Decisions and Higher-Order Knowledge

Moritz Schulz, University of Tübingen

Penultimate draft. Forthcoming in *Nôus*.

Abstract

A knowledge-based decision theory faces what has been called the *prodigality problem* (Greco 2013): given that many propositions are assigned probability 1, agents will be inclined to risk everything when betting on propositions which are known. In order to undo probability 1 assignments in high risk situations, the paper develops a theory which systematically connects higher level goods with higher-order knowledge.

1 Knowledge and Decision

It is not unlikely that the prime reason for seeking knowledge is to make good decisions. When Sally checks the timetable to see when the bus is coming, she will typically do so because she can make a better decision about when to leave the house with the relevant information at hand. Needless to say, this does not imply that it is the only reason. Knowledge surely has some intrinsic value, at least in certain cases, and knowledge can also serve other goals one might have, for example, becoming famous by winning the Nobel prize.

If there is some truth to this conjecture, one would expect that knowledge plays a key role in decision making. There is evidence that this is indeed the case. For instance, Williamson (2005a: 231) suggests that a proposition p "is an appropriate premise for one's practical reasoning" iff one knows p. In a similar vein, Hawthorne and Stanley (2008) argue that one should treat the proposition that p as a reason for acting only if one knows that p (similar claims have been defended by Fantl and McGrath 2002, Hyman 1999 and Unger 1975: ch. 5).

If decisions should be based on what one knows, knowledge should play a fundamental role in decision theory. Assuming that the right decision to make is one which maximizes (in one way or other) expected utility, the question would be of what kind the probabilities are which figure in the calculation of expected utility. If rational decisions are based on knowledge, a minimal requirement should be that the probabilities are taken conditional on the knowledge the agent has. This has the consequence that what the agent knows always has probability 1 (because $P(p|K) = 1$ if $K \models p$). Williamson (2000) has developed a picture of knowledge and probability which accommodates a decision theory of this kind very naturally.[1] According to Williamson, there is an *evidential probability function* P measuring, roughly speaking, the objective degree of justification a body of evidence E lends to a proposition p.[2] The evidential probability for a proposition p relative to an agent S and a time t is then given by $P(p|K)$, where K is the knowledge S has at t.[3] On this picture, knowledge always has evidential probability 1.

The decision theory sketched so far conflicts with standard interpretations of Bayesian decision theory. According to many Bayesians, probability 1 should only be given away very sparingly, to logical truths and perhaps to propositions reporting the contents of our subjective experiences (cp. Jeffrey 1965 and Lewis 1981). One prominent reason for this sparingness is decision-theoretic: if a proposition p has probability 1, betting on p has positive expected utility no matter how daunting the odds are (as long as the utility of winning the bet is itself positive), for the chances of losing the bet will be 0.[4] In the form of a popular slogan: one would have to bet one's life for a penny (see sec. 2 for a more detailed analysis). Given that a knowledge based decision theory without any skeptical inclinations allows assignments of probability 1 across a wide range of cases, including substantial claims about the external world, the question how such a theory deals with high-loss low-gain bets of the aforementioned kind is particularly pressing. In light of this problem, Greco (2013) and Kaplan (2009) have argued that a knowledge based decision theory is ultimately untenable.

There are various possible responses available to a defender of a knowl-

[1] See also Levi (1980), who takes knowledge to be the standard of (serious) possibility.

[2] This exposition is more directed at conditional probabilities of the form $P(p|E)$. Unconditional probabilities of the form $P(p)$ are explained by Williamson (2000: ch. 10) as measuring the intrinsic plausibility of p.

[3] For ease of exposition, I use symbols like 'p' ambiguously both as schematic names for propositions and as propositional schema letters in sentence position. This ambiguity should be harmless as the context will always make clear of which logical type the relevant symbols are.

[4] Another reason has to do with the *dynamics* of belief: assignments of probability 1 cannot be undone by conditionalization, the standard update rule in Bayesian epistemology.

edge based decision theory. She could go contextualist, relativist or sensitive invariantist and hold that higher stakes increase the standards for knowledge, so that one would no longer count as knowing in the problematic scenarios. For instance, Hawthorne and Stanley (2008) adopt sensitive invariantism as part of their defense strategy (but see Greco 2013 for a critical assessment of contextualist and sensitive invariantist solutions of the problem). In a similar spirit, a proponent of an insensitive invariantism about knowledge could explore the possibility that higher stakes require higher-order knowledge (cp. Williamson 2005a, but also Hawthorne and Stanley (2008) consider this as a possible way out). Greco (2013) has objected to this strategy that—as things currently stand—it is unclear whether such a strategy can be turned into a systematic decision theory. In this paper, I try to close this gap by drawing a more systematic picture of how high stakes decisions might be connected with higher levels of knowledge. If successful, this would make room for a knowledge based decision theory which is compatible with an anti-skeptical insensitive invariantism about knowledge.

Here is a brief overview. Section 2 takes a closer look at the prodigality problem. Section 3 introduces the higher-order approach. Section 4 is concerned with how different stakes are reflected in the preferences of an agent. Section 5 presents a problem for the approach so far undertaken. Section 6 develops the idea that symmetry considerations might be the key for solving the prodigality problem. Section 7 turns this suggestion into a precise formulation of a possible decision theory. Section 8 closes with a discussion of a number of anticipated objections.

2 The Prodigality Problem

In order to set up the problem for a knowledge based decision theory, a few basic notions have to be introduced. Let us represent the agent's basic preferences by a utility function $u : \mathcal{W} \to \mathbb{R}$ from possible worlds into the reals (cp. Lewis 1981). Given two worlds w and w', an inequality like $u(w) > u(w')$ reveals that w is preferred over w'. More generally, the value $u(w)$ is supposed to represent how much the agent would like to see w actualized. To keep things simple, I assume that $\mathcal{W}$ is finite. In this case, the *expected value* (what Jeffrey (1965) calls the 'news value') of a proposition A relative to a probability function P and a body of knowledge $K \subseteq \mathcal{W}$ over the algebra generated by $\mathcal{W}$ can be defined as follows:

$$V(A) := \sum_{w \in \mathcal{W}} P(\{w\}|A \wedge K)u(w).$$

The expected value of A is a weighted average of the values $u(w)$ of worlds in A. It can be seen as an expression of how much a salient agent would welcome the news that A is true. According to *evidential decision theory* (=EDT), an agent should choose an action A with maximal news value. More precisely, if $A_1, \ldots, A_n$ are the available actions, then the agent should perform an action A_i such that $V(A_i) = \max_j V(A_j)$. The presence of K in this formula is due to the fact that P merely represents the evidential probabilities which have not yet been updated by the knowledge the agent has.

Although I shall work with this simple version of EDT, let me say that the problem to be discussed does not hinge on this choice. It applies to all broadly Bayesian decision theories including causal decision theory, for it concerns cases in which the available actions are both evidentially and causally independent of the relevant states of affairs. (With hindsight, it will become evident that it applies to any decision theory which allows the possible outcomes of an action to be restricted by the agent's knowledge.) For this reason, I shall often speak of 'decision theory' without any qualifications and I shall, where admissible, ignore the fact that EDT requires to calculate the conditional probabilities $P(X_i|A)$ and simply work with $P(X_i)$.

A consequence of an anti-skeptic insensitive invariantism about knowledge is that even in high stakes decision situations, many contingent propositions are known. If it is further assumed that knowledge is the standard for what should count as a possibility in a decision situation, then even when the stakes are high, many contingent propositions receive epistemic probability 1 in the sense relevant for decision making. Let us call this the *prodigality assumption*.[5]

Intuitively, the prodigality assumption leads to implausible verdicts about certain decision problems. Suppose an agent S knows that she owns a pair of Nikes. She is now offered a bet according to which she stands to gain \$100 if she does own a pair of Nikes but is going to lose her life otherwise. Should she

[5]The label is derived from Greco's 2013 discussion of the 'prodigality problem'. See below.

bet?

If, for the purposes of this example, S's utilities are identified with money, betting seems superior to not betting. The expected utility of betting would be (with the obvious abbreviations):

(1) $V(\text{BETTING}) = P(\text{NIKES}) \cdot 100 + P(\text{NO NIKES}) \cdot u(\text{LIFE})$
$= 1 \cdot 100 + 0 \cdot u(\text{LIFE}) = 100.$

So, the expected value of betting would be 100 (note that it is independent of how much S values her life). Not betting keeps up the status quo: no money is gained or lost. Hence, given that utility was identified with monetary value, not betting has expected utility 0. With the prodigality assumption in place, decision theory recommends taking the bet. Intuitions seem to disagree. One should not be forced to accept low-gain high-loss bets about any piece of ordinary knowledge. Call this, following Greco (2013), the *prodigality problem*.

The quick calculation of the expected value of taking the bet involved a simplification. In order for the expected value of betting to be 100, the agent needs not only to know that she owns a pair of Nikes, but she also has to know that the bookie is fully reliable. For instance, she needs to know that if she takes the bet and owns a pair of Nikes, she will be given a $100 and runs no risk of losing her life. Following Greco (2013), we may call such conditionals *bridge principles*. In addition to an ordinary piece of knowledge, the set-up of the problem requires knowledge of the bridge principles. But on a robust anti-skepticism about knowledge, there seems to be no reason to assume that such knowledge, even if occasionally harder to obtain, is generally unavailable.

3 The Higher-Order Approach

Various authors have worried that a knowledge based decision theory cannot be upheld in the light of the prodigality problem (see e.g. Kaplan 2009 and-Greco 2013). If knowledge had probability 1, one would have to bet one's life for a penny on any known proposition. But refraining from doing so does not seem irrational. Is there anything which could be said in response?

Suppose I am offered a bet of $100 against my life on a proposition p that I know, say that I own a pair of Nikes. What kind of reasoning may lead me to refuse the bet? Let us set aside potential doubts about the integrity of the bookie and a possible aversion against gambling. Then the first question I would probably ask myself is whether I actually own a pair of Nikes. I might answer "Yes" at this point having a clear memory of my pair of jogging shoes and may even conclude that I know that I own a pair of Nikes. But given that so much is at stake, I might also ask whether this knowledge is afflicted by any possible defects. I might ask myself whether I know that I know that I own a pair of Nikes. I might continue to ask these questions for some iterations of knowledge more. Should I come to the conclusion that for some iteration of knowledge, I probably fail to know, the bet may no longer feel sufficiently safe. Acknowledging the lack of knowledge at some level of iteration is a way of acknowledging the fallibility of our epistemic methods. Mere first-order knowledge is not an entitlement to absolute certainty, for what looks like knowledge from the inside can sometimes turn out to be a false belief. In the light of this fact, a bet on p need not appear risk-free even if p constitutes knowledge.

If this is roughly on the right track, then one may consider the idea that higher stakes require higher iterations of knowledge for a corresponding decision to be rational. Williamson makes a suggestion along such lines in the context of a discussion of insensitive invariantism versus sensitive invariantism or contextualism about knowledge:

> The insensitive invariantist could try to build variation in the required number of iterations of knowledge into appropriateness itself [...]: in some cases q would be appropriate iff one knew q, in others iff one knew that one knew q, and so on, depending on the stakes. (Williamson 2005a: 232)

According to such a view, knowledge would be an appropriate premise for one's practical reasoning only if the stakes are sufficiently low. If the stakes get high, one should be epistemically more careful by resting one's decisions on more secure knowledge, that is on some iteration of knowledge. To make the idea more precise, imagine an order of stakes $s_1, s_2, \ldots, s_n$, so that n iterations of knowledge are required when the stakes are of degree n.

There is a certain picture of knowledge required for such a proposal to have any plausibility. Knowing that one knows something should be substantive knowledge which does not come for free. By knowing p, one should not automatically be in a position to know that one knows that p. Thus, one has to assume frequent failures of the so-called *KK*-principle as put forward, for example, in Williamson 2000. As a matter of fact, one may even assume that for human beings the *KK*-principle fails for any proposition at some level of iteration. For no proposition, all possible iterations of knowledge are known.

A second feature of higher-order knowledge assumed here is that it does not merely add a piece of knowledge about one's own mind to the first-order knowledge. Knowing that one knows p requires more than knowledge of p and something like knowledge that one believes p. I shall assume that a certain kind of safety condition is a necessary condition for knowledge: if one knows p, then one does not falsely believe p in appropriately similar situations. Consequently, if one knows that one knows p, one does not falsely believe that one knows p in appropriately similar situations. It follows that one does not falsely believe p in all situations appropriately similar to the original class of appropriately similar situations. Hence, second-order knowledge is *safer* than first-order knowledge: if one knows that one knows, one does not falsely believe in a larger class of similar situations.

Williamson himself is somewhat skeptical about the higher-order approach to decision theory and favors a slightly different solution to the prodigality problem. Part of the reason is a more general skepticism about the prospects of a systematic decision theory:

> It is important to realize that no decision theory based on expected utility, calculated according to the standard axioms and definitions of mathematical probability theory, will be everywhere consistent with what pretheoretic common sense predicts a sensible person would do. (Williamson 2005b: 480)

In particular, he mentions the case of logical theorems and measure-one sets due to non-regular probability distributions. Both are examples of assignments of probability 1 which are in some sense unavoidable. That probability 1 is assigned to logical theorems is already required by the axioms of probability theory. But even contingent propositions are sometimes to be assigned probability 1. Think, for instance, of an infinite lottery where each ticket has the same chance of winning. Then the proposition that the first ticket does not win has probability 1 despite being contingent (for more discussion, see Williamson 2007). But one should at least sometimes be in a position to rationally turn down a bet of one's life on such a proposition or on a (moderately complex) logical theorem. If so, then just about any Bayesian decision theory will face a variant of the prodigality problem.

A second point Williamson makes is more immediately relevant to a higher-order approach to the prodigality problem. He notes that any decision theory will run occasionally into certain application problems. With respect to a knowledge based decision theory, failures of the *KK*-principle will lead to cases in which one does not know whether one has correctly applied the theory. Suppose the theory tells us to make a certain decision if we have knowledge that p. Given that knowledge is not always discriminable from false belief—we are not always in a position to know that we falsely believe when we happen to do so—we are prone to sometimes misapply the theory by acting as if we knew p although in fact we falsely believe p. In a situation of life and death, the consequences of misapplying the theory might be severe. For this reason, Williamson suggests that the risk of misapplying the decision theory needs to be taken into account in making a good decision. One should not bet one's life for a penny on a piece of knowledge, for this would open one up to bet one's life for a penny on a false belief in a relevantly similar situation. To require higher-order knowledge when a lot is at stake is a "good cognitive habit" aimed at minimizing the risk of misapplying the decision theory. It need not be part of the decision theory itself:

> Some will be tempted to build an alternative account of rationality around the 'good cognitive habits'. But such habits are too loose and contingent on the accidents of human psychology to provide a systematic decision theory. Nor do they solve the underlying problem, for they are just as vulnerable as anything else to the anti-luminosity argument. Sometimes good cognitive habits would lead one to do something although one is not in a position to know that good cognitive habits would lead one to do it. (Williamson 2005b: 483)

According to Williamson, the relation between decisions and higher-order knowledge cannot be expected to be systematic and even if it were, it would still run into the same kind of application problem. The first feature is the reason why Greco (2013) is dissatisfied with such an implementation of a knowledge based decision theory. In the present paper, I wish to show that a higher-order approach to the prodigality problem has more systematic features than one may initially expect. I will leave it open whether the account I develop should be seen as describing a decision theory proper or as revealing some hidden systematicity in the realm of the "good cognitive habits". I should perhaps say, though, that I take an association of higher risks with more epistemic care to be a very natural idea, which is why I am inclined to make it part of the decision theory proper. Let me stress that I grant Williamson's point that a higher-order approach to the prodigality problem will still face occasional application problems, but I hope less so.

4 Preferences and Higher Stakes

The prodigality problem seems to have two crucial features: (a) particularly valuable goods are at stake and (b) already excluded possibilities have some kind of comeback. If one stood to lose only a $1000, say, one would probably take a bet for a $100 on a proposition one knows to be true. It is the fact that losing the bet would come with a particularly high loss—one's life—that generates the prodigality problem. Due to the potentially devastating loss of the bet, possibilities which are excluded by one's knowledge become relevant. Given that one generally has less second-order knowledge than first-order knowledge (KKp implies Kp but not *vice versa*), considering higher iterations of knowledge does indeed open up previously excluded possibilities. As alluded to earlier, one may suspect that there is some correspondence between the stakes and the number of iterations of knowledge which is required for the decision. Is there any way to model this systematically?

Perhaps the first question is how to represent higher stakes. In standard frameworks, preferences are represented by a real-valued utility function u over the possible outcomes of the agent's actions, with higher values representing better outcomes. By itself, this does not allow to discriminate between various stakes, though it is clear that an extreme value for a possible outcome means that a lot is at stake. In order to identify individual stakes, the image of the utility function could be divided into different levels. So we may consider a partition $\mathbb{R} = E_n \cup \ldots \cup E_2 \cup E_1$ of the real numbers into a number of intervals ordered by their values, that is if $i < j$, then the values in E_i are smaller than the values in E_j. I assume $0 \in E_1$. For simplicity, I also assume that there are only finitely many levels. Note that the higher the levels get, the worse are the outcomes. Outcomes in level E_1 may be seen as the consequences of our everyday decisions. It is the lowest stake. Outcomes at higher levels may concern things like the loss of our job or our house, bankruptcy, our physical health, or the well being of those close to us.[6]

It should be stressed that the levels are as subjective as the original preferences. Different people may set the levels differently and may also have different numbers of levels. The levels in the preference structure can be seen as an expression of how much epistemic care the agent wishes to use to protect certain outcomes from happening.

It might have been noticed that levels were introduced only in one direction: from good to bad. In principle, one could also consider levels which represent particularly high goods—getting rich or winning the *Tour de France*. Partly because I am less sure about cases of this kind and partly because I would like to keep things simple, I will set this option aside. But let me briefly comment on how one could accommodate such cases. One may extend the partition of levels in the other direction by considering divisions of the form $\mathbb{R} = E_n \cup \ldots \cup E_2 \cup E_1 \cup E_2 \cup \ldots \cup E_m$. On the theory to be developed in more detail below, this would allow one to bet against something one knows but does not know that one knows, say, on the off-chance of becoming a millionaire.

Now, following up on the idea that higher stakes require greater epistemic care, let us associate each level with the corresponding iteration of knowledge. So, one would correlate E_1 with knowledge, E_2 with knowledge that one knows, and generally E_n with n iterations of knowledge. The claim would then be that a decision which concerns outcomes of level n should be based on n iterations of knowledge. For ordinary decisions, mere first-order knowledge would be enough, but when a lot is at stake, higher iterations of knowledge might be required. Let us write K^n for the body of information for which a salient agent possess n iterations of knowledge. More succinctly, the idea can then be expressed thus: if P is a probability function (the evidential probability function on Williamson's account), then a decision concerning actions with possible outcomes of level n (and no higher-level outcomes) should be made by using the probabilities $P(\bullet|K^n)$ which come from P by conditionalizing on K^n. In other words, the standard of possibility relevant to the decision would

[6] There is a point of contact with the debate about incommensurability here (see e.g. Broome 1998, 1999b). Although different levels are not incommensurable in the sense of being incomparable—one of them is always better than the other—, they share with incommensurability the general feature that outcomes may be arranged into decision-theoretically significant classes.

be K^n: what is a possible consequence of one's actions is what is compatible with K^n. For instance, when one's life is at stake, one may no longer rely on all one's first-order knowledge but only on those pieces of information for which one possesses sufficiently high iterations of knowledge.

5 How Does One Know what is at Stake?

On the picture developed so far, preference orders are given a richer structure by subdividing them into different levels. By considering the possible outcomes of one's actions, one would decide what is at stake in a given decision situation. One would then move to a corresponding iteration of knowledge and calculate the expected utilities of one's options from this perspective, i.e. by using a probability function of the form $P(\bullet|K^n)$.

Unfortunately, this picture is too simple. It presupposes that there is an independently available answer to the question of what is at stake. The problem can be illustrated by considering a few possible resolutions. Let us start with an objective, non-epistmic notion of a stake in a decision situation. Say that what is at stake is at least of level n if an outcome of level n is an objectively possible consequence of one of the agent's available actions. We do not have to decide about how exactly the notion of objective possibility should be understood in order to see that this does not give the right verdict. In an instance of the prodigality problem, the agent knows the proposition p she is offered a bet on and nothing in the prodigality problem requires that the truth value of p is not already settled at the time the decision is made. Hence, there will be instances of the prodigality problem in which the truth value of p is already settled. Consequently, it will not be an objective possibility that the undesired outcome comes about. Objectively speaking, the bet cannot be lost.

The problem persist if we consider the first-order epistemic possibilities of the salient agent. Given that the agent knows p, she knows or is at least in a position to know that taking the bet will not lead to the loss of her life. From the perspective of the agent's first-order knowledge, the bet cannot be lost either. It would seem as if there is actually nothing at stake.

Of course, the potential risk will show up for some iteration of knowledge at which the agent either does no longer know p or has lost her knowledge of the required bridge principles. For this reason, we might consider the following proposal: an outcome of level n is at stake if at some iteration of knowledge, such an outcome is a possible consequence of one of the actions available to the agent. This might give a more plausible verdict about instances of the prodigality problem, but only at the cost of having implausible consequences for almost all other cases. To see this, note that almost any piece of knowledge is lost at some level of iteration. At very high levels of iteration, one probably knows close to nothing (as mentioned earlier, it might even be that there is to any proposition and agent some level of iteration at which this proposition is no longer known). For this reason, any action has at a high level of iteration possibly very bad consequences. I might get struck by lightning if I step out of the house. I might break my neck by going down the stairs. At very high levels of iteration, one can no longer exclude such possibilities. The present proposal would recommend to evaluate almost any decision problem from an epistemic perspective at which one knows close to nothing.

In general, the problem seems to be this. The intuitions about what is at stake do not concern the objective possibilities. They are in some sense about the epistemic possibilities of the agent. However, the first-order knowledge of the agent does not always, as witnessed by the prodigality problem, provide the right standard of possibility. On the other hand, if we consider higher and higher iterations of knowledge, more and more possible consequences of one's available actions open up, so that at some point, almost any action will possibly lead to disaster. In sum, the right iteration of knowledge for determining what is at stake cannot always be the lowest nor always be a very high one. Is there some non-arbitrary way of picking something in-between?

6 Symmetry as a Guideline

We can take from the previous section that the challenge for the higher-order approach is to pair a decision problem with an appropriate iteration of knowledge in a non-circular and non-arbitrary way. The iteration of knowledge from the perspective of which the decision problem is evaluated should not be too high and should vary with the problem at hand. In ordinary decision situations, no further iterations of knowledge should have to be appealed to. Only when the stakes get high, iterations of knowledge should become relevant. The approach I shall explore suggests that symmetry considerations may help to find an appropriate iteration of knowledge. Low stake situations may be distinguished from high stake situations by recognizing that high stake situations have an asymmetric risk profile: some options come with risks visible from a corresponding level of knowledge which are unmatched by some of the alternatives. Making this notion of symmetry precise will be the main task of the present section.

Some terminology will be helpful. In the current framework, possible outcomes of an action are modeled as possible worlds (or small sets thereof which do not differ in any respect the agent cares about; cp. Lewis 1981). The preferences of an agent are defined over possible worlds, which are further subdivided into levels $E_1, E_2, \ldots$. Call a possible outcome an *n-level outcome* or an *outcome of level n* if it is situated in the preference order at level E_n. Given the agent's knowledge and its iterations $K = K^1, K^2, \ldots$, we can define a graded notion of possibility. Call an outcome o an *n-possible consequence of an available action a* if the conjunction of a and o is compatible with K^n, i.e. if performing the action a and yielding the outcome o is compatible with the body of information for which the agent possesses n iterations of knowledge. For $n = 1$, this simply means that it is an epistemic possibility for the agent that the action will result in outcome o. Derivatively, let us say that an outcome o is *n-possible* if it is an n-possible consequence of at least one of the available actions. It is important not to conflate the notions of an n-level outcome and an n-possible outcome. The former notion deals with the level of goodness or badness of the outcome, whereas the latter concerns the question what the agent knows about the possible consequences of her available options.

To get a first grip on how symmetry considerations might come into play, let us take a closer look at the initial example giving rise to the prodigality problem. I know that I own a pair of Nikes and I also know the relevant bridge principles, i.e. that if I bet and win, I will be given \$100 and if I I bet and lose, I will lose my life. To fix ideas, suppose that protecting my life is very important to me, so I might assign it something like level 5 intending to base decisions where my life is at stake only on information for which I have five iterations of knowledge. In the terminology introduced above, losing my life is a 5-level outcome. As far as possibilities are concerned, losing my life is not 1-possible, for it is incompatible with my knowledge that I am going to take the bet and lose. But things might be different when it comes to five iterations of knowledge. Perhaps I do not have five iterations of knowledge for owning a pair of Nikes. More likely even is that I do not have five iterations of knowledge for the relevant bridge principles. Be that as it may, let us assume that losing my life, an outcome of level five, is a 5-possible consequence of taking the bet. What about not taking the bet? Here nothing happens and we can assume that even at five iterations of knowledge it does not become possible that I will lose my life. Thus, from the perspective of five iterations of knowledge, the number of iterations which is associated with the outcome of losing my life, the risk profile of the two available actions is *asymmetric*: taking the bet comes with a high level risk whereas not taking the bet does not.

Compare this with a case in which I have to choose between two paths taking me through dangerous territory. I have been told by reliable sources that both of them are safe, so we may assume that I know that they are safe. But we may also assume that my knowledge is not particularly robust: already after a few iterations, I no longer know that the two routes are safe. As a result, losing my life is an 5-possible outcome. In this case, however, the risk profile of the two available actions is symmetric: a 5-level outcome is a 5-possible consequence of both of the actions. Thus, even when we consider iterations of knowledge, none of the two actions reveals a hidden risk not shared by the other. It is tempting to conclude that this might be the reason why I can make my decision in this case from the perspective of my first-order knowledge, while I should move to a higher iteration of knowledge in the former case illustrating the prodigality problem due to the fact that it reveals an asymmetry in the risks the available actions come with.

Let us see how the idea of risk symmetry can be made more precise. This is a somewhat more complicated task than one might initially expect, for one has to simultaneously keep track of the levels of the possible outcomes in the preference order and of the unfolding hierarchy of iterations of knowledge. To cope with the interplay of these two dimensions, it is helpful to start by focussing on a fixed level m of outcomes and a fixed iteration n of knowledge (the previous examples assumed $m = n = 5$). If we take a clue from the example of the two paths, the risk profile of the two available actions seems symmetric from the fifth iteration of knowledge because both actions have a level-5 outcome which is an equally probable consequence of these actions. There is a one-to-one mapping between the level-5 outcomes of the two actions which preserves probability. In general, the situation may be a little more complex, though. The available actions may have more than one possible outcome of the relevant level, the number of such outcomes may also be different for various actions, and these outcomes may not always be equally likely consequences of the actions. Still, if actions with more problematic outcomes lead to them with a lower probability, they could be counted as coming essentially with the same risk. If this is correct, the notion we need is one of *expected risk*. We may define such a notion for an action A, relative to a level of outcome m and a number n of iterations of knowledge, thus:

$$EV(A, m, n) := \sum_{w \in E_m} P(w|A \wedge K^n) \cdot u(w).$$

This mirrors the definition of expected utility with the difference that it considers only outcomes of a certain level and evaluates the probabilities based on possibly higher-order knowledge. With this definition in place, we can then say that a family of available actions A_i are *symmetrical relative to outcomes of level m and evaluated from the perspective of n iterations of knowledge*, for short (m, n)*-symmetrical*, if $EV(A_i, m, n) = EV(A_j, m, n)$ for all i, j.

For instance, the paths example is $(5, 5)$-symmetrical, whereas the problematic case illustrating the prodigality problem is one where the available actions are not $(5, 5)$-symmetrical: as I described the case, losing my life becomes possible from the higher-order perspective and so $EV(\text{BET}, 5, 5)$ will be negative while $EV(\text{NOT BET}, 5, 5)$ remains zero. The asymmetry in the latter case only shows up for a higher iteration of knowledge: the two available actions are still $(5, 1)$-symmetrical. From the perspective of first-order knowledge, none of them comes with any high-level risk.

In order to make progress in finding the right epistemic perspective from which a given decision should be made, we should also focus on the interplay between the level of an outcome and higher degrees of knowledge. Consider, for instance, a hypothetical situation in which the available actions are not $(2, 3)$-symmetrical. Should this be a reason to count three iterations of knowledge as being relevant for making the decision? If we go back to the initial proposal of how levels of outcomes might be connected with iterations of knowledge, the suggested answer is negative, because the levels of outcomes were introduced with the idea that they indicate how much epistemic care should be invested into decisions where corresponding outcomes are at risk. A level 2 outcome is one for which the agent has decided that decisions where it is at stake should be epistemically protected by two layers of knowledge. But this means that an asymmetry showing up at three iterations of knowledge should not count as relevant, for this is more epistemic protection than the agent takes to be necessary. Consider, for instance, a decision situation where the agent does not have three iterations of knowledge guaranteeing that one of the actions will not lead to a loss of her job, while from the perspective of two levels of knowledge losing her job is not a possibility. If job loss is a level 2 outcome for her, then according to the present line of thought, the asymmetry for three iterations of knowledge should not force her to move to this level of knowledge in order to make her decision.

Generalizing, one may consider the thought that (m, n)-asymmetry is only a factor if $m \geq n$. The idea would be that an asymmetry at a certain layer of knowledge should only be taken into account if it concerns outcomes the layer of knowledge is supposed to protect. More precisely, we may call a layer n of knowledge *symmetrical* if it is (m, n)-symmetrical for all $m \geq n$. Going back to the examples, it may be the fact that the fifth layer of knowledge is not symmetrical in the prodigality case which makes us steer away from betting, while in the paths example the symmetry at higher iterations is the reason why we should base the decision on our first-order knowledge.

As defined, symmetry from the perspective of a certain layer of knowledge is a restricted kind of indifference. In the case of first-order knowledge, it actually coincides with indifference. To see this, recall that the first level is a symmetrical layer of knowledge if it is $(m, 1)$-symmetrical for all $m \geq 1$. But the latter condition is satisfied by all m, and so the first level being symmetrical means that it is symmetrical regarding all levels of outcomes. But now note that the expected value of an action is just the sum of the expected values for the various levels of outcomes, which derives from the fact that they form a partition:

(2) $$V(A) = \sum_w P(w|A \wedge K^1) \cdot u(w)$$
$$= \sum_m \sum_{w \in E_m} P(w|A \wedge K^1) \cdot u(w) = \sum_m EV(A, m, 1).$$

Hence, if two actions A_1 and A_2 are such that $EV(A_1, m, 1) = EV(A_2, m, 1)$ for all m, they will have the same utility, i.e. $V(A_1) = V(A_2)$.

Symmetry for higher iterations of knowledge will no longer imply this, however, for symmetry is defined so that it only "looks up" but not "down". Still, a symmetrical layer of knowledge comes with a limited recommendation of indifference: as far as higher-level goods are concerned, there is no reason to prefer one of the actions over the others. Each of the actions comes with the same amount of risk as far as higher-level goods are concerned (to be precise: goods which belong to a level at least as high as the layer of knowledge one is looking at).

This observation might help to narrow down the situations in which a decision should be made from a higher-order perspective of knowledge. This is an important constraint, for a decision made from a higher-order perspective is a decision which avails itself of less information than is available at a lower level of knowledge, for one generally has more first-order knowledge than second-order knowledge, and so on. Moving to a higher layer of knowledge thus comes with a certain price: one is going to make a decision which will, objectively speaking, tend to be suboptimal. Of course, one also gains something in return: more epistemic care for goods which are particularly valuable. But one should only pay the price of moving to a higher layer of knowledge when it is really necessary. Here is where symmetry might come in: a symmetrical layer of knowledge tells us that the goods it is supposed to protect are unaffected by the choice we make, which means that we can go down to a lower iteration of knowledge and make our decision on the basis of a larger, though somewhat less secure information base.

As a consequence, this suggests that the right epistemic perspective from which a decision should be approached is never a symmetrical layer of knowledge unless all levels are symmetrical including the first one, in which case we can be indifferent. For one can always improve one's decision by going down at least one level without ignoring any kind of hidden risks. Hence, an intermediate conclusion is this: *a decision should be made from an asymmetrical level of knowledge.* The question thus simplifies to: which asymmetrical layer of knowledge should guide our decisions?

If the levels in one's preference order are finite as assumed, then there will always be a highest asymmetrical layer of knowledge. Take m to be the highest preference level and suppose $n > m$. Then a layer of n iterations of knowledge is trivially symmetrical, for there are no outcomes of a level at least as high as n, for which an asymmetry in risks among the actions could then be revealed.

The highest asymmetrical layer reveals an asymmetry in risks for one's available actions concerning a level of outcome in need of more epistemic care than any of the outcomes at lower levels. It is therefore a natural candidate for being the perspective from which one's decision should be made. Were one to make it at a lower level, one might ignore the risk revealed at the higher layer of knowledge which was chosen to protect the goods of the corresponding level. On the other hand, it does not seem necessary to go higher than the highest asymmetrical level, for as we have seen, symmetrical layers of knowledge always suggest to go down a layer. Thus, this is the proposal I would like to make: *the right epistemic perspective from which to evaluate a decision problem is the highest asymmetrical level of knowledge.*

7 A Decision Theory

The previous section contains a sketch of a possible decision theory which connects iterations of knowledge with higher-level goods. What is still missing is a clear statement of the formal structure of the theory. To this end, it proves helpful to start with a description of a *decision problem* as it is conceived of within such a theory. A decision problem can be identified with an ordered tuple $D = (W, P, (K^n)_{n\in\mathbb{N}}, u, (E_m)_{m=1,2,\ldots,k}, (A_i)_{i=1,2,\ldots,j})$. W is a set of possible worlds—the possible outcomes—, which is assumed to be finite for the sake of simplicity. P is a probability function defined over W (to be fully precise: it is defined over the algebra generated by the power set of W). The function P represents the evidential probabilities. As far as its role in the theory is concerned, it can be compared to an initial probability function in standard frameworks of Bayesianism. $(K^n)_{n\in\mathbb{N}}$ is a family of nested non-empty sets of possible worlds $K_1 \subseteq K_2 \subseteq \ldots$, where K^n represents the information for which the salient agent possesses n iterations of knowledge. As defined earlier, u is a utility function from the set W of possible worlds into the reals and $(E_m)_{m\in\mathbb{N}}$ is the ordered (and finite) partition of W into different levels. Finally, the family $A_1, A_2, \ldots, A_j$ is a finite partition of W representing the actions available to the agent.

The evidence an agent has for a proposition will be taken to be the probability conditional on the agent's knowledge, i.e. as given by $P^1(\bullet) := P(\bullet|K^1)$. More generally, let us define $P^n(\bullet) := P(\bullet|K^n)$. The probabilities P^n represent the likelihoods of propositions from the perspective of n iterations of knowledge. Moreover, given a decision problem D, let $f(D)$ be the highest asymmetrical layer of knowledge. If all iterations are symmetrical, set $f(D) := 1$ (as explained earlier, this means that the agent can be indifferent). With this in place, we can now calculate the expected utility from the perspective of the asymmetrical layer of knowledge, to which we may refer as the *level adjusted expected utility*:

$$V^*(A) := \sum_{w \in W} P^{f(D)}(w|A) \cdot u(w).$$

The claim would then be that the agent should choose an action A_i for which $V^*(A)$ is maximal.

The form of the definition for expected utility is the same as in standard explications of (evidential) decision theory. So rational agents are still taken to be expected utility maximizers. What is added to standard decision theory is a claim about what the standard for possibility is relative to a given decision problem: sometimes it is simply knowledge, at other times a layer of higher-order knowledge. This means that the probabilities figuring in the calculation of expected utility may vary from decision problem to decision problem depending on how much is at stake. This is the crucial (but only) difference to standard decision theories. In ordinary situations in which all higher levels are symmetrical, so that the highest asymmetrical level of knowledge is simply 1, the proposal converges to the standard theory as outlined earlier, i.e. one finds $V^*(A) = V(A)$.

There is a special case for which the theory can be improved a little further. Suppose $f(D) > 1$, i.e. the decision problem is approached from a higher layer of knowledge. Assume in addition that there is more than one action A for which $V^*(A)$ is maximal. Typically, decision theory recommends in such cases to be indifferent between these options. But in the present kind of case, there might still be a reason to prefer one of the options over the others, for one can avail oneself of more information by going down a layer of knowledge. With more information at hand, there could be a reason to favor one of the options.

So, the idea would be that one calculates

$$V^{-1}(A) := \sum_{w \in W} P^{f(D)-1}(w|A) \cdot u(w)$$

for those actions with maximal $V^*(A)$ and choose one with maximal $V^{-1}(A)$. This might filter out some actions, but it could still leave one with more than one. If one has not already reached the level of first-order knowledge, one may simply repeat the procedure. So one would successively avail oneself of more and more knowledge to see whether one can find a reason to prefer one of the options with maximal expected utility at the highest asymmetrical layer of knowledge. The procedure terminates if either one has reached level 1 or there is only one option left with maximal expected utility. If level 1 is reached with more than one action having maximal utility, we have a real case of indifference. Starting with the set of actions $\mathcal{A}_0^*$ of those actions A for which $V^*(A)$ is maximal, this defines a finite nested sequence $\mathcal{A}_0^* \supseteq \mathcal{A}_1^* \supseteq \mathcal{A}_2^* \supseteq \ldots \supseteq \mathcal{A}_j^*$. The more refined proposal then reads: an agent should pick an action out of $\mathcal{A}_j^*$. This will be an action A for which $V^*(A)$ is maximal, but it may have the additional advantage of having greater expected utility from the perspective of a more informative layer of knowledge.

8 Objections and Replies

There are still a number of questions open and a lot of further issues to be addressed. I conclude with a discussion of some anticipated objections, with a focus on those which strike me to be the most immediate ones.

1. *Does it really solve the prodigality problem?* The worry here might be that the theory does not predict that one should turn down the bet under all circumstances. Suppose we have sufficiently many iterations of knowledge for the proposition to be bet upon. Then the expected value of betting will be the same as in the original problem for first-order knowledge and the theory recommends taking the bet. And even if we do not have enough iterations of knowledge, there will still be cases in which the probability of the proposition is so extremely high that the expected value of betting is higher than the expected value of not betting. So, the theory seems to solve the prodigality problem only for a limited range of cases.

I disagree. To begin with, the prodigality problem was not that a bet of \$100 against our life should never be taken. The problem was that one should not *always* be forced to take the bet if the target proposition is known (Greco (2013) makes the same point). This latter problem is solved by the theory. But perhaps the former still is a problem even if it is not to be identified with the prodigality problem. This is an interesting question, but I am inclined to think that it has a negative answer. If one has enough iterations of knowledge, the bet will be very safe, much safer than if one possessed only first-order knowledge. Taking this observation from the realm of betting to more ordinary decision situations, this means that one can ignore certain potential risks if one has strong enough evidence to suggest that they are not actual. Ordinarily, we do not take into consideration outlandish possibilities ruled out by information for which we have sufficiently many iterations of knowledge, like our partner being the much wanted ax murderer or our house falling apart by tomorrow. But even a stronger conclusion seems to be justified. In practice, we sometimes (not always) seem to tolerate a very small probability of a fairly big loss for a moderate gain. On the present theory, this kind of tolerance is appropriate if the probability is sufficiently small when evaluated from a sufficiently high level of knowledge. This means that the low probability will be particularly resilient because it is based on a body of information for which we possess a high number of iterations of knowledge.

To further defend this consequence, it should also be noted that any Bayesian decision theory according to which agents should maximize expected utility has the consequence that an action with a particular bad possible outcome should sometimes be favored when the probability of the outcome is sufficiently small.[7] If one is unhappy about this result, one would probably have to give up utility maximization and add an element to decision theory which ranks actions without particularly bad consequences over actions with such consequences independently of the respective probabilities. In principle, it seems possible to add such an element to the present theory by placing those

[7] The reason is that Bayesian decision theories typically validate a version of the *continuity axiom*. This axiom states: If A is preferred to B and B to C, then there are probabilities $p, q \in (0, 1)$ such that ApC is preferred to B and B to AqC, where for any probability r, ArC means the prospect of getting A with probability r and C with probability $1 - r$. So, the axiom requires that our preferences for prospects do not change if one adds possible outcomes with very small probabilities, even if the outcome is arbitrarily bad like the loss of one's life.

actions which possibly lead, from the perspective of the relevant iteration of knowledge, to high level outcomes below actions whose worst outcome belongs to a lower level in the preference ordering.[8]

2. *Doesn't ignoring our first-order knowledge by moving to a higher iterations of knowledge make our decisions worse, not better?* This is an objection which comes very naturally. The set of propositions for which we have higher iterations of knowledge will in general be smaller than the corresponding set of first-order knowledge. But this means, as remarked earlier, that availing ourselves only of higher-order knowledge effectively ignores some information we have. And shouldn't we bring to a decision situation, particularly when highly valuable goods are at stake, all the knowledge we have?

There is a sense in which the objection is valid. Objectively speaking, decisions based on first-order knowledge will tend to be better than decisions based on higher-order knowledge. This is a consequence of the factivity of knowledge: known propositions are true, and so a decision based on first-order knowledge avails itself of more true information about the world. But this is not necessarily the sense of a decision being good usually targeted in decision theory. A decision would be perfect if based on all (relevant) truths, but a decision theory merely making this recommendation would be of little guidance. This is why decision theory typically works with information the agent has epistemic access to. But regarding our epistemic access, it is a non-trivial questions which epistemic perspective is appropriate.

Still, the appeal to higher-order knowledge instead of first-order knowledge is in need of some justification. The best way I can think of is in comparison with an insurance fee. Here we pay a small price to protect ourselves against particularly bad events. In a similar way, we can think of retreating to a higher level of knowledge as a way of making our decisions safer. The price we pay is that our decisions will become a little less than optimal overall compared to how they would have been had they been based on first-order knowledge only.

It should also be noted that our first-order knowledge will often leave traces still visible from the perspective of higher-order knowledge. Suppose

[8] Effectively, this would amount to adopt what is sometimes called the *leximin rule* with respect to the level of outcome the actions have.

I see a skyscraper from a certain distance close enough to know that it is between 70 and 120 yards high. However, suppose further that I am not close enough to know that I know that it is 70–120 yards high. Still, from the perspective of my second-order knowledge, it will still be highly likely that the skyscraper is 70–120 yards high. What is certain from the perspective of knowledge is very often at least probable from the perspective of higher-order knowledge. Although information might be lost at higher levels, it can still have an impact on the respective probabilities. Exceptions to this general rule are cases of inferential knowledge based on multiple known premises.[9] If the premises cease to be known on a higher level, the probability of the inferred conclusion may drop radically due to the accumulating possibility of falsity in the set of premises. But this is not an unwelcome result: it accords well with pre-theoretic intuitions to trust information which is inferred from multiple premises significantly less in decision situations.

3. Is the theory really any better than a more straightforward decision theory based on first-order knowledge only? Such a decision theory would construe the prodigality problem as an application problem, as suggested by Williamson (2005a). Although taking the bet would be predicted by the theory to be the right decision, the problem of knowing that the conditions of its proper application are satisfied becomes more and more relevant the higher the stakes get. As I said at the outset, the present theory can be reinterpreted along such lines. The idea would be that symmetry considerations guide us in identifying the potential risk of a decision and thereby constrain how cautiously we should proceed in applying the theory. Taken this way, the theory would primarily be a theory of the correct application of a straight knowledge based decision theory.

Independently of whether the theory is taken as a decision theory proper or merely as an afterthought about the correct application of a more simple basic theory, it is open to the objection, voiced by Williamson (2005a) and mentioned earlier, that it faces a similar kind of application problem. Just like one may not know that one knows when trying to apply the simple theory, one will not always be in a position to know whether a given iteration of knowledge is symmetrical. When taken as a decision theory proper, there will still be cases in which one is not in a position to know whether one has correctly applied the

[9]Many thanks to Martin Smith for helping me getting clear about this point.

theory. When taken as a theory of application, on the other hand, these cases are such that one is in no position to know whether one has correctly applied what is supposed to be the theory of application. On either of the two possible construals, one may wonder whether progress has been made.

I agree that one cannot always know whether one has correctly applied the proposed theory. I also agree with Williamson (2005a) that virtually any decision theory is going to face a problem of this kind, for whichever conditions a theory imposes on good decisions, it is very likely that there will be cases in which one cannot know whether these conditions are satisfied. If this is correct, having a decision theory with transparent application conditions is an unobtainable goal. But this does not mean that no progress can be made. What I would like to suggest is that the present theory is more informative and thereby makes more explicit what should guide us when making high stakes decisions. It describes what kind of considerations—primarily considerations of symmetry—one should be sensitive to in making such decisions. This does not mean and should not be taken to require that it provides a fail-proof decision procedure.

4. Higher-order knowledge is partly about our own minds. Why should such knowledge matter? If knowledge is a state of mind, then knowing that one knows is knowing that one is in a certain state of mind. If knowledge implies belief, then knowing that one knows puts one in a position to know that one is in a certain belief state. Why should such knowledge be relevant for decisions which usually have to do only with what happens in the external world?

In response, one might say that thinking about what one knows can in many cases be naturally associated with the phenomenology of high stakes decisions. It seems reasonable to reflect on whether one knows a certain proposition in situations in which a lot depends on its truth. Although a decision concerns the consequences of our actions in the world, it is still our epistemic condition which partly determines which decision is best.

Alternatively, it is possible to grant the objection. However, this does not mean that the core of the theory has to be given up. Recall that the formal details of the theory only require a sequence $K^1 \subseteq K^2 \subseteq \ldots$ of sets of possible worlds. These can be taken to represent the information for which one has knowledge, knowledge that one knows, etc. but this is just one out of various possible interpretations. The sequence of sets can also be taken to represent information for which one has knowledge to which stronger and stronger safety conditions apply. In more traditional terms, it would be knowledge for which one has justification of stronger and stronger kinds (for more, see also the next discussion point).

5. Can the theory also be formulated in terms of belief? Some people may think that a decision theory should not be formulated in terms of knowledge but rather in terms of belief. Bayesian decision theories are usually cast in terms of subjective probabilities, which are degrees of (partial) belief. Will something similar be possible with the present theory?

The discussion of the previous objection already points to a positive answer. It is also possible to interpret the sequence $K^1 \subseteq K^2 \subseteq \ldots$ of sets of possible worlds as pertaining to information one only believes to be true. The first level could be described as outright belief and higher levels as higher degrees of outright belief (cp. the discussion of degrees of outright belief in Williamson 2005a).[10] Degrees of outright belief would have to be sharply distinguished from degrees of partial belief. That a coin comes up heads may deserve a credence of 1/2, but its degree of outright belief would be 0. Structurally, degrees of outright belief are similar to iterations of higher-order knowledge, but they do not amount to mere iterations of higher-order belief. They are outright first-order beliefs the subject takes to be justified with variable levels of strength.

6. Introducing levels into the preference order postulates arbitrary cutoff points. Doesn't that force undesirable discontinuities in our decision behavior?[11] Moving an outcome at the border of two levels from one level to the next can make for a fairly big difference in certain decision situations. Although the utility of the outcomes has only changed slightly, a further layer of knowledge suddenly becomes relevant.

In response, let me point out that modeling the preference of an agent in terms of a single utility function defined over all possible worlds is generally acknowledged to gloss over various ways in which our preferences may be partial, imprecise, vague or indeterminate. In the same way, the sharp cutoff points of the introduced levels can be seen as glossing over the fact that

[10] Degrees of outright belief have been extensively studied in the framework of ranking theory. See, for instance, Spohn 2012.

[11] Thanks to Andy Egan and Alexander Steinberg for pressing me on this.

real-world agents' preferences may not be that precise. It may be a vague or indeterminate matter where one level ends and another begins.

This might help to soften the consequences of introducing levels into the preference order of an agent. I should admit, though, that the theory is committed to certain discontinuities when a lower stakes decision is compared to a higher stakes decision. This is an immediate consequence of the very point of the theory: its core idea is to approach the prodigality problem in terms of a flexibility in the epistemic perspective one applies to a decision problem. It remains to be seen whether the envisaged discontinuities actually lead to any problematic predictions concerning concrete examples, where the problematic features cannot be accounted for by the inevitable imprecision (or vagueness or indeterminacy) which besets the exact location of the cutoff points.[12]

7. Couldn't a dutch book be made against an agent following the proposed decision theory? The reason why one may suspect that something like that could be possible is that the theory sometimes recommends to evaluate two decision situations in terms of different probability functions even though the agent possesses exactly the same evidence. A clever bookie knowing this might be able to design a series of bets which lead to a sure loss by exploiting the difference between the probability functions. The change in probability functions could be assimilated to a belief change unprompted by any new evidence to which the standard diachronic dutch book arguments for conditionalization would apply (described and attributed to David Lewis in Teller 1973). Let me emphasize that this is only a sketch of a possible counterargument. But there is some evidence that such a strategy could be carried out. Greco (2013) has constructed a dutch book against sensitive invariantist solutions of the prodigality problem, according to which high stakes raise the standards for knowledge.

Now, it is a controversial issue to what extent unexploitability by a clever bookie constrains rationality. For instance, Williamson (2000: ch. 10) suggests that rational agents sometimes cannot avoid being exploitable.[13] As an in-

[12]Ralph Busse also suggested to me in discussion that it might be an option to drop the discrete order of levels and rewrite the theory in terms of a continuous version of ranking theory. Whether this idea can be brought to work and which modifications would have to be made is something I have to leave for further research.

[13]In a similar vein, Douven (2013) has recently argued that unexploitability by a dutch book may just be one out of many legitimate and possibly competing goals an agent might have, which can be outweighed by faster convergence to truth, for example.

depth discussion of dutch book arguments is beyond the scope of this paper, I will only put forward one consideration which is more specifically concerned with the proposed decision theory. It has already been acknowledged that decisions made from a higher-order perspective are in a certain sense overall slightly less than optimal. The agent pays a certain price in order to make safer choices. But if it is agreed that an increase in safety is worth a certain price, then it may simply be that the exploiting bookie collects a price worth paying. From the point of view of the agent, being offered a high stakes bet would not be a neutral event. The world has subjectively speaking become more dangerous, for a wrong decision could have disastrous consequences (the agent could lose her life). When the agent now pays a certain price for protecting herself against making the wrong decision by passing, for instance, on a sure $100, she may not act less rationally than someone who buys insurance against robbery after having been moved to an unsafe neighborhood.

9 Conclusion

On an anti-skeptical insensitive invariantism about knowledge, a knowledge-based decision theory faces the prodigality problem: due to the fact that plenty of propositions are assigned probability 1, an agent has to positively evaluate a bet of her life on the truth of any such proposition for an arbitrarily small reward. In order to solve this problem, the paper develops the idea that high level goods might require higher-order knowledge when they are at stake in a decision situation. It proves crucial to invoke considerations of symmetry in fixing which level of higher-order knowledge is relevant for a given decision situation. Various issues are still open. In particular, a detailed comparison of the present account with alternative solutions of the problem, which would either reject an anti-skeptical insensitive invariantism about knowledge or else would not construct decision theory around the concept of knowledge, is still missing. But given that it raises much more general questions in epistemology and decision theory, it has to wait for another occasion.[14]

[14] **Acknowledgements.** Earlier drafts of this paper have been presented at the *Institutskolloquium* in Hamburg, a workshop in Mainz, the *PERSP Final Conference* in Barcelona and a workshop on norms of reasoning in Bochum, all 2014. I would like to thank all the discussants for their valuable comments. Various issues raised in section 8 were prompted by their questions. Thanks are also due to an anonymous referee of this journal.

References

Broome, J. 1998: *Is Incommensurability Vagueness?* In Broome 1999a.

—— 1999a: *Ethics out of Economics*. Cambridge: Cambridge University Press.

—— 1999b: *Incommensurable Values*. In Broome 1999a.

Douven, I. 2013: 'Inference to the Best Explanation, Dutch Bookds, and Inaccuracy Minimisation'. *The Philosophical Quarterly* 63, pp. 428–44.

Fantl, J. and M. McGrath 2002: 'Evidence, Pragmatics, and Justification'. *The Philosophical Review* CXI, pp. 67–94.

Greco, D. 2013: 'Probability and Prodigality'. In Gendler, T. S. (ed.): *Oxford Studies in Epistemology*, vol. 4, Oxford University Press, pp. 82–107.

Hawthorne, J. and J. Stanley 2008: 'Knowledge and Action'. *The Journal of Philosophy* 105, pp. 571–90.

Hyman, J. 1999: 'How Knowledge Works'. *Philosophical Quarterly* XLIX, pp. 433–51.

Jeffrey, R. 1965: *The Logic of Decision*. Chicago: University of Chicago Press.

Kaplan, M. 2009: 'Williamson's Casual Approach to Probabilism'. In Greenough, P. (ed.): *Williamson on Knowledge*, chap. 8, Oxford University Press, pp. 122–139.

Levi, I. 1980: *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.

Lewis, D. 1981: 'Causal Decision Theory'. *Australasian Journal of Philosophy* 59, pp. 5–30.

Spohn, W. 2012: *The Laws of Belief*. Oxford University Press.

Teller, P. 1973: 'Conditionalization and Observation'. *Synthese* 26, pp. 218–58.

Unger, P. 1975: *Ignorance: A Case for Scepticism*. New York: Oxford University Press.

Williamson, T. 2000: *Knowledge and Its Limits*. Oxford: Oxford University Press.

—— 2005a: 'Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge'. *The Philosophical Quarterly* 55, pp. 213–35.

—— 2005b: 'Replies to Commentators'. *Philosophy and Phenomenological Research* 70, pp. 468–491.

—— 2007: 'How Probable is an Infinite Sequence of Heads?' *Analysis* 67, pp. 173–180.