (will) do A is the value I attach to doing A . Similarly, for the rival option B . The shape of what is at stake on the truth of A involves the pattern of values I attach to A and to its rivals. Variation in the shape of the stakes involves variation in the pattern, variation in the relative values I attach to the options.

Deliberation and instability. In the context of extended deliberation, beliefs and evaluations may interact. In some deliberations the interactions are complex, and the process of deliberation

delivers contextual changes all by itself; this is particularly so when the decision problem has a form that leads to *decision instability*.

Decision instability can arise in a variety of settings.⁸ Allan Gibbard and William Harper illustrated it with the example of the man who met death in Damascus:

Consider the story of the man who met death in Damascus. Death looked surprised, but then recovered his ghastly composure and said, 'I am coming for you tomorrow'. The terrified man that night bought a camel and rode to Aleppo. The next day, death knocked on the door of the room where he was hiding and said, 'I have come for you'.

'But I thought you would be looking for me in Damascus,' said the man.

'Not at all,' said death "that is why I was surprised to see you yesterday. I knew that today I was to find you in Aleppo."

Now suppose the man knows the following. Death works from an appointment book which states the time and place; a person dies if and only if the book correctly states in what city he will be at the stated time. The book is made up weeks in advance on the basis of highly reliable predictions. An appointment on the next day has been inscribed for him. Suppose, on this basis, the man would take his being in Damascus the next day as strong evidence that his appointment with death is in Damascus, and would take his being in Aleppo the next day as strong evidence that his appointment is in Aleppo.⁹

To find his best option by using causal decision theory, as I will assume he should, the man can use the possible states 'Damascus is inscribed' and 'Aleppo is inscribed.' He takes those states to

be outside his causal influence, and they are sufficiently specific, given his present interests, to form a K -partition. Call them K_D and K_A , respectively.

If, for example, he believes the appointment book states Damascus rather than Aleppo, and his belief $pr(K_D)$ is greater than $\frac{1}{2}$, then causal decision theory endorses going to Aleppo.¹⁰ But that decision seems unstable: Since he takes death's appointment book to be based on reliable predictions, when the man comes to believe he is about to go to Aleppo, he also comes to believe that Aleppo is inscribed after all, so $pr_n(A)$ and $pr_n(K_A)$ are greater than $\frac{1}{2}$, and $pr_n(K_D)$ is less, which makes Damascus the better option. But then the man's preference for Damascus again influences his beliefs; he will probably go there, and Damascus is probably inscribed, $pr_{nn}(D)$ and $pr_{nn}(K_D)$ exceed $\frac{1}{2}$, which makes Aleppo the better option; ... and so on.

Deliberation dynamics. What should the man do?¹¹ The sense of instability in problems like *Death in Damascus* arises when he can reevaluate his options in light of the course of his deliberations. This is a good setting for the theory of deliberation dynamics, where your comparisons of your options influence your beliefs about whether you will do them, leading to new evaluations and comparisons of your options, which in turn further influence your beliefs. The theory can be applied to deliberations about many sorts of decision problems, simple and complex.¹²

Deliberation dynamics applies to deliberation that takes place over time. At each moment, your beliefs about states of the world (*e.g.* Damascus is inscribed) underwrite your current

assessments of your options. Those assessments conform, let us suppose, to subjective rational decision theory—throughout this discussion, to causal decision theory. If we assume that the values you attach to outcomes (*e.g.* life, death) are not shifting, then when your beliefs are stable, so are your assessments of your options. But as in the case of the man who met death, your beliefs may not be stable. At a given moment, your current assessments may influence your beliefs about what you will do, as well as your beliefs about states of the world that matter to the outcome of doing it. Your newly current beliefs then underwrite new assessments, and we can entertain the trajectories of your shifting beliefs and assessments over time. Sometimes those trajectories may display oscillations, as we imagined in the case of the man who met death.¹³

The trajectory of your beliefs will depend on the details of your dynamics. How do your beliefs about what you will do depend on the values you attribute to each option?¹⁴ At time t , you regard one action A as better than its alternative B . (Your $U_t(A)$ is greater than $U_t(B)$ on some specific U_t scale that represents your preferences at t .) How does that influence your belief that you will do A , your $pr_{t+}(A)$, at t_+ ? The answers to such questions throughout your deliberation might be given by a dynamical rule. Many such rules are possible; among them are rules that *seek the good*, according to which your probabilities that you will perform actions rise exactly when you currently see those actions as better than their alternatives, or more precisely, as better than the *status quo*, which is your current expectation of the outcome of the problem you are deliberating about.¹⁵ If at some moment t your beliefs lead you to regard your options as equally good, then your assessments give you no basis, under dynamics that seek the good, for

altering those beliefs and assessments. You are at an *equilibrium* of the dynamics. Since you then regard each action as equally good, to choose one or the other is to break a tie, or ‘pick’ among the tied options.

Returning to Death in Damascus, then, suppose that during his deliberation, the man’s evaluations of his options influence his beliefs about what he will soon do, and about what is inscribed in death’s appointment book. Suppose that at time t_1 during his deliberation, he regards going to Aleppo as the better action, so that $U_{t_1}(A) > U_{t_1}(D)$; this raises his belief that he will go to Aleppo, $pr_{t_2}(A) > pr_{t_2}(D)$, and also that death will be waiting for him there, $pr_{t_2}(K_A) > pr_{t_2}(K_D)$.¹⁶ Then, when he reevaluates his options at t_2 with those beliefs, he sees Damascus as the better action, $U_{t_2}(D) > U_{t_2}(A)$, which influences his beliefs again. Under plausible assumptions, in the Death in Damascus problem, deliberation that seeks the good will eventually lead the man’s beliefs to a stable equilibrium, where he sees neither act as better than the other.¹⁷ At that point, his tied evaluations give no basis for further adjustments in the beliefs that underlie them. At the equilibrium state in the original Death in Damascus problem, $pr_{eq}(K_A)$ and $pr_{eq}(K_D)$ are both $\frac{1}{2}$, as are $pr_{eq}(A)$ and $pr_{eq}(D)$. His action will be the outcome of some way of dealing with the tie. So, one rational solution to Death in Damascus is to break a tie between A and D ; just pick one of the equally (un)attractive options. Our present focus, however, is on the beliefs at work in rational deliberation, rather than solutions to the problem.

A general feature of equilibrium states, whether your deliberation leads you to them or not, is that you see your available options as equally choice-worthy, as having equal expected utility. It

may also happen that you then believe that you are as likely to do one act as the other, but that need not be so in problems that lack the symmetry of Death in Damascus.¹⁸

Why should the man embark on this deliberative journey? There is at least this reason: a rational choice should be based on all of your relevant beliefs at the time you make it. So, *if* you believe at time t that death is more likely to go to Aleppo than to Damascus, $pr_t(K_A) > pr_t(K_D)$, your evaluations at t of your options, $U_t(A)$ and $U_t(D)$, must use those beliefs. Or, to put it another way, a rational decision theory such as causal decision theory is properly used only when those evaluations do so. What about the beliefs? What sorts of influences on beliefs, what sort of adjustments to beliefs, have we been talking about?

Regarding A as better, and believing you will do it. Consider the man's shifting belief that he will go to Aleppo, $pr_t(A)$, as he deliberates about what to do. At one point, when going to Aleppo looks better, he thinks he will go there; at another point, when it looks worse, he thinks he will not. His state of belief about going to Aleppo appears to be sensitive to the changing shape of the stakes that A carries for him.

Or is it? Is he instead learning new beliefs, stimulated by changes, say, in his belief that ' A is better than D '? If learning experiences are revisions that replace old belief states with new ones, then as we noticed earlier, they give no reason to think the states are stake-sensitive. To explore this, we could turn our attention to the rational acquisition and updating of his belief that A is better than D , and others like it. Might *that* belief, if he has it, be sensitive to what is at stake on

'A is better than D'? Maybe, but this is getting a bit intricate, and there are other paths of influence to consider.¹⁹

A deliberator assesses and compares his options; does rationality require that he also maintain up-to-date beliefs about how he assesses them, about how much better he finds one than another? Suppose instead that the shifts in $pr_i(A)$ are directly stimulated by his higher or lower regard for A, rather than by what he believes about that regard. How would his assessments, the shape of the stakes on A, influence his belief? Setting aside learning *via* other beliefs, two possibilities come to mind.

One possibility is that he has *experience* of his comparative regard for his options, and the experience provides grounds for a change of belief about what he will do. Perhaps his current comparative assessments, including that A looks better than D, are sources of experience, akin to perceptual experiences that alter his beliefs in the elements of a Jeffrey-shift (looks like rain), or akin to perceptual experiences endowing him with certainties on which to conditionalize (left knee hurts). If his new belief state replaces his old one, this amounts to a possible way of learning from experience, rather than stake-sensitive belief. That is one possibility. The second is the one we have been pursuing: As his current regard for his options alters the shape of the stakes on A, his stake-sensitive belief that he will go to Aleppo shifts too. The influence is direct, and not mediated by what he believes about his regard for going to Aleppo vs. staying in Damascus, nor by an experience of that regard, other than having it.

Either of the two possibilities is compatible with deliberation dynamics, a theory about the rational interaction between your beliefs about what you will do and your assessments of the value of doing it. It takes the influence of your assessments on your beliefs to be systematic, for you, and characterizable by dynamical rules, but it is otherwise silent about the nature of that influence. A dynamical rule is a function from your assessments—of *A*, for example—to your beliefs about whether you (will) do *A*. The inputs to the rule are your assessments, not your *beliefs* about your assessments, if you have them. The outputs are your beliefs about what you do.

So it is compatible with deliberation dynamics that the man's belief that he will go to Aleppo is stake-sensitive, and given a reasonable dynamical rule for making such shifts—a rule that seeks the good—there need be nothing irrational about the dependence of that belief on the value of its truth. Here, then, is a scenario for a rational stake-sensitive belief: In extended deliberation, my belief that I will do *A* sometimes reasonably depends on the shape of what is at stake on its truth.

During our deliberations we sometimes entertain beliefs about, and often have experiences of, assessing and comparing. But how often? To the extent that we suspect that deliberation may at times proceed without them, we leave room for the possibility of rationally stake-sensitive belief. We can investigate whether it is a rational requirement that, throughout your deliberation, you maintain a live stream of beliefs about how you currently evaluate your options, or that you monitor all your evaluations in experience. We can also look for empirical evidence

that you do one or the other whenever you deliberate. The theory of deliberation dynamics neither expects, requires, nor excludes that you do either.

Conclusion. I will close with a line of thought about distinguishing among the possibilities; it has to do with the ways rational deliberation may end, in problems like Death in Damascus where deliberation is unstable. We have seen that deliberation may settle into, and remain in, an equilibrium state where going to Aleppo and staying in Damascus are seen as equally good. Given sufficient time, in fact, dynamical rules that seek the good in Death in Damascus are bound to lead to an equilibrium. But time for a given activity, deliberation included, is sometimes in short supply. Before reaching equilibrium, you may get tired, you may have other things to do, or the world may interrupt you. Your deliberation may be truncated for many sorts of reasons.²⁰

A question, then: When ongoing deliberation is truncated, why is it truncated where it is? You waver between *A* and *D*; suppose you become tired before reaching the point of thinking the options are equally good, or equally bad. What makes you tired? Is it the effort of evaluating the options yet again, based on your newly adjusted beliefs? Or is it the effort of adjusting beliefs, based on your just-determined values of the options? Perhaps there is no clear answer to the questions. But if the point at which fatigue typically sets in is when yet another utility estimate is called for, that suggests that the adjustments of belief may require less effort, may be more immediate. If on the other hand, fatigue more often leads one to quit at the point of yet another consideration and adjustment of what you will probably do, that suggests the opposite, that the belief adjustments may be less immediate than the revised evaluations. The former case is more

favorable to the idea that the belief adjustments arise from stake-sensitivity; the latter case is less so.

Subjective rational decision theory interests us from both third-person and first-person points of view: as a theory that explains the choice-worthiness of actions for a decision-maker in light of his relevant beliefs and desires, and as a tool that can help us ascertain what to do, as we reflect on our own relevant beliefs and desires. Insofar as it is easier to notice what occurs in our own deliberations than in the deliberations of others, the questions of the previous paragraph implicitly appeal to what we discern from the perspective of first-person deliberation. Present or future empirical methods for studying the deliberations of others might yield third-person evidence too.

Suppose we had empirical support for the idea that your rational belief about what you will do, during your deliberation, is a stake-sensitive response to your regard for doing it. What, if anything, would that tell us about other beliefs of yours? That surely depends on the nature of the empirical support, but the stake-sensitivity of some rational beliefs would be significant, leading us to wonder whether the means by which it occurs also appear in other situations, affecting other beliefs.

In conclusion, we have seen that extended deliberation is a setting in which some rational beliefs are influenced by what is at stake on their truth. If that influence is direct, and does not destroy and create new belief states, the beliefs are stake-sensitive as well as rational. If influence must

always be mediated by other beliefs or by experience, they need not be. The latter alternative seems to me doubtful, but for all I have said, it may be so. Our actual deliberative practices are one source of evidence, among others, that may shed more light on the topic of rational stake-sensitive belief.

Acknowledgement

Thanks to Brian Kim for his very helpful comments and excellent suggestions for improving this paper.

References

- Armendt, Brad (2008) "Stake-invariant Belief." *Acta Analytica* 23: 29-43.
- Armendt, Brad (2010) "Stakes and Beliefs." *Philosophical Studies* 147: 71-87.
- Armendt, Brad (2013) "Pragmatic Interests and Imprecise Belief." *Philosophy of Science* 80: 758-768.
- Armendt, Brad (manuscript) "Causal Decision Theory and Decision Instability."
- Gibbard, Allan and Harper, William (1978) "Counterfactuals and Two Kinds of Expected Utility." In C. Hooker, *et al* (eds.) *Foundations and Applications of Decision Theory*. Dordrecht: Reidel. Reprinted in W. Harper, *et al.* (eds.), *Ifs*, 153-190, Dordrecht: Reidel.
- Harper, William (1986) "Mixed Strategies and Ratifiability in Causal Decision Theory." *Erkenntnis* 24: 25-36.
- Harper, William (forthcoming) "Decision Dynamics and Rational Choice," in festschrift for Allan Gibbard.
- Jeffrey, Richard C. (1965, 1983) *The Logic of Decision*. Chicago: University of Chicago Press.
- Joyce, James M. (2012) "Regret and Instability in Causal Decision Theory." *Synthese* 187: 123-145.
- Richter, Reed (1984) "Rationality Revisited." *Australasian Journal of Philosophy* 62: 392-403.
- Shah, Idries (1967) (Ed.) *Tales of the Dervishes*. London: Jonathan Cape.
- Skyrms, Brian (1982) "Causal Decision Theory." *Journal of Philosophy* 79: 695-711.
- Skyrms, Brian (1990) *The Dynamics of Rational Deliberation*. Cambridge: Harvard University Press.
- Weirich, Paul (1985) "Decision Instability." *Australasian Journal of Philosophy* 63: 465-472.

¹ See Armendt (2008, 2010, 2013).

² Since there is little cost to any answer you give, what you say may only loosely indicate how strongly you believe. A better example might have the radio station allow you to bid for a chance to answer the question. The story raises the possibility of stake-sensitive belief, on which I focus here. Elsewhere I consider similar examples and give more attention to other things that may be happening as well, or instead. See Armendt (2008, 2010).

³ This indefiniteness might be seen as a source of imprecision; I have explored that idea elsewhere (Armendt 2013) but do not focus on it here.

⁴ Your friend offers to place your bet, along with his own, at the racetrack. You tell him to bet \$20 on the horse you like, Apparition. You are anticipating a good race and prospects of victory when your friend returns with your bet on Whirlwind. No doubt this affects your hopes about the outcome of the race; if it also affects your beliefs that Apparition will win, or that Whirlwind will, those beliefs seem sensitive to the shape of the stakes on their truth.

⁵ Taking on a new commitment, say by accepting a new bet on p at odds that fit your degree of belief in p , would alter the part p plays in your new current interests. That might in turn influence your new degree of belief in p . Does making a bet on p ever lead you to become more confident, or less confident, in the truth of p ?

⁶ There is some abuse of notation here and below; I sometimes use A to refer to an action, and sometimes to the proposition 'I do A '.

⁷ The offer expires day after tomorrow; the election is on Tuesday; a plan must be in place by next week; I have one minute to make a move... Since the time of Richard Jeffrey's development of evidential decision theory (Jeffrey 1965), and before, there have been objections to the idea that I may deliberate about doing A , and simultaneously have some belief about whether I will do A . Evidential decision theory expects an agent to assess A 's desirability using conditional beliefs of the form $pr(-/A)$. Without argument, I assume here that present deliberation and present belief about the future action are compatible, and related.

⁸ Such settings often appear in games among Bayesian players; Battle of the Sexes with a Twin is an example.

⁹ Gibbard and Harper (1978, pp. 185-186). The story in their example is a variant of one told in W. Somerset Maugham's *Sheppey*, 1933, and alluded to in the title of John O'Hara's *Appointment in Samarra*, 1934; earlier versions of the story are far older, dating to the ninth century and probably earlier; for example, see 'When Death came to Baghdad,' in Idries Shah (1967). Different cities appear in various earlier versions of the story.

¹⁰ Here I use the K -expectation version of causal decision theory due to Skyrms (1982). Causal decision theory endorses going to Aleppo in the sense that going to Aleppo's expected utility U is maximal, and U represents his preferences. To say that causal decision theory endorses action A is to rely on a principle that endorses actions that maximize U .

¹¹ A plausible answer is toss a coin, or adopt some internal way of performing a mixed act (Harper 1986). However, mixed acts might be ruled out or heavily penalized (Weirich 1985). We might question the legitimacy of such a restriction, but I save that for another occasion. In what follows, I neglect beliefs and assessments concerning mixed acts, but it is straightforward to include them, if a decision-maker takes them to be among his options.

¹² See Skyrms (1982). In later work, Skyrms (1990) developed an important connection between dynamic deliberation and well-founded solution concepts for noncooperative games among Bayesian players. In decision problems with more than two options, deliberation can be significantly more complex than in Death in Damascus.

¹³ When you deliberate about different sorts of problems, your shifting beliefs and evaluations follow different sorts of paths. When causal decision theory is used to evaluate your options in a *Newcomb Problem*, for example, deliberation may yield straightforward convergence to high confidence that you will take both boxes, and that the opaque box will be empty, since an increasing confidence that you will take both boxes does not lead you to think it would be better to do otherwise.

In connection with our larger interests, I do not suggest that all deliberations are clarified by deliberation dynamics; in many deliberations our assessments carry no significant evidence about the state of the world (e.g. what is inscribed). Examples of decision instability such as Death in Damascus provide a setting where beliefs about what you will do matter in deliberation, and the theory of deliberation dynamics gives a framework for saying how they matter, and how they are affected by deliberation.

¹⁴ Also, how *much* do your beliefs about what you will do depend on your assessments of those values? An interesting question; here I suppose that the belief shifts are driven only by your shifting assessments of your options.

¹⁵ The idea of dynamics that seek the good is Skyrms' (1990, p. 30). Such dynamics must also raise the sum of the probabilities of all the actions better than the *status quo*. Nash, Darwin, and Bayes dynamics are examples of dynamics that seek the good. Dynamics that seek the good may fail to represent the deliberations and actions of some agents. An akratic agent, for example, may have little reason to raise beliefs that at the moment of action he will do what now looks better. (Thanks to Brian Kim here.) A self-destructive agent may have reason to lower those beliefs.

¹⁶ The dynamic rule expresses the shift in his beliefs about what he will do, $pr_{t_2}(A)$ and $pr_{t_2}(D)$. Accompanying shifts in his beliefs about what is inscribed, $pr_{t_2}(K_A)$ and $pr_{t_2}(K_D)$, will satisfy Jeffrey-conditionalization if his conditional probabilities such as $pr(K_A/A)$ are stable, and there are no further complexities.

¹⁷ In general, given sufficient time, we can expect convergence to equilibrium from reasonable starting points. For Death in Damascus, continuous deliberation dynamics that seek the good are guaranteed to converge to equilibrium, but they will not display the oscillations I have described. Discrete-time dynamics that seek the good may well display oscillations; with plausible properties such as a dampening in the learning over time, convergence to equilibrium can also be guaranteed. See Skyrms (1990), Joyce (2012), and William Harper (forthcoming).

¹⁸ The original version of Death in Damascus is a symmetric problem, but asymmetric versions are easily given; just add an incentive against travel that makes the outcomes of staying in Damascus a little better than the corresponding outcomes of traveling to Aleppo. Or, imagine that death's appointment book more reliably predicts the traveler's presence when he is in one city than when he is in the other. See Reed Richter (1984).

¹⁹ If a belief that A is better than D really plays a part, it could be influenced by A 's being seen as better than D in either of the two ways discussed in the next paragraphs. One of the two would be stake-sensitivity.

²⁰ See Armendt (manuscript) for more extensive discussion of truncated dynamical deliberation.