Agency and Evidence

Berislav Marušić and John Schwenkler*

It is a special and hard problem to understand how evidence figures into the reasoning of an agent. It is a problem for philosophers but also one we all encounter in our daily lives. Our aim in this chapter is to identify the problem and outline a possible solution to it.[1]

1. The Problem

In a nutshell, the problem that agents face is how to consider evidence in their decision-making without taking it as grounds for a prediction, that is, without turning their decision into an *estimate* of what they are going to do.[2]

Let us illustrate this with an example. Imagine that you are considering running a marathon with a limited number of starting places.[3] The starting places are assigned through a lottery. Now suppose that consideration of your chances reveals the following: Your chance of getting a starting place is pretty good—say 80%. Your chance of actually running the marathon, conditional on getting a starting place and the world cooperating with your plans, is also pretty good—say also 80%. If a well-informed bookie were offering bets on whether (a) you would get

---

* Both authors contributed equally to this work.

[1] For further discussion, see Marušić (2012; 2013; 2015), Marušić and Schwenkler (2018), and Schwenkler (2019, Chs. 5-7). The present discussion synthesizes and builds upon that work.

[2] This use of "estimate" is due to Anscombe (1963/2000, 2). The problem is sometimes discussed in decision theory in terms of the question of whether deliberation crowds out prediction. See Spohn (1977), Levi (1997; 2007); Rabinowicz (2002), Joyce (2002), Hájek (2016), Vavova (2016), Liu and Price (2019). We think that, with the exception of Vavova (2016), those discussions don't quite get to the heart of our problem, because they fail to distinguish the problem as it arises for decision theory from the problem as it arises for the ordinary agent. Our topic here is the latter.

[3] This example organizes the argument of Marušić (2015).

a starting place in the lottery, and whether (b) you would actually run the marathon, conditional on getting a starting place and the world cooperating, they would offer the very same odds for both.

The problem is that *you* cannot consider the question of whether you will actually run the marathon in the same way as that of whether you will get a starting place through the lottery. As an agent, you must face the respective uncertainties in two very different ways. That is because whether you get a starting place in the race is not up to you, and so you can't decide to get a starting place. You have to make your decision to run the marathon conditional on winning a starting place in the lottery—and, in general, conditional on the world's cooperation. But whether you actually run the marathon is, we may suppose, entirely up to you. (Let us ignore the possibility of getting injured, or sick, or hit by a car or a meteorite. Though some of these things are to some extent up to you, they also involve ways in which the world must cooperate.) And if it is entirely up to you whether you actually run the marathon, then you cannot make your decision to run it conditional on—your running the marathon! You have to regard the uncertainty that arises from the possibility of deciding not to run the marathon in a different way than the uncertainty that arises from the possibility of not getting a starting place, or of the world not cooperating with your choice. The fact that it is up to you whether you actually run the marathon makes a difference to how you should think of the evidence concerning whether you will actually do this.

However, it would be a mistake to conclude that, since it is up to you to run the marathon, you needn't consider evidence concerning the difficulty of following through. You cannot make a good decision if you ignore how difficult it is to actually run a marathon. You must consider the fact that running a marathon will require resolve and persistence. So, if you

make a good decision, you will take the difficulty into account in considering both which ends you set for yourself and which means you choose to achieve them. That is to say, you will consider whether this difficult project is really worth the effort. And, if you decide that it is, you will consider how best to pull it off: You will train hard in advance of the race, and you will be mindful of the need for resolve during the race. In this way, you will not ignore the evidence of difficulty. But you will also not use it as the basis for prediction.

The problem is to say how you should consider this evidence, without using it as grounds for *prediction* of your future behaviour. In what way can you reason about the possibility that you will fail to do something that it is up to you to do?[4]

2. Answer I: Decision Theory

A first answer might be suggested by decision theory: The rationality of a decision is a function of both preferences and probabilities. In making a decision, an agent has to take into account both how much they prefer the possible outcomes of their actions and how probable those outcomes are, conditional on performing those actions.

However, we do not think that this is an adequate answer, for two reasons. First, this view does not really address the problem we have identified. Second, it leads to incoherence between an agent's practical and theoretical conclusions.[5]

[4] The question here bears on the extensive discussion of the case of Professor Procrastinate from Jackson and Pargetter (1986). See Goldman (1976), Portmore (2011; 2019) Ross (2012) and Jackson (2014). The view defended here has affinities with Thomason (1981) and Liu and Price (2019).

[5] A further problem is that, on this view, there is no resting place for practical thinking, because the fact that one is making a decision (on certain grounds) affects the probability assignment of the outcomes. Since the outcomes are not independent of one's decisions, it is not clear how the decision-theoretic approach applies. We won't pursue this problem here. The structure of the problem is nicely brought out by Bok (1998), albeit not in probabilistic terms.

Let us illustrate this with an example. Suppose you dramatically prefer to run a marathon, on the grounds that it will secure you abiding admiration from those you care about the most. You then set out to consider *how* to do it. And suppose you have three ways of pursuing your goal: First, you could set out to train on your own. Second, you could join a training group at a local gym. Third, you could join a training group and hire a personal trainer at the same local gym. (Note: You can only hire a personal trainer if you join the training group!) Needless to say, joining the local gym is expensive, and joining the gym as well as hiring a personal trainer is extremely expensive.

If you went by your preferences alone, you would pursue your goal of running the marathon by training by yourself, thus saving yourself the money for a gym membership and a personal trainer. However, since you have to take into account the difficulty of running a marathon, decision theoretic calculation suggests that one of the other courses of action is the more rational one. Let us suppose that it is, in fact, joining a training group and foregoing a personal trainer, as the increase in costs for hiring the trainer do not justify the increase in odds of success. Finally, let us suppose that, in arriving at this conclusion, you have assigned probability 0.8 to the outcome of successfully running the marathon, conditional on joining a training group at the local gym.

But now we can see the problem. In your decision theoretic calculations, you are doing nothing other than making predictions about your future actions—you are assigning probabilities to outcomes that are up to you. And our challenge was to articulate how to consider the uncertainty that arises about matters that are up to the agent is *different* from the uncertainty that arises about matters that are not up to the agent. The decision-theoretic approach has not done that.

Moreover, the decision-theoretic approach has also given rise to incoherence between your practical thought and your theoretical thought. On practical grounds, you have reached a conclusion that you might express by saying, "I will run the marathon by joining a training group." But on purely theoretical grounds, you have reached a conclusion that you might express by saying, "There is an 80% chance that I will run the marathon by joining a training group."[6] And that is odd. You are, in effect, committed to the conclusion that you might express by saying, "I will run the marathon, though there is a 20% chance that I won't." Your view of the future is incoherent.[7]

3. Answer II: Practical Reasoning

Suppose that you deliberate about whether to run the marathon in a way we might represent as follows:

(1) There's a marathon being held in my home city six months from now.
Running the marathon will secure me abiding admiration from those I care about the most.
If I enter to run, there's an 80% chance that my name will be drawn in the lottery.
It will also cost about $150 to register for the race if my name is drawn.
In order to complete the entire race I'll need to do a lot of training.

[6] Assuming something like Lewis's Principal Principle (e.g. Lewis 1994).

[7] One might reply by defending non-cognitivism about practical reason, according to which the conclusions of practical reasoning are not beliefs (e.g. Bratman 1987). However, even on a non-cognitivist view, there are coherence requirements on one's intentions and beliefs, so that one's overall view will be irrational. See Marušić (2013; 2015, Ch. 3.1) and Marušić and Schwenkler (2018) for further discussion.

I will be much better able to accomplish this if I join a training group rather than attempting to train on my own.

However, joining the training group will be expensive, since it requires joining a gym.

My household finances are in good enough shape to cover the $150 fee and the gym expenses, though it will be tight.

So I'll enter to run, pay the $150 fee if my name is drawn, join a training group to help me prepare, and then complete the race on the day it is run.

Here, the "so" of your conclusion in the final sentence is an instance of what Jonathan Dancy (2018, 15) helpfully calls the "practical 'so'": It is part of an act in which a decision to do something is drawn from a body of considerations that one takes to rationalize or support it. The decision is a conclusion of *practical reasoning*.

The hard question is: *How* are the considerations listed as premises in (1) supposed to support your drawing this conclusion?

The first thing we should note is that the expression of your conclusion in (1) is shorthand for a judgment with a lot more internal complexity. That is, the decision you reach when your reasoning concludes is not the decision that, if your name is drawn in the lottery, you will do the things that it describes *no matter what*—that you will do them even if, say, your spouse has a medical emergency and your checking account is overdrawn, or the training group meets during your child's weekly soccer games, or on the day of the race you have a seriously injured ankle. These conditional aspects of your intention, which have been explored in detail in important work by Luca Ferrero (2009), are left implicit in (1) but might be stated explicitly if they became contextually relevant—if someone were to ask, say, whether you really are going to run the

marathon if your injured ankle isn't healed.[8] This is what we earlier tried to capture by saying that the world must cooperate with your plans.

The conclusion that you reach in (1) is not, however, even implicitly conditional on anything like your still *wanting* to run the marathon on the day it happens or to train for it during the months leading up to it. If the conclusion of your reasoning were conditional on things like those, then your conclusion would be better expressed by saying, e.g., that "I'll think about running the marathon," or "I'll start training for the race and then see how I feel," or something of that sort. But what accounts for this difference? Looking forward to your future, you can appreciate that there's a non-zero chance both that (a) when the day of the marathon comes you'll have sprained your ankle training and will decide not to run the race in order to avoid making the injury more serious; and that (b) when the day of the marathon comes you'll strongly desire to sleep in and will decide not to run the race on those grounds. We may suppose that neither your injured ankle, nor your desire to skip the race, will *prevent* you from running in the way that, say, being kidnapped would do. Then why is the injury, but not the desire to bow out, something your decision can be conditional on?

The answer is not simply that you are able to *control* your desires in a way that you cannot control your bodily health—for in fact the reasoning in (1) can be supplemented with further practical reasoning aimed at doing both of these things. For example, in anticipation of getting cold feet, you might reason as follows:

(2) I'm likely to want to skip the race if I haven't prepared mentally for it.

[8] For further development of our view about this, see Marušić and Schwenkler (2018, §3.1).

So I'll spend the week before the race watching films like *Chariots of Fire* and *Prefontaine*.

Similarly, in anticipation of a possible injury, you might reason along these lines:

(3) I'm likely to injure my ankle if I play pickup basketball.
So I'll also avoid playing pickup basketball until the marathon is past.

In both cases, you reason practically in light of the conclusion you reached in (1). Indeed, it seems like a commitment to reasoning in ways like these is a rational requirement of being seriously committed to the conclusion of (1). In each case, however, it seems clear that your strategy for preventing the undesired outcome *might* fail to do the trick. We are back with the question asked earlier: Why is the potential injury, but not the potential desire to bow out, something that your decision to race the marathon can be conditional on?

Here is a better answer: The difference between an injured ankle and a desire not to run is that only the former gives you any kind of *reason* not to run the race. If on the day of the race your ankle is injured, this gives you reason to sit the race out, because completing it would cause you severe pain, or risk making your injury more serious. By contrast, if on the day of the race you simply want to skip it, then this gives you *no reason at all* not to run, but only a reason to suck it up. Of course, if your desire not to run the race is itself well grounded, say because there's been an emergency at home that needs your attention, *then* you have a reason not to run the race. But in that case the reason you have is not your desire itself, but rather the situation that

it is a response to. *Simply* wanting not to do something you've decided to do, gives you no reason at all to abandon your decision in favour of this desire.

Let's return to the question we asked originally. In virtue of what do considerations listed as premises in (1) put you in a position to draw the conclusion you do? That conclusion has the form of a statement about the future: You'd express it by telling someone, "I am going to enter the lottery to run the marathon, and train for the race and then run it if my name is drawn." Yet how can you conclude *this* on the basis of the considerations that you reason from?

To see how this is possible, let us compare the reasoning in (1) with reasoning that might proceed from the following considerations:

(4) All the evidence indicates that Candidate X is going to enter the race.
If she enters the race she will raise millions of dollars in campaign funds.
If she raises millions in campaign funds she will have an 80% chance to win the primary.
If she wins the primary she will win the general election.

Imagine that on the basis of the considerations in (4) you draw a conclusion that you might express by saying, "Candidate X is going to enter the race, raise millions of dollars in campaign funds, and win the general election if she wins the primary." This conclusion, just like your conclusion in (1), is a description of future events. Yet whereas the conclusion in (4) will be reasonable only if it is supported by sufficient *evidence* that this is what Candidate X is going to do, it seems reasonable for you to conclude as you do in (1) even if the considerations you reason from, or any others you might be able to appeal to, do not provide sufficient evidence that you'll

act in the way described.[9] That is because in reasoning as you do in (1), in contrast to (4), *the conclusion that you draw is such as to depend on the reasoning in light of which you draw it*. That is, when you reason as you do in (1), the question "Am I going to run the marathon?" is not a question that you treat as having an answer that it is the aim of your reasoning to discover. This is rather the way that you reason in (4), where Candidate X's likely decisions, and her prospects in the election depending on what she decides, are matters that you treat as fixed, and that you aim in your reasoning to reflect.[10] By contrast, when you reason about what *you* are going to do, the answer that you reach is supposed to be settled by the reasoning that leads you to it, not by some independent matters which that reasoning aims to reflect. Indeed, we can now see the error of predicting matters that are up to us: The error is to treat the subject matter as to be *reflected* in our reasoning rather than to be *made so* by it.[11]

4. Practical Knowledge

This final section will relate the position we have developed to a further thesis, frequently identified with the work of G.E.M. Anscombe, namely that practical reasoning about what to do gives an agent practical *knowledge* of what she is doing or is going to do.

The version of this thesis that Anscombe herself explicitly endorses extends only to present action, insofar as this action is intentional: it says that practical reasoning is the ground of

[9] You will, however, need to have adequate evidence for many of the premises you reason from—e.g. that the race is being held on such-and-such a date, that you have the money to pay the registration fee, etc.—but that is another issue.

[10] There is, of course, a general problem about how to understand rational belief and knowledge of the future, if the future is genuinely open. This is an issue we set aside.

[11] Velleman's distinction between conclusions we accept "so as to reflect the truth" and conclusions we accept "so as to create the truth" informs our present discussion (Velleman 1996, 195, n.55). However, we do not adopt the entirety of Velleman's view. For differences, see Marušić (2015, Ch. 6.2).

practical, non-observational knowledge of what one is *presently* in the process of doing (Anscombe 1963, 11-12 and *passim*). But Anscombe brings her account of practical reasoning to bear on the topic of future action as well, claiming that statements of the form, "I am going to φ," where they express intentions rather than estimates of one's future behaviour, are grounded in reasons that do not suggest "what is probable, or likely to happen," but rather in reasons "suggesting what it would be good to make happen with a view to an objective" (Anscombe 1963/2000, 4). And other philosophers, most prominently David Velleman (1989/2007) and Kieran Setiya (2007; 2008) have treated belief about what one *will do* as a central element in their accounts of an agent's practical knowledge.

The account of practical reasoning that we have developed here suggests a way around a common objection to positions like Velleman's and Setiya's, namely that in positing an essential connection between intention and belief they demand that an agent make an epistemically unjustified "leap of faith" in forming the intention to do something.[12] By contrast, a central element in our account is that beliefs about one's future actions are necessarily *not* subject to the evidential norm governing beliefs about other matters. To the extent that your future actions are up to you to decide, they are such as to depend on the reasoning in light of which you decide this, which means there is not a requirement that they be supported by adequate evidence.

However, we also hold that an agent's practical reasoning does not always gives her practical knowledge that she will do what she decides.[13] When, through practical reasoning, you conclude that you will run the marathon, you do not necessarily have practical knowledge that you will do so. That is because the fact that it is difficult to run the marathon may be evidence

[12] See esp. Langton (2003) for this objection in response to Velleman (1989/2007).

[13] We take a similar position in regard to present action as well. For relevant arguments see Schwenkler (2015, 17-25) and (2019, Ch. 6.4).

that defeats your practical knowledge, preventing you from *knowing* that you will act as you have decided.

To see this, consider a very different kind of case, involving mathematical knowledge. Suppose you calculate something by longhand. Your calculation is accurate and you have made no mistake. It seems that you have *a priori* knowledge of the result of your calculation: Your knowledge of the result does not rest on evidence or observation but is grounded in the calculation alone. However, if you acquire evidence that there is a good chance that you made a mistake—say, because you often make mistakes when you perform such calculations, or because people who are in many respects like you often make mistakes, or because a bookie will offer bets that reflect a significant chance that you have made an error—this will defeat your a *priori* knowledge, even if your calculation was perfect. Analogously, the fact that there is a 20% chance that you will fail to run the marathon serves as a defeater for any potential practical knowledge you might have had. Your situation would be analogous to a case in which you proved something mathematically but had empirical evidence that you made a mistake.

It is instructive to consider why this must be so.[14] If our practical reasoning always had to conclude in practical knowledge in order to be good reasoning, there would be a constraint on the strength or significance of our reasons: reasons could speak in favour of φ-ing only if they put us in a position to know that will φ if we decide to φ. The resulting view of the strength or significance of our reasons would be implausible, because we may have overwhelming reasons to do something that is exceedingly difficult. We may strongly want, immensely value or even simply be obliged to do something, and this can make it practically rational to do it—but the fact that we desire or value or are obliged to do it need not put us in a position to know that we will

[14] See Marušić (2015, esp. 115-116) and Schwenkler (2019, Ch.6.4) for further discussion.

do it, since the difficulty of doing it would be a defeater for our purported knowledge.

The argument of the preceding section also shows that an agent is in a different position with respect to evidence about the difficulty of her intended actions than with respect to evidence concerning matters that are not up to her. For example, the fact that it will be difficult to run the marathon, due to the physical demands of the race as well as the challenge of keeping yourself sufficiently confident and motivated, is something you can account for in your practical reasoning by forming plans that will help you to overcome these difficulties. Yet plans like these cannot eliminate entirely the possibility of failing to do what you intend. In particular, they cannot eliminate the need for *resolve* in sticking to a decision in the face of temptations to abandon it.

What is the right response to this limitation? It does *not* lie in saying "I'll do it … unless I don't" in any situation where the possibility of failure becomes salient—for as we have argued, this is to regard one's intention as conditional on one's future state of mind (cf. Anscombe 1963/2000, §52). The moral is rather that, in making up our mind, we sometimes must be prepared to be resolute in the face of difficulty—even when we know that no bookie will think in the way that we do.

5. Conclusion

Kant's famous dictum is that we act under the idea of freedom. In the present paper, we do not mean to make headway on the question of what freedom is. However, we hope to have shown that freedom has an important epistemological significance: It licenses us to think of the future in a distinctive way. In particular, when something is up to us, then no matter how hard it is for us to do it, what will happen is something for us to decide—and deciding is a *practical*

matter. We evade our responsibility if in those circumstances we seek to predict what will happen—and sometimes such evasion is couched in claims of theoretical rationality. As agents, we should forego such comfort.

References

Anscombe, G.E.M. (1963/2000). *Intention*, 2nd ed. Cambridge, MA: Harvard University Press, 2000.

Bok, Hilary. (1998). *Freedom and Responsibility.* Princeton: Princeton University Press.

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

Dancy, Jonathan. (2018). *Practical Shape: A Theory of Practical Reasoning.* Oxford: Oxford University Press.

Ferrero, Luca. (2009). "Conditional Intentions." *Noûs* 43, 700–41.

Goldman, Holly. (1976). "Dated Rightness and Moral Imperfection." *The Philosophical Review* 85, 449–87.

Hájek, Alan. (2016). "Deliberation welcomes prediction." *Episteme* 13, 507–528.

Jackson, Frank and Robert Pargetter. (1986). "Oughts, Options, and Actualism." *The Philosophical Review* 95, 233–55.

Jackson, Frank. (2014). "Procrastinate Revisited." *Pacific Philosophical Quarterly* 95, 634–47.

Joyce, James. (2002). "Levi on causal decision theory and the possibility of predicting one's own actions." *Philosophical Studies* 110, 69–102.

Langton, Rae. (2003). “Intention as Faith.” In *Action and Agency*, edited by Helen Steward. Cambridge, UK: Cambridge University Press.

Levi, Isaac. (2007). “Deliberation Does Crowd out Prediction.” In *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*, edited by Toni Rønnow-Rasmussen, Björn Petersson, Jonas Josefsson and Dan Egonsson, Lund University.

Levi, Isaac. (1997). *The Covenant of Reason: Rationality and the Commitments of Thought*. Cambridge, UK: Cambridge University Press.

Lewis, David. (1994). “Humean Supervenience Debugged.” *Mind*, 103: 473–90.

Liu, Yang and Huw Price. (2019). “Heart of DARCness.” *Australasian Journal of Philosophy* 97, 136–50.

Marušić, Berislav and John Schwenkler. “Intending is Believing: A Defense of Strong Cognitivism.” *Analytic Philosophy* 59.3 (2018): 1–32.

Marušić, Berislav. (2015). *Evidence and Agency: Norms of Belief for Promising and Resolving.* Oxford, UK: Oxford University Press.

Marušić, Berislav. (2013). “Promising against the Evidence.” *Ethics* 123, 292–317.

Marušić, Berislav. (2012). “Belief and Difficult Action.” *Philosophers' Imprint* 12, 1–30.

Portmore, Douglas. (2019). *Opting for the Best: Oughts and Options.* New York: Oxford University Press.

Portmore, Douglas. (2011). *Commonsense Consequentialism.* New York: Oxford University Press.

Rabinowicz, Wlodek. (2002). “Does Practical Deliberation Crowd out Self-Prediction?” *Erkenntnis* 57, 91–122.

Ross, Jacob. (2012). "Actualism, Possibilism, and Beyond." *Oxford Studies in Normative Ethics* 2, edited by Mark Timmons. Oxford: Oxford University Press.

Schwenkler, John. (2019). *Anscombe's* Intention: *A Guide.* New York: Oxford University Press.

Schwenkler, John. (2015). "Understanding 'Practical Knowledge'." *Philosophers' Imprint* 15, 1–32.

Setiya, Kieran. (2007). *Reasons Without Rationalism*. Princeton: Princeton University Press.

Setiya, Kieran. (2008). "Practical Knowledge." *Ethics* 118, 388–409.

Spohn, Wolfgang. (1977). "Where Luce and Krantz Do Really Generalize Savage's Decision Model." *Erkenntnis* 11, 113–34.

Thomason, Richmond. (1981). "Deontic Logic and the Role of Freedom in Moral Deliberation." *New Studies in Deontic Logic: Norms, Actions and the Foundations of Ethics* 152, 177–86.

Vavova, Katia. (2016). "Deliberation and Prediction: It's complicated." *Episteme* 13, 529–38.

Velleman, David. (1989/2007). *Practical Reflection*, 2nd ed. Stanford: CSLI Publications.

Velleman, David. (1996). "The Possibility of Practical Reason." In *The Possibility of Practical Reason.* Oxford: Oxford University Press, 2000, 270-299.