

UNCERTAINTY AND CONTROL

– Sven Ove Hansson –

Abstract. In a decision making context, an agent’s uncertainty can be either epistemic, i.e. due to her lack of knowledge, or agentive, i.e. due to her not having made (full) use of her decision-making power. In cases when it is unclear whether or not a decision maker presently has control over her own future actions, it is difficult to determine whether her uncertainty is epistemic or agentive. Such situations are often difficult for the agent to deal with, but from an outsider’s perspective, they can have sensible pragmatic solutions.

Keywords: uncertainty, control, agentive uncertainty, Newcomb’s paradox, death in Damascus.

1. Two types of uncertainty

When we adopt words from everyday language for technical usage, there are invariably changes in meaning. New distinctions are introduced, and connotations are removed or disregarded. However, connotations from common usage can have an influence on how we use and understand such terms, despite being unmentioned in the technical definitions. The term “uncertainty,” as used in decision theory, is a case in point. Technically, it is standardly defined as lack of knowledge,¹ but in actual usage it can also refer to something else. Consider the following two examples:

- (1) I am uncertain whether that book is translated or it was originally written in English.
- (2) I am uncertain whether I will read the book she gave me.

Let us test the interpretation of uncertainty as a lack of knowledge:

- (1′) I lack knowledge on whether that book is translated or it was originally written in English.
- (2′) I lack knowledge on whether I will read the book she gave me.

(1′) is an ordinary sentence, roughly synonymous with (1). However, (2′) is peculiar. The question whether I will read a book in my possession is not a matter of whether I have or lack some knowledge or information. It is a matter of what deci-

¹ Knight (1921/1935): 19–20; Luce, Raiffa (1957): 13.

sions I make (and carry out). To further clarify this, let us try that interpretation in both cases:

(1'') I have not decided whether that book is translated or it was originally written in English.

(2'') I have not decided whether I will read the book she gave me.

Here (1'') is anomalous. Whether a particular book was translated into English or originally written in that language is (after the fact) not something that I or anyone else can decide. It can be known but (no longer) decided.

Unfortunately, this distinction is far from perfectly clear in ordinary English. We sometimes use the term "know" in the sense of having decided what to do in the future. ("I do not know if I will read that book." - "Now I know what I will recommend her to do.") However, the distinction will usually come out clearly if we substitute "have knowledge" for "know" and "lack knowledge" for "do not know".

Hence, the word "uncertainty" covers two meanings. That something is uncertain in an agent's perspective can mean that it is not *known* by the agent. We can call this epistemic uncertainty. Yet uncertainty can also mean that something has not been *decided* by the agent, or not decided in the right way (e.g., with sufficient determination). We can call this agentive uncertainty. The distinction between epistemic and agentive uncertainty is agent-relative. Something that is uncertain for two persons can be so in the epistemic sense for one of them and in the agentive sense for the other. For instance, a teacher and her student can both be uncertain about what grade the student will receive. For the student this is unknown; for the teacher it is undecided.

In addition to our first-order uncertainty we can be uncertain about the nature, extent etc. of our own (or someone else's) uncertainty. In particular, we can be uncertain whether our own uncertainty in a particular matter is epistemic or agentive. Such situations are often difficult for decision makers to deal with. It is the purpose of this article to show that agentive uncertainty is practically important, theoretically interesting, and in need of explicit decision-theoretical treatment.

Section 2 shows that agentive uncertainty is ubiquitous in practical decision-making. Section 3 makes it clear that although agentive uncertainty is not much discussed in decision theory, it is an important (but often unrecognized) factor in some of the problems discussed in that discipline. Section 4 is devoted to the analysis of agentive uncertainty from the agent's own perspective and Section 5 to its analysis from an outsider's point of view. In Section 6 we return to the agent's perspective and draw some general conclusions from the discussion.

2. Agentive uncertainty in practice

Decision theory has traditionally been devoted to well-defined decision problems in which the options and the potential outcomes are known. In decision making “under risk,” the probabilities of the outcomes (given each of the options) are known. In decision making “under uncertainty” they are unknown or only partially known. Real-life decisions usually start out with problems that are much less well-defined than decision-theoretical problems “under uncertainty”. In such decisions under “great uncertainty” (also called “deep uncertainty”) information can be missing about a wide range of aspects of the decision, including:

- what issues are going to be settled in the decision;
- whether these issues will all be settled in a single act of decision making, or divided up between subdecisions;
- the appropriate time for decision making;
- whether additional information should be gathered before the decision is made;
- what options can be chosen in the decision;
- what aspects of the decision outcome should be included in the evaluation of the options.²

The issues on this list have to be determined in the course of decision making, typically in preparatory decisions that lay down the structure of the decision. Traditional decision theory assumes that all this has been settled, thereby excluding many of the concerns of real-world decision making from its considerations.

The second item on the list is highly illustrative. It usually concerns sequential decisions, i.e. decisions on related topics that are made at different points in time. A decision maker can plan in detail beforehand how she will act throughout a whole series of decisions. But how sure can she be that she will follow through with her plans at all later decision points?³ Or, to put it in another way: When she makes these plans, is she in full control over her future actions? If she is in full control, then she can treat the whole series of decisions as a single, one-shot decision: First she makes a plan for how to act at each future decision-point, and then she just follows that plan. If she does not consider herself to be presently in control over what she will do in the future, then it would be more sensible to base each decision in the sequence on an assessment of the various ways in which she may come to act in the future. (Such an assessment may include the assignment of

² Hansson (1996, forthcoming-a); Hansson, Hirsch Hardon (2016).

³ Rabinowicz (2002); Ledwig (2005).

probabilities to her own alternative future courses of action, or it may employ non-probabilistic decision rules such as the maximin rule.) Often, she will be uncertain about which of these two types of situation she is in; in other words she does not know whether to treat her uncertainty as agentive or epistemic. We have probably all asked ourselves questions such as: Can I open the box of chocolates and take just one single piece? If I join my friends at the pub, will I return home sufficiently early and sober to finish the work that I promised to deliver early tomorrow morning? If I put off this tedious work until just before it has to be finished, will I actually complete it in time? If we are in full control over the future actions that are relevant in these cases, then our uncertainty is purely agentive, and we can resolve it by making a decision. In other words, we can safely open the chocolate box, join our friends at the pub, and postpone the uninspiring task. If we have little or no control over these future actions, then we presumably had better refrain from doing so. But what should we do in the (arguably typical) case when we are uncertain about whether we have (sufficient) control over our future actions?

These examples alone should be sufficient to show that agentive uncertainty needs to be treated explicitly in decision theory. There is also another reason for this, namely that some of the problems discussed in decision theory are in fact concerned with agentive uncertainty.

3. Agentive uncertainty in decision theory

Agentive uncertainty is highly relevant in the type of intriguing but practically rather irrelevant decision problems in which the agent has access to highly reliable predictions about her own future decisions. The most famous of these problems was conceived by the physicist William Newcomb, but first published by Robert Nozick:⁴ In front of you there are two boxes. One of them is transparent, and you can see that it contains \$ 1000. The other is covered, so that you cannot see its contents. It contains either \$ 1,000,000 or nothing. You have two options to choose between: Either you take both boxes, or only the covered box. A predictor, who has infallible (or almost infallible) knowledge about your psyche, has put the million in the covered box if he predicted that you will only take that box. Otherwise, he has put nothing in it. If you rely on the predictor, then it makes sense to treat this as a knowledge problem and consequently take only one box. On the other hand, if you persevere in seeing yourself as a decision maker who is still able, after the prediction, to choose either option, then it makes more sense to take both boxes.

⁴ Nozick (1969).

Gibbard and Harper have contributed an example in this tradition that is commonly referred to as “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 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...

If... he decides to go to Aleppo, he then has strong grounds for expecting that Aleppo is where death already expects him to be, and hence it is rational for him to prefer staying in Damascus. Similarly, deciding to stay in Damascus would give him strong grounds for thinking that he ought to go to Aleppo...⁵

In this case as well, the agent has a difficult time because it is unclear whether he should treat his uncertainty about his own future actions as epistemic or agentic. In the former case, he should assume that the predictor is right and presumably cannot be outsmarted; in the latter he should instead treat the problem as resolvable with his own decision-making power. Paradoxes like these are unrealistic since they rely on the existence of some external intelligence that can predict the agent's actions to such a high degree that the agent lacks control over her actions in a situation when we would normally expect her to have such control. Yet, as we have already seen, the self-control problem also arises in common everyday situations.⁶

⁵ Gibbard, Harper (1978/1988): 373–374.

⁶ The chocolate box problem mentioned in Section 2 may provide a better vantage point for attacking it than Newcomb's two-box problem.

4. The agent's perspective

The following is such a realistic, everyday example:

The gymgoer's dilemma

„I have decided to go to the gym twice a week from now on. Should I buy a 12 month gym membership, or should I pay for each visit? If I carry through my plan and go twice a week, then the 12 month card will be less expensive than paying each time. But if I fail to do as I plan, then the 12 month card will, of course, be a waste of money.”

„Well, I read an article in the newspaper last week by someone working at a consumer information bureau. She said that most gym beginners do not fulfil their ambitious plans, and consequently they tend to buy expensive long-term cards that they do not use much. Her advice was not to buy such a card until you have established regular training habits.”

„I can see the point. But on the other hand, buying the 12 month card can be a way to convince myself to actually do as I have planned.”⁷

The gymgoer in this example vacillates between the control and no-control approaches to her own future decisions and behaviour. If she assumes that she is in control, then it is reasonable for her to treat the uncertainty as agentive. In other words, she can then assume that if she decides with sufficient determination to go to the gym twice a week, then this is what she is going to do. Consequently, she will buy the 12 months card. If, on the other hand, she applies the no-control approach, then her uncertainty is epistemic. She can then look at her future decisions much in the same way as she would consider the corresponding decisions by another person in the same situation. That might lead her to refrain from buying the 12 month card, at least for the time being.

The obvious way to structure this decision is to first make up one's mind on whether one is going to treat the uncertainty as epistemic or agentive, and then make the actual decision with a decision rule that is appropriate for the chosen type of uncertainty. The problem is, of course, that it may be difficult, perhaps even practically impossible, to determine the type of uncertainty. The following two cases illustrate the same form of metauncertainty.

⁷Hansson (forthcoming-a).

Torture perseverance

Pat is a secret agent. She decides that she will never, even under the worst imaginable torture, give away any information that the enemy can have use for. Based on that, she decides to volunteer for a mission with a high risk of being caught, although she knows that she has much more sensitive information than other agents who might have been selected for this particular mission.

Too much wine

Pat is determined to be faithful to her partner. She knows from experience that she is too easily seduced under the influence of alcohol. One evening after work she finds herself alone in a bar with Richard, who offers her one drink after the other. She says to herself: "What happens, happens. Unfortunately, I cannot control myself in a situation like this. It is meaningless even to try." The next morning she wakes up in the wrong bed, remorseful and miserable.

Torture perseverance shows that in some cases when it is unclear whether uncertainty is epistemic or agentive, it appears more sensible to treat it as epistemic. *Too much wine* illustrates that in other such cases, it seems more reasonable to treat the uncertainty as agentive. But *the gymgoer's dilemma* does not belong to any of these categories. In that case, there are reasonable justifications for both approaches.

5. The outsider's perspective

Up to now we have discussed decision making from the decision maker's perspective. However, decisions can also be discussed from the perspective of an outsider, who may for instance be an observer or an adviser. Since an outsider cannot control the agent's (future or present) decisions, she will always treat the uncertainty as epistemic, but when reasoning in that perspective she may very well conclude that she should recommend the agent to treat it as agentive.

The consistent physician: First consultation

PHYSICIAN: I will give you a referral to the Smoking Cessation Clinic for support sessions and perhaps a nicotine replacement.

PATIENT: People say it is very difficult to quit for someone who smokes as much as me.

PHYSICIAN: It may be difficult, but it is certainly possible and, as I said, due to your heart condition it is very important in your case. I am convinced that if you make a good effort and follow their advice, then you will be a non-smoker when I see you next time three months from now.

Second consultation (three months later):

PHYSICIAN: Have you been to the Smoking Cessation Clinic?

PATIENT: Yes, and they were very friendly and tried to help me. But the cravings were too strong. I had eight smoke-free days, but then I couldn't resist the urge any more. This is really bad of me.

PHYSICIAN: I couldn't agree more. This is certainly bad of you. And as I told you last time we met, if you had made a real effort you would have been almost sure to succeed. It is therefore your own fault that the disease has worsened since last time you were here.

What our hypothetical physician said at the second visit is perfectly consistent with what he said at the first visit. However, this is not how we would expect a physician to interact with patients. At the second visit, we would expect him to avoid statements that unnecessarily burden the patient with responsibility for the failure. Instead, he is supposed to help the patient to find out how external circumstances can be improved in ways that will increase the chances to succeed in a new attempt to quit smoking. I propose that such an approach, rather than that exhibited in the above dialogue, corresponds to actual practice in health care.⁸ When patients need to make difficult changes in their habits, health care personnel tend to encourage them to take a control approach to these habits, "you can do it if you try hard." However, if a patient fails in her attempts to implement the recommended changes in her life, then she is no longer told that doing so was totally within her control. Instead, the focus is shifted to how the circumstances can be changed in ways that will increase her chances of control in a new attempt to achieve the recommended changes in her way of life. Such shifts can be justified in terms of the effects that the message is expected to have on the patient. Telling a person that she is able to do something difficult appears to strengthen her chances of success. On the other hand, if she fails, then a message that lays the blame on her lack of willpower or stamina would seem to have a negative effect on her chances to succeed if she tries again. Therefore, the best way to promote the desired outcome seems to involve a shift from a control to a no-control message in situations like this. This can be described as a *pragmatic* approach. It avoids self-defeating control ascriptions that could reduce the chances of achieving that which the use of control would expectedly aim at achieving.⁹

⁸ Hansson (forthcoming-b).

⁹ Hansson, Edvardsson Björnberg, Cantwell (2016).

6. The agent's perspective reconsidered

Since the pragmatic approach, at least in principle, maximizes the chances of success, it would seem appropriate for the agent to apply it to herself. This would mean that she tries to look at herself from the outside, and choose the control or no-control approach in the way that a benevolent observer would have recommended. However, that can be a cognitively difficult operation to perform. It is frequently (all too) easy to make others believe something that one does not have sufficient grounds for believing to be true, but we are often unable to apply that same operation to ourselves.¹⁰

Unfortunately, the alternatives do not seem very promising either. We have already seen that neither a consistent control approach nor a consistent no-control approach will yield desirable results in all cases.

But even if we are unable to apply a pragmatic approach to ourselves, we may have other mechanisms with approximately the same effects. Psychologists have described an illusion of control, which consists in overestimating the chances of succeeding in what one tries to do.¹¹ Another way to express this is that in some situations we tend to treat epistemic uncertainty as agentic. Perhaps that illusion is an evolutionary advantage? If it is, then the best way to deal with second-order uncertainty on whether our first order uncertainty is epistemic or agentic may be a non-rational way of thinking that it is rational for us to indulge in.

References

- Gibbard A., Harper W.L. (1978/1988), "Counterfactuals and Two Kinds of Expected Utility," [in:] *Decision, Probability, and Utility*, P. Gärdenfors, N.-E. Sahlin (eds), Cambridge University Press, Cambridge: 341–376.
- Hansson S.O. (1996), "Decision-Making Under Great Uncertainty," *Philosophy of the Social Sciences* 26 (3): 369–386.
- Hansson S.O. (2013), "Contraction, Revision, Expansion: Representing Belief Change Operations," [in:] *Kristen Segerberg on Logic of Actions*, R. Trypuz (ed.), Springer, Dordrecht: 135–151.
- Hansson S.O. (forthcoming-a), "Decision Structuring for Risky Decisions," *Ethical Theory and Moral Practice*.
- Hansson S.O. (forthcoming-b), "The Ethics of Making Patients Responsible," *Cambridge Quarterly of Healthcare Ethics*.

¹⁰ On the difficulties of epistemic voluntarism, see Hansson (2013): 143–145, and the references given there.

¹¹ Langer (1975); Yarritu, Matute, Vadillo (2014).

- Hansson S.O., Edvardsson Björnberg K., Cantwell J. (2016), "Self-Defeating Goals," *Dialectica* 70 (4): 491–512.
- Hansson S.O., Hirsch Hadorn G. (2016), "Ten Core Concepts for the Argumentative Turn in Policy Analysis," [in:] *The Argumentative Turn in Policy Analysis. Reasoning about Uncertainty*, S.O. Hansson, G. Hirsch Hadorn (eds), Springer, Dordrecht: 347–353.
- Knight F.H. (1921/1935), *Risk, Uncertainty and Profit*, London School of Economics and Political Science, London.
- Langer E.J. (1975), "The Illusion of Control," *Journal of Personality and Social Psychology* 32 (2): 311–328.
- Ledwig M. (2005), "The No Probabilities for Acts-Principle," *Synthese* 144 (2): 171–180.
- Luce R.D., Raiffa H. (1957), *Games and Decisions: Introduction and Critical Survey*, Wiley, New York.
- Nozick R. (1969), "Newcomb's problem and two principles of choice," [in:] *Essays in Honor of Carl G. Hempel*, N. Rescher et al. (eds), Reidel, Dordrecht: 114–146.
- Rabinowicz W. (2002), "Does Practical Deliberation Crowd Out Self-Prediction?" *Erkenntnis* 57 (1): 91–122.
- Yarritu I., Matute H., Vadillo M.A. (2014), "Illusion of Control. The Role of Personal Involvement," *Experimental Psychology* 61 (1): 38–47.