Philosophia Scientiae

Travaux d'histoire et de philosophie des sciences

19-1 | 2015

Logic and Philosophy of Science in Nancy (II)

Decisions without Sharp Probabilities

Paul Weirich

Electronic version

URL: <http://journals.openedition.org/philosophiascientiae/1062>

DOI: 10.4000/philosophiascientiae.1062

ISSN: 1775-4283

Publisher

Éditions Kimé

Printed version

Date of publication: 1 March 2015

Number of pages: 213-225

ISBN: 978-2-84174-707-8

ISSN: 1281-2463

Electronic reference

Paul Weirich, « Decisions without Sharp Probabilities », *Philosophia Scientiae* [Online], 19-1 | 2015, Online since 23 April 2015, connection on 03 November 2020. URL : <http://journals.openedition.org/philosophiascientiae/1062> ; DOI : <https://doi.org/10.4000/philosophiascientiae.1062>

Tous droits réservés

Decisions without Sharp Probabilities

Paul Weirich

University of Missouri (USA)

Résumé : Adam Elga [Elga 2010] fait valoir qu'aucun principe de rationalité ne mène de probabilités imprécises à des prises de décisions. Il conclut qu'un agent parfaitement rationnel n'a pas de probabilités imprécises. Cet article défend les probabilités imprécises. Il montre comment les probabilités imprécises peuvent justifier des décisions rationnelles.

Abstract: Adam Elga [Elga 2010] argues that no principle of rationality leads from unsharp probabilities to decisions. He concludes that a perfectly rational agent does not have unsharp probabilities. This paper defends unsharp probabilities. It shows how unsharp probabilities may ground rational decisions.

Unsharp probabilities arise from sparse or unspecific evidence. For example, meteorological evidence, because unspecific, often does not suggest a sharp probability that tomorrow will bring rain. An agent may assign to rain a range of probabilities going from, say, 0.4 to 0.6. A. Elga argues that unsharp probability assignments may lead an agent to forgo a sure gain [Elga 2010]. In this event, a dilemma arises: the agent may have either unsharp probability assignments that accurately represent evidence, or sharp probabilities that prevent forgoing a sure gain. Should an agent's probability assignments be faithful to the evidence, or should they promote practical success? This paper maintains that an agent's probability assignments can attain both goals. It defends a principle of choice that uses imprecise probabilities.

1 Arbitrage

Elga's argument against sharp probabilities begins with a case involving a sequence of choices. In the case, H is a hypothesis. Sally's evidence concerning

H is scant, and she therefore assigns to H the probability interval $[0.1, 0.8]$. Gamble A wins \$15 if H does not hold and loses \$10 if H holds. Gamble B wins \$15 if H holds and loses \$10 if H does not hold. Sally knows that she will receive an offer of A and then an offer of B . She may accept or reject A and afterwards accept or reject B . She also knows that her evidence concerning H will not change during the sequence of offers. The two offers give Sally an opportunity for arbitrage, that is, an opportunity to make a sequence of transactions that ensures a gain; accepting both gambles ensures a gain of \$5. The following payoff table shows Sally's net gain from the two gambles if H is true and if it is false.

	H	$\sim H$
A	-10	15
B	15	-10
Net	5	5

According to a permissive decision principle that Good proposes, if a set of pairs of probability and utility assignments represent an agent's mental state, then in a decision problem the agent may choose any option that maximizes expected utility according to some pair of probability and utility assignments in the set [Good 1952, 114]. Amounts of money settle Sally's utility assignment. So according to the principle, Sally may, when offered the gambles A and B , choose any option that maximizes expected utility according to some probability in the interval that she assigns to H . In particular, the principle may calculate a gamble's expected utility using either of the interval's two endpoints. Suppose that for A , a calculation uses the interval's upper endpoint, whereas for B it uses the interval's lower endpoint. Using P to stand for probability and EU to stand for expected utility, the results are:

$$(1) \quad \begin{array}{l} \text{If } P(H) = 0.8, \\ \text{then } EU(A) = (0.8 \times -10) + (0.2 \times 15) = -5, \text{ and } EU(\sim A) = 0. \end{array}$$

$$(2) \quad \begin{array}{l} \text{If } P(H) = 0.1, \\ \text{then } EU(B) = (0.1 \times 15) + (0.9 \times -10) = -7.5, \text{ and } EU(\sim B) = 0. \end{array}$$

According to Good's principle, Sally may first reject A because this option maximizes expected utility according to (1) and then reject B because this option maximizes expected utility according to (2). However, if she rejects both gambles, she wastes her opportunity for arbitrage. Sally cares about money only, and so if rational does not reject both gambles. The permissive principle is wrong to suggest otherwise, Elga argues.

A decision principle that is *strict* in Elga's terminology goes from unsharp probabilities to decisions using sharp representatives of unsharp probabilities. For example, the principle MIDPOINT uses the midpoint of a probability interval to reach a decision. Elga rejects such principles because they depart

from the rationale for unsharp probabilities. They treat a probability interval as a sharp probability, such as the interval's midpoint. An acceptable principle of rational choice has to be both permissive about choices taken one by one and also strict about sequences of choices.

After rejecting strict decision principles, Elga examines three principles going from unsharp probabilities to decisions that prohibit rejecting both *A* and *B* but allow rejecting *B* in isolation. Although the principles achieve the correct mixture of permissions and obligations, they are defective. The principle NARROW adjusts probability intervals in light of past choices to prevent sure losses in sequences of choices. However, no change in evidence grounds the adjustments in probability intervals, so the adjustments are unwarranted. The principle PLAN states that an agent who begins a sequence of choices should plan coherent choices for the sequence and should then stick to the plan. If Sally rejects *A*, planning to accept *B*, then she should accept *B* after rejecting *A*. However, as Elga observes, rejecting *B* after rejecting *A* has the same monetary consequences as rejecting *B* when it is offered in isolation. He contends that because Sally cares about money only, if she may reject *B* offered in isolation, she may also reject *B* when it is offered after rejecting *A* even if she planned to accept *B*. The principle SEQUENCE says that a sequence of choices may be irrational despite the rationality of each choice in the sequence. It concedes that for Sally either response to each offer is rational but maintains that rejecting both *A* and *B* is irrational. Elga claims that SEQUENCE errs because it permits Sally to reject *B* offered in isolation, but prohibits her rejecting *B* if it is part of a sequence of choices that includes prior rejection of *A*; Sally cares about money only, and rejecting *B* has the same monetary consequences in both situations.

Elga's argument against decision principles using unsharp probabilities relies on the case of Sally, but this case leaves some important features unsettled. First, Sally's goals are unclear. Elga specifies that for Sally the utility of money is linear and that Sally cares only about money. The first assumption simplifies calculations of expected utility. The second assumption grounds Elga's objections to PLAN and to SEQUENCE. I interpret the second assumption so that it attributes to Sally, besides any goals that rationality itself may require, the single basic goal of gaining money and derivative goals concerning means of and opportunities for gaining money.

Second, the example does not specify whether Sally decides without forming preferences between her options, or forms preferences and then acts on them. Suppose that she forms preferences that yield her choices and that she rejects both *A* and *B*. To represent her preferences, take the symbol for a gamble to stand for having the gamble and the tilde before the symbol for a gamble to stand for not having the gamble. Because Sally rejects *A*, she prefers some $\sim A$ -result, either $(\sim A \ \& \ B)$ or $(\sim A \ \& \ \sim B)$, to all *A*-results, both $(A \ \& \ B)$ and $(A \ \& \ \sim B)$. Sally prefers arbitrage to the status quo, so she prefers $(A \ \& \ B)$ to $(\sim A \ \& \ \sim B)$; she does not prefer $(\sim A \ \& \ \sim B)$ to $(A \ \& \ B)$. So she prefers $(\sim A \ \& \ B)$ to both $(A \ \& \ B)$ and $(A \ \& \ \sim B)$. Because

she rejects B , she prefers $(\sim A \ \& \ \sim B)$ to $(\sim A \ \& \ B)$. Therefore, she prefers $(\sim A \ \& \ \sim B)$ to $(\sim A \ \& \ B)$, prefers $(\sim A \ \& \ B)$ to $(A \ \& \ B)$, and prefers $(A \ \& \ B)$ to $(\sim A \ \& \ \sim B)$. These are cyclical and so incoherent preferences.

If Sally rejects both A and B , the blame for her mistake can be laid on her incoherent preferences rather than on her imprecise probabilities. To target her imprecise preferences, Elga's argument should assume that Sally makes choices without first forming preferences that direct her choices. I grant this assumption about the example.

Third, can Sally predict her choices about the sequence of offers? If she can, backwards induction applies to her choices, and rationality leads her to reject A only if she knows that she will accept B . She has no reason to reject A if she will reject B . However, suppose that Sally cannot predict her choices. Then her rejecting A prior to her rejecting B may be excused; it may arise from her ignorance that she will reject B . To prevent excuses for rejecting both gambles, the argument should assume, and I grant, that when deciding about A , Sally correctly predicts her choice about B . She does not know the outcomes of the gambles but knows her future choice about B .

2 Coherent choices

Standard decision theory evaluates a sequence of choices by evaluating its components. The sequence is rational if its components are rational. Rationality is compositional in this sense. A case for compositional evaluation appeals to the consistency of a theory of rationality's standards for single choices and its standards for sequences. A consistent theory does not maintain the rationality of the single choices and the irrationality of a sequence they constitute. It does not permit the choices in the sequence but prohibit the sequence.

Elga's remarks about the decision principle SEQUENCE show his acceptance of the standard method of evaluating sequences of choices. Elga disputes the rationality of single choices resting on imprecise probabilities rather than the sufficiency of the rationality of single choices for the rationality of a sequence they constitute. Nonetheless, Elga's example challenges rationality's compositionality. In Sally's case an irrational sequence of choices apparently arises from choices that are each rational.

Suppose that Sally rejects A and then rejects B . These choices are incoherent because they forgo a sure gain. How does the rationality of each choice in a set of choices ensure the coherence of the choices in the set, granting that relevant circumstances are constant and that agents are ideal? Decision principles in the literature use various techniques to achieve coherence.

The maximin principle selects for an agent an option that maximizes the agent's security level. It bypasses probabilities both sharp and imprecise. In the following decision tree, double lines mark the options that maximize Sally's security level in her sequence of decision problems.

	//B	5
//A	/H	-10
	\~B	
	\~H	15
	/H	15
	/B	
	\~H	-10
\~A		
	\~B	0

Suppose that Sally applies the maximin principle and knows that she will. When offered *A*, given each of her options, she knows what she will do when offered *B*. She will accept *B* if she accepts *A*, and will reject *B* if she rejects *A*. When deciding about *A*, she chooses the option that maximizes her security level. So she accepts *A*. Afterwards, she accepts *B*, as she foresaw.

Although strict principles of decision, such as the maximin principle, solve Sally's decision problems using backwards induction, permissive rules do not because they do not settle future choices. A permissive principle leaves open an agent's exercise of the permissions it grants.

Suppose that an agent is indifferent between options *A*, *B*, and *C*. Breaking ties between these options may lead to incoherent choices. An agent may pick *A* over *B*, *B* over *C*, and *C* over *A* to form a cycle of choices. Also, consider (1) a choice between incomparable options *A* and *B*, (2) a choice between incomparable options *A* and an improvement of *B* called *B+*, and (3) a choice between *B* and *B+*. A permissive principle may allow choosing *A* over *B+*, and also choosing *B* over *A*, although it requires choosing *B+* over *B*. Taken together, the two permitted choices and the required choice are cyclical.

Given indifference or an incomplete preference ranking of options, incoherence threatens for reasons independent of imprecise probabilities. Elga's argument assumes that in Sally's case imprecise probabilities generate any indifference or suspension of preference that leads to incoherent choices. I grant this assumption so that his argument attacks imprecise probabilities rather than only indifference or incomplete preference rankings.

A way for a permissive decision principle to ensure coherence in a sequence of decisions is for it to acknowledge that earlier decisions may affect the consequences of later decisions. Because of earlier decisions, the later decisions may generate an incoherent sequence of decisions and its bad effects.

Suppose that rationality is permissive concerning doxastic attitudes formed in light of evidence. Rationality may impose coherence on these attitudes by attending to the consequences of forming new attitudes. Suppose that it is rational to believe or to disbelieve that the universe is infinite. However, if an agent adopts one of these attitudes toward the universe's size, then rationality prohibits also adopting the other. Even if the evidence allows either attitude, rationality does not permit both because together they are incoherent. Holding a belief that *p* creates a reason not to hold a belief that $\sim p$. The coherence

requirements for beliefs become requirements for single beliefs in the context of other beliefs.

Similarly, rationality requires permissible preferences to form a coherent group. Given two preferences, rationality bars a third that creates a cycle. The coherence requirements for preferences become requirements for single preferences in the context of other preferences, assuming their rationality.

The challenge of Elga's example is to derive coherence requirements for Sally's sequence of choices from the requirements for each choice in the sequence. The next two sections show that because Sally makes each choice rationally, considering all its consequences, her choices form a coherent sequence. These sections show that Good's decision principle for single choices in finite decision problems does not yield incoherent choices in Sally's case. Despite its permissiveness, it rules out rejecting both gambles.

3 Coordination

A person facing a sequence of choices has opportunities to coordinate choices to improve the results of the sequence. A theory of rationality, to hold an agent responsible for making good use of opportunities for coordinating choices, evaluates a choice in a sequence with an eye on the choice's consequences for opportunities later in the sequence. When evaluating a chess player's move in a game, rationality considers whether it gains a winning position. A rational chess player whose only basic goal is winning cares derivatively about putting herself in position to win. She may use moves early in a game to achieve a position from which she can checkmate her opponent's king. Such consequences of the early moves may make them rational.

Rationality may allow a person to perform one but not both of two acts. Having done the first, the second becomes impermissible. Exercising a permission may affect the consequences of a subsequent act, and the change in consequences may render it impermissible. Imagine that someone may press button *A* or press button *B* but may not press both buttons because that triggers an unwanted explosion. After pressing button *A*, pressing *B* is irrational because it produces the explosion. Pressing *B* after pressing *A* yields a sequence that rationality prohibits. The bad consequences of the sequence accrue to the component that completes the sequence. Because past acts affect the consequences of current acts, evaluation of a current decision looks to the past as well as to the future. Although the decision's evaluation considers only its consequences, its consequences may depend on the past.

Suppose that Jane has only one basic goal, namely, the goal to gain money, and has a choice between a pair of dollar bills today and another choice between another pair of bills tomorrow. The bills available have the serial numbers 1, 2, 3, and 4. Jane is indifferent between the bills. However, she gains an extra dollar if the bills she picks today and tomorrow have adjacent serial numbers,

and she knows that today she chooses between bill 1 and bill 2, and tomorrow she chooses between bill 3 and bill 4. Assuming that she picks bill 2 today, when she picks a bill tomorrow, she is not indifferent between bill 3 and bill 4. Picking bill 3 brings an extra dollar. Her choice today affects the consequences of her options tomorrow.

Next, suppose that a weather forecaster assigns to rain the probability interval $[0.4, 0.6]$. The forecaster sells for \$0.40 a gamble that pays \$1 if it rains and otherwise nothing. Then she buys back the gamble for \$0.60, thereby losing \$0.20. Each transaction seems justified by Good's decision principle although rationality, assuming that the forecaster is averse to losing money, prohibits the pair of transactions because they result in a sure loss. In fact, Good's decision principle does not permit the pair of transactions because, when applied to the second transaction, it considers all the relevant consequences of buying back the gamble. The forecaster wants to avoid a sure loss and can do this by not buying back the gamble. So she should not buy it back. The context of the purchase affects its relevant consequences. That buying the gamble concludes a sequence of transactions that guarantees a sure loss is a relevant consequence of the purchase. The consequence, although not monetary, matters to the forecaster if she is rational.

In Sally's case rejecting both *A* and *B* does not incur a loss but instead forgoes a gain. Does it make a difference whether Sally ends up wasting an opportunity for arbitrage or, as the weather forecaster, ensuring a loss? The weather forecaster sells a gamble before buying it back. She is not in her original monetary situation after selling the gamble. Sally is in her original monetary situation if she rejects *A*, and so rejecting *B* then appears to be equivalent to rejecting *B* if it is offered in isolation. Rejecting *B* has the same monetary consequences in the two contexts. In both contexts rejecting *B* yields the status quo, no gain or loss. However, strictly speaking, Sally is not in her original situation after rejecting *A* because she has then forgone the opportunity for arbitrage. Moreover, rejecting *A* and then rejecting *B* ensures a loss if possession of the opportunity for arbitrage counts as an advantage equivalent to \$5. Rejecting *A* relinquishes the opportunity for arbitrage, and then rejecting *B* eliminates the prospect of compensation for relinquishing the opportunity. Framing rejection of *A* and then rejection of *B* as forgoing a sure gain rather than as incurring a sure loss does not affect evaluation of this sequence of choices.

Sally cares about gaining money. Because she is rational, she cares derivatively about opportunities to gain money. A principle of rationality requires an agent with an end to care about means of achieving the end. The expected-utility principle explicates the requirement. An option's utility equals its expected utility, as computed from the probabilities and utilities of its possible outcomes. The principle evaluates options as means to ends.

Elga's objections to PLAN and to SEQUENCE assume a principle of separability. Separability is often defined using conditional preferences. To ac-

commodate Sally's not forming preferences between her options, its definition may use choices. Sally's choice about B is separable from her choice about A , if no matter how she settles her choice about A , her choice about B is the same. Because the consequences of rejecting B depend on her choice about A , rationality does not require this type of separability for Sally's choice about A and her choice about B .

Sally, being rational and desiring money, has an aversion to wasting opportunities to gain money. Because she has this aversion, rationality does not require that her choices be independent. Her choice about B may depend on her choice about A . Rejecting A loses the opportunity for arbitrage, and then rejecting B eliminates the prospect of compensation for the loss. Rejecting B does not have this unwanted consequence if accepting A precedes it.

A rational agent abandons objectives that serve a basic goal after attaining the basic goal. It may seem that an agent's evaluation of a possible world should not consider attainment of derived goals. Consequently, it may seem that Sally should be indifferent between any two worlds in which she has the same amount of money. However, Sally, being rational, is averse to a series of decisions that without compensation loses an opportunity to gain \$5. A world in which she stays at the status quo without squandering an opportunity for arbitrage is better in her lights than a world in which she stays at the status quo by squandering an opportunity for arbitrage.

An objection claims that an evaluation of options that uses worlds as possible outcomes should evaluate worlds using only realizations of basic goals and basic aversions, thus ignoring desires and aversions concerning means. When evaluating rejection of B after rejection of A , the objection holds that Sally should not consider an aversion to wasting her opportunity for arbitrage because it is a derived, not a basic, aversion. The reply to this objection invokes a basic aversion of rational agents, namely, an aversion to wasting opportunities to reach basic goals. Because for Sally gaining money is a basic goal, any world in which Sally wastes an opportunity for arbitrage realizes her basic aversion to wasting opportunities to reach basic goals. An evaluation of the world should consider its realization of this basic aversion.

Let $\sim B$ by itself stand for rejecting B if it is offered alone. Sally does not rank together the world that results from realization of $(\sim A \ \& \ \sim B)$ and the world that results from realization of $\sim B$. She ranks the first world below the second because she squanders an opportunity to advance her basic goals if she rejects both A and B . Rejecting B when offered in isolation does not have the consequence of squandering an opportunity for arbitrage. The difference in the consequences of rejecting B in the two contexts may justify a difference in Sally's decisions regarding B in the two contexts. Rejecting B has different consequences for coordination in the two contexts. In contexts where an agent, whose only basic goal is to gain money, has opportunities to coordinate acts to gain money, an act's monetary consequences do not include all its relevant consequences. A relevant non-monetary consequence is achieving a position

to gain money. The agent, if rational, pays attention to such consequences. Should Sally decide about B without regard for her decision about A ? Should she consider only the monetary consequences of her decision about B ? No, that she squanders an opportunity for arbitrage if she rejects B after rejecting A is a relevant consequence for her.

In Elga's example, Sally's opportunity for arbitrage is an opportunity to coordinate her decisions about A and about B so that her sequence of decisions is sure to gain money. Because Sally knows that she will receive an offer of A and then an offer of B , she should exercise her opportunity for arbitrage, unless forgoing it yields prospects at least as good as arbitrage.

In Sally's case, rejecting A and then accepting B are parts of a coordinated sequence of choices that stakes her chance for money on H 's being true. Rejecting A puts Sally in a position to complete this sequence by accepting B . If instead she rejects B , she thwarts coordination of her choices. Rejecting B is less appealing after rejecting A than when B is offered in isolation because Sally cares about her position to gain money. Rejecting B after rejecting A closes her opportunities to gain money.

Assuming that for Sally rejecting A is permissible and the sequence rejecting A and then rejecting B is irrational, coordinating her choices by rejecting A and accepting B is worth more to her than rejecting B after rejecting A . Consequently, rejecting B after rejecting A does not maximize expected utility according to any pair of a probability assignment and a utility assignment in her set of such pairs.

Using the lowest probability in the interval for H drives down a calculation of the expected utility of B . If $P(H) = 0.1$, then $EU(B) = (0.1 \times 15) + (0.9 \times -10) = -7.5$. Because of Sally's aversion to squandering opportunities to gain money, $EU(\sim B) < -7.5$. Good's rule therefore prohibits rejecting B .

Attention to consequences besides amounts of money rules out Sally's rejecting both gambles given a strong aversion to wasting her opportunity for arbitrage. However, some versions of the example weaken this aversion. After rejecting A , Sally may forgo the chance for money that gamble B offers because of a weak aversion to wasting her opportunity for arbitrage. She may without irrationality forgo prospects for gains from accepting gamble B to squash the risk of losing \$10 if she accepts this gamble.

Attention to a choice's comprehensive consequences is insufficient to ensure coherent choices in all versions of Sally's case. The inequality $EU(\sim B) < -7.5$ must hold to prevent Sally's rejecting B . However, Sally may reasonably be no more averse to squandering her opportunity for arbitrage than she is averse to losing \$5, the monetary value of the opportunity. Moreover, changing Sally's probability interval for H may lower the value of $EU(\sim B)$ required to prevent B 's rejection. A revision of the case may use for H the probability interval $[0, 1]$ to make B 's expected utility go from -7.5 to -10 when calculated using the interval's lower endpoint. Sally's aversion to squandering her opportunity for arbitrage may not be strong enough to make her accept gamble B if its ex-

pected utility is as low as -10 . Also, the payoff table for gambles A and B may change to reduce the value of arbitrage and to increase the loss from gamble B if H is false. Sally's aversion to squandering her opportunity for arbitrage may not be strong enough to eliminate incoherent choices after adjusting the payoffs. The consequences of rejecting B after rejecting A prohibit rejecting B if Sally has a strong aversion to wasting her opportunity for arbitrage but does not prohibit rejecting B if her aversion is weak. Rejecting A does not make rejecting B violate Good's rule in all versions of the case.

4 Prediction

Rationality's compositionality and its permissiveness leave few methods of showing that Sally will not reject both A and B . Because of rationality's compositionality, the coherence of her sequence of choices must derive from the rationality of the choices in the sequence. Rationality's permissiveness, as expressed by Good's rule, prevents rationality's single-handedly grounding predictions of steps in her sequence of choices for applications of backwards induction. Nonetheless, an ideal agent, such as Sally, has resources for preventing incoherence. Her self-knowledge may ground predictions of steps in her sequence of choices, and her predictions may prevent incoherence.

In Elga's example, Sally, being ideal, may predict her choices, her exercises of rationality's permissions. If she predicts acceptance of B , then rejecting A amounts to putting her chance for money on H rather than on arbitrage. If she predicts rejection of B , then rejecting A is the first step toward wasting her opportunity for arbitrage. Her prediction of her response to the offer of B may ground her response to the offer of A .

The consequences of rejecting B depend on prior acts. If Sally rejects A , she thereby forgoes an opportunity for arbitrage. Forgoing the opportunity is not irrational if Sally then accepts B to put her chance for money on H 's being true. She may rationally do this instead of exercising her opportunity for arbitrage. However, if she also rejects B , then she not only rejects arbitrage but also rejects it without gaining as compensation a chance for money if H is true. Rejecting A trades a position from which Sally can guarantee gaining \$5 for a position from which she can gain more money if H is true. Then rejecting B trades that position for a position in which no gain is possible. Rejecting both gambles is a mistake because it wastes an opportunity for arbitrage.

Not all mistakes are irrational. Some mistakes are excused and so not irrational. Suppose that an agent rejects both gambles and thereby wastes the opportunity for arbitrage. This mistake occurs in two steps. Either rejecting A or rejecting B may be a mistaken step that circumstances excuse. However, Sally is an ideal agent without excuses and rationally avoids mistakes.

A compelling argument against Sally's rejecting both gambles uses her special traits. Given that Sally is a rational ideal agent, and so accurately

predicts her choice about *B*, she does not reject both *A* and *B*. After rejecting *A*, rejecting *B* may be rational, but the idealizations prevent Sally's rejecting *B*. Rejecting *A* requires as justification a prediction of *B*'s acceptance. Rejecting *A* is rational only if Sally foresees accepting *B*. Rejecting both *A* and *B*, given Sally's rationality, implies a failure to predict her rejection of *B*, and so a respect in which Sally is not ideal. Given that Sally is rational and ideal, she does not reject both *A* and *B*.

Suppose that Sally rejects *A*, accurately predicting that she will accept *B*. If Sally were to reject *B*, contrary to her prediction, then her rejecting *A* would have been a mistake. However, in that case she would have had an excuse for rejecting *A*. She would not have foreseen her rejection of *B*. Her lack of foresight, assuming that it is excused, would excuse her incoherent sequence of choices by excusing one of its components. Rationality demands coherence of ideal agents, but accepts excuses for a nonideal agent's mistakes; its standards take account of an agent's abilities. Sally's uncompensated loss of her opportunity for arbitrage would be the price she pays, not for having imprecise probabilities, but for not knowing how she will exercise rationality's permissions.

Although Sally is rational and ideal, her hypothetical realization of an incoherent sequence of choices changes her traits. If she retains her rationality while rejecting both *A* and *B*, she loses her power to predict accurately her choice about *B*. However given that Sally is rational and ideal, she correctly predicts her choice about *B* and so does not reject both *A* and *B*. Because Sally is an ideal agent, following Good's rule does not lead her into this mistake.

Good's permissive decision principle survives Elga's objections. The principle does not bring unsharp probabilities to grief. Its attention to all an act's consequences prevents mishaps in some cases, and an agent's predictions of her own acts prevent mishaps in other cases.

5 Objections and replies

Standard decision theory does not prescribe a way of moving from unsharp probabilities to decisions in sequences of decisions. Elga's case against unsharp probabilities claims that, in fact, no principle of rationality governs the move. Elga supports this claim by refuting contenders drawn from the literature. Sections 3 and 4 defend Good's permissive decision principle. Does the defense withstand objections?

First objection. Applications of Good's principle in sequences of choices requires keeping track of past decisions and their effect on the consequences of current decisions. This record keeping is demanding. An agent with limited cognitive capacities has a reason to assign sharp probabilities because they eliminate the need to keep track of past decisions. This point counts against a defense of Good's principle for nonideal agents. However, Elga's objections

treat ideal agents who lack excuses for failing to comply with familiar decision principles, and sections 3 and 4 defend Good's principle for ideal agents.

Second objection. Section 3's points about ends and means make Good's decision principle look to the past. Rational decisions look to the future. The future is separable from the past. Ranking options according to their futures produces the same ranking as ranking options according to their comprehensive outcomes—the past does not influence a rational choice among options. Sunk costs do not count, this objection claims.

Sunk costs do not count in most cases, however sunk costs count in some cases. Rational deliberation looks ahead to an act's consequences, but an act's consequences may depend on past acts and their effect on present opportunities. Suppose that John is indifferent between tea without milk and coffee with milk, but prefers his tea without milk, and prefers his coffee with milk. John pours some milk into a cup. Should he add tea or coffee? Before pouring milk, John is indifferent between pouring tea and pouring coffee. Pouring milk produces a preference for pouring coffee afterwards. Although pouring milk is permissible, and pouring tea is permissible, pouring milk affects the consequences of pouring tea and thereby lowers the utility of pouring tea. An act's consequences depend on its history, and its consequences affect its present utility.

6 Against sharpness

This paper has defended unsharp probabilities against objections but has not argued positively in support of unsharp probabilities. To close, it offers a brief argument that in some cases rationality not only permits but also requires unsharp probabilities.

Although Elga, because of his arguments against decision principles using unsharp probabilities, recommends that Sally assign a sharp probability to H , he does not recommend a particular sharp probability. He maintains just that Sally should assign some sharp probability to H so that she does not waste her opportunity for arbitrage. Rationality does not tolerate an imprecision in Sally's probability assignment, but tolerates an imprecision about the sharp probability she should assign.

Sharp probabilities, conforming to the probability laws, prevent preferences leading to Dutch books but may still generate irrational preferences. They may license some preferences that should not be formed. Suppose that $P(R) = 0.80$ and $P(S) = 0.81$, but these are arbitrarily precise probabilities, and evidence does not support the judgment that S is more likely than R . Then the evidence does not warrant preferring a gamble on S to a gamble on R although the sharp probabilities require that preference. Probabilities that are sharper than the evidence warrants lead to preferences that the evidence does not support.

Probabilities do not function properly as a guide to action unless they reflect the character of the evidence on which they rest.

Although requiring a particular sharp probability given a body of evidence prevents preferences that the evidence does not support, in some cases scant and sparse evidence does not support a particular sharp probability assignment, and fidelity to evidence requires imprecision. In these cases careful application of decision principles prevents their authorizing incoherent sets of choices. Fidelity to evidence and action guidance are compatible goals for probability assignments.

Acknowledgement

I benefitted from discussions with Arif Ahmed, Roger Clarke, Adam Elga, Nils-Eric Sahlin, and Teddy Seidenfeld.

Bibliography

- ELGA, A. [2010], Subjective probabilities should be sharp, *Philosophers' Imprint*, 10(5), 1–11, URL www.philosophersimprint.org/010005/.
- GOOD, I. J. [1952], Rational decisions, *Journal of the Royal Statistical Society, Series B*, 14(1), 107–114.